

UNIVERSITY OF SOUTHAMPTON

# Longterm Lightcurves of X-ray Binaries

Will Clarkson

Submitted for the degree of Doctor of Philosophy

SCHOOL OF PHYSICS AND ASTRONOMY  
FACULTY OF SCIENCE

March 19, 2004

This work is dedicated to William J. Davis, for providing an example and inspiration.

## Abstract

The X-ray Binaries (XRB) consist of a compact object and a stellar companion, which undergoes large-scale mass-loss to the compact object by virtue of the tight ( $P_{orb}$  usually hours-days) orbit, producing an accretion disk surrounding the compact object. The liberation of gravitational potential energy powers exotic high-energy phenomena; indeed the resulting accretion/outflow process is among the most efficient energy-conversion machines in the universe.

The Burst And Transient Source Experiment (BATSE) and RXTE All Sky Monitor (ASM) have provided remarkable X-ray lightcurves above 1.3keV for the entire sky, at near-continuous coverage, for intervals of 9 and 7 years respectively (with  $\sim 3$  years' overlap). With an order of magnitude increase in sensitivity compared to previous survey instruments, these instruments have provided new insight into the high-energy behaviour of XRBs on timescales of tens to thousands of binary orbits. This thesis describes detailed examination of the long-term X-ray lightcurves of the neutron star XRB X2127+119, SMC X-1, Her X-1, LMC X-4, Cyg X-2 and the as yet unclassified Circinus X-1, and for Cir X-1, complementary observations in the IR band.

Chapters 1 & 2 introduce X-ray Binaries in general and longterm periodicities in particular. Chapter 3 introduces the longterm datasets around which this work is based, and the chosen methods of analysis of these datasets. Chapter 4 examines the burst history of the XRB X2127+119, suggesting three possible interpretations of the apparently contradictory X-ray emission from this system, including a possible confusion of two spatially distinct sources (which was later vindicated by high-resolution imaging). Chapters 5 and 6 describe the characterisation of accretion disk warping, providing observational verification of the prevailing theoretical framework for such disk-warping. Chapters 7 & 8 examine the enigmatic XRB Circinus X-1 with high-resolution IR spectroscopy (chapter 7) and the RXTE/ASM (chapter 8), establishing an improved orbital ephemeris and suggesting the system may be in a state of rapid post-supernova evolution. In chapter 8 we follow this up with a direct search for the X-ray supernova remnant expected from such a system, concluding that with present observations the diffuse emission from Cir X-1 is indistinguishable from scattering by dust-grains in the interstellar medium.

### **Statement of Originality**

The work contained in this text is original unless otherwise stated. Material in chapters 4-8 has been published in the form of papers, which are cited in the text.

# Contents

1 X-ray Binaries	1
1.1 Close Binaries and XRB	1
1.1.1 LMXB and HMXB	2
1.1.2 The Ultracompacks	4
1.1.3 XRB Populations	4
1.2 Accretion and the Roche Lobe	5
1.2.1 Accretion Power	9
1.2.2 Accretion Disks	10
1.3 X-ray Behaviours and the Emission Regions	12
1.3.1 Standard Model Predictions	12
1.3.2 X-ray Pulsations	13
1.3.3 Type I X-ray Bursts	15
1.3.4 Inclination Dependent Dipping Behaviour and the ADC	16
1.3.5 Observed X-ray spectra of XRBs	18

---

1.3.6	Variability on the Sub-Day timescale . . . . .	20
1.3.7	Orbital Variability in $\dot{M}_2$ . . . . .	24
1.3.8	Interstellar Scattering and Extinction . . . . .	24
1.4	The properties of XRB at other wavelengths . . . . .	26
1.4.1	Ellipsoidal variation in optical brightness . . . . .	26
1.4.2	Optical Reprocessing . . . . .	26
1.4.3	X-ray heating of the donor . . . . .	27
1.4.4	Axial outflow from the system . . . . .	27
1.5	Summary . . . . .	28
<b>2</b>	<b>Superorbital Periodicities in X-ray Binaries</b>	<b>29</b>
2.1	Observations of Superorbital Periodicities . . . . .	29
2.1.1	A survey of superorbital periodicities . . . . .	30
2.1.2	Disk Precession and/or warping . . . . .	30
2.1.3	Longterm Changes in $\dot{M}_1$ . . . . .	32
2.1.4	Variations with Uncertain or Low Significance . . . . .	33
2.1.5	Undetermined Mechanism . . . . .	35
2.1.6	Superorbital Variations from External Galaxies . . . . .	35
2.1.7	Mechanisms for the Superorbital variations . . . . .	36
2.2	Non-planar, precessing accretion disks . . . . .	36

---

2.2.1	Time Evolution of a Warped Disk	37
2.2.2	Mechanisms of Warp Excitation	39
2.2.3	Radiation-Driven Warping	40
2.2.4	Wind-driven Warping	46
2.2.5	Lense-Thirring Precession: the Bardeen-Petterson effect	47
2.2.6	Magnetic Warping	47
2.3	Update: the Direct Approach	49
<b>3</b>	<b>Datasets</b>	<b>50</b>
3.1	Long-Term X-ray Datasets	50
3.1.1	RXTE/ASM	51
3.1.2	CGRO/BATSE	53
3.2	Timing Analysis Techniques	55
3.2.1	Lomb-Scargle Power Spectrum	55
3.2.2	Dynamic Power Spectrum	58
3.2.3	Dynamic Light Curve	60
3.2.4	Period Detection beyond the Power Spectrum	62
3.2.5	Phase Dispersion Minimsation (PDM)	62
3.2.6	Phase Scatter Minimisation	64
3.3	Significance Testing and Determination of Uncertainties	66

---

3.3.1	Lomb-Scargle Periodogram: Statistical Significance . . . . .	66
3.3.2	LS Periodogram: Uncertainty in the Detected Period . . . . .	68
3.3.3	Red Noise . . . . .	69
3.3.4	Phase Minimisation . . . . .	70
3.3.5	Non-Gaussian $p(x)$ . . . . .	70
<b>4</b>	<b>X-ray Variability in the Core of M15</b> . . . . .	<b>73</b>
4.1	Introduction . . . . .	73
4.2	RXTE/ASM Light Curve . . . . .	75
4.2.1	Period Search . . . . .	75
4.2.2	X-ray burst candidates . . . . .	78
4.2.3	The clockwork burster X1826-238: a comparison source . . . . .	79
4.3	The October 1988 burst and its implications for the ASM search . . . . .	80
4.4	The X2127+119 burst candidates . . . . .	81
4.4.1	ASM sampling and the true burst rate . . . . .	83
4.5	HST UV Imaging: a Variability Study . . . . .	85
4.5.1	The Nature of the Variables V1-V3 . . . . .	88
4.6	Discussion . . . . .	89
4.7	Update: M15 with Chandra & XMM . . . . .	91
<b>5</b>	<b>The Variable Third Period in SMC X-1</b> . . . . .	<b>92</b>

---

5.1	SMC X-1	92
5.1.1	Accretion Disk and Superorbital Variation	93
5.1.2	X-ray Spectral Behaviour	94
5.2	Dynamic Power Spectra of RXTE and BATSE datasets	96
5.3	Spectral Behaviour of the Superorbital Period	97
5.4	Discussion	100
5.4.1	Mechanisms of Superorbital Variation	102
5.4.2	Stability of Precessing, Warped Accretion Disks	105
5.4.3	How the precessing warp manifests itself	105
5.5	Conclusion	108
5.6	Update: Further Spin Evolution of SMC X-1	109
<b>6</b>	<b>Stability of XRB to Warping</b>	<b>111</b>
6.1	Introduction	111
6.1.1	Accretion Disk Warping	111
6.2	Source Selection	113
6.2.1	Discarded Sources	113
6.3	Dynamic Power Spectrum	115
6.4	Results	117
6.4.1	Her X-1	117

---

6.4.2 LMC X-4 . . . . .	119
6.4.3 Cygnus X-2 . . . . .	119
6.5 Discussion . . . . .	121
6.5.1 Her X-1: Disk Inclination Change? . . . . .	121
6.5.2 LMC X-4: Stable variation . . . . .	125
6.5.3 Cyg X-2: Multimode Variation? . . . . .	127
6.5.4 Stability of Accretion Disks to Warping . . . . .	131
6.6 Conclusion . . . . .	133
6.7 Update: examining warpers in more detail . . . . .	134
<b>7 Near-IR Spectroscopy of Circinus X-1</b>	<b>137</b>
7.1 Circinus X-1 . . . . .	137
7.2 IR observations of X-ray Binaries . . . . .	139
7.3 Observations . . . . .	140
7.4 Near-periastron spectra . . . . .	143
7.5 Spectra Near Apastron . . . . .	145
7.6 Can $M_2$ be detected from its radial velocity? . . . . .	148
7.7 Update: Further AAT/IRIS2 Observations of Circinus X-1 . . . . .	150
<b>8 X-ray Dips as a Diagnostic of Periodic Behaviour in Cir X-1</b>	<b>151</b>
8.1 Summary . . . . .	151

---

8.2	Introduction . . . . .	152
8.3	RXTE/ASM Light Curve . . . . .	152
8.3.1	Methods of analysis . . . . .	152
8.4	X-ray Behaviour . . . . .	154
8.4.1	Longterm Decline . . . . .	155
8.4.2	X-ray Maxima . . . . .	156
8.4.3	X-ray Dips . . . . .	157
8.5	Ephemerides . . . . .	159
8.5.1	Ephemeris based on X-ray Dips . . . . .	160
8.5.2	The Duration of the Dips . . . . .	160
8.6	Discussion . . . . .	163
8.6.1	A simple wind model . . . . .	164
8.6.2	The X-ray maxima . . . . .	166
8.6.3	The Accretion disk and the Maxima . . . . .	167
8.6.4	X-ray dips as tracers of $\dot{M}(t)$ . . . . .	168
8.6.5	Using the dips to investigate the binary parameters . . . . .	169
8.6.6	A simple model . . . . .	171
8.6.7	Is Cir X-1 a very young system? . . . . .	175
8.6.8	The External Magnetic Field of the Neutron Star . . . . .	177
8.7	Using Chandra to search for an associated SNR . . . . .	179

---

8.7.1	Spatially-resolved spectra . . . . .	180
8.7.2	Lateral Profiles . . . . .	183
8.7.3	Radial Profiles . . . . .	186
8.8	Conclusion . . . . .	188
8.9	Update: Radio Ejections from Cir X-1 . . . . .	189
<b>9</b>	<b>Conclusions</b>	<b>190</b>
9.1	Future Work . . . . .	190
9.2	Future All-Sky Monitors . . . . .	191

# List of Figures

1.1	Observational Characteristics of LMXBs ( <b>top</b> ) and HMXBs ( <b>bottom</b> ). Visualisation software developed by Rob Hynes. . . . .	3
1.2	<b>Left:</b> Roche equipotentials for the system configuration $M_1=0.85$ , $M_2=0.17$ and system separation $a=0.72 R_\odot$ . The axes are in units of $a$ . From Carroll & Ostlie (1996). <b>Right:</b> $\Phi_R$ plotted as a function of distance along the line joining the two centres of mass. From Carroll & Ostlie (1996). . . . .	7
1.3	Schematic of X-ray pulsar geometry, from NASA/HEASARC. . . . .	14
1.4	Type I X-ray bursts from the “clockwork burster” X1826-268. <b>Left:</b> RXTE/PCA observation of a typical Type-I X-ray burst. <b>Right:</b> burst recurrence time plotted against persistent flux. The dashed line is the trend expected for constant accumulated mass (so $t_{recur} \propto \dot{M}_1^{-1}$ ). From Strohmayer & Bildsten (2003). . . . .	16
1.5	EXOSAT-ME lightcurves of the prototypical ADC source X 1822-371 and the prototypical dipper EXO 0748-676. Notice that for X1822-371 the ellipse is partial, indicating a scattered component. Reproduced from Charles & Seward (1995). . . . .	17
1.6	The standard model of inclination dependence. At low inclination no dips are visible; at high inclination the disk rim shields the central source from direct sight, producing a low $L_X/L_{Opt}$ and ADC behaviour. At intermediate inclinations the disk rim leads to dips and/or eclipses. Reproduced from Frank, King & Lasota (1987). . . . .	18

---

1.7 Schematic of emission components as suggested by X-ray spectra of the dipping sources. <b>Left:</b> system geometry. The ADC is confined to a thin region above and below the accretion disk, producing a power-law component in the observed spectrum. As vertically extended disk matter obscures the inner disk, the power-law component drops rapidly in strength. <b>Right:</b> size of the ADC. By measuring the time taken for the power law component to disappear, the radius of the ADC itself is estimated. From Church & Balucińska-Church (2004) . . . . .	20
1.8 Sample Z-track from the IMXB Cyg X-2, based on EXOSAT observations.. From (Kuulkers et al., 1996). $\dot{M}_1$ increases from the Horizontal Branch (HB) through to the Flaring Branch (FB). . . . .	22
1.9 Colour-colour diagram of the Atoll LMXB 4U 1608-52, taken with RXTE/PCA. Soft colour is defined from RXTE/PCA channels as (3.5-6.0keV)/(2.0-3.0 keV), Hard colour defined as (9.7-16.0 keV) / (6.0-9.7 keV). From van Straaten et al (2003). . . . .	23
1.10 Schematic of a Be-X binary. The eccentric orbit of the neutron star takes it close enough to the donor to pass through overdense material flung out at the equator by the donor rotation (and observed in $H\alpha$ line profiles). Enhanced mass transfer at periastron powers intense optical and X-ray outbursts, and a (generally) short-lived accretion disk can form after periastron passage. Reproduced from Charles & Seward (1995). . . . .	25
2.1 Stability to warping for $\alpha = 0.3, \eta = 0.1$ in $(r_b, q)$ space. $r_b$ (the binary radius) is in units of $GM_1/c^2$ , while $q$ is the mass ratio of the system. For $r_b < 0.7$ , beneath the lowest solid line, no warp is possible. The first solid line corresponds to matter input at the outer Lindblad radius; every source above this line for which matter is input in this way can sustain mode 1 and higher warps. If matter is input at the disk circularisation radius, this region still corresponds to no warp. In the next region up, simple mode 0 warping is possible but cannot persist. Between the upper boundary of this region and the dotted line, mode 0 warping persists. Finally in the uppermost region, mode 1 and higher warping modes become possible. With $r_b \sim 9$ and $q = 10.8$ , SMC X-1 is on the boundary between stable mode 0 and higher modes of warping. From Ogilvie & Dubus (2001) . . . . .	45

---

2.2 Still frame from SPH simulation of a strongly irradiated accretion disk, showing a prominent warp, in this case for mass ratio $q = 0.075$ . The warp is stable and precesses in the binary frame. The prototypical warpers Her X-1 and SMC X-1 have $q > 1$ , a regime which will be the target of future work (Foulkes et al., 2004). Colours represent viscous dissipation. Reproduced with permission. . . . .	48
3.1 RXTE/ASM schematic. From Levine et al (1996). . . . .	52
3.2 <b>Left:</b> Raw RXTE/ASM source and Background lightcurves for the active but faint LMXB LMC X-4, binned at 1.5-day intervals to reduce scatter for the plot. <b>Right:</b> Count Rate-Frequency histogram for the same source, after filtering in the manner described in the text. Note the count rate distribution extends both sides of zero net count rate: to remove the negative points thus skews the distribution towards higher count rates. . . . .	53
3.3 Schematic of an individual LAD of CGRO/BATSE. From Harmon et al (2002). . .	54
3.4 RXTE/ASM lightcurves (top) and Lomb-Scargle periodograms (bottom) for the galactic bulge sources GX13+1 (Atoll source, donor unconstrained) and GX17+2 (Z-source LMXB). The recurrent sampling gap at 180d is clearly present in the lightcurves from both sources. . . . .	58
3.5 Illustration of the Dynamic Power Spectrum. <b>Left:</b> Fold profile of Her X-1, used as the template for a 43-day variation in the synthetic lightcurve. Also present in the synthetic lightcurve is frequency-independent white noise, and the sampling was taken from the true lightcurve of Her X-1. <b>Right:</b> Dynamic Power Spectrum of the synthetic lightcurve. The 43-day fundamental is most strongly detected. Also strongly present is a 21.5-day harmonic, due to the presence of a secondary as well as a primary peak in the fold (left). Also clearly visible is a weaker detection of a higher harmonic. . . . .	59

---

3.6 <b>Left:</b> Section of a simulated lightcurve with abrupt jump in the frequency of variation. The frequency of the input sinusoid switches instantaneously from 60 days to 43 days but maintains the same amplitude. The switch takes place at MJD 50500, and switches back at MJD 51300. <b>Right:</b> Dynamic Power Spectrum of this lightcurve. For an interval of four windows, no dominant periodicity is detected during the period switches. . . . .	60
3.7 Illustration of apparent period shift due to an incorrectly determined periodicity. Dataset used consists of white noise superimposed on the fold profile of Her X-1, with a fold repetition time of 10.00 days. <b>Left:</b> Dynamic Lightcurve of the fake dataset, using the input period of 10.00 days as the cycle length. <b>Right:</b> The same dataset, this time folded on a period of 10.02 days, i.e. 2% discrepant from the true periodicity in the dataset. . . . .	61
3.8 LS and PDM periodograms of a simulated 80.0-day variation. <b>Left:</b> Fold profile of the simulated white-noise + 80.0d variation (top), and the lightcurve with the 3000-d sinusoid added over this periodicity (bottom). <b>Right:</b> PDM (top) and LS (bottom) periodograms of the entire simulated datasets. The solid line denotes the 80-d variation over white noise, the dotted line includes the 3000-d variation. The horizontal dashed line corresponds to $> 99.7\%$ significance over white-noise, from Monte-Carlo simulations of noise simulations lasting 2500-d each. . . . .	64
3.9 Empirically determined CDF for 100,000 synthetic, white-noise lightcurves lasting 400 days each. Overplotted is are the fits using the exponential approximation (3.14) and the high-z approximation (3.15). . . . .	67
3.10 Period histogram for 100,000 synthetic datasets, with a 60-d periodicity input over white-noise. The shift in peak detection is likely to be due to spectral leakage to or from other periods due to the non-sinusoidal nature of the variation. . . . .	69
3.11 <b>Left:</b> Frequency chart of dip sampling for the RXTE/ASM lightcurve of Cir X-1. Note that in the majority of cycles, no dip is sampled. <b>Right:</b> Sample histogram of period detection trials for the phase minimisation technique. . . . .	71

---

3.12 <b>Left:</b> Illustration of the rejection method (from Press et al 1992). See text for details. <b>Right:</b> Frequency chart of random number dataset generated with the rejection method, using an example <i>non-gaussian</i> function $p(x)$ as probability distribution function. . . . .	72
4.1 RXTE ASM lightcurve and background before cleaning. . . . .	76
4.2 The RXTE ASM dataset, but with all points with background levels above 5 counts per second removed, all points with an error greater than 0.75 or a chisquared statistic greater than 1 removed. Data has been binned to 5 days per bin to highlight extended activity. . . . .	76
4.3 Periodograms over a range of periods from 2 to 1000 days. There is no sign of a 37d period. . . . .	77
4.4 Period analysis of X2127+119 performed on the lower two channels. . . . .	78
4.5 Burst candidates from X2127+119. The top plot represents the total ASM count rate, beneath this are the three ASM channels individually, in increasing order of hardness. At the bottom is the background behaviour. All plots are to the same scale. The broken lines on the plots represent 0 and 1 for each plot. . . . .	80
4.6 PCA and ASM lightcurves of the LMXB burster x1735-444. <b>Top:</b> RXTE/PCA lightcurve, 1.9-24.8 keV, 17th June 1999. <b>Bottom:</b> RXTE/ASM lightcurve near this observation (vertical dotted line) . . . . .	82
4.7 Count-rate vs spectral softness for high points in the lightcurves of X2127+119 and X1826-238 . . . . .	83
4.8 The 2000 burst recorded by the RXTE/PCA, from Smale (2001). Overplotted is our simple representation of the burst profile and the cases of overlap with the 90-s RXTE/ASM dwell, reading from top to bottom in the set of conditions (3.2). . . . .	84

---

4.9	Peak ASM count rate plotted against hardness ratio, for the Crab Nebula (top, offset by $65\text{ cs}^{-1}$ ), brightest points from the known <i>non</i> -burster X1700-37 (offset by $50\text{ cs}^{-1}$ ), the burst candidates from X1826-238 (offset by $30\text{ cs}^{-1}$ ) and burst candidates from X2127+119 itself (bottom, no offset). We find that this analysis does indeed differentiate bursting sources from simple variabilities, strengthening our confidence that the candidates selected for X2127+119 do indeed indicate snapshots of type I-like X-ray bursts. . . . .	86
4.10	HST WFPC2 U-band images of the core of M15, obtained in April 1994 (left) and October 1994 (right). The positions of AC211 and the three new variables are indicated. Each image is 11.5 arcsec across; North points to the top left hand corners, East to the bottom left corners. . . . .	87
4.11	Discovery of a second LMXB in M15 (from White & Angelini 2001). <b>Left:</b> ACIS/HETG zero-order image of M15, showing the X-ray source associated with AC211 and the new source, denoted X2. <b>Right:</b> ACIS/HETG lightcurves for the two sources. . . . .	91
5.1	<b>Left:</b> Optical lightcurve of SMC X-1, showing clear evidence for an extensive accretion disk. From van Paradijs & Kuiper 1984. <b>Right:</b> X-ray pointings of SMC X-1, from Wojdowski et al (1998). GINGA, ASCA1, ROSAT1, ROSAT3 and RXTE pointings occurred during the high state of the superorbital cycle, with ROSAT2, ROSATC and ASCA2 taking place during the low state. . . . .	94
5.2	Schematic of SMC X-1, from the established masses and inclinations. Visualisation software by Rob Hynes, UT Austin. . . . .	95
5.3	Section of the ASM and BATSE datasets for SMC X-1. Data has been binned to 7.78 days, or two orbital cycles, per bin. . . . .	96
5.4	Dynamic power spectrum for the 20-100 keV lightcurve of SMC X-1 observed by BATSE over 9 years (MJD 48361-51690; right) and for the 1.3-12.1 keV lightcurve provided by RXTE over 6 years (MJD 50083-52312; left). Contours spaced at 4-unit intervals of LS power for BATSE and 40-unit intervals for RXTE. . . . .	98

---

5.5	Periodogram of simulated dataset, in which the initial periodicity is abruptly replaced by a second periodicity at 43 days, to test the hypothesis that the apparent smoothness of the long period variation might be an artefact of the 50-d overlap of adjacent data windows. The DPS has clearly resolved the two periodicities and there is little apparent smoothing (see discussion in Chapter 2, in which this figure is also used). . . . .	99
5.6	ASM countrate, ASM hardness ratio, and scaled BATSE count rate, binned by phase in the quasi-stable superorbital period of 46 days that persists from MJD 50550 - 50800. Spectral Hardness is defined in this paper as the ratio of counts at 3.0 - 12.1 keV to those at 1.3 - 3.0 keV. . . . .	99
5.7	ASM spectral hardness plotted against total count rate, for the dataset of figure 5.6. 100	
5.8	ASM total countrate, together with channel - by - channel behaviour, on the same phase scale as figure 5.6. The low energy channels show smaller amplitude variations and smoother morphology of variation. . . . .	101
5.9	Upper limits on the inclination of third body orbits, as a function of the outer orbital eccentricity, for the three scenarios discussed in the text. Scenario 1 (solid line): $T_{mod} = 50\text{d}$ , $P_3 = 3.89\text{d}$ . Scenario 2 (dashed line): $P_3 = 50\text{d}$ , $T_{mod}$ unconstrained. Scenario 3 (dotted line): $T_{mod} = 7\text{years}$ , $P_3 < 43.9\text{d}$ from $M_3 > 0.08M_\odot$ . . . . .	104
5.10	Fit to spin periods of SMC X-1 since the beginning of the BATSE mission. <b>Top:</b> Spin periods measured by Ginga, ROSAT, ASCA and RXTE, as tabulated by Wojdowski et al (1998). Dashed line represents quadratic fit to the spin periods. <b>Bottom:</b> residuals (in ms) from the fit, with a 2583-day sinusoid overplotted (compare with figure 5.4) . . . . .	108
5.11	Spin periods and their residuals from quadratic fit, with the addition of the new points reported by Naik & Paul (2004). Crosses represent points tabulated by Wojdowski (1998), diamonds are the new points. The quadratic has been re-fit to take the new points into account; as can be seen the residuals show variability on a similar timescale to the variation in superorbital period noted in this chapter. . . . .	110

---

6.1	Sections of the RXTE/ASM lightcurves of the sources examined in this work, binned appropriately to emphasise their long-term behaviour. . . . .	116
6.2	DPS of Her X-1. <b>Left:</b> RXTE/ASM dataset, contours spaced at intervals of LS power 100. Day numbers run from the beginning of the RXTE/ASM dataset, i.e. MJD 50088. <b>Right:</b> CGRO/BATSE dataset, contour spacing LS powers of 20. Day numbers run from the beginning of the CGRO/BATSE dataset, i.e. MJD58369. . . . .	118
6.3	Superorbital lightcurve ( <i>left</i> ) and count rate-hardness relation ( <i>right</i> ) for Her X-1. The spectral hardness varies with the ASM count rate and there is a clear relation between the two quantities. . . . .	118
6.4	DPS of LMC X-4. Contours spaced in intervals of LS power 10 (by comparison, LS power 9 corresponds to 99.9% significance over noise - chapter 3) . . . . .	120
6.5	Superorbital lightcurve ( <i>left</i> ) and count rate - hardness relation ( <i>right</i> ) for LMC X-4. Though a more noisy dataset, there is a correlation between the variations of the two quantities. . . . .	120
6.6	DPS of Cyg X-2. Contours spaced at LS powers of 200. . . . .	122
6.7	Superorbital lightcurve ( <i>left</i> ) and count rate - hardness relation ( <i>right</i> ) for Cyg X-2, over day numbers 600-900. In contrast to Her X-1 and LMC X-4, the spectral hardness appears to have a negative correlation with the count rate. . . . .	122
6.8	The 35-day RXTE/ASM and CGRO/BATSE dynamic lightcurves of Her X-1. The phasing of the peaks undergoes a shift before the onset of the Anomalous Low State (ALS), and the relative significance of the peaks alters. When the ALS ends, the phasing is unaltered from the early years of the dataset. . . . .	123
6.9	Her X-1 Hardness-Intensity relation for the normal state ( <b>Top</b> ) and the ALS ( <b>Bottom</b> ), using 1-day data bins to improve S/N. There is no significant difference between the low-level emission in the two states. . . . .	124

6.10 Comparison of 1994 and 2001 Anomalous Low States. <b>Left:</b> RXTE/ASM (top) and CGRO/BATSE (bottom) lightcurves of the transition into the 2001 ALS, each binned at intervals of 3.4 days (i.e. two binary orbits). <b>Right:</b> CGRO/BATSE lightcurve at the time of the 1994 ALS as determined by Vrtilek et al (1994), at two timescales (top and bottom). Horizontal line represents the interval of the Vrtilek et al observations, at MJD 49209-49214.5. . . . .	126
6.11 Dynamic lightcurve of LMC X-4. Data window length chosen to include 5 cycles per window - each bin in this plot is then the average of typically 30 datapoints. . . . .	126
6.12 DPS of simulated Cyg X-2 lightcurves. <b>Top:</b> input data consists of the superorbital lightcurve from day number 600-900, which undergoes period change in a manner approximating the fundamental from figure 6.6. <b>Bottom:</b> As for top, but with 40-day periodicity added that is slaved to the fundamental, to approximate a non-sinusoidal shape of the variation. . . . .	128
6.13 Input lightcurve as for figure 6.12 (top), but with 40-day periodicity kept steady as the fundamental varies (i.e. two independent modes of variation). . . . .	129
6.14 Cyg X-2: DPS of the Ariel-5 dataset, binned to one-day averages, with contours at LS powers of 9 (corresponding to 99.9% significance over white-noise). . . . .	130
6.15 Schematic bifurcation diagram for radiation-driven warping. As the control parameter $r_b$ is increased, the number of stable precession solutions increases. Initially there are none, then stable mode-0 precession rising in frequency as $r_b$ increases (solid line). Near the mode region border, the solution becomes marginally unstable as short-lived mode-1 instabilities form (thick line). Finally in the mode 1+ region, two or more steady solutions are possible, and the system precesses at a combination of these warping modes. When the positions of the sources Her X-1, LMC X-4, SMC X-1 and Cyg X-2 are overlaid on this plot, we find that their behaviour is fully commensurate with their system parameters, as shown by their behaviour in the DPS. . . . .	132

---

6.16 (O-C) diagram of the main-on in the superorbital cycle of Her X-1, assuming constant ephemeris, from BATSE & ASM data (Still & Boyd 2004 in prep). The ALS is characterised by apparent step-changes in the superorbital periodicity, which are small enough that the DPS approach was not sensitive to them. . . . .	135
7.1 Schematic of the mysterious XRB Cricinus X-1. Observational properties are either ambiguous or contradict each other. See text for details. Figure prepared using visualisation software by Rob Hynes. . . . .	139
7.2 IRIS2 spectra of Cir X-1, taken on the nights 26-28 July 2002 in shared-risk service mode. The Br( $\gamma$ ) and He-I lines are clearly present, but so are extremely bright counterpart lines, displaced by $\sim 3100 \text{km s}^{-1}$ . . . . .	142
7.3 RXTE/ASM lightcurve of Cir X-1 during the K-band IRIS2 spectra taken in July 2002. The times of observation are denoted by the vertical solid lines on the plot. . .	143
7.4 Mean K-band spectrum of Cir X-1, along with the IR spectrum taken by Johnston, Fender & Wu (1999), the comparison Be X-ray binary XTE 1946+274 (Verrecchia et al 2002) and the microquasar GRS1915+105 (Harlaftis et al., 2001) . . . . .	145
7.5 RXTE/ASM lightcurve at the time of the JHK near-apastron spectra. Observation intervals are denoted by the vertical solid lines. . . . .	146
7.6 K-band IRIS2 service spectra, taken Sep 4 2003, i.e. just after apastron. Preliminary reductions courtesy S. Ryder, AAO. . . . .	147
7.7 H-band IRIS2 service spectra, taken Sep 5 2003, i.e. just after apastron. A strong set of nearly symmetrical double-peaked emission lines are present, which we identify as Br(10) - Br(14). Preliminary reductions courtesy S. Ryder, AAO. . . . .	147
7.8 J-band IRIS2 service spectrum, taken 4th September 2003.. Preliminary reductions courtesy S. Ryder, AAO. . . . .	148

---

7.9	Computed radial velocity profile for Cir X-1 assuming $e=0.6$ , $i^\circ = 60^\circ$ , and angle of observation such that the line-of-sight velocity is zero at periastron, producing the most pessimistic estimate. Here $M_1 = 1.4M_\odot$ (solid line) and $M_2 = 5.0M_\odot$ (dashed line). . . . .	149
8.1	RXTE/ASM dataset used in this chapter, binned at 5 days per bin. Throughout the figures in this chapter time is shown in units of days elapsed since MJD 50000, i.e. 10th October 1995. <b>Top:</b> RXTE/ASM count rate. <b>Bottom:</b> RXTE/ASM spectral hardness (see Section 8.2). . . . .	153
8.2	RXTE/ASM lightcurve of Cir X-1, normalised to a running average evaluated over intervals of 200 days (section 3.1). Count rate errors, not overplotted in this figure for clarity, are typically $0.3 \text{ cs}^{-1}$ for the “high” state ranging to $\sim 3 \text{ cs}^{-1}$ for the “declined” state. . . . .	153
8.3	Distribution in phase of the X-ray maxima, folded on the radio ephemeris (Stewart et al 1991). . . . .	154
8.4	Top: RXTE/ASM spectral hardness for a 50d stretch of the lightcurve starting at day number 490, during the “high” state. Bottom: RXTE/ASM lightcurve during the same state, illustrating the anticorrelation. For one cycle, phase zero according to the radio ephemeris is plotted as a vertical dotted line. . . . .	155
8.5	50-day sections from the RXTE/ASM lightcurve of Cir X-1 corresponding to the three “states” of the lightcurve. Top: starting at day number 490, count rate offset $340 \text{ cs}^{-1}$ . Middle: starting day 1425, offset $150 \text{ cs}^{-1}$ . Bottom: starting day 2350, no offset. For one cycle, phase zero according to the radio ephemeris is plotted as a vertical dotted line. . . . .	156
8.6	Top: Frequency chart for the entire RXTE/ASM dataset indicates the existence of a separate population of low-count rate “dips.” Bottom: ASM hardness-count rate plot for the RXTE/ASM dataset. . . . .	157

---

8.7 <b>Top:</b> Dynamic Power spectrum of the full RXTE/ASM dataset. <b>Middle:</b> Dynamic Lightcurve showing the X-ray maxima, folded on the radio ephemeris. <b>Bottom:</b> Dynamic lightcurve showing X-ray maxima, folded on the SP03 ephemeris. Note that for the dynamic power spectrum the day numbers represent the beginning of each 400-day data window in the power spectrum. . . . .	158
8.8 <b>Top:</b> Dynamic lightcurve of X-ray Dips, folded on the SP03 ephemeris. <b>Middle:</b> X-ray dips folded on the radio ephemeris. <b>Bottom:</b> X-ray dips folded on our ephemeris. In this figure the grayscale is inversely proportional to the RXTE/ASM count rate. . . . .	161
8.9 RXTE/ASM lightcurve at two epochs (diamonds), with $3\sigma$ bounds on the timing of X-ray dips, as predicted by the three ephemerides (section 8.5). In each plot the ephemerides, from top to bottom, are based on the X-ray maxima (SP03), X-ray dips (this work) and radio flares (Stewart et al 1991). Ephemeris predictions offset for clarity; no prediction of count rate during dipping is made. . . . .	162
8.10 CDF of the duration of the X-ray dips (in days) as determined from the RXTE/ASM dataset. The upper limit represents the time interval between the first “high” points before and after each dipping interval caught by the ASM, while the lower limit represents the interval between the first and last recorded point in the dipping interval. Points with zero minimum duration (i.e. only one measurement in a dipping interval) have been removed. . . . .	163
8.11 RXTE/PCA lightcurve of Cir X-1, at periastron interval as predicted by the X-ray dip ephemeris (vertical line). This lightcurve is not background subtracted, however the presence of a deep, long X-ray dip at phase 0, and subsequent and prior shorter dips, is evident. . . . .	164
8.12 Example wind-driven accretion model following section 8.6.1. . . . .	166
8.13 X-ray dip occurrence profile as a function of (presumed orbital) phase. Overplotted is a sample prediction of the BB84 model, in this case $M_2=2$ , $e=0.4$ , $H_p=7.5 \times 10^4$ km. . . . .	174

---

8.14 Combinations of orbital eccentricity and donor mass that fit the observed dip occurrence rate profile. From left: $T_4=0.1, 0.5, 1.0, 5.0, 10.0, 50.0$ and $100.0$ respectively.	175
8.15 Regions of extraction for Chandra ACIS-S/HETG spectra of Cir X-1. <b>Left:</b> ATCA radio map of the Cir X-1 region, indicating the presence of an extended $\sim 5 \times 10$ arcmin synchotron nebula. From Stewart et al (1993). <b>Middle:</b> On a larger scale, the extraction regions for the spatially-resolved X-ray spectra of Cir X-1. The regions corresponding to the point source, nebula, halo and background, superimposed on a low-brightness, high-contrast image of an original ACIS-S/HETG observation (in this case 1906). <b>Right:</b> The extent of the radio nebula to the same scale, calculated from the co-ordinates given in the radio map of Stewart et al (1993).	180
8.16 Spectra of the four extraction regions, with $\sim 2''$ separation of the nebula/halo/background regions from the dispersion pattern. The halo region contains significant contamination from the central point source, as can be seen from the similarity in the two spectra below 2 keV.	181
8.17 Illustration of the method used to obtain rectangular regions for each lateral profile. <b>Left:</b> a region of interest is first selected from the overall grating image. <b>Right:</b> Annuli are constructed with large radii and at large displacement from the point source. The result is nearly rectangular regions, parallel to one of the grating dispersion patterns. The annuli show appreciable curvature over the second arm of the "X-pattern," so this arm is not used for lateral profile analysis.	182
8.18 Regions of extraction for the five observations. The same area was used for observations 1700, 1905, 1906 and 1907, translated and rotated as appropriate, but for observation 706 the same area cannot be used.	183
8.19 Point Source (black), Nebula (red), Halo (green) and Background spectra (blue), for observation 1700 (top left), 1905 (top right), 1906 (bottom left) and 1907 (bottom right). All plots are to the same scale.	184

---

8.20 Lateral profiles as a function of $nH$ . <b>Left:</b> Cir X-1 during outburst, over the best-fit profile for the low- $nH$ source Cyg X-2. As expected, the profile is consistent with a single point source. <b>Right:</b> The same profile applied to the high- $nH$ source GX 13+1. The dispersion is clearly more broad than for the simple point-source, suggesting either increased background due to halo scattering or variation in dispersion optical path due to smearing of the point source. . . . .	185
8.21 The best-fit profile from GX13+1, scaled to the count rate of Cir X-1 during dipping. There is marginal evidence for an additional broad component. . . . .	185
8.22 Section from the exposure map image of Cir X-1 during dipping. Clear variations both within each chip and across the chipset can be seen. . . . .	187
8.23 Radial profiles for the low- $nH$ source Cyg X-2, the archetypal X-ray halo source GX 13+1, and Cir X-1 during intervals of dipping. The radial profiles for GX 13+1 and Cir X-1 are entirely consistent with each other. . . . .	187

# List of Tables

4.1	RXTE ASM Burst candidates . . . . .	86
4.2	HST/UV Variables in the Core of M15 . . . . .	88
6.1	Neutron star XRB exhibiting superorbital modulations . . . . .	115
6.2	Predictions of Radiation-Driven Warping model . . . . .	115
6.3	Detected periodicities . . . . .	121
7.1	Strongest emission lines from K-band spectra, 25-7 July 2002 . . . . .	144

# Chapter 1

## X-ray Binaries

We provide a brief introduction into the X-ray Binaries (XRB), their energetics and the phenomena which play out on orbital- or sub-orbital timescales. Section 1.1 provides a broad overview of the classes of object, while section 1.2 briefly introduces the physics at work in the accretion process and the formation of an accretion disk, central to most of the phenomena observed in this thesis. Section 1.3 outlines the characteristics of X-ray phenomena, while section 1.4 briefly introduces the characteristics of XRB at other wavelengths. The properties of long-term variations (on timescales greater than orbital) are deferred to Chapter 2.

### 1.1 Close Binaries and XRB

Most stars in the universe are members of binary systems, in which two stars orbit their mutual centre of gravity. Close binaries are a subclass of stellar binary, in which a compact object may accept substantial mass transfer ( $\gtrsim 10^{-10} M_{\odot} \text{ yr}^{-1}$ ) from a stellar donor, powering intense X-ray emission ( $L_X \sim 10^{36} \text{ erg s}^{-1}$  is typical in active neutron star binaries; van Paradijs & McClintock 1995). If the compact object is a neutron star or black hole, the energy available through accretion is generally similar, as both objects have similar radii ( $\sim 1 - 10 \text{ km}$ ), whereas for the same accretion rates the energy available through accretion is  $\sim 100$  times lower for white dwarf accretors (radii  $\sim 10000 \text{ km}$ ). This can be illustrated by considering the ratio  $\eta$  of the accretion luminosity to the rest-mass energy of infalling matter ( $\eta = L_X / \dot{M}c^2$ ); for neutron stars  $\eta \sim 0.1$ , for black

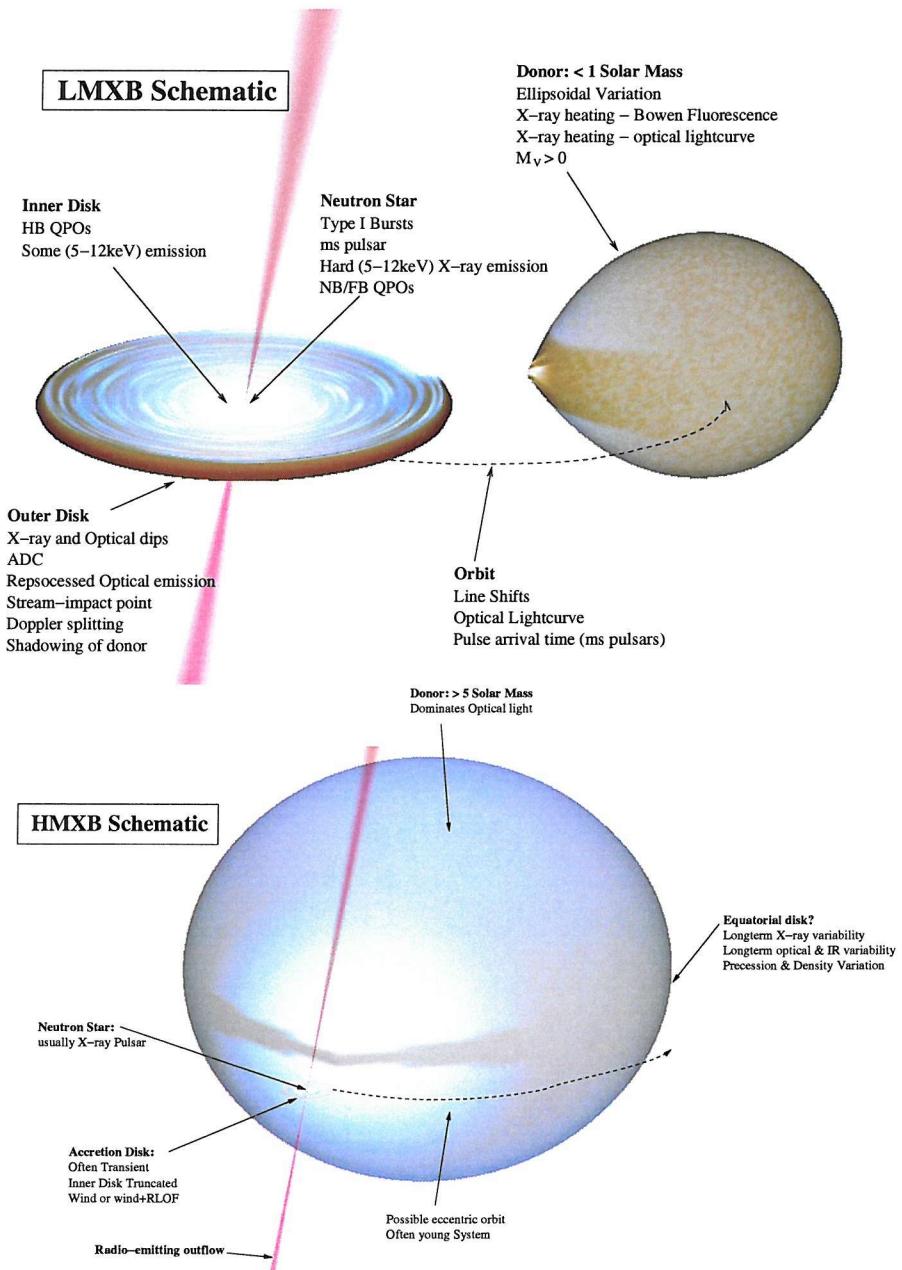
holes  $\eta \sim 0.05 - 0.4$  and for white dwarves  $\eta \sim 0.001$ . By comparison, nuclear reactions produce  $\eta \sim 0.001 - 0.01$ ; thus for the white dwarves non-accretion processes such as nuclear shell-burning can be significant or even dominant contributors to the persistent emission of the system (Warner, 1995).

By virtue of its angular momentum about the accretor, infalling matter will often not impact the accretor directly but settle into orbit about the object, spreading to form a disk through viscous interactions. Observations of this *accretion disk* provide a test of our understanding of the accretion/outflow process and in some cases vital constraints on the system properties.

For the remainder of this thesis we follow standard convention and label those close binaries with a neutron star or black hole accretor the X-ray Binaries (XRB), and treat systems with a white dwarf accretor as a separate class of object. For completeness we remind the reader that the white dwarf systems are classified according to their lightcurves and the source of principal behaviour; for example, the lightcurves of dwarf novae are dominated by powerful outbursts across all wavelengths, due to a global state transition in the accretion disk (Warner, 1995), while the Supersofts exhibit powerful ( $\sim 10^{37} \text{ erg s}^{-1}$ ) X-ray emission dominated by quasi-stable shell burning on the surface of the white dwarf (Kahabka & van den Heuvel, 1997).

### 1.1.1 LMXB and HMXB

X-ray binaries are classified according to the mass of the stellar donor, as this broadly divides the mechanism for mass transfer (see figure 1.1). The Low Mass X-ray Binaries (LMXBs) contain a  $\lesssim 2M_{\odot}$  donor, which transfers mass via Roche Lobe Overflow (see section 1.2). Donor types range from white-dwarves, through late-type main-sequence stars to A-type donors and F-G-type subgiants (Liu et al., 2001). Donors that are currently F-G subgiants may have started as *Intermediate* mass stars ( $1.5 \lesssim M_2 \lesssim 4M_{\odot}$ ), leading to a new class of XRB (Podsiadlowski & Rappaport, 2000). The High Mass X-ray Binaries (HMXBs) contain a  $\gtrsim 5M_{\odot}$  O- or B-type donor, where the mode of accretion depends further on the donor mass and system configuration; a  $\gtrsim 10M_{\odot}$  O-type donor produces a powerful stellar wind, some fraction of which may be gravitationally focused onto the accretor. HMXBs with a B-type donor usually power accretion by virtue of an eccentric orbit; during close passage at periastron mass transfer from donor to accretor is enhanced in a time-varying analogue of Roche Overflow. Roughly 70% of these systems, the Be-X-ray binaries, show strong  $H\alpha$  emission lines from an equatorial disk on the donor, from which matter



**Figure 1.1:** Observational Characteristics of LMXBs (top) and HMXBs (bottom). Visualisation software developed by Rob Hynes.

is taken at periastron passage (the reader is urged not to confuse these with B[e]-binaries, which show emission from a distribution of dust in the system (e.g. Clark et al., 2000).

### 1.1.2 The Ultracompacts

LMXBs with orbital periods  $\lesssim 60$  minutes are designated the Ultracompacts, as here the orbital period is so short that a main-sequence donor star is considered unlikely (section 1.2). The Ultracompact LMXBs are thought in most cases to have a donor that is certainly evolved and also at least partially degenerate, and indeed there is strong evidence that the accretion flow is virtually devoid of Hydrogen (Cumming, 2003). We remind the reader that the white dwarf ultracompacts are somewhat different in their accretion behaviour (classified AM CVn systems after the prototype of this class); here a variety of scenarios are suggested by observations of the various AM CVn systems; although the majority show disk evidence, a further sub-class exists in which the accretion stream may impact the donor directly without forming a disk first (see e.g. Warner 1995, Marsh & Steeghs 2002).

### 1.1.3 XRB Populations

At present, 280 XRBs are known to exist in the Milky Way and satellite galaxies, of which 150 are LMXBs and 130 HMXBs (Liu et al., 2001). In the Milky Way galaxy the vast majority are found in the galactic plane or the central bulge, though there are a small number of systems found in the halo. The sub-arcsecond imaging of the *Chandra* satellite, combined with the high effective area of *XMM* is bringing within reach the study of XRB populations in external spiral galaxies, in particular M31; there the XRB population is also associated with star-forming regions (e.g. Kong et al. 2002 et seq.). The Galactic Bulge sources not on the near-side of the Milky way (often designated with the prefix *GX*) suffer from high interstellar reddening ( $nH$  typically  $\gtrsim 10^{22}\text{cm}^{-2}$ ). Under these circumstances, direct detection of the donor is usually very difficult, and indeed the nature of the system components is for most systems a matter of some debate (e.g. Cornelisse et al., 2003). These sources are starting to reveal their secrets through investigations on two fronts: (i) the better studied systems at lower interstellar absorption are allowing system diagnostics based on X-ray emission alone, and (ii) the development of higher sensitivity instruments and telescopes of larger collecting area (e.g. VLT/ISAAC) is for the first time making direct detection of these objects possible.

The XRB population is remarkably less well-represented in the field as compared to the globular clusters and galactic tidal structures such as the Small Magellanic Cloud (SMC). The Milky Way contains  $\sim 280$  active XRBs out of  $\sim 10^{12}$  stars; scaling by the number of stars in the SMC thus predicts perhaps  $\sim 3$  XRBs in the SMC, in fact there are now over 60 HMXBs discovered in the SMC with no known LMXBs (Haberl et al., 2000; Liu et al., 2001), whereas the LMC contains  $\sim 30$  HMXBs (Haberl & Pietsch, 1999; Kahabka, 2002) and two LMXBs (LMC X-2 & RX J0532.7-692.6; Liu et al. 2001). The LMC and SMC are thought to have undergone tidal interaction  $\sim 0.2$ – $0.4$  Gyr ago (Gardiner et al., 1994), triggering star formation in both bodies and producing the massive overpopulation of HMXBs compared to the Milky Way. That the SMC has a higher population of HMXBs than the LMC supports this model as the smaller of the two interacting bodies will undergo the more severe disruption.

The LMXB population in globular clusters is also overdense compared to the galactic field by an order of magnitude (e.g. Verbunt & Hut, 1987), suggesting formation channels are at work in clusters that are not important in the galactic disk. The obvious candidate is interactions between stellar systems; the collisional timescale for a star in a globular cluster is several orders of magnitude shorter than for the galactic disk (e.g. Binney & Tremaine, 1987); as a result, each XRB may undergo several stellar encounters over the cluster lifetime of  $\sim 10^{9}$ – $10^{10}$  years. Two- and three-body encounters can thus drive the evolution of close binary systems, through either the capture by an isolated compact object of an encountered star through tidal dissipation (Fabian et al., 1975) or by three-body encounters in which a star may be exchanged for the encountered star (Hut et al., 1992).

## 1.2 Accretion and the Roche Lobe

We introduce here the physics of accretion in close binaries; further details may be found in earlier papers on this subject (e.g. Lubow & Shu, 1975) or graduate texts (e.g. Frank et al., 2002). The majority of stellar binaries are comparatively wide, with periods of order  $\sim 1$  year, which corresponds to separations  $\sim 1$  A.U. through Kepler's third law

$$a = 2.9 \times 10^{11} (M_1 + M_2)^{1/3} P_{day}^{2/3} \text{ cm} \quad (1.1)$$

where  $M_1$  and  $M_2$  are expressed in Solar Masses. Representing the equilibrium condition of each star by steady fluid flow (i.e.  $\partial \mathbf{v} / \partial t = 0$ ) and assuming circular orbits, conservation of momentum

for stellar material in a frame rotating with the binary then demands

$$\frac{1}{\rho} \nabla P + (\mathbf{v} \cdot \nabla) \mathbf{v} = -\nabla \Phi_R - 2(\boldsymbol{\Omega} \times \mathbf{v}) \quad (1.2)$$

where the *Roche Potential*  $\Phi_R$  is the sum of gravitational and centrifugal potentials, i.e.

$$\Phi_R = \frac{GM_1 M_\odot}{|\mathbf{r} - \mathbf{r}_1|} - \frac{GM_2 M_\odot}{|\mathbf{r} - \mathbf{r}_2|} - \frac{1}{2}(\boldsymbol{\Omega} \times \mathbf{r})^2 \quad (1.3)$$

We wish to evaluate the equilibrium configuration of the binary by seeking pressure contours. (We ignore here the momentum flux due to strong irradiation from one body on the other, which can be significant for very luminous accretors; see e.g. Phillips & Podsiadlowski 2002). We wish to evaluate the equilibrium configuration of the binary by seeking pressure contours. We assume the gas to be ideal so that

$$\nabla P = \frac{k}{m_H} (\rho \mu^{-1} \nabla T + T \mu^{-1} \nabla \rho + \rho T \nabla \mu^{-1}) \quad (1.4)$$

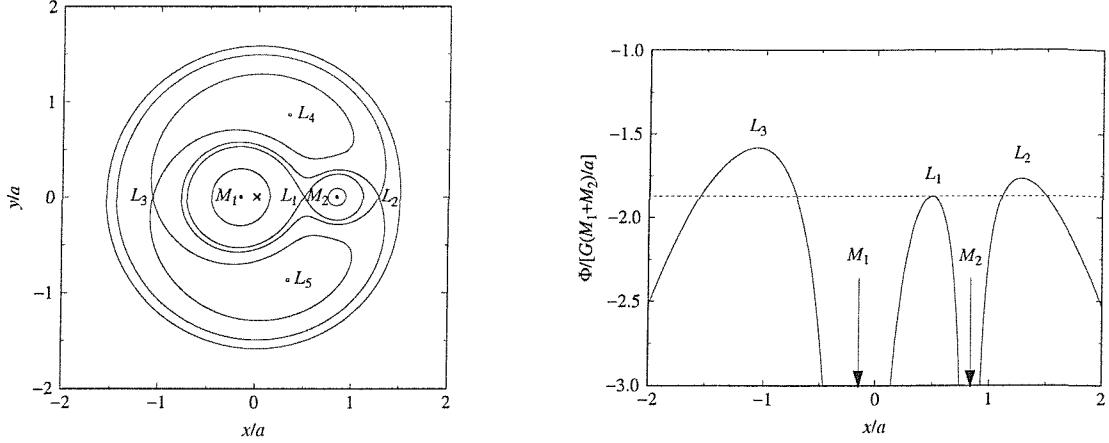
and make the further assumption that the chemical composition is small constant (so  $\nabla \mu^{-1} \simeq 0$ ). It is then straightforward to show using vector identities that if  $\nabla T$  and  $\nabla \rho$  are not parallel to each other, a flow  $\mathbf{v}$  relative to the binary frame must be established:

$$\frac{k}{\mu m_H} (\nabla \rho \times \nabla T) = \rho \nabla \times [2\boldsymbol{\Omega} \times \mathbf{v} + (\nabla \times \mathbf{v}) \times \mathbf{v}] \quad (1.5)$$

(If we allow the chemical composition to vary there is the additional possibility that nonzero heating may be balanced by gradients in composition  $\mu$ ). Conversely, in the absence of strong heating ( $\nabla \rho \times \nabla T = 0$ ) we have  $\mathbf{v}=0$  and thus hydrostatic equilibrium

$$\nabla P = -\rho \nabla \Phi_R \quad (1.6)$$

In this case  $\nabla P \times \nabla \Phi_R = 0$  and  $\nabla P \times \nabla \rho = 0$ , so surfaces of constant potential are then surfaces of constant pressure and density. In hydrostatic equilibrium, the surfaces of the stars can thus be obtained from plotting the equipotentials of the Roche Potential  $\Phi_R$  in the binary plane (figure 1.2). For small component separation relative to the stellar radii, the potential surfaces deform towards a teardrop shape. The *Roche Lobe* is the equipotential surface passing through the Lagrangian Point  $L_1$ , at which the force on a test mass is zero.  $L_1$  is a saddle point in the potential: in the direction along the line connecting the centres of mass (the  $x$ -direction in Figure 1.2) it is a local maximum, whereas in the  $y$ -direction it is a minimum. Thus matter from one star passing through the neighbourhood of  $L_1$  will tend to be channeled towards the other star rather than escaping the critical surface altogether. The radius of a Roche Lobe-filling star is approximated by the radius of a sphere filling the same volume; tabulations of this radius as a function of mass ratio (Kopal,



**Figure 1.2:** Left: Roche equipotentials for the system configuration  $M_1=0.85$ ,  $M_2=0.17$  and system separation  $a=0.72 R_\odot$ . The axes are in units of  $a$ . From Carroll & Ostlie (1996). Right:  $\Phi_R$  plotted as a function of distance along the line joining the two centres of mass. From Carroll & Ostlie (1996).

1959) have been fitted to better than 1% accuracy by the formula of Eggleton (1983):

$$\frac{R_2}{a} = \frac{0.49q^{2/3}}{0.6q^{2/3} + \ln(1+q)^{1/3}} \quad (1.7)$$

with the labeling convention

$$q = \frac{M_{\text{donor}}}{M_{\text{accretor}}} = \frac{M_2}{M_1} \quad (1.8)$$

For mass ratios ( $0.1 \lesssim q \lesssim 0.8$ ) an alternative form for  $R_2$  may be used (Plavec & Kratochvil, 1964)

$$\frac{R_2}{a} = 0.462 \left( \frac{q}{1+q} \right)^{1/3} \quad (1.9)$$

which allows  $a$  to be eliminated from Kepler's Third Law, and under the assumption that the low-mass donor has evolved slowly and is thus still on or near the main sequence, results in the period-mass relation

$$M_2 \sim 0.11 P_{hr} \quad (1.10)$$

with  $M_2$  measured in solar masses. Indeed, for low-mass Roche Lobe-filling X-ray Binaries, orbital periods of order 1 hour are frequently observed.

Roughly 95% of the binaries in the Universe consist of two widely separated, main-sequence stars, in which the stellar radii are tiny fractions of the separation between them. For example, a system containing two main-sequence stars at 1 A.U. separation have  $R/a \sim 10^{-3}$  for each star, so from figure 1.1 we expect the binary rotation to be unimportant as the gravitational terms in (1.3) dominate and the stars assume spherical configurations. However, once the more massive star evolves off the main sequence into the Asymptotic Giant Branch phase, it can easily attain

$R_2/a \sim 0.2$ , enough for the star to fill its Roche Lobe and start matter transfer to its companion.

Taking the orbital angular momentum of the binary to be that of two point masses

$$J = M_1 M_2 \left( \frac{G a}{M_1 + M_2} \right)^{1/2} \quad (1.11)$$

and assuming this mass transfer is conservative (i.e.  $\dot{M}_1 = -\dot{M}_2$ ), the effect of this transfer on the Roche Lobe of the donor is then

$$\frac{\dot{R}_2}{R_2} = 2 \frac{\dot{J}}{J} - 2 \frac{\dot{M}_2}{M_2} \left( \frac{5}{6} - q \right) \quad (1.12)$$

where  $\dot{M}_2 < 0$ . Thus in the majority of initial configurations ( $q > 5/6$ ) the mass transfer results in the Roche Lobe shrinking, resulting in further mass transfer and a runaway process. Mass transfer from more massive  $M_2$  to less massive  $M_1$  redistributes matter away from the centre of mass, so the remaining mass  $M_2$  moves in to conserve angular momentum; as a result the binary separation decreases along with the runaway increase in accretion rate. This process takes place on the thermal timescale of the donor ( $\sim 10^{3-7}$  yr depending on the donor radius), so corresponding systems must be very rare, representing an intermediate step to the final state of the binary. The runaway spin-together is only halted if sufficient mass is transferred to bring the mass ratio beneath the critical value  $5/6$  before system coalescence. For  $q \lesssim 5/6$ , mass transfer will expand the Roche Lobe, which would oppose accretion as the donor will no longer overflow the Roche Lobe; however observationally the majority of XRBs are indeed in this mass ratio regime. In wide ( $a \gtrsim$  few  $R_\odot$ ) binaries, post main-sequence donor evolution may cause the donor expansion to keep pace with expansion of the Roche Lobe. However the majority of LMXBs have much more compact orbits than this ( $a \sim R_\odot$ ), so accretion must be driven by angular momentum loss from the system; common mechanisms are magnetic braking and (for  $P_{orb} \sim$  few hours) gravitational radiation (see e.g. Verbunt 1993 or Frank et al. 2002 for discussion).

An evolving binary has a wide variety of evolutionary paths open to it, depending on the masses at the initial configuration. Throughout this work we focus on systems in which the end product is a neutron star in orbit with some less-compact star than itself. The significant mass overflow along the evolutionary paths, combined with heating and stripping from the supernova event producing the neutron star, can drive the evolution of the companion in several directions; as a result, a wide range of donor masses are possible in observed systems. However, the universal system configuration of neutron star X-ray binaries consists of a neutron star in a compact orbit with some stellar companion. More specifically, in the X-ray Binaries, significant mass transfer takes place onto the neutron star ( $\dot{M}_1 \sim 10^{-8} M_\odot \text{yr}^{-1}$ ), powering observable high-energy phenomena through the accretion of transferred matter.

### 1.2.1 Accretion Power

The total power available in the accretion process is just the rate of conversion of gravitational potential energy when matter falls to the surface of the compact object from large distances, i.e.

$$W_{acc} = \frac{GM_1\dot{M}_1}{R_1} \quad (1.13)$$

So for a neutron star ( $M_1 \sim 1.4M_\odot$ ,  $R_1 \sim 10\text{km}$ ), the power output of the accretion process is then

$$W_{acc} \simeq 8.2 \times 10^{37} \left( \frac{\dot{M}_1}{10^{-7}} \right) \text{ergs}^{-1} \quad (1.14)$$

with  $\dot{M}_1$  in units of  $10^{-7}M_\odot\text{yr}^{-1}$ . The nature of the mass transfer broadly follows the division of the X-ray binaries into the mass classes (section 1.1.1).

A rough upper limit to the sustainable accretion luminosity may be derived by calculating the luminosity at which radiation pressure due to the radiation produced by accretion is sufficiently strong to halt accretion. Photons from the central source undergo Thomson (free-free) scattering off incoming electrons in the accretion flow; electrostatic coupling between protons and electrons thus produces a pressure opposing the infall of matter. The *Eddington Limit* is the value of this luminosity for spherical accretion, and is given by

$$L_{Edd} = 1.3 \times 10^{38} M_1 \text{ergs}^{-1} \quad (1.15)$$

(or alternatively  $\dot{M}_{1,Edd} \simeq 1.5 \times 10^{-8} R_{10} M_\odot \text{yr}^{-1}$ , where  $R_{10}$  is the accretor radius in units of 10km).

However, there are at least two ways for the accretion flow to sidestep this limit. *Einstein* observations of seven bright spiral galaxies detected 16 point sources with observed  $L_X > 10^{39} \text{erg s}^{-1}$  (Fabbiano, 1989); assuming such emission to be isotropic and Eddington-limited then produces the estimates  $M_1 \gtrsim 50 - 100 M_\odot$  and  $\dot{M} \gtrsim 10^{-6} M_\odot \text{yr}^{-1}$  (e.g. Makishima et al 2000); this is the Intermediate Mass Black Hole (IMBH) model. Observations with ROSAT and ASCA yielded the discovery that these Ultraluminous X-ray sources (ULX) are preferentially, but not exclusively, distributed in star-forming regions (Roberts & Warwick, 2000), suggestive of an XRB association and resulting in a considerable resurgence of interest in the nature of these objects (Makishima et al., 2000). The main objection to spherically symmetric models is the difficulty of forming a large number of intermediate mass black holes in a  $\sim 1\text{yr}$  orbit with an evolved companion (King et al., 2001); however individual objects may still contain IMBH accretors. The alternative model involves deviation from spherical accretion, with mild beaming of emergent radiation (for example

through warping of the accretion disk) enhancing the observed  $L_X$ . The high  $\dot{M}$  required could be achieved during the thermal-timescale stage of mass transfer inevitable in HMXBs, thus fitting the association with starburst regions.

The second method to overcome  $L_{Edd}$  is to decrease the efficiency with which the radiation pressure couples with incoming matter; for example, the HMXB X-ray pulsar A0538-66 frequently exhibits episodic outbursts in which  $L_X > 10^{39}$  erg s $^{-1}$  (Charles et al., 1983), but this may be due to reduction of the scattering cross section below the Thomson cross section at the high magnetic field  $\gtrsim 10^{11}$  G (Boerner & Meszaros, 1979). Thus in the X-ray pulsars we may expect supereddington accretion rates to be sustained (see also King et al., 2001).

An important caveat to note is that the Eddington limit applies only to the X-ray emission which interacts with the accretion flow; it is thus irrelevant to X-rays generated in regions *outside* the central object. For example, spectral fits of the neutron star LMXB Sco X-1 during flaring show peak  $L_X$  at  $\sim 1.1 \times 10^{39}$  erg s $^{-1}$ , well in excess of its Eddington limit (Barnard et al., 2003a). The majority of this emission arises in an extended region well beyond the inner accretion flow (section 1.3)

### 1.2.2 Accretion Disks

Matter flowing through  $L_1$  freefalls towards the accretor, attaining a locally supersonic ballistic trajectory (Lubow & Shu, 1975). As this matter has high specific angular momentum, the matter stream often will not impact directly onto the neutron star but instead go into orbit around it, at radius approximately

$$R_c = a(1+q) \left( \frac{R_{L1}}{a} \right)^4 \sim 0.2a \quad (1.16)$$

where  $R_{L1}$  is the distance from the neutron star to the lagrange point  $L_1$ . The result is an accretion disk, in which matter builds up to some steady state around the neutron star. We introduce the standard model of such disks here, which is based on the thin-disk model of Shakura & Sunyaev (1973). Angular momentum is transferred outwards by viscous processes, resulting in the gradual inward drift of matter towards the acceptor. The outward transfer of angular momentum results in the accretion disk spreading beyond  $R_c$ , in some cases approaching  $R_{out} \sim 0.8R_{L1}$ . If the disk obeys the thin-disk condition, i.e. its scale height  $H$  obeys

$$H \simeq \frac{c_s R}{v_K} \ll R \quad (1.17)$$

where  $v_K$  is the Keplerian velocity, then the vertical structure of the disk decouples from its horizontal structure, whose surface density  $\Sigma$  is described by the nonlinear diffusion equation of Pringle (1981):

$$\frac{\partial \Sigma}{\partial t} = \frac{3}{R} \frac{\partial}{\partial R} \left( R^{1/2} \frac{\partial}{\partial R} (\nu \Sigma R^{1/2}) \right) \quad (1.18)$$

where  $\nu$  is the kinematic viscosity. To make further progress, a prescription must be found for the viscosity  $\nu$ ; although there is not yet a consensus on the form of this prescription, there has been some promising recent work. Hydrodynamic turbulence alone fails; globally the disk satisfies Rayleigh's criterion for stability against axisymmetric perturbations, while local perturbations (after artificially stirring up the disk) result in the transport of angular momentum *inwards* (Balbus et al., 1996; Tout, 2000, e.g.). Whether convection can produce viscosity of the right magnitude and sign is currently an open question, as it depends on the model and computational approach used (e.g. Ryu & Goodman, 1992; Kumar et al., 1995). Magneto-Hydrodynamic (MHD) turbulence is the currently favoured possibility, as it appears to predict viscosity of reasonable magnitude (Frank et al., 2002). Shear flow (with angular velocity  $\Omega$  increasing outwards from the centre) will tend to cause vertical field lines to become bent; this is opposed by the tendency of field lines to straighten out, which thus attempts to transport angular momentum outwards to enforce corotation (Balbus & Hawley 1991; see also Tout 2000). The low magnetic diffusivity of astrophysical plasmas allows angular momentum transport outwards even for weak vertical magnetic fields. Whichever mechanism is physically active in the viscosity, the standard prescription is to take  $\nu$  to be linearly related to the sound speed  $c_S$  and pressure scale height  $H$  through the dimensionless parameter  $\alpha$ :

$$\nu = \alpha c_S H \quad (1.19)$$

with  $\alpha \lesssim 1$ . This dependence appears broadly sensible as the size of turbulent eddies are unlikely to exceed the disk thickness  $H$ , and the velocity of their motion will likely be below  $c_S$ , otherwise shocks would thermalise their motion (Frank et al., 2002).

Dynamical equilibrium will be established in an accretion disk on roughly the timescale for particle orbit  $t_{dyn} \sim R/v_k$ , while the disk is expected to adjust to changes in  $\Sigma$  on the viscous timescale  $t_{visc} \sim (R^2/\nu) \sim (R/v_R)$  where  $v_R$  is the radial drift velocity. Thermal balance is reestablished in the disk on the thermal timescale, defined as the heat content divided by the dissipation rate, i.e.  $t_{th} \simeq (\Sigma c_s)/D(R)$ . For the  $\alpha$ -disk solutions of Shakura & Sunyaev (1973), typical values for these quantities are hours for  $t_{dyn}$  &  $t_{th}$ , and weeks-months for  $t_{visc}$ .

In the case of wind-driven HMXBs, matter is only captured if its velocity relative to that of the neutron star is low enough for capture, corresponding to a capture distance  $R_{cap} \sim 2GM_1/v_{rel}^2$ .

The corresponding angular momentum for captured matter is a factor  $\sim (R_{cap}/a)^2$  lower than for matter accreted from the Roche stream (Frank et al., 2002). Thus the formation of a disk is a more doubtful outcome for the HMXBs than for LMXBs (however not unknown, as disks are emphatically observed in the HMXBs SMC X-1 and LMC X-4; chapters 5 & 6).

### 1.3 X-ray Behaviours and the Emission Regions

The two populations of X-ray binaries, LMXBs and HMXBs, show markedly different X-ray behaviours, resulting largely from the difference in the modes of mass transfer, mass accretion onto the neutron star, and the age of the system (all of which relate to the donor mass  $M_2$ ). The properties of the LMXBs are consistent with them being a more evolved population of objects; the  $\sim 150$  known systems are mostly distributed in the galactic bulge and in globular clusters, consistent with Population II objects. Furthermore, the orbits of LMXBs are all very nearly circular, and the neutron stars themselves show evidence for a weak external magnetic field. By contrast, the  $\sim 150$  known HMXBs are concentrated in the galactic plane (or with star formation regions in external galaxies; see section 1.1.1), often have eccentric orbits, and show evidence for a strong neutron star magnetic field. We introduce here the X-ray characteristics of the two classes of object, first introducing very briefly the predictions of the standard “thin-disk” models (Shakura & Sunyaev, 1973). The reader is cautioned that the division between the classes is not ironclad, as there are a handful of systems that show properties from both classes.

#### 1.3.1 Standard Model Predictions

In the Newtonian limit, the binding energy of the innermost possible Keplerian orbit of mass  $m$  of gas is  $\frac{1}{2}(GM_1m)/R_1$ . Matter falling from  $L_1$  initially has negligible binding energy, so this must be the energy released by mass  $m$  as it spirals inwards through the disk; therefore the total power output of this disk in a steady state must be  $L_{disk} = \frac{1}{2}(GM_1\dot{M}_1)/R_1$ , i.e.  $L_{disk} = \frac{1}{2}W_{acc}$ . Half the total accretion luminosity is thus released between the inner accretion disk and the surface of the compact object, in an accretion column (for X-ray Pulsars) or thin boundary layer.

The location of the X-ray generation in the disk may be estimated for thin disks from the basic parameters of the binary (Pringle 1981). Taking the rotational velocity of disk material to be

Keplerian, then making the assumption that the viscosity adjusts itself to maintain steady  $\dot{M}_1$  (so substituting  $\nu\Sigma(r, R_1, M_1, \dot{M}_1)$  for the dependence of the dissipation rate  $D(r)$  on  $\nu$ ) gives  $D(r)$  in a *steady* disk as (Pringle 1981):

$$D(r) = \frac{3GM_1\dot{M}_1}{4\pi R_1^3 f_r^3} \left[ 1 - f_r^{-1/2} \right] \quad (1.20)$$

where  $f_r = \frac{R}{R_1}$ . From this we see that the vast majority of the dissipation occurs within  $\sim 3R_1$  from the compact object. For a neutron star XRB, for which the disk tidal truncation radius is typically  $\sim 0.9R_{L1} \sim 10^6$ km, the dissipation thus takes place in the inner  $\sim 0.01\%$  of the disk radius. In regions where the disk is optically thick to X-ray emission, in the steady state disk elements may be expected to radiate as blackbodies, predicting a temperature profile  $\sigma T^4(r) = D(r)$ . Early X-ray spectra of LMXB's were well-fit with a model including just this component (Mitsuda et al., 1984), however recent progress suggests this blackbody emission should be regarded as seed photons for Compton scattering off high-temperature electrons, producing the power-law Compton component that is preferentially observed (section 1.3.5).

### 1.3.2 X-ray Pulsations

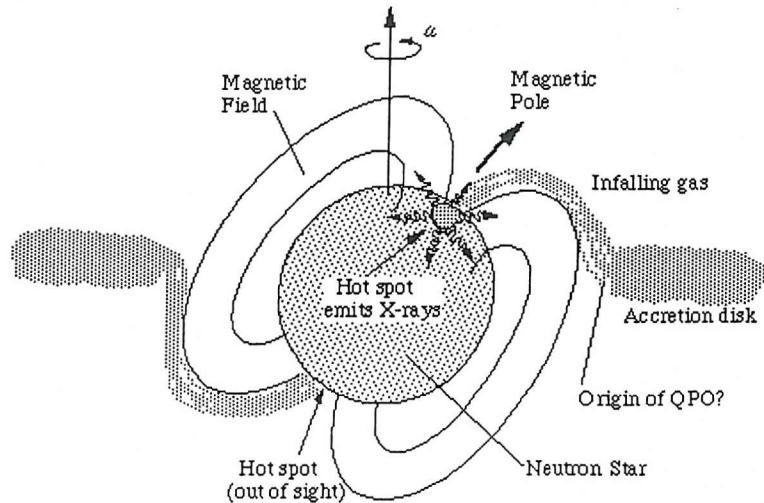
The HMXBs are often discovered as X-ray pulsars, in which the neutron star magnetic field is strong enough to disrupt the inner disk, channeling the accretion flow towards the poles (figure 1.3)<sup>1</sup>. The disruption radius  $r_M$  at which this occurs (for spherical accretion) is found by balancing the magnetic energy density with the kinetic energy density of infalling matter, resulting in

$$r_M = 1.49 \times 10^8 B_S^{4/7} R_1^{12/7} M_1^{-1/7} \dot{M}_1^{-2/7} m \quad (1.21)$$

with  $R_1$  in units of  $10^4$ m,  $M_1$  in units  $M_\odot$ ,  $\dot{M}_1$  in units  $10^{-7} M_\odot \text{yr}^{-1}$  and the surface field strength  $B_S$  in units  $10^{12}$ G. Observed spin periods in the HMXBs are typically  $\sim 0.1 - 10$ s (Liu et al., 2000), suggesting the compact object must be a neutron star from breakup considerations, with a surface field strength  $B_s \gtrsim 10^9$ G and accretion rate  $\dot{M}_1 \sim 10^{-7} M_\odot \text{yr}^{-1}$ . X-ray pulsars have been seen with much longer pulse periods  $\sim 100 - 1000$ s (e.g. Corbet et al., 2004).

Depending on its angular momentum about the neutron star, accreting matter causes the spin rate of the pulsar to evolve. The *corotation radius*  $r_C$  is defined as the radius at which matter in a Keplerian orbit corotates with the surface of the neutron star. At the inner disk edge  $r_M$ , the magnetosphere forces disk matter to corotation. If this process imparts angular momentum to the

<sup>1</sup>figure from <http://lheawww.gsfc.nasa.gov/users/white/xrb/xrb.html>



**Figure 1.3:** Schematic of X-ray pulsar geometry, from NASA/HEASARC.

material (i.e.  $r_M > r_C$ ), accretion cannot occur. Alternatively, if angular momentum is taken from the material (so  $r_M < r_C$ ), accretion occurs with a corresponding spin-up torque on the neutron star. For a 10km,  $1.4 M_\odot$  neutron star, then, accretion will only occur if

$$\frac{r_M}{r_C} \simeq 88.7 B_S^{4/7} \dot{M}_1^{-2/7} P_s^{-2/3} < 1 \quad (1.22)$$

(with the pulsar spin period  $P_s$  measured in seconds).

Therefore by regulating  $\dot{M}_1$ , accretion disks regulate the angular momentum transport onto the neutron star. X-ray pulsars with persistent accretion disks thus show persistent trends in the spin period; SMC X-1, the prototypical disk-fed X-ray pulsar, has been in a state of near-continuous spin-up for as long as it has been observed (Chapter 5). In the disk-fed LMXB pulsar GX1+4, the pulsar spins *down* at a rate correlated with the X-ray luminosity, suggesting the accretion disk may rotate in the opposite sense to the pulsar (Chakrabarty et al., 1997). Despite the likely high variation in instantaneous angular momentum transfer onto the neutron star, the X-ray pulsars were until very recently thought to be free of the “glitches” (instantaneous changes in spin-frequency) which are often seen in the radio pulsars (Lorimer, 2001). However it now appears that the Be-XRB KS 1947+30 has undergone a step-change in frequency without the increase in accretion luminosity expected from an episodic change in  $\dot{M}_1$ . This object appears to be the X-ray analogue of the glitches seen in radio pulsars (Galloway et al., 2003).

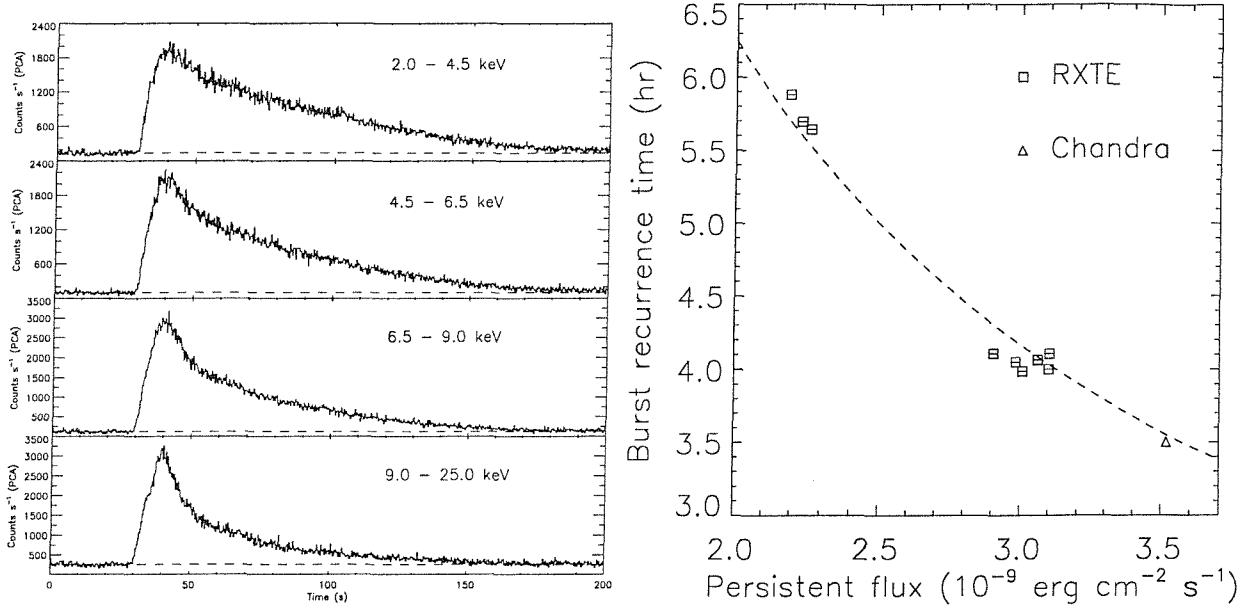
The millisecond *radio* pulsars occur almost exclusively in binaries (80% of the observed sample compared to  $\lesssim 1\%$  for “normal” pulsars; Lorimer 2001), suggesting their existence as a separate

population from the isolated radio pulsars (which generally show monotonic spin-down). These systems are thought to be the result of spin-up in an accreting LMXB, consistent also with lower ( $B_S \lesssim 10^8$  G) magnetic fields (Frank et al., 1992; Verbunt, 1993). This scenario was confirmed by the RXTE/PCA discovery of millisecond X-ray pulsations from the LMXB SAXJ1808.4-3658 (Wijnands & van der Klis, 1998). Only four millisecond X-ray pulsars are currently known, three of which have  $P_{orb} \sim 40$ -44 minutes and one (SAXJ1808.4-3658) with  $P_{orb}$  2.01 hours. All such systems are transients, with peak brightness only  $\sim 0.01L_{Edd}$ . Millisecond X-ray pulsars are the only LMXBs to show X-ray pulsations.

### 1.3.3 Type I X-ray Bursts

LMXBs are often observed to show dramatic, rapid increases in X-ray luminosity, with a rise to maximum luminosity  $\sim 10^{38}$  ergs $^{-1}$  in  $\sim 1$ -10s followed by exponential decay to the quiescent level with a timescale of  $\sim 90$ s (Lewin et al., 1995). Such *Type I X-ray bursts* are not seen from the HMXB pulsars, but are seen in the LMXB millisecond pulsars. They are thought to be due to explosively unstable nuclear processing of accreted material built up on the surface of the compact object. The X-ray spectrum of Type I bursts is that of an isothermal blackbody, allowing the evolution of the blackbody temperature (up to  $\sim 3$ keV) and radius (typically  $\sim 20\%$ ) to be charted.

The observed temperatures and the requirement of a surface on which to build up the burning layer suggests the compact object must be a neutron star. All bursters have magnetic fields  $\lesssim 10^9$  G, as at higher fields, explosive fronts cannot propagate. Observations of EXO 0748-673 showed anticorrelation between the persistent luminosity and the average interval between bursts (Gottwald et al., 1986), as well as an inverse relation between the peak flux and the persistent luminosity and a strong dependence of the blackbody radius on the peak flux. More recently, RXTE and Chandra studies of the regular “clockwork burster” X1826-268 show the burst recurrence time  $t_{recur} \propto \dot{M}_1^{-1}$ ; this suggests, remarkably, a constant mass is accumulated between bursts (see Galloway et al. 2004 and figure 1.4). In the case of the luminous bursters 1820-303 and GX3+1, bursting ceases altogether once the persistent luminosity increased above a critical level (Lewin et al., 1995). These behaviours are common amongst Type-I bursters; see Strohmayer & Bildsten (2003) for a recent review. In the case of the luminous bursters 1820-303 and GX3+1, bursting ceases altogether once the persistent luminosity increased above a critical level (Lewin et al., 1995). During bursting, localisation of nuclear burning produces variability on millisecond

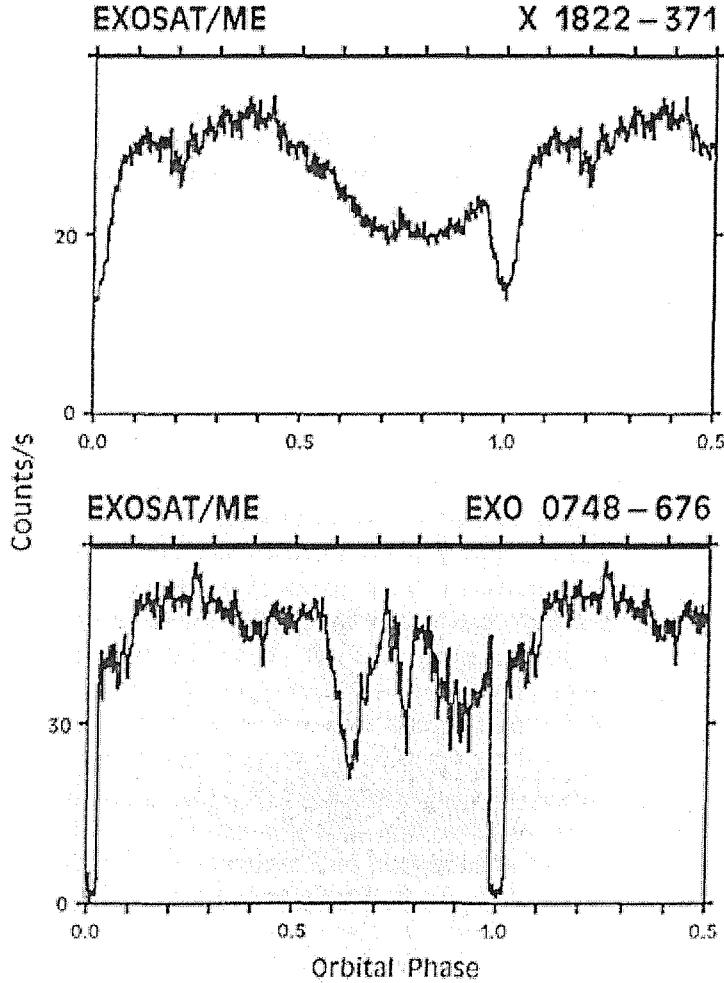


**Figure 1.4:** Type I X-ray bursts from the “clockwork burster” X1826-268. **Left:** RXTE/PCA observation of a typical Type-I X-ray burst. **Right:** burst recurrence time plotted against persistent flux. The dashed line is the trend expected for constant accumulated mass (so  $t_{recur} \propto \dot{M}_1^{-1}$ ). From Strohmayer & Bildsten (2003).

timescales, which has recently been shown in at least one case (SAX J1808.4-3658) to correspond to the neutron star spin frequency (see Strohmayer & Bildsten 2003 for discussion and references).

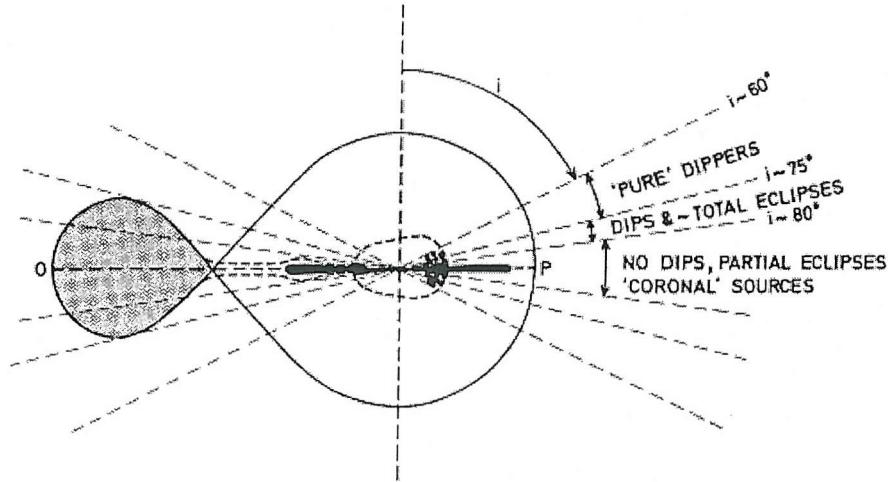
### 1.3.4 Inclination Dependent Dipping Behaviour and the ADC

In many systems from both HMXB and LMXB class, the X-ray flux (and also flux in other wavebands) is periodically cut off at or near the orbital period. For inclinations  $i^\circ \gtrsim 60^\circ$ , elevated structure in the disk rim occults the compact object, leading to the removal of its flux across all wavebands. In the case of the HMXBs, inclinations  $\gtrsim 50^\circ$  give rise to sharp eclipses of the compact object by the giant companion, allowing the radial extent and therefore mass of the donor to be estimated and providing a reliable system clock. However, eclipses in LMXBs are more rare; indeed, by 1970, X-ray lightcurves of some 25 sources had been collected, with *none* showing full eclipses (Charles & Seward, 1995), suggesting some mechanism was preventing the observation of eclipses by the donor. No such selection effect is observed in the HMXBs, which as a class have much smaller accretion disks (if even present) and more extended donors. The natural interpretation is that some feature of the accretion disk, probably outer rim structure, extends so far out of the plane that, from the point of view of the point source, it subtends a greater solid angle than the donor (Frank et al., 1987). This indeed appears to be the case; the X-ray lightcurve



**Figure 1.5:** EXOSAT-ME lightcurves of the prototypical ADC source X 1822-371 and the prototypical dipper EXO 0748-676. Notice that for X1822-371 the eclipse is partial, indicating a scattered component. Reproduced from Charles & Seward (1995).

of EX0748, the first eclipsing LMXB discovered, shows eclipses lasting just eight minutes (figure 1.5), corresponding to finely-tuned inclination that allows the line of sight to pass through high donor latitudes. Thus for LMXBs the nature of dipping behaviour is highly inclination dependent; at moderate inclinations ( $i^\circ \gtrsim 60^\circ$ ) dipping is due to structure in the mid-outer disk, while for a narrow range of inclinations ( $75^\circ \lesssim i^\circ \lesssim 80^\circ$ ) the donor may also eclipse the compact object (figure 1.6). At higher inclinations ( $\gtrsim 80^\circ$ ), the outer accretion disk obscures the compact object for most (if not all) of the orbit. Variation in the height of this disk rim structure then produces eclipses on or near the orbital period, but these eclipses are *partial*, suggesting an extended region is producing the observed emission: the Accretion Disk Corona (ADC). Such observations of the prototypical ADC source X1822-371 provided the first evidence for the ADC; the component *not* removed by the accretion disk rim must be the spatially extended corona (figure 1.6). In the ADC



**Figure 1.6:** The standard model of inclination dependence. At low inclination no dips are visible; at high inclination the disk rim shields the central source from direct sight, producing a low  $L_X/L_{Opt}$  and ADC behaviour. At intermediate inclinations the disk rim leads to dips and/or eclipses. Reproduced from Frank, King & Lasota (1987).

sources, the inner disk and neutron star are never directly visible, as suggested by the absence of Type I burst behaviour from any ADC source. (The luminous X-ray binary in M15, once thought to be a single LMXB that exhibits both ADC behaviour and bursts, has recently been resolved into two separate sources; see Chapter 4). ADC sources are thus underluminous in the X-ray by a factor  $\sim 25$ , with X-ray to optical ratio  $L_X/L_{Opt} \sim 15$  compared to the usual  $\sim 200$  (White et al., 1981). For other systems, the ADC is still an integral part of the X-ray emission, as suggested by spectral fits, to which we turn next.

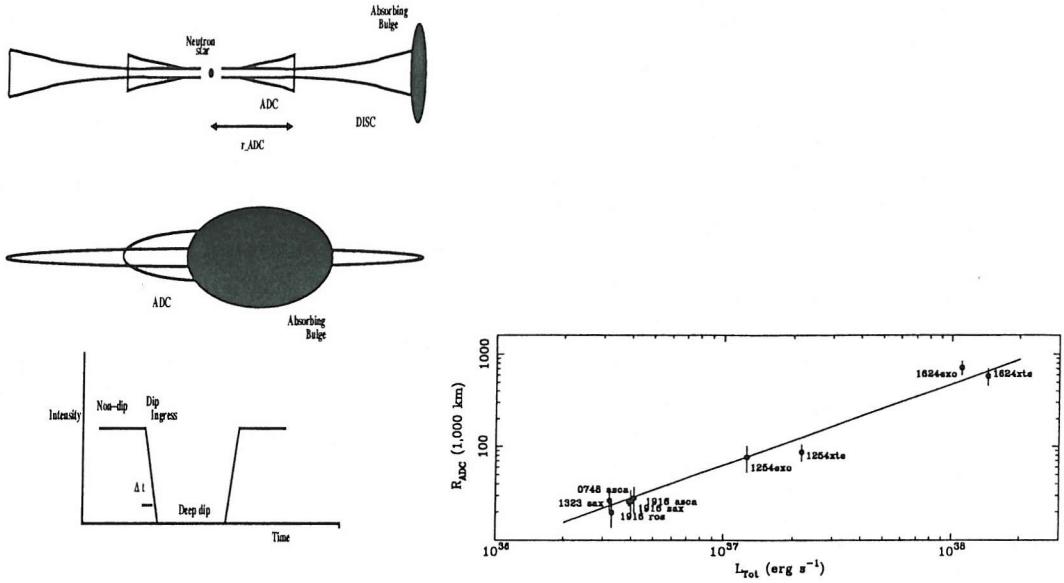
### 1.3.5 Observed X-ray spectra of XRBs

In each class of system there are several emission regions which contribute to the total radiation recorded in the X-ray lightcurves discussed in this thesis. We thus introduce the emission regions here to aid understanding of the lightcurves. Spectra of the LMXBs are best fit with two component models, the interpretation of which has not been unique. This uncertainty has arisen largely because the  $\sim 1 - 20$  keV bandpass of most X-ray missions does not allow cast-iron discrimination between mechanisms (Christian & Swank, 1997). Interstellar absorption removes significant flux below  $\sim 0.2$  keV, while the equivalent area of X-ray detectors usually drops off dramatically above  $\sim 10$  keV. A survey of four LMXBs with Tenma (1-37 keV) resulted in a two-blackbody model, with a stable soft ( $\lesssim 1.5$  keV) component comprised of the multi-temperature disk blackbody ( $T \propto r^{3/4}$ )

mentioned above, plus a time-varying isothermal blackbody ( $\gtrsim 2\text{keV}$ ), taken to represent the deposition of energy at the surface of the neutron star (Mitsuda et al., 1984). The “Eastern Model” (Mitsuda et al., 1989) was a generalisation of this to replace the neutron star blackbody component with a compton component at the inner disk, with the seed photons taken from the blackbody emission at the surface of the neutron star. The “Western Model” (White et al., 1988), based initially on EXOSAT observations in the 2-25 keV bandpass, replaces the disk blackbody with a single comptonised component, with the spectrum of seed photons unconstrained; a blackbody component is required only for the highest luminosity sources, when it is identified with the inner disk/neutron star surface boundary layer. Beyond this the geometry of the emission regions is largely unconstrained.

More recently, observations of the dipping LMXBs have allowed the geometric extent of the Comptonised and blackbody regions to be uncovered. During an X-ray dip, flux is removed across all energies as the accretion disk rim blocks the inner disk region from view (section 1.3.4). Comparison of dip ingress lightcurves with X-ray spectra shows the blackbody component to be rapidly blocked, while the more gradual removal of the comptonised component suggests it must originate in a region  $\gtrsim 50,000\text{km}$  (or  $\sim 15\%$  the accretion disk radius) in extent (Church et al., 1997). This ADC is not spherical but instead has a small height-to-radius ratio (Smale et al., 2001). This model was shown to fit a sample of ten LMXBs including both atoll- and Z-sources observed with ASCA (Church & Balucińska-Church, 2001), and has since been adopted by many authors. In both this two-component model and the Western model, no *direct* X-ray emission from the accretion disk itself is observed, instead the hot thermal plasma of the disk provides the seed photons for the observed comptonised component (Barret, 2001). This model appears to fit the Atolls as well as the Z-sources (Church et al 1997). We caution the reader that, while persuasive, this model is not yet accepted by all authors; some still prefer to use the generalised Western Model when fitting the Atolls (e.g. Gierliński & Done, 2002).

The X-ray spectra of HMXBs are in general phenomenologically similar to that of the LMXBs, in that a disk and accompanying corona are present as a power-law; however for the X-ray pulsars there is of course a further component due to the channelled-accretion process taking place at the surface of the pulsar, which can dominate the observed spectrum. A combination of the beaming properties of the localised emission poles, with the rotation of the neutron star, gives rise to the pulsed emission observed. While the geometry is then perhaps straightforward (section 1.3.2), the balance of mechanisms producing the continuum emission are still something of a mystery;



**Figure 1.7:** Schematic of emission components as suggested by X-ray spectra of the dipping sources. **Left:** system geometry. The ADC is confined to a thin region above and below the accretion disk, producing a power-law component in the observed spectrum. As vertically extended disk matter obscures the inner disk, the power-law component drops rapidly in strength. **Right:** size of the ADC. By measuring the time taken for the power law component to disappear, the radius of the ADC itself is estimated. From Church & Bałucińska-Church (2004).

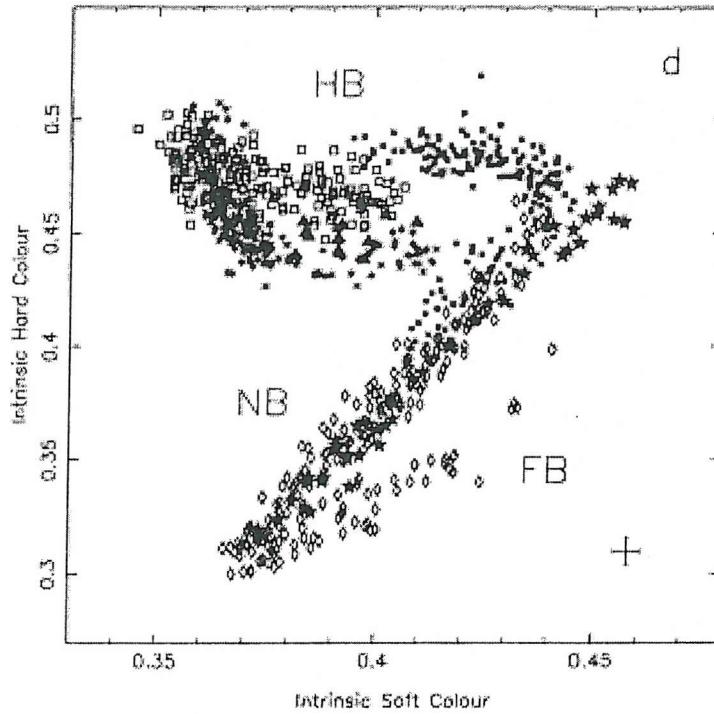
indeed there is currently *no* convincing model for the continuum shape of the X-ray spectra of the accreting pulsars (Coburn et al., 2002) and refs therein. The magnetic field of the pulsar gives rise to cyclotron emission features, the energies of which (fundamental and harmonics) are predicted from the magnetic field of the accretor (equation 1 of Coburn et al. 2002), which are in some cases observed (Heindl et al., 2001).

### 1.3.6 Variability on the Sub-Day timescale

The LMXBs have traditionally been classified into two categories, based on their behaviour in the X-ray colour-colour diagram. Sources trace out locii of different states as the central  $\dot{M}_1$  changes; the typically high luminosity ( $\gtrsim 10^{37} \text{ erg s}^{-1}$ ) sources such as Sco X-1 and Cyg X-2 trace out a Z-shaped path (figure 1.8), and were thus dubbed the *Z-sources* (Hasinger & van der Klis, 1989). Sources at low luminosities tended to cluster around an island-shape region in the colour-colour diagram, and were dubbed *Atolls* (figure 1.9). Studies of the few atolls showing large excursions in  $L_X$  has revealed a further curved branch of the atoll colour-colour behaviour, dubbed the “Banana” state. The advent of RXTE/PCA, with its comparatively small spatial errorboxes ( $\sim 1^\circ$ ) and high

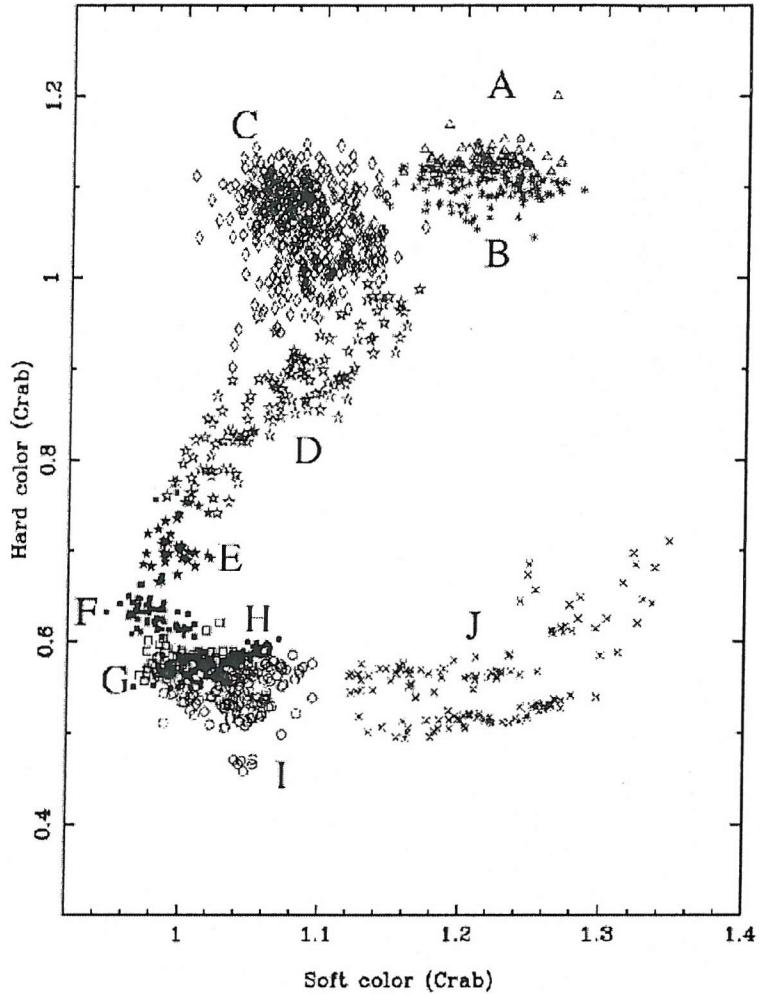
time resolution (down to  $3\mu\text{s}$ ) has allowed the X-ray variability of the XRB to be studied over the 2ms - 250s range in periods. At low accretion rates, the power spectra of the Z-source and Atoll neutron star XRB, and the black hole candidates, show very similar morphology (Wijnands & van der Klis, 1999a). At frequencies  $\nu \gtrsim 1\text{Hz}$  the underlying morphology consists of a power law component with  $P_\nu \propto \nu^{-1}$ , which evolves into a flat noise-type component below some break frequency  $\nu_b$ . In some atolls the power law also breaks into flat noise above  $\sim 100\text{Hz}$ . Above the broad morphology a broad bump is usually present at high amplitude ( $\sim 50\%$  rms), sometimes called a Quasi-Periodic Oscillation (QPO), typically at frequencies ( $1 \lesssim \nu_{QPO} \lesssim 50$ )Hz. Also present in both Atolls and Z-sources are much more rapid *Kilohertz* QPOs.

The physical mechanisms behind the power-law broadband variability, the  $\sim 10\text{Hz}$  QPOs and the kHz QPOs are probably linked, as  $\nu_{QPO}$  is tightly correlated with both the kilohertz QPOs (Psaltis et al., 1999) and the break frequency  $\nu_b$  (Wijnands & van der Klis 1999). The  $\nu_b - \nu_{QPO}$  correlation is not associated with the surface of the neutron star or its magnetosphere, as it persists in the black-hole candidates. Furthermore, both  $\nu_b$  and  $\nu_{QPO}$  are correlated with the instantaneous  $\dot{M}_1$  (van der Klis, 1994); usually this relationship is one-to-one, but not in the case of the millisecond X-ray pulsar SAXJ1808 (Wijnands & van der Klis, 1999b). As these QPOs are observed from systems at a wide range of inclinations, their identification with obscuration in a similar manner to the X-ray dips is ruled out. The large rms amplitude suggests the QPOs originate in the innermost  $\sim$ few  $R_1$  from the compact object, as this is where the majority of dissipation occurs (section 1.3.1). The correlation of  $\nu_{QPO}$  with  $\dot{M}_1$  suggests the QPOs probably originate through modulation of the mass flow close to the neutron star itself. The Z-source QPOs (in their low  $\dot{M}_1$  state) show a slightly different relationship between  $\nu_{QPO}$  and  $\nu_b$ , with the same slope in  $(\log \nu_b - \log \nu_{QPO})$  space (Wijnands & van der Klis 1999). The correlation between  $\nu_{QPO}$  and  $\nu_{kHz}$  is the same in the Z-sources as in the Atolls, suggesting  $\nu_b$  may correspond to different timescales in the two systems. The Z-sources typically show a larger variation in  $\dot{M}_1$  than the Atolls (though there are exceptions, e.g. the Atoll source 4U1608-52, which shows a factor  $\sim 100$  variation in  $L_X$ ; e.g. Mitsuda et al 1989). “Low” accretion rates ( $\lesssim 0.5L_{Edd}$ ) correspond to the *horizontal branch* (HB) QPOs, with typical frequencies  $\lesssim 0.5\text{Hz}$ . As  $\dot{M}_1$  increases, the source moves from the HB to the *Normal* (NB) and *Flaring* (FB) branches. In these branches  $\nu_{QPO}$  clusters near  $\sim 5\text{Hz}$ . The HB and FB QPOs are thought to represent variations in  $\dot{M}_1$ ; feedback between radiation pressure through the high  $L_X \gtrsim 0.5L_{Edd}$  and  $\dot{M}_1$  is thought to support variations in optical depth on this timescale (van der Klis 1995). The HB QPOs have traditionally been explained with the *beat frequency* model (Alpar et al., 1992), in which the spin of the misaligned magnetosphere leads to quasi-periodic



**Figure 1.8:** Sample Z-track from the IMXB Cyg X-2, based on EXOSAT observations.. From (Kuulkers et al., 1996).  $\dot{M}_1$  increases from the Horizontal Branch (HB) through to the Flaring Branch (FB).

variations in the magnetically induced pressure on material orbiting with frequency  $\nu_{Kep}$ . The frequency for such QPOs would then be the beat frequency  $\nu_{QPO} = \nu_{Kep} - \nu_{sp}$ . This model has been popular because it predicts results of the right order of magnitude (e.g. van der Klis 1995) and indeed is supported by the observation that the Z-sources as a class have higher magnetic fields ( $\gtrsim 10^9 G$ ) than the Atolls ( $\lesssim 10^8 G$ ); however it has recently been thrown into doubt. For example, the correlation of  $\nu_{QPO}$  with  $\nu_b$  is similar to the atolls; thus if the HB QPOs are produced by the same mechanism, the beat frequency model cannot apply as no magnetosphere is required (Wijnands & van der Klis, 1999a). The Z-sources typically trace out the three branches of their colour-colour behaviour on a timescale of hours-days (e.g. GX340+0; (Jonker et al., 2000). By contrast, evolution within the “Island” state of the Atolls takes place on a timescale of days-weeks (e.g. 4U1728-34; Di Salvo et al. 2001), with the banana track traced out on a timescale of hours-days. Few atolls show large variations in  $L_X$ , suggesting the apparent difference in colour-colour behaviour might be a selection effect due to the largely different  $\dot{M}_1$  regimes. Indeed, colour-colour diagrams plotted from large RXTE/PCA datasets including the three Atoll transients 4U1608-52, 4U 1705-44 and Aql X-1, suggested the atolls also trace out three-branch patterns, implying the Z- and Atoll sources might simply be different  $\dot{M}_1$  regimes of the same behaviour (Muno et al.,



**Figure 1.9:** Colour-colour diagram of the Atoll LMXB 4U 1608-52, taken with RXTE/PCA. Soft colour is defined from RXTE/PCA channels as  $(3.5\text{-}6.0\text{ keV}) / (2.0\text{-}3.0\text{ keV})$ , Hard colour defined as  $(9.7\text{-}16.0\text{ keV}) / (6.0\text{-}9.7\text{ keV})$ . From van Straaten et al. (2003).

2002). However, recent analysis of the timing behaviour of 4U 1608-52, and its comparison with the Atolls 4U 0614+09 and 4U 1728-34, suggests the variability of the Atolls is still distinct from the Z-sources (van Straaten et al., 2003).

Power-law variability  $P_\nu \propto \nu^{-1}$  has also been seen in the X-ray active AGN, on timescales which, when scaled with the central black hole mass, are comparable to the power-law variability found in the X-ray binaries. Equivalent variability is also seen at optical and IR wavelengths, showing that longer timescale variation corresponds to lower temperatures and thus larger radii. The variation is then likely related to the propagation time for inwards disturbances, as then the  $\nu^{-1}$  dependence is expected (Vaughan et al., 2003). This mechanism is also plausible for the XRB, though verification through observations at other wavelengths will be extremely difficult as the timescales in which to collect sufficient photons are so much shorter. Another plausible mechanism is shot/flicker noise,

particularly given the presence of turbulent, highly energetic mass flow; such noise is also predicted to produce a  $\nu^{-1}$  power spectrum (e.g. Lyubarskii, 1997). Not all authors accept the shot-noise model; some groups contend these models allow fine-tuning of parameters to produce the  $\nu^{-1}$  dependence, and are largely phenomenological rather than physically based (King et al., 2004).

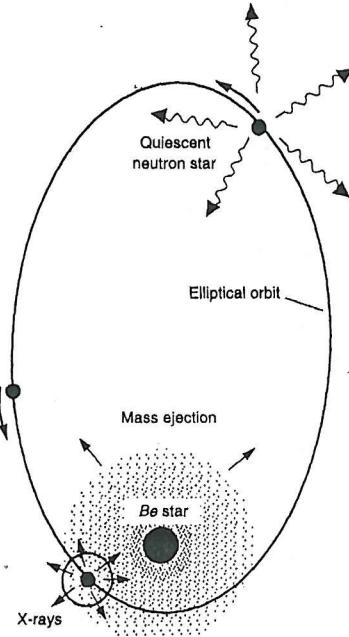
### 1.3.7 Orbital Variability in $\dot{M}_2$

The majority of LMXBs show little variability in the mass transfer rate from the donor, as expected from Roche lobe overflow from a circular orbit, especially with the accretion disk smoothing out changes on its viscous timescale. Indeed, the circular-orbit LMXB X1826-268 shows a remarkably stable recurrence rate in Type-I X-ray bursts, suggesting little change in  $\dot{M}_1$  (that this is a good diagnostic of the stability of  $\dot{M}_2$  is seen from the correlation between an increase in burst frequency with the orbital decay of the system (Galloway et al., 2004)).

The HMXBs, however, often show dramatic variation of  $L_X$  with orbital phase. This is naturally explained by the eccentric orbits (confirmed from timing studies of the high-inclination X-ray pulsars) as arising from variations in the local wind density and relative velocity throughout the orbit. In the most violent systems, periastron passage provokes tidal overflow from the donor (the time-dependent analogue to Roche overflow), resulting in an additional component to  $\dot{M}_2$  and thus a dramatic X-ray outburst accompanying periastron passage, with best-fit spectral models suggesting an accompanying change in radius. (e.g. A0538-66; Charles et al 1983; see also figure 1.10).

### 1.3.8 Interstellar Scattering and Extinction

X-ray emission is modified by the ISM in two main ways (we ignore polarisation in this work). Photoelectric absorption of K-shell electrons causes a reduction in flux with cross section  $\sigma(E) \simeq N_H E^{-2.6}$  (Ryter et al., 1975), producing a measurable low-energy cutoff in measured emission. Secondly, Rayleigh scattering off interstellar dust grains gives rise to an X-ray “halo,” as was first pointed out by (Overbeck, 1965) and substantiated by (Martin, 1970). For a given dust-grain distribution, the observed fractional halo intensity may be related to the column density  $N_H$



**Figure 1.10:** Schematic of a Be-X binary. The eccentric orbit of the neutron star takes it close enough to the donor to pass through overdense material flung out at the equator by the donor rotation (and observed in  $H\alpha$  line profiles). Enhanced mass transfer at periastron powers intense optical and X-ray outbursts, and a (generally) short-lived accretion disk can form after periastron passage. Reproduced from Charles & Seward (1995).

through the scattering optical depth (Mathis & Lee, 1991):

$$\frac{I_{halo}}{I_X + I_{halo}} = (1 - \exp[-KN_H E^{-2}]) \quad (1.23)$$

The first such halo was detected with Einstein in 1983 (Rolf, 1983). In total only six sources showed halos prominent enough for detection with Einstein (Catura, 1983), which were later used to establish consistency with common grain models (Mauche & Gorenstein, 1986). A full calculation shows that the scattering angle will be energy-dependent (Alcock & Hatchett, 1978):

$$\theta_{sc} \simeq 4'' \times \left( \frac{a}{0.25\mu m} \right)^{-1} \left( \frac{E}{1keV} \right)^{-1} \quad (1.24)$$

which leads to the (perhaps surprising) result that variability in the source may also be observed in any halo. The extra travel time  $\Delta t$  between scattered and unscattered paths may be given by (Bode et al., 1985):

$$\Delta t \simeq 0.22\text{days} \times \left( \frac{a}{0.25\mu m} \right)^{-2} \left( \frac{E}{1keV} \right)^{-2} \left( \frac{d}{2.5kpc} \right) \quad (1.25)$$

Thus an X-ray halo will add a component to X-ray lightcurves and spectra which varies in a complex way. The time delay between source variation and its corresponding halo variation will be larger at lower energies, as will the observed region in the halo to which this corresponds, with clear implications for spectral fitting of any observed halo.

The existence of the X-ray halo allows an independent check of the interstellar reddening law. The high spatial resolution of ROSAT allowed a systematic analysis of some 25 point sources showing halos, which were used to determine a tight correlation between the optical depth to scattering and  $N_H$  (Predehl & Schmitt, 1995), resulting in the relation

$$A_V = 3E_{B-V} = 5.6 \times 10^{-22} N_H \quad (1.26)$$

which is in agreement with the correlation between  $N_H$  and  $E_{B-V}$  derived from  $L_\alpha$  measurements of hot stars (Savage & Jenkins, 1972) and observations of soft X-ray cutoffs in supernova remnants (Gorenstein, 1975). In the absence of other information, then, the extinction  $A_V$  may thus be estimated from measured values of  $N_H$ , for example from radio maps (Dickey & Lockman, 1990).

## 1.4 The properties of XRB at other wavelengths

X-ray binaries are modern astronomical objects in the sense that their nature could not be uncovered without information across many energy regimes. For completeness we include here briefly the major effects outside the X-ray window, which are used (if interstellar absorption allows) to determine the physical parameters of the systems.

### 1.4.1 Ellipsoidal variation in optical brightness

The tidal deformation of the donor is in many cases measurable as a variation in optical brightness on the orbital period. For a spherical star the projected emitting area is constant, but when the star assumes the shape of its Roche Lobe, this area becomes a function of orbital phase. Fitting this ellipsoidal variation allows constraints to be placed on the inclination of the system (Kuiper et al., 1988).

### 1.4.2 Optical Reprocessing

The reprocessed emission from the extended accretion disk in the LMXBs usually dominates the optical light in the system. If the system is at high inclination, the orbit of the system can be traced from observations of line shifts in this disk, yielding the mass function. Furthermore, the accretion

disk rotation may be visible as a double-peaked feature in optical spectra, allowing constraints to be set on the disk temperature for assumptions about the location of the emitting region (Haswell & Shafter, 1990). The extent of the accretion disk can also be estimated from observations of the delay between occurrence of strong X-ray variation (e.g. a Type I burst) and its counterpart in optical lightcurves, in a small-scale analogue of the reverberation mapping common in the study of AGN (O'Brien et al., 2002; Hynes et al., 2003).

#### 1.4.3 X-ray heating of the donor

Optical spectra of LMXBs rarely show stellar absorption features from the donor due to the dominance of the reprocessed optical emission from the disk. However, in some cases the surface of the donor is strongly heated by X-ray emission from the neutron star and inner disk, producing Bowen fluorescence which allows the orbital motion to be traced as a spectroscopic and photometric variation (Casares et al., 2003). By contrast, in HMXBs the optical emission is dominated by the  $\gtrsim 5M_{\odot}$  donor. However, in the case of tidal overflow at periastron, the donor can still be affected by the accretion process; in A0538-66, for example, the spectral type of the donor varies in spectral class from a B9 giant at apastron to B9 supergiant at periastron (Charles et al., 1983).

#### 1.4.4 Axial outflow from the system

In HMXBs and LMXBs with high  $\dot{M}_1$ , outflow can be launched along a bipolar trajectory. This outflow is usually observed by radio telescopes as spatially resolved synchotron emission (Fender, 2001), but may be observed in IR, optical and X-ray spectra as a pair of narrow emission lines (e.g. the prototype SS 433; Gies et al. 2002). The broadband spectrum of sources in the “low-hard” (i.e. jet-producing) state differs for neutron star and black hole binaries; black hole systems are totally dominated by synchotron emission from the jet, while for neutron star systems the jet is much less important energetically (Fender et al., 2003). As the authors demonstrate, this removes the need for significant energy advection across the event horizon to explain the factor  $\sim 100$  discrepancy in quiescent luminosity between black hole and neutron star systems in quiescence. This discrepancy is thus not an indicator of the event horizon, as had been previously suggested (Narayan et al., 1997)

## 1.5 Summary

The X-ray Binaries clearly represent the frontier of modern astrophysics in several respects. Matter undergoing accretion in these systems may encounter extremes of magnetic and gravitational fields, temperature and ionisation states. Furthermore, understanding these systems requires the development of new techniques in astronomy, in sensitivity, time resolution and spectral resolution, to provide the necessary information to follow the energetics of such systems fully.

## Chapter 2

# Superorbital Periodicities in X-ray Binaries

The phenomena outlined in Chapter 1 are used to establish the nature of the system components in X-ray binaries. The last two decades have extended the coverage of X-ray datasets to uncover new behaviour on timescales from a few to thousands of orbital periods. The study of these variations forms a large part of the work in this thesis, so we introduce them here.

### 2.1 Observations of Superorbital Periodicities

In about 20 X-ray binaries, investigations at X-ray and other wavelengths have shown variations on timescales of tens - hundreds of days, the prototypes for which are the well-studied XRBs Her X-1, LMC X-4 and SS433. These variations divide broadly into two classes, the first being those for which variation is periodic and stable, such as Her X-1 and LMC X-4, (which show periodicities of  $\sim 35$ d and  $\sim 30$ d respectively). The second class consists of sources for which modulation is “quasi-periodic,” in that even long-term monitoring datasets do not yield precise periods in their power spectra, but a broad peak, often superposed on “red-noise” power. Examples of these include Cyg X-2, with  $P \sim 70$ - $80$ d (Paul et al., 2000); GX354-0 with  $P \sim 70$ d (Kong et al., 1998); X1916-053 with  $P \sim 83$ d (Homer et al., 2001) and X1820-30 with  $P \sim 200$ d (Chou & Grindlay, 2001). The longest-standing superorbital periodicities have usually been claimed in sources which are

bright enough to have been detected by the early survey instruments Ariel 5 and Vela-5 (Smale & Lochner, 1992); however recently the unprecedented long X-ray lightcurves afforded by the RXTE All-Sky Monitor (hereafter RXTE/ASM) and the CGRO Burst And Transient Source Experiment (hereafter CGRO/BATSE) have led to claims of superorbital periodicities in sources that have been comparatively poorly studied.

### 2.1.1 A survey of superorbital periodicities

The number of sources for which a superorbital periodicity has been detected is in a state of rapid growth: eleven years ago a review paper (Schwarzenberg-Czerny, 1992) listed 10 such sources across all categories of compact object: as a result of the long-timebase X-ray datasets, well over thirty such sources are now documented between the black hole candidates and neutron star systems alone, with claimed superorbital periodicities ranging up to 1400 days.

### 2.1.2 Disk Precession and/or warping

Her X-1, LMC X-4 and SMC X-1 are prototypical disk-warping systems, in which a substantial body of independent evidence now exists for a precessing, warped accretion disk, for example optical variations on the beat between the orbital and precessional cycles or cycle-dependent variations in column density  $nH$  suggestive of periodic obscuration (see Chapters 5 & 6 for more on these sources).

Cyg X-2 shows significant ( $\sim 50\%$ ) X-ray brightness variation on a timescale  $\sim 78$  day, observed in longterm X-ray datasets since 1975 (Wijnands et al 1996). The burst activity and QPO-state are *not* correlated with the X-ray cycle (Wijnands et al., 1997; Kuulkers et al., 1996), suggesting this variation does not represent variation in  $\dot{M}_1$ . Periodic obscuration by a precessing disk has instead been suggested (Wijnands et al 1996); indeed recent HST-STIS spectra of Cyg X-2 can be well-fit with a warped-disk geometry (Vrtilek et al., 2003). This source is investigated further in Chapter 6.

SS 433 is one of the canonical superorbital systems, with extensive relativistic ( $\beta \simeq 0.26$ ) jets, which precess on a 162-day cycle (see Margon, 1984, for a review). Consideration of the angular momentum balance of the system components suggests that jet precession must be driven by

precession of at least the inner disk, and indeed recent analysis of long-term X-ray lightcurves may show direct evidence for such disk precession taking place (Clarkson et al 2004 in prep).

The galactic bulge source GX354-0 (4U1728-34) is almost certainly a LMXB as it shows type-I X-ray bursts (Basinska et al., 1984) and 363 Hz QPOs (Strohmayer et al., 1996). However, the high extinction and unknown distance have prevented the direct identification of a donor, and no signatures of orbital motion have been found. Analysis of the 6-year Ariel-5 dataset and the first 2 years of the RXTE/ASM lightcurve shows evidence for a quasi-stable variation at  $\sim 63$  and  $\sim 72$  days respectively. The apparent instability of this period suggest it is not orbital, but instead is ascribed to possible disk precession (Kong et al., 1998).

X-ray lightcurves of the Be-binary A0535+26 show outbursts on a period 111.4 days (Motch et al., 1991), but analysis of optical lightcurves over an interval of 15 years (Larionov et al., 2001) show a distinct optical period  $\simeq 103$  days. The authors interpret this difference as a product of precession of accretion disk or a Be-equatorial disk, on the beat period between the optical and X-ray periods of  $\sim 1400$  days.

KS1731-260 is an LMXB with no identified donor, showing bursts (Smith et al., 1997) and  $\sim 524$  Hz QPOs (Muno et al., 2000). Analysis of RXTE/ASM data during a state of high activity suggests a 38-day periodicity, which is not detectable during the subsequent low state (Revnivtsev & Sunyaev, 2003). This periodicity cannot be orbital in nature, as a Roche-lobe filling donor would then have a surface temperature  $\lesssim 1600$  K to reproduce the observed apparent magnitude  $m_J \gtrsim 17$  (Barret et al., 1998). The variation is instead ascribed to accretion disk precession.

The LMXB 4U0614+09 has a faint, blue ( $\sim 18$  mag) optical counterpart, suggesting  $M_2 \lesssim 1M_\odot$  and thus an orbital period  $\lesssim 10$  hours (Machin et al., 1990). However, Ariel-5 lightcurves show strong modulation on a  $\sim 5$ -day timescale (Marshall & Millit, 1981), while optical observations shows a possible 10-day periodicity (Machin et al 1990). Anticorrelation between optical and X-ray flux suggest disk reprocessing is the dominant source of optical radiation, leading the authors to suggest disk precession as the mechanism for the possible long-term variation.

LMC X-2 is a bright LMXB ( $L_X \sim 0.3 - 6 \times 10^{38}$  erg s $^{-1}$ ), for which a faint ( $m_V \sim 18.8$ ) counterpart has been identified (Bonnet-Bidaud et al., 1989) but the orbital period is still unknown. Optical photometry performed by two groups suggested divergent periods at 8.15 hours (Callanan et al., 1990) and 12.54 days (Crampton et al., 1990); to attempt to resolve this dispute, the 6-

year MACHO lightcurve of LMC X-2 was searched to place limits on either periodicity (Alcock et al., 2000). While the 12.54-d periodicity was not detected at the previously claimed level, the search was insensitive to the 8.15-hour modulation. The authors of this study claim variations on a timescale  $\gtrsim 10$  days, which they identify with possible unstable disk precession.

### 2.1.3 Longterm Changes in $\dot{M}_1$

The BHC LMC X-3 was shown to exhibit optical and X-ray variability on a  $\sim 100$ -200d timescale, which at first was identified with a possible precessing accretion disk (Cowley et al., 1991). A recent analysis of 6 years of optical lightcurves show that there is no stable period, but instead a timescale of variation (Brocksopp et al., 2001). The X-ray spectral behaviour of LMC X-3 was monitored over several cycles in this variation; fitting a two-component power law + disk blackbody model (as appropriate for black hole candidates; see e.g. McClintock & Remillard 2003) suggests the inner disk temperature scales with the X-ray flux, consistent with recurring transitions between the canonical low/hard and high/soft X-ray states of black hole candidates (Wilms et al 2001). The authors suggest a precessing accretion disk is not responsible for the variation as fitted column density variations along the cycle are far smaller than for the canonical precessing disk system Her X-1 (Kuster et al., 1999).

For the LMXB 4U1957+11, periodograms of the first two years of RXTE/ASM long-term X-ray lightcurve (Nowak & Wilms, 1999) and changes in the optical lightcurve (Hakala et al., 1999) were interpreted as requiring a warped accretion disk, which precesses with a  $\sim 117$ -d period. However, more recent analysis of the full X-ray lightcurve, combined with short-term pointings (Wijnands et al., 2002, hereafter WMK2002), show instead that the long-term variation is unstable with time, with detected periods ranging from 100-260 days. Spectra taken from the short-term pointings indicate no evidence for varying absorption from intervening matter (WMK2002), suggesting the variation is instead due to changing  $\dot{M}_1$ .

Analysis of the Vela-5B lightcurve of the soft X-ray transient 4U1608-52 shows evidence (once the large outbursts are subtracted) of two periodicities, at 4.10 or 5.29 days (Lochner & Roussel-Dupre, 1994). The orbital period of this system has never been constrained, due to the faintness of the companion (Wachter, 1997), but the system is known to be a neutron star from the discovery of Type-I X-ray bursts (Tananbaum et al., 1976). The 4.10-day period has been conventionally interpreted as the orbital period of the system, with the longer period thought to be due to time-

varying occultation by a precessing accretion disk (Lochner & Dupre 1994). This identification was based upon analogy with the canonical superhumper SU Uma (Haswell et al., 2001), in which here the accretion rate varies on the beat between the orbital and disk precessional period, which would imply a disk precessional period of  $\sim 18$  days. The superhump scenario is unlikely for 4U1608-52, as the resonance condition on superhumps restricts them to systems in which  $q \lesssim 0.3$  (Haswell et al 1999). The period-mass relation for LMXBs then implies the orbital period should be  $\lesssim 30$  hours for a neutron star accretor, which is inconsistent with the 4.10 days assumed by the superhump scenario itself. A more likely scenario is that the longterm variation found in RXTE/ASM corresponds to the tracing of the Atoll track of 4U1608-52, and thus corresponds to fluctuations in the mass transfer rate through the accretion disk (see e.g. van Straaten et al. 2003).

Two of the Ultracompact neutron star X-ray binaries have long shown evidence for superorbital variation. X1820-303 shows a spectacular  $\sim 170$ -d cycle, which is almost certainly due to variation in mass transfer rate; the burst recurrence rate scales inversely with the count rate, and indeed bursting ceases altogether once the count rate rises above a certain level (Chou & Grindlay, 2001). Until recently, the favoured interpretation for this system involved a third body in a wide orbit, which would drive precession of the eccentricity of the inner orbit on the 170-d timescale, bringing about large variations in mass transfer (c.f. Mazeh & Shaham, 1979); see Chapter 5 for more discussion on this mechanism). Our exploratory analysis with dynamic power spectra (see chapters 3,5,6 & 8 for examples) suggested that this may not be the case due to the instability in the “precessional” period (Clarkson 2003); coincident work suggested that in fact this variation is most likely due to disk instabilities (Šimon, 2003).

#### 2.1.4 Variations with Uncertain or Low Significance

The other ultracompact suspected of superorbital variation, 4U 1916-053, has been something of a cornerstone of precessing disk work, as the suspected 199-day precession period (Smale & Lochner, 1992) is consistent with *prograde* precession, thus establishing such behaviour as disk-related. However, recent studies with higher quality datasets (Clarkson 1999, Homer et al 2001), have shown that no such periodicity persists. The strongest remaining candidate variation, a periodicity at  $\sim 80$ days, is of low ( $\sim 2\sigma$ ) significance, as in this source the combination of behaviours at many frequencies has produced a continuum that is frequency dependent (chapter 3), and it is over this continuum that the significance of power spectrum peaks must be evaluated for this source.

Examination of the Vela-5B lightcurve of 4U1907+09 by Priedhorsky & Terrell (Priedhorsky & Terrell, 1984) with fourier transform methods suggests a periodicity at  $\sim 41.6$  days, which however has not been reported by any group since, and indeed does not appear to be present in the RXTE/ASM lightcurve.

Recent analysis of the RXTE/ASM lightcurve of GX13+1 has led to a claim of a periodicity at  $24.065 \pm 0.018$  days (Corbet, 2003). The orbital period has not been unambiguously identified, though optical photometry yields maximum Fourier power at 12.6d (Groot et al., 1996), but the donor has been directly identified as a K5 III star (Bandyopadhyay et al., 1999), suggesting a  $\sim 5M_\odot$  donor. By computing the number of ASM observations made per day of GX13+1, a periodicity was found in the sampling of 52.6 days. Corbet (2003) presents power spectra of GX13+1 with data prepared in five different ways, with only the weighted and unweighted power spectra of the daily average datapoints showing the 24.065-day periodicity to any significance. The author does not give significance levels, and has normalised the power spectra by the flux, which makes evaluation of the significance from the values of the peaks nontrivial. Furthermore, the claimed periodicity is close to twice the orbital period, and half the sampling period, so it is unclear at present how seriously this detection should be taken.

Analysis of a campaign of RXTE/PCA lightcurves of the galactic Black Hole Candidates 1E1740.7-2942 and GRS1758-258 suggests orbital periods 12.73 and 18.45 days (Smith et al., 2002a). Using the BATSE lightcurves and extrapolating the spectra out to the 50keV range, the authors claim superorbital periodicities of  $\sim 600$  days in both sources. For 1E1740.7-2942 the periodicity is visible in the lightcurve itself, and is interpreted by the authors as a possible recurring, transient disk warp. The GRS1758-258 periodicity is much less robust, as the periodicity appears indistinguishable from frequency-dependent red-noise (Chapter 3).

4U 1907+09 is a pulsar XRB with massive companion (van Kerkwijk et al., 1989) in a wide orbit of moderate eccentricity ( $e \sim 0.16 \pm 0.14$ ; Makishima et al., 1984), that has shown a monotonic *increase* in pulse period since its discovery (Mukerjee et al., 2001). This source is considered in more detail in the introduction to Chapter 6; it is included in this section because two theoretical review papers of precessing, warped accretion disks include this source because of a suggested 42-day periodicity. However, this periodicity is not to be found in any of the literature on this object, so it is unclear where the assertion of the long-term periodicity originates.

Analyses of longterm X-ray lightcurves of the X-ray binary in M15, i.e. X2127+119, have in the

past given rise to suggestions of a 37-day period in this system (Corbet et al., 1997). However, it has recently been shown to consist of two sources which cannot be resolved by RXTE (White & Angelini, 2001). While resolving the mystery of how a source can show evidence for both an ADC and Type I bursts, it is not possible to relate components in the longterm X-ray lightcurves exclusively to either of these systems. This source is examined in more detail in Chapter 4.

### 2.1.5 Undetermined Mechanism

Cen X-3 shows a significant variation on a  $\sim 180$ - day timescale, as evidenced in the long-term lightcurves of Ginga, BATSE, RXTE/ASM, Ariel-V and Vela-5B. As noted by Bildsten et al. (1997), this variation is not correlated in any way with the spin period of (and therefore accretion torque acting on) the neutron star, which experiences intervals of spin-up and spin-down with no clear pattern.

The Be Star LSI  $+61^\circ 303$  shows variation in X-ray, optical and radio emission on the 26.5-day orbital period (Taylor & Gregory, 1982), thought to be due to enhanced mass transfer at periastron passage ( $e \sim 0.6$ ; Hutchings & Crampton, 1981). Also seen in X-ray, optical and radio is a 1584-day variation in the brightness at peak outburst (Gregory et al., 1989; Zamanov et al., 1999; Apparao, 2001). Two mechanisms have been proposed; the first, precession of jets (and therefore accretion disk), is perhaps uncertain because the variation is not stable enough (compared to the known jet-precesser SS 433). The preferred model involves density enhancements in the outflow from the Be-donor, which produces the observed phase delay between  $H\alpha$  and radio peaks through the interval taken for such an enhancement to dissipate and reduce its optical thickness (Apparao 2001).

### 2.1.6 Superorbital Variations from External Galaxies

In addition to the sources listed in the table, observations of external galaxies are beginning to yield X-ray lightcurves of individual sources. Seven of the nine Ultraluminous X-ray Sources (ULX) in the Antennae galaxies (NGC 4038/9) are seen to show variability in  $L_X$  by factors  $\sim 3 - 10$  on a timescale of months (Fabbiano et al., 2003), which may represent disk/jet precession as seen in SS 433 (Clarkson et al 2004 in prep) or recurrent transitions between low/hard and high/soft states as for the black hole binary Cyg X-1. Analysis of 20 years' observations with Einstein, ROSAT,

BeppoSAX and ASCA shows that the ULX source X-9 in M31 is highly variable on a timescale of months, but no strictly periodic variability is present (La Parola et al., 2001). Eight Chandra ACIS-I observations of the Andromeda Galaxy (M31) show one source in the M31 globular cluster Bo86 with a possible 200-d periodicity, plus 99 further sources showing variability on a timescale  $\gtrsim$  months (Kong et al., 2002). Indeed, of the sources resolved,  $> 55\%$  show this variability.

### 2.1.7 Mechanisms for the Superorbital variations

Following intense observations of spectral and temporal variability of Her X-1 and LMC X-4 at many wavelengths, the most popular explanation for superorbital periodicities involves a tilted or warped accretion disk. Precession of this disk under internal and external torques leads to progressive occultation of the inner disk and neutron star itself, thereby producing the longterm variation in X-ray brightness observed. The nature of the torque experienced by a warped accretion disk is a subject of intense recent theoretical investigation; the most promising candidates at this stage are accretion disk-wind driving (Schandl, 1996), interaction of sheet currents on the accretion disk with the magnetosphere of the donor or the neutron star (Lai 1999), and the torque produced by the pressure of reprocessed radiation (Wijers & Pringle 1999, Ogilvie & Dubus 2001).

We argue in Chapters 5 & 6 that radiation-driven warping is the most persuasive framework for the luminous XRB Her X-1, SMC X-1, LMC X-4 and Cyg X-2, but of course this mechanism is not necessarily at work for *all* binaries with superorbital periods. Isolated neutron stars are understood to undergo free precession about their axis (Jones & Andersson, 2002), so it is perhaps reasonable to expect neutron stars in X-ray binaries to do so as well, provided the accretion process does not overcome the tendency to precess (Truemper et al., 1986). The presence of a third body in a hierarchical orbit may also cause large variations in  $\dot{M}_1$  and thus  $L_X$ . The third body causes the inner XRB to undergo nodal precession, leading to mass transfer variations even if the inner binary is eccentric only by a small amount (Chou & Grindlay 2001, Mazeh & Shaham 1979).

## 2.2 Non-planar, precessing accretion disks

We introduce here the physics of precessing, warped accretion disks, as inspection of the above survey of superorbital variations shows it to be the most common interpretation for the superorbital

periodicities. In many of the systems most amenable to study, a precessing disk is indeed the natural explanation; for example, in SS 433 the precession of the jets demands disk precession, as the angular momentum evolution of the jets must be driven by the much larger angular momentum of the disk. Another example is Her X-1, which shows both an X-ray superorbital period of 35 days, stable to within 10% (Oegelman et al., 1985) and a 1.6-d photometric period. The photometric period is the beat between Her X-1's 1.7d orbital period and the 35-d superorbital cycle, leading to the “clockwork model” (Deeter et al., 1976; Gerend & Boynton, 1976), in which the superorbital cycle of on-off states is due to the body of a retrogradely precessing accretion disk periodically obscuring the neutron star from the point of view of the *observer*, and the 1.6-d variation corresponding to the time taken for the accretion disk to present the same area to the *neutron star* for X-ray reprocessing.

This is easier to visualise when one considers the similar phenomenon of negative superhumps; in these systems, were the accretion disk to be perfectly circular, no superhump variability would exist. However, because of *retrograde* precession of an eccentric disk, the system takes just *under* an orbital cycle to return to the same donor-disk geometry, leading to periodic modulation of the dissipation on the beat period between the orbit and precession; e.g. Haswell et al (2001).

### 2.2.1 Time Evolution of a Warped Disk

Assuming an accretion disk is forced into a warped configuration by some mechanism, tracing its subsequent evolution in the general case is an extremely difficult three-dimensional problem of magnetohydrodynamics, full numerical treatment of which is unfeasable at present (Ogilvie & Dubus, 2001). A short-cut method has traditionally been used, of developing a one-dimensional description of the warped disk and following its evolution on the long, viscous timescale (thus assuming shorter-term variations are averaged out through the cycle). This approach was taken in deriving the first expressions for the radiation torque (Petterson, 1977; Hatchett et al., 1981); however, these early models failed to conserve angular momentum in the absence of external torques and are therefore incorrect (Pringle, 1992). In the absence of a driving force maintaining a warped configuration, we expect dissipative viscous interactions within the disk to lead to settling back into the binary plane on  $\sim$  a few viscous timescales. Analytic derivation of the time dependence of  $\alpha$ -disks in the limit of small warps and thin disks suggested resonant modes arising in the disk leading to restoration of planar configuration on faster timescales  $< t_{visc}$  (Papaloizou & Pringle,

1983). A later revisit (Pringle 1992) of the basic derivation relaxed the requirement of small tilt angles, which we summarise here as it forms the basis for many subsequent treatments of warped disks in the literature.

The accretion disk is divided into concentric annuli of small width  $\Delta R$ . The tilt of each annular element is described by the angle from the vertical  $\beta(R, t)$ , with the line of intersection of the annulus and the binary plane making angle  $\gamma(R, t)$  with the line connecting the centres of mass of the binary components. The tilt of each annulus is thus described by the unit normal vector

$$\hat{\mathbf{l}} = (\cos \gamma \sin \beta, \sin \gamma \sin \beta, \cos \beta) \quad (2.1)$$

The angular momentum  $\Delta L$  and mass  $\Delta M$  associated with the annulus at  $(R \pm \Delta R)$  are then

$$\Delta M = 2\pi R \Delta R \Sigma \quad (2.2)$$

$$\Delta \mathbf{L} = 2\pi R \Delta R \Sigma R^2 \Omega \hat{\mathbf{l}} \quad (2.3)$$

with  $\Omega(R)$  the angular velocity of material in the annulus (assumed to follow Keplerian orbits),  $\Sigma(R)$  the surface density and  $\hat{\mathbf{l}}$  the unit normal vector to the annular element, describing annular tilt. in the direction normal to the annular element. In the case of a planar disk, the angular momentum and surface density conservation relations are derived by considering the flow of each quantity across the annular element  $\Delta R$ . Mass flows inwards with velocity  $V_R$ , carrying with it mass flux ( $V_R 2\pi R \Sigma$ ) and angular momentum flux ( $V_R 2\pi R \Delta R \Sigma R^2 \Omega$ ). Evaluating these quantities at  $R$  and  $R + \Delta R$ , and including the viscous torques  $G_{visc}$  acting between neighbouring annuli, gives (see e.g. Frank, King & Raine 2002):

$$\frac{\partial \Sigma}{\partial t} + R^{-1} \frac{\partial}{\partial R} (R \Sigma V_R) = 0 \quad (2.4)$$

and

$$\frac{\partial}{\partial t} \left( \Sigma R^2 \Omega \hat{\mathbf{l}} \right) + R^{-1} \frac{\partial}{\partial R} \left( R \Sigma V_R R^2 \Omega \hat{\mathbf{l}} \right) = \frac{1}{2\pi R} \frac{\partial \mathbf{G}_{visc}}{\partial R} \quad (2.5)$$

The next step is to evaluate the torque  $G_{visc}$  between neighbouring annuli. This can be expressed as two components; one describing the torque in the plane between annuli rotating at different angular velocities, with its associated kinematic viscosity  $\nu_1$ :

$$\mathbf{G}_1 = 2\pi R \Sigma R^2 \nu_1 \left( \frac{\partial \Omega \hat{\mathbf{l}}}{\partial R} \right) \quad (2.6)$$

and a component associated with the tendency of neighbouring annuli to attempt to equalise their tilt angles, with its associated kinematic viscosity  $\nu_2$ :

$$\mathbf{G}_2 = 2\pi R \Sigma R^2 \Omega \frac{1}{2} \nu_2 \left( \frac{\partial \hat{\mathbf{l}}}{\partial R} \right) \quad (2.7)$$

allowing us to re-write (2.4) as

$$\frac{\partial}{\partial t} \left( \Sigma R^2 \Omega \hat{\mathbf{l}} \right) + R^{-1} \frac{\partial}{\partial R} \left( R \Sigma V_R R^2 \Omega \hat{\mathbf{l}} \right) = \frac{1}{R} \frac{\partial}{\partial R} \left( \nu_1 \Sigma R^3 \frac{\partial \Omega}{\partial R} \hat{\mathbf{l}} \right) + \frac{1}{R} \frac{\partial}{\partial R} \left( \frac{1}{2} \nu_2 \Sigma R^3 \Omega \frac{\partial \hat{\mathbf{l}}}{\partial R} \right) \quad (2.8)$$

Solving for  $V_R$  with (2.3) yields (after some manipulation) the following relation for the evolution of  $\Sigma$ :

$$\frac{\partial \Sigma}{\partial t} = R^{-1} \frac{\partial}{\partial R} \left[ \frac{\frac{\partial}{\partial R} \left( -\nu_1 \Sigma R^3 \frac{\partial \Omega}{\partial R} \right)}{\frac{\partial}{\partial R} \left( R^2 \Omega \right)} \right] + R^{-1} \frac{\partial}{\partial R} \left[ \frac{\frac{1}{2} \nu_2 \Sigma R^3 \Omega \left| \frac{\partial \hat{\mathbf{l}}}{\partial R} \right|^2}{\frac{\partial}{\partial R} \left( R^2 \Omega \right)} \right] \quad (2.9)$$

From this the usual evolution of  $\Sigma$  in the case of a flat disk can be recovered by setting  $\frac{\partial \hat{\mathbf{l}}}{\partial R} = 0$  (Pringle, 1996). The first term on the right in (2.9) thus describes the planar diffusion of angular momentum through the usual viscous interaction of neighbouring annuli. The second term describes the result of interaction between neighbouring twisted annuli: annuli at  $R$  and  $R + \Delta R$  will equalise their tilts to share angular momentum through the vertical stress governed by  $\nu_2$ . To do so, both must dissipate the energy associated with moving in the plane normal to the new tilt  $\hat{\mathbf{l}}'$ , and so both annuli move inwards to restore Keplerian balance. Thus the second term in (2.9) represents *inward* flow associated with local untwisting of the twisted disk.

By further manipulation, Pringle (1992) was able to obtain the evolution of the disk tilt  $\hat{\mathbf{l}}$ , which forms the basis for the more recent analyses of disk-warps:

$$\frac{\partial \hat{\mathbf{l}}}{\partial t} + \left[ V_R - \nu_1 \Omega^{-1} \frac{\partial \Omega}{\partial R} - \frac{1}{2} \nu_2 \frac{\frac{\partial}{\partial R} \left( \Sigma R^3 \Omega \right)}{\Sigma R^3 \Omega} \right] \frac{\partial \hat{\mathbf{l}}}{\partial R} = \frac{\partial}{\partial R} \left( \frac{1}{2} \nu_2 \frac{\partial \hat{\mathbf{l}}}{\partial R} \right) + \frac{1}{2} \nu_2 \left| \frac{\partial \hat{\mathbf{l}}}{\partial R} \right| \hat{\mathbf{l}} \quad (2.10)$$

That this framework does conserve angular momentum has been shown by re-writing (2.10) explicitly in terms of the angular momentum density (equation 2.6 of Pringle 1992). Note that up to this point there is no description of any torque from driving processes; the relations thus far can be used to describe the response of the accretion disk to a sudden imposition of a warped configuration. In the absence of a driving mechanism, the disk will thus relax towards the plane.

### 2.2.2 Mechanisms of Warp Excitation

The clockwork model can fit the superorbital periods in SS 433, Her X-1 and the 30-d cycle in LMC X-4 quite well, which led to the general acceptance that a precessing disk was indeed responsible for the superorbital periods in at least these bright X-ray binaries. Explaining exactly how such a tilted, precessing disk arose was more of a challenge. The donor will exert a tidal force on an accretion disk, encouraging retrograde precession, but as this force is a strong function of location

in the disk (Katz, 1973), the difficulty lies in explaining how the entire disk might precess with a single period. Furthermore, such a warp must be continually driven, as viscous torques between neighbouring disk elements should lead the disk to settle back to the binary plane.

An early suggestion was the slaving of the accretion disk to a donor whose spin axis is misaligned with the binary axis (Roberts, 1974); however, tidal dissipation is expected to synchronise the donor with the binary motion before circularisation is achieved (Mazeh & Shaham 1979), so the near-circular orbits of Her X-1 and the majority of the systems with precessing disks argues strongly against this model. Petterson (1977) considers a number of scenarios in which a small tilt in the accretion disk is set up by a slight misalignment of the rotation axis of the compact object (through the quadrupole moment of its gravitational field and the much more significant frame dragging of the inner disk), forced precession of the donor (through misalignment of its rotation axis) and precession of disk elements through the influence of the donor (c.f. Katz 1973). In general, models invoking the donor lead to instability in the outer region of the disk, while compact object mechanisms affect the inner disk (Petterson 1977). We introduce the warp-driving mechanisms here, focusing particular attention on radiation-driving.

### 2.2.3 Radiation-Driven Warping

For our purposes the most significant result of the investigation of Petterson (1977) was the realisation that, *if* a warp can be generated, it will be exacerbated through radiation pressure exerted on the unshadowed disk area from the central X-ray emission. Iping & Petterson (1990) attempted to use the early framework to numerically model the longterm periodicity in Cen X-3; although the base equations were at fault (Pringle 1992), their qualitative conclusion that a disk warp can indeed be maintained by strong irradiation, has been borne out by subsequent studies. Indeed, as pointed out by Pringle (1996), for a range of accretion luminosities, radiation pressure alone is sufficient to drive and maintain a disk-warp. Optically thick disk material absorbs incident radiation from the central accretor, re-radiating it along the local normal and producing a backreaction on the disk. Any small net asymmetry in the re-radiation pattern in the disk will lead to substantial warping provided the central luminosity is high enough to overcome viscous torques and drive a warping mode which grows with time. We introduce here the analysis leading to this conclusion. Armed with the appropriate description of the disk evolution (Pringle 1992 and section 2.2.1), the action of radiation-driving is incorporated by adding a term describing its torque  $dG_{rad}$  to the

right of equation (2.10) of the form

$$+ \left( \frac{1}{\Sigma R^2 \Omega} \right) \frac{1}{2\pi R} \frac{d\mathbf{G}_{\text{rad}}}{dR} \quad (2.11)$$

with the term in brackets included to re-cast the radiative torque contribution in terms of the tilt vector  $\hat{l}$  rather than the angular momentum density  $\Sigma R^2 \Omega \hat{l}$ . We introduce here the derivation of  $dG_{\text{rad}}$  following Pringle (1992,1996). We make the assumption that the disk is optically thick to incident radiation everywhere, so that all incident radiation from the point source on a disk element is absorbed by the element. Each element is then assumed to reradiate the incident flux isotropically. Small tilt of an annulus out of the binary plane then produces corresponding  $dG_{\text{rad}}$  on the annular disk element  $dR$ :

$$d\mathbf{G}_{\text{rad}} = \oint \mathbf{r} \times d\mathbf{F} \quad (2.12)$$

where we consider only those parts of the annulus not shadowed by inner disk material, and  $dF$  is the reaction force on area element  $dA$ :

$$d\mathbf{F} = \frac{2}{3c} \frac{L}{4\pi r^2} \hat{\mathbf{r}} \times d\mathbf{A} \quad (2.13)$$

We may shortcut the full analysis of Pringle (1996) by anticipating the result  $dG_{\text{rad}} \sim R dF$  (e.g. Frank, King & Raine 2002), which we can see will give

$$\frac{1}{2\pi R} \frac{\partial G_{\text{rad}}}{\partial R} \sim \frac{L}{6\pi R c} \quad (2.14)$$

(Note that the interception of radiation from the point source will exert a corresponding radial radiation pressure on each disk element, but this does not of course contribute any torque as there is no azimuthal component.) Investigating the evolution of initially small warps, Pringle (1996) makes the approximation of small vertical tilt  $\beta$ , so that  $\hat{l} = (\beta \cos \gamma, \beta \sin \gamma, 1)$ , to allow the evolution of  $\hat{l}$  to be re-cast in terms of

$$W = \hat{l}_x + i\hat{l}_y = \beta e^{i\gamma} \quad (2.15)$$

Using reasonable assumptions on  $V_R(\Omega, \nu_2)$  and  $\nu_1(\Sigma, R^3 \Omega)$  and assuming constant  $\nu_2$ , this allows conservation of  $\hat{l}$  (2.10,2.11, and the full version of 2.14; see Pringle 1996) to be re-written as

$$\frac{\partial W}{\partial t} = -i\Gamma \frac{\partial W}{\partial R} + \frac{1}{2} \nu_2 \frac{\partial^2 W}{\partial R^2} \quad (2.16)$$

where the radiation torque is parameterised as

$$\Gamma = \frac{L}{12\pi \Sigma R^2 \Omega c} \quad (2.17)$$

We can now carry out a simple perturbation analysis to examine the behaviour of a disk initially perturbed to some (small) tilt from the binary plane. Seeking wavelike solutions to 2.16 of the

form

$$W \propto e^{i(\sigma t + kR)} \quad (2.18)$$

(where  $kR \gg 1$  yields finally the dispersion relation

$$\sigma = i\left(\frac{1}{2}\nu_2 k^2 - \Gamma k\right) \quad (2.19)$$

This relation correctly predicts that without driving (i.e.  $\Gamma = 0$ ), any warp will decay back to the plane. Initially small warps will grow provided  $\text{Im}(\sigma) < 0$ , i.e. for wavenumbers

$$0 < k < 2\Gamma/\nu_2 \quad (2.20)$$

or explicitly,

$$L \gtrsim 12\pi^2\nu_2\Sigma v_\phi c \quad (2.21)$$

(There is a subtlety here in that  $\phi$  is the azimuthal co-ordinate on the surface of the disk rather than in the binary plane; however as we have assumed small tilt the distinction is at present unimportant).

Note this result is similar in form to that expected from a simple physical argument; as pointed out by Frank, King & Raine (2002) we might ask simply under what conditions the radiative torque  $dG_{rad}$  exceeds the opposing viscous torque  $dG_2$ ; setting  $dG_{rad} > dG_2$  then leads immediately to the condition  $L \gtrsim 6\pi\nu_2\Sigma v_\phi c$ . However this argument says nothing about the growth or decay of warping, for which the perturbation analysis outlined above is preferred.

Simplifying the expression for  $\Gamma$  with Pringle's assumptions ( $\dot{M}_1 = 3\pi\Sigma\nu_1$ ,  $R/R_{Sch} \simeq 2v_\phi^2/c^2$ ) gives the approximate analytic stability criterion that has been used in all subsequent studies of radiation-driven disk warps. A disk annulus at radius  $R$  will experience growth of warping modes provided its radius  $R$  satisfies

$$\frac{R}{R_{Sch}} \gtrsim \frac{8\pi^2\eta^2}{\epsilon^2} \quad (2.22)$$

where we define  $\eta = \nu_2/\nu_1$  and the accretion efficiency  $\epsilon = L/(M_1 c^2)$ . Once the disk warp grows, self-shadowing becomes important and the stability of the disk to the various warping modes must be calculated numerically. While providing a useful rule of thumb, the above analysis suffers from the weakness that the simplest modes are expected to dominate, when  $kR \simeq 1$ , far from the region  $kR \gg 1$  assumed by the dispersion-relation analysis above (Maloney et al., 1996). Global analysis under the assumption of an isothermal disk produces results that are qualitatively similar, and were shown to be consistent with both the  $\sim 164$ -day precession timescale of the SS 433 jets (Margon 1984) and the shape of the warped maser disk in NGC 4258 (Greenhill et al., 1995).

Motivated by this success, a survey of the superorbital periodicities in X-ray binaries was undertaken under the linear framework of Pringle (1996), to examine the appropriateness of radiation-driven warping and precession as an explanation for the superorbital periodicities as a whole (Wijers & Pringle, 1999). In particular, the stability to warping was estimated analytically by linearising the expression for the evolution of the element normal vector  $\hat{\mathbf{i}}$  (Pringle 1996) and searching for normal-mode behaviour. The behaviour of a disk-warp was also simulated numerically, through direct integration of the earlier equations of Pringle (1992). The evolution of the accretion disk was characterised as a competition between forced precession due to the donor, the growth of instabilities due to the irradiation, and viscous dissipation of these instabilities. For the parameters of the 12 systems studied, the predicted timescales for forced precession were found to differ substantially from those observed, whereas the timescale of precession from radiation-driving was found to agree with the observed superorbital periodicities to within a factor  $\sim 10$ .

Furthermore, these simulations made predictions as to the time behaviour of the disk-warp, suggesting that beneath a critical central luminosity  $L_{crit}$ , the warping is not stable with regular precession, but instead evolves in a chaotic manner. Assuming warping and precession to be stable and regular, the timescale for disk precession at radius  $r$  was estimated in terms of the luminosity emitted by the disk  $L_d$  as

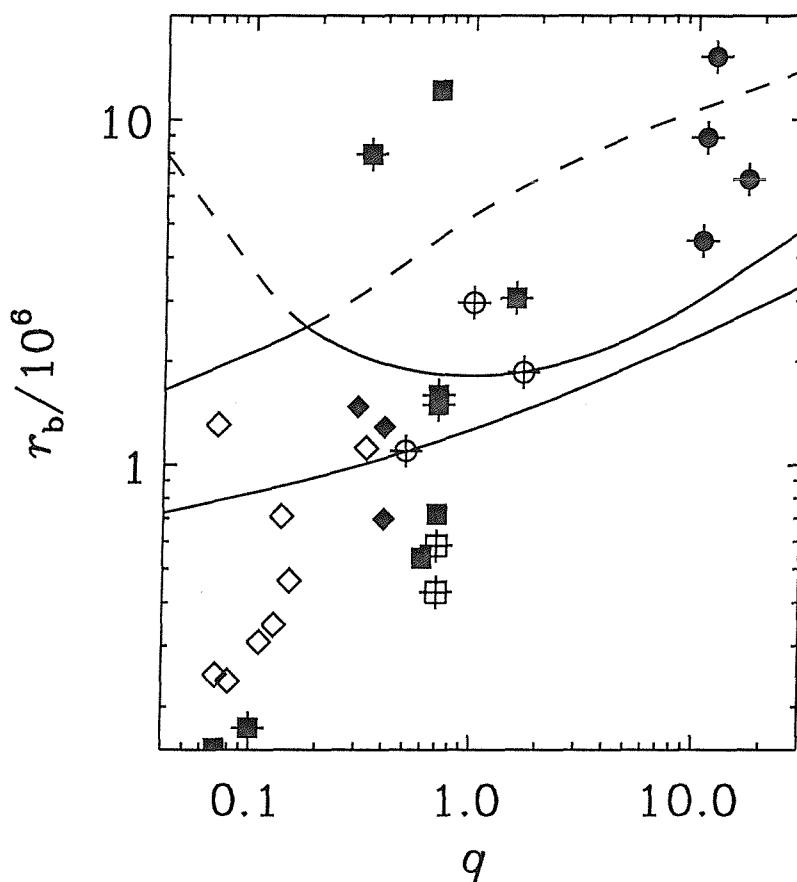
$$P_{prec} \simeq \frac{\Sigma r}{L_d} \quad (2.23)$$

where  $\Sigma$  and  $r$  are expressed in multiples of characteristic values  $\Sigma_0$  and  $r_0$ , to be determined in the numerical modeling. Since in the strongly irradiated limit  $L_d \sim L_x$ , this predicts the precessional timescale should be inversely proportional to the luminosity of the central source. This is *not* in general observed (Chapters 5 & 6), which may come as no surprise to the reader; as the precessional timescale varies with position in the disk and the complex self-shadowing will complicate the  $L_d/L_x$  relation, a global analysis of warping modes is necessary to trace the true disk behaviour.

This approach suffered from a somewhat coarse numerical grid and laborious treatment of self-shadowing (Ogilvie & Dubus 2001). The expressions governing the evolution of the disk surface were recently formally re-derived from the Navier-Stokes equation (Ogilvie, 1999). A set of coordinates following the warp were defined, and the conservation equations then derived for thin disks under the approximation of isotropic viscosity. (Due to its mathematical complexity, this analysis is not reproduced here). It is assumed that the resonant bending-waves found by Papaloizou & Pringle (1983) can be neglected (i.e.  $\alpha(r) \gtrsim H/r$ ; Ogilvie 1999). The effective viscosities  $\nu_1$  and  $\nu_2$  were also linked to the general kinematic viscosity. This approach has two features

which allow confidence of its applicability to warped, fluid disks: (i) it has been derived “from the ground up,” and (ii) is fully nonlinear, so is also valid for large warps. In a followup paper (Ogilvie, 2000), viscous dissipation and energy transport under the full Rosseland approximation were included, with the fluid consisting of an ideal gas and the Kramers or Thompson opacities included as appropriate. The resulting theoretical framework “may be considered as the logical extension of the alpha-disk theory to time-dependent, warped accretion disks” (Ogilvie & Dubus 2001). In a complementary study to that of Wijers & Pringle (1999), the complex eigenvalues of the nonlinear bending modes of the accretion disk were computed for a wide variety of system parameters, this time using the full evolutionary equations of Ogilvie (1999,2000) and assuming realistic boundary conditions (as compared to e.g. Maloney, Begelman & Pringle 1996). A more rigorous derivation of self-shadowing was integrated into the analysis, as well as correct expressions for the torque due to radiation pressure, the addition of matter in the accretion stream, tidal torque and the internal torque of disk elements. Bifurcation analysis was used to determine the long-term stability of the resulting warping modes. For reasonable values of the global disk viscosity  $\alpha$  and accretion efficiency  $\eta$ , the stability of an accretion disk to warping is determined by the binary separation  $r_b$  and mass ratio  $q = \frac{M_2}{M_1}$ . This provides a second-order prediction on how disk-warppers will behave as a class of systems, which we put to the test in chapters 5 and 6.

This analysis does not make easily-used predictions as to the disk precessional timescale in the manner of Wijers & Pringle (1999). This is partially due to the particular nondimensionalization used (in particular the precessional timescale is quoted in units dependent on the instantaneous value of  $\dot{M}_1$ ), but also due to the authors’ lack of confidence in the applicability of such predictions. When declining to apply their predictions to the prototypical stable warper Her X-1, for example, the authors point out (i) the lack of information as to instantaneous  $\dot{M}_1$ , (ii) the lack of full exploration of the contribution of the tidal torque and (iii) the lack of full exploration of the dependence on  $\alpha, \epsilon, q$  and the true location of the mass-addition point  $r_{add}$ . The reader is thus left with the difficulty of interpreting the two papers of Wijers & Pringle (1999) and Ogilvie & Dubus (2001); while the former makes the more useful predictions for individual sources in terms of the time-dependent behaviour that may be expected, the latter is based on a fuller development of the physics involved. This author believes the best middle-ground is to use Ogilvie & Dubus (2001) for comparisons between sources (especially as the binary radius  $r_b$  is now the control parameter) and to use Wijers & Pringle (1999) when estimating the precessional lightcurve of individual sources beyond the global stability predictions of Ogilvie & Dubus (2001).



**Figure 2.1:** Stability to warping for  $\alpha = 0.3, \eta = 0.1$  in  $(r_b, q)$  space.  $r_b$  (the binary radius) is in units of  $GM_1/c^2$ , while  $q$  is the mass ratio of the system. For  $r_b < 0.7$ , beneath the lowest solid line, no warp is possible. The first solid line corresponds to matter input at the outer Lindblad radius; every source above this line for which matter is input in this way can sustain mode 1 and higher warps. If matter is input at the disk circularisation radius, this region still corresponds to no warp. In the next region up, simple mode 0 warping is possible but cannot persist. Between the upper boundary of this region and the dotted line, mode 0 warping persists. Finally in the uppermost region, mode 1 and higher warping modes become possible. With  $r_b \sim 9$  and  $q = 10.8$ , SMC X-1 is on the boundary between stable mode 0 and higher modes of warping. From Ogilvie & Dubus (2001).

A key difference between the full analysis of Ogilvie & Dubus (2001) and that of Wijers & Pringle (1999) and indeed earlier work, is the value for the viscosity ratio  $\eta$  in the accretion disk. As pointed out by Pringle (1996), the assumption that disk viscosity is isotropic ( $\eta \sim 1$ ) implies that for realistic accretion efficiencies the disk in an LMXB should be unstable to warping outside  $\sim 3 \times 10^8$  cm. This predicts that even the ultracompacts such as 4U 1916-053 should be unstable to radiation-driven warping. However, the full analysis of Ogilvie & Dubus (2001) prescribes  $\eta \simeq 10$  for realistic approximations (see their equation 42). Under this parameterisation the critical radius increases by a factor  $\sim 100$ , which means it is now never reached in the ultracompacts. Evidence at present appears somewhat uncertain as to the resolution of this difference; while the limits of Ogilvie & Dubus (2001) do appear to correctly predict the stability of the bright XRB to warping (Chapters 5 & 6), evidence is now growing that the ultracompact 4U 1916-053 may be undergoing warping and precession. Observations have been proposed to follow the spectral evolution of the inner disk region with the precessional phase (suggested by the orbital and superhump lightcurve); such observations will uncover if a precessing warp is indeed at work in this system.

#### 2.2.4 Wind-driven Warping

The strong irradiation of the accretion disk by the central source is thought to lead also to the generation of a strong disk-wind. This then drives the warp through essentially the same mechanism as the radiation-driving, except that the driving particles are in this case ions rather than photons. Under the right conditions this mechanism is expected to be comparable to the effects of radiation-driving, as the specific momentum carried by each driving ion is higher (Pringle, 1996). The wind-driving mechanism has been studied in some detail (Schandl & Meyer, 1994), with authors even going so far as to fit the profile of the disk rim to the long-term lightcurve, assuming the structure thus fit would be produced by wind-driven warping. At the time of writing, however, no physically consistent connection has been made between wind-driving and the shape of the disk rim. In particular the lopsided structure of the rim in these simulations appears to require at least some of the disk matter to orbit a point above the plane of the disk, which is not physical (Billington et al., 1996). However, the broad basis for the wind-driving mechanism is probably correct; it is perhaps surprising, therefore, that investigation into warp driving from the disk-wind appears to have virtually halted since 1997. The major difficulty appears to be the derivation of a consistent expression for the radiation torque; as several authors have commented, the most recently-used version is based on an earlier, incorrect equation of motion for the disk (e.g. Wijers & Pringle

1999). We conclude that this mechanism is clearly deserving of further study.

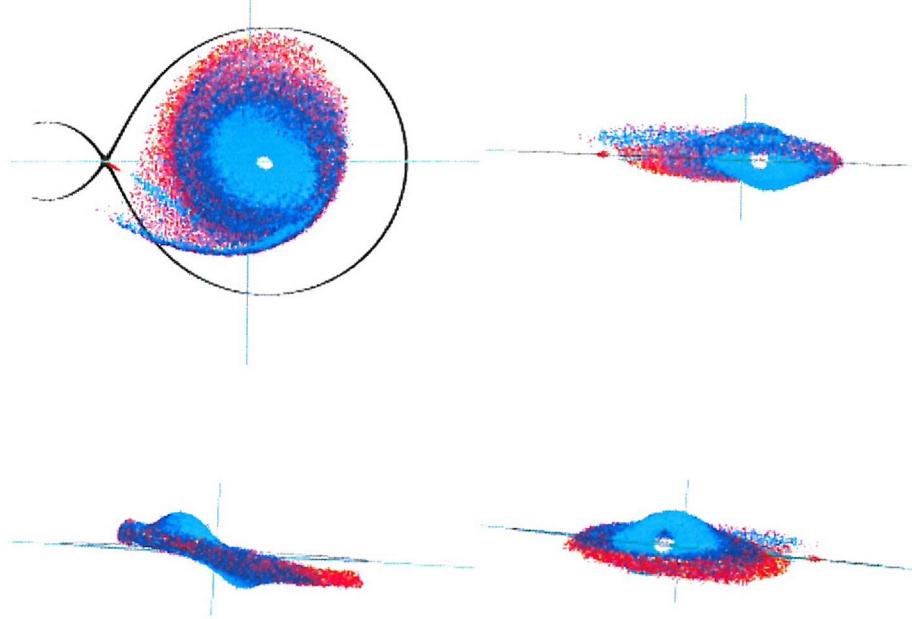
### 2.2.5 Lense-Thirring Precession: the Bardeen-Petterson effect

As pointed out by Bardeen & Petterson (1975), frame-dragging by a spinning, misaligned compact object should drive precession of the *inner* accretion disk. However the viscosity of course tends to smooth out such precession, resulting in an equilibrium state in which the inner disk is aligned with the spin axis of the compact object out to some radius  $r_{BP}$ . Application of the evolutionary equations of Papaloizou & Pringle (1983) suggests this picture is qualitatively correct, with subsequent studies suggesting that over a long timescale the entire disk should tend to align with the compact object (Scheuer & Feiler, 1996; Natarajan & Pringle, 1998). More recent analysis of wave-like propagation of bending modes (rather than diffusive as suggested by earlier work) suggests that if  $\alpha \lesssim H/r$ , a warped configuration may persist indefinitely within  $r_{BP} \sim 200r_{Sch}$  (Demianski & Ivanov, 1997; Lubow et al., 2002). However, as the authors point out, SPH simulations of misaligned accretion disks find no such warps (Nelson & Papaloizou, 2000), so the extension of this result into the full nonlinear regime is not at present clear. If robust, however, this result implies that the slaving of the jet to the inner accretion disk may lead to any bipolar outflow not necessarily taking a direction perpendicular to the accretion disk as a whole. Indeed such misalignment has been observed in at least one X-ray binary with bipolar outflow (see Butt et al., 2003, for a discussion).

### 2.2.6 Magnetic Warping

The disk material is usually highly ionised, which raises the obvious possibility that interaction with the magnetic field of the donor, compact object or indeed the disk itself may provide a non-planar force to drive warping. The possibility of warping due to the magnetic field of the compact object was investigated in detail by (Lai, 1999). Replacing the radiation torque with the magnetic torque under the possible configurations, an analysis similar to that of Pringle (1996) was carried out to examine if such driving could overcome viscous dissipation. Linearising for small tilt angle and examining the resulting dispersion relation for steadily growing solutions, a condition on the magnetic torque term  $\Gamma_m$ :

$$\Gamma_m > \frac{1}{2}\nu_2 k^2 \quad (2.24)$$



**Figure 2.2:** Still frame from SPH simulation of a strongly irradiated accretion disk, showing a prominent warp, in this case for mass ratio  $q = 0.075$ . The warp is stable and precesses in the binary frame. The prototypical warpers Her X-1 and SMC X-1 have  $q > 1$ , a regime which will be the target of future work (Foulkes et al., 2004). Colours represent viscous dissipation. Reproduced with permission.

where the kinematic viscosity  $\nu_2$  and warp-mode wavenumber  $k$  take the same meaning as for sections 2.2.1 and 2.2.2; Lai 1999 writes out the definition of  $\Gamma_m$  and related terms; also compare with equation 2.9 for the radiation-driven case. Application of this work to the superorbital periodicities in XRB suggests a characteristic timescale for precession of the inner accretion disk (equation 6.13 of Lai 1999);

$$P_{prec} \propto B^{0.8} \left( \frac{r\eta}{R_M} \right)^{4.9} \quad (2.25)$$

with  $R_M$  the magnetospheric radius and  $\eta$  the viscosity ratio  $\frac{\nu_1}{\nu_2}$ . However, as the predicted precession period varies with radius in the disk, global analysis is needed to determine how the disk actually behaves. This mechanism has been applied to explain the *millihertz* QPOs seen in several XRBs (Boroson et al., 2000). In particular, HST observations of the LMXB 4U 1626-67 show a strong QPO around 1mHz which is not present in high time-resolution XTE/PCA lightcurves of the source, and indeed is strongest in the UV waveband (Chakrabarty et al., 2001). This variation is naturally interpreted as the precession of a warped inner disk, however at  $P_{orb} \simeq 42$  minutes, this system is beneath the radius limit at which radiation-driven warping is predicted (Ogilvie & Dubus 2001). Application of sensible values of  $B$ ,  $L_X$  and  $\dot{M}_1$  does predict an inner disk precessional timescale commensurate with observation (Shirakawa & Lai, 2002b,a).

## 2.3 Update: the Direct Approach

At the time the work in this thesis was performed, investigation of disk-warps was largely limited to the numerical solution of the disk evolutionary equations and the search for oscillatory modes as described above. Recent work by a team at the Open University and Swinburne is starting to yield consistent 3D simulations of the evolution of the accretion disk with the physics of Pringle (1996) input to the mode. These SPH simulations are showing disk-warps may form and persist for a range of mass ratios (Foulkes et al., 2004). As these simulations develop they will provide a direct numerical tool to characterise the warping of accretion disks and thus provide observational predictions (see figure 2.2).

## Chapter 3

# Datasets

This thesis discusses data taken with the All-Sky Monitor on board the Rossi X-ray Timing Explorer (hereafter RXTE/ASM), the Burst And Transient Source Experiment on board the Compton Gamma-Ray Observatory (hereafter CGRO/BATSE), the Chandra satellite and the InfraRed Imaging Spectrometer on the Anglo Australian Telescope (AAT/IRIS2). In the case of Chandra and AAT/IRIS2, data was reduced and analysed using standard software. These datasets are introduced in the relevant chapters. Analysis of CGRO/BATSE and RXTE/ASM represent the backbone of the work described here, so we introduce these datasets and their analysis in this chapter.

### 3.1 Long-Term X-ray Datasets

The large amount of effort put in to analysis/reanalysis of satellite data has seen the advent of a remarkable set of near-continuous monitoring datasets of nearly the entire sky, with coverages  $\gtrsim 7$  years. Such a long coverage has allowed the variation of long periodicities to be measured. The two instruments that have done most to advance understanding of long periods in XRB are RXTE/ASM and CGRO/BATSE.

### 3.1.1 RXTE/ASM

Launched in December 1995, the Rossi X-Ray Timing Explorer (RXTE) carries an All Sky Monitor (ASM), which gives near-continuous coverage of the entire sky. It consists of three Scanning Shadow Cameras (SSC's), fitted on a rotating platform, each with a field of view of  $(90^\circ \times 6^\circ)$ . Each SSC has a coded slit mask, consisting of 12 subsections with 30 slit elements each, which are held at open or closed positions according to chosen pseudorandom patterns. The displacements and strengths of the shadow patterns thus produced from astronomical X-ray sources are inferred using a position-sensitive proportional counter (PSPC) in each camera. For each SSC/PSPC combination, position information is of course recorded in the direction perpendicular to the slit pattern. Each PSPC contains eight signal anodes suspended in 95% Xenon gas. An X-ray event near any of the signal anodes produces a pair of pulse heights used to calculate the lateral displacement and energy of the event. Spurious events caused by charged particles are vetoed by 12 further anodes. Further filtering of ambient events is achieved through a  $50\mu\text{m}$ -thick Beryllium window coating each PSPC, which low-energy events are unable to penetrate. Further information about the detector setup for the RXTE/ASM instrument can be found in (Levine et al., 1996). Two onboard Event Analysers (EAs) reduce the event data from the PSPC anodes for insertion into the telemetry stream. The spacecraft and ASM assembly maintain a fixed attitude for a period of  $\sim 90$  seconds (a 90-second “dwell”). During this 90s period, one of the EA's accumulates histograms of counts vs position for each anodes, in three energy bands corresponding to  $\sim 1.5\text{-}3.0$ ,  $3.0\text{-}5.0$  and  $5.0\text{-}12.0$  keV. The other EA produces count rates for both X-ray and background events, with time resolution 0.125s and 1.0s respectively. It also produces pulse-height spectra from 0-20 keV with 64-channel resolution, over a time interval of 64 seconds. Typically 5-10 “dwells” are taken per day of each active source.

Before release from MIT, preliminary analysis is performed on the telemetered data to extract source fluxes from the pulse-height data. Intensities are calculated for sources listed as active in a master catalogue, with additional sources located and characterised subsequently. Model shadow patterns are computed for each active source within the field of view, along with patterns representing non-X-ray and diffuse X-ray background. The position histograms obtained by the PSPC/EA1 are then fit to these patterns using linear least-squares minimisation. This is then iterated twice, weighting by the reciprocal of the variance expected in each bin from the unweighted initial minimisation. Photon-counting statistics from the best fit are then used to produce an estimate of the error on the final count rate. Fitted intensities are normalised to the on-axis

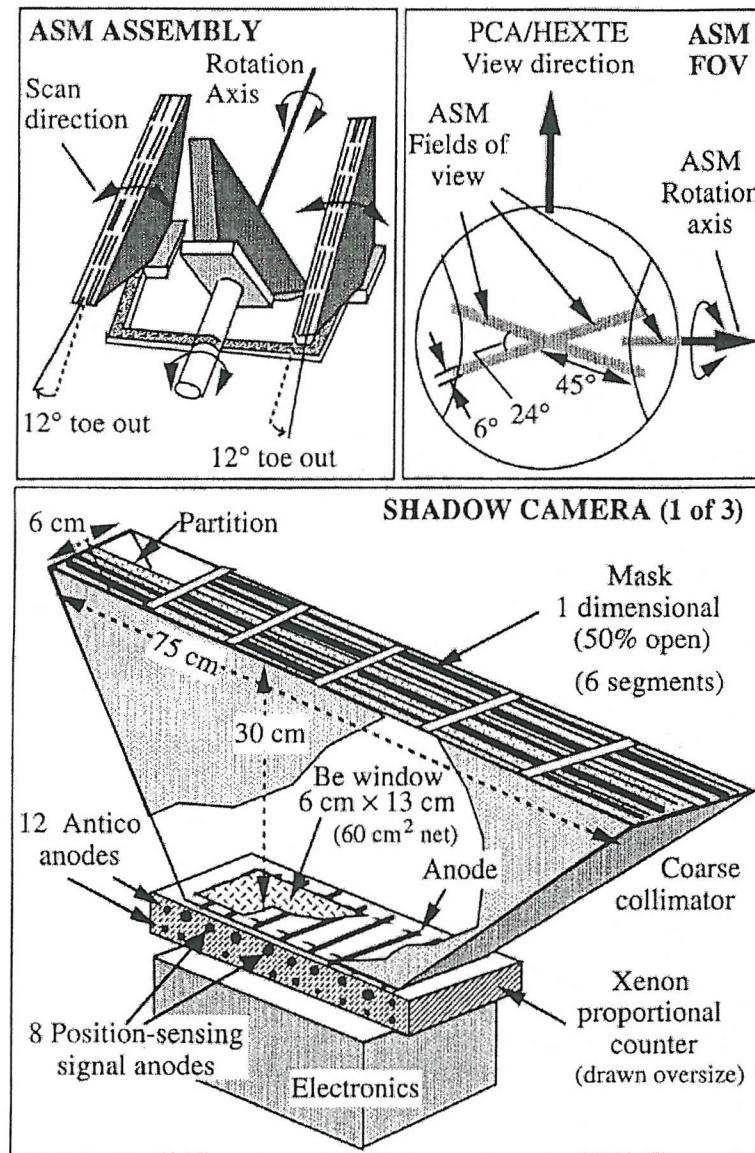
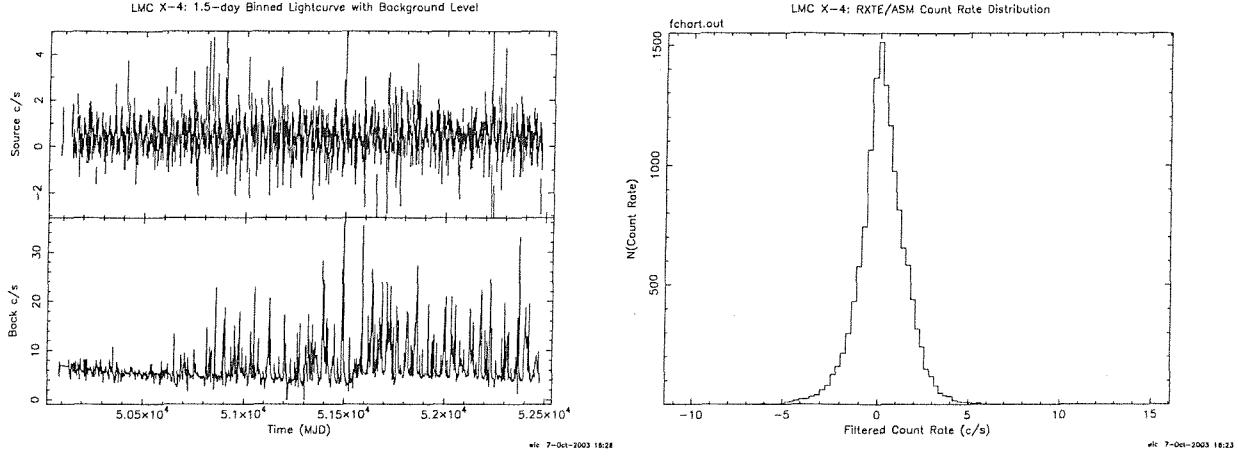


Figure 3.1: RXTE/ASM schematic. From Levine et al (1996).

count rate of SSC1, with corrections made to the effective area variation at large elevations (e.g. a source on-axis in SSC2 may be at high elevation in SSC1), and to absorption from the leak-protection coating on SSC2 and SSC3. The first of these corrections is determined empirically from observations of the Crab Nebula. At this point the data is released to the astronomical community. In the case of faint sources (count rates  $\lesssim 3\text{cs}^{-1}$ ), background contamination is comparable to the source count rate; background fluctuations can then lead to negative source count rates after background subtraction (see figure 3 of Levine et al 1996). In this way, negative ASM count rates are *not* unphysical but instead represent intervals in which the fitted source brightness is lower



**Figure 3.2:** Left: Raw RXTE/ASM source and Background lightcurves for the active but faint LMXB LMC X-4, binned at 1.5-day intervals to reduce scatter for the plot. Right: Count Rate-Frequency histogram for the same source, after filtering in the manner described in the text. Note the count rate distribution extends both sides of zero net count rate: to remove the negative points thus skews the distribution towards higher count rates.

than the size of fluctuations in the background (see figure 3.2). Episodes of particularly high background can cause source flux to be misidentified with background flux in the model fitting. For all analyses reported in this thesis, ASM data taken during intervals of background levels above  $10 \text{ cs}^{-1}$  (1.3-3.0kev) were rejected (see Levine et al 2000)<sup>1</sup>. Furthermore, all points with a reduced- $\chi^2$  statistic  $< 1.5$  were rejected as part of an algorithm modeled after (Levine et al., 2000). In general, however, the background filtering also removed the points with high scatter; subsequent filtering on reduced- $\chi^2$  as suggested removed only a handful of further points.

### 3.1.2 CGRO/BATSE

The Burst And Transient Source Experiment (BATSE) carried on the Compton Gamma Ray Observatory (CGRO) provided near continuous monitoring at energies  $> 20\text{keV}$  for a period of just over 9 years from JD 24448361 - 24451690. Incoming photons are detected by the Large Area Detectors (LADs), which each have collecting area  $2025 \text{ cm}^2$ . The eight LADs are mounted on the corners of the CGRO with normal vectors perpendicular to the faces of a regular octahedron. Any point on the sky can thus always be viewed by four detectors at angles less than  $90^\circ$  to the source direction. This gives CGRO/BATSE a capability of obtaining crude locations to within a few degrees using the count rates from a combination of detectors and the known response of each detector from photons directly from the source, and source photons scattered off the atmosphere

<sup>1</sup>See also [http://heasarc.gsfc.nasa.gov/docs/xte/ASM/asm\\_events.html](http://heasarc.gsfc.nasa.gov/docs/xte/ASM/asm_events.html)

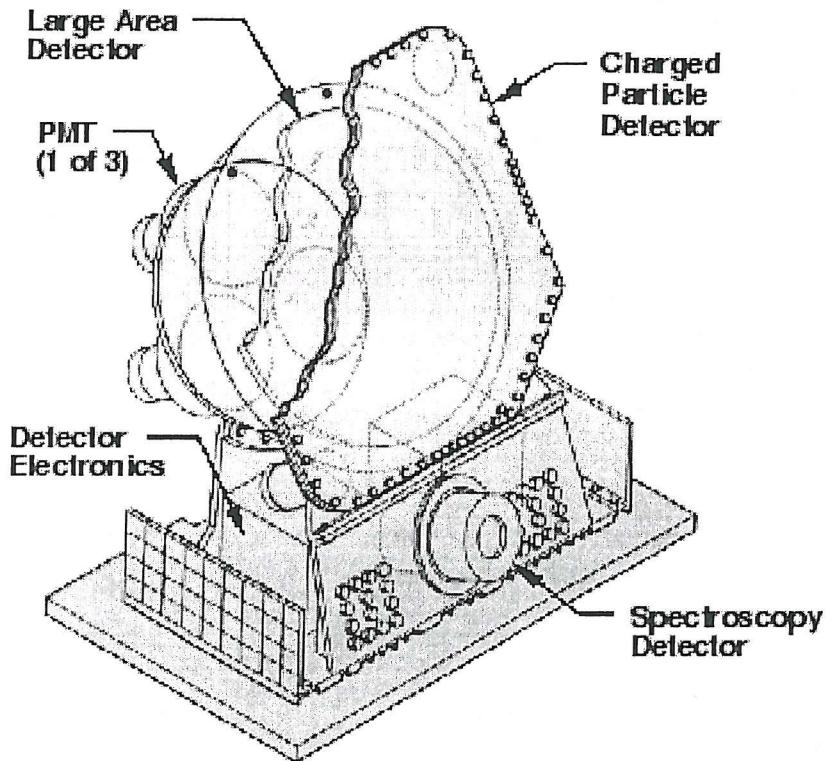


Figure 3.3: Schematic of an individual LAD of CGRO/BATSE. From Harmon et al (2002).

(Pendleton et al., 1995). The Earth Occultation Technique (Harmon et al., 2002), uses the Earth's limb as an occulting mask to further constrain the source position to within  $1^\circ$  and to measure point-source emission. During each 90-minute orbit appropriately positioned point-sources will be occulted by the earth twice. The location of each LAD detector is precisely known, allowing the flux history of any occulting point-source to be obtained from the pattern of step-like changes in flux measured by each LAD. In practice measurements are not made at both occultations, usually due to spacecraft passage through the South American Anomaly (SAA) or out of line-of-sight contact with ground-based tracking systems. Furthermore the spacecraft orbit precesses on a  $\sim 50$ -day period, meaning  $\sim 2$ -day intervals repeatedly occur during which sources at high declination ( $|\delta| > 41^\circ$ ) are not occulted, resulting in coverage gaps. Count rates for a point source are estimated by fitting occultation step features simultaneously for the detectors facing the source, taking into account contributions from background, atmospheric transmission and other sources visible within the interval of occultation, and fitting each channel independently. Fitted count rates for each channel are then averaged over the interval of observation to produce count spectra. Models for these spectra are then folded through the detector response matrices (Pendleton et al 1995), with the best-fit model determined by  $\chi^2$  minimisation of the response-folded counterpart. The form of the spectral model depends on the type of source and is chosen by trial and error.

Finally the flux is calculated from the best-fit spectral model. To increase signal-to-noise the resulting data may be binned over time; the 1-day average 20-100 keV fluxes were used for this analysis.

## 3.2 Timing Analysis Techniques

Chapters 5-8 relate the evolution of long-term periodicities in both RXTE/ASM and CGRO/BATSE datasets. The details of these analyses are given in this section. When seeking periodicities and their harmonic content in time series data, the commonly used methods may be broadly divided into two categories: (i) reconstructing the data from an optimal combination of periodic functions such as sine waves, (e.g. Fourier transforms) and (ii) searching for periods on which features in the dataset recur (e.g. Phase Dispersion Minimisation). We employ examples from both in our analysis.

### 3.2.1 Lomb-Scargle Power Spectrum

Reconstruction of the dataset from sinusoids may proceed in two ways: the Fourier Transform breaks the data down into its representation as sinusoids, with the power corresponding to the relative contribution of sinusoids at each period. Alternatively, sinusoids may be fit to the data and the least-squares fit examined for the strongest contribution. The two approaches move towards the same approximation from opposite directions, so in principle are quite similar to each other (Lomb, 1976) in fact, with the appropriate modifications (Scargle, 1982), the two methods are identical.

The Discrete Fourier Transform (hereafter DFT), is defined for arbitrary sampling as

$$DFT_X(\omega) = \sum_j^{N_0} X(t_j) e^{-i\omega t_j} \quad (3.1)$$

where  $t_j$  is the time of the  $j$ th measurement of  $X$  and  $N_0$  is the number of measurements. The periodogram is then defined as

$$P_{X,DFT}(\omega) = \frac{1}{N_0} \left| \sum_j^{N_0} X(t_j) e^{-i\omega t_j} \right|^2 \quad (3.2)$$

The DFT has been popular for two main reasons. Firstly, in the case of even sampling, then taking  $\Delta t = 1$ ,  $t_j = j$  and  $X(t_j) = X_j$  gives the form

$$P_{X,DFT}(\omega) = \frac{1}{N_0} \left| \sum_{j=1}^{N_0} X_j e^{-i(j\omega)} \right|^2 \quad (3.3)$$

which can be quickly evaluated at the set of  $N_0$  frequencies  $2\pi n/T$ . Secondly, in the evenly sampled case the statistics are well-understood, with the probability that a peak of power between  $z$  and  $z+dz$  could be generated from frequency-independent white noise alone is given by the exponential relation

$$P_z dz = e^{-z} dz \quad (3.4)$$

making the setting of confidence levels straightforward.

However, in the case of uneven sampling, the DFT displays serious drawbacks. Large spurious peaks are produced from the uneven sampling, which can easily be mistaken for a real periodic signal. More critically, the noise prediction (3.4) above does not hold for the general case (3.2), making the evaluation of confidence levels extremely difficult. Motivated by a desire to re-create (3.4) for uneven sampling, Scargle (1982) modified the periodogram to

$$P_{X,LS}(\omega) = \frac{1}{2} \left[ \frac{\left[ \sum_j X_j \cos(\omega(t_j - \tau)) \right]^2}{\sum_j \cos^2(\omega(t_j - \tau))} + \frac{\left[ \sum_j X_j \sin(\omega(t_j - \tau)) \right]^2}{\sum_j \sin^2(\omega(t_j - \tau))} \right] \quad (3.5)$$

with

$$\tan(2\omega\tau) = \frac{\sum_j \sin(2\omega t_j)}{\sum_j \cos(2\omega t_j)} \quad (3.6)$$

As pointed out by Cumming, Marcy & Butler (1999; see also Horne & Baliunas 1986), if the noise variance  $\sigma_0$  is known in advance (for example when the noise is dominated by a known effect such as photon counting statistics) then with the normalisation

$$P_N(\omega) = \frac{P_{X,LS}(\omega)}{\sigma_0^2} \quad (3.7)$$

the exponential probability distribution is indeed recovered. Detrending the data in any way other than subtracting the mean level, for example by removal of a quadratic, is thus *not* appropriate in principle (though if it has little effect on the variance this procedure is allowable in practice). Note that *in general* the noise variance is not known in advance, but must be estimated from the data. In this case the exponential distribution (3.4) above is *not* appropriate and its use must be replaced by Monte-Carlo simulations (see section 3.3.1). However, the LS periodogram is less

prone to spurious peaks due to uneven sampling, so we use it here for analysis of the unevenly sampled RXTE/ASM data.

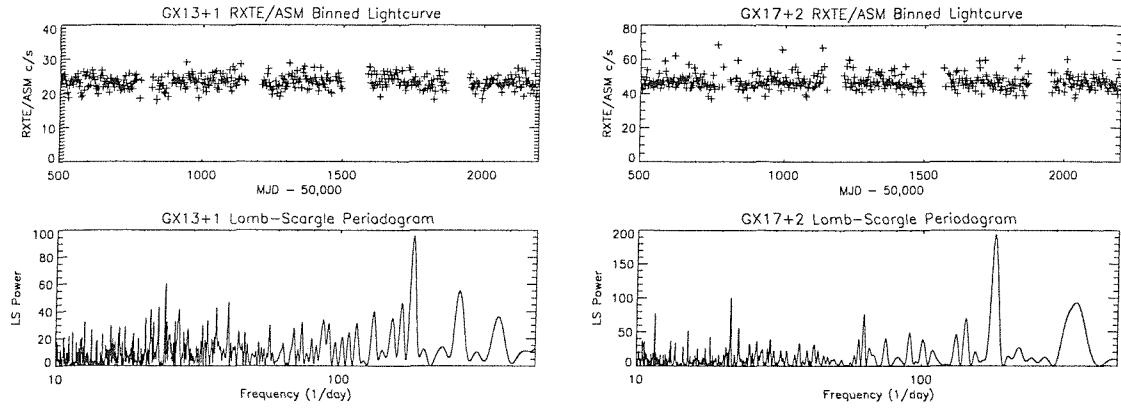
The correspondence between variation amplitude and LS power may be estimated from comparison to the classical DFT: for a given signal amplitude and frequency, powers  $P_{X,LS}$  are higher by only  $\sim 10\%$  than powers  $P_{X,DFT}$  (Scargle 1982), so that  $P_N(\omega) \simeq 1.1 P_{X,DFT}/\sigma_0^2$ . From the definition of the periodogram (3.2), an ideal sinusoidal variation with amplitude  $A$  and frequency  $\omega_1$  produces spectral power  $P_{X,DFT}(\omega_1) = A^2/N_0$ . This then yields the relation

$$A(\omega_1) \simeq \sigma_0 (N_0 P_N(\omega_1))^{1/2} \quad (3.8)$$

from which an equivalent sinusoidal amplitude of variation may be estimated (Again we have limited ourselves to the case where  $\sigma_0$  is known in advance.)

Armed with a method that is able to handle uneven sampling, we may now take advantage of the properties of unevenly-sampled data to overcome two significant drawbacks of even sampling. The first is simply that even sampling imposes a maximum frequency that can be searched of  $f_{Ny} = 1/2\Delta$ , where  $\Delta$  is the data spacing; this is because no information is available in the data gap to discriminate between frequency  $f_{Ny}$  and  $nf_{Ny}$ . Therefore the maximum frequency that can be detected from evenly sampled data is  $f_{Ny}$ . The second, related problem is that Fourier power at  $f > f_{Ny}$  is spuriously moved into the range  $f < f_{Ny}$  through the process of *aliasing*: two signals with  $f_1$  and  $f_2$  will produce the same two points with separation  $\Delta$  if the signal frequencies differ by  $n/\Delta$  (e.g. Press et al 1992). Uneven sampling breaks the period degeneracy by sampling between the gaps; the result is that aliasing is significantly reduced (e.g. Scargle 1982). By the same token, uneven sampling allows frequencies higher than even-sampled  $f_{Ny}$  to be probed.

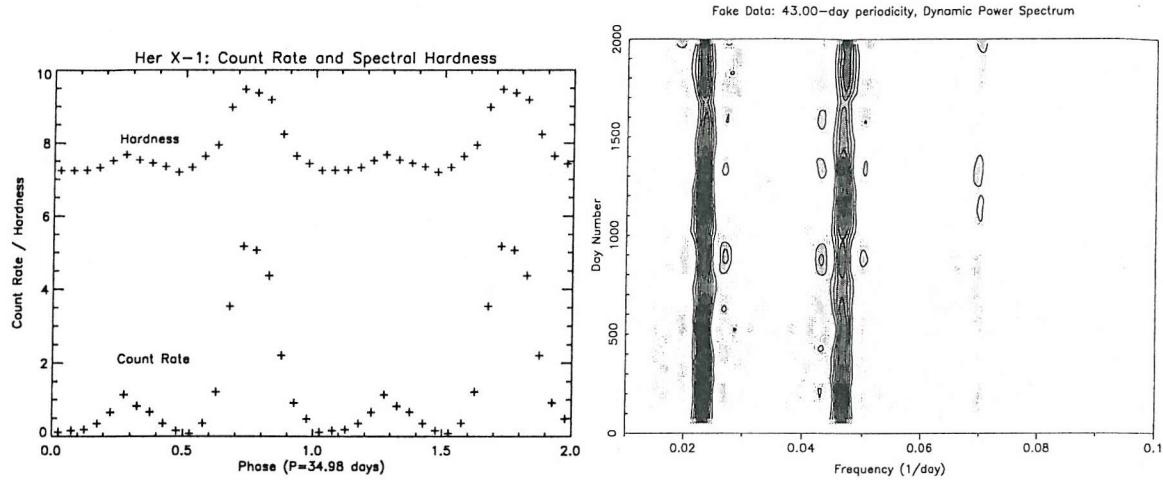
Care must still be taken when determining periodicities from this method, however, as spurious peaks can still arise, through three main mechanisms. The first is *spectral leakage* of power from the fundamental of some periodicity  $\omega_0$  to other frequencies. For example, sidelobes may appear separated by the reciprocal of the total observation interval. Secondly, if the data contains a periodic signal outside the frequency range of interest, power from this signal may leak into this region. In some cases the near-random sampling occurs over a true periodicity, which produces a strong peak in the periodogram. For example, sources near the galactic centre are unobservable from late November to late January, producing a significant spurious detection (figure 3.4). Finally, if the signal is significantly non-sinusoidal, the periodogram will produce several subordinate peaks associated with the fundamental.



**Figure 3.4:** RXTE/ASM lightcurves (top) and Lomb-Scargle periodograms (bottom) for the galactic bulge sources GX13+1 (Atoll source, donor unconstrained) and GX17+2 (Z-source LMXB). The recurrent sampling gap at 180d is clearly present in the lightcurves from both sources.

### 3.2.2 Dynamic Power Spectrum

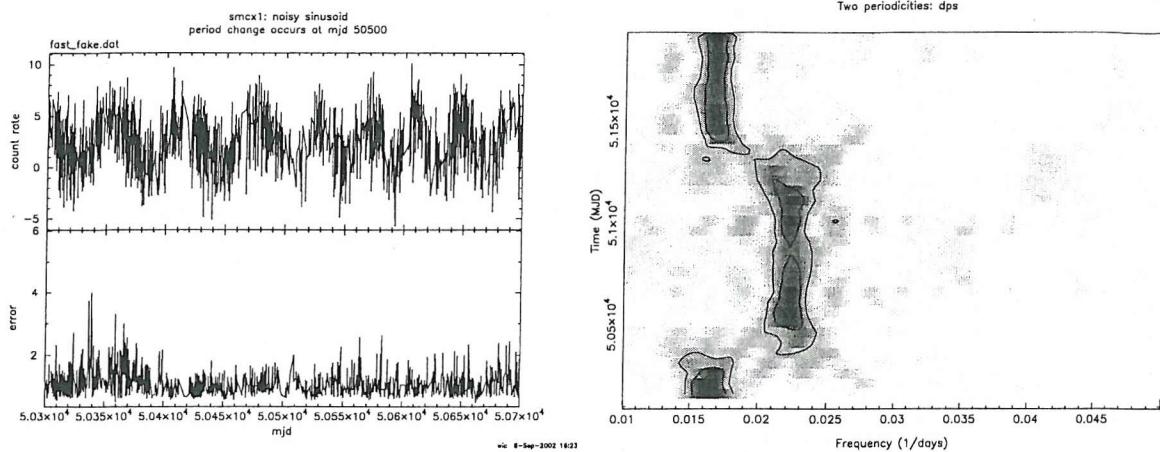
At  $\sim 9$  years and  $\gtrsim 7$  years respectively, the CGRO/BATSE and RXTE/ASM datasets have a long enough baseline that the evolution of periodic behaviour can be traced. This is accomplished by dividing the single long datatrain into subregions of user-specified length and evaluating the power spectrum for each region. The result is the *Dynamic Power Spectrum*, visualised by constructing 3-dimensional maps in (Time, Frequency, Power) space with spectral power being represented jointly by grayscale intensity and contours (figure 3.5). The grayscale represents the raw periodograms and the contours provides a smoothed interpolation over the data. The period ranges used in searches in this work ranges from 10-100 days in most instances, with a typical period spacing of 0.01d. The Lomb-Scargle periodogram algorithm (Scargle, 1982, 1989) implemented in Starlink software PERIOD was adapted in conjunction with a sliding ‘data window’ to produce power density spectra (PDS) for a series of overlapping stretches of the time series. Adjustable parameters in the analysis were the length of the data window and the amount of time by which the window was shifted each time to obtain the overlapping stretches of data. The code accounted for variations in the number of datapoints per interval such that the power spectrum resolution was identical for each interval. Choice of data window length was influenced by the expected minimum number of cycles needed for a significant detection of modulation from the point source in question. Of course, a longer window has the effect of reducing sensitivity to changes on timescales shorter than the window length. The amount of shift between consecutive intervals influences the degree to which changes can be followed, for example setting the shift to equal the window length, results in zero overlap and no ability to determine whether an observed change in the PDS was gradual or sudden. Conversely,



**Figure 3.5:** Illustration of the Dynamic Power Spectrum. **Left:** Fold profile of Her X-1, used as the template for a 43-day variation in the synthetic lightcurve. Also present in the synthetic lightcurve is frequency-independent white noise, and the sampling was taken from the true lightcurve of Her X-1. **Right:** Dynamic Power Spectrum of the synthetic lightcurve. The 43-day fundamental is most strongly detected. Also strongly present is a 21.5-day harmonic, due to the presence of a secondary as well as a primary peak in the fold (left). Also clearly visible is a weaker detection of a higher harmonic.

excessive overlap reduces the statistical independence of consecutive intervals. In any case there is no gain in sensitivity once the barrier of the periodogram resolution is approached, this being a significant factor when there are only a few signal cycles in the window.

When using a sliding window in this fashion, there are just under 6 fully independent 400 day data windows in the  $\sim 2500$  days of data. Thus the 50 day overlap between windows could smooth out abrupt variation, leading us to mistake an abrupt change in periodicity, or the replacement of one periodicity by an unrelated periodicity at a different frequency, for a single smoothly varying periodicity. We illustrate first that for well-separated periodicities, an abrupt jump is not confused as a smooth variation provided that the periodicities are sufficiently distinct. We estimate the confusion limit of such period jumps as follows: white-noise datasets are generated with representative sampling from example RXTE/ASM lightcurves. The fold profile of a known *non-sinusoidal* variation (in this case the superorbital periodicity of Her X-1) is then used to generate a variation on top of the noise of period  $P_1$ . The period of this variation is then switched abruptly by a trial amount to  $P_2$ , and kept there for sufficient length of time for the final data windows at this periodicity not to include any cycles of the original periodicity. The period is then shifted back to its original value (see figure 3.6 for an example). Period detections are then denoted as “confused” if the contour at the approximate FWHM of the peaks links the two period detections. The confusion limits thus reached are approximate, as the LS powers detected by the analysis are affected by spectral leakage from the sampling.



**Figure 3.6:** Left: Section of a simulated lightcurve with abrupt jump in the frequency of variation. The frequency of the input sinusoid switches instantaneously from 60 days to 43 days but maintains the same amplitude. The switch takes place at MJD 50500, and switches back at MJD 51300. Right: Dynamic Power Spectrum of this lightcurve. For an interval of four windows, no dominant periodicity is detected during the period switches.

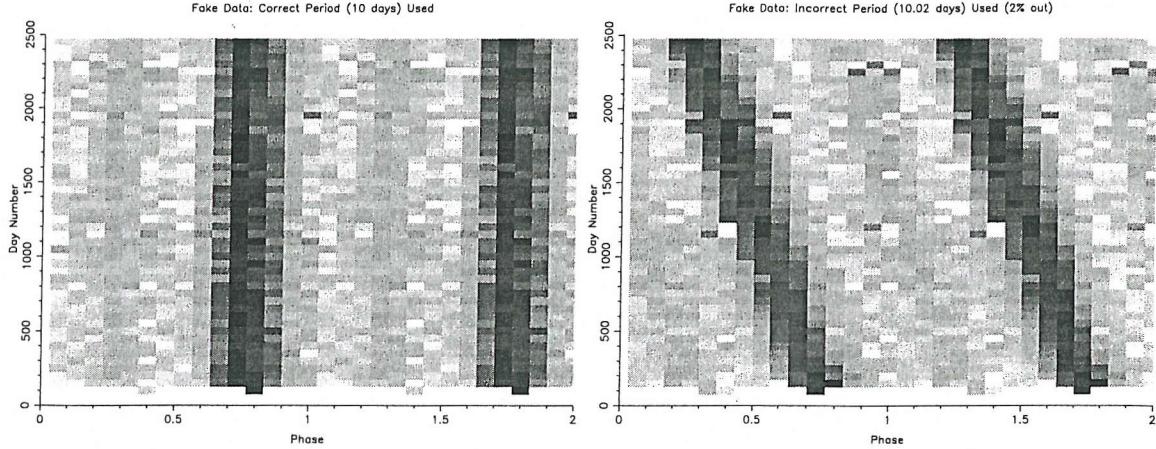
The above simulations were performed for input periodicities at 30-90 days, at 10-day intervals, with the second period introduced varied at 2-day intervals. Dynamic Power Spectra were then obtained of the resulting datasets for a range of analysis window lengths and overlap. Using 400-day windows, the degree of overlap was varied from 25 to 250 days, and the minimum detectable period shift  $\Delta P$  measured, where

$$\Delta P = \left(1 - \frac{P_2}{P_1}\right) \quad (3.9)$$

For all window spacings tested,  $(0.15 \lesssim \Delta P \lesssim 0.25)$ , with  $\Delta P$  increasing slightly with  $P_1$ , though it is not clear at present if this trend is an artifact of the complex interaction between power spectrum peaks. The process was then repeated for window overlap of 50 days and the window length varied from 100 to 800 days, in 100-d intervals.

### 3.2.3 Dynamic Light Curve

A counterpart technique to the Dynamic Power Spectrum, the Dynamic Lightcurve (DL) splits the filtered lightcurve into sub-regions of user specified length. This is useful in the case where a periodic or quasi-periodic cycle is present, and certain features (e.g. X-ray maxima) are expected to shift the phase in this cycle at which they occur. The lightcurves for each region are then visualised as a 2-dimensional grayscale map in (Time, Phase, Rate) space. Because this technique deals with raw count rates, the user has a virtually free hand in determining the length of each sub-region, limited only by the counting statistics for each phase bin. Specific regions of the cycle



**Figure 3.7:** Illustration of apparent period shift due to an incorrectly determined periodicity. Dataset used consists of white noise superimposed on the fold profile of Her X-1, with a fold repetition time of 10.00 days. **Left:** Dynamic Lightcurve of the fake dataset, using the input period of 10.00 days as the cycle length. **Right:** The same dataset, this time folded on a period of 10.02 days, i.e. 2% discrepant from the true periodicity in the dataset.

may be plotted, and the software producing the plot contains a front-end that allows filtering by count rate, should a specific feature of the lightcurve (e.g. x-ray maxima) be desired.

Due to the binning method used, the current version of the software producing the dynamic lightcurve does not correctly take into account gaps in the data, from cycles during which no desired events are present. For the work described in this thesis, the workaround used was to input a “clock pulse” of events with a recurrence period identical to the sub-region length, but with a count rate such that the clock pulses would not be present in the dynamic lightcurve. An example dynamic lightcurve is plotted in figure 3.7. The data used in this figure is a synthetic lightcurve composed of a 10.00-day periodicity (with the double-peaked shape of the Her X-1 longterm lightcurve; see Chapter 6), on top of frequency-independent white Noise (section 3.3.), with the random sampling of a real RXTE/ASM dataset (in this case Her X-1). As can be seen on the right of figure 3.7, if the fold period is identical to the periodicity in the dataset, the result is a line parallel to the time axis. If, however, the fold period is discrepant from the intrinsic periodicity, the structure is more complex. This situation may occur if the intrinsic behaviour evolves over the  $\gtrsim 7$ -year datasets (Chapter 8), or alternatively if the behaviour is stable but the period has been incorrectly chosen.

### 3.2.4 Period Detection beyond the Power Spectrum

The power spectrum functions on the principle that any periodic function satisfying the Dirichlet Conditions (Press et al., 1992) can be made up of a combination of sinusoidal components. The weighting of these components does of course depend on the shape of the variation: for a sinusoidal signal over zero noise, clearly only one component is needed at the frequency of the sinusoid, with all other components weighted to zero. However, if the signal is significantly non-sinusoidal, several components are needed to reproduce it; the classic example is the ideal square wave, which requires an infinite number of components to fully reproduce the instantaneous state transitions. For example, if the fundamental of some highly non-sinusoidal variation is outside the limits of the period search, then the extra components needed to describe the variation may still be within the period search, producing spurious detections. For this reason (and others), peaks in power spectra should never be used to claim periodicities without some other check on whether the variation is a real effect; for the RXTE/ASM, for example, the very least that should be done is to fold the lightcurve on the detected periodicity and also check the spectral behaviour, to determine if the variation is physically realistic for the system of interest. In this work we employ two further methods of period determination that are not strongly misled by non-ideal shapes of variation, Phase Dispersion Minimisation, in which the scatter in count rate is minimised to fit, and our own implementation of Phase Scatter Minimisation, in which the scatter in phase of salient features is desired. The former is useful if the count rate is expected to vary in a regular, highly non-sinusoidal way with time, while the latter is useful if there are few datapoints, and no information is known about the expected count-rate variation, or the sampling is too poor to determine it.

### 3.2.5 Phase Dispersion Minimisation (PDM)

Phase Dispersion Minimisation (PDM) detects periodicities by folding the data on trial periods, producing a measure of the scatter about the average lightcurve at each fold. The detected period is then the value for which the scatter in the fold is a minimum. The measure of the scatter is the *PDM Statistic*

$$\Phi = \frac{s^2}{\sigma^2} \quad (3.10)$$

where

$$\sigma^2 = \frac{\sum_j^{N_0} (X_j - \bar{X})^2}{N_0 - 1} \quad (3.11)$$

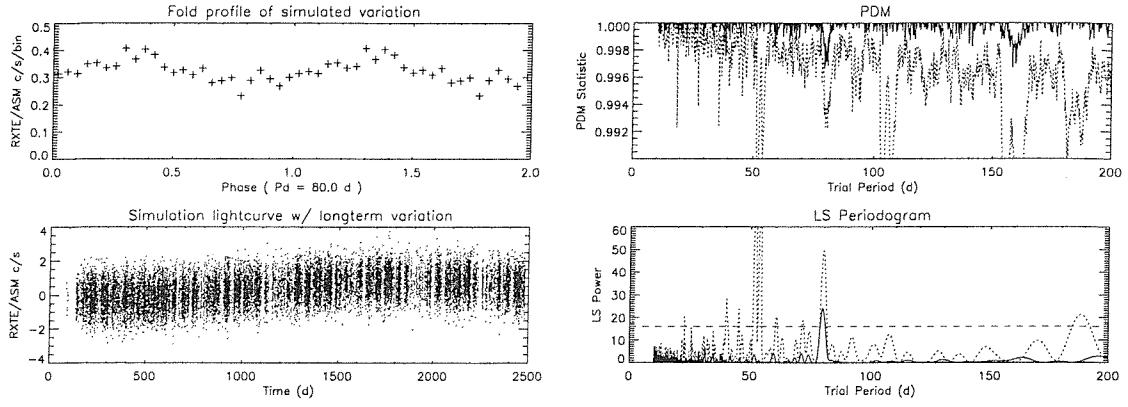
is the variance of  $X$ . Each trial fold produces  $M$  samples of this dataset, each with its own variance  $s_j^2$  and number of points  $n_j$ . The total variance of the trial fold is then given by (Stellingwerf, 1978):

$$s^2 = \frac{\sum_j^M (n_j - 1)s_j^2}{\sum_j^{N_0} n_j - M} \quad (3.12)$$

On an incorrect trial period, we expect  $\Phi \simeq 1$ , as the fold gives no reduction in scatter. As the trial period approaches the fundamental of real variation,  $\Phi$  reaches a local minimum.

Any periods found by this method do not depend on the form of the underlying modulation. However, as the method works only by minimising the scatter of the fold, it is extremely susceptible to harmonics of its own - a fold at twice the period of the fundamental also produces a PDM statistic much lower than nearby trials. Furthermore, as the existence of harmonics does not depend on the shape of the lightcurve, the harmonic structure of PDM periodograms does not carry information as to the shape of the variation itself (see figure 1 of Stellingwerf 1978). Setting confidence levels in periodicities determined with PDM techniques is something of an involved process; the confidence levels appear to depend on the shape of the variation in a complex way. For our purposes, the key point to note about the PDM is that its harmonics appear from a different mechanism than those in the power spectrum. Thus, peaks occurring in both the power spectrum and the PDM are likely to be real variations in the signal and not noise or harmonic structure from some other periodicity. The PDM is thus an invaluable check on any period detections in the power spectrum, which really comes into its own when the power spectrum gives several peaks, the strongest of which is at borderline (i.e.  $\lesssim 99\%$ ) statistical significance.

We illustrate the appearance of harmonics and spurious peaks through spectral leakage by running LS and PDM periodograms on a simulated noisy dataset. We begin with a white-noise lightcurve, with time values taken from a well-sampled RXTE/ASM lightcurve (in this case the Crab Nebula). We then superimpose a sine wave with an 80.0-d period, with variation amplitude chosen to produce peak LS power  $\sim 20$ . This detection power is well clear of white-noise significance limits (see below). The LS and PDM periodograms of this variation + white noise is computed. Subsequently, a sinusoidal variation at 3000.0 days is superimposed (the Galactic Bulge burster GX3+1 shows this type of variability). Evaluation of the periodograms of the new lightcurve show the effect of the longer-term variation on the previously strongly detected periodicity. The 80.0d periodicity is still strongly detected, but is no longer the strongest peak; indeed there is significant leakage of power



**Figure 3.8:** LS and PDM periodograms of a simulated 80.0-day variation. **Left:** Fold profile of the simulated white-noise + 80.0d variation (top), and the lightcurve with the 3000-d sinusoid added over this periodicity (bottom). **Right:** PDM (top) and LS (bottom) periodograms of the entire simulated datasets. The solid line denotes the 80-d variation over white noise, the dotted line includes the 3000-d variation. The horizontal dashed line corresponds to  $> 99.7\%$  significance over white-noise, from Monte-Carlo simulations of noise simulations lasting 2500-d each.

from lower frequencies (see figure 3.8).

### 3.2.6 Phase Scatter Minimisation

In several circumstances it is necessary to fit an ephemeris to periodic behaviour in the lightcurves of astronomical sources. For small deviations from a strictly stable periodicity, the  $N$ th occurrence of some regular feature is given by the Taylor expansion

$$MJD_N = MJD_0 + PN + \frac{1}{2}P\dot{P}N^2 + \frac{1}{6}P\ddot{P}N^3 \dots \quad (3.13)$$

(valid for all  $N$ ), where  $MJD_0$  is the time of some reference event,  $P$  the recurrence period and  $\frac{1}{2}P\dot{P}$  (hereafter denoted by  $C$ ) is the quadratic correction representing the period evolution. Higher terms in the ephemeris are assumed to be negligible for the sources examined in this thesis: the ephemerides in this thesis can all be adequately fit without appealing to higher than quadratic terms. As the above section illustrates, the phase-folded lightcurve can be a sensitive discriminant between ephemerides. In at least one bright XRB (Circinus X-1: Chapter 8), the lightcurve is dominated by extremely variable and erratic features, but there are regular X-ray dips, thought to be orbital in origin. In this case the bulk of the lightcurve is only quasi-periodic, and the section expected to be periodic is markedly non-sinusoidal. In this circumstance the LS power spectrum fails to determine an appropriate ephemeris, for two reasons: (i) the sharp deviation from sinusoidal variation leads to a combination of periods detected as the power spectrum reconstructs the true

variation from harmonics, and (ii) the lightcurve is dominated by the erratic variability throughout the majority of the cycle, which leads to large error in the determined period. If the data is selected such that only the periodic feature of interest is retained, the power spectrum fails because the dataset has been so reduced that the strongest detected period is uncertain. Furthermore, in the often-encountered case in which the period itself evolves with time, the LS periodogram fails if the period evolution is dramatic enough to cause the period to change over the observation interval (Chapter 8), and there are too few events to allow accuracy in period determination for the Dynamic Power Spectrum.

When fitting an ephemeris to a dataset, then, the LS power spectrum is only the first approximation. Having obtained a trial value from this method, we then employ the more direct technique of phase minimisation to fit the true ephemeris. The lightcurve feature expected to be the best system clock is isolated into a new dataset. This dataset is then folded with a range of trial ephemerides, and the range of phases at which the feature appears compared between the trials. The best approximation can then be read off and the search repeated with a finer grid of trials. This process can be repeated with finer and finer grids until the limit is reached at which the ephemerides cease to differ in the goodness of their fit to the data. Once the optimal combination of  $P$  and  $C$  has been found, the best value of the offset  $MJD_0$  can then be determined. Unlike the periodogram, for which the width of a peak carries information as to the degree to which the data is truly periodic, the relationship between the shape of the phase minimisation profile and the data periodicity is somewhat less certain, mainly because of two factors: (1) the selection of datapoints for this technique usually leaves a small number of points ( $\sim 200$ ), therefore the relation of the peak detection to nearby detections does not vary smoothly with trial period, and (2) for a quadratic ephemeris, both  $P$  and  $C$  are being fit to the data ( $MJD_0$  being a constant offset), where  $P$  and  $C$  are not independent but relate through the phase sampling of the data - which of course is different for each trial period selected. Instead, the error on the ephemeris components is determined from Monte Carlo simulations (next section). Note that this period detection method is entirely insensitive to the count rate, but instead focuses exclusively on the timing behaviour of the lightcurve.

### 3.3 Significance Testing and Determination of Uncertainties

#### 3.3.1 Lomb-Scargle Periodogram: Statistical Significance

Astronomical time series consist of some signal in which one is interested, superimposed onto other signals perhaps not of interest, along with noise of some description. If the majority of the signals are not of interest, the signal of interest must be separated from the rest of the dataset, which often leads to the selection out of so much data that the remaining dataset is not well enough sampled for Fourier techniques to be useful (section 3.2.4). For many of the sources examined in this work, the feature of interest is a long-term periodicity, and dominates the lightcurve to such an extent that power spectrum analysis of the total lightcurve is not significantly contaminated from other lightcurve features. Furthermore, for SMC X-1, Her X-1, LMC X-4 and M15, the underlying noise is well-described by white noise in which the amplitude of the noise contribution is independent of frequency (frequently called Gaussian white-noise, as the noise time series is a random deviate with gaussian probability distribution). This is determined by plotting the LS power spectrum, binned coarsely so as to smooth out signal features and bring out the underlying noise. A flat binned power spectrum is indicative of frequency-independent white-noise.

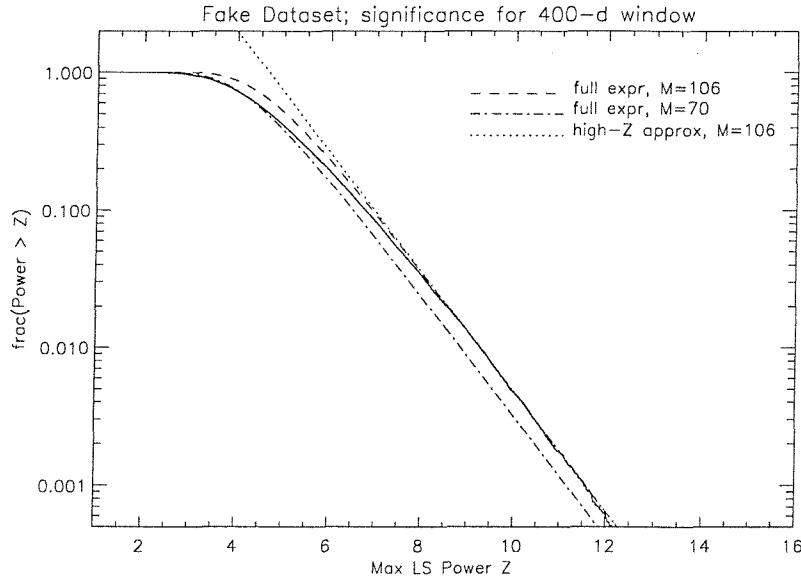
We introduce here the setting of confidence levels for the even-sampled traditional DFT and then discuss its extension to uneven sampling. Under the assumption of white-noise, the statistical significance of detected periodicities with *even* sampling depends exponentially on the power itself (equation 3.4). If the power spectrum is evaluated only at a set of  $M$  independent frequencies (more precisely a set of  $M$  frequencies such that the corresponding powers  $P_X(\omega)$  are independent) then the *False-Alarm Probability*, i.e. the probability that a power  $z$  or higher would result from white-noise alone, is given by

$$\text{Prob}_{FA}(z) = 1 - [1 - e^{-z}]^M \quad (3.14)$$

In practice, power spectra are usually strongly oversampled, to produce power at the frequencies of interest rather than just the  $M$  classically independent frequencies. In the limit of high  $z$  (i.e.  $e^{-z} \ll 1$ ), the False Alarm Probability is expected to be given by

$$\text{Prob}_{FA}(z) \simeq M e^{-z} \quad (3.15)$$

In the case of unevenly sampled data, this procedure breaks down in two ways: (i) the value of  $M$  is difficult to estimate and even interpret (Cumming et al., 1999), and more importantly (ii) the exponential probability distribution breaks down through uncertainty in determining  $\sigma_0$



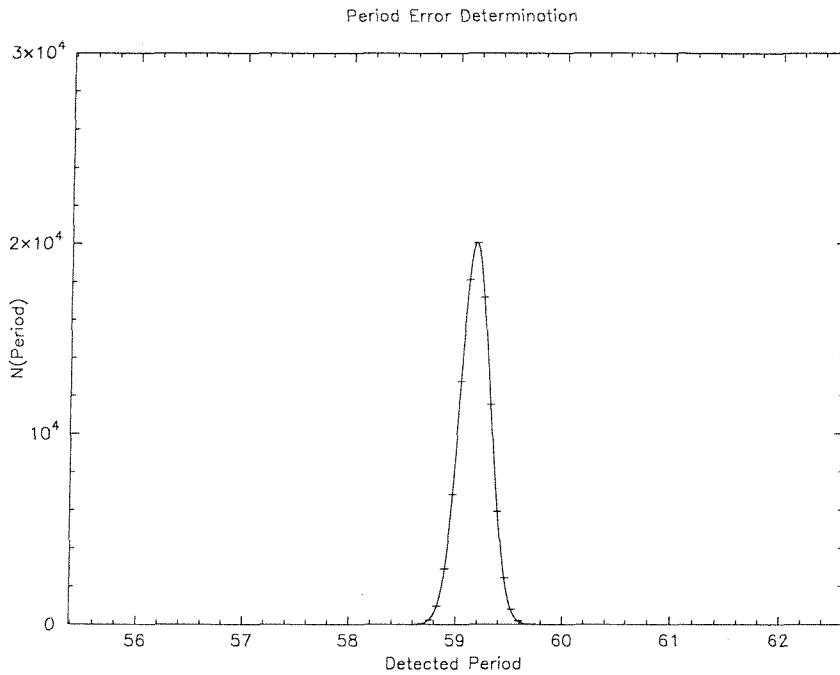
**Figure 3.9:** Empirically determined CDF for 100,000 synthetic, white-noise lightcurves lasting 400 days each. Overplotted is are the fits using the exponential approximation (3.14) and the high- $z$  approximation (3.15).

(Schwarzenberg-Czerny, 1998, 1999). The significance of detected peaks must thus be determined empirically. Synthetic lightcurves are created, using a gaussian random number generator to calculate the count rates, with standard deviation equal to that of the original dataset. The time value for each datapoint is offset from the value that would be expected from even sampling on the average sampling rate, with the offset calculated as a uniform random deviate, with allowed range equal to the true range of deviations from even sampling. For the filtered RXTE/ASM and CGRO/BATSE lightcurves, coverage is continuous enough for this to represent the statistical behaviour of the data. (If data is selected so that only a certain lightcurve feature is followed, the large data gaps that result must be taken into account; see section 3.3.2 for more information.) For each synthetic lightcurve, the LS power spectrum is computed, and the peak LS power recorded. This process is repeated for a large number of datasets (100,000 synthetic datasets were used).  $\text{Prob}_{FA}(z)$  is then measured from the CDF as the fraction of synthetic datasets for which the maximum detected LS power is  $z$  or higher. In figure 3.9 an example synthetic CDF is plotted. The synthetic lightcurve used was based on a 400-day stretch of data from SMC X-1, with frequencies between  $0.01\text{d}^{-1}$  and  $0.1\text{d}^{-1}$  searched at an interval of  $10^{-4}\text{d}^{-1}$ , representing accurately the analysis corresponding to a single window in the dynamic power spectrum. We attempt to fit  $\text{Prob}_{FA}(z)$  with its expressions (3.14) and (3.15; see figure 3.9), and find that indeed the distribution does not quite follow the exponential prescription (3.4) for all powers  $z$ . For high  $z$  ( $\gtrsim 10$ ), however, an equivalent  $M \sim 106$  produces a good approximation. The reader is cautioned that

the exponential assumption has propagated through the literature; see for example the otherwise excellent handbook of Wall & Jenkins (2003). We notice immediately that this figure is markedly less than the 900 grid frequencies and also far lower than the number of datapoints ( $N \sim 6000$ ) in the 400-d window; this contrasts sharply with Horne & Baliunas (1986), who assert  $M \sim N$ . We also note immediately from the simulations that a false-alarm probability of 0.1% corresponds to a LS power of 11. This figure is slightly lower than the  $\sim 13$  usually quoted in studies with the LS power spectrum, and is a result of the much shorter time base in each data window, and the resulting lower number of independent frequencies (which scales with the total observation interval).

### 3.3.2 LS Periodogram: Uncertainty in the Detected Period

The low number of statistically independent frequencies leads us to determine the statistical spread in the period detection empirically. The method is similar in principle to the significance determination: LS power spectra are computed for a large number (100,000) of synthetic datasets. This time the datasets contain a user-input periodicity above the white noise. To correctly take into account the behaviour of the power spectrum resulting from deviations from sinusoidal variation, the lightcurve is folded on the determined period and this fold is used as a template for the variation. If this variation is obvious in the filtered lightcurve, the scatter about this variation is estimated and used as the standard deviation for the white-noise model; otherwise (in the limit of small-amplitude variation) the standard deviation of the dataset is again used. The sampling is generated using a linear random deviate generator as for the significance test. For each dataset the frequency corresponding to the peak detection is recorded, allowing a histogram to be built up of the fraction of trials for which the peak LS power occurs at frequency  $\nu$ . This histogram is well-modeled by a gaussian, allowing the uncertainty in period to be expressed in units of  $\sigma$  (figure 3.10). A  $3\sigma$  Period uncertainty of  $\pm 0.7$ d is typical. Curiously, in some cases the peak detection can be shifted from the input period of variation by  $\sim 0.5$ d; this is likely due to spectral leakage to or from other periods due to the non-sinusoidal shape of variation. The variation of this small shift from source to source leads us to repeat the simulations for each source, with the shift incorporated into the figure for the period uncertainty for each source.



**Figure 3.10:** Period histogram for 100,000 synthetic datasets, with a 60-d periodicity input over white-noise. The shift in peak detection is likely to be due to spectral leakage to or from other periods due to the non-sinusoidal nature of the variation.

### 3.3.3 Red Noise

In the case of Cyg X-2 and Cir X-1, both of which are active and X-ray bright sources (Chapters 6 & 8), the binned power spectrum is not flat but instead can be well fit by a power law of index  $\sim 1.3$  - such a contribution to the continuum is termed *red-noise*. This power-law variability is reminiscent of that exhibited by XRB at frequencies  $\gtrsim 0.1\text{Hz}$ , and by the AGN at similar timescales. The same detector was used to record the lightcurves of the white-noise and the red-noise sources, and furthermore there does not appear to be any physical property shared by Cyg X-2, Cir X-1 and the prototypical red-noise XRB 4U1916-053 (Homer et al 2001), which is not shared by the white-noise binaries. It thus appears that the red-noise contribution is not noise in the same sense as the white-noise, but instead is due to confusion from the intrinsic variation in the lightcurve, which contributes LS power preferentially at lower frequencies. The result of red-noise is that the above procedure of setting significance levels fails when white-noise models are used, as the significance contours are no longer straight lines. For peaks in the power spectrum for which no *a priori* knowledge of variation exists, each dataset must be modeled with the best-fit red-noise power spectrum, leading to a separate CDF calculation for each frequency bin of interest. This is a laborious process (see e.g. Homer et al 2001), which we can sidestep for the red-noise sources examined in this work (Chapters 6 & 8) by noting that for both Cyg X-2 and Cir X-1, the statistical

significance of the variation is well established, and indeed our interest is primarily in tracing the evolution of observed variation.

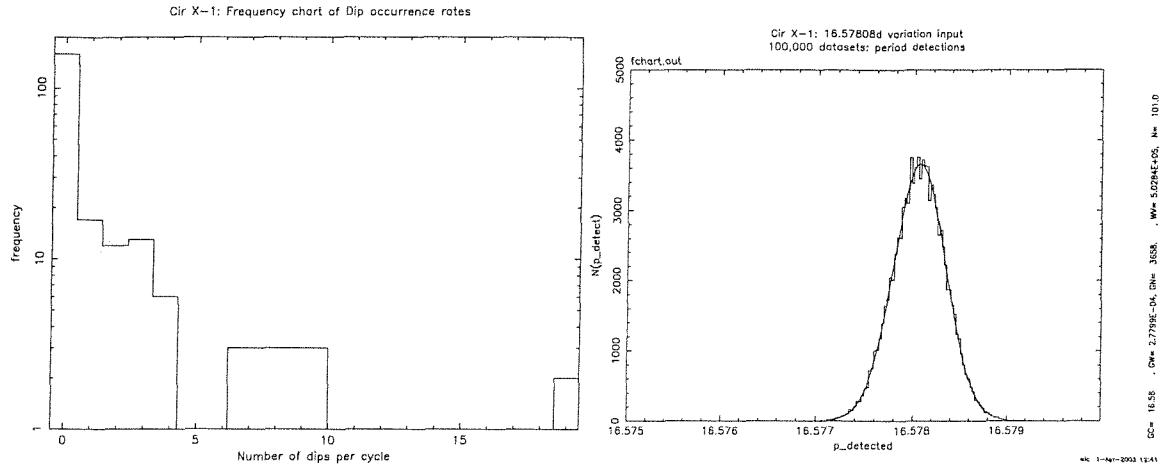
### 3.3.4 Phase Minimisation

In Chapter 8 we fit a quadratic ephemeris to the  $\sim 400$  X-ray dip events in the RXTE/ASM lightcurve of Cir X-1. Determination of the uncertainty in the detected ephemeris components proceeds in a broadly similar manner to the Monte Carlo simulations for LS methods, with two important differences. Firstly, the coverage with RXTE/ASM of dips is sparse enough to leave large gaps in data coverage. For the majority of the 16.6-day cycles, no dips are sampled by the RXTE/ASM (figure 3.11). The addition of offsets to uniformly spaced data is therefore not valid for this situation; instead, the 16.6d cycles are populated with dips, where the number of dips per cycle is random but approximates the measured sparsity of dips. Secondly, the intrinsic occurrence rate of the dips is a strong function of phase in the 16.6 day cycle. When determining the effect of instrumental sampling, the form of this function fit to first approximation by a Gaussian. The random deviate used as the time of dip measurement can then be approximated with the standard gaussian-distributed random deviate generator GASDEV (see Press et al 1992 for more details).

A large number (100,000) of synthetic datasets are thus built up in such a way as to preserve (i) the observed distribution of dip detection frequency and (ii) the intrinsic dip occurrence rate, on a user-defined periodicity. The uncertainty in each of the detected ephemeris components is then read from histograms of the number of detections of each trial value. As with the LS power spectrum determinations, the histogram is well fit with a gaussian, allowing errors to be expressed in terms of the standard deviation  $\sigma$ . As with the determination of the ephemeris components themselves, considerable computation time can be saved by noting that  $MJD_0$  is just a constant offset, while  $C$  is best fit for each period  $P$ .

### 3.3.5 Non-Gaussian $p(x)$

The above procedure models the sampling of the RXTE/ASM to adequate precision when determining the uncertainty in determined periods. However, when attempting to fit the observed dip occurrence rate with model predictions (Chapter 8), the model predictions are not exactly gaussian but follow a more complex profile. In this case, when determining significances for the



**Figure 3.11:** Left: Frequency chart of dip sampling for the RXTE/ASM lightcurve of Cir X-1. Note that in the majority of cycles, no dip is sampled. Right: Sample histogram of period detection trials for the phase minimisation technique.

model fits, the gaussian random number generator GASDEV must be replaced with one that uses a user-defined probability distribution  $p(x)$ . This is accomplished using the *rejection method* of random number generation. I summarise the method here; the interested reader is referred to Press & Teukolsky for more details. The great advantage of this method is that  $p(x)$  need not be continuous or even computable, but may be input as a table of numbers. This allows us to import the observed dip occurrence profile straight from the RXTE/ASM lightcurve. Consider some random deviate described by the probability distribution  $p(x)$ , so that the probability of retrieving a value between some  $x_1$  and  $x_2$  is given by  $\int_{x_1}^{x_2} p(x)dx$ . The rejection method generates random numbers uniform in  $\int p(x)dx$ , thus creating a random dataset following the desired probability distribution. A comparison function  $f(x)$  is chosen, which has finite area and lies everywhere above  $p(x)$ , for which the indefinite integral is known analytically. Furthermore this comparison function must be analytically invertible to give the value  $x$  as a function of  $F(x) = \int f(x)dx$ . A uniform random deviate is then chosen for  $F$ , which gives a corresponding value of  $x$ . A second uniform random deviate  $y$  is then created, resulting in a two-dimensional point  $(x, y)$ . The set of points  $(x, y)$  is thus uniformly distributed in the area under the curve  $f(x)$ . Finally, this point is rejected if it lies above the  $p(x)$  curve. The  $x$  component of this point is then uniformly distributed in area under  $p(x)$ , and thus uses  $p(x)$  as its probability distribution. To illustrate this method, we plot in figure 3.12 an example random number set using a model prediction of the dip occurrence profile, with coefficients fit to match the observed profile (Chapter 8).

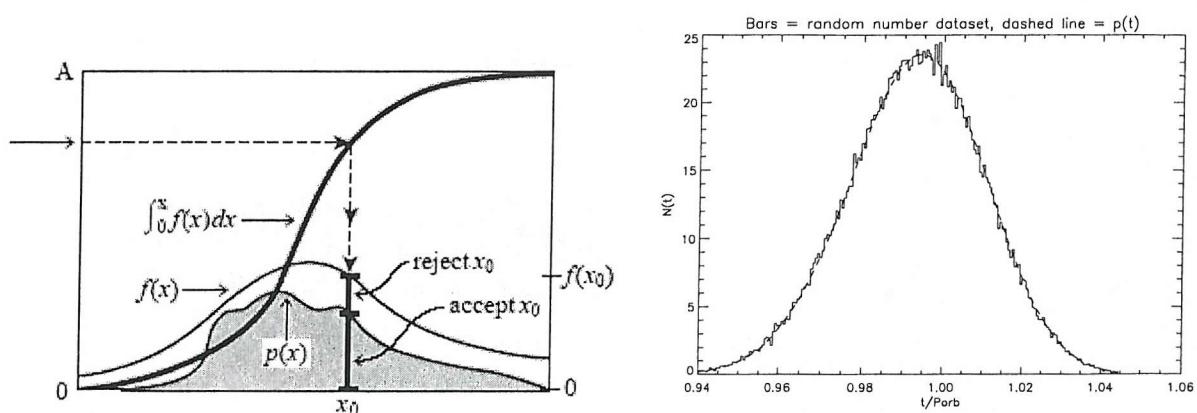


Figure 3.12: Left: Illustration of the rejection method (from Press et al 1992). See text for details. Right: Frequency chart of random number dataset generated with the rejection method, using an example non-gaussian function  $p(x)$  as probability distribution function.

## Chapter 4

# X-ray Variability in the Core of M15

The prototypical post-core-collapse globular cluster M15 provides an ideal environment for the formation of exotic binaries, it being already known to contain a luminous LMXB (X2127+119, optically identified with the 17.1 hour eclipsing binary AC211) and six millisecond pulsars in the core. However, the X-ray properties of X2127+119 are strange in that it appears to be a high inclination accretion disc corona source (ADC) from which we only see scattered X-radiation, and yet it has produced extremely luminous type I X-ray bursts, which may or may not have come from AC211. We have therefore examined the  $\sim$ 5-year RXTE/ASM light curve of this object in order to search for long-term modulations which might shed light on this unusual behaviour. Furthermore we have used archival HST UV images of the M15 core to search for other variable objects which might indicate the presence of a second LMXB. From these we have found one highly variable ( $> 5$  mag) object, which we interpret as a dwarf nova. Work in this chapter has been published in Charles, Clarkson & van Zyl (2002).

### 4.1 Introduction

The dense cores of globular clusters are now well established as breeding grounds for the creation of exotic interacting binaries (e.g. Verbunt & van den Heuvel 1995). In particular, it has been

noted for almost 20 years that luminous low-mass X-ray binaries (LMXBs) are found in  $\sim$ 12 globular clusters, an overdensity relative to the field of a factor 100, indicating a greatly enhanced formation mechanism in the dense cluster environment (both 2-body (tidal) and 3-body (exchange) mechanisms are still actively considered, see e.g. (Elson et al., 1987) and references therein). And the descendants of LMXBs, the millisecond radio pulsars (MSPs), are also found in substantial numbers in clusters, with 8 of them in M15 alone (of which 6 are within 8 arcsecs of the core; see e.g. Lyne et al. 1996). Furthermore these LMXBs have a range of properties more extreme than those in the field. All produce type I X-ray bursts (e.g. Charles, 1989, and references therein), and hence are accreting neutron stars, but their period distribution is skewed to include 2 (possibly 3) extremely short period binaries which require both members of the binary to be degenerate (Rappaport et al., 1987). Recent UV surveys with HST (de Marchi & Paresce, 1994, 1995, 1996) have revealed a population of very blue stars (VBS) in M15, the nature of which is as yet undetermined, but their very high density in the core indicates a population segregation that must be related to other properties of the cluster.

Against this background, M15 houses one of the more enigmatic of these high luminosity LMXBs, X2127+119 in that it was identified 18 years ago (Charles et al., 1986) with by far the most luminous optical counterpart (AC211) of any galactic LMXB. This 15th magnitude extremely blue, highly variable ( $>1$  mag in U) object was photometrically monitored by (Ilovaisky et al., 1993) who showed that it was eclipsing with a 17.11 hour period. The broad eclipse feature (seen in both optical and X-rays) suggested that X2127+119/AC211 belonged to the accretion disc corona (ADC) class, in which the high inclination means that the compact object is obscured by the disc rim and hence not directly visible. This accounts for the high optical luminosity associated with a relatively faint X-ray source.

However, this model suffered a severe blow when (Dotani et al., 1990) and (van Paradijs et al., 1990) observed M15 with *Ginga* and discovered an extremely luminous type I X-ray burst, thereby inferring that the X-ray source must be directly visible (at least for part of the time). In which case the persistent, low apparent X-ray luminosity ( $3 \times 10^{36}$  erg s $^{-1}$  for an assumed distance of 10.5 kpc; Djorgovski 1993) was then attributed to a genuinely low accretion rate ( $\sim 0.01 L_{Edd}$ ), rather than being the result of obscuration in an ADC, and indeed there are other comparable luminosity directly-viewed LMXBs in globular clusters (Sidoli et al., 2001). However, these are all optically much fainter and hence the remarkable optical brightness of AC211 remains unexplained.

We therefore decided to use archival X-ray and UV data on M15 in order to investigate the strange

properties of X2127+11/AC211 in more detail. The on-line database of the RXTE ASM provides a unique, long-term ( $\sim 5$  yrs), continuous X-ray light curve which we could examine for behaviour that might be related to the occurrence of the bursting events (it already having been suggested by Corbet et al (1997) that there might be a 37 day modulation in earlier samples of the ASM data). Furthermore it was not clear to us that the very compact core of M15 was restricted to containing only one luminous LMXB, and so we extracted the HST archival images of the core region in order to search for variable UV sources, a study that might also shed light on the nature of the VBS.

## 4.2 RXTE/ASM Light Curve

The particulars of the RXTE/ASM dataset have been described in Chapter 2. X2127+119 is a weak source even for the ASM, averaging only  $\sim 1\text{cs}^{-1}$ , therefore particular attention must be paid to problems with the RXTE environmental background (see figure 4.1). As with the other RXTE/ASM sources in this thesis, we utilised the software provided by MIT (Levine et al 2000) which sets levels of background activity for rejection of data which are significantly better at removing contaminated data than the standard data product provided on-line. However, given the faintness of the source, the filtering criteria were necessarily stricter than for e.g. SMC X-1, for which, with twice the ASM count rate the signal is clear of the noise (chapter 5). Data were kept only if the fit had a reduced- $\chi^2$  value between 0.05 and 1, and the background at the time of the dwell was less than  $5\text{ cs}^{-1}$ . Finally, all points were rejected whose error was greater than three quarters of the data value itself. This produces a significant improvement in the data quality. To identify longterm trends in the ASM lightcurve of a weak source such as X2127+119, it is instructive to bin the data, and we chose a binning of five days per bin and required a minimum of five datapoints per bin. The resulting lightcurve and rms error are shown in figure 4.2. We see regions of extended activity at about day numbers 100 and 400, lasting about a month.

### 4.2.1 Period Search

Because the ASM data is unevenly sampled, it was necessary to replace the Fourier Transform with the Lomb-Scargle periodogram (LS) when searching for periodicities (Chapter 3). The LS periodogram is a Discrete Fourier Transform modified to take account of unevenly sampled data.

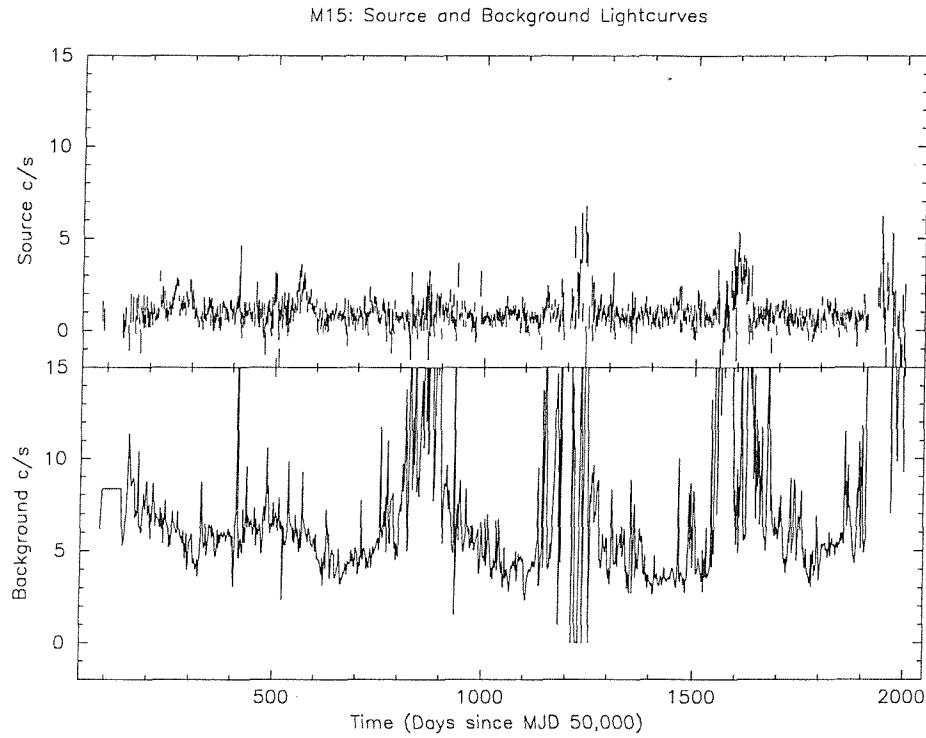


Figure 4.1: RXTE ASM lightcurve and background before cleaning.

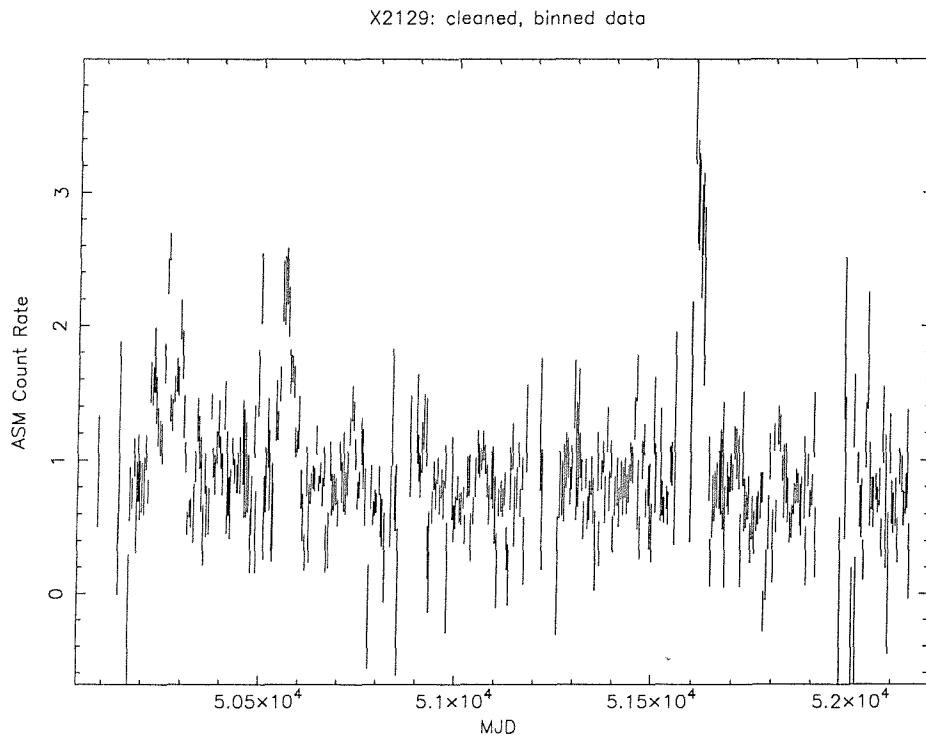
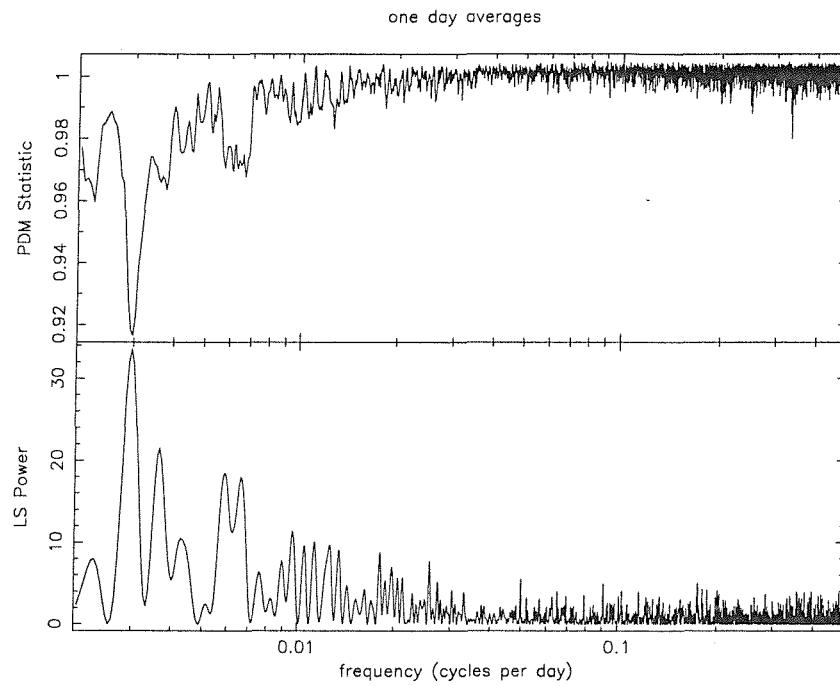


Figure 4.2: The RXTE ASM dataset, but with all points with background levels above 5 counts per second removed, all points with an error greater than 0.75 or a chisquared statistic greater than 1 removed. Data has been binned to 5 days per bin to highlight extended activity.



**Figure 4.3:** Periodograms over a range of periods from 2 to 1000 days. There is no sign of a 37d period.

Phase Dispersion Minimisation (PDM) detects periodicities by folding the data on trial periods, producing a measure of the scatter in the resulting fold. Because any periods found by the PDM do not depend on the form of the underlying modulation, it is a useful check on periodicities found with LS techniques and is used here to provide an independent measure of the period. The resulting periodograms are shown in figure 4.3. There is no sign from this analysis of the 37d period reported by Corbet et al (1997). By adding sinusoidal signals to our simulated light curves, a rough upper limit was set for the amplitude of a 37d periodic signal that might go undetected with the time sampling of the ASM dataset. This upper limit occurs at about 0.08 c/s, or roughly 8 percent of the average quiescent level. However, it must be borne in mind that this analysis was carried out for the full energy range of the ASM; whereas Corbet et al point out that this modulation is only prominent in ASM Channel 1. With this in mind, power spectra were computed for the individual channels (figure 4.4. No significant periodicity at or near 37 days is present.

Finally, a similar period analysis performed over a range of shorter periods confirms that the 17.1 hour orbital period continues to appear in the RXTE ASM dataset (see Homer & Charles 1998 for a detailed analysis of the first 2 years of ASM data).

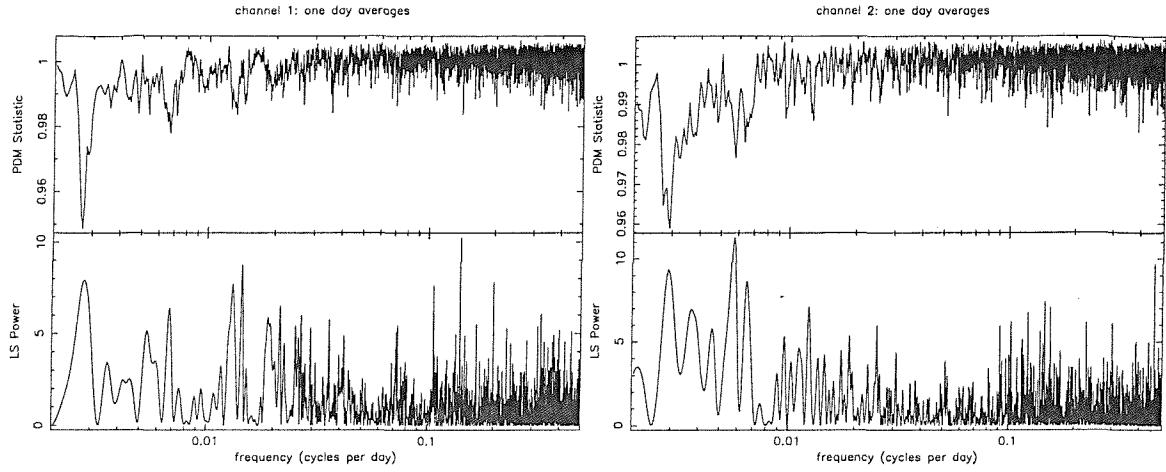


Figure 4.4: Period analysis of X2127+119 performed on the lower two channels.

#### 4.2.2 X-ray burst candidates

In order to determine the visibility of the neutron star, it was desired to search the ASM lightcurve for further Type-I X-ray bursts with similarities to the October 1988 event. The coverage of the RXTE ASM is such that a few individual dwells (of 90s each) are obtained each day, with an average recurrence timescale of  $\sim 3$  hours. This sampling rate is an order of magnitude longer than typical type I X-ray burst durations - the October 1988 burst, for example, lasted less than 200s. However, we might still expect to see a burst event in the RXTE ASM lightcurve, as a single, high datapoint. Because any burst picked up by the RXTE ASM will be measured for a window of 90s, at a random point in the burst profile, the detected level of a burst may well take a value significantly lower than the peak X-ray brightness of the event.

When searching for single high datapoints, the filtering requirements differ from those necessary to plot clean lightcurves for period searches. In particular, beyond simply selecting by background, the cleaning procedures detailed in Chapter 3 select the data based on minimising a least squares fit to the average behaviour (the reduced  $\chi^2$  fitting). While appropriate for most analyses of RXTE data, this tends to reject single high points - the target of this search for bursts! We therefore decided to refilter the raw data based on behaviour of the background levels; all points with a background level greater than 20 counts per second were rejected. A burst is then recognised if the burst profile differs from the background profile, whatever the absolute value of the background. All points with no background information were rejected.

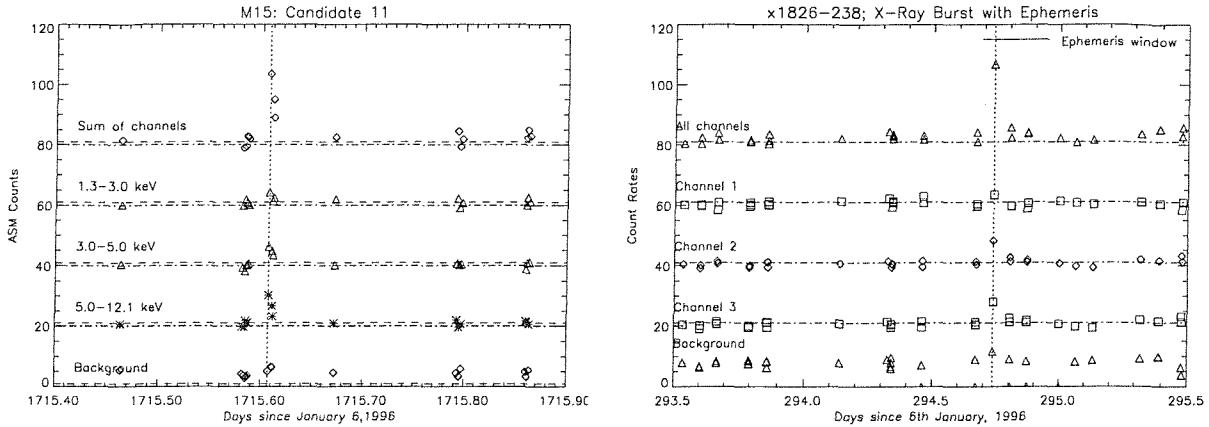
Examining the shape of the lightcurve before and after a burst candidate is useless when attempting

to discern bursts. At  $\sim 1\text{cs}^{-1}$ , the quiescent output of X2127+119 is so low that, with the exception of burst candidates, which with RXTE have  $\sim 5$  times the signal-to-noise ratio of the quiescent state, variations in the shape of the lightcurve are indistinguishable from noise. This is especially true of data for the individual channels, with values typically  $\sim 1/3\text{cs}^{-1}$ . Another measure must thus be found that uses only the burst candidate events themselves. The method chosen to examine candidates was to compare their approximate peak photon energies.

The classical type I X-ray burst profile consists of a fast rise to peak luminosity, (rise time typically  $\sim 1 - 10\text{s}$ ), followed by a more gradual decay to quiescence (decay time typically from  $\sim 10\text{s}$  to  $\sim 1$  minute). Decay times are somewhat energy dependent; for example, burst lightcurves at several energies show decay times  $\sim 20\text{s}$  at  $\gtrsim 6\text{keV}$  but  $\gtrsim 60\text{s}$  below this energy (Hoffman et al., 1980). While the fast rise is detected predominantly at energies  $> 5\text{keV}$ , the energy dependence of the cooling rate produces a softening of the burst profile during decay, ascribed to the cooling of the neutron star photosphere after the initial detonation (Lewin et al, 1995). This allows us to reject burst candidates with activity in one channel only. A further implication is that we should see a rough anticorrelation between total luminosity and hardness ratio (the ratio  $(3\text{-}12\text{ keV})/(1\text{-}3\text{keV})$ ; see Chapter 3) as the more luminous points would arise from earlier points in the profile with higher photon energies.

#### 4.2.3 The clockwork burster X1826-238: a comparison source

To demonstrate the validity of our search strategy, the X-ray burster X1826-238 was examined with the ASM. At  $\sim 2$  ASM c/s, its quiescent level is quite similar to X2127+119, making it a good comparison source, but more importantly, it is also a prolific burster. X1826-238 displays a remarkable near-periodicity in its burst behaviour, with seventy bursts observed by the BeppoSAX Wide Field Camera over 2.5 years. Analysis by (Ubertini et al., 1999) revealed a significant quasi-periodicity in the bursts of  $5.76 \pm 0.6\text{hr}$ . RXTE/PCA analysis of 24 further bursts shows the burst lightcurves to be remarkably uniform, making 1826-238 an ideal object to test burst models (Galloway et al. 2003). Moreover, the burst rate has been steadily increasing (recurrence time reduced by 50% in  $\sim 7$  years (Galloway et al., 2004). The quasi-steady burst rate of this source allows us to predict possible windows in which we might expect to see a burst. A burst candidate for X1826-238 is plotted from the ASM lightcurve in figure 4.5: as can be seen, it fits within the expected window very well. The energy distribution is weighted towards the higher energies, which



**Figure 4.5:** Burst candidates from X2127+119. The top plot represents the total ASM count rate, beneath this are the three ASM channels individually, in increasing order of hardness. At the bottom is the background behaviour. All plots are to the same scale. The broken lines on the plots represent 0 and 1 for each plot.

supports the standard model for type I bursts given that the candidate occurs early in the predicted burst window. If the method of searching for a rough luminosity-hardness ratio anticorrelation is valid, then X1826-238, with its dependable burst activity, should be the ideal calibration source. Plotting the ASM count rate against hardness ratio for RXTE/ASM burst candidates from this source does indeed show such a relation; the points with the highest count rate correspond to those with the hardest photon energies (see section 4.4).

### 4.3 The October 1988 burst and its implications for the ASM search

We must also consider profiles that are a little more varied than the classical type I profiles exhibited by X1826-238, as the October 1988 burst exhibited by X2127+119 was by no means typical. To begin with, after ten years, this burst is still one of the most luminous X-ray bursts ever seen, which at  $\sim 4.5 \times 10^{38}$  ergs $^{-1}$  is near-Eddington for this system. (By considering gravitational redshift at the neutron star surface, van Paradijs et al (1990) obtain a figure of  $\sim 5 \times 10^{38}$  ergs $^{-1}$  for the Eddington limit of this system). At  $\sim 150$ s, the burst was unusually long for a classical type I burst, and there was also a precursor burst  $\sim 6$ s before the main burst. Finally, the energy profile was unusual for a type I burst, with output rising throughout the first half of the profile, at successively harder energies, causing a *hardening* of the spectrum. These properties are thought to be hallmarks of photospheric radius expansion on the neutron star as a result of the burst (Lewin

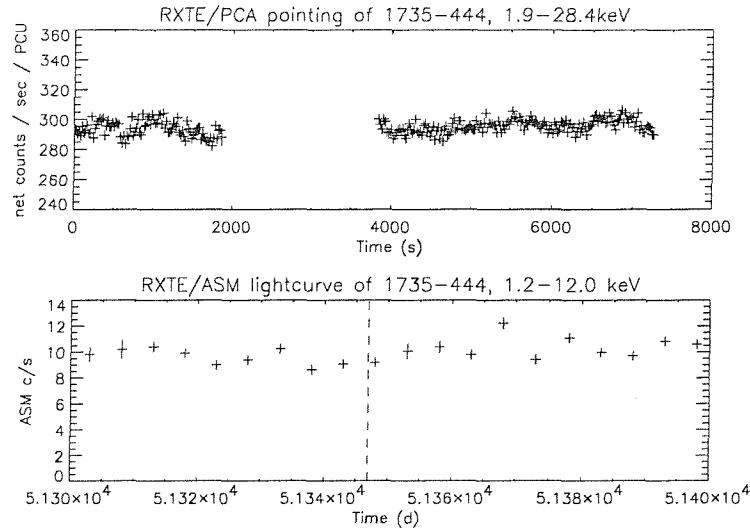
et al., 1995)

## 4.4 The X2127+119 burst candidates

Table 4.1 lists the candidate bursts that satisfy our search criteria. The time interval from day number 1500 to 1600 was rejected on the grounds that any true bursts will be indistinguishable from the extremely high background behaviour. An event was selected as a burst if its ASM count rate was above  $10 \text{ cs}^{-1}$  and the hardness ratio between 0.5 and 5. In each case the bursts take place in the absence of background events of similar scale, and all contain contributions from each energy channel.

To compare candidate burst fluxes with previous bursts, it is necessary to relate PCA count rates to ASM count rates. This was accomplished through examination of publicly available ASM data and archival PCA observations of the steady output of the known burster X1735-44 (figure 4.6). The steady output was chosen over a burst for the reason that one can be confident that the steady output measured by the PCA is due to the same source behaviour as the steady output recorded by the ASM at a different point on the same day. Furthermore, the spectral properties of these sources are similar. Based on the steady output from X1735-44, an ASM count rate of  $1 \text{ cs}^{-1}$  was found to correspond to an average PCA count rate (per PCU) over 90s of  $32.2 \pm 3.6 \text{ cs}^{-1}$ . We estimate the total PCA counts recorded during the September 2000 burst (Smale 2001) at 370000  $\pm 5\%$ , which would then correspond to an ASM count rate of  $127 \pm 15.6 \text{ cs}^{-1}$  over 90 seconds. This does assume that the 90s dwell overlaps perfectly the  $\sim 80$ s burst profile (we examine the more general situation in section 4.4.1).

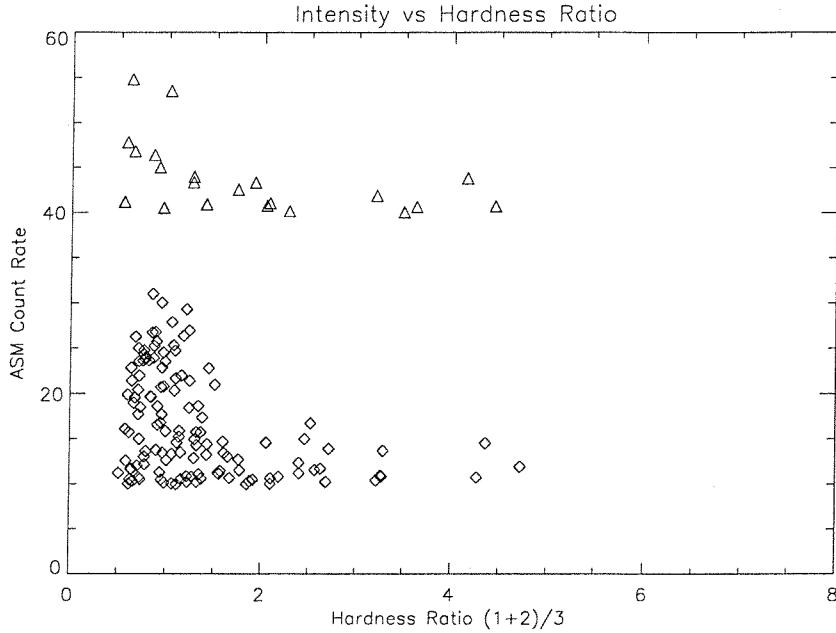
The flux values given in the table are fractions of the flux of the September 2000 burst ( $\sim 6.5 \times 10^{38} \text{ erg s}^{-1}$ ). These values depend on both the source behaviour and the amount of overlap between dwell and burst, so cannot be taken as indicators of either quantity on its own. The hardness ratio used here is defined as the ratio of the count rate in channels 1 and 2 to that in channel 3 (strictly speaking, a “softness ratio”). The maximum duration of outburst was estimated from the data sampling; we know that the source went from quiescence, to outburst, to quiescence, in at most the time interval between the two quiescent points. Galactic bulge sources have been discovered recently to exhibit X-ray “superbursts” lasting several hours (Cornelisse et al., 2000). The fact that we can only set rough limits on our burst candidate durations allows the possibility that an



**Figure 4.6:** PCA and ASM lightcurves of the LMXB burster x1735-444. **Top:** RXTE/PCA lightcurve, 1.9-24.8 keV, 17th June 1999. **Bottom:** RXTE/ASM lightcurve near this observation (vertical dotted line)

event may last longer than the few minutes of the October 1988 burst; indeed, if neutron star photospheric radius expansion is a recurrent phenomenon this might be expected. With the time resolution afforded us by the RXTE/ASM, we can do little more than admit the possibility that some of our candidates may be consistent with long bursts.

The fluxes hover around 10 to 20% of the September 2000 burst, and there is quite a wide range of hardness ratios and durations for the candidates. The coverage of candidate event 11 is such that there are actually three points recorded that may be part of the same profile; this candidate is plotted alongside a probable burst from X1826-238 in figure 4.5. Figure 4.7 shows the count rate-softness plots for both sources in the high count-rate regime. The key point to note is the qualitative similarity of the count rate/hardness plots of the burst candidates from the two sources. This gives us some confidence that we are indeed seeing burst profiles sampled at random points from X2127+119; all points for which the low energy channels dominate also show low count rates. This is further strengthened by the spectral properties displayed by sources known not to burst; the Crab SNR does not of course show bursts, nor does the supergiant/X-ray Pulsar system X1700-37 (figure 4.9).



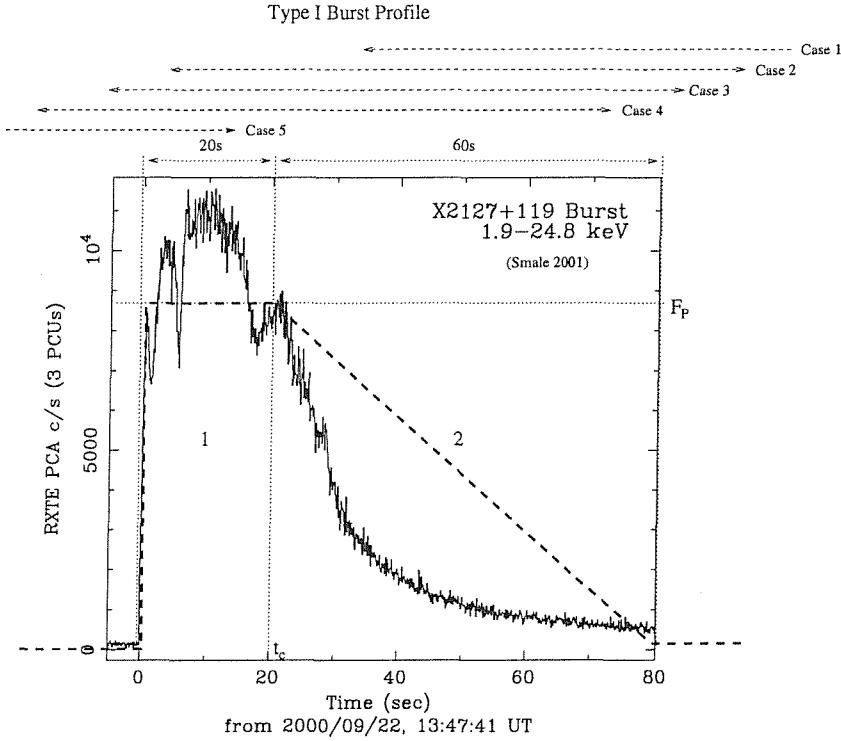
**Figure 4.7:** Count-rate vs spectral softness for high points in the lightcurves of X2127+119 and X1826-238

#### 4.4.1 ASM sampling and the true burst rate

Each dwell measurement lasts 90s, and there are of order 10 dwells per day. To be detected, a significant fraction of the  $\sim$ 80s burst profile must overlap with a 90s dwell, so that the average count rate of the 90s dwell is high enough to stand out from noise. Making the assumption that the burst profile of September 2000 (Smale, 2001) is typical, the probability of burst detection can be calculated and related back to the number of bursts measured to make an estimate of the true burst recurrence time at the source. We model the profile of a Type I burst as a  $\sim$ 20s interval at the near-peak flux  $F_P$ , switching to a steady decline back to quiescence over the remaining  $\sim$ 60 seconds (figure 4.8). Denoting  $t=0$  as the onset of the burst and then  $t_d$  as the commencement time for the 90-s dwell, the average count rate recorded in a 90-s RXTE/ASM dwell is therefore

$$F = \frac{F_P}{90} \times \begin{cases} \frac{(80-t_d)^2}{120} & : t_d > 20 \\ 50 - t_d & : 0 < t_d \leq 20 \\ 50 & : -10 < t_d \leq 0 \\ 50 - \frac{(-10-t_d)^2}{120} & : -70 < t_d \leq -10 \\ (t_d + 90) & : t_d \leq -70 \end{cases} \quad (4.1)$$

The value  $F_P$  is  $\sim 0.7 \times$  the peak brightness recorded by the RXTE/PCA during a burst. At  $\sim$



**Figure 4.8:** The 2000 burst recorded by the RXTE/PCA, from Smale (2001). Overplotted is our simple representation of the burst profile and the cases of overlap with the 90-s RXTE/ASM dwell, reading from top to bottom in the set of conditions (3.2).

$1.1 \times 10^4 \text{ cs}^{-1} \text{ PCU}^{-1}$ , this level, if maintained over 90 seconds, would correspond to an RXTE/ASM count rate of  $\sim 342 \text{ cs}^{-1}$ . Thus we take  $F_P \sim 239 \text{ cs}^{-1}$ . When discriminating candidate burst events, we set our minimum threshold for burst detection at  $10 \text{ RXTE/ASM cs}^{-1}$ . Thus, for a burst to be detected, the degree of overlap must be such that the average count rate over the dwell is  $\gtrsim 0.04F_P$ . This count rate is reached if either (i) the end of the dwell period occurs 4s or more into the burst profile, or (ii) the beginning of the dwell occurs 49s or more before the end of the burst profile, i.e.  $(-85 < t_d + 41)$ .

We approximate the distribution of the times of burst events and dwell measurements as independent, constant random variables. The probability of dwell occurrence is then given by the average duty cycle of ASM observations,  $p_d(t) = \frac{T_{dw}}{T_d}$ , where  $T_d$  is the average time interval between dwell measurements and  $T_{dw}$  the duration of each dwell. The probability that, should a dwell occur at time  $t$ , a burst occurs with sufficient overlap to be detected, is then  $p_b(t) = \frac{\Delta T}{T_b}$ , where  $\Delta T$  is the interval about each burst for which detection is possible and  $T_b$  the average time interval between bursts. The number of burst detections  $N_{det}$  expected is then given by

$$N_{det} = N_b \frac{\Delta T}{T_b} \frac{T_{dw}}{T_d} \quad (4.2)$$

The number of bursts produced by the source over the observation interval  $T_{obs}$  is  $\frac{T_{obs}}{T_b}$ . The average interval between bursts is then given from the parameters of the observation as

$$T_b^2 = \frac{T_{obs} \Delta T T_{dw}}{T_d N_{det}} \quad (4.3)$$

We have estimated  $\Delta T$  above as  $\sim 126$ s. From the parameters of the RXTE/ASM dataset used for this analysis, then (i.e.  $T_{obs} = 2000d$ ,  $T_d = 1.75h$ ,  $N_{det} = 15$ ,  $T_{dw} = 90$ s), the *average* burst interval  $T_b$  is  $\sim 5$  hours.

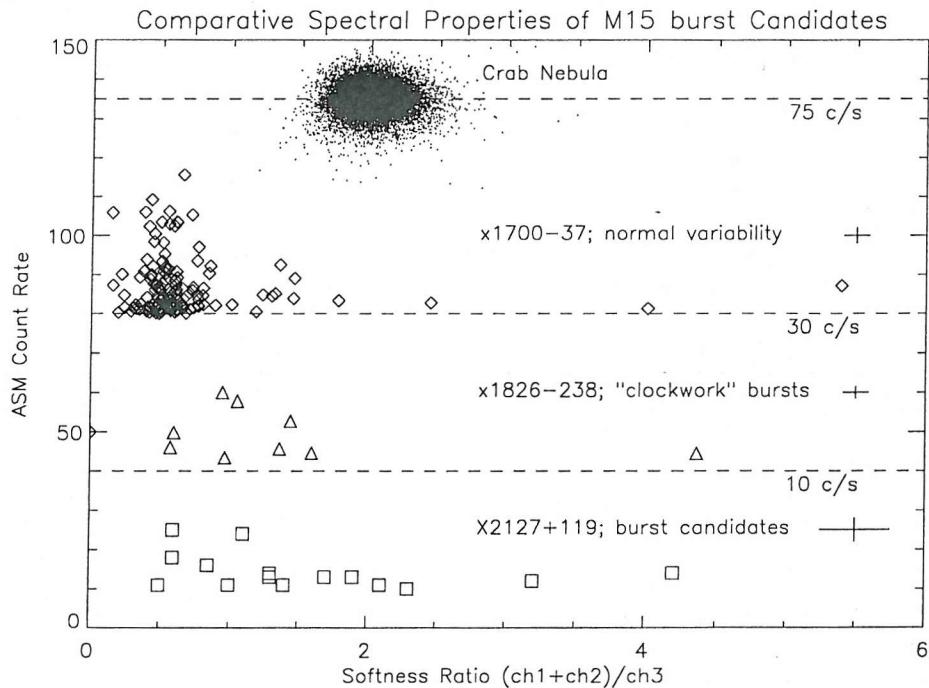
This figure is an estimate only, as we have assumed (i)  $p_d(t)$  &  $p_b(t)$  are truly independent, and (ii) that  $T_b$  is constant over the 2000-day RXTE/ASM dataset. The first assumption is an approximation because the method of fitting dwell detections means that the chance of a datapoint being rejected by the ASM is dependent on the source count rate. As burst candidates are significantly brighter than the quiescent lightcurve, the fraction of datapoints rejected for burst candidates will thus be significantly lower than for the quiescent lightcurve, meaning that  $p_d(t)$  should really be replaced by the conditional probability  $p_{d|b}(t)$ . Assumption (ii) arises because the above method makes no allowance for possible clustering of bursts. Should M15 really contain two separate X-ray sources, then variations in the accretion rate of the burster might be hidden due to contamination from the other source. In this case, we might expect the burst rate to vary significantly (as does x1820-303, for example; Chou & Grindlay 2001). As can be seen in figure 4.2, the ASM sampling has many gaps even at 5-day binning. The number of bursts detected by RXTE/ASM does not therefore represent the number produced by M15 in so simple a manner as we have stated here. As the RXTE/ASM sampling improves with count rate, we suppose at this point that, were bursts more frequent than we have estimated here, they would have been picked up by the RXTE/ASM by virtue of their high count rate, preserving the fraction of bursts generated to bursts detected. This immediately implies that our determination is a rough upper limit on the burst frequency, i.e.  $\bar{T}_b \gtrsim 5$  hours.

## 4.5 HST UV Imaging: a Variability Study

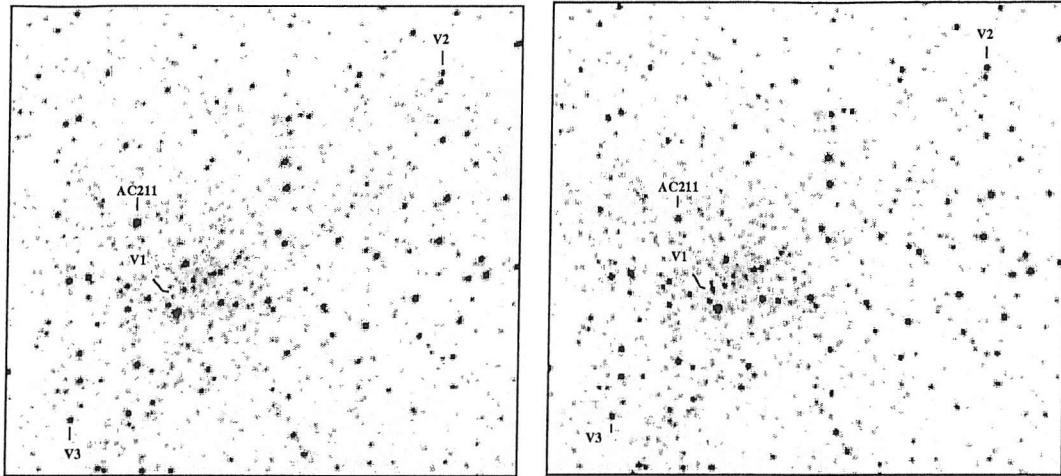
With its extremely bright and crowded central regions, only the brightest members of the core of M15 have been studied with ground-based telescopes (even then requiring the best natural seeing conditions e.g. Aurière & Cordoni 1981). Consequently M15 has been a regular target for the pre- and post - COSTAR corrected HST. The most detailed optical/UV HST studies of

Table 4.1: RXTE ASM Burst candidates

Event	Day <sup>1</sup>	Flux <sup>2</sup>	softness ratio <sup>3</sup>	Max. burst duration (mins)
1	587.946632	0.14	0.5	749.6
2	647.897373	0.14	1.4	96.5
3	652.492373	0.20	0.85	562.9
4	1240.353114	0.16	1.7	96.0
5	1338.510521	0.14	1.0	430.0
6	1367.671632	0.16	1.9	92.3
7	1373.595335	0.31	0.6	93.9
8	1374.19441	0.22	0.6	89.6
9	1375.124595	0.17	4.2	95.7
10	1375.255521	0.12	2.3	4.8
11	1715.607188	0.30	1.1	34.9
12	1726.483298	0.14	2.1	93.6
13	1857.959595	0.15	3.2	860.0
14	1945.611071	0.17	1.3	95.7
15	1984.320503	0.16	1.3	96.3

<sup>1</sup> Days after January 6th, 1996<sup>2</sup> In units of the peak flux of the September 2000 burst<sup>3</sup> Defined as (1.3-5 keV/5-12.1 keV)

**Figure 4.9:** Peak ASM count rate plotted against hardness ratio, for the Crab Nebula (top, offset by  $65 \text{ cs}^{-1}$ ), brightest points from the known *non*-bursting X1700-37 (offset by  $50 \text{ cs}^{-1}$ ), the burst candidates from X1826-238 (offset by  $30 \text{ cs}^{-1}$ ) and burst candidates from X2127+119 itself (bottom, no offset). We find that this analysis does indeed differentiate bursting sources from simple variabilities, strengthening our confidence that the candidates selected for X2127+119 do indeed indicate snapshots of type I-like X-ray bursts.



**Figure 4.10:** HST WFPC2 U-band images of the core of M15, obtained in April 1994 (left) and October 1994 (right). The positions of AC211 and the three new variables are indicated. Each image is 11.5 arcsec across; North points to the top left hand corners, East to the bottom left corners.

the central stellar populations of M15 are presented by De Marchi & Paresce (1994; 1996) and Guhathakurta et al (1996) (hereafter DMP94, DMP96 and G96 respectively). DMP94 used far-UV and U band FOC images to derive colour-magnitude diagrams (CMDs) and demonstrate that there are a significant number (15) of very blue stars (hereafter VBS) in the M15 core. DMP96 show that these are located beyond the blue-straggler section of the CMD, and have properties similar to sdB or extreme horizontal branch stars (based mainly on colour and luminosity), but note that detailed spectroscopic study is necessary in order to elucidate their true nature.

However, the VBS have properties similar to active LMXBs (indeed one of them is actually AC211 itself), and the fainter ones could be accreting white dwarf binaries (i.e. CVs). Since these interacting binaries are known to be highly variable, we decided to examine the archive of these images and select for further study those obtained of the same regions with the same instrumental configurations, but separated in time by long intervals. In order to find blue variables in the core of M15, we analysed two sets of HST WFPC2 narrow-U archival images. The two sets of observations were made in April 1994 (UBVRI observations) and in October 1994 (U-band only).

After careful alignment of the images we used *iraf* to form difference images in order to highlight any variable objects. We detected four blue variables in an 8-arcsec radius from the central cusp. One is the known LMXB AC211, but none of the other three (which we have called V1 - V3; see Table 4.2) is a previously noted VBS (which can be most clearly seen in the G96 composite colour image). We used the Vegamag photometric zeropoints for the WFPC2 to determine the

**Table 4.2: HST/UV Variables in the Core of M15**

Star	Apr 1994			Oct 1994		Location*
	U	B	V	U		
V1	>22.5	-	-	17.60 ± .03		~0.3" NE
V2	16.49 ± .01	16.20 ± .01	15.86 ± .02	15.70 ± .01		~3.5" due E
V3	17.20 ± .02	16.97 ± .02	16.56 ± .04	16.07 ± .02		~7.5" due W

\* with respect to cluster centre

listed magnitudes. The limiting U-magnitude in these exposures (400 s in the F336W filter) is about 22.5. The variables appeared in each of the 8 U-band exposures taken in October 1994 and are therefore not cosmic rays or chip blemishes. These 3 sources (which varied by >0.8 mag between the observation epochs), are marked in figure 4.10. We have not corrected for interstellar reddening, which in the direction of M15 is a very small effect; E(B-V) for M15 is 0.09 (Rosenberg et al., 2000).

#### 4.5.1 The Nature of the Variables V1-V3

On careful examination of the neighbourhood of V1 in the two images, it became clear that the source is only present in the October 1994 U-band image (figure 4.10). There is no trace of it in the April 1994 images and so the variability range is actually >5 mags. There is also no sign of it in any other HST images of M15 (August 1994, December 1998 or August 1999). The variability of V1 is thus characteristic of a dwarf nova outburst. Dwarf Novae (DNe) are faint CVs ( $M_V$  typically 6 to 8 in quiescence, with a large scatter; see Warner 1995) which undergo periodic outbursts of 3 to 8 magnitudes. X-ray transients are X-ray binaries which have similar quiescent absolute magnitudes to DNe, and which undergo infrequent outbursts of 6 to 9 magnitudes (see for example Charles & Coe 2003). At 17.6 mag, the transient V1 is well within the expected range of peak magnitudes for both DNe and X-ray transients at the distance of M15 (though of course there is no way to know how close to the true peak this value represents), which makes it an event long sought after but rarely seen in the high density environments of globular cluster cores. Moreover, the quiescent magnitude of < 22.5 is well within the 21-24 mag expected with the 10.2kpc distance to M15. Dwarf Nova outbursts typically repeat on a timescale of weeks to months (Buat-Ménard et al., 2001), whereas X-ray transients recur on timescales of tens of years (Šimon, 2002). Thus the chance of a randomly placed HST observation picking up a Dwarf Nova outburst is far higher than an X-ray transient. The expected detection rate with HST may be estimated from CV population studies. For example, with its higher core mass and central density, M15 is expected to produce more than

the  $\sim 200$  CVs predicted for 47 Tuc (Di Stefano & Rappaport, 1994), through a higher rate of the tidal capture and three-body formation mechanisms (Elson, Imagaki & Hut 1987). Applying the 47 Tuc results to M15, we also expect  $\sim 50$  DNe. With a mean recurrence time of 50 days and mean outburst length 10 days, we might thus expect HST to see  $\sim 10$  DNe in outburst in each visit. These predictions may be overoptimistic, however they do indicate the chances of detecting a DN should be very good. In the absence of further information about the nature of the accretor associated with V1, we thus conclude V1 is a Dwarf Nova. We caution, however, that we cannot rule out the possibility that HST might have caught a rare LMXB transient outburst: indeed, several millisecond radio pulsars (the descendants of neutron star LMXBs) are detected in the core of M15 (Phinney, 1996). At  $\sim 1.2$ mag, the variability range of V2 and V3 between the April and October observations are consistent with less dramatic systems. The absolute magnitudes ( $U \sim 15 - 17$ mag) and colours are typical of RR Lyrae variables (Allen, 1973).

## 4.6 Discussion

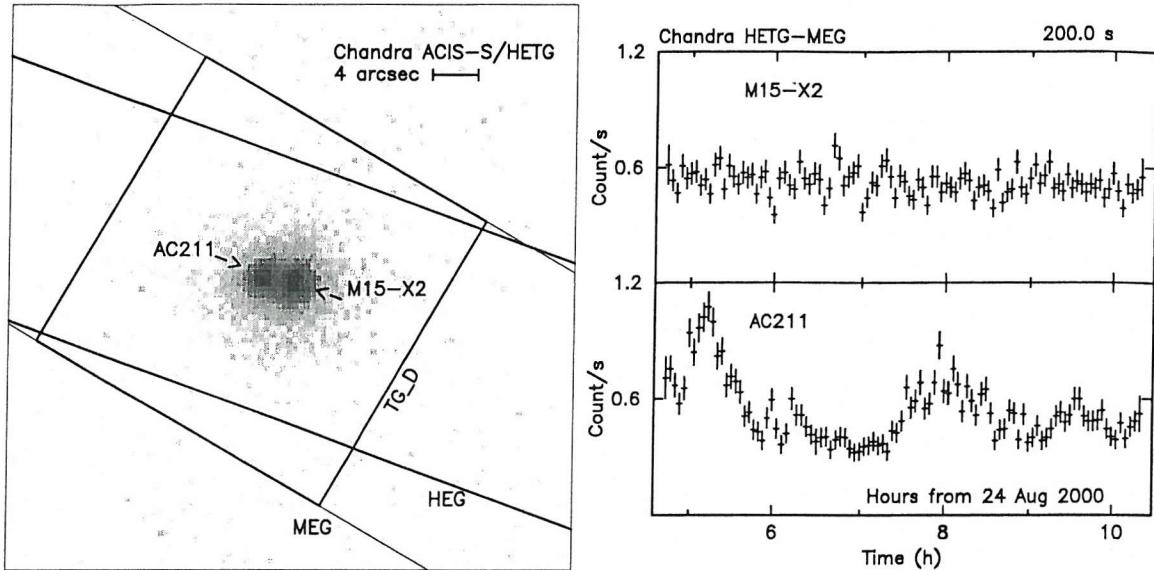
The optical and X-ray light curves of X2127+119/AC211 (Homer & Charles, 1998; Ioannou et al., 2002) contain the very stable (and extended) eclipse feature which requires a high inclination ( $\gtrsim 70^\circ$ ) for this system. With its very low  $L_X/L_{opt}$  ratio ( $\sim 20$ ) this has provided very strong support for the ADC interpretation. But the van Paradijs et al (1990) X-ray burst from M15 is the major key factor that does not fit into the ADC model and overall scenario. Note also that the distance to M15 is well constrained at 10.5 kpc (Rosenberg et al 2000), thereby providing accurate luminosities for both the quiescent and burst X-ray emission. We consider then that there are only 3 possible explanations available to resolve this situation:

- (a) the X-ray burst is actually scattered in the ADC but remains detectable as a coherent event. We believe that this is unlikely because the time profile is typical for a type I X-ray burst, suggesting differences in scattered paths to the observer less than a few light seconds (the approximate Roche Lobe radius of the neutron star in X2127+119). Moreover, the peak luminosity is extremely high (at  $L_X \sim 4.5 \times 10^{38}$  erg s $^{-1}$  assuming isotropic emission; van Paradijs et al 1990). Scattered bursts are unlikely to reach this luminosity because of the considerable X-ray scattering and reprocessing that would be occurring (which restricts ADC sources to only a few percent of their intrinsic luminosities; White & Holt 1982);

(b) there exists a long-term modulation in the X-ray light curve that could arise from a warped accretion disc (as a result of an instability to X-ray irradiation; see chapters 2 and 6), allowing an occasional direct view of the inner disc/compact object. The prototypical source for this behaviour is Her X-1, for which the disk is believed to block the neutron star from view during its 35-day disk precession period; see chapter 6. We believe that this is also unlikely to occur for X2127+119. If the X-ray source were to become directly visible at certain times, then, assuming an average  $L_X/L_{opt}$  ratio for LMXBs (van Paradijs & McClintock 1995), the increase in X-ray brightness would be much more than the factor 2 or 3 visible in the ASM lightcurves (figure 4.2). Indeed, when Her X-1 enters its Anomalous Low States, the source drops out of the bottom of the range detectable by RXTE/ASM, corresponding to at least a factor 10 difference in X-ray brightness (chapter 6). Furthermore, the source would need to be bursting frequently in order for an event to occur during a period of direct viewing; Frequent bursts are simply not present in the extensive ASCA and RXTE monitoring campaigns (Ioannou et al 2002);

(c) there exists another LMXB (the burster) within the core region of M15 and also very close to AC211 (within 5 arcsecs). This is both expected (on the basis of the compact, high density nature of the core combined with the effects of mass segregation) and required (in order for the source not to have been resolved by previous imaging X-ray missions, particularly ROSAT). Of course, the X-ray burster could also be a chance alignment with the direction of M15, but this is extremely unlikely because of the high galactic latitude of M15 ( $b_{II} \sim -27^\circ$ ) and the necessarily very precise alignment required (arcseconds) with the core of M15.

We consider therefore that a second active LMXB in M15 is the most likely explanation and also suggest that such a source might be associated with one of the UV variable stars we have found in the archival HST images. For the observed X-ray flux levels (now indicative of the intrinsic luminosity of the source) we may scale with respect to similar LMXB bursters in the plane (e.g. X1735-444, which has  $M_V$  of 2.2 but is ten times more luminous in X-rays; see van Paradijs & McClintock 1995) to infer  $V \sim 19$ . This additional source could then be responsible for the variations present in the ASM light curve (figure 4.1), and hence is likely to be comparable to the currently observed flux from X2127+119, but of course must be much fainter than the intrinsic luminosity of the ADC source AC211. Note that the observed X-ray, optical and UV 17.1hr modulation of X2127+119/AC211 demonstrates that a significant fraction of the total M15 X-ray flux must originate from this system, but if it is contaminated by another source then this implies that the X2127+119 eclipse must be deeper. There would also be implications for the spectral properties



**Figure 4.11:** Discovery of a second LMXB in M15 (from White & Angelini 2001). **Left:** ACIS/HETG zero-order image of M15, showing the X-ray source associated with AC211 and the new source, denoted X2. **Right:** ACIS/HETG lightcurves for the two sources.

of the source (Ioannou et al 2001)

## 4.7 Update: M15 with Chandra & XMM

While writing the paper on which this chapter is based, we became aware of a Chandra observation by White & Angelini (now published in ApJ, 2001) that had serendipitously resolved X2127+119 into 2 separate sources. This clearly shows that our preferred interpretation of the RXTE/ASM and HST/WFPC2 datasets, in which the burster and ADC source are separate systems, is the correct one. We further note that the second source is *not* at the position of V1, strengthening our interpretation of V1 as a dwarf nova. Indeed, from our analysis of the archival HST images, M15/X2 is instead associated with a faint blue star ( $U \sim 18.6$ ), consistent with this type of system (figure 4.11).

## Chapter 5

# The Variable Third Period in SMC X-1

Long term X-ray monitoring data from RXTE/ASM and CGRO/BATSE reveal that the superorbital period in SMC X-1 is not constant, but varies between 40-60 days. A dynamic power spectrum analysis indicates that the third period has been present continuously throughout the five years of ASM observations. This period changed smoothly from 60 days to 45 days and then returned to its former value, on a timescale of approximately 1600 days. During the nearly 4 years of overlap between the CGRO & RXTE missions, the simultaneous BATSE hard X-ray data confirm and extend this variation in SMC X-1. Our discovery of such an instability in the superorbital period of SMC X-1 is interpreted in the context of recent theoretical studies of warped, precessing accretion disks. We find that the behaviour of SMC X-1 is consistent with a radiation-driven warping model. The counterpart paper to this chapter has appeared in print (Clarkson et al., 2003b).

### 5.1 SMC X-1

SMC X-1 is a luminous ( $L_X \sim 6 \times 10^{38} \text{ erg s}^{-1}$ ) X-ray binary, with eclipsing pulsar behaviour. The spin period of the pulsar is 0.71 s and the orbital period of the system is 3.89 days, with orbital decay on a timescale of  $10^5$  years (Reynolds et al., 1993). Pulse arrival times from the neutron star and spectroscopic studies of the donor show the orbit to be circular ( $e < 0.00004$ ; Levine et al.,

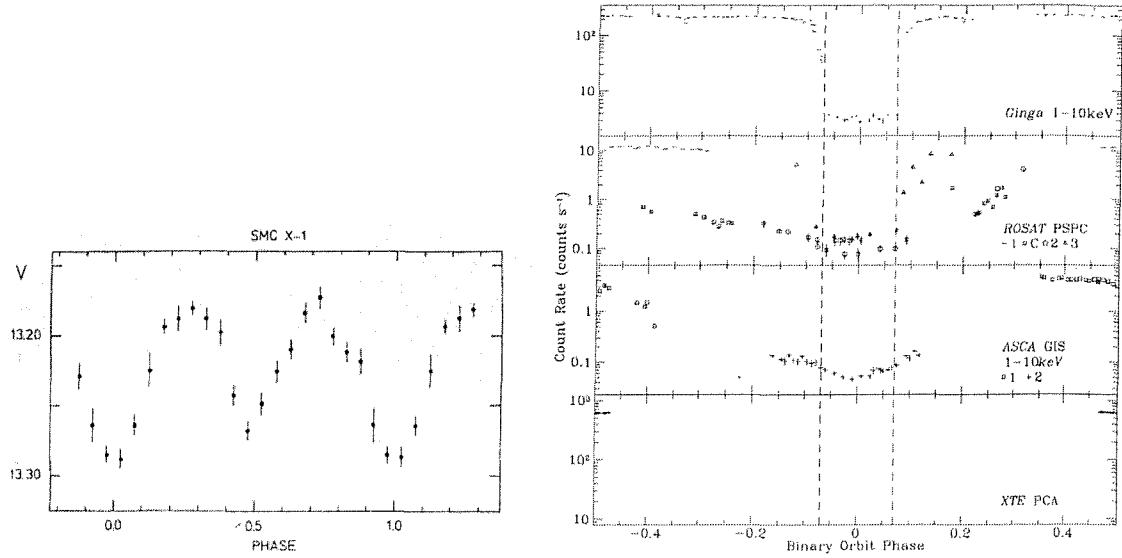
1993; Wojdowski et al., 1998), with neutron star and B0 I donor masses  $1.6 \pm 0.1 M_{\odot}$  and  $17.2 \pm 0.6 M_{\odot}$  respectively, corresponding to ( $56.2^\circ < i^\circ < 58.6^\circ$ ; Hutchings et al., 1977).

### 5.1.1 Accretion Disk and Superorbital Variation

Evidence for an extended accretion disk has persisted as long as this system has been observed. Accumulation of optical lightcurves from 1979-1982 shows two remarkable features: (i) at phase 0.5 in the X-ray orbital lightcurve, i.e. when the neutron star crosses the face of the donor, the optical light drops by  $\sim 0.07$ mag, and (ii) at phase  $\sim 0.0$ , when the neutron star passes behind the donor, there is a larger drop of  $\sim 0.1$ mag (van Paradijs & Kuiper, 1984). The drop at phase 0.0 is well fit by a combination of ellipsoidal variation due to donor deformation, and emission due to X-ray heating at the face of the donor. The dip at phase 0.5 cannot be fit by these features, however, but is instead well-fit by an extensive accretion disk, which the authors assumed to fill its Roche Lobe in order to make estimates on the X-ray albedo of the disk (figure 5.1). Furthermore, wind-driven accretion models cannot provide the observed accretion luminosity (Hutchings et al., 1977), suggesting Roche overflow and disk formation. Finally, as shown from X-ray pointings taken by Ginga, ROSAT, ASCA and RXTE over a period of nearly seven years, the pulsar has been spinning up constantly since its discovery, with  $P_{\text{spin}}(t)$  described by a quadratic fit (Wojdowski et al., 1998). This suggests that any accretion disk must be persistent, to regulate the sense of the spin change, and also large enough to regulate  $\dot{M}_1$  so that the propellor effect never causes accretion to switch off, as often occurs in purely wind-fed systems (c.f. White, Nagase & Parmar 1995).

Taking the maximum extent of the accretion disk  $R_{\text{disk}}$  to be the neutron star's Roche Lobe, and assuming the donor fills its Roche Lobe, we can use elementary trigonometry to estimate the range of system inclinations necessary to eclipse the pulsar but not the disk (figure 5.2), as ( $54.1^\circ < i_g^\circ < 61.0 \pm 0.2^\circ$ ). Remarkably, the measured  $i^\circ$  (see above) is within this range. However, the upper limit for  $i_g^\circ$  is reduced to  $58.0^\circ$  and  $56.6^\circ$  for disk radius  $0.5R_{L1}$  and  $0.3R_{L1}$  respectively, so for the disk not to be occulted by the donor, it must have  $R_{\text{disk}} \gtrsim 0.55R_{L1}$ . We urge the reader to bear in mind that this calculation assumes the disk to be flat and in the binary plane; as we shall see later in this chapter, the disk probably deviates significantly from this geometry, making it more likely that the donor occults the compact object fully but not the disk.

Two early sets of optical lightcurves by competing teams of investigators give contradictory evidence as to the existence of any superorbital variation: observations by the Dutch group show no

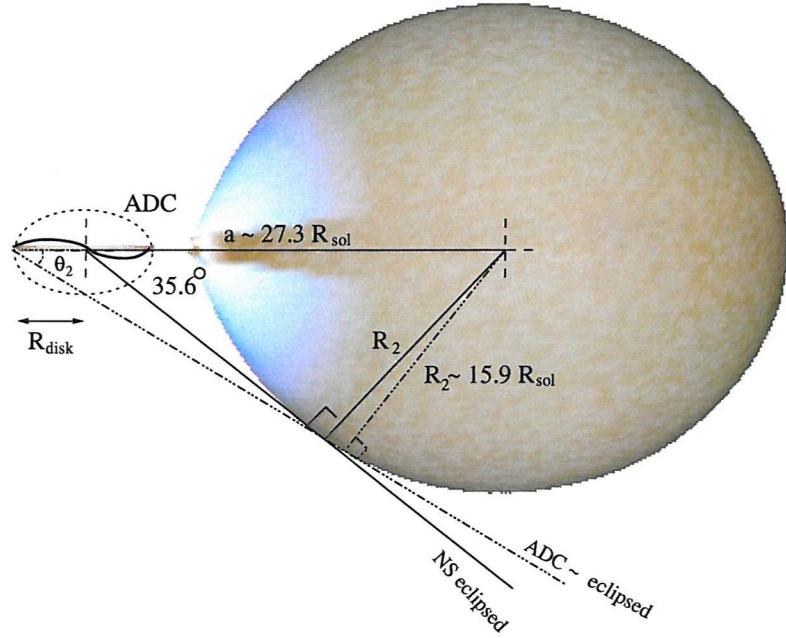


**Figure 5.1:** Left: Optical lightcurve of SMC X-1, showing clear evidence for an extensive accretion disk. From van Paradijs & Kuiper 1984. Right: X-ray pointings of SMC X-1, from Wojdowski et al (1998). GINGA, ASCA1, ROSAT1, ROSAT3 and RXTE pointings occurred during the high state of the superorbital cycle, with ROSAT2, ROSATC and ASCA2 taking place during the low state.

evidence for any coherent optical superorbital period between 25 and 150 days (van Paradijs & Kuiper 1984). Observations by the Russian group over the same interval, claim an optical superorbital periodicity of amplitude  $\sim 0.04\text{mag}$  and period 35.25 days (Cherepashchuk, 1982), but as their raw data has not been published, this claim is impossible to test. However, a superorbital periodicity is emphatically present in the X-ray emission, on a timescale first measured as a few months (Rothschild, 1981), then refined to  $\sim 60$  days (Wojdowski et al 1998 and references therein). The arrival times of the high and low states in this superorbital cycle are seen to vary by  $\sim 10$  days (Wojdowski et al 1998). This disk is presumed to periodically occult the X-rays from the central source, thereby creating a modulation in the X-ray lightcurve in a manner analogous to that of Her X-1 and LMC X-4 (chapter 2).

### 5.1.2 X-ray Spectral Behaviour

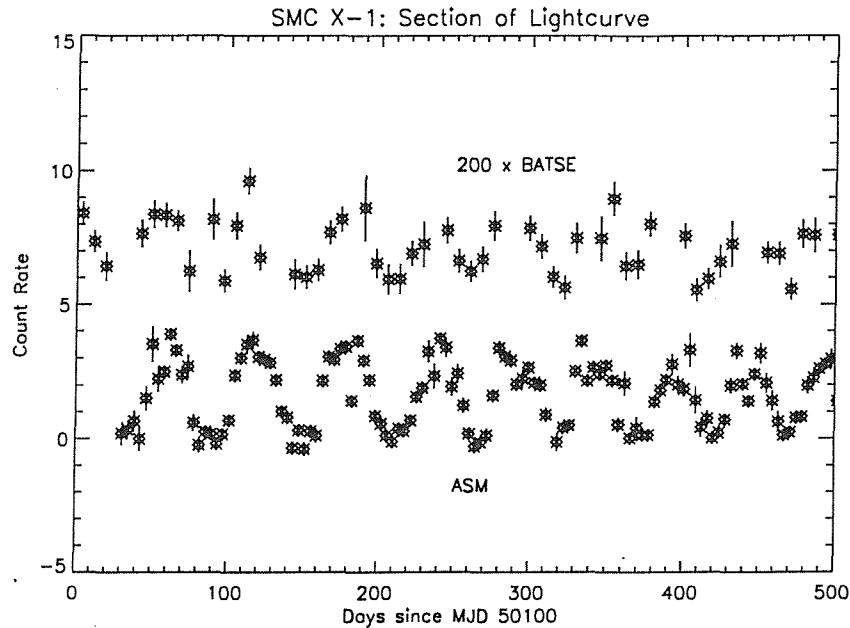
As part of a simultaneous Chandra-ACIS/HST study (Vrtilek et al., 2001), the 1-10keV spectrum of SMC X-1 was sampled during two low states and one high state in the superorbital cycle. In each case the spectrum was sampled at four intervals in the orbital cycle, allowing spectral behaviour to be followed with phase in both cycles. During the high state of the superorbital cycle, the out-of-eclipse spectra are dominated by continuum emission from the neutron star, which is successfully



**Figure 5.2:** Schematic of SMC X-1, from the established masses and inclinations. Visualisation software by Rob Hynes, UT Austin.

modeled with the Wojdowski, Clark and Kallman (2000; hereafter WCK00) cut-off power-law model. During eclipse, the continuum level drops dramatically, replaced by broad emission lines of O, Ne, Mg, S, Ca and Fe. During the low state in the superorbital cycle, all spectra are dominated by the emission lines. These lines were not detected by previous missions because of their lower sensitivity and spectral resolution. ROSAT pointings show that during the low state the orbital lightcurve has no abrupt transition into and out of eclipse but rather a gradual transition, suggesting the emission originates in an extended region. The broad emission lines detected by Chandra are interpreted as recombination features from X-ray heated wind from the donor; these lines are certainly too broad to arise at the surface of the neutron star. Furthermore, the observations at orbital phase 0.5 show variation in count rate along the superorbital cycle of a factor  $\sim 10$ , whereas pointings taken at eclipse show only a factor  $\sim 2$  variation. Alternatively, restricting attention to the high state in the superorbital cycle shows a factor  $\sim 50$  difference between the points in and out of eclipse. The ROSAT observations show similar behaviour, with the same factor  $\sim 10$  difference between high and low superorbital phases out of eclipse, and the same factor  $\sim 2$  difference during eclipse, with the same factor 50 difference in and out of eclipse during the high state (Wojdowski et al., 1998).

The published UV spectra from the Chandra/HST campaign show little variation with orbital phase. This would suggest the wind dominates the UV emission from the system. However, only



**Figure 5.3:** Section of the ASM and BATSE datasets for SMC X-1. Data has been binned to 7.78 days, or two orbital cycles, per bin.

spectra near phase 0.0 and 0.5 have been published. In comparison, the optical brightness is similar at phase 0.5 to 0.0, but these points represent the extrema of the optical variation. Thus it is not clear from the published work how the UV output really varies with orbital phase, leaving identification with a system component uncertain. In this chapter we present the results of an analysis of both ASM and BATSE data, that allows the variation in the superorbital period to be followed continuously over more than a decade (Chapter 3). We also extend the timeline of observations by 3 years since (Wojdowski et al., 1998), which allows us to see that the variation of the superorbital period might itself be periodic, as suggested by (Ribó et al., 2001). We show that this variation in the long period must be due to interaction of modes in a warped, precessing accretion disk.

## 5.2 Dynamic Power Spectra of RXTE and BATSE datasets

Dynamic Power Spectra were computed for the BATSE and RXTE/ASM datasets of SMC X-1 as described in Chapter 3. The lightcurves can be seen in figure 5.3, with the power spectra in figure 5.4. Not only is there a clear superorbital period on a period similar to that reported by Wojdowski et al (1998), but it varies with time, on a scale of  $\sim 7$  years. Because the data windows used in this analysis are not independent but have significant overlap, the possibility arises that

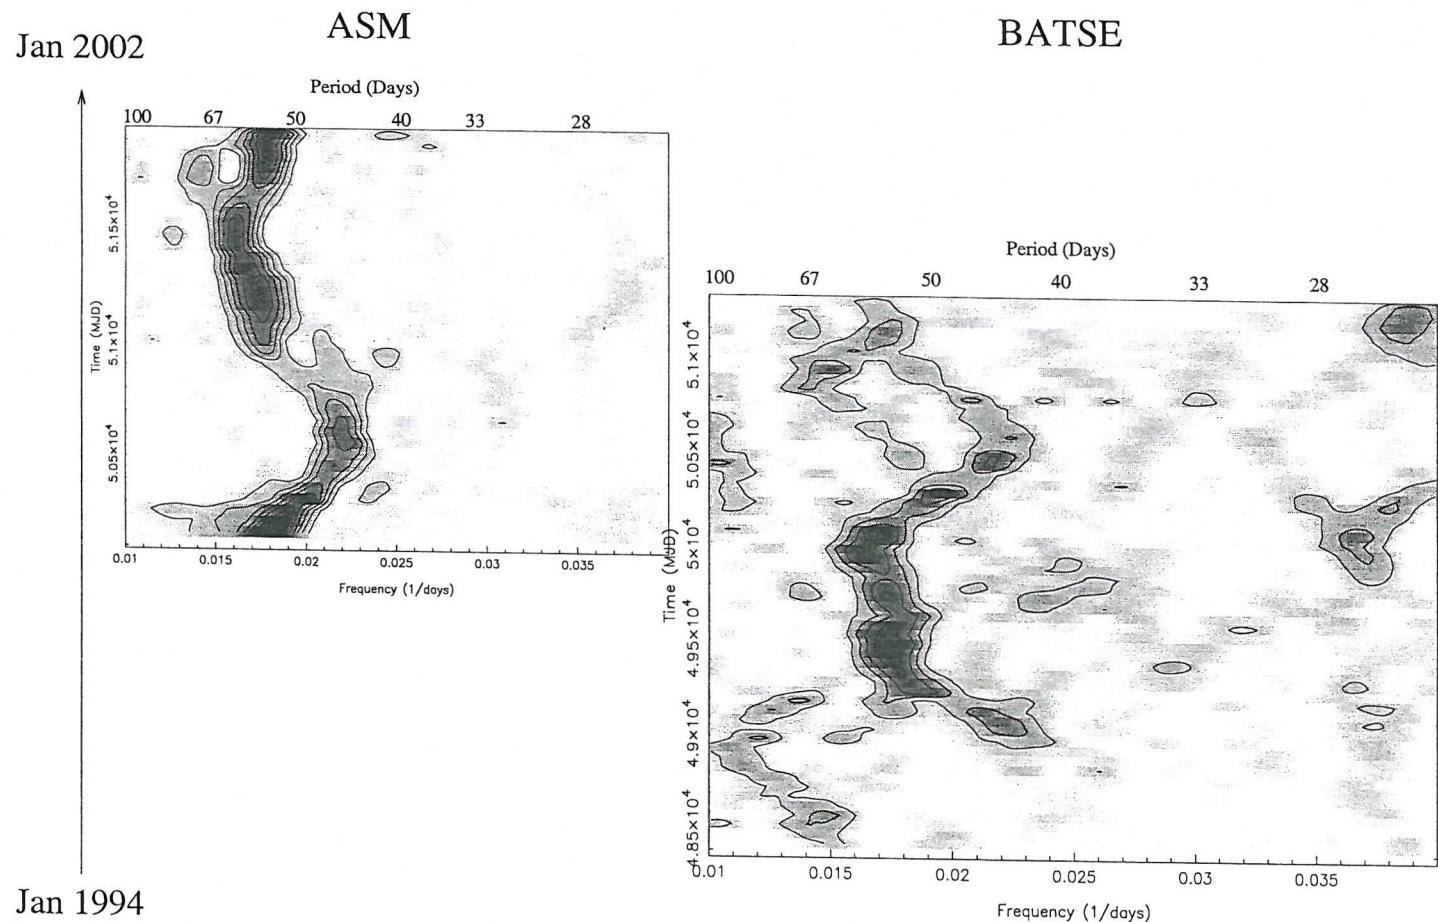
this apparently smooth variation might instead be a result of this overlap. For example, if one periodicity were to suddenly switch off or be replaced by another, entirely unrelated periodicity at a different frequency, one might expect the observed DPS behaviour to be reproduced. This possibility was tested for by simulating an abrupt switch in periodicities in the manner described in chapter 3. The jump from  $\sim 65$  to  $\sim 45$  days exhibited by SMC X-1 is large enough that, were it to be an abrupt change, the DPS would clearly not confuse this for smooth variation (see e.g. figure 5.5). We thus believe we are dealing with smooth variation of a single periodicity.

Some additional structure is seen in the BATSE lightcurve at roughly twice the frequency of the main structure: this is thought to be a harmonic of the main variation. The remarkably close agreement between the datasets suggests that whatever process is responsible for the variation affects both wavebands. In the case of BATSE occultation data, interaction between the spacecraft precession period of  $\sim 51$  days and any detected periodicity in SMC X-1 might be expected. The orbital precession period was plotted for the duration of the BATSE dataset and found not to correlate with the variation in SMC X-1's superorbital period. Furthermore, no bright X-ray sources are within the Earth Occultation errorbox of  $2^\circ$  from SMC X-1. However the key point to note is the close correlation between the ASM and CGRO results during the period of overlap. This is convincing in itself because RXTE has a different orbital precession period and the ASM works on an entirely different principle to BATSE. Data from the ASM and BATSE provide a 4.4 year stretch of simultaneous coverage in two separate regions of the X-ray spectrum. During this period of overlap (MJD 50083-51690), the behaviour of the superorbital period is closely mirrored in the two, completely independent, datasets. This can be seen clearly in figure 5.4. Nor is this effect seen in any other bright X-ray source observed by the CGRO.

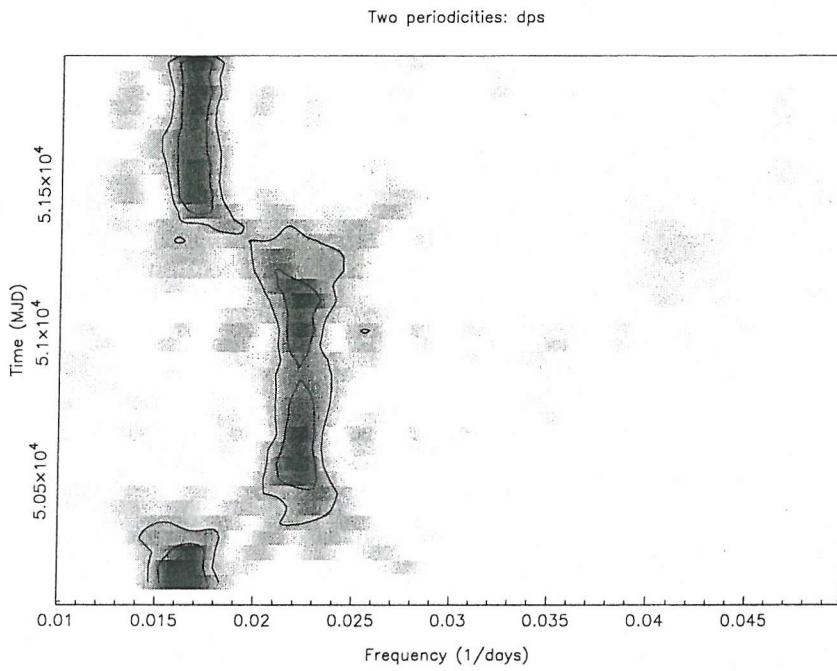
### 5.3 Spectral Behaviour of the Superorbital Period

We use the spectral information on the variability which is provided by the multi-channel nature of the ASM dataset, in order to discriminate between possible superorbital period mechanisms. To increase signal to noise, binning the data by phase is desired. However, the variation in superorbital period rules out folding over the entire dataset. Examination of figure 5.4 shows that over the region MJD 50550 - 50800, the superorbital period is roughly stable for both ASM and BATSE datasets.

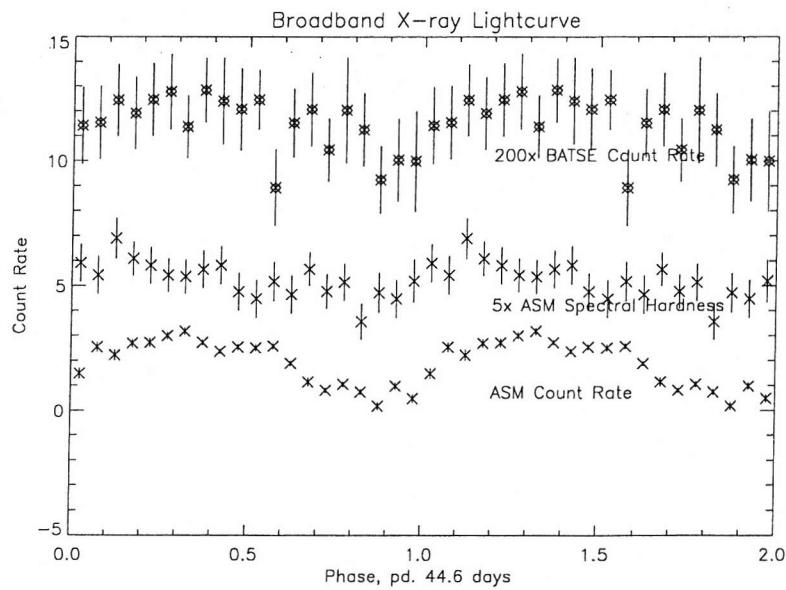
Figure 5.6 shows the phase-binned lightcurve of the SMC X-1 datasets during this interval. We



**Figure 5.4:** Dynamic power spectrum for the 20-100 keV lightcurve of SMC X-1 observed by BATSE over 9 years (MJD 48361-51690; right) and for the 1.3-12.1 keV lightcurve provided by RXTE over 6 years (MJD 50083-52312; left). Contours spaced at 4-unit intervals of LS power for BATSE and 40-unit intervals for RXTE.



**Figure 5.5:** Periodogram of simulated dataset, in which the initial periodicity is abruptly replaced by a second periodicity at 43 days, to test the hypothesis that the apparent smoothness of the long period variation might be an artefact of the 50-d overlap of adjacent data windows. The DPS has clearly resolved the two periodicities and there is little apparent smoothing (see discussion in Chapter 2, in which this figure is also used).



**Figure 5.6:** ASM countrate, ASM hardness ratio, and scaled BATSE count rate, binned by phase in the quasi-stable superorbital period of 46 days that persists from MJD 50550 - 50800. Spectral Hardness is defined in this paper as the ratio of counts at 3.0 - 12.1 keV to those at 1.3 - 3.0 keV.

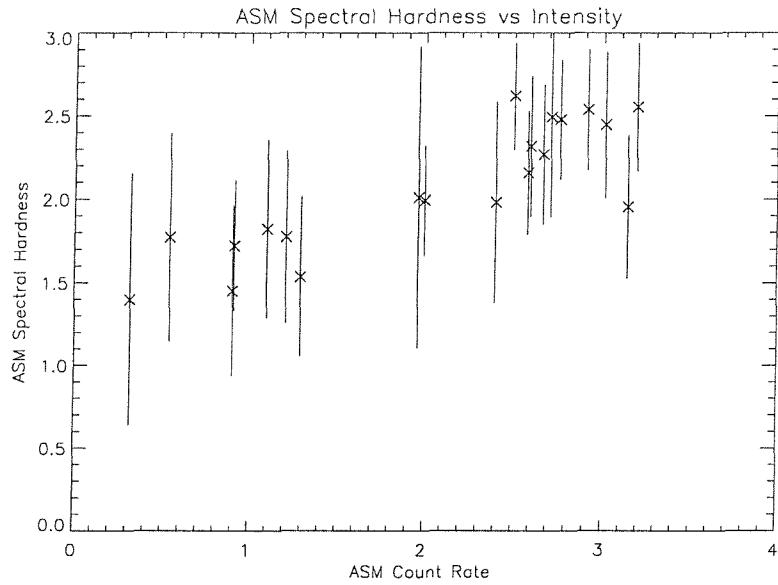


Figure 5.7: ASM spectral hardness plotted against total count rate, for the dataset of figure 5.6.

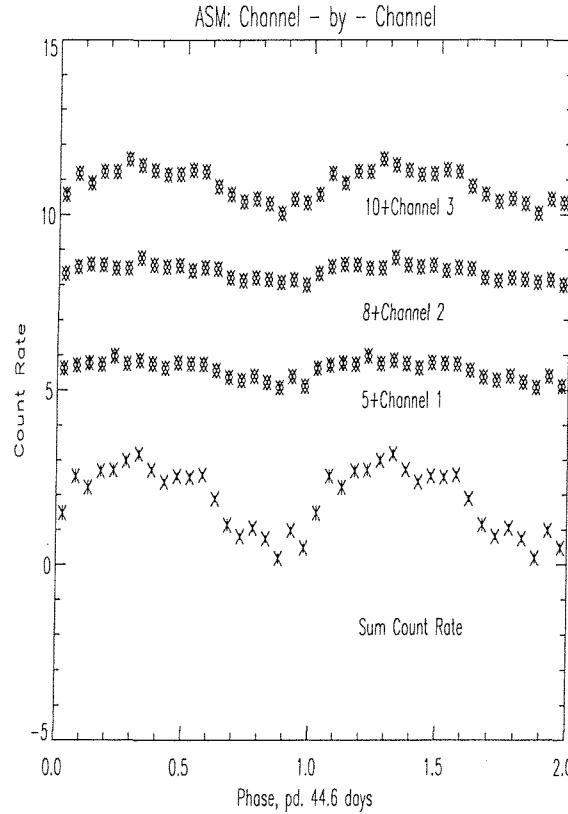
see immediately that the ASM variations mirror those in of BATSE. Furthermore, the shape of the ASM total count rate variation is almost exactly reproduced in the variation of the hardness ratio. The relation between total count rate and spectral hardness is shown explicitly in figure 5.7.

When the ASM phase-binned lightcurve is split into its individual channels (figure 5.8), we see that the amplitude of variation is highly energy dependent. The variation at energies 1.3 - 3.0 keV, represented by ASM channels 1 and 2, has amplitude  $\sim 1 \text{ cs}^{-1}$  and the variation is smooth, whereas at energies 5.0 - 12.1 keV the variation is about twice this amplitude.

## 5.4 Discussion

Our analysis of the RXTE and CGRO archival datasets has not only confirmed the presence of the superorbital period in SMC X-1, it has clearly demonstrated that this itself is varying on an even longer timescale (conceivably periodic, though as yet there is no evidence for this). Here we shall discuss the superorbital period, along with reasons why it might vary with time. Four mechanisms have been put forward to explain the superorbital period, and we shall discuss these for the case of SMC X-1.

Broadly speaking, mechanisms for generating a superorbital period must do so by a variation in the X-ray intensity of the source, by variation in uncovered area as seen from the observer, or by



**Figure 5.8:** ASM total countrate, together with channel - by - channel behaviour, on the same phase scale as figure 5.6. The low energy channels show smaller amplitude variations and smoother morphology of variation.

a varying absorption along the line of sight to the observer. Varying absorption does not however look promising as the prime mechanism for the variation observed in SMC X-1. This is because the BATSE lightcurve varies in step with the ASM variation, showing that the variation occurs across the entire energy range 1.3 - 100 keV. This rules out wind absorption: although the wind density in SMC X-1 can vary spatially by factors of up to  $\sim 1000$  (Blondin & Woo, 1995; Wojdowski et al., 2000; Vrtilek et al., 2001), this would have no effect on the BATSE lightcurve. Variation of uncovered area may be brought about by varying occultation by the accretion disk during the precessional cycle, a possibility we examine in section 5.4.1.

In itself, the relation between ASM intensity and spectral hardness does not act as a good discriminant between mechanisms, as the lightcurve represents a multiple component source (Chapter 1). The WCK00 model ascribes the hard component to the region of the central source, which is also entirely consistent with the simultaneous Chandra/HST observations, in which, during the low state in the superorbital cycle, the continuum vanishes  $\gtrsim 5$  keV but is still present at a low level beneath this approximate cut-off (Vrtilek et al. 2001). Emission at softer energies is thus ascribed

to outflow from the accretion disk, which may not be fully occulted due to its elevation above the disk surface.

### 5.4.1 Mechanisms of Superorbital Variation

Precession of the neutron star's magnetic axis was invoked to explain the  $\sim 35$ d periodicity in Her X-1 (by changing the accretion geometry; see Trümper et al. 1986). However this mechanism can be ruled out in the case of SMC X-1 for two reasons. Firstly, the precession torque is expected to be of the same order as the spin-up torque (Lamb et al., 1975). But pointed X-ray observations of SMC X-1 show the spin period to be monotonically decreasing (Wojdowski et al., 1998), suggesting that precession should either also monotonically decrease in period, or if the sign of the torque changes, produce more violent behaviour in the dynamic power spectrum over time. This is not observed. Secondly, the precession of the neutron star would result in a superorbital variation of the pulse profile, whereas pointed observations suggest that the pulse profile is in fact stable. Forced precession has recently been ruled out in the case of Her X-1 for exactly this reason (Scott et al., 2000).

Another possibility is variation in mass transfer from donor to neutron star on the superorbital period. Giant pulsators, most notably RR Lyrae variables, are known to change their pulsation periods (Jurcsik et al., 2001). Thought to be driven by horizontal branch evolution, these pulsation period changes are typically monotonic and of order 1 d per 0.1 GYr. However there are cases in which the variation is more chaotic and on a faster timescale (Jurcsik et al., 2001), so we must address this mechanism. The fact that such pulsations have never been detected in the donor of SMC X-1, given the high sensitivity of searches for such behaviour, leads us to strongly doubt the existence of pulsations at all. Furthermore, we expect companion star pulsations to produce large changes in bulk mass transfer rate from donor to accretion disk, manifesting themselves as large changes in the magnitude, or even direction (Nelson et al., 1997) of spin of the neutron star. As noted by Wojdowski et al (1998), however, the changes in neutron star spin rate are always in the same sense.

More gentle variation in mass transfer rate can be brought about if the XRB is part of an hierarchical triple. The third body causes the inner XRB to undergo nodal precession, leading to mass transfer variations even if the inner binary is eccentric only by a small amount (Chou & Grindlay 2001, Mazeh & Shaham 1979). Such a mechanism can be ruled out for SMC X-1, however, by

consideration of the orbit of the centre of mass of the inner XRB about the system centre of mass, and the variation in pulse arrival time this would produce. The maximum discrepancy in pulse arrival time  $(\Delta t)_m$  due to the orbit of the centre of mass of the binary, is the light travel time between the points in the orbit marking the maximum and minimum distance from the observer to the inner binary CM. Taking  $a_3$  as the outer binary separation,  $e_3$  the eccentricity of the outer orbit,  $i_3^\circ$  its inclination to the observer and  $\theta_3$  the angle in the third body orbital plane between the semimajor axis and the line of sight, this distance is then given by

$$2c(\Delta t)_m > a_3 \frac{\mu_{12,3}}{M_1 + M_2} \sin(i_3^\circ) f(e_3, \theta_3) \quad (5.1)$$

with  $(\Delta t)_m$  the maximum discrepancy in pulse arrival time and  $\mu_{12,3}$  the reduced mass of the hierarchical triple formed by SMC X-1 and its outer companion. Ellipse geometry tells us that the projection factor  $f(e_3, \theta_3)$  is given by

$$f(e_3, \theta_3) = (1 - e_3^2) \left( \frac{\sin(\theta_3 + \theta_+)}{1 + e_3 \cos \theta_+} - \frac{\sin(\theta_3 + \theta_-)}{1 + e_3 \cos \theta_-} \right) \quad (5.2)$$

with  $\theta_+$  and  $\theta_-$  the two solutions to the relation

$$\tan(\theta + \theta_3) = -\frac{\frac{1}{e_3} + \cos \theta}{\sin \theta} \quad (5.3)$$

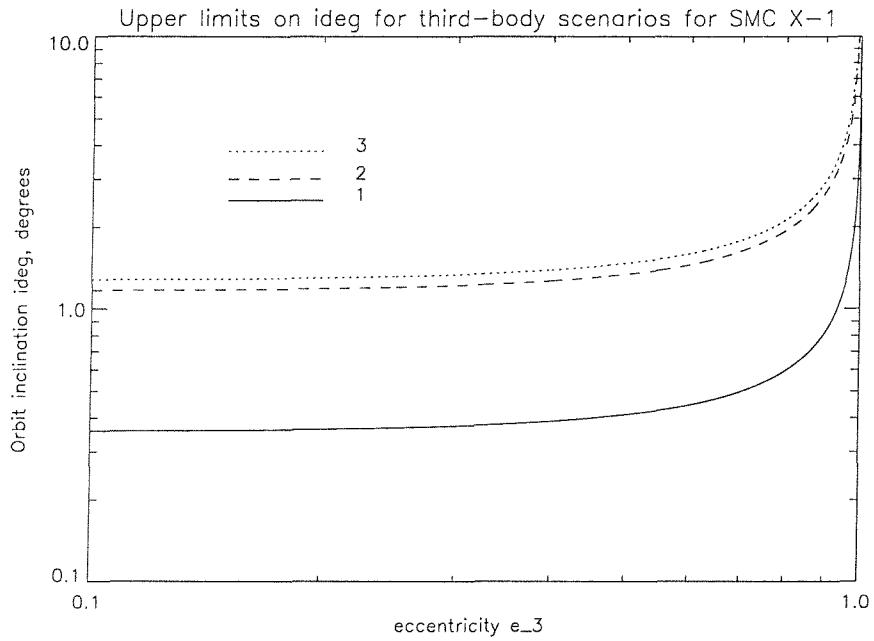
Evaluating  $f(e_3, \theta_3)$  thus adds considerable effort to the calculation. However, as our interest is in finding the *maximum* range of  $e_3$  and  $i_3^\circ$  for which a third body is consistent with observations, we can simplify matters greatly by noting that the minimum of  $f(e_3, \theta_3)$  is just the ratio of the minor and major axes of the binary, i.e.  $b_3/a_3 = (1 - e_3^2)^{1/2}$ . Measured discrepancies between pulse arrival delays and those predicted from a two-body system amount to less than  $\sim 15$ ms in any X-ray pointing since Ginga (Wojdowski et al 1998). Therefore, to go undetected in the pulse arrival times, the orbital period of any third body must satisfy

$$3.10 \times 10^{-3} \gtrsim \frac{M_3}{(18.8 + M_3)^{2/3}} P_3^{2/3} \sin(i_3^\circ) (1 - e_3^2)^{1/2} \quad (5.4)$$

with  $P_3$  the orbital period of the third body measured in days and  $M_3$  the third body mass in solar masses. The consistent range of  $e_3$  and  $i_3^\circ$  now depends on the manifestation of the third body orbit on the inner binary. The outer body orbit is unlikely to produce the superorbital periodicity on  $P_3$ , because of course this superorbital periodicity varies by  $\sim 10\%$ . A more likely mechanism is precession of the eccentricity of the inner binary, which occurs on a timescale  $T_{mod}$  given by (Mazeh & Shaham 1979):

$$T_{mod} = \left( \frac{M_1 + M_2}{M_3} + 1 \right) \frac{P_3^2}{P_{21}} \quad (5.5)$$

We classify third-body scenarios into three groups depending on the identification of SMC X-1 components with the 50d and  $\sim 7$ -year periodicities.



**Figure 5.9:** Upper limits on the inclination of third body orbits, as a function of the outer orbital eccentricity, for the three scenarios discussed in the text. Scenario 1 (solid line):  $T_{mod} = 50\text{d}$ ,  $P_3 = 3.89\text{d}$ . Scenario 2 (dashed line):  $P_3 = 50\text{d}$ ,  $T_{mod}$  unconstrained. Scenario 3 (dotted line):  $T_{mod} = 7\text{years}$ ,  $P_3 < 43.9\text{d}$  from  $M_3 > 0.08M_\odot$ .

(1) Identifying  $T_{mod}$  with the  $\sim 50$ -day superorbital period then constrains the allowed range of  $M_3$ , as the condition  $P_3 > P_{21}$  requires  $M_3 \gtrsim 1.5M_\odot$ . The outer body orbit then requires a high degree of fine-tuning, with  $i_3^\circ > 2^\circ$  only if  $e_3 > 0.98$ .

(2) Alternatively, we might identify  $P_3$  with the  $\sim 50\text{d}$  superorbital cycle if some mechanism, probably disk-related, is invoked to translate a stable orbital periodicity into an unstable variation in inner disk/boundary layer X-ray brightness). Ignoring the seven-year variation for the moment, we can estimate that for  $M_3 > 0.08M_\odot$ , for any  $i^\circ > 2^\circ$ , the system must still have  $e_3 \gtrsim 0.8$ . If we further identify  $T_{mod}$  with the  $\sim 7$ -year variation in superorbital periodicity, equation (5.5) then fixes  $M_3$  at  $8.2M_\odot$ , which requires even more severe fine-tuning; then  $e_3 < 0.98$  only if  $i_3^\circ < 0.4^\circ$ !

(3) Even if  $T_{mod}$  is identified with the  $\sim 7$ -year variation in superorbital periodicity and the  $\sim 50$ -day periodicity itself is left to another mechanism, the allowed range of  $e_3$  and  $i_3^\circ$  is still quite restrictive: for example, if we assume  $M_3 > 0.08M_\odot$ , then for  $i^\circ \gtrsim 2^\circ$  the orbit must have  $e_3 \gtrsim 0.76$ .

Thus any third body scenario requires an extraordinary degree of fine-tuning to reproduce the observed results, especially given that the  $\sim 58^\circ$  inclination of the inner binary orbit of SMC

X-1 would then require a  $\sim 32^\circ$  angle between the two orbital planes (figure 5.9). The warped, precessing accretion disk has two clear advantage over the other mechanisms considered here: it can support variations in superorbital period, and does not require such fine-tuning. To show this, we first briefly remind the reader of the phenomenology of warped, precessing disks; the reader is referred to the discussion in Chapter 2 for more detail.

#### 5.4.2 Stability of Precessing, Warped Accretion Disks

Modelling of disk behaviour by Wijers & Pringle (1999) and subsequent more rigorous treatment by Ogilvie & Dubus (2001: hereafter OD01) has shown that an accretion disk is unstable to radiation-driven warping when the accretion efficiency of the central source exceeds a critical value. The most pertinent of the OD01 results for this discussion is the classification of behaviour into modes of warping (Chapter 2,6). The stability of the accretion disk to warping is predicted as a function of the mass ratio  $q$  and the binary separation  $r_b$ . SMC X-1 is then placed near the border between the region supporting stable mode 0 warping and that supporting mode 1 and higher modes (see figure 2.1 and the associated discussion in 2.2.3). We would thus expect the disk to form a precessing warp, giving rise to a long periodicity. However, this may not be a stable monotonic warp; instead, there may be an interaction between modes by virtue of the unique location of SMC X-1 near the border between the two regions. As two or more modes compete, then, a longer term variation may be set up in which the sum of modes itself changes slowly with time. This may manifest itself as the observed variation in the superorbital period of SMC X-1, which must represent an important constraint on future development of such models.

#### 5.4.3 How the precessing warp manifests itself

Systems in which a warped, precessing disk brings about long-term behavior are usually thought to have as their primary observational effect a change in uncovered emitting area in the inner disk (c.f. Gerend & Boynton 1976). The Chandra/HST study of Vrtilek et al (2001) shows a clear spectral distinction between the high and low states of the superorbital cycle, in that the continuum from the pulsar dominates during the high in the cycle, while during the low the X-ray spectra are dominated by an additional component, too faint to stand out in the high state. This provides further support to the scenario of variation in uncovered emitting area: as the disk precesses it blocks the neutron star. The BATSE flux is also modulated; although the cross-section

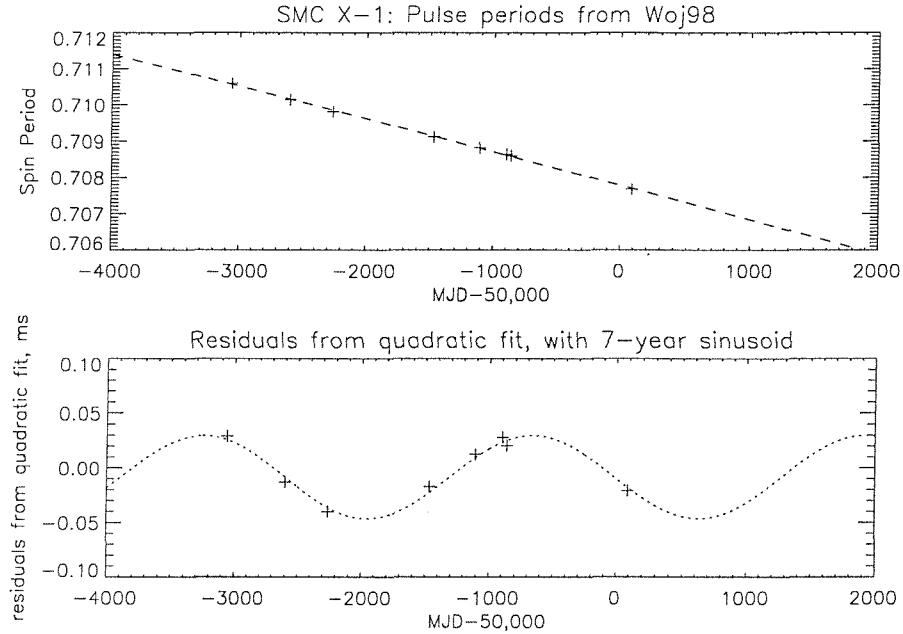
for photoelectric absorption is smaller by a factor  $\sim 100$  at energies  $> 20\text{keV}$  than at  $\sim 2\text{keV}$  (Eisberg & Resnick, 1985; Zombeck, 1990), the passing of the line of sight across the plane of the accretion disk may still allow absorption at these wavebands by virtue of column densities many orders of magnitude above those above the surface of the disk. Note also that this extensive disk structure out of the plane may have led van Paradijs & Kuiper (1984) to infer a much thicker disk than is really the case.

The location of the emission region of the broad emission lines seen in X-ray eclipse is not immediately clear. One obvious possibility is excitation in wind material from the giant donor. The rate of heating of such material depends on the instantaneous accretion rate at the point source, the ionisation parameter  $\log(L/nr^2)$  (WCK2000; recognising also that there may be other sources of X-ray luminosity than just the point source), and also weakly on the angular distribution of the resulting X-radiation. Hydrodynamic simulations of the ion density  $n^2$  in the donor wind (Blondin & Woo 1995) suggest that for  $L_X \sim 6 \times 10^{38}\text{erg s}^{-1}$ , significant amounts of gas with this range of ionisation parameters exist in the orbital plane, with a small component above the donor (see figure 11 of Wojdowski et al 2000 for qualitative comparison). The exact location of this material in the orbital plane is unknown, as the BW95 simulations do not account for shielding by any accretion disk, which would bring these regions nearer the neutron star than the simulations would suggest. As the binary rotates, the observed flux from the recombination region varies by virtue of changing its projected area along the line of sight and also from intrinsic structure in the wind (e.g. Wojdowski et al 2000). This may also explain the apparent asymmetry in the low state eclipse profile, which is steeper after eclipse than before it (figure 5.1, right). The ACIS count rates recorded during eclipse appear to vary by a factor  $\sim 2$  between the high and low states in the superorbital cycle. The count rates of the lines vary between high and low states but not the spectral shape; in particular there does not appear to be any absorption cut-off in the eclipse spectrum in the low state. That the ROSAT and Ginga observations are consistent with this factor 2 difference suggests it is associated with the superorbital variation, and not attributable to random short-term variation in the wind. This variation might be brought about by the precession of the accretion disk, by variation of disk shielding of the wind material in the equatorial plane.

However, the optical lightcurve is dominated by the accretion disk, ellipsoidal variations and heated face of the donor, with no need for a wind density variation. It is thus unclear if the wind can produce the observed variation in X-ray line flux. Another site for the recombination lines might be an accretion disk wind (possibly the “circumsource material” suggested in Wojdowski

et al 1998). This would probably also require some shielding by the accretion disk to bring the ionisation parameter  $\log(L/nr^2)$  down to the range 1-3 required for production of the lines; in the neighbourhood of the pulsar, values are typically a factor  $\sim 100$  higher (Blondin & Woo 1995). During the low superorbital state, line fluxes are a factor  $\sim 20$  lower during eclipse than out of eclipse (Vrtilek et al 2001), which is consistent with observed ADC sources (e.g. X1822-371; White & Holt 1982). The physical extent of the accretion disk in the prototypical ADC source X1822-371 may be estimated from its system parameters ( $P_{orb} = 0.232$ d and  $M_2 \sim 0.4M_\odot$ , therefore  $a \simeq 1.9R_\odot$ ), with the result that if the disks in SMC X-1 and X1822-371 fill their Roche lobes by the same fraction, the disk in SMC X-1 should be larger than that in X1822-371 by a factor  $\sim 6$ . Thus a significant disk wind should be expected in SMC X-1. If  $R_{disk} \gtrsim 0.55$ , then the disk will not be occulted by the donor during pulsar eclipse. However, the disk must have significant structure out of the orbital plane, widening the range of  $i^\circ$  for which the line-emitting disk wind is not obscured. Thus, the observed source inclination of  $\sim 58^\circ$  probably does allow the disk wind to be unobscured during eclipse of the pulsar. What challenges the disk-wind mechanism is the factor 2 variation in eclipse count rates with superorbital cycle; if the same disk wind is viewed in both states there should be no difference in the recombination flux measured. Variation in projected area of the emitting area with disk precessional phase might be a possible mechanism for this difference. A further challenge is then to explain the asymmetry in the orbital profile during the low state (figure 5.1, right).

As noted above, the *variation* in the superorbital cycle is probably due to a change in the shape of the warp, by virtue of the likely competition of warping modes. This would lead to a variation in the shadowing and therefore heating of the surface of the accretion disk. This would change the rate of reprocessed emission from the disk, thus changing the radiation-induced torque, bringing about the observed change in accretion disk precession rate. A further consequence of this heating would be to change the viscosity parameter  $\alpha$ , and thus change  $\dot{M}$  through the disk. This is not in conflict with the spin-up of the neutron star, as there are deviations from a purely monotonic increase (Wojdowski et al 1998). Indeed if one looks at the residuals in the quadratic fit to the pulse periods measured by Ginga, ASCA, ROSAT and RXTE (Wojdowski et al 1998), the residuals show variation of similar shape and timescale to the variations of the superorbital X-ray period! The published fit in table 6 of Wojdowski et al (1998) does not match the observed spin periods, which is probably due to a typographical error, as the fit in the *figures* does reproduce the spin period evolution. We have therefore re-fit the published spin periods from the beginning of the



**Figure 5.10:** Fit to spin periods of SMC X-1 since the beginning of the BATSE mission. **Top:** Spin periods measured by Ginga, ROSAT, ASCA and RXTE, as tabulated by Wojdowski et al (1998). Dashed line represents quadratic fit to the spin periods. **Bottom:** residuals (in ms) from the fit, with a 2583-day sinusoid overplotted (compare with figure 5.4)

BATSE observation interval to a quadratic, yielding the following ephemeris:

$$P_{sp} = 0.707779 - 9.33267 \times 10^{-7}t - 7.49741 \times 10^{-12}t^2 \text{ s} \quad (5.6)$$

where  $t$  is the time in days from the reference MJD 50000.00. The residuals of this fit are modeled with a sinusoidal variation with a seven-year period, i.e. the timescale for the variation in superorbital period. These residuals are well represented by a sinusoid at 2583 days (figure 5.10). This cycle was based on pointings up to the beginning of the RXTE/ASM dataset; if pointings taken during the interval since 1996 affirm this cycle, it would provide direct evidence linking the variation in third period to variation in accretion flow onto the neutron star.

## 5.5 Conclusion

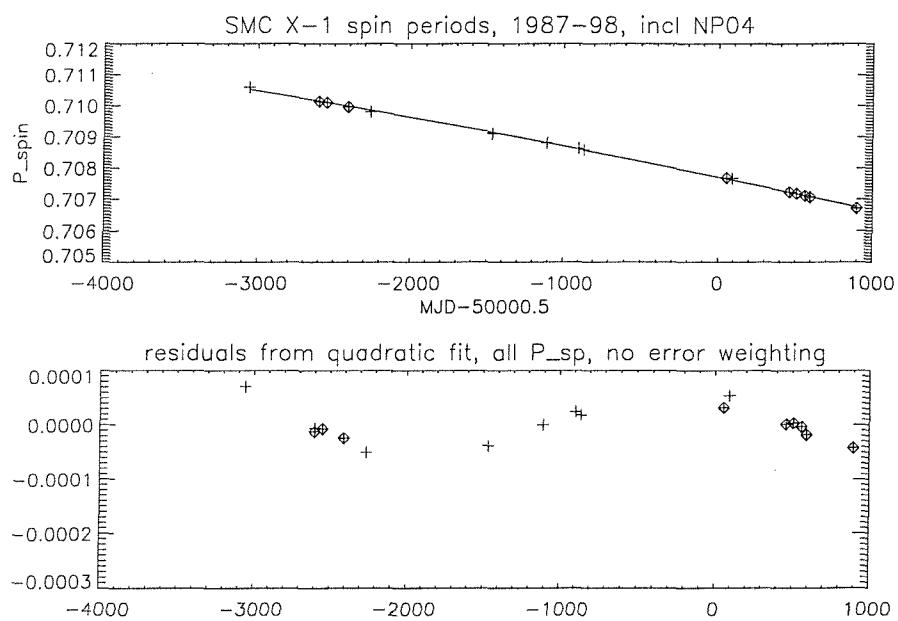
We have used the dynamic power spectrum to show that the superorbital period of SMC X-1 varies in a coherent, apparently almost sinusoidal, way. In conjunction with the spectral behaviour of this variation, we have seen that a precessing, warped accretion disk is likely to be responsible for the superorbital variation. However the influence of the warp may be felt not only through varying occultation, as is supposed for Her X-1, but through varying accretion at the neutron star

boundary layer, permitted because the disk precession is not steady but quasi-steady. Finally we saw that the variation of the superorbital period is consistent with current stability theory of warps in accretion disks. An interaction of warp modes may give rise to the variation in superorbital period, allowing its variation to change direction several times in a few years.

## 5.6 Update: Further Spin Evolution of SMC X-1

The remarkable apparent correlation between the spin period of the neutron star and the value of the superorbital period in the long-term X-ray lightcurves is open to investigation and confirmation. The spin periods used for the identification of the spin-period cycle were taken from the RXTE/PCA investigation and literature review in Wojdowski et al (1998), and terminate at the beginning of the RXTE era. Recently, seven more points have been added to the spin-evolution of SMC X-1 through BeppoSAX observations and further archival observations not previously tabulated (Naik & Paul, 2004). When the spin evolution of SMC X-1 is re-fit with the new points taken into account, we see a similar sinusoidal fit to that given in this chapter, with a slightly longer period - but still consistent with the superorbital period evolution timescale noted above (figure 5.11).

Clearly the next step is to follow the evolution of the spin period into the interval covered by RXTE/ASM to establish if a direct correlation between the superorbital period and the spin-up of the neutron star really exists. Since 1996 there have been at least 30 RXTE/PCA pointings of SMC X-1 at various phases in the superorbital cycle, all of which have been archived at HEASARC. Furthermore, the USA satellite, which provided similar capability to the RXTE/PCA and ASM instruments at 2-10keV for two years, has obtained both long-term observations and pointings of SMC X-1. Initial work has been performed by the USA team analysing torque variations in SMC X-1 (Wolff et al., 2000), however due to funding termination upon the cessation of scientific activities with the satellite, the effort necessary to reduce and publish this data has been lacking. Initial steps have been made to gain access to the USA pointings of SMC X-1, which, combined with the RXTE/PCA archives, should allow the period since 1996 to be populated with  $\sim$ 10-20 spin period measurements. Derivation and analysis of spin periods from RXTE/PCA data has begun with collaborators at Southampton, while ARGOS/USA data will be made available to us for spin-period derivation (Bandyopadyay 2004 priv. comm.) This will allow the relationship between the accretion disk warp behaviour and the mass transfer rate to be definitively established



**Figure 5.11:** Spin periods and their residuals from quadratic fit, with the addition of the new points reported by Naik & Paul (2004). Crosses represent points tabulated by Wojdowski (1998), diamonds are the new points. The quadratic has been re-fit to take the new points into account; as can be seen the residuals show variability on a similar timescale to the variation in superorbital period noted in this chapter.

- or disproven.

## Chapter 6

# Stability of XRB to Warping

### 6.1 Introduction

In Chapter 5 we used a dynamic power spectrum (DPS) analysis of long-term X-ray datasets (CGRO/BATSE, RXTE/ASM), focusing on the remarkable, smooth variations in the superorbital period exhibited by SMC X-1. Here we use a similar DPS analysis to investigate the stability of the superorbital periodicities in the neutron star X-ray binaries Cyg X-2, LMC X-4 and Her X-1, and thereby confront stability predictions with observation. We find that the period and nature of superorbital variations in these sources is consistent with the predictions of warping theory establishing an observational basis on which to develop models further.

We also use a dynamic lightcurve analysis to examine the behaviour of Her X-1 as it enters and leaves the 1999 Anomalous Low State (ALS). This reveals a significant phase shift some 15 cycles before the ALS, which indicates a change in the disk structure or profile leading into the ALS. The counterpart paper to this chapter has been published in MNRAS (Clarkson et al., 2003a).

#### 6.1.1 Accretion Disk Warping

As established in Chapter 5, SMC X-1, shows a clear variation in its superorbital period. The dynamic power spectrum (or DPS) was used to chart the behaviour of the superorbital period. An



apparently cyclic variation (on an even longer timescale of  $\sim 7$  years) in the superorbital period was evident in the DPS. It was suggested that the behaviour of SMC X-1 is consistent only with a precessing, warped accretion disk, where a competition of warping modes leads to variation in the warp, giving rise to the shifting superorbital period.

This interpretation gains support from the theoretical stability analysis of Ogilvie & Dubus (OD01), which extends  $\alpha$ -disk theory to radiation-driven warping, with correct analytic expressions for the torque on a disk element. Solutions are examined numerically to determine stability to warping for initially flat disks, steadily precessing disks and nonlinear warps. They find that for reasonable values of the global disk viscosity  $\alpha$  and accretion efficiency  $\eta$ , the stability of an accretion disk to warping is determined by the binary separation  $r_b$  and mass ratio  $q = \frac{M_2}{M_1}$ . A further variable is the location of the mass input to the disk, whether close to the Lindblad resonance point or the circularisation radius  $r_c$ . For most X-ray binaries, the majority of the torque exerted by the mass input is close to  $r_c$ , as even in LMXB with large accretion disks, the matter stream is believed to continue to approximately the circularisation radius (for example, the SPH simulations of Kunze et al., 2001, showed a substantial fraction of the matter stream hitting the disk near  $r_c$  for a wide range of binary mass ratios). Accordingly, we follow OD01 in taking the location of mass addition to be  $\sim r_c$ . Figure 2.1 shows the division of XRB into regions of stability to warping from this analysis. The location of SMC X-1 at the boundary between regions stable to simple mode 0 warping and combinations of modes supports our assertion that the long-term X-ray behaviour of SMC X-1 is due to a competition of warping modes.

The location of the boundaries between stability regions in  $r_b$  -  $q$  space depends strongly on the global disk viscosity parameter,  $\alpha$  and the neutron star accretion efficiency  $\eta$ . Uncertainty in these values translates into uncertainty as to the location of the boundary between regions and thus the stability classification of a system. However, the adopted values of  $\alpha = 0.3$  and  $\eta = 0.1$  are usually assumed to be appropriate for neutron star XRB (Frank, King & Raine 2002, OD01 and references therein).

In this chapter we examine the long-term behaviour of a selection of X-ray persistent neutron star XRB in which superorbital periodicities are well-established. We apply the techniques from Chapter 3, in particular the DPS, to explore the phenomenology of these variations. We focus on neutron star XRB for which persistent superorbital periods have previously been reported. Our goal is to relate the long-term behaviour of the sources to their position in the  $r_b$  -  $q$  diagram, so as to directly confront the stability predictions of the radiation-driven warping framework with

observation. If successful, it will establish the time evolution of periodicities as a powerful indicator of XRB properties in systems whose properties are less well-understood, for example the Ultracompact X-ray Binaries.

## 6.2 Source Selection

Each source used in this analysis was selected on the basis of three criteria. Firstly, independent estimates of  $r_b$  and  $q$  must be available, which restricts our search to optically identified systems with full orbital solutions. Secondly, the sources must be persistent X-ray emitters at a high enough level to provide long-term datasets of sufficient quality to detect long-term modulations, which favours neutron star XRB. Third, each system used must show independent evidence for a persistent accretion disk.

Table 6.1 shows the sources selected for this study, with the warping behaviour predicted by OD01 in table 6.2. The strong dependence of the stability boundaries on  $\alpha$  and  $\eta$  leads us to choose systems with similar compact objects in the hope that for such systems these parameters will take similar values. Poor coverage of the  $r_b$  -  $q$  diagram forces us to select systems with both high and low - mass companions.

Figure 6.1 shows sections of the RXTE/ASM lightcurves from the sources used here, binned appropriately to improve signal-to-noise. The lightcurve of Cyg X-2 shows similar morphology to that of SMC X-1, though at count rates about a factor of 20 higher. Its morphology is also more complex than a single periodic cycle, which leads to the detection of more than one periodicity (Paul et al., 2000). The ASM lightcurve for Her X-1 shows a clear variation on the well-known 35 day cycle, with a secondary peak half a cycle after the main period of extended activity. LMC X-4 also shows evidence of a cycle that is superficially similar to that of SMC X-1, although the lightcurve shows more scatter.

### 6.2.1 Discarded Sources

Several well-constrained, persistent neutron star XRB are within the regions of interest in the  $r_b$  -  $q$  diagram, but we choose not to analyse them here as they do not show significant long-term

cyclic phenomena in the RXTE/ASM. Circinus X-1 is well within the region of the stability diagram corresponding to mode 1 and higher warping, and is a luminous neutron star XRB (Tennant 1986). However, its orbit is strongly suspected to be highly eccentric, among many other properties the system shares with A0538-66 (Charles et al. 1983). As pointed out in OD01, the stability analysis as it stands is *not* valid for non-circular orbits (indeed, in such systems we would expect any superorbital behaviour to be dominated by the orbit through tidal effects).

4U 1907+09 is a pulsar XRB with massive companion (e.g. van Kerkwijk et al 1989) in a wide orbit of moderate eccentricity ( $e \sim 0.16 \pm 0.14$ ; Makishima et al 1984), that has shown a monotonic *increase* in pulse period since its discovery (Mukerjee et al 2001). The X-ray brightness of this source increases by a factor  $\sim 4$  at a probable interval of 0.5 orbits (in't Zand et al. 1998, Makishima et al 1994, Mukerjee et al 2001), at orbital phases 0.195 and 0.695 after periastron, remaining at the high count rate for  $\sim 0.3$  orbits. QPOs have been discovered at frequencies  $\sim 10 - 400$  mHz, but only during these “flaring” intervals. Formation of a transient accretion disk has been suggested (the 18s QPO, for example, corresponds to a Keplerian orbit with radius roughly 10% of the corotation radius; in 't Zand et al. 1997). This accretion disk, if rotating in the opposite sense to the pulsar, could provide the observed rate of spin-down if it persists for a few percent of the orbit (a retrograde disk has also been posited to explain the correlation between luminosity and spin-down in the disk-fed accreting X-ray pulsar GX1+4; Chakrabarty et al. 1997). The disk must be retrograde if accretion is to occur, as the 18s QPO suggests matter must pass through the corotation radius before accretion. The disk forms twice per orbit as the neutron star passes through the equatorial plane of the donor, where the density of ejecta is greatest, as suggested from the observation that the eccentric wind-fed pulsar GX301-2 also shows two prominent X-ray flares in each orbit (Koh et al 1997). Such a transient disk is unlikely to persist for long enough to produce radiation-driven disk warps measurable with our RXTE/ASM analysis.

X2127+119 has recently been shown to consist of two sources which cannot be resolved by RXTE (White & Angelini 2001; Charles, Clarkson & van Zyl 2002). While resolving the mystery of how a source can show evidence for both an ADC and Type I bursts, it is not possible to relate components in the observed RXTE/ASM X-ray flux to a single system.

OAO1657-415 is a 38s X-ray pulsar in a 10.4d orbit with a B-type supergiant companion (Chakrabarty et al., 2002). It shows marginal evidence for a superorbital periodicity of 16 days (Laycock 2002 priv. comm.). It also shows spin-up and spin-down trends without any correlation with X-ray luminosity  $> 20$ keV (Baykal, 2000), suggesting no *persistent* accretion disk is present. Transient

**Table 6.1:** Neutron star XRB exhibiting superorbital modulations

Name	$r_b$ ( $\times 10^6 GM_1/c^2$ )	$q$	$P_{sup}$ <sup>1</sup> (d)	$M_X$ ( $M_\odot$ )	$P_{orb}$ (d)	$-\dot{M}_2$ ( $\times 10^{-8} M_\odot yr^{-1}$ )
Her X-1	3.1	1.56	35	1.36	1.700	1.4
LMC X-4	4.5	10.6	30	1.35	1.408	3.5
SMC X-1	8.9	11.0	50	1.6	3.890	2.3
Cyg X-2	8.0	0.34	78	1.78	9.844	2

<sup>1</sup> Claimed superorbital period

**Table 6.2:** Predictions of Radiation-Driven Warping model

Name	Prediction <sup>1</sup>
Her X-1	Persistent mode 0
LMC X-4	Persistent mode 0
SMC X-1	Mode 0 dominant but other modes possible
Cyg X-2	Complex combination of modes 0,1 and higher

<sup>1</sup> Based on position in figure 2.1.

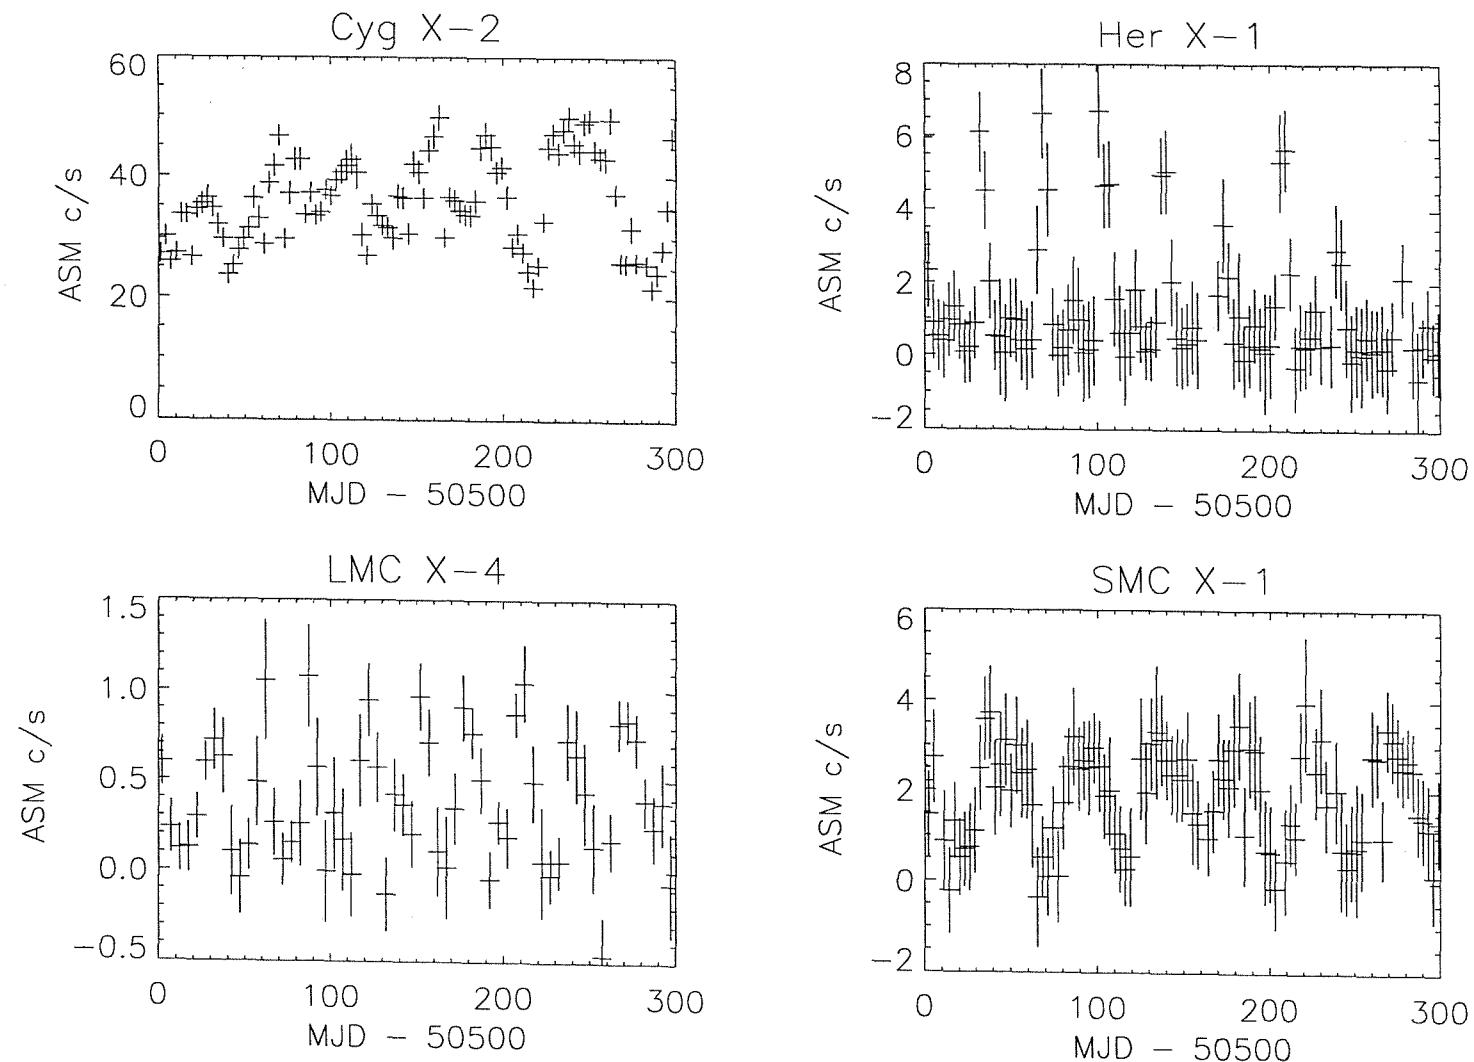
episodes of X-ray brightness increase, coupled with spin-up, are thought to arise through formation and destruction of a transient, prograde accretion disk, in the manner implied by Bildsten et al (1997). As in the case of 4U1907+09, such a disk is unlikely to support radiation-driven disk warps measurable by RXTE/ASM.

Finally, as the purpose of this work is to confront the radiation-driven warping framework with observation, we further restrict our analysis to systems of sufficient separation to allow such warping. The ultracompacts X1820-303 and X1916-053 are thus deferred to a future publication (Clarkson et al 2004 in prep).

### 6.3 Dynamic Power Spectrum

Data were analysed using our DPS approach, designed to be sensitive to long-term quasi-periodic variations (Chapter 3). Choice of data window length was influenced by the limits of reliable period search with the PDS; we choose two complete cycles as the minimum needed for a reliable detection. We choose 400-d windows, the maximum length that meets our criteria because this brings about the sharpest period detections.

The error in the detected period was determined empirically from Monte Carlo simulations in the



**Figure 6.1:** Sections of the RXTE/ASM lightcurves of the sources examined in this work, binned appropriately to emphasise their long-term behaviour.

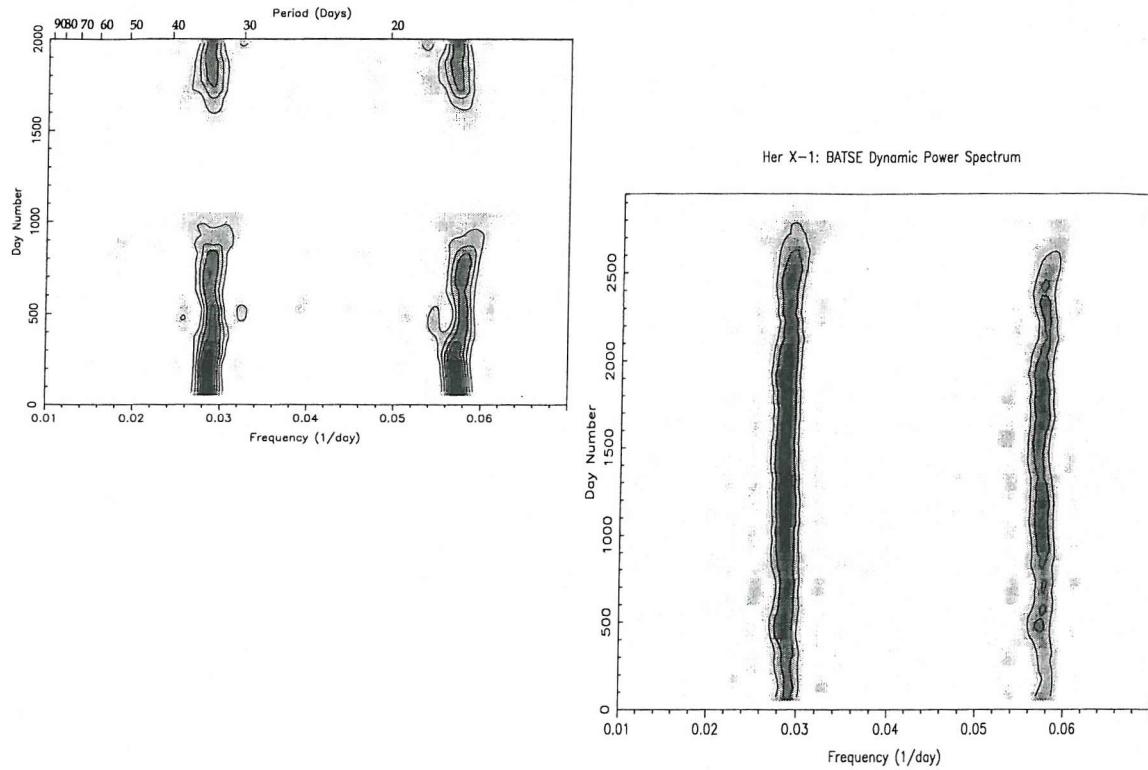
manner described in Chapter 3. The detection accuracy is quoted in units of  $3\sigma$  of the fit to the resulting distribution in period detections. In one case (Cyg X-2) the peak period detection was shifted from the input period by an amount corresponding to the beat between the detected period and  $\sim 8$  times the time interval of a data window, which has been incorporated into the quoted error on the period detection (table 6.3).

## 6.4 Results

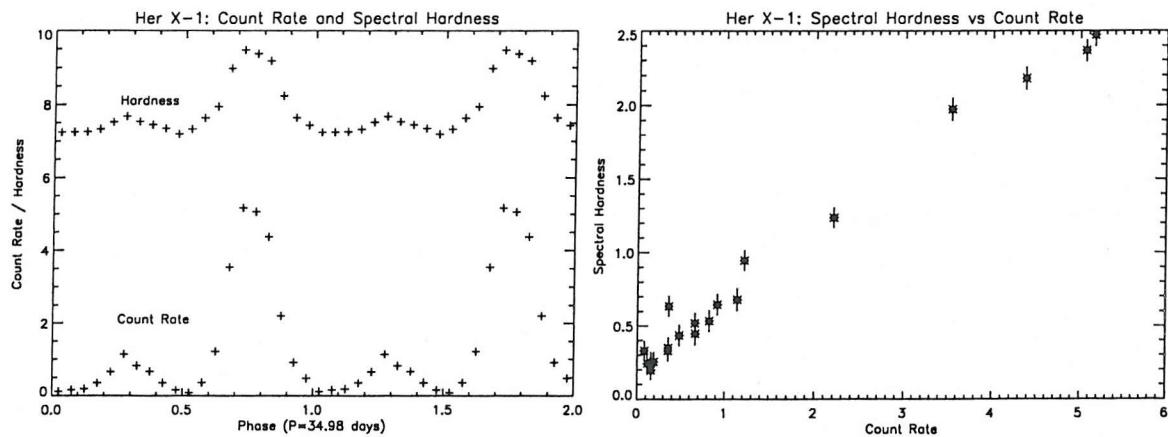
### 6.4.1 Her X-1

Her X-1 is the prototypical disk warping system, with a  $\sim$ 35-day periodicity detected at optical and X-ray wavelengths. This periodicity is almost certainly due to a precessing, warped accretion disk, as first suggested by Gerend & Boynton (1976) to explain the optical lightcurve, and corroborated by the evolution of the pulse profile (Scott et al., 2000; Ramsay et al., 2002), and variations in X-ray illumination of the companion HZ Her (Leahy et al., 2000; Still et al., 2001). The accretion disk is thought to be tilted and twisted (Gerend & Boynton 1976, Schandl & Meyer 1994, Schandl 1996). With X-ray pulsations at 1.24s, eclipses at 1.7 days and an A7 Roche lobe-filling companion (Middleditch & Nelson, 1976) of  $\sim 2 M_{\odot}$ , the mass ratio and binary radius are constrained (but see Leahy & Scott, 1998, for uncertainties due to possible non-synchronous donor rotation). Geometric arguments based on optical lightcurves suggest  $1^{\circ} \gtrsim 74.9^{\circ}$  (Bahcall & Bahcall, 1972). On three occasions, the source has failed to reach a high state at the time expected from the 35-day cycle (Parmar et al., 1985; Vrtilek et al., 1994; Oosterbroek et al., 2000). These Anomalous Low States (ALS) last for several 35-day cycles each, and are thought to be the result of a state change in the accretion disk (Still et al 2001).

The DPS of Her X-1 (figure 6.2) shows an extremely significant ( $>>99.99\%$ ) periodicity at 34.98 days, along with a less significant secondary peak at 17.5 days. This secondary peak is not an artefact of the search technique; as figure 6.3 shows, the DPS has picked up the secondary peak at phase 0.3 in the 35-day cycle. The system entered the third observed ALS in January 1999; the effect on the ASM emission can be clearly seen in figure 6.2. Her X-1 is the only system suspected of disk precession to show this secondary feature so significantly. Her X-1 also shows the clearest relation between ASM count rate and spectral hardness of all the sources studied here (figure 6.3).



**Figure 6.2:** DPS of Her X-1. Left: RXTE/ASM dataset, contours spaced at intervals of LS power 100. Day numbers run from the beginning of the RXTE/ASM dataset, i.e. MJD 50088. Right: CGRO/BATSE dataset, contour spacing LS powers of 20. Day numbers run from the beginning of the CGRO/BATSE dataset, i.e. MJD58369.



**Figure 6.3:** Superorbital lightcurve (left) and count rate-hardness relation (right) for Her X-1. The spectral hardness varies with the ASM count rate and there is a clear relation between the two quantities.

### 6.4.2 LMC X-4

LMC X-4 is a  $\sim 13.5$  s X-ray pulsar with a  $15 M_{\odot}$  O-type companion that eclipses the X-ray source every 1.4 days (Hutchings et al., 1978). QPO's have been found at 2-20 mHz during large X-ray flares (Moon & Eikenberry, 2001), confirming the presence of orbiting matter at least  $\sim 0.25 r_{circ}$  from the neutron star. Disk modeling of the orbital and precessional optical lightcurve (Heemskerk & van Paradijs, 1989) suggest the accretion disk may extend out to about  $4r_c$ , and is very likely to show both precession and warping. Recent analysis of ASCA spectra show that the region producing the 6.64 keV iron-line complex varies with the superorbital cycle, suggesting warping of the inner disk region (Naik & Paul, 2003).

The DPS of LMC X-4 (figure 6.4) shows the simplest behaviour of the sources analysed here: a periodicity at 30.28 days that is both stable and persistent throughout the entire dataset. The DPS is not sensitive to the  $\sim 0.02\% yr^{-1}$  systematic variation of this periodicity reported recently by (Paul & Kitamoto, 2002). At LS power  $\sim 38$ , this 30-d cycle is somewhat less significant than for Her X-1 (LS Power  $\sim 300$ ), but the power still corresponds to  $> 99.9\%$  significance. There appears to be slight jitter in the peak period, but at  $\pm \sim 0.05$  days this is not significant when compared with the  $3\sigma$  uncertainty for this source of  $\pm 0.46$  days. LMC X-4 shows a relationship between RXTE/ASM spectral hardness and count rate (figure 6.5) that is somewhat less steep than for Her X-1.

### 6.4.3 Cygnus X-2

Although not a pulsar, Cyg X-2 nevertheless has a well-constrained binary orbit, with a 0.5 - 0.7  $M_{\odot}$  companion in a 9.84-day period (Casares et al., 1998), unusually long for a LMXB (Liu et al., 2001). Type I X-ray bursts (Kahn & Grindlay, 1984) established Cyg X-2 as a neutron star XRB. Modeling of IUE spectra (Vrtilek et al., 1990) yields a best-fit outer accretion disk radius of  $\sim 0.6 R_L$  (where  $R_L$  is the Roche lobe radius), while modeling of optical lightcurves (Orosz & Kuulkers, 1999) suggests the disk radius may be closer to  $0.9 R_L$ . A pattern of high- and low-intensity X-ray states has been seen in early RXTE/ASM, Vela 5B and Ariel-5 observations (Wijnands et al., 1996), on timescales of  $\approx 77.7$ ,  $\approx 77.4$  and  $\approx 77.7$  days respectively, while the Ginga/ASM lightcurve shows the strongest periodicities on periods of 53.7 and 61.3 days (Paul et al., 2000). There is still debate over the nature of this variation; the burst activity and QPO-state are *not* correlated

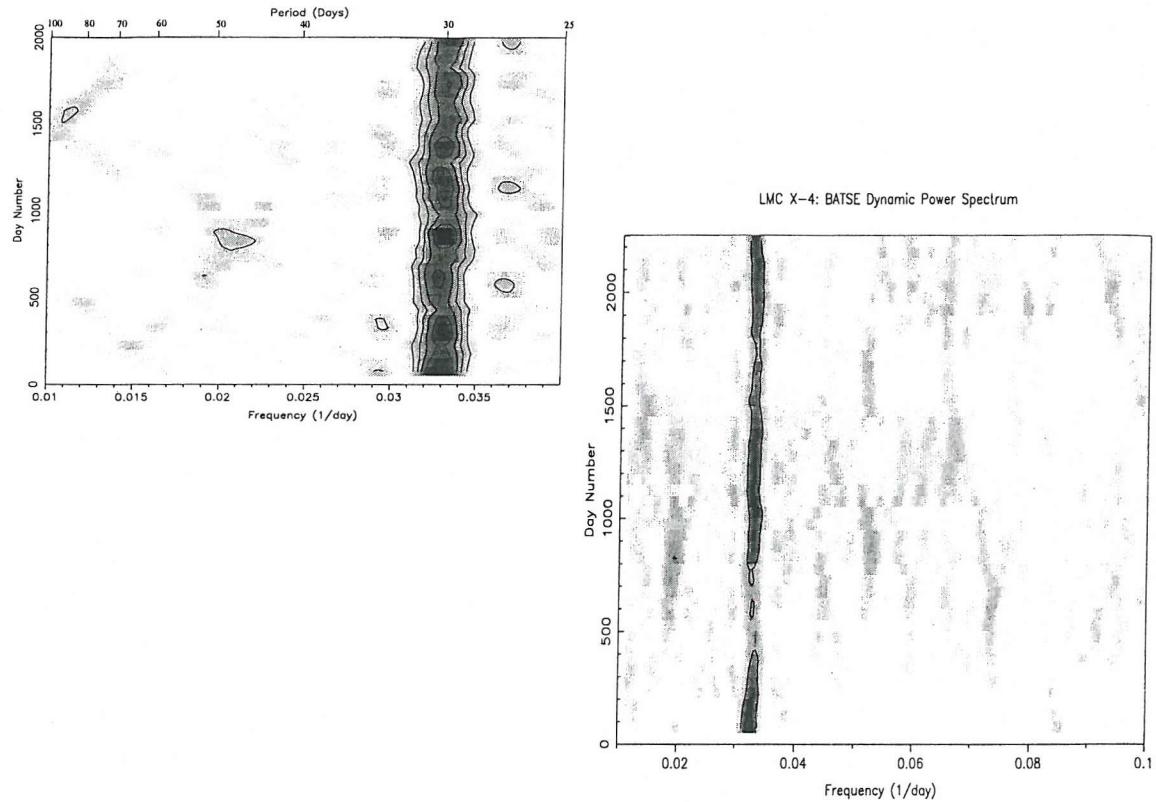


Figure 6.4: DPS of LMC X-4. Contours spaced in intervals of LS power 10 (by comparison, LS power 9 corresponds to 99.9% significance over noise - chapter 3)

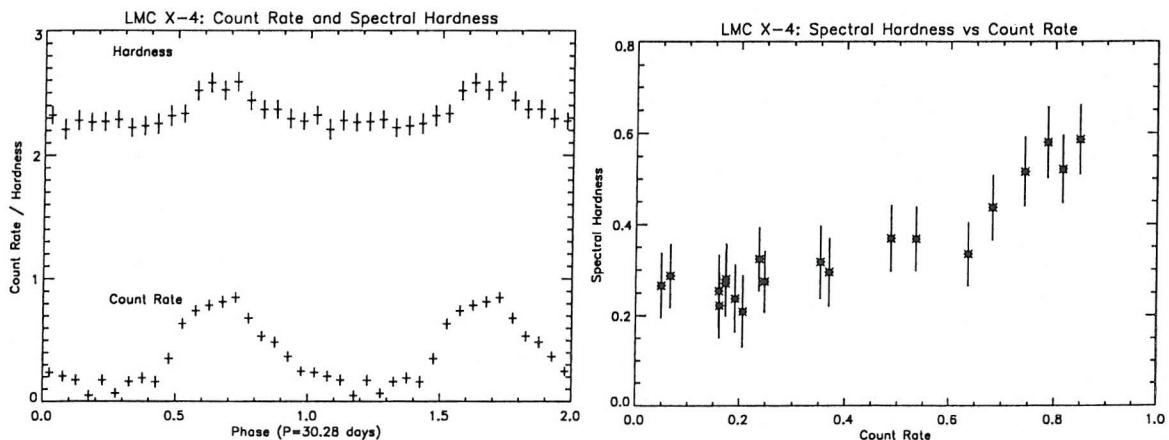


Figure 6.5: Superorbital lightcurve (left) and count rate - hardness relation (right) for LMC X-4. Though a more noisy dataset, there is a correlation between the variations of the two quantities.

**Table 6.3:** Detected periodicities

Name	$P_{sup}$ (d)	$\Delta P_{sup}$ (d) <sup>1</sup>
Her X-1	35.0	0.16
LMC X-4	30.3	0.46
Cyg X-2	59.0	0.45

<sup>1</sup>  $3\sigma$ 

with the X-ray cycle (Wijnands et al 1997, Kuulkers et al 1996), suggesting this variation does not represent variation in  $\dot{M}_1$ ; however the Z-pattern of variability does vary with the  $\sim 78$ d variation. The variation has instead been attributed to accretion disk warping and precession, though there is still debate about this (Vrtilek et al., 1997). Cyg X-2 is not detected by CGRO/BATSE, unlike Her X-1, LMC X-4 and SMC X-1.

Cyg X-2 shows complex behaviour in the DPS, with an extremely significant ( $LS > 900$ ) periodicity at period  $\sim 59$  days, seen over the interval from day number 600-900 (figure 6.6). This feature seems to form as two separate periodicities converge, then diverge as the feature vanishes. There are also significant recurring features at  $\sim 40$  days. The  $\sim 77$ -day periodicity noted previously may be present during the beginning of the RXTE/ASM dataset, though is of equal or lower significance than its harmonic. In any case the longterm variation of Cyg X-2 is not the reliable clock suggested by Wijnands et al (1996). In contrast to Her X-1 and LMC X-4, the binned superorbital lightcurve of Cyg X-2 shows a negative relation between the RXTE/ASM spectral hardness and count rate (figure 6.7), as expected if the superorbital lightcurve is due to an obscuration effect (section 6.3).

## 6.5 Discussion

### 6.5.1 Her X-1: Disk Inclination Change?

The clear hardness-intensity relation in the RXTE/ASM lightcurve (figure 6.3) is not by itself an indication of the physical mechanism (Tanaka, 1997). In fact we find no feature of the lightcurve of this source that contradicts the precessing, warped accretion disk scenario of Gerend and Boynton (1976). As the warp precesses, it could uncover regions deeper towards the central accretor, predicting the correspondence between spectral hardness and count rate observed. The asymmetry in the superorbital lightcurve can be brought about by a combination of the system inclination and vertical extent of the warp (Gerend & Boynton 1976, Schandl & Meyer 1994, Schandl 1996).

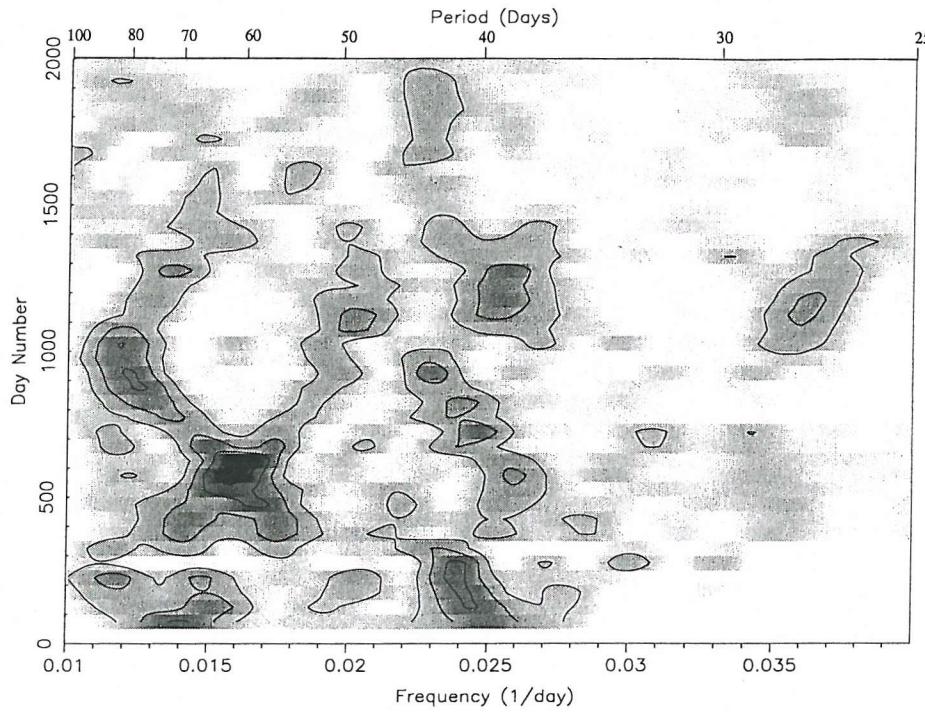


Figure 6.6: DPS of Cyg X-2. Contours spaced at LS powers of 200.

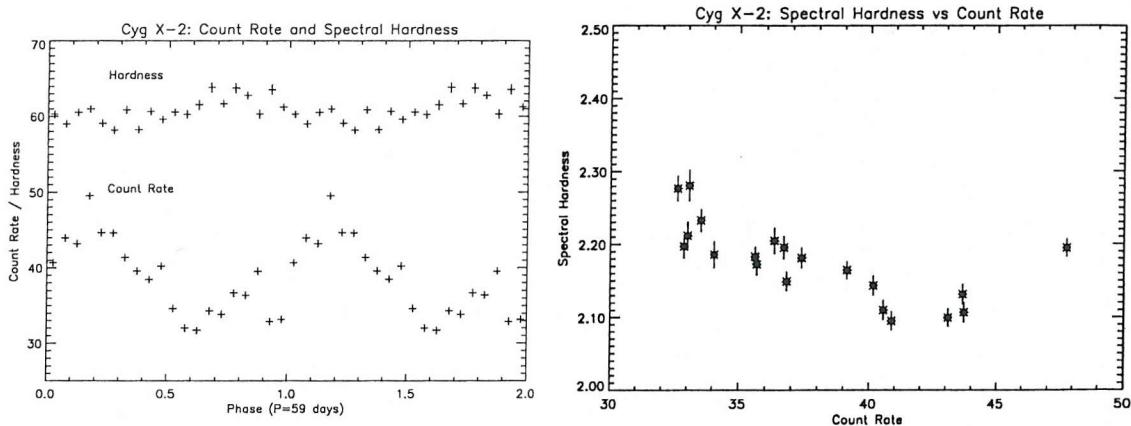
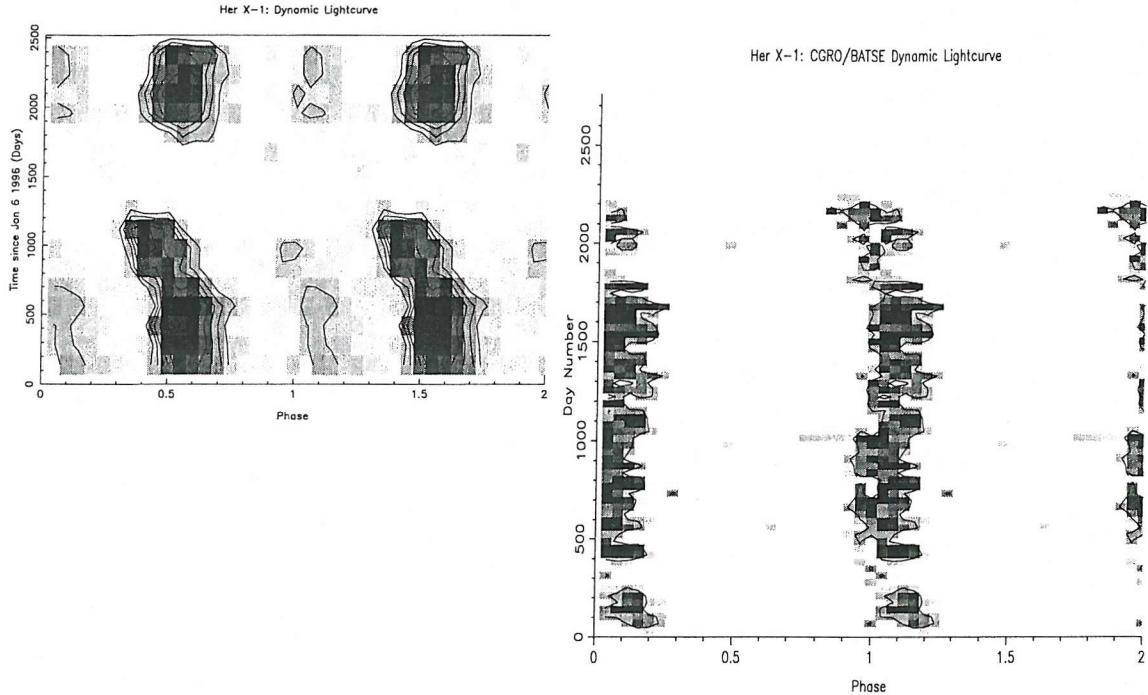


Figure 6.7: Superorbital lightcurve (left) and count rate - hardness relation (right) for Cyg X-2, over day numbers 600-900. In contrast to Her X-1 and LMC X-4, the spectral hardness appears to have a negative correlation with the count rate.

The superorbital periodicity is completely absent from a significant fraction of the data, marking the third ALS detected in this source to date (Oosterbroek et al. 2001). This is not an effect of data sampling, as might be expected for a source with such a low count rate; on the contrary, the persistent emission at  $\sim 1 \text{ cs}^{-1}$  is still well sampled. Measurements by BeppoSAX (Oosterbroek et al 2001) show that the most recent ALS is accompanied by a rapid spin-down of the neutron star, in which the spin-down rate is some nine times faster than the spin-up rate during “normal” X-ray output. At the same time, optical observations during the ALS show that the companion is still strongly irradiated (Margon et al., 1999). Spectral fits and lightcurve modeling of RXTE/PCA

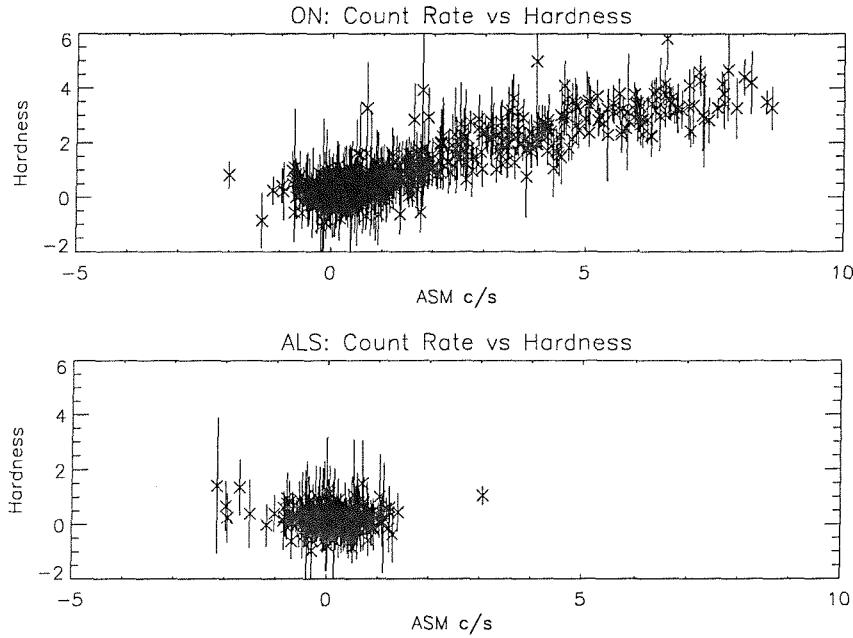


**Figure 6.8:** The 35-day RXTE/ASM and CGRO/BATSE dynamic lightcurves of Her X-1. The phasing of the peaks undergoes a shift before the onset of the Anomalous Low State (ALS), and the relative significance of the peaks alters. When the ALS ends, the phasing is unaltered from the early years of the dataset.

pointings during the ALS suggest the X-ray output during the ALS consists of X-ray reflection from the companion (Still et al 2001). The current interpretation for this behaviour is that the line of sight to the central source has been obscured during the ALS, causing the turn-off, by an increase in line-of-sight column density due to the disk.

This would impact the superorbital lightcurve in two ways. Firstly, the asymmetry between the peaks in the fold would be altered, but whether as an increase or decrease would depend on the current inclination and the precise disk profile. Secondly, a phase shift in the superorbital fold would be seen during the change of inclination. The twist of the disk brought about by the warp means that any given feature in the warp would be seen at different disk azimuth angles for different heights above the nominal flat disk equator.

In figure 6.8 we plot the *dynamic lightcurve* (DL) of Her X-1, in both RXTE and BATSE wavebands. From the DL we see that there is indeed a phase shift taking place over some fifteen cycles of the RXTE/ASM lightcurve, before the ALS. This behaviour seems to also be present in the BATSE lightcurve. The secondary peak in the 35-day cycle appears to decrease in relative amplitude just before onset of the ALS. The cut-off in the superorbital period detection, however, is



**Figure 6.9:** Her X-1 Hardness-Intensity relation for the normal state (Top) and the ALS (Bottom), using 1-day data bins to improve S/N. There is no significant difference between the low-level emission in the two states.

abrupt, taking place over three cycles at most. By the end of the ASM dataset, the system shows behaviour similar to the beginning of the dataset, and the secondary feature is recovered.

This phase shift and profile change is most naturally interpreted as a change in the line-of-sight disk profile. There are two good candidates for the nature of this change. It has been suggested (Wijers & Pringle 1999, van Kerkwijk et al. 1998) that if a disk is sufficiently warped, the inner disk may actually flip over and start to precess in the opposite direction to the outer disk. This would produce the extremely rapid spin-down observed during the ALS as the direction of flow of accreting matter would then oppose the spin of the neutron star. Following this interpretation, we would expect to be observing the neutron star through the accretion disk for at least some of the ALS, which would produce a harder spectrum for the ALS than for the normal state. However, the continuum X-ray emission is not spectrally harder for the ALS than for the emission prior to and after the ALS (figure 6.9). A further difficulty is that one would expect to see some manifestation of the change of disk orientation in the overall ASM lightcurve.

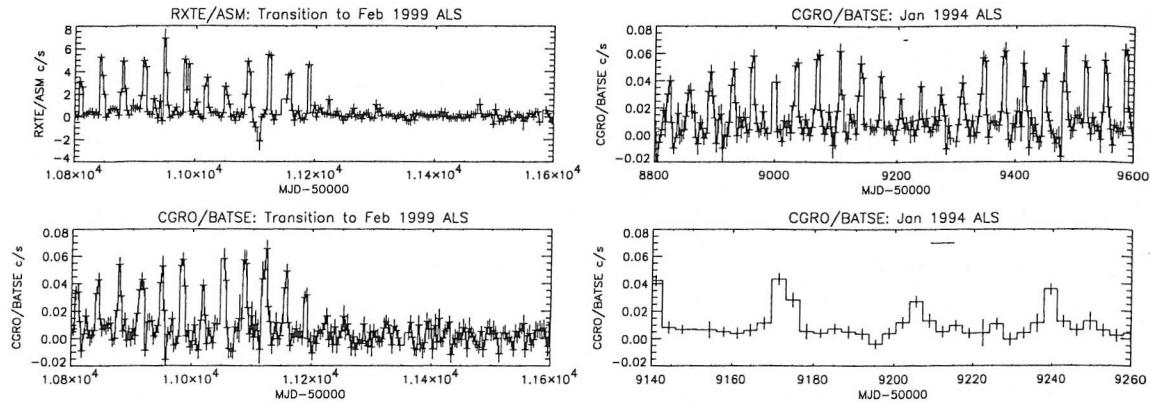
A second interpretation involves the disk changing its inclination with respect to the observer. We suppose here that the X-ray intensity from a twisted, warped disk varies as the amount of emitting area uncovered in the inner disk region varies. If such a disk were to change its inclination with respect to the observer, the uncovered area at each point in the cycle would also change with the

inclination, were the warp configuration to remain constant over the plane of the disk. A twisted disk would naturally produce a variation in the phase of features in the profile. That the ALS is accompanied by rapid spin-down of the neutron star presents a difficulty for this scenario, unless either (i) the accretion flow at the magnetosphere/inner disk boundary changes direction to act against the neutron star rotation (in effect replicating the flipped-disk scenario), or (ii) the spin of the neutron star causes ejection of matter through the propellor effect (c.f. Chapter 1); SMC X-1 is certainly in the approximate range of  $L_X$  for which such an effect might occur. The cessation of accretion at the inner disk edge would have two effects; firstly, as the donor will transfer matter to the disk regardless of the inner disk circumstances, the continuation of matter transfer will cause the amount of matter in the disk to increase, and secondly the luminosity intercepted by the disk will be reduced. We suggest here that the cessation of mass removal in the inner disk in fact leads to a change in the angle of inclination of the accretion disk. A further consequence could be the variation with the warp of the area open to X-ray heating, which might bring about a reduction in mass throughput depending on the timescale for such heating (Dubus et al., 1999, see also section 5.4.3).

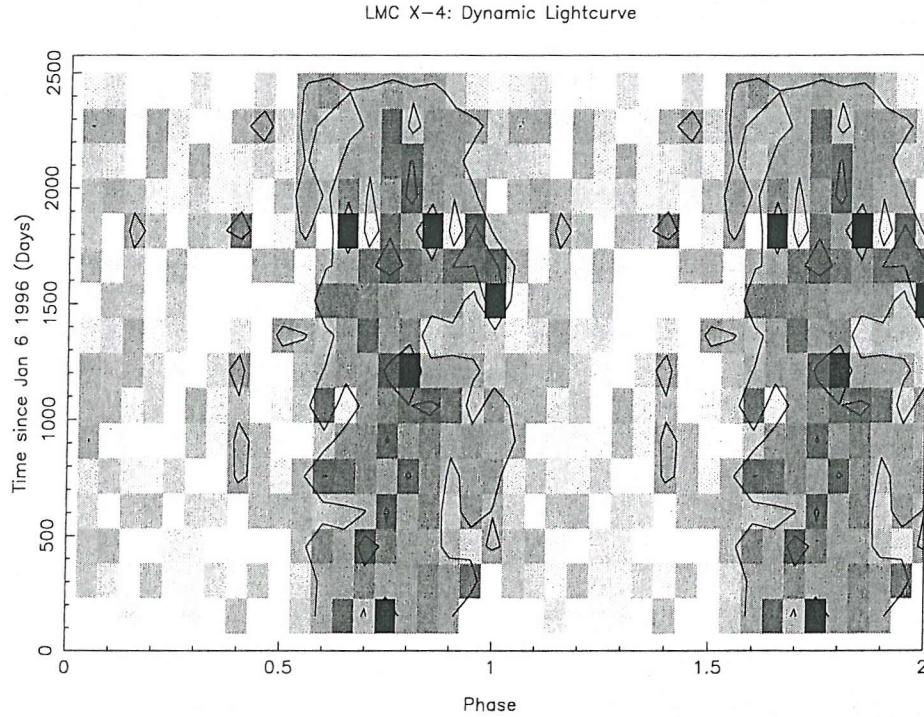
The 2001 ALS is clearly different from the 1994 ALS in its hard X-ray behaviour. In the case of the 2001 ALS, the RXTE/ASM turnoff is mirrored in the BATSE lightcurve, further strengthening the standard interpretation of this ALS as a reorientation of the accretion disk body along the line of sight. However, examination of the CGRO/BATSE lightcurve at the time of the 1994 ALS shows that the 35d cycle was manifestly present throughout, reaching  $\sim 45\%$  of the peak brightness attained at any high state. Furthermore, while the phasing of the 35d cycle varies on the run-up to the 1994 ALS, no such variation is apparent in the 1994 ALS. Thus, the 1994 ALS is clearly different in character to the 2001 ALS. The cleanest explanation for this difference occurs when the times of the Vrtilek campaign are examined on the BATSE lightcurve; it appears that the observations may have just missed the primary high state in the superorbital cycle. Whether the secondary high is absent is not clear from the lightcurve.

### 6.5.2 LMC X-4: Stable variation

As the most distant source in this study by a factor  $\sim 5$ , LMC X-4 exhibits comparatively low count rate and hence lower S/N, as can be seen in its 31-day dynamic lightcurve (figure 6.11), which furthermore undergoes no significant shift in phase or overall intensity over the ASM dataset. The



**Figure 6.10:** Comparison of 1994 and 2001 Anomalous Low States. **Left:** RXTE/ASM (top) and CGRO/BATSE (bottom) lightcurves of the transition into the 2001 ALS, each binned at intervals of 3.4 days (i.e. two binary orbits). **Right:** CGRO/BATSE lightcurve at the time of the 1994 ALS as determined by Vrtilek et al (1994), at two timescales (top and bottom). Horizontal line represents the interval of the Vrtilek et al observations, at MJD 49209-49214.5.



**Figure 6.11:** Dynamic lightcurve of LMC X-4. Data window length chosen to include 5 cycles per window - each bin in this plot is then the average of typically 30 datapoints.

superorbital lightcurve itself is somewhat different in shape to that of Her X-1 (figures 6.3 and 6.5), with the primary peak in the profile less symmetric, and no significant secondary peak is detected. A secondary peak might be present, however: assuming a similar relative amplitude as that in Her X-1, such a peak would occur at only  $\sim 0.2 \text{ cs}^{-1}$ .

This suggests that if the accretion disk in this source does shift its inclination, as suggested for

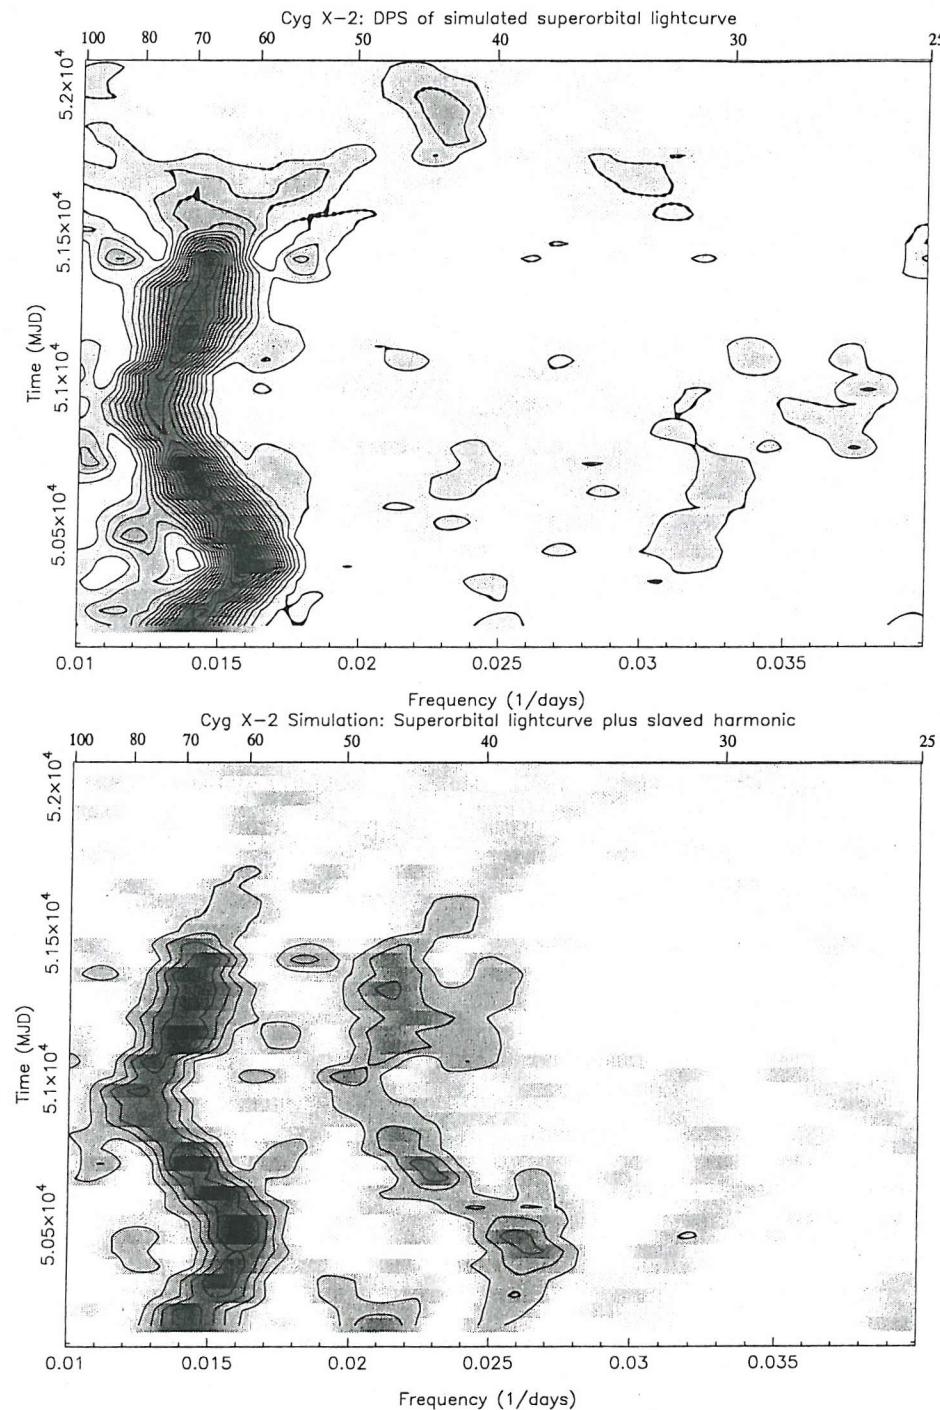
Her X-1, the LMC X-4 accretion disk either does not do so at all on the timescale detectable with the DL, or the disk profile and inclination are such that any existing wobble does not impact the superorbital fold to the same extent. We can place a rough upper limit on superorbital phase variation of features in the dynamic lightcurve of  $\sim \pm 0.1$  orbit.

### 6.5.3 Cyg X-2: Multimode Variation?

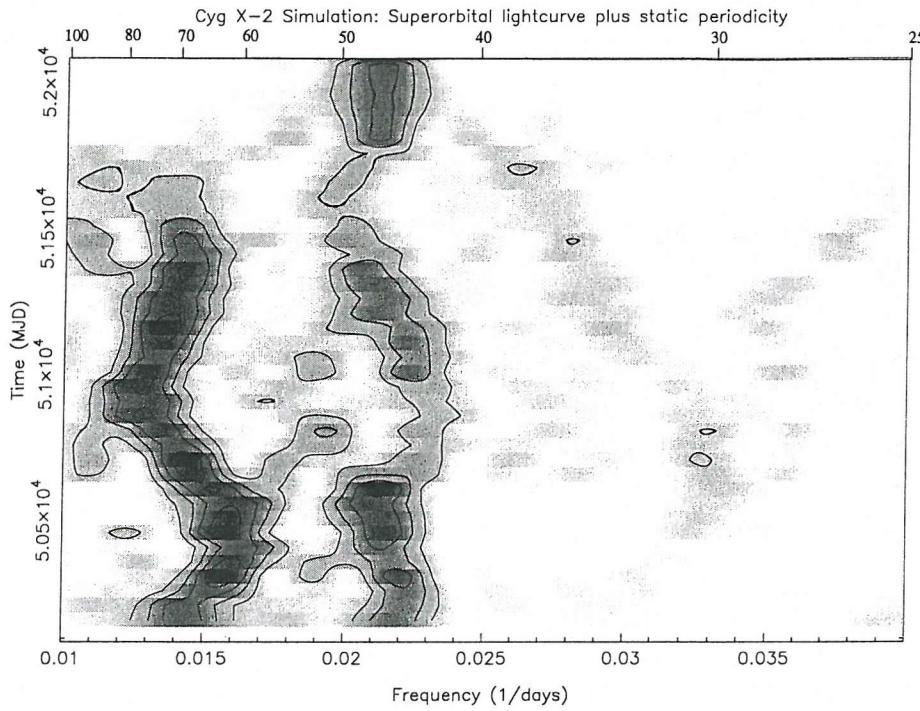
For Cyg X-2, in which the DPS behaviour does not show a single stable periodicity, such analysis is not possible. Indeed, its behaviour is highly complex, with significant periodicities forming and shifting over the 7 years of the ASM dataset. This explains the range of modulations reported for this source over the years (e.g: Smale & Lochner 1992, Paul, Kitamoto & Makino 2000). Of particular interest is the apparent convergence, reinforcement and subsequent divergence of periodicities near April 1996 - January 1999 (figure 6.6). Because the accretion disk is a fluid entity and the DPS overlap between data windows of the same order as the warp precession periods, we do not expect too discontinuous a change in the variation shape: in other words we expect the pattern (primary peak+single harmonic) to persist over a timescale of several data windows. At the point at which the strong DPS peak appears to split into two modes, we interpret the mode evolving towards lower frequencies as the dominant frequency, as it is more significant. It is thus easy to interpret this feature as a single periodicity with ill-defined phasing (hence the broad peaks in the DPS), with a period that changes in a manner superficially similar to SMC X-1, only on a much shorter period:  $\sim 2$  years compared to the  $\sim 7$  years of SMC X-1.

However, the behaviour may be more complex. Closer examination of the secondary peak at  $\sim 40$  days suggests it is not a simple harmonic of the primary period, as its period is some 20% longer than the expected location of the harmonic (see also Paul, Kitamoto & Makino 2000). Furthermore, this structure is still present over the last 1500 days of the dataset, even though the fundamental it would be a harmonic of has all but disappeared. This feature may represent either (i) a slightly shifted harmonic that represents a deviation from a purely sinusoidal variation profile, or (ii) a mode of variation independent of the  $\sim 59$ -day variation.

We can test which scenario describes Cyg X-2 better through simulations of varying periodicities. We take the superorbital lightcurve over day number 600-900 (from Jan 6 1996), then scale the period over time in a manner approximating our identified dominant periodicity. The DPS of this simulation is quite different in character to the real DPS (figure 6.12); the most obvious



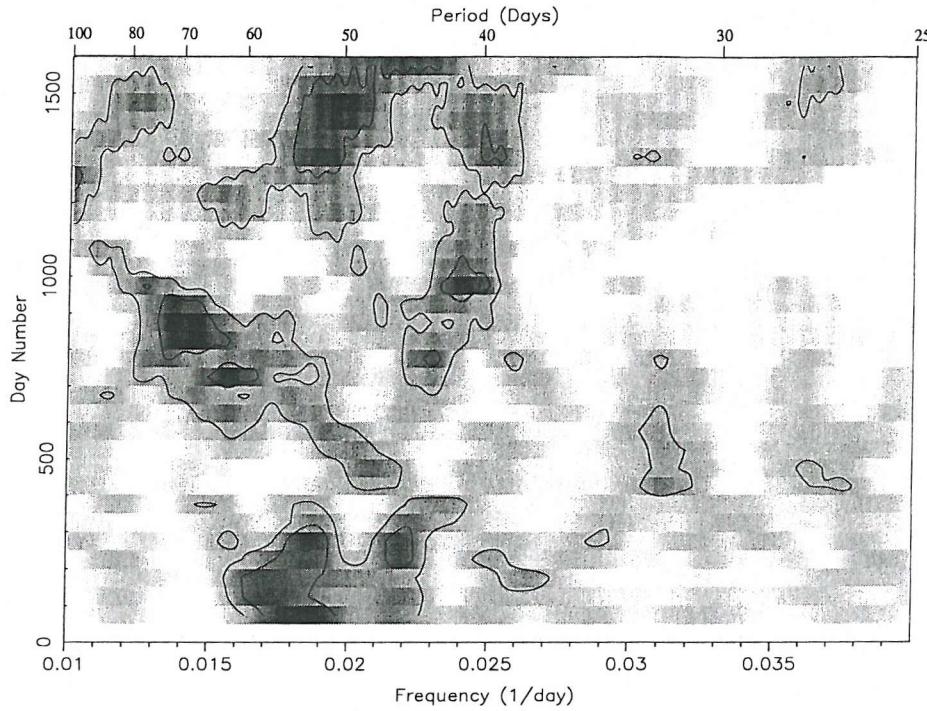
**Figure 6.12:** DPS of simulated Cyg X-2 lightcurves. **Top:** input data consists of the superorbital lightcurve from day number 600-900, which undergoes period change in a manner approximating the fundamental from figure 6.6. **Bottom:** As for top, but with 40-day periodicity added that is slaved to the fundamental, to approximate a non-sinusoidal shape of the variation.



**Figure 6.13:** Input lightcurve as for figure 6.12 (top), but with 40-day periodicity kept steady as the fundamental varies (i.e. two independent modes of variation).

differences are its far lower significance and frequency than the  $\sim 40$ -day periodicity actually seen. A 40-day periodicity was then artificially added but slaved to the fundamental to simulate a non-sinusoidal variation. The DPS structure that results varies more closely with the dominant periodicity than that actually seen. If the  $\sim 40$  day periodicity is made completely independent of the fundamental, however, by being held steady as the fundamental varies), the DPS that results is similar in morphology to that actually seen (figure 6.13).

We thus suggest here that the apparently variable periodicity at  $\sim 59$  days and the more stable periodicity at  $\sim 40$  days represent independent variations, rather than a single variation of non-sinusoidal shape. This suggestion is in agreement with the conclusion of an earlier analysis based on the static periodogram of the early years of the RXTE/ASM (Paul et al., 2000). This suggestion gains some support when we notice that the 40-day periodicity is close to the expected harmonic of the  $\sim 77$ -day periodicity reported in the archival Vela 5B and Ariel-5 datasets. To attempt to constrain the properties of this behaviour over a longer baseline, the Ariel-5 and Vela-5B datasets were subjected to DPS analysis. The Vela-5B dataset was used in the filtered, binned form provided by NASA HEASARC. Enough points were present to use the static periodogram to confirm the 77.4-day period already noted elsewhere (e.g: Smale & Lochner 1992), but not enough to follow its time evolution in a statistically significant way. However the Ariel-5 dataset has higher S/N,



**Figure 6.14:** Cyg X-2: DPS of the Ariel-5 dataset, binned to one-day averages, with contours at LS powers of 9 (corresponding to 99.9% significance over white-noise).

allowing less coarse binning and thus providing enough datapoints to permit DPS analysis. As figure 6.14 shows, the periodicities are highly variable over time, producing the highly complex structures seen in static periodograms (Paul et al., 2000). The  $\sim 77$ -day periodicity is strongly detected over the middle  $\sim 400$  days of the dataset, while its harmonic at  $\sim 40$  days is detected throughout most of the ASM dataset. However, as with the RXTE/ASM dataset, there is a further periodicity which appears variable and transient. Thus the apparent combination of two or more separate variations does not appear to be a new behaviour. Furthermore, this complex behaviour has now been confirmed using two independent datasets separated by 20 years.

The X-ray manifestation of the precessing warp may be different from that evidenced from the canonical disk precessing systems SMC X-1, Her X-1 and LMC X-4. Cyg X-2 shows a decrease in spectral hardness with count rate (figure 6.7), which at first sight would be inconsistent with a precessing warp. However, the variation of uncovered emitting area invoked to explain the superorbital periods in Her X-1 and LMC X-4 requires the column density through which the bulk of the X-rays reach RXTE to remain the same as the warp precesses. If, however, the emitting area is permanently obstructed by the outer accretion disk, assumed not to be a significant emitter in the ASM bandpass, the variation can be brought about by changing the column density through which the majority of the X-rays travel. In this case, the warped disk acts to subtract flux from

the central source, rather than contributing significantly to any emission. This removal of flux will be seen as X-ray obscuration, which thereby causes the observed spectrum to harden.

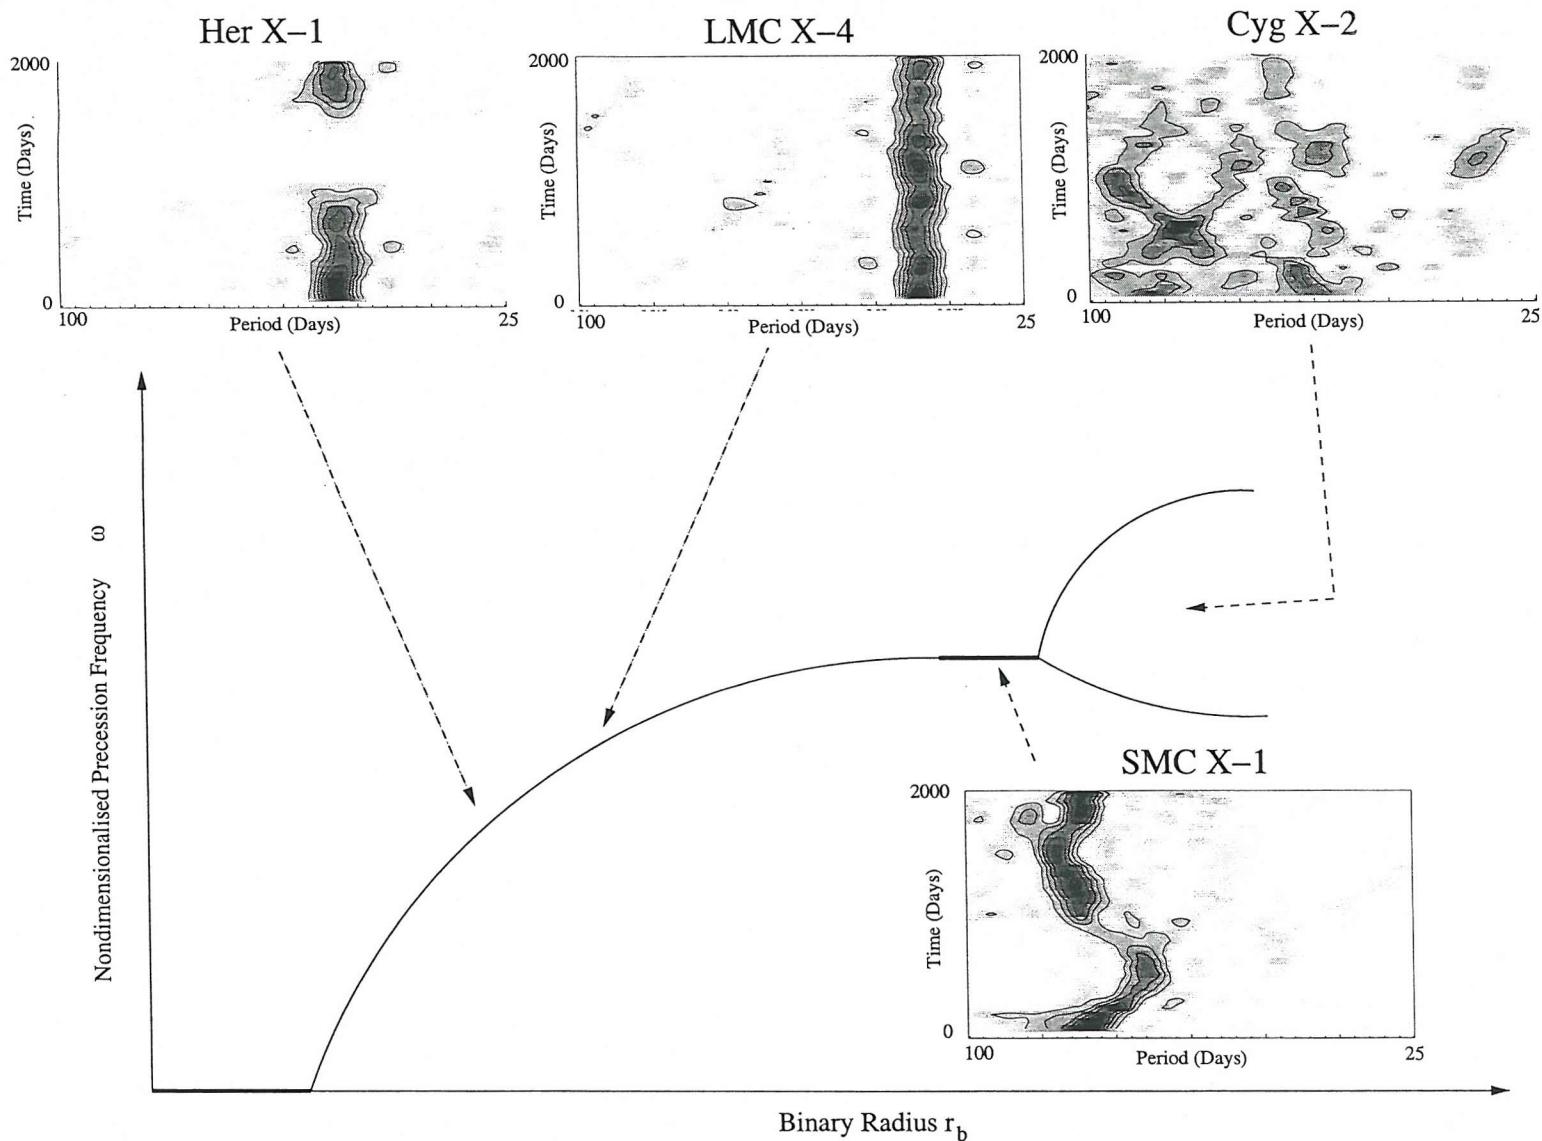
#### 6.5.4 Stability of Accretion Disks to Warping

We now use our results for the other systems showing superorbital modulations to test the stability framework of OD01. We first attempt to identify the warping modes corresponding to the superorbital variations identified in this work. With their stable behaviour in the DPS and simple superorbital lightcurves, Her X-1 and LMC X-4 show clear evidence for mode 0 warping. SMC X-1 shows a simple superorbital fold that varies over time, probably due to disk instabilities brought about as other modes begin to form (Chapter 5). Cyg X-2 shows at least two independent modes of warping, thought to represent mode 0,1 and possibly higher warp modes.

The sources thus behave in ways qualitatively predicted by the stability analysis of OD01 (table 6.2). Her X-1 and LMC X-4 are both predicted to show stable mode 0 warping, which is indeed observed. SMC X-1 is thought to show warping commensurate with its position near the mode 0 and mode 1 regions. Cyg X-2 is deep within the region predicting mode 1 and higher modes, which is consistent with the appearance and interaction of at least two warping modes.

We illustrate the evolution of disk warping with binary radius in a schematic bifurcation diagram, shown in figure 6.15 (modeled after figure 4 of OD01). This charts the evolution of precession solutions as the control parameter  $r_b$  is increased. According to the OD01 predictions, for low values of the control parameter  $r_b$ , no *radiation*-driven disk warping is predicted at present. In the region corresponding to stable mode 0 predicted warps, such warps are indeed found, with the precession frequency increasing with  $r_b$  (Her X-1 and LMC X-4 type variations). Near the boundary between regions, modes begin to compete, producing an instability in the superorbital period (SMC X-1 type variations). Finally, once the boundary into full multi-mode variations is crossed, strong periodicities form and interact (Cyg X-2 type variations). When we overplot the location of these sources on the schematic, along with their DPS results, we find that the long-term behaviour of the sources agrees qualitatively with the predictions of table 6.1, as can be determined from the DPS of the sources.

This result establishes radiation-driven warping as the most likely mechanism at work in the high- $L_X$  systems Her X-1, LMC X-4, SMC X-1 and Cyg X-2, as only this mechanism predicts the



**Figure 6.15:** Schematic bifurcation diagram for radiation-driven warping. As the control parameter  $r_b$  is increased, the number of stable precession solutions increases. Initially there are none, then stable mode-0 precession rising in frequency as  $r_b$  increases (solid line). Near the mode region border, the solution becomes marginally unstable as short-lived mode-1 instabilities form (thick line). Finally in the mode 1+ region, two or more steady solutions are possible, and the system precesses at a combination of these warping modes. When the positions of the sources Her X-1, LMC X-4, SMC X-1 and Cyg X-2 are overlaid on this plot, we find that their behaviour is fully commensurate with their system parameters, as shown by their behaviour in the DPS.

stability behaviour illustrated here. Warping driven by a disk-wind is not yet a likely candidate explanation, as no correct expression for the torque exerted by the disk-wind has yet been advanced (see discussion in Wijers & Pringle). However it should be noted that an accretion disk wind is expected in the systems examined in this chapter, and indeed may be inevitable in those systems in which radiation-driven warping plays a significant part. The radiation reaction from emitting protons rather than photons is greater by a factor  $\sim (c/v_\phi)$  (see Pringle 1996), which suggests this effect may even dominate under certain circumstances. However, the characteristics of wind emission are still far from clear; it is uncertain whether such a wind would be emitted instantaneously or locally, or whether a wind would act with or against the radiation-driven framework. What *does* appear to be clear is that the wind-driving scenario likely predicts different stability behaviour than is seen in the observed systems. The other promising candidate is magnetic warping, in which interaction between sheet currents in the accretion disk and the magnetosphere of the neutron star or the donor produce a torque out of the plane which in turn drives warping behaviour (Lai, 1999; Murray et al., 2002). This mechanism fails to explain the stability results presented in this Chapter because the magnetic field of the neutron star is a free parameter with respect to  $r_b$  and  $q$  (indeed the neutron star magnetosphere in Cyg X-2 is likely several orders of magnitude smaller than for the other three systems studied here), as is the magnetic field of the donor as experienced by the accretion disk.

## 6.6 Conclusion

We have analysed the time variation of superorbital periodicities present in three persistent neutron star XRB: Her X-1, LMC X-4 and Cyg X-2. Her X-1 shows a shift in phasing of the superorbital period shortly before entering its third Anomalous Low State, which we interpret as evidence for a change in disk inclination. LMC X-4 shows stable mode 0 warping, with little change over the period of the RXTE/ASM dataset. Cyg X-2 shows a combination of at least two separate modes of variation, suggesting more complex warping shape than for the other sources. Furthermore this behaviour has persisted for at least  $\sim 1000$  binary orbits, as shown by new analysis of the Ariel-5 dataset.

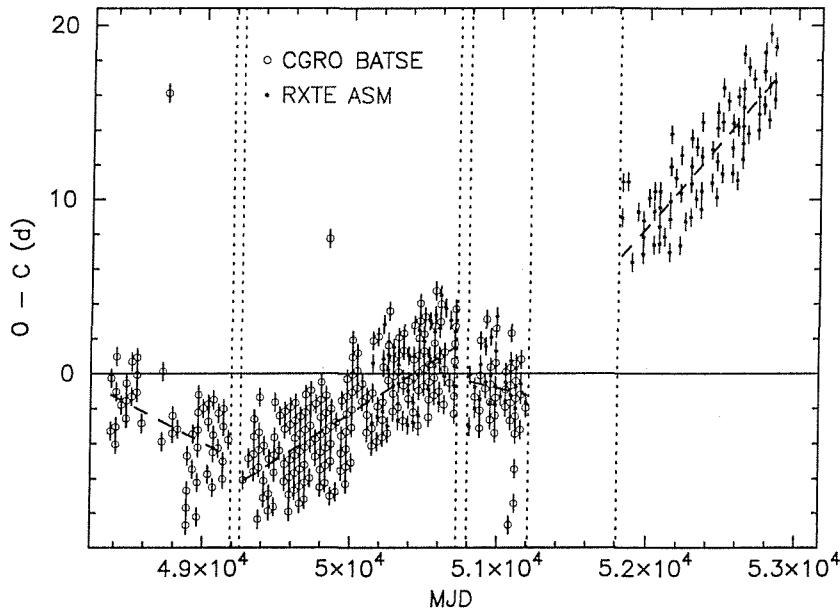
We have identified the most likely warping modes to produce the observed variations in Her X-1, LMC X-4 and Cyg X-2, and together with SMC X-1 have constructed an observational stability sequence based on their behaviour in the DPS. We find that this qualitatively agrees with the

predictions of the stability analysis of OD01 and establishes the DPS as a powerful tool to probe XRB properties.

## 6.7 Update: examining warpers in more detail

In the interval since the work reported here was accepted for publication, several developments have been made known to this author that promise to extend this work further. The most obvious way to proceed further in the identification of radiation-driven warping as the dominant warp mechanism, is to further populate the  $r_b$ - $q$  diagram. However, until the truly large (30m-class) ground-based telescopes such as CELT come online some time in the next 20 years, it will be difficult to establish the binary parameters of the as yet unconstrained superorbital period systems in our galaxy due to severe interstellar reddening (a significant fraction being in the highly obscured bulge). However, we note here that *all* the sources in this work have luminosities significantly above the sensitivity limit of modern examinations of external galaxies (Zezas et al., 2002). For example, of the  $\sim$ 200 XRB's discovered with Chandra in M31, well over half show variability on timescales of days or longer (Kong et al 2002). We may thus assume the number of observable disk-warpers to jump by a factor  $\gtrsim 3$  as more X-ray data is gathered on these sources. Remarkably, constraints on the binary parameters of XRB in external galaxies is in some cases already possible; for example in some HMXBs the donor and orbital period may be identified from *optical* eclipses (Pietsch et al., 2004). Another example, at least in the black-hole binaries, is the possibility of using the break frequency in the power spectrum to estimate the mass of the compact object (Barnard et al., 2003b)

Also well within the capabilities of modern X-ray instruments is the examination of the nature of the superorbital variations in more detail. For example, ASCA investigations of LMC X-4 showed varying column depth from the 6.64 Fe-line complex, but little continuum variation, implying the warp primarily occurs in the inner accretion disk (Naik & Paul 2003). Further investigation with the higher-sensitivity instruments on XMM should offer much more detailed insight into the nature of the warp, including phase-dependent absorption and emission features. The interpretation advanced for the ALS in Her X-1, in which the accretion disk changes its tilt, is not of course the only possibility; another interpretation is simply that the height of the warp might be changing (one might imagine that if the propellor effect inhibits central accretion, the resulting matter buildup in the accretion disk would lead to an increase in the warp extent; in a sense this is simply a different response of the disk to the same phenomenon). The extent of the disk warp may be



**Figure 6.16:** (O-C) diagram of the main-on in the superorbital cycle of Her X-1, assuming constant ephemeris, from BATSE & ASM data (Still & Boyd 2004 in prep). The ALS is characterised by apparent step-changes in the superorbital periodicity, which are small enough that the DPS approach was not sensitive to them.

varying from cycle to cycle anyway, as evidenced by the cycle-to-cycle variation in peak brightness reached by Her X-1. This hypothesis is in principle testable, as one would expect the superorbital lightcurve features to broaden as the warp changes its vertical extent; however, the RXTE/ASM lightcurve by itself probably does not offer sufficient statistics to place limits on this behaviour. However, the USA satellite was also operating during the egress from the 1997 ALS, which should enable enough points to be added to the superorbital lightcurve to constrain its evolution out of the ALS.

Finally, the nature of the ALS in Her X-1 now appears to be somewhat different to that suggested in this chapter. As discussed in Chapter 3, when a source changes its periodicity abruptly, the DPS is insensitive to this change due to the smoothing, provided the instantaneous period change is small. Recently Still & Boyd (2004) have traced the evolution of the superorbital cycle to high time resolution by plotting the time of occurrence of the superorbital high states and plotting their phase residuals from the best-fit linear ephemeris onto a superorbital (O-C) diagram.

When this is done (figure 6.16) it appears that the superorbital cycle in Her X-1 has a different period before and after each ALS, with the period change ( $\sim 0.2$ d out of 35d, i.e.  $\sim 0.5\%$ ) occurring entirely during the ALS. Due to the smoothing from the sliding windows of the DPS, the DPS was not sensitive to this period change, which appears instead as a periodicity of diminish-

---

ing significance and broader peak in the power spectrum. Thus it now appears that what was interpreted in our work as a variation in *phasing* at constant period, was in fact a variation in period after a short ALS. This casts our interpretation of the ALS into doubt, and raises again the possible interpretation of the ALS as an increase in the vertical extent of the warp rather than a disk-tilt. Investigation of the relative phasings of the main- and secondary-high states in the superorbital cycle should help to resolve this issue. Her X-1 appears to be entering a new ALS at present (early 2004); future scheduled XMM observations (Still, private communication) should allow the spectral evolution both during and on the egress from the ALS to be followed, yielding more insight into the nature of the ALS. In particular, the high effective area of the EPIC-PN instrument should allow the disk geometry to be probed during the ALS itself, thus establishing if the ALS represents a different accretion state or further disk precession.

## Chapter 7

# Near-IR Spectroscopy of Circinus X-1

In this short chapter I introduce the properties of the eccentric (in many senses of the word) X-ray binary Circinus X-1, and present results of the highest-resolution IR spectra yet taken of this object in order to directly constrain the donor. Part of this chapter is based on previously published work (Clark, Charles, Clarkson & Coe 2003).

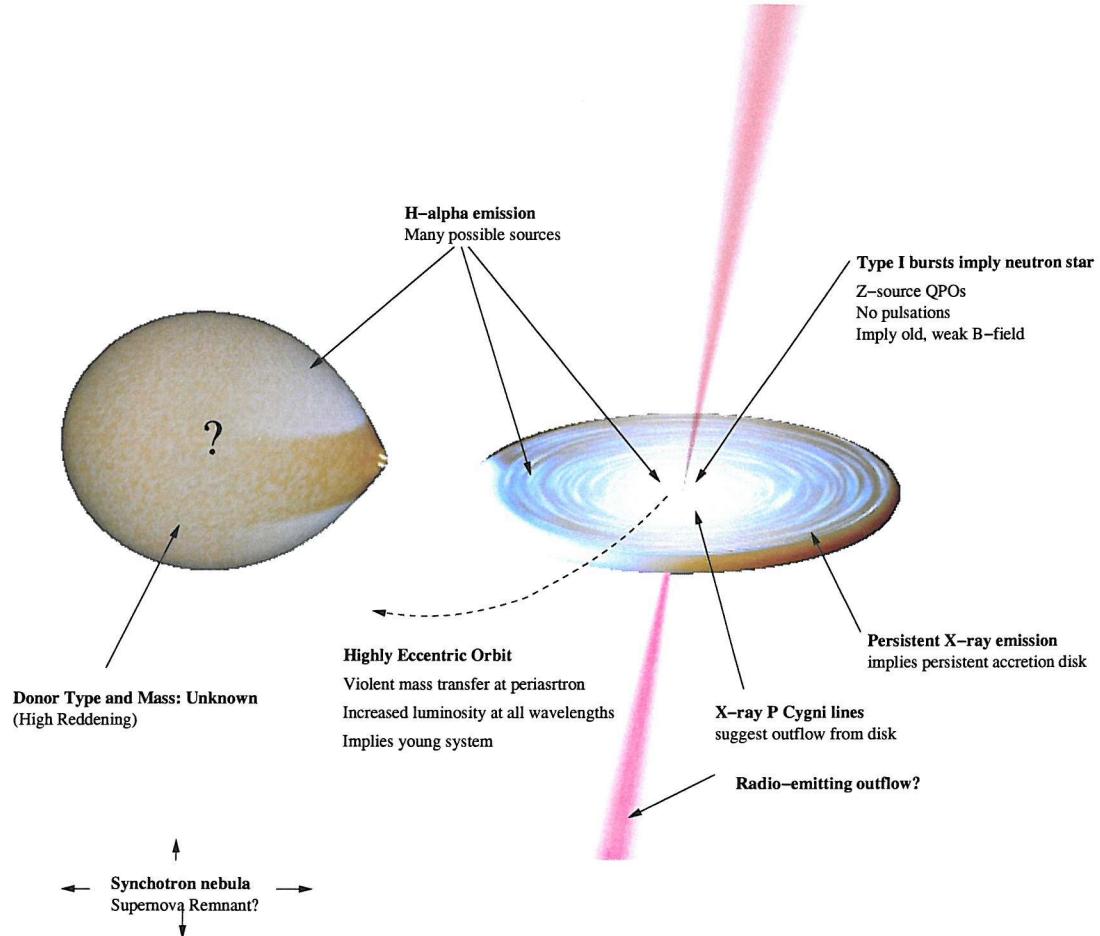
### 7.1 Circinus X-1

Cir X-1 is a galactic X-ray binary located in the galactic plane at a distance of 6-8 kpc (Goss & Mebold, 1977). Its most well-known feature is strong X-ray variability on a (presumed orbital) cycle of 16.6 days (Kaluzienski et al., 1976). The source distance and X-ray lightcurve imply variation of X-ray luminosity  $L_x$  of a factor  $\sim 100$  throughout this cycle (e.g. Kaluzienski et al. 1976). Radio flares were found at peak levels of 1Jy, on a similar period to the X-ray variations (Haynes et al., 1978), and a similar cycle is evident in the IR activity (Glass, 1994). An arcsec-scale asymmetric jet has been imaged (Fender et al., 1998), raising the possibility that outflow from the system is relativistic. During an interval of low X-ray activity, EXOSAT detected eleven X-ray bursts, and their identification with Cir X-1 strongly suggests that the compact object is a neutron star (Tennant et al., 1986). Quasi-periodic oscillations (QPO's) have been detected and characterised

on two occasions (Tennant, 1987; Shirey et al., 1998). The earlier EXOSAT observations suggested Atoll-like behaviour on the X-ray colour-colour diagram, whereas RXTE/PCA observations in the late 1990's show Z-source behaviour (Shirey et al., 1999a), implying an external magnetic field strength  $\gtrsim 10^9 G$  (van der Klis 1995). X-ray maxima are often preceded by dips in the X-ray lightcurve. Spectral fits with ASCA, RXTE and BeppoSAX during these dips are consistent with a partial covering model, in which the X-ray emission consists of a bright component undergoing varying absorption, plus a fainter component not attenuated by absorbing matter (Brandt et al., 1996; Shirey et al., 1999b; Iaria et al., 2001).

Very little is known about the mass-losing star in this system, due to a high but uncertain amount of extinction to the source ( $5 < A_v < 11$  mag; Glass 1994). Early classification of this system as an HMXB was based on an optical identification (Whelan et al. 1977) which was later disproven as the single object originally identified was subsequently resolved into three distinct stars (Moneti, 1992). The best optical spectra to date are likely dominated by the accretion disk, with the donor an insignificant component (Johnston et al., 1999).

Indirect methods of determining the mass of the donor are also indeterminate. Modeling the donor mass and eccentricity from the optical spectra (e.g. Johnston, Fender & Wu 1999) assumes identification of spectral components with the donor, which is clearly uncertain. The apparent radio association of Cir X-1 with the supernova remnant (SNR) G 321.9-0-3 suggested an extremely high transverse velocity of  $\sim 450 \text{ km s}^{-1}$  due to an asymmetric supernova kick (Stewart et al., 1993). Subsequent exploration of binary parameter space with this kick velocity as a constraint suggested both an extremely high eccentricity ( $0.90 < e < 0.94$ ) and low donor mass ( $M_2$  likely below  $1 M_\odot$ ; Tauris et al. 1999). However, a recent HST study (Mignani et al., 2002) finds no measurable proper motion, implying an upper limit on the transverse space velocity of  $\sim 200 \text{ km s}^{-1}$ . This rules out the association with SNR G 321.9-0-3, thereby leaving  $M_2$  and  $e$  undetermined. The Chandra X-ray Observatory (CXO) detected a remarkable set of strong X-ray P Cygni profiles through periastron passage, corresponding to highly ionised states of Ne, Mg, Si, S and Fe (Brandt & Schulz, 2000). The breadth of the lines ( $\sim 200 - 1900 \text{ km s}^{-1}$ ) and their time variability suggest identification of these lines with highly ionised outflow from the inner region of the accretion disk (Schulz & Brandt, 2002). Taken with the radio observations of an expanding jet from the source (Stewart et al 1993, Fender 2003 priv. comm), the emerging picture of Cir X-1 is of an extremely unusual neutron star analogue to the galactic black hole candidate microquasars such as SS433 (see Mirabel & Rodríguez 1999 for a review of stellar jet sources, and figure 7.1). To attempt to



**Figure 7.1:** Schematic of the mysterious XRB Circinus X-1. Observational properties are either ambiguous or contradict each other. See text for details. Figure prepared using visualisation software by Rob Hynes.

make a direct detection of the donor in this system, Circinus X-1 was observed in the JHK band with the IRIS2 spectrograph mounted on the Anglo Australian Telescope (AAT).

## 7.2 IR observations of X-ray Binaries

The mounting of IR and near-IR spectrographs on 4m- class and larger telescopes is opening up new regions of the galaxy in which to probe the properties of XRB, due to the lower susceptibility interstellar absorption. Cir X-1 is an ideal object for such investigation, as it exhibits visual magnitude  $M_V \sim 21.5$  (Mignani 2002), but in the IR has appeared as bright as  $K \sim 7$  (Glass 1994)! With this in mind, we review briefly here the recent IR and near-IR observations of XRB.

In the case of massive donors, with  $T_{surf} \sim 10^4$ K, the parameters of the outer atmosphere/circumstellar

environment can be probed by identification of emission line series and the line strengths of the constituents (Oudmaijer et al., 1995; Clark et al., 1999). Just as optical spectra from various components of XRB are often dominated by the Balmer series, in the near-IR the Brackett series dominates. With  $T \sim 10^4$ K, an equatorial disk about a Be- donor gives rise to significant levels of free-free and bound-free IR emission, which results in excess emission in bands redwards of  $V$ . This (J-K) colour excess is often used in confirming the identity of a Be star in the absence of optical information, and may also be used to estimate the size of the equatorial disk (Charles & Coe 2003). The IR spectral properties of *OB*-stars have been charted in a systematic study of 180 such stars that are also well-studied optically, with the result that the  $2-2.2\mu\text{m}$  behaviour of such stars allow discrimination between spectral sub-classes (Hanson et al., 1996).

Turning to lower mass systems, several class of donor may be identified from their IR emission features. K-band spectra of isolated (G-M) stars exhibit neutral metal lines and the  $^{12}\text{CO}$  bandhead in absorption (e.g. Mould, 1978; Kleinmann & Hall, 1986). Furthermore the luminosity class within the spectral type can be inferred from the observation that the ratio of line strengths [ $^{12}\text{CO}/^{13}\text{CO}$ ] varies from  $\sim 10$  in giants up to  $\sim 90$  in dwarfs (Campbell, Maillert & Lambard 1990 and references therein). These features are also seen in LMXBs, notably GX13+1 (Bandyopadhyay et al 1999). A product of the accretion-outflow process is strong  $Br(\gamma)$  and He-I emission lines, which are often broad and may have P Cygni profiles. This, coupled with the line velocity and breadth, suggests in several systems (e.g. Sco X-1, GX13+1; Bandyopadhyay et al 1999 & Bandyopadhyay et al. 1997) that the lines originate in an accretion disk wind. By contrast, narrow line profiles suggest instead origin in the outer accretion disk and/or heated face of the secondary (e.g. Sco X-2; Bandyopadhyay et al 1997). This raises the possibility of carrying out radial velocity searches (Steeghs & Casares, 2002) in the IR; as we show in this chapter, Cir X-1 may be the ideal candidate for such an investigation.

### 7.3 Observations

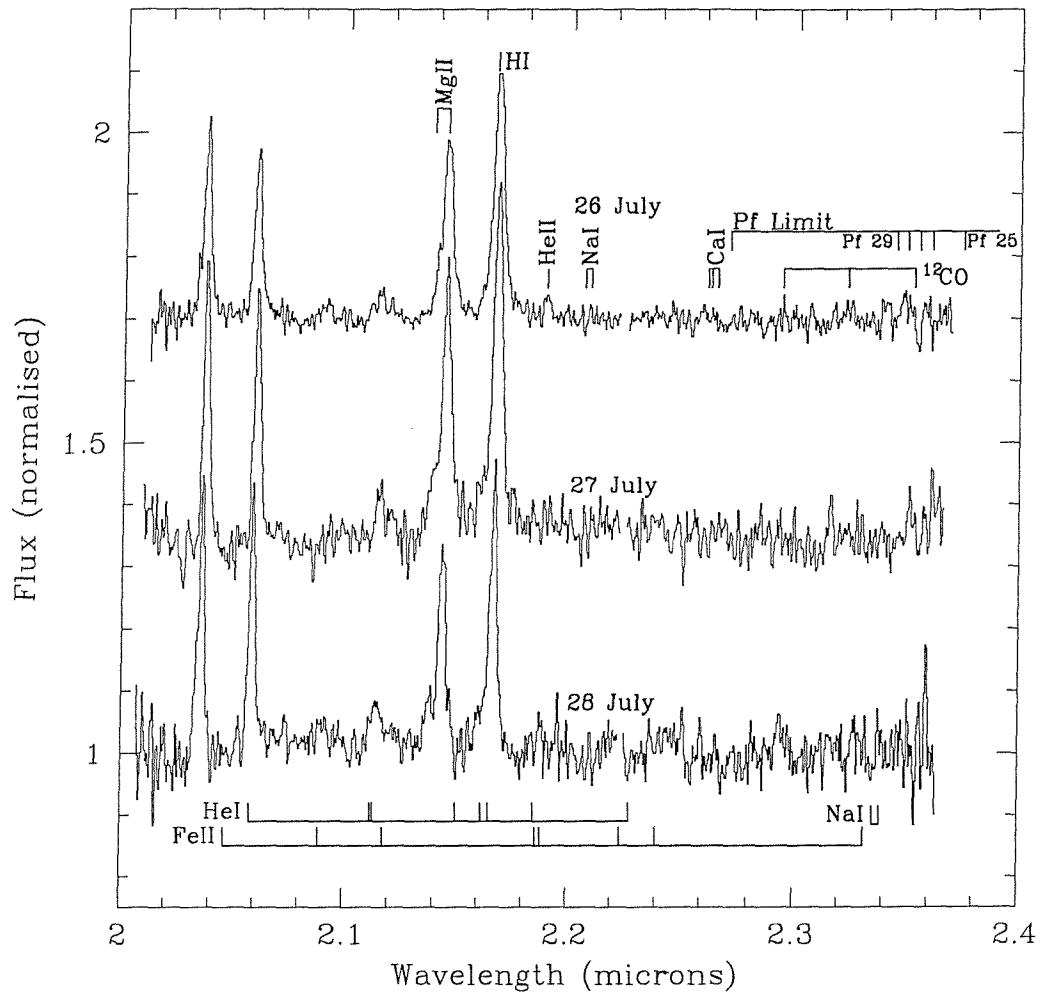
Commissioned between October 2001 and July 2002, the IRIS2 imager and spectrograph mounted on the 3.92m AAT provides long-slit near-IR spectra with typical spectral resolution  $R \sim 2400$  (Smith & Churilov, 2003). At its heart is a 1024x1024 Rockwell HAWAII-1 HgCdTe infrared detector, illuminated by a f/8 to f/2.2 focal reducer. Spectroscopy is achieved by the implementation of a slit wheel allowing the insertion of a 1" wide slit (or multi-slit masks), and a wheel

in the collimated beam containing grisms. These grisms use 50mm 45-degree prisms of sapphire to deviate the beam and feed replica transmission gratings applied to their surfaces. The spatial scale of IRIS2 is 0.446" per pixel, giving a field of view of 7.7' x 7.7'. The 1" slit for spectroscopy corresponds to limiting spectroscopic resolution of 2.2 pixels on the detector, yielding the  $\sim 2400$  figure quoted above. This value turns out to be within  $\sim 10\%$  of the spectral resolution after reduction.

A set of three K-band spectra were taken near periastron in the 16.6-day cycle, during shared-risk service runs as part of the commissioning process, on the nights of 26-28 July 2002 (see figure 7.2). On each night, six exposures were taken at 5 minutes each, yielding a spectrum with  $S/N=40$ . The high sky brightness in the near-IR makes standard flat-fielding, as performed in the optical, inappropriate as the resulting systematic errors are significant. Instead, the target is nodded along the spectrograph slit to perform direct sky subtraction. This process is susceptible to error due to variations in e.g. OH emission over time; however these are subsequently removed by subtracting along the slit itself. Arc calibrations and observations of telluric standards O8V (night 1) and A0V (nights 2 & 3) were also made.

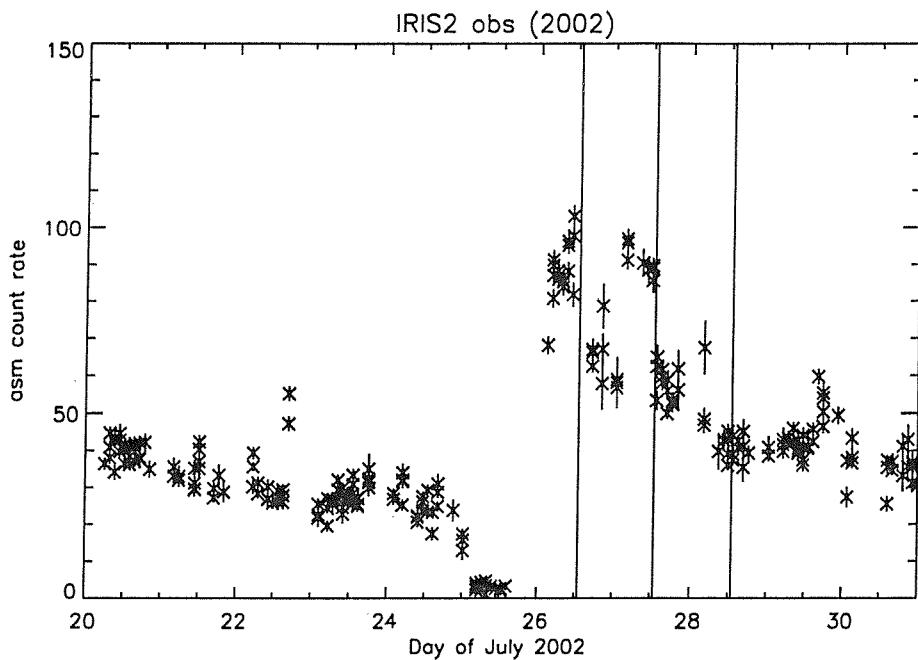
There are clear detections of strong  $Br(\gamma)$  and He-I lines in emission (equivalent widths  $\sim 30$  and  $\sim 15\text{\AA}$  respectively), as well as lines of  $He-II$  and  $Fe-II$ . However, what immediately grabs attention are the two lines at 2.0364 and 2.144  $\mu\text{m}$ ; these lines are of comparable significance to the  $Br(\gamma)$  and  $He-I$  lines. Identification based on known transitions proved impossible; there is a  $Mg-II$  line at 2.144  $\mu\text{m}$ , however for optically thin emission one expects the  $Mg-II:2.138:2.144$  ratio to be  $\sim 2$ , which is clearly not observed. The alternative explanation for these lines was their origin in high-velocity outflow, both with apparent velocity separation  $\sim 3100\text{km s}^{-1}$ . The absence of outflow at longer wavelengths could then be explained by the transverse doppler effect for highly relativistic material, i.e. matter moving away from the observer is redshifted out of the K-band, while a combination of jet orientation and doppler shift allows retention of the blue jet (Kaiser et al., 2000). These lines would thus represent the first direct detection of highly relativistic outflow from *any* neutron star, a highly significant discovery.

Unfortunately, the “jet” lines are in fact artefacts of the IRIS2 optics. The grisms were created with the plane of birefringence  $90^\circ$  out of alignment, with the result that the optical path corresponding to the birefringent second image was sent to the detector, rather than harmlessly to one side as was made clear in the design (Boyle 2003, private communication). As a result, a double image of each spectral line was produced.



**Figure 7.2:** IRIS2 spectra of Cir X-1, taken on the nights 26-28 July 2002 in shared-risk service mode. The Br( $\gamma$ ) and He-I lines are clearly present, but so are extremely bright counterpart lines, displaced by  $\sim 3100 \text{ km s}^{-1}$ .

As a result of this error on the part of the contractors, we were awarded two nights of “payback” service time, which were used to obtain a set of JHK spectra near apastron in the 16.6-day cycle, on the nights 4-5 September 2003. Because these observations were performed some ten days before a PATT telescope proposal deadline, we were able to obtain preliminary reductions made by the observer, Stuart Ryder, who is warmly thanked for his help. Owing to time constraints on the part of all the collaborators in this work, we base our analysis on these reductions.



**Figure 7.3:** RXTE/ASM lightcurve of Cir X-1 during the K-band IRIS2 spectra taken in July 2002. The times of observation are denoted by the vertical solid lines on the plot.

## 7.4 Near-periastron spectra

The July 2002 K-band observations took place just after the source underwent a periastron passage, as evidenced by the RXTE/ASM lightcurve during the observations (figure 7.3). The system was thus in the usual state of high X-ray activity encountered after the quasi-regular X-ray “dips” (Chapter 8). During this state, the system is thought to undergo extensive (possibly quasi-spherical) mass transfer onto the compact object, with the result that outflow-driven phenomena are expected to dominate the IR spectrum. Excising the regions corresponding to the “ghost” lines from the result, the remaining line features were compared to model atmospheres to attempt to identify the spectra with known stellar features. The results are summarised here; the reader is referred to Clark et al (2003) for more details on the modeling and identification of the spectral features.

We plot the mean K-band spectrum of Cir X-1 in figure 7.4, along with the IR spectrum taken by Johnston, Fender & Wu (1999), the comparison Be X-ray binary XTE 1946+274 (Verrecchia et al., 2002) and the microquasar GRS1915+105 (Harlaftis et al 2001). The strongest lines, and their probable identifications, are presented in Table 7.1 (taken from Clark et al 2003). No evidence is seen for the Na-I doublet (2.20624/2.20897  $\mu\text{m}$ ) the Ca-I triplet (2.26141/2.26311/2.26573  $\mu\text{m}$ ) or the  $^{12}\text{CO}$  and  $^{13}\text{CO}$  molecular bandheads longwards of 2.29  $\mu\text{m}$ , which would suggest emission

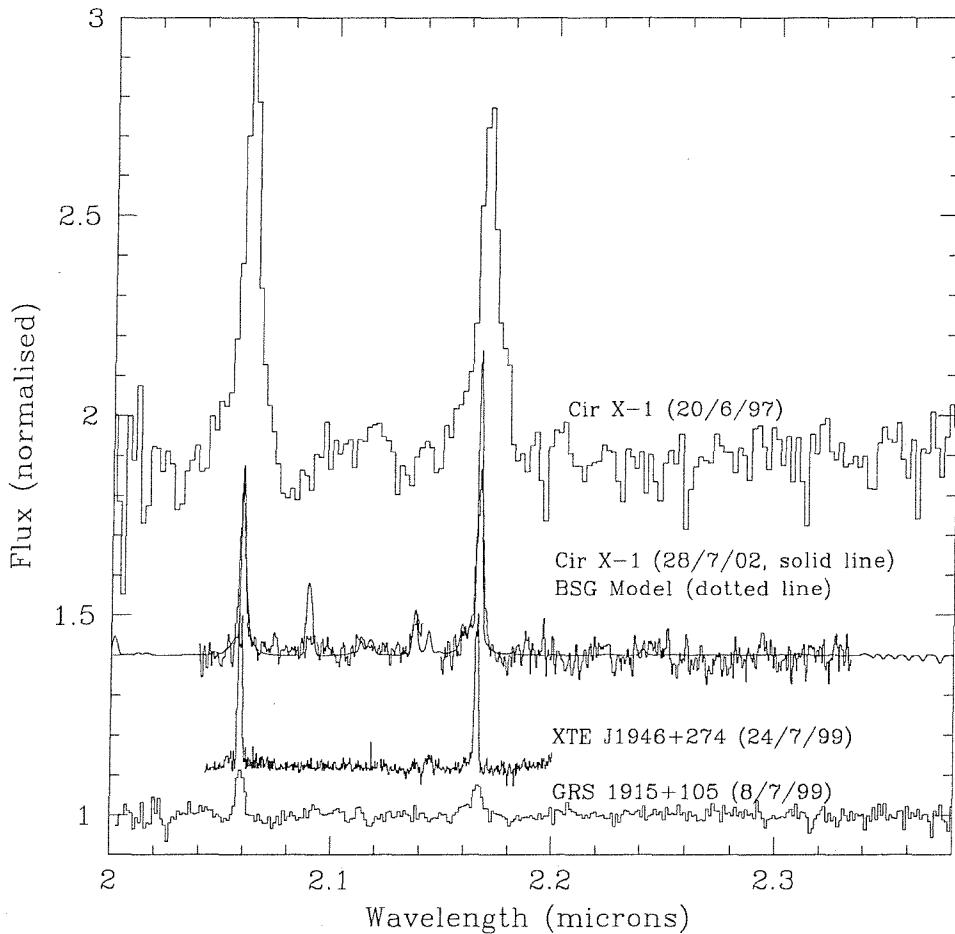
**Table 7.1:** Strongest emission lines from K-band spectra, 25-7 July 2002

Central Wavelength ( $\mu m$ )	Suggested Transition	EW(Å)		
		N1	N2	N3
2.0597	He-I ( $2s^1S$ $2p^1P^0$ )	2.058	11.8	15.3
2.0907	Fe-II( $z^4F_{3/2}$ $c^4F_{3/2}$ )	2.089	1.8	-
2.1665	$Br(\gamma)$	2.1661	24.2	32.2
2.1887	He-II(10-7)	2.189	1.3	-

from late (G-M) spectral types for the donor. Conversely, spectra of stars earlier than  $\sim O6$  show evidence for strong  $C - IV$  and  $N - III$  emission, which is absent from the spectra, while later O stars are excluded as a result of the  $Br(\gamma)$  and  $He - II$  emission (Hanson et al., 1996). The strength of the  $Br(\gamma)$  emission rule out classification as a Wolf-Rayet star (Figer et al., 1997). The remaining possibilities for standard massive stars are thus a Be star or B supergiant. The spectrum is indeed superficially similar to the Be X-ray binary XTE 1946+274; however, the primary emission lines are uniformly less than 20% of the strength observed in Cir X-1 (Clark et al 1999), so this classification is discarded. The final remaining possibility, i.e. a B-supergiant, is broadly consistent with synthetic spectra of e.g. a mid-B supergiant (c.f. Hillier & Miller, 1998). However, the significant and rapid time variability of the  $Br(\gamma)$  and  $He - II$  emission lines probably also rules out this classification. If the emission is primarily due to uncontaminated donor emission, then, the only possible identification would be some unusual F- or earlier type star, suggesting IMXB classification.

A crude constraint on the donor mass might be expected through the non-detection of any shift in central wavelength over the course of the three nights for which the K-band spectra were obtained. Examination of figure 7.2 suggests that any intrinsic central line shift must be less than  $\sim 0.005\mu m$ , corresponding to an upper limit on line shifts of  $\sim 700\text{km s}^{-1}$  over two days. This provides a useful upper limit on the donor mass only for unrealistically eccentric orbits (simulations suggest  $e \gtrsim 0.95$  is needed to produce this size velocity variation over such a short time period).

As suggested from the high X-ray output, the contamination from the accretion process is likely to be highly significant. In the case of the microquasar GRS1915+105, photospheric features of the (K-M) donor were unidentifiable (Greiner et al., 2001), and indeed the emission features are also entirely consistent with disk or disk-wind origin (c.f. Bandyopadhyay et al., 1999). To unearth the nature of the donor by direct detection thus requires comparison with observations at apastron, when the contamination from the accretion/outflow process is minimal.

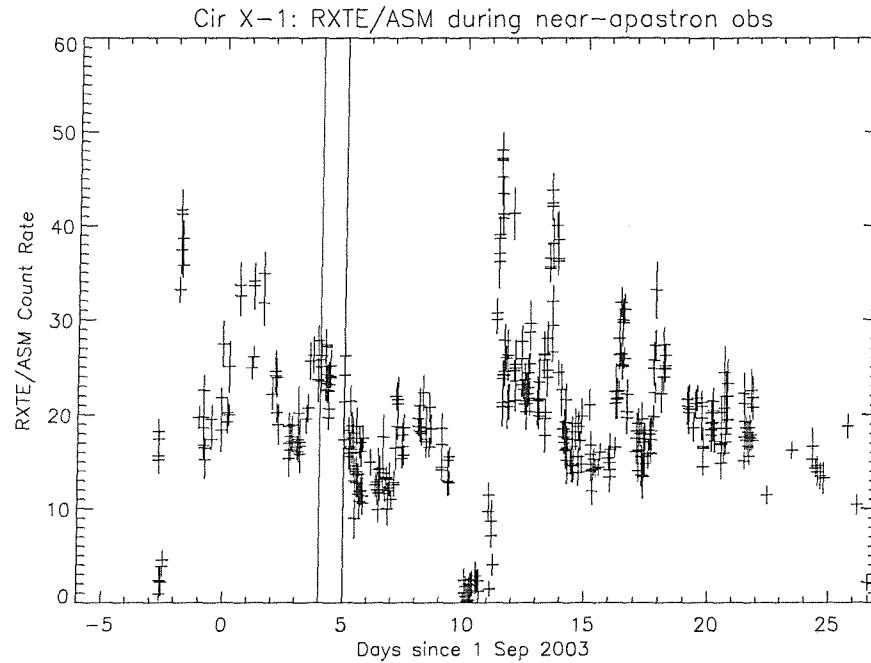


**Figure 7.4:** Mean K-band spectrum of Cir X-1, along with the IR spectrum taken by Johnston, Fender & Wu (1999), the comparison Be X-ray binary XTE 1946+274 (Verrecchia et al 2002) and the microquasar GRS1915+105 (Harlaftis et al., 2001)

## 7.5 Spectra Near Apastron

In contrast to the three nights of K-band spectra, the observations near apastron (figure 7.5) consist of a single set of JHK spectra, spread over two nights (figures 7.6, 7.7 & 7.8). The RXTE/ASM lightcurve shows that the observations occurred during the (somewhat erratic) decline in X-ray output following the main X-ray maximum. Furthermore, the source has continued its longterm decline (Chapter 8), with an average X-ray brightness a factor  $\sim 2$  lower in September 2003 than for July 2002.

The contrast between the apastron and post-periastron spectra is remarkable. Both show strong Br( $\gamma$ ) emission, but at apastron it is double-peaked. Strong He-I  $1.08\mu\text{m}$  emission is also double-peaked, as are Br(10) - Br(14), which suggests doppler splitting due to (i) an accretion disk and/or

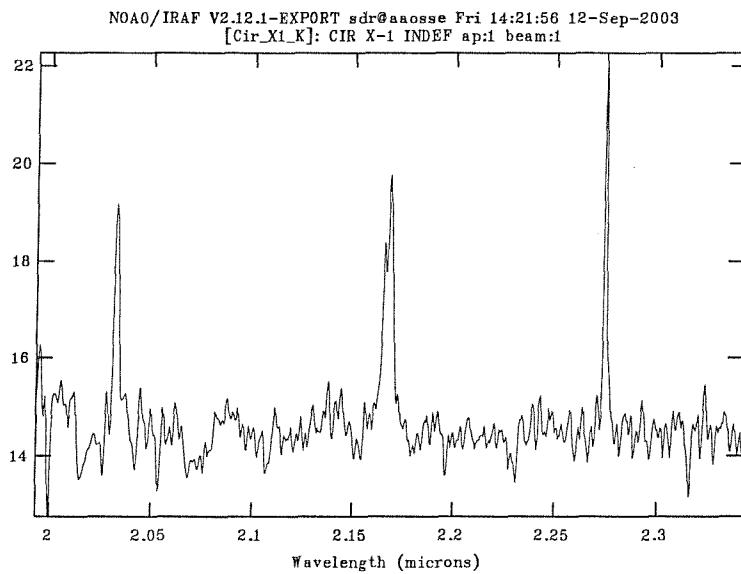


**Figure 7.5:** RXTE/ASM lightcurve at the time of the JHK near-apastron spectra. Observation intervals are denoted by the vertical solid lines.

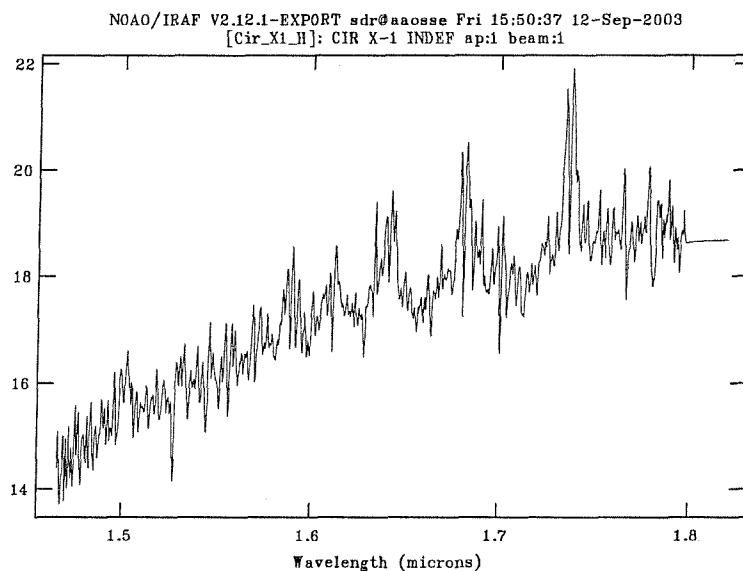
bipolar outflow at the accretor, or (ii) rotation in a Be-X circumstellar disk. From initial reductions we estimate the velocity separation of these lines to be  $\sim 600 \text{ km s}^{-1}$ , which when identified with the keplerian velocity of rotating matter is consistent with rotation rates observed in Be-disks (Negueruela et al., 1998). Alternatively, assuming the majority of this emission occurs in the outer  $\sim 10\%$  of the accretion disk (Shafter et al., 1986), the outer radius derived is then  $\sim 2 \times 10^{11} \text{ cm}$ , or roughly 10% the Roche Lobe radius of the accretor, which is fully consistent with the transient disk expected in such a violent eccentric-orbit binary Johnston et al. (1999). Whichever explanation is posited, the existence of such a feature appears to depend on binary phase.

A dominant feature of the post-periastron spectrum is the strong He-I line at  $2.058 \mu\text{m}$ . Objects undergoing high mass loss rates usually show this feature in conjunction with the  $1.08 \mu\text{m}$  line. Its apparent absence in the apastron spectrum, for which the  $1.08 \mu\text{m}$  line is extremely strong, is therefore puzzling. Furthermore a strong emission line is present at apastron at  $2.03 \mu\text{m}$ , without an obvious double-peaked structure. This line is presently unidentified: the He-II(15-8) transition does occur at  $2.0379 \mu\text{m}$ , but it is not expected in isolation.

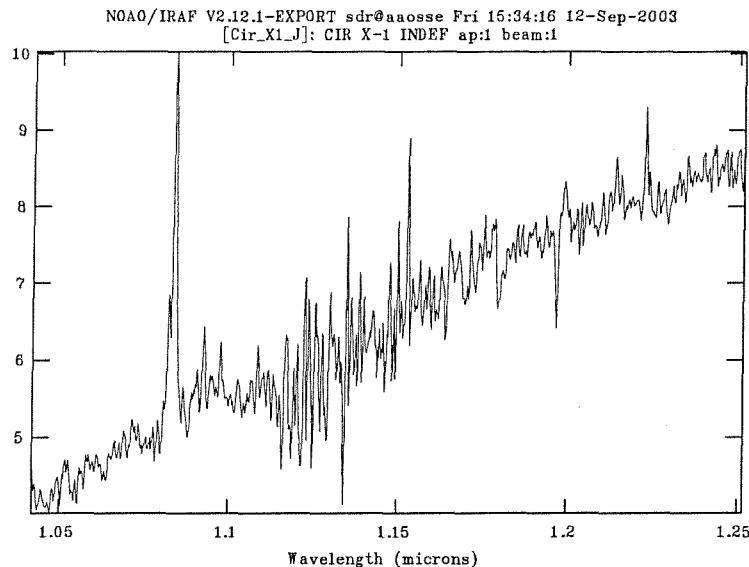
A further surprise is the strong, narrow emission feature at  $\sim 2.275 \mu\text{m}$ , present only at apastron (compare figure 7.6 with figure 7.2). This is unlikely to be an artifact of the reduction process, as it is present in the raw spectra. As yet, this line has no unambiguous identification, but may



**Figure 7.6:** K-band IRIS2 service spectra, taken Sep 4 2003, i.e. just after apastron. Preliminary reductions courtesy S. Ryder, AAO.



**Figure 7.7:** H-band IRIS2 service spectra, taken Sep 5 2003, i.e. just after apastron. A strong set of nearly symmetrical double-peaked emission lines are present, which we identify as Br(10) - Br(14). Preliminary reductions courtesy S. Ryder, AAO.



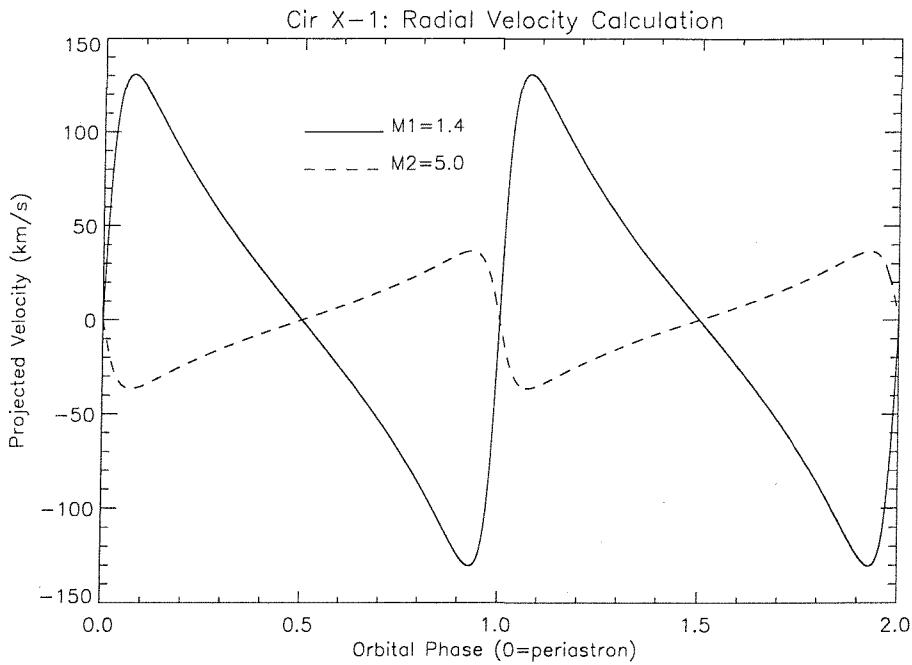
**Figure 7.8:** J-band IRIS2 service spectrum, taken 4th September 2003.. Preliminary reductions courtesy S. Ryder, AAO.

indicate the nature of the donor, since it is also present in the black hole candidate B[e] system CI Cam (Clark et al., 2000). Absolute flux calibration of these spectra have not yet been performed, largely through time constraints, so it is unclear at present if this feature is also present during the 2002 observations. We might estimate flux calibration by assuming the  $Br(\gamma)$  feature is associated entirely with the accretion-outflow process, in which case it would scale with the X-ray luminosity of the source. We might thus expect the  $Br(\gamma)$  line strength to be a factor  $\sim 2 - 3$  lower for the apastron spectrum, which means the  $2.275\mu\text{m}$  line would have been detectable. However, this estimate fails because the continuum is also expected to scale with the X-ray luminosity, making any such ad-hoc estimate extremely unreliable.

If this line were associated with the donor, it might still be *intrinsically* variable throughout the orbit: for example, the recurrent transient A0538-66 shows remarkable variation in its spectral type over the course of a 16.6-day X-ray outburst cycle from B2-III at quiescence to B8-9I in outburst, or a  $\sim 30\%$  jump in  $T_{eff}$  (Charles et al 1983)!

## 7.6 Can $M_2$ be detected from its radial velocity?

With the new combinations of high-spectral resolution IR spectrographs and large telescopes, uncovering the origin of these line features, and therefore the nature of Cir X-1, should be well



**Figure 7.9:** Computed radial velocity profile for Cir X-1 assuming  $e=0.6$ ,  $i^\circ = 60^\circ$ , and angle of observation such that the line-of-sight velocity is zero at periastron, producing the most pessimistic estimate. Here  $M_1 = 1.4M_\odot$  (solid line) and  $M_2 = 5.0M_\odot$  (dashed line).

within the realm of possibility. For example, the long-unresolved issue of the value of  $M_2$  should be amenable to solution through radial velocity studies. The system inclination is likely moderate-high, given the existence of the prominent X-ray dips at periastron, so the projection factor  $\sin i^\circ$  should not reduce the detectability of radial velocity features too severely. Using a spectrograph on an 8m- class instrument, with spectral resolving power  $R \sim 8000$  (e.g. VLT/ISAAC), would lead to centroid velocity resolution  $\sim 20 \text{ km s}^{-1}$ . Assuming from the double-peaked features  $i^\circ \sim 60 \text{ deg}$ , then for eccentricity 0.6, motion of the neutron star will be detectable for  $M_2 \gtrsim 0.4 M_\odot$ , while detection of donor motion requires  $M_2 \lesssim 10 M_\odot$  (see figure 7.9 for an example profile). (We have assumed in these estimates that radial velocity detection corresponds to six distinguishable measurements of the velocity, i.e. the range of projected velocity values must be at least  $6 \times$  the expected velocity resolution.) This shows that the motion of *at least* one component will be measurable for *all* donor masses, while the motion of *both* will be measurable for  $(0.4 \lesssim M_2 \lesssim 10) M_\odot$  (figure 7.9). A lack of any detectable radial velocity shift from either body would thus immediately imply that the accretion/outflow process dominates the IR emission.

## 7.7 Update: Further AAT/IRIS2 Observations of Circinus X-1

To attempt to resolve the motion of  $M_2$ , we proposed for and were awarded follow-up observations of Cir X-1 with the AAT/IRIS2 combination. These observations will take place between Jun 28 - Jul 05, 2004, and are timed to occur through and after apastron passage, i.e. when the contribution from the accretion/outflow process is at a minimum. As outlined in the previous section, these observations should allow limits to be placed on the donor mass.

## Chapter 8

# X-ray Dips as a Diagnostic of Periodic Behaviour in Cir X-1

### 8.1 Summary

We examine the periodic nature of detailed structure (particularly dips) in the RXTE/ASM lightcurve of Circinus X-1. The significant phase wandering of the X-ray maxima suggests their identification with the response on a viscous timescale of the accretion disk to perturbation. We find that the X-ray dips provide a more accurate system clock than the maxima, and thus use these as indicators of the times of periastron passage. We fit a quadratic ephemeris to these dips, and find its predictive power for the X-ray lightcurve to be superior to ephemerides based on the radio flares and the full archival X-ray lightcurve. Under the hypothesis that the dips are tracers of the mass transfer rate from the donor, we use their occurrence rate as a function of orbital phase to explore the (as yet unconstrained) nature of the donor. The high  $\dot{P}$  term in the ephemeris provides another piece of evidence that Cir X-1 is in a state of dynamical evolution, suggesting the possibility that the system might be in a state of post-supernova circularisation. We further suggest that the radio “synchotron nebula” immediately surrounding Cir X-1 is in fact the remnant of the event that created the compact object, and discuss briefly the evidence for and against such an interpretation. The counterpart paper to this chapter is available from MNRAS Clarkson et al. (2004).

## 8.2 Introduction

The various parameters and behaviours of Cir X-1 are detailed in section 7.1. The 16.6 day cycle has long been fit with a quadratic ephemeris, based on radio flares measured at HARTRAO between 1978 and 1988 (Stewart et al., 1991). This ephemeris shows a comparatively high quadratic correction, which implies a remarkably short characteristic timescale for the period evolution ( $P/2\dot{P}$ ) of  $\sim 5600$  years. The longterm X-ray lightcurve of Cir X-1 over the three decades since its identification (Margon et al., 1971) was recently examined (Parkinson et al., 2003, hereafter SP03), using Fourier techniques to determine the first ephemeris for this system based on X-ray data alone. Period determination with this method is dominated by the X-ray maxima (section 8.4), whose large intrinsic phase scatter leads to systematic errors in the resulting ephemeris.

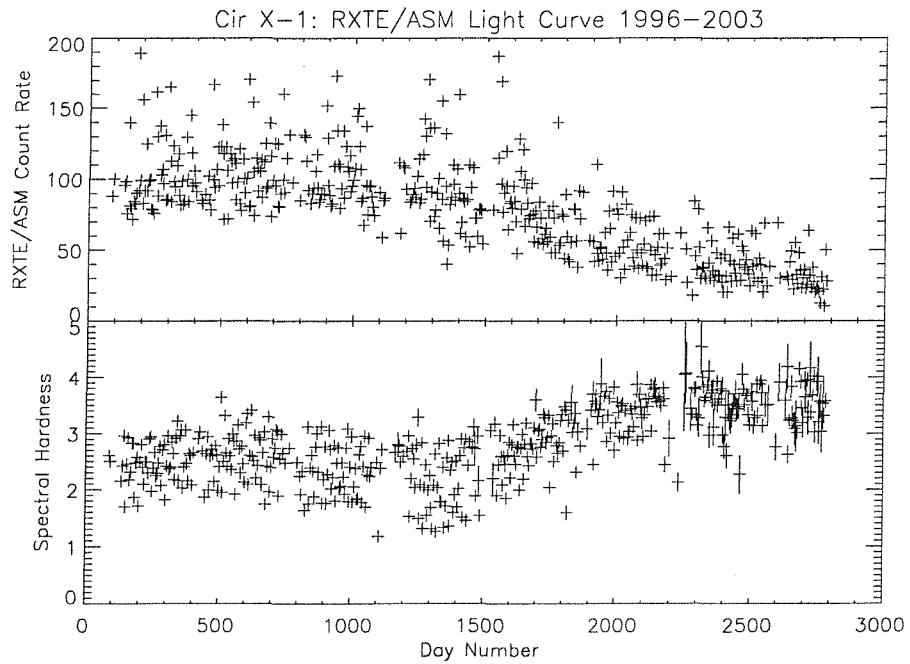
In this chapter we examine in detail the RXTE/ASM lightcurve of Cir X-1, and find the periastron X-ray dips provide a superior system clock to the X-ray maxima. Possible explanations for the phase wandering of the X-ray maxima are discussed. We use the resulting measurement of the intrinsic phase profile of the dips to extract information about the orbital eccentricity  $e$  and discuss the implication of the X-ray results for the age of Cir X-1 and its future determination.

## 8.3 RXTE/ASM Light Curve

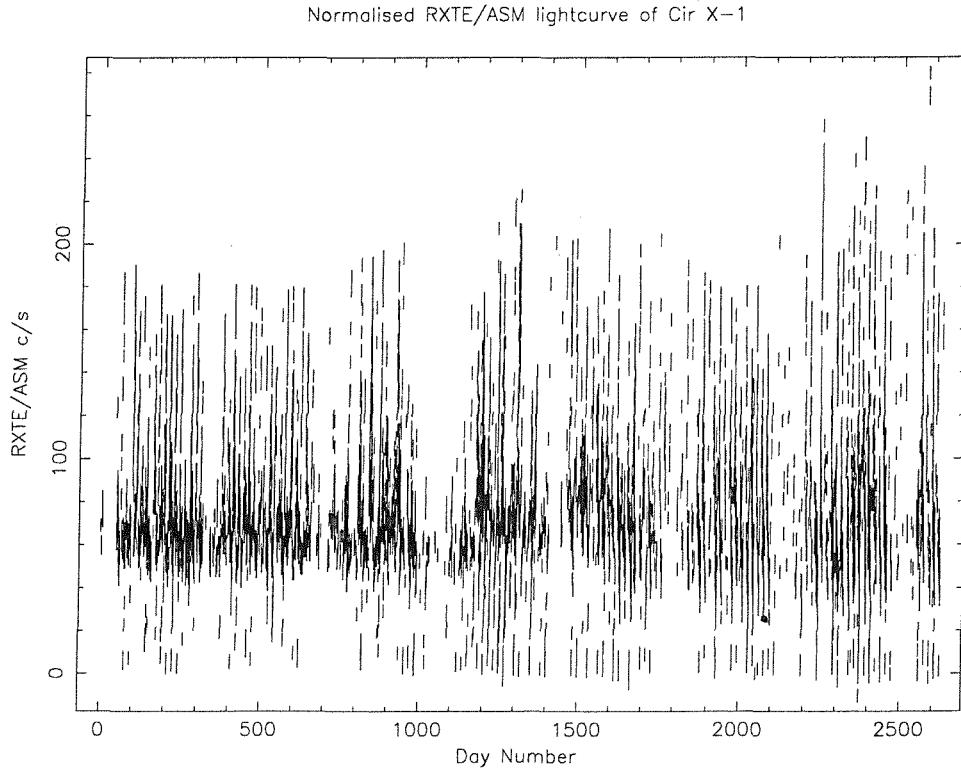
### 8.3.1 Methods of analysis

To illustrate the phase variation of features in the lightcurve, and how this can be used as a tool for testing an ephemeris under the hypothesis of steady behaviour, we used the *dynamic lightcurve*. This breaks the dataset into windows of sufficient length to give good statistics, then folds these windows on the ephemeris specified to show the evolution of phase-dependent behaviour as a function of time. One orbital cycle per data window is sufficient to give good signal to noise to follow the light curve structural changes. Unless otherwise stated, the intensity of the grayscale is directly proportional to the RXTE/ASM count rate.

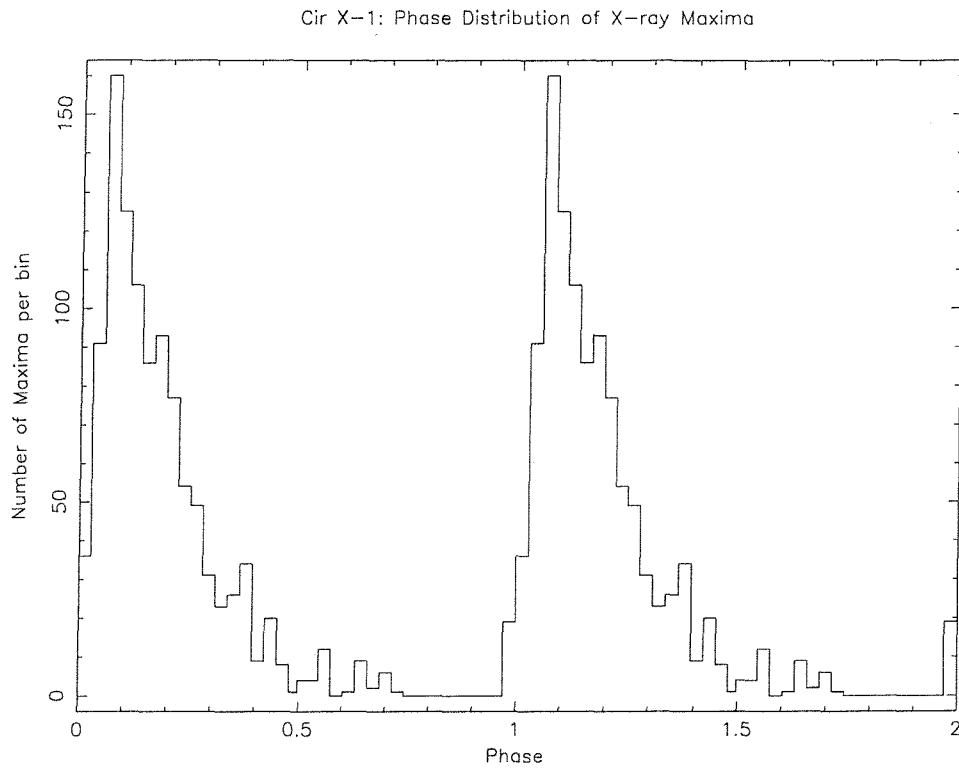
We examine the usefulness of the pure RXTE/ASM lightcurve for period determination by using a dynamic power spectrum approach (see chapter 3). Beyond Fourier techniques, much of the period determination in this chapter is performed under the hypothesis that a certain type of behaviour



**Figure 8.1:** RXTE/ASM dataset used in this chapter, binned at 5 days per bin. Throughout the figures in this chapter time is shown in units of days elapsed since MJD 50000, i.e. 10th October 1995. **Top:** RXTE/ASM count rate. **Bottom:** RXTE/ASM spectral hardness (see Section 8.2).



**Figure 8.2:** RXTE/ASM lightcurve of Cir X-1, normalised to a running average evaluated over intervals of 200 days (section 3.1). Count rate errors, not overplotted in this figure for clarity, are typically  $0.3 \text{ cs}^{-1}$  for the “high” state ranging to  $\sim 3 \text{ cs}^{-1}$  for the “declined” state.

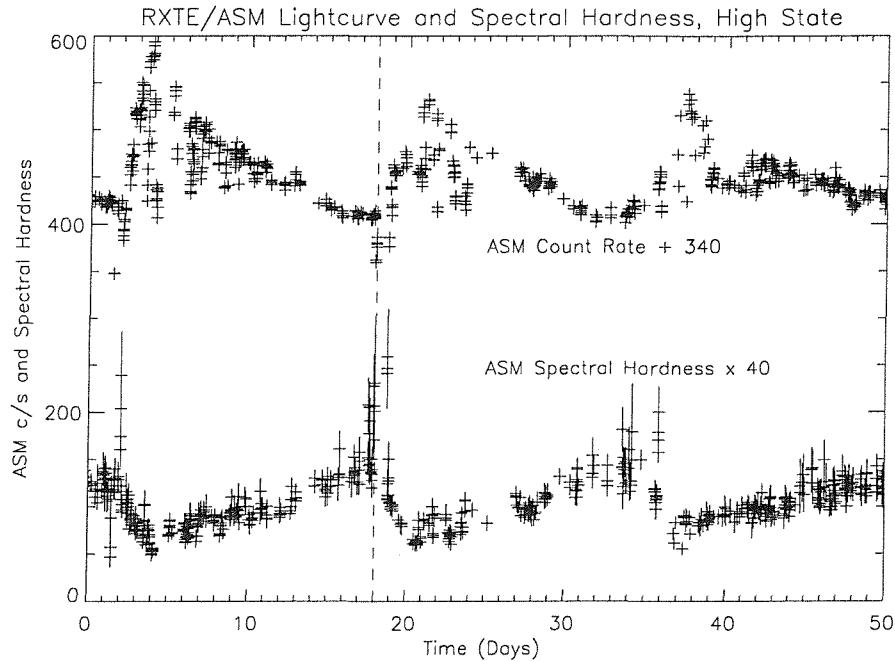


**Figure 8.3:** Distribution in phase of the X-ray maxima, folded on the radio ephemeris (Stewart et al 1991).

should (i) be stable over time and (ii) be minimally scattered in phase at the true period. The best-fit ephemeris is then defined as that which minimises the phase scatter of these events over the entire lightcurve. To determine this ephemeris, the start time, period and quadratic correction are varied and for each combination of parameters the resulting phase scatter is determined. The period search was performed over a range of parameters much wider than the variation in ephemerides to avoid bias towards any set of values. The range of values used for the start time  $MJD_0$ , period  $P$  and quadratic correction,  $C \simeq \frac{1}{2}P\dot{P}$ , respectively, was  $\pm 5$  d,  $\pm 0.2$  d and  $\pm 10^{-3}$  d. Solutions were then re-examined in finer detail, iterating to the accuracy of the ephemeris determination.

## 8.4 X-ray Behaviour

The X-ray lightcurve can broadly be divided into three separate behaviour types: quasi-regular dips, near phase zero of the radio ephemeris; a “steady” average level beneath which the output rarely drops, except during the dips; and periods of peak activity, during which the output varies erratically from  $(1-3) \times$  the steady level. Furthermore, this “steady” level actually evolves in the



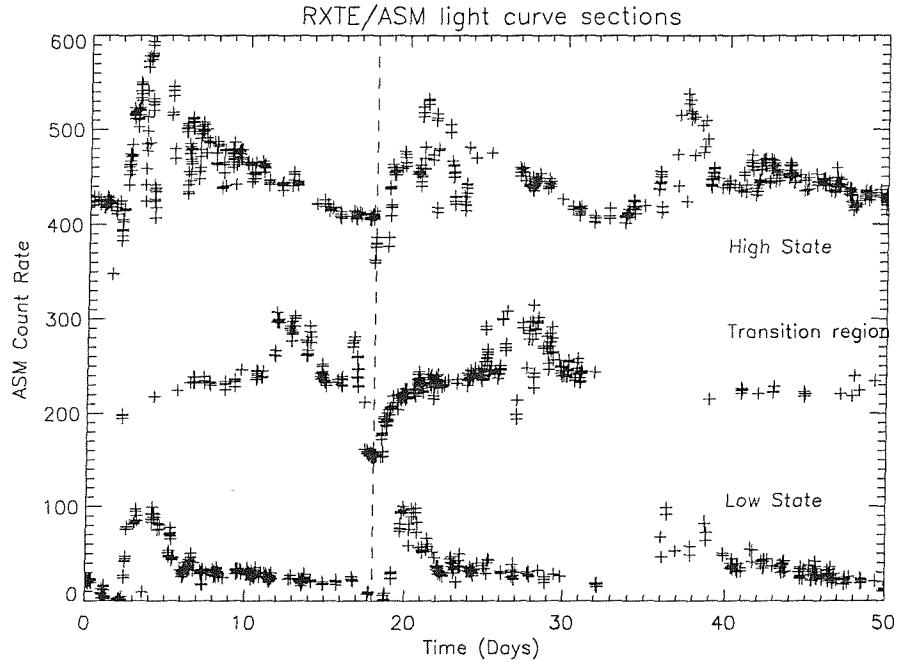
**Figure 8.4:** Top: RXTE/ASM spectral hardness for a 50d stretch of the lightcurve starting at day number 490, during the “high” state. Bottom: RXTE/ASM lightcurve during the same state, illustrating the anticorrelation. For one cycle, phase zero according to the radio ephemeris is plotted as a vertical dotted line.

long-term lightcurve (see figure 8.1) and so, for convenience in this chapter, we divide the lightcurve into three “states” depending on the steady value.

We shall refer to day numbers 0-700 as the “high” state of the source, day numbers 1800-present as the “declined” state, and the interval inbetween as the “transitional” region. We urge the reader not to confuse these denominations with the low/hard and high/soft states often used when describing black hole microquasars (Mirabel & Rodríguez, 1999).

#### 8.4.1 Longterm Decline

The steady level has decreased dramatically over the 7 years of the RXTE/ASM dataset (figure 8.1). During the “high” state this level corresponds to at least  $0.7 L_{Edd}$  for a  $1.4 M_{\odot}$  neutron star (assuming the lower limit 6 kpc distance), but currently represents perhaps only  $0.3 L_{Edd}$ . As pointed out in SP03, the decline represents an apparent return to the X-ray levels of some thirty years previously. This decrease is accompanied by a corresponding increase in spectral hardness. We attempt to normalise the longterm lightcurve by the “steady” level, by calculating the average count rate over bins of 200 days, boxcar smoothing the result and dividing the instantaneous

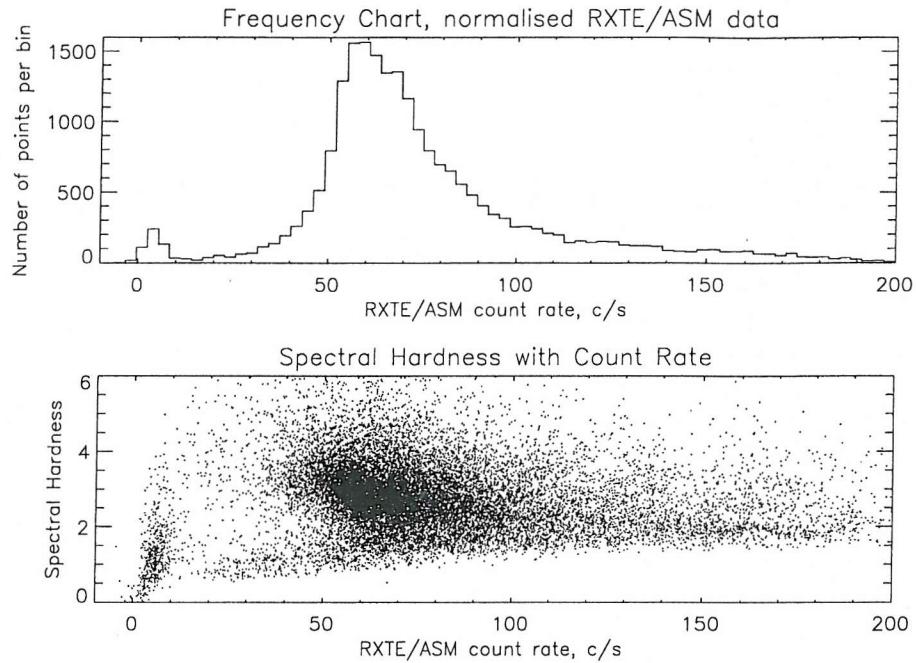


**Figure 8.5:** 50-day sections from the RXTE/ASM lightcurve of Cir X-1 corresponding to the three “states” of the lightcurve. Top: starting at day number 490, count rate offset  $340 \text{ cs}^{-1}$ . Middle: starting day 1425, offset  $150 \text{ cs}^{-1}$ . Bottom: starting day 2350, no offset. For one cycle, phase zero according to the radio ephemeris is plotted as a vertical dotted line.

lightcurve by this running average lightcurve. We notice immediately that to within  $\sim 20\%$  the maximum count rate achieved at outburst scales with the running average level (figure 8.2).

#### 8.4.2 X-ray Maxima

The X-ray behaviour at peak activity is extremely erratic and variable, and to some extent appears dependent on the source state. During the “high” state the X-ray output varies erratically between the steady level and an envelope described by the histogram in figure 8.3. Also present during a few cycles is a smaller secondary maximum, which appears to be similar to a secondary maximum in the radio light curve during this state (Fender, 1998). During the “high” state the spectral hardness is anticorrelated with the X-ray activity (see figure 8.4). The “transitional” interval shows similar evolution of the spectral hardness to the “high” state. The X-ray maxima do not appear to be accurate system clocks, as they occur across a large range of phases in the 16.6 day cycle (figure 8.3). Furthermore, the shifting during this interval of the occurrence of the maxima to phase 0.5 and later has uncovered an apparent gradual recovery in X-ray output from a low level at phase zero (figure 8.5, middle). The lightcurve during the “declined” state appears less erratic

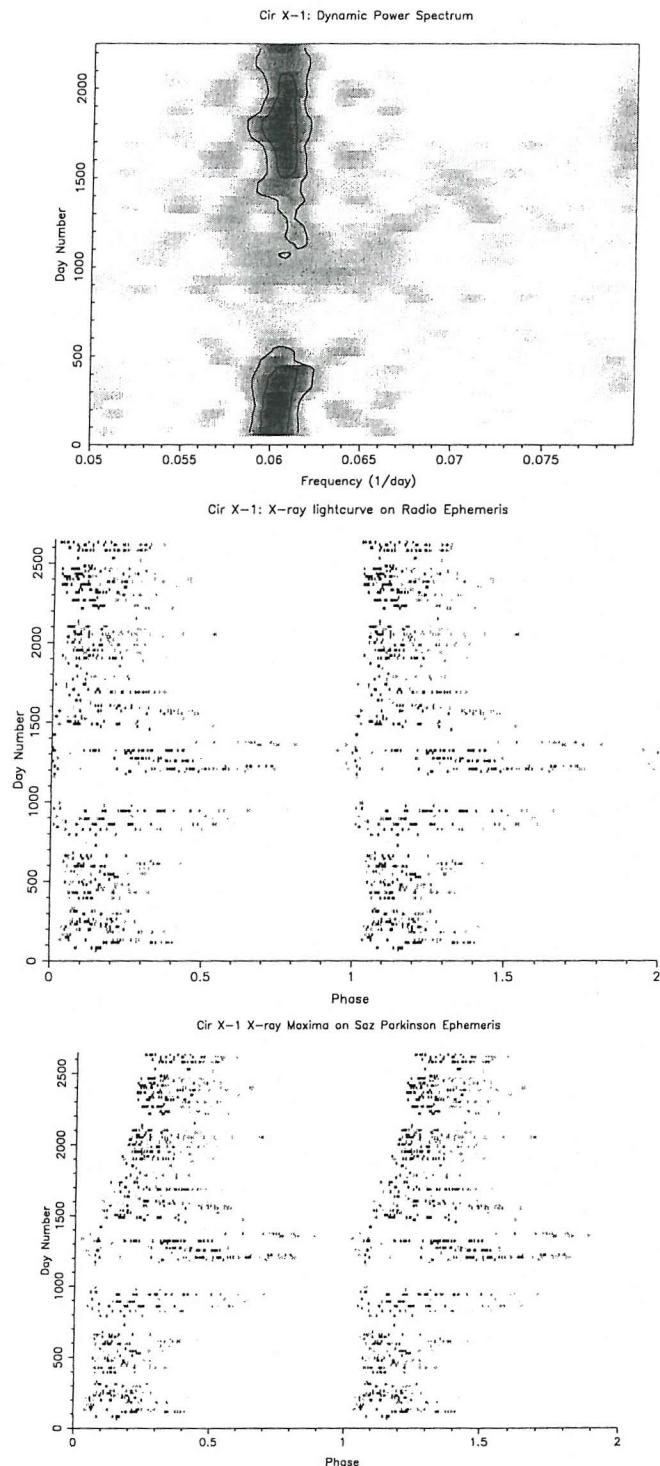


**Figure 8.6:** Top: Frequency chart for the entire RXTE/ASM dataset indicates the existence of a separate population of low-count rate “dips.” Bottom: ASM hardness-count rate plot for the RXTE/ASM dataset.

than the previous two states, with fewer maxima and less phase wandering (figure 8.5, bottom).

#### 8.4.3 X-ray Dips

Phase zero in the radio ephemeris approximately predicts the quasi-regular dips in X-ray intensity (Shirey 1998). These dips form a separate population of datapoints, as can be seen from the frequency-count rate chart of the ASM dataset (figure 8.6). When folded on the radio ephemeris, we see that the dips show far less scatter about phase zero than do the maxima (figures 8.7 & 8.8). X-ray dip spectra are best fit with a two-component model, with a strongly absorbed bright component and an unabsorbed faint component (Shirey 1998). Similar behaviour is seen during X-ray dips in RXTE/PCA observations of black hole candidates GRO J1655-40 and 4U1630-47 (Kuulkers et al., 1998).



**Figure 8.7:** Top: Dynamic Power spectrum of the full RXTE/ASM dataset. Middle: Dynamic Lightcurve showing the X-ray maxima, folded on the radio ephemeris. Bottom: Dynamic lightcurve showing X-ray maxima, folded on the SP03 ephemeris. Note that for the dynamic power spectrum the day numbers represent the beginning of each 400-day data window in the power spectrum.

## 8.5 Ephemerides

As the best X-ray system clocks, the dips can be used to test the various published expressions for the ephemerides of Cir X-1. Until very recently, the most commonly used was the radio ephemeris based on HARTRAO and ATCA observations taken in 1989 (Stewart et al 1991), which, when the zero point is shifted to the RXTE/ASM epoch, can be expressed as (Shirey, 1998)

$$MJD_N = 50082.04 + (16.54694 - 3.53 \times 10^{-5}N)N \quad (8.1)$$

where  $N$  is the cycle number. This ephemeris shows an r.m.s. scatter over the interval of the original radio observations of 0.06 days (Stewart et al 1991), or a  $3\sigma$  scatter in flare occurrence of  $\sim 0.36$  days when extrapolated to the beginning of the RXTE era. As this ephemeris is based on radio data now more than 15 years old, its deteriorating predictive power makes an updated ephemeris desirable. The recent FFT-based analysis of all long-term X-ray datasets of Cir X-1 (SP03) provided the first ephemeris based on X-ray activity alone. This ephemeris retains  $MJD_0$  from the radio ephemeris, but other than this constant offset the two are independent. Shifting to the RXTE/ASM epoch, this X-ray ephemeris is expressed as:

$$MJD_N = 50082.04 + (16.5389 - 1.354 \times 10^{-4}N)N \quad (8.2)$$

with  $3\sigma$  errors on the period  $P$  and quadratic correction  $C$  of  $\sim 1.7 \times 10^{-2}$  d and  $\sim 1.39 \times 10^{-6}$  d respectively.

We test the suitability of the x-ray maxima for period determination by plotting the dynamic power spectrum of the entire RXTE/ASM lightcurve of Cir X-1 with the dynamic lightcurve (figure 8.8). Error in the period determinations is manifested as a systematic drift in phase from cycle to cycle. The transitional period, during which the maxima wander considerably in phase, corresponds to the interval noted by SP03, during which the 16.6d cycle is suppressed. This suggests that the outbursts dominate the period detection from such methods, which is why the transitional region was excised from the RXTE/ASM lightcurve by SP03 when determining their ephemeris.

The radio ephemeris (figures 8.7 & 8.8) shows less of a drift, particularly in the dynamic lightcurve of the X-ray maxima, but a systematic drift in the occurrence of the dips is still visible from day number 1000 onwards. The extrapolated scatter in the radio ephemeris (section 8.5) leads to  $3\sigma$  phase error of  $\sim \pm 0.025$  by the end of the RXTE dataset, but this is still small compared to the systematic shift in phase of  $\sim 0.1$  that we observe. Consequently, we do not believe that this shift can be explained by the extrapolated uncertainty in the radio ephemeris alone. The radio

ephemeris is currently quite accurate in predicting the onset of dip events, but is expected to become less accurate as the events drift farther backwards in phase.

### 8.5.1 Ephemeris based on X-ray Dips

When folded on the radio ephemeris, we see that the dips show far less scatter about phase zero than do the maxima (figures 8.7 & 8.8). Based on the hypothesis that the dips near phase zero on the radio ephemeris are in some way related to the binary orbit, and thus should provide a good system clock, we have attempted to fit an ephemeris on the X-ray dips alone (chapter 3). The radio flares to some extent correlate with the X-ray dips (Shirey 1998), suggesting a causal link, but at present little is known about the mechanisms producing either the radio flares or X-ray dips, which makes it impossible to exploit their relative timing. We minimise assumptions by keeping the time of phase zero as a free parameter in the fit. (If instead the best fit ephemeris is calculated holding  $MJD_0$  at the same value as the radio ephemeris, the smallest value of the phase scatter with respect to phase zero is increased by about 15%). The result of dip phase minimisation is the following ephemeris:

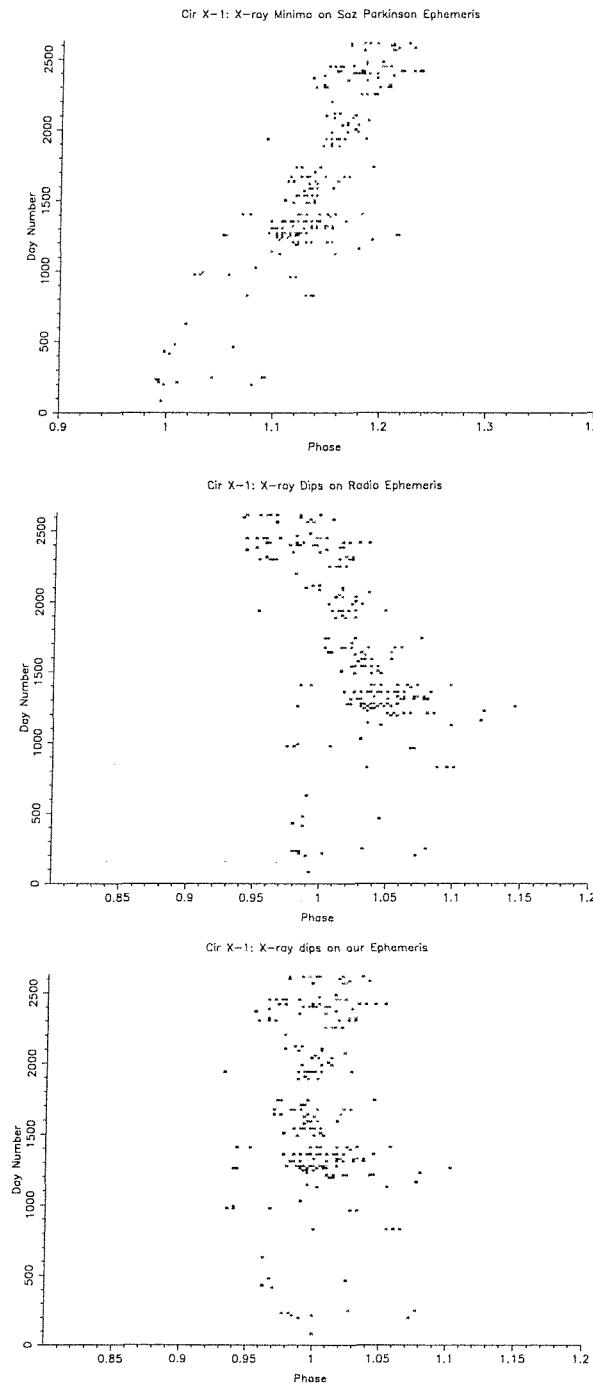
$$MJD_N = 50081.76 + (16.5732 - 2.15 \times 10^{-4}N)N \quad (8.3)$$

with  $3\sigma$  errors on the ephemeris components  $MJD_0$ ,  $P$  and  $C$  of  $8.7 \times 10^{-2}$ d,  $3.2 \times 10^{-3}$ d and  $2.6 \times 10^{-5}$ d respectively (c.f. chapter 3).

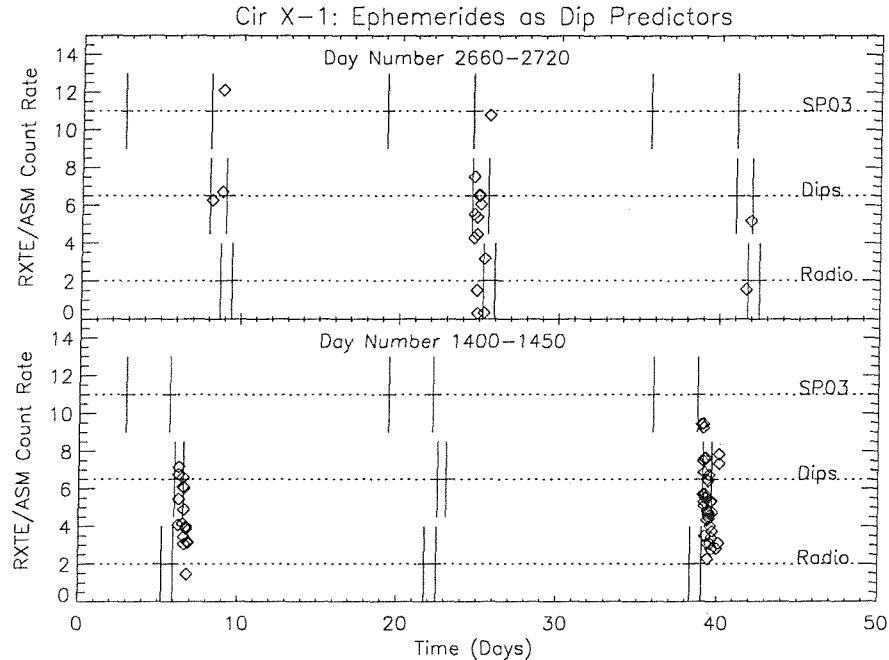
We test the predictive power of this X-ray dip ephemeris by plotting the occurrence of dips against its predictions. As can be seen in figures 8.8 & 8.9, the dip ephemeris predicts the occurrence of dips with a higher degree of accuracy than the radio or bulk X-ray ephemerides. In broad agreement with the SP03 ephemeris, we find a quadratic correction term almost an order of magnitude above that suggested from the radio ephemeris. It appears that  $MJD_0$  for the X-ray dips occurs approximately 7 hours *earlier* than for the radio flares, but this lag is within the extrapolated error of the radio ephemeris of  $\sim 0.36$  days (8.5 hours).

### 8.5.2 The Duration of the Dips

The uneven sampling of the RXTE/ASM makes detailed constraints on the dips as a population somewhat uncertain. For example if two RXTE/ASM points are found at a count rate suggestive



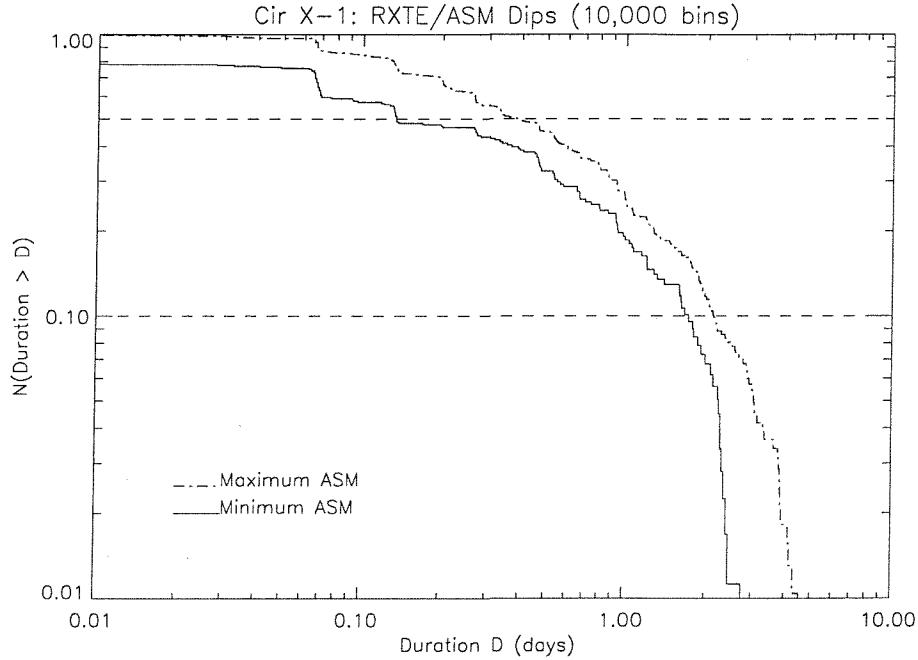
**Figure 8.8:** Top: Dynamic lightcurve of X-ray Dips, folded on the SP03 ephemeris. Middle: X-ray dips folded on the radio ephemeris. Bottom: X-ray dips folded on our ephemeris. In this figure the grayscale is inversely proportional to the RXTE/ASM count rate.



**Figure 8.9:** RXTE/ASM lightcurve at two epochs (diamonds), with  $3\sigma$  bounds on the timing of X-ray dips, as predicted by the three ephemerides (section 8.5). In each plot the ephemerides, from top to bottom, are based on the X-ray maxima (SP03), X-ray dips (this work) and radio flares (Stewart et al 1991). Ephemeris predictions offset for clarity; no prediction of count rate during dipping is made.

of dipping, only upper and lower limits can be set on the duration of the dip. The lower limit on the duration is the interval between the two points (assuming this duration is lower than some sensible limit, thus avoiding the suggestion of 2-week long dips). The upper limit is the time interval between the two out-of-dip points (identified by count rate) directly before and after the dipping points. Between these limits the true behaviour and thus dip duration is unconstrained. As seen in figure 8.10,  $\sim 50\%$  of dips have durations  $\lesssim 0.2$  days while  $90\%$  have durations  $\lesssim 1.7$  days.

However, we note that, of the small number of high time-resolution X-ray lightcurves of Cir X-1 through periastron passage, *all* have shown indications of two separate populations of dips, with several short dips lasting  $\lesssim 5$ ks each, and a single much longer dip ( $\gtrsim 35$ ks). It is unclear if the shorter dips should follow the same phase distribution as the longer dips; we have proceeded on the hypothesis that they are due to the same process of obscuration by accreted matter, and thus have assumed they do. As illustration, we plot in figure 8.11 an example RXTE/PCA lightcurve we obtained in Jun 2002, taken at periastron passage as predicted by the X-ray dip ephemeris and with very high coverage ( $\sim 60\%$ ). The division into dip populations based on duration seems clear; moreover similar behaviour is seen in the PCA lightcurves at and near phase 0 from observations

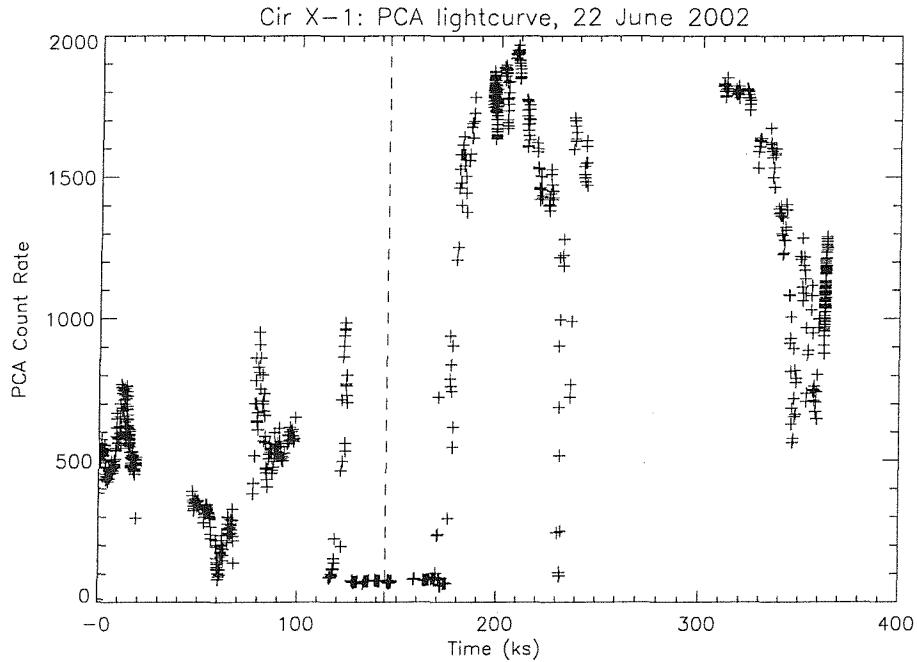


**Figure 8.10:** CDF of the duration of the X-ray dips (in days) as determined from the RXTE/ASM dataset. The upper limit represents the time interval between the first “high” points before and after each dipping interval caught by the ASM, while the lower limit represents the interval between the first and last recorded point in the dipping interval. Points with zero minimum duration (i.e. only one measurement in a dipping interval) have been removed.

of Cir X-1 in the high state at the beginning of the RXTE era (Shirey 1998). Without more high time-resolution lightcurves of Cir X-1 it is unclear if this population division is robust; it is not reproduced in the RXTE/ASM lightcurve (see e.g. figure 8.10, which shows no clear break in dip populations).

## 8.6 Discussion

The central issue with Cir X-1 has been the nature of the companion star, and the ASM lightcurve does not appear immediately to resolve the issue. Furthermore, the system eccentricity and inclination have not been constrained (we note that the eccentricity is also only estimated for A0538-66, the prototypical Be XRB, whose X-ray lightcurve, when active, is strikingly similar to Cir X-1; Charles et al 1983). We examine here the use to which the RXTE/ASM lightcurve can be put to address these issues.



**Figure 8.11:** RXTE/PCA lightcurve of Cir X-1, at periastron interval as predicted by the X-ray dip ephemeris (vertical line). This lightcurve is not background subtracted, however the presence of a deep, long X-ray dip at phase 0, and subsequent and prior shorter dips, is evident.

### 8.6.1 A simple wind model

We attempt to model the X-ray lightcurve by using a simple spherical wind model, on the grounds that during the high state accretion may be quasi-spherical at periastron. Taking the Bondi-Hoyle result, we consider that matter is accreted only if its specific kinetic energy is less than its gravitational potential energy. In the vicinity of the neutron star, wind material moves nearly radially from the donor at  $r < a(\theta)$ , thus matter satisfying this condition lies within a cylinder of radius

$$r_{acc} \simeq \frac{2GM_x}{v_{rel}^2} \quad (8.4)$$

where  $v_{rel}^2$  is the velocity of the wind material relative to the neutron star. The mass accretion rate then depends on the binary separation  $a(t)$  as

$$\dot{M}_1 \simeq \pi r_{acc}^2 v_{rel} \rho = \frac{4\pi G^2 M_x^2 \rho(a)}{v_{rel}^3(a)} \quad (8.5)$$

The X-ray lightcurve is then simulated from estimates of  $\rho$  and  $v_{rel}$ . We assume azimuthal density variations to be unimportant (i.e. along a radial path, matter deflected off through rotation is

replaced by matter deflected on from adjacent paths), and approximate the density as a radial power law

$$\rho = \rho_0 \left( \frac{r}{R_2} \right)^{-n} \quad (8.6)$$

The velocity of the wind relative to the accretor,  $v_{rel}$ , consists of three components; the radial wind velocity in the frame of the donor  $v_{r,w}$ , the tangential wind velocity in the same frame  $v_{\phi,w}$  and the relative velocity of the centres of mass of the donor and accretor  $v_{12}$ . Continuity yields the radial component

$$v_{r,w} = v_0 \left( \frac{r}{R_2} \right)^{n-2} \quad (8.7)$$

and donor rotation with period  $P_{sp}$  (in the frame of the observer) produces the azimuthal component from angular momentum conservation

$$v_{\phi,w} = \frac{2\pi R_2^2}{P_{sp} r} \quad (8.8)$$

The positions of the two stars are computed numerically, producing the relative velocity  $v_{12}$  from the time increments in the computation. This toy model thus produces wind lightcurves for estimates of the unknowns  $M_2$ ,  $R_2$ ,  $P_{sp}$ ,  $v_0$ ,  $e$  and the wind power-law index  $n$ . The resulting lightcurve can be quite complex: the contribution from the density power law always peaks at periastron, but the occurrence of the trough in  $v_{rel}(\phi)$  (c.f. eq 8.5) varies with the velocity balance.

We use a  $15M_{\odot}$  donor in order to generate a strong stellar wind, with  $e = 0.7$ . The power-law expressions for  $v_{w,r}(r)$  and  $\rho(r)$  are approximate fits to more realistic derivations. For these fits, the index  $n$  is estimated from the results of standard treatment of steady, adiabatic, spherically symmetric gas flow (e.g. Frank et al., 2002; Choudhuri, 1998), which yields five classes of solution for  $v_{r,w}(r)$ . We assume  $\partial v_{r,w} / \partial r > 0$  for all  $r$ , corresponding to class  $V$  in figure 6.9 of Choudhuri 1996 or class 1 in figure 2.1 of Frank, King & Raine 2002. For this class of solution,  $\partial^2 v_{r,w} / \partial r^2 < 0$ , therefore our power-law fit must have ( $2 \lesssim n \lesssim 3$ ). The spin period of the donor is unknown, but a reasonable assumption is that the angular velocity of the neutron star is within a factor  $\sim$ a few of the angular velocity of the accretor-donor pair about each other at closest approach, which would imply  $P_{sp} \sim 3$ d. For  $v_0 \gtrsim 0.1$ km s $^{-1}$ , spin periods  $\gtrsim 1$  day have little effect on the shape of the lightcurve. The closest fit to the observed lightcurve that can be produced by wind alone is given in figure 8.12, corresponding to  $v_0 \simeq 50$ km s $^{-1}$  and  $n = 2.8$ . A steep power law index is required to produce the rapid rise in  $\dot{M}_1$  at periastron. This is clearly dissimilar to the observed lightcurve, as the accretion rate always drops nearly to zero within a few days of the periastron outburst. Attempts to remove the abrupt drop to quiescent level after periastron passage by invoking fast rotation (to reduce the relative velocity of wind and accretor during the appropriate part of the

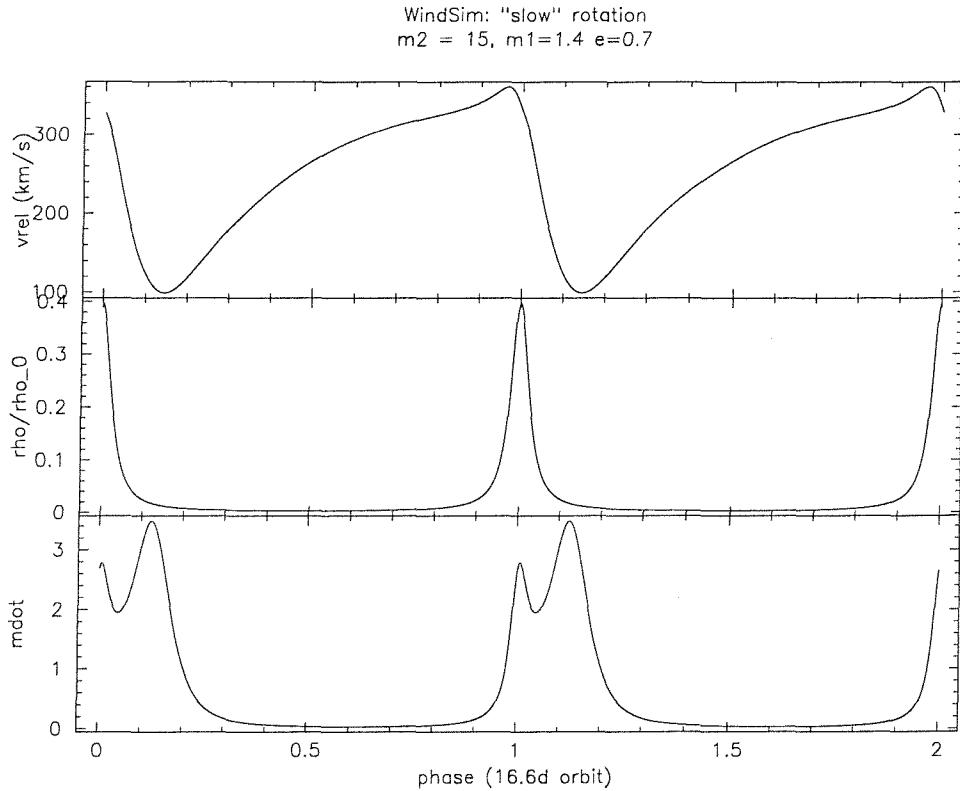


Figure 8.12: Example wind-driven accretion model following section 8.6.1.

orbit) require donor spin periods far higher than breakup for any reasonable stellar companion; thus, as expected, stellar wind alone cannot reproduce the shape of the X-ray lightcurve.

### 8.6.2 The X-ray maxima

The significant phase wandering of the maxima during the “transitional” state of the lightcurve suggests that they are associated with some feature of the system that is not entirely locked into the orbital clock. This feature is not completely independent of the orbital clock either, as there are still *no* cycles for which an X-ray peak appears between orbital phase 0.8 and 0.9, although during the “transitional” interval there are cycles in which the peak occurs slightly *before* phase zero. The X-ray dips appear to be a better system clock than the outbursts, and we identify their occurrence with changing mass transfer from the donor during periastron passage. Stellar wind accretion cannot account for the phase wandering of the maxima; for example, in the case of GX301-2, an eccentric ( $e = 0.462$ ) X-ray binary in which the accretion takes place via wind plus a small accretion stream Leahy (2002), the phase wandering of maxima is tiny compared to that for Cir X-1.

### 8.6.3 The Accretion disk and the Maxima

We are thus left with the accretion disk to provide the source of material for the maxima, as it is capable of doing so in a manner that could show sufficient variation in the time of occurrence. The mechanism to do so is not immediately clear. The obvious analogy would be the dwarf nova cycles seen in cataclysmic variables (Warner 1995), however this interpretation suffers obvious limitations. The irradiation from the neutron star is likely to keep the inner disk in a permanent high state, suggesting such cycles would require extremely fine tuning (e.g. Frank, King & Raine 2002). Furthermore, this model does not account for the scaling of the X-ray output at maximum with the light curve mean.

We prefer to argue that the X-ray maxima are due intervals of increased mass transfer onto the compact object as a result of a combination of disk perturbations and varying mass transfer rate, both driven to some extent by the eccentric orbit. These variations will occur at or near the viscous timescale of the disk as this is the timescale for disturbances to propagate through the disk. This interpretation has the attraction that it allows the outbursts to occur over a quasi-steady base level of accretion, as is observed with RXTE/ASM: such response to perturbations will occur preferentially in the outer disk, where material is closer to the impact region of the varying mass transfer stream. This model also has the benefit of predicting how the amplitude of the maxima will vary with the running average: both should scale with the global density of matter in the accretion disk.

It is still unclear what causes the phase of the X-ray maxima to vary so wildly during the “transitional” region, but this interval does immediately precede the longterm return to low levels of X-ray output, so therefore may be related to it. However, it is also possible that the longterm decline is part of an even longer-term pattern of variation in the mass transfer from the donor, possibly due to tidal evolution of the system.

An alternative model for both the longterm variation outburst level and the phase wandering thus involves Cir X-1 as a Be/X-ray transient, in which the reformation and destruction of the equatorial disk provides the mechanism for the variation of the X-ray lightcurve on a timescale of decades. Variations in the phase of X-ray maxima on a timescale of years are then produced by the precession of the equatorial disk and/or evolution of 2D structure in this disk. It is known that in several Be/X-ray transients with similar orbital periods to Cir X-1, the equatorial disk is tilted with respect to the plane of the orbit and precesses, as seen in for example V0332+53

(Negueruela et al., 1999) and 4U 0115+63/V635 Cas; (Negueruela et al., 2001, hereafter N01). This disk may also be destroyed and reformed on a timescale of months - years, as seen in V635 Cas (N01). Furthermore, 2D structure in equatorial decretion disks through global one-armed oscillations is predicted theoretically (Okazaki, 2001; Okazaki & Negueruela, 2001) and indeed suggested observationally from V/R variability in several Be/X-ray transients on a timescale of years to decades (Hanuschik et al., 1995). The high and unknown reddening towards Cir X-1 makes constraining such models extremely difficult for the foreseeable future (chapter 7).

#### 8.6.4 X-ray dips as tracers of $\dot{M}(t)$

Whichever ephemeris is used (we are assuming each is a different approximation to the true orbital ephemeris), the dip occurrence rate depends strongly on orbital phase, as all dips are clustered around  $\pm \sim 1$  day of phase zero. That this is in sharp contrast to the maxima used by SP03 to determine the ephemeris does not invalidate their discussion, as we are examining a different feature of the lightcurve. We examine here the information that can be extracted about the binary components from this property.

Absorption dips have been used as a diagnostic in the RXTE/ASM lightcurve of the near-circular orbit black hole candidate Cyg X-1 (Bałucińska-Church et al., 2000). In this case the occurrence rate of X-ray dips was also shown to vary strongly with orbital phase, in a manner dependent on the line of sight to the binary: the dips occurred with greatest frequency when the donor passed directly in front of the accretor. In this model the X-ray dips are due to photoionisation in the wind of the donor. Applying this model to Cir X-1, however, the near-coincidence of the dips with times of maximum activity at other wavelengths (and thus periastron passage) would suggest the line of sight to the observer passes through the densest parts of a donor wind at periastron, i.e. the semimajor axis of the binary orbit would have to be almost aligned with our line of sight to the binary. We reject this as requiring fine tuning.

However, subsequent studies of Cyg X-1 with RXTE/PCA divide the X-ray dips in Cyg X-1 into two populations based on spectral behaviour: “Type A” dips, concentrated about donor interposition phases as above, and “Type B,” which occur throughout the orbit (Feng & Cui, 2002). One suggestion put forward for these dips is absorption of an X-ray emitting region by optically thick “clouds” in the accretion flow near the accretion disk. The number density of such clouds in the vicinity of accretion stream impact in the outer disk scales strongly with  $\dot{M}$  (equation

10 of Frank, King & Lasota 1987, hereafter FKL87). In an eccentric system such as Cir X-1,  $\dot{M}$  will be a strong function of orbital phase (see Charles et al 1983 or Brown & Boyle 1984). We thus identify the X-ray dips exhibited by Cir X-1 as “Type B” dips, which here occur only during episodes of enhanced mass transfer at periastron passage. This has the attraction of removing the need for fine-tuning of the line of sight to the system.

The accretion disk will respond to changes in  $\dot{M}$  on its viscous timescale. We can estimate this value from the phase spread in the outbursts (figure 8.3) or alternatively from assumptions on the mass transfer rate and viscosity (Frank, King & Raine 2002), to be of order 10 days. The inner disk can respond more quickly but is likely to be highly ionised (Brandt & Schulz 2000), thus will not figure in the formation of the X-ray dips. The response of the accretion disk to varying  $\dot{M}$  thus occurs on too long a timescale to affect the relationship between  $\dot{M}(t)$  and the occurrence rate of the dips.

The velocity of the gas flow in the accretion stream will be greater than the adiabatic or isothermal sound speeds, allowing internal shocks to be neglected in the direction of the flow (c.f. Frank, King & Raine 2002; we assume further that any variation in sound speed along the stream due to changing  $\dot{M}$  is insufficient to bring the sound speed above the freefall velocity at any point in the stream). The sound travel time across the stream will be long compared to the transfer time to the accretion disk (c.f. Lubow & Shu 1975), so we neglect the internal response of the gas in the accretion stream. The impact of the accretion stream on the disk results in shock formation and the condensation of cold absorbing clouds on a timescale similar to the recombination timescale  $t_{rec}$  (FKL87),  $\sim 1\text{-}30$  minutes for the parameters of Cir X-1. The evaporation timescale for such clouds is of order a few dynamical timescales  $t_{dyn}$  (FKL87, Krolik, McKee & Tarter 1981), so they will persist for  $\gtrsim 5$  hours. The formation rate of clouds leading to X-ray dips will thus respond quickly to changes in  $\dot{M}$ , and furthermore such dips will persist for a long enough interval to affect the X-ray lightcurve. The  $\gtrsim 5$  hour timescale for cloud evaporation suggests reduction in dip occurrence rate may lag shortly behind the reduction in  $\dot{M}$ , which is fully consistent with the slight apparent asymmetry in the dip population (section 8.6.6).

### 8.6.5 Using the dips to investigate the binary parameters

We attempt in this section to relate the rate of occurrence of X-ray dips to the donor mass  $M_2$ , the eccentricity  $e$  and the donor temperature  $T$ . To do so we require a model predicting  $\dot{M}(t)$

as a function of these parameters. A preliminary theoretical investigation of  $\dot{M}(t)$  due to tidal lobe overflow in an eccentric-orbit binary was undertaken (Brown & Boyle, 1984, hereafter BB84) (hereafter BB84), which we summarise here.

The surface of the donor is described by an “effective Roche Lobe radius”  $R$ , which is the value the Roche Lobe radius would take if the orbit were circular with separation equal to the separation in the orbit at time  $t$ . Mass transfer towards the accretor is then envisaged as free thermal diffusion along the effective Roche eqipotentials, i.e.:

$$\dot{M}_2 = -m_P \int_R^\infty 2\pi r \int_0^\infty f(v, r) v dv dr \quad (8.9)$$

where  $f(v, r)$  is the number density of protons at altitude  $r$  moving with velocity  $v$  in one particular direction, chosen in this case as the direction of the accretor  $x$ . Describing the outer layers of the donor as an exponential density decrease with constant pressure scale height  $H = \left(\frac{\partial P}{\partial \rho}\right)_T = (kTr^2)/(Gm_P M_2)$ , and  $f(v, r)$  by a Maxwellian distribution then gives

$$f(v, r) = \frac{\rho_0}{m_P} \left(\frac{m_P}{2\pi kT}\right)^{1/2} \exp\left[\frac{-(r - R_0)}{H}\right] \exp\left[-\frac{1}{2} \frac{v^2}{c_S^2}\right] \quad (8.10)$$

with  $c_S^2 = kT/(m_P P)$ . Assuming  $H/R \ll 1$  then the mass diffusion rate is then the usual approximation of the density at the Roche surface  $\rho$ , multiplied by the Roche nozzle area  $\sim RH$  (Warner, 1995), i.e.

$$\dot{M}_2 = -(2\pi)^{1/2} RH \rho_0 \exp\left[\frac{-(R - R_0)}{H}\right] \quad (8.11)$$

As the orbit is eccentric, the equivalent Roche Lobe radius  $R = \alpha d$ , (where the binary separation  $d$  and correction from rotation  $\alpha$ ) both depend on orbital phase  $\phi$ . Joss & Rappaport (1984) fitted an expression for  $\alpha(\phi)$ , the exact form of which depends in a somewhat complex manner (eqn 5 of BB84) on the system masses and the quantity

$$K \equiv \frac{\omega_{rot}^2}{\omega_{orb}^2} \frac{(1+e)^4}{(1+e \cos \phi)^3} \quad (8.12)$$

( $\omega_{orb}$  and  $\omega_{rot}$  being the orbital and rotational angular velocities respectively of the donor at periastron, with  $\phi$  the angular co-ordinate around the orbit and  $e$  the system eccentricity). Assuming the instantaneous Roche lobe is stationary during periastron passage, this allows  $\dot{M}_2$  to be modelled through insertion of  $R(\phi)$  into the expression for  $\dot{M}_2$ . An innovation of BB84 was to model the effect of the instantaneous variation in the effective Roche lobe radius, by assuming only particles moving with positive velocity relative to the Roche Lobe can be accreted. This results in multiplication of the Maxwellian integral by a factor

$$\gamma(\epsilon) = 2 \int_0^\infty \exp[-(u + \epsilon)^2] u du \quad (8.13)$$

where  $\epsilon = (\dot{R})/(2^{1/2}c_S)$ , and  $\dot{R}$  is found by carrying out the differentiation of  $R = \alpha d$  (equations 13-14 of BB84). The mass transfer rate is then predicted by the expression

$$\frac{\dot{M}(\phi)}{\dot{M}_{peri}} = \frac{\alpha}{\alpha_{peri}} \frac{1+e}{1+e \cos \phi} \gamma(\epsilon) \exp \left[ -\alpha \beta e \left( \frac{1-\cos \phi}{1+e \cos \phi} \right) \right] \quad (8.14)$$

with

$$\beta = \frac{a(1-e)}{H} \quad (8.15)$$

Thus the  $\dot{M}_2(\phi)$  profile is predicted from the system masses  $M_1$  &  $M_2$ , the eccentricity  $e$  and the temperature  $T$  of the envelope material.

This semianalytic approach was followed up by numerical simulation (Boyle & Walker, 1986), exploring the range of parameter space for which capture of transferred matter might occur, in an attempt to model the X-ray activity of A0538-66. Only a single binary orbit was followed in these simulations, so it is not possible to determine the fate of material not immediately captured by the neutron star. Earlier simulations (Haynes et al., 1980) did not take into account hydrodynamic effects, but in covering more orbits suggested that matter not immediately accreted in one periastron passage may be accreted in subsequent cycles, complicating the accretion profile. We thus reject the capture condition espoused by Boyle & Walker (1986), in which gas is only captured if its specific potential energy is greater in magnitude than its specific kinetic energy. More modern simulations of eccentric orbit accreting systems suggest this subsequent accretion may be extremely important in determining the X-ray output of an eccentric binary (Podsiadlowski 2003 priv. comm.). However, as our interest is not in the X-ray generation rate per se but in absorption due to the matter stream, we take the BB84 model for lack of current alternatives. The reader is therefore cautioned that the model on which our constraints will be based is in some sense a preliminary one.

### 8.6.6 A simple model

We use the  $\dot{M}$  predictions of BB84 to attempt to fit the profile of dip occurrence rate as a function of  $M_2$  and  $e$ . We take this opportunity to remind the reader of our core assumption; that the rate of occurrence of the X-ray dips does indeed represent the instantaneous accretion rate. One can readily imagine a model in which this is not (quite) the case, for example if the dips are due to obscuration in jets produced at periastron (consistent with the low inclination suggested by the radio jets; Fender et al 2004); in this case the radio ephemeris would still represent the orbital motion of the system. There would then be an offset between true periastron and the interval of

most frequent dipping, due to the time taken for the system to travel through its orbit to the point at which the line of sight passes through the jet, and our assumption would be violated. However, were this offset to be large, we would expect a significant peak in X-ray emission before the onset of dipping; that this is not observed suggests it should be a good approximation to assume the rate of dipping does trace the accretion rate (and therefore the orbital phase  $\phi$ ). We estimate  $\beta$  by assuming the pressure scale height  $H$  is constant over the outer layers of the donor, and thus its instantaneous Roche Lobe radius  $R$  for the altitude  $r$ . The standard Roche Lobe prescription (Egginton 1983) is used for  $R$ , which combined with the orbital period of Cir X-1 gives

$$\beta \simeq 106 \left( \frac{q}{f^2(q)(1+q)^{1/3}} \right) \left( \frac{1}{T_4(1-e)} \right) \quad (8.16)$$

where  $f(q)$  is the tidal lobe radius at periastron as a fraction of the semimajor axis,  $q = M_2/M_1$  and  $T_4$  is the temperature at the surface of the donor, at periastron, in units of  $10^4$  K. As the X-ray dips occur over a very limited phase range, we neglect any radius variation that may occur in the donor over the course of the periastron passage, beyond the variation of the tidal lobe itself. Furthermore, we follow Charles et al (1983) in assuming the surface of the donor corotates with the orbit at periastron (i.e.  $\omega_{rot} = \omega_{orb}$  in (8.12) above). It is currently unclear what the quantitative effect of non-synchronous donor rotation will be on the  $\dot{M}(t)$  profile. We expect the change in the amount of matter ejected from the surface of the donor to compete with an opposing change in the amount of matter that can be captured by the accretor (Petterson, 1978), however the balance of the effects is far from obvious. In the absence of a model for the phase dependence of this behaviour, we retain the assumption of corotation at periastron. We merely caution here that if Cir X-1 is a dynamically young system, as suggested by its apparent eccentricity and other behaviours (section 8.6.5), non-synchronous rotation at periastron is highly possible.

The final step when making predictions is to translate  $\dot{M}(\phi)$  into an expression for  $\dot{M}(t)$ . This is readily accomplished by evaluating the time taken for the orbital phase to change from  $\phi_1$  to  $\phi_2$ :

$$t_{12} = \frac{P_{orb}}{2\pi} (1-e^2)^{3/2} \int_{\phi_1}^{\phi_2} (1+e \cos \phi)^{-2} d\phi \quad (8.17)$$

To save computation time, it was elected not to evaluate the integral directly for each  $(\dot{M}, e)$  combination, but instead replace it with an analytic expression from a table of indefinite integrals (Petit-Bois 1961):

$$\int (1+e \cos \phi)^{-2} d\phi = \frac{2}{(1-e)^{3/2}} \tan^{-1} \left[ \left( \frac{1-e}{1+e} \right)^{1/2} \tan \frac{\phi}{2} \right] - \frac{e}{1-e^2} \frac{\sin \phi}{1+e \cos \phi} \quad (8.18)$$

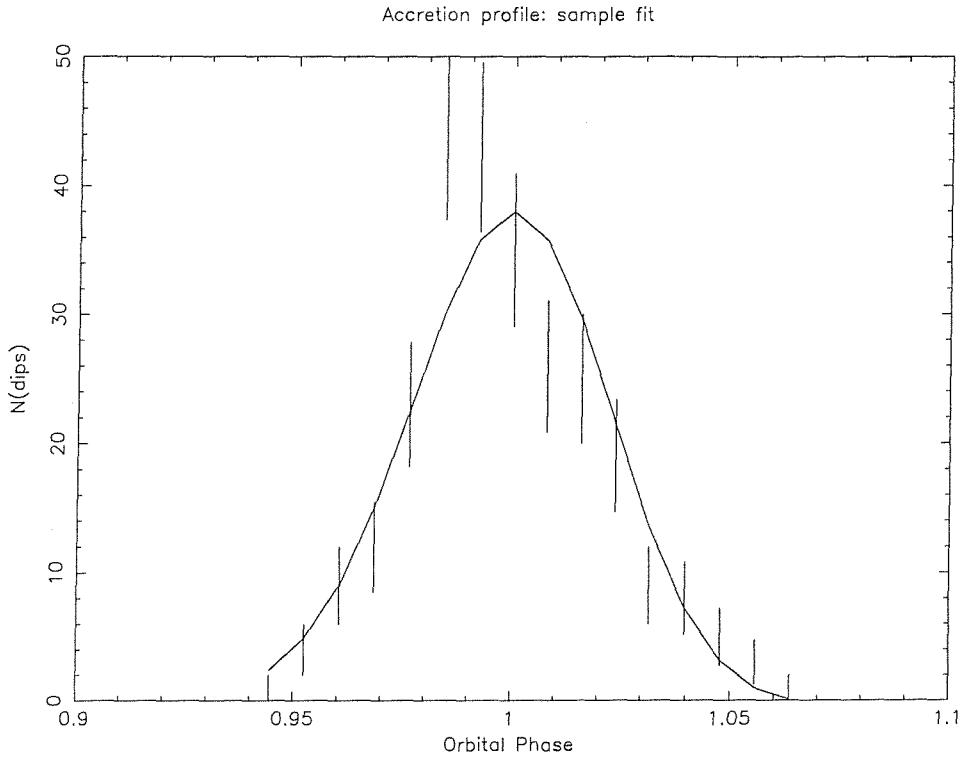
with  $\phi = 0$  at periastron as is usually taken. In the cornerstone paper establishing Cir X-1 as a low-mass system based on its kinematics and identification with SNR G321.9-0-3, the approximation

used is inappropriate for all but high eccentricity, as setting  $e = 0$  gives a different answer from the  $P_{orb}/(\phi_2 - \phi_1)$  demanded from symmetry (equation 10 of Tauris et al 1999). Expression (8.18) may cease to be valid for high eccentricity; it clearly breaks down for all  $\phi$  at  $e=1$ , while (8.17) is still be valid for restrictive ranges of  $\phi$ . We thus check the accuracy of our imported expression (8.18). The integral (8.17) was evaluated numerically for a given set of parameters and the approximation 8.18 overplotted for the same parameter set. The direct numerical evaluation and analytic expression are found to differ by less than 10%.

We plot the dip occurrence profile, along with an example  $\dot{M}(t)$  prediction, in figure 8.13. As can be seen, the profile can be fit quite readily with the tidal lobe model, within the errors (assumed gaussian) between the measured X-ray dip rate and the intrinsic cloud formation rate. For the following analysis, dips occurring in the phase range (0.95 - 1.05) are used, as outside this region the errors are comparable to the observed dip rate. Clearly the fit is degenerate, in that in principle any combination of  $q, T_4$  and  $e$  can reproduce the observed dip profile. We narrow the parameter space somewhat by noting that the donor mass is unlikely to exceed the  $\sim 15M_\odot$  necessary to form a neutron star (Tauris & Heuvel, 2003) or else it would have noticeably evolved. More stringent constraints on the donor mass through direct detection of the system are not possible due to the uncertain location of the IR-emitting regions (chapter 7) and the uncertain level of reddening (Glass 1994).

A large range of  $T_4$  is in principle possible, as there are several sources of energy dissipation that could heat the surface of the donor significantly above its unperturbed value of  $T_4 \sim 0.5 - 1$ . The accretor has been among the brightest X-ray sources in the sky, so might be expected to irradiate the donor to  $T_4 \gtrsim 10$  (e.g. van Teeseling & King 1998; we neglect here the enhanced mass loss such a mechanism might produce). The supernova itself may have deposited a large amount of energy in the donor leading to significant heating of the stellar envelope to  $T_4 \sim 10$  (Marietta et al., 2000; Podsiadlowski, 2003). Finally, if the system really is eccentric, tidal circularisation will give rise to internal heating of the donor, also bringing  $T_4$  to  $\sim 10$  or higher (see Podsiadlowski, 1996). To account for these possibilities, we repeat the  $\dot{M}(t)$  profiles for several input values of  $T_4$  in the range ( $0.1 \leq T_4 \leq 100$ ).

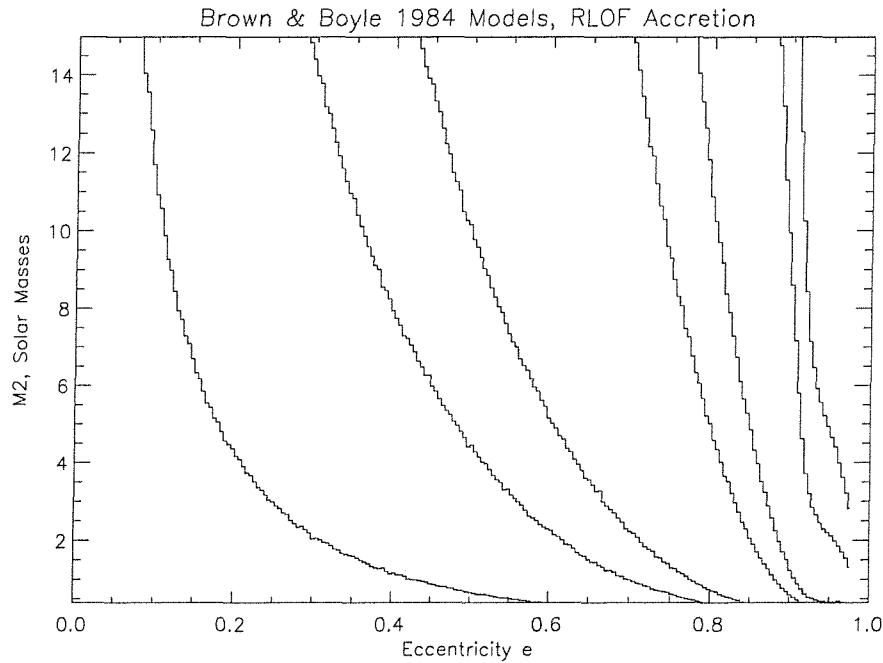
Initial sets of simulations were fit to the observed dip profile through comparison of the FWHM of the prediction with the observations (figure 8.13). In further simulations the predictions were evaluated using the least-squares fit. The results are plotted in figure 8.14. We note that the condition of tidal lobe overflow at periastron leads to larger scale heights than is usually the case



**Figure 8.13:** X-ray dip occurrence profile as a function of (presumed orbital) phase. Overplotted is a sample prediction of the BB84 model, in this case  $M_2=2$ ,  $e=0.4$ ,  $H_p=7.5 \times 10^4$  km.

( $\sim 10^4$  km), though still small compared to the stellar radius. The reader is urged to bear in mind that the low number of dips recorded per bin by RXTE/ASM leads to significant uncertainty in the true dip occurrence profile (figure 8.13). Monte Carlo simulations suggest that the values of  $M_2(e)$  calculated from fitting these dips are thus uncertain to a  $3\sigma$  error of at least 50%. Above  $T_4 \sim 5$  the factor  $\beta$  becomes largely insensitive to the donor mass  $M_2$ , and *all* combinations occur at  $e \gtrsim 0.6$ . We note further that the values of  $M_2(e)$  should be considered lower limits, as the true  $\dot{M}$  profile may be narrower than the X-ray dip occurrence profile, which would shift the curves upwards in  $M_2$ . Indeed, the  $\gtrsim 5h$  persistence timescale for the stream-impact “clumps” suggests the dip profile should be asymmetric, with more dips occurring after the peak than before it (see figure 8.13). We thus make no attempt to determine the value of  $M_2$  from our simulations, opting instead to determine regions of  $e$  for which solutions exist for reasonable values of  $T_4$ .

As expected, for the same width of  $\dot{M}$  profile, larger eccentricities correspond to lower values of  $M_2$  and thus lower donor radius. Perhaps surprisingly, low eccentricities are not in principle ruled out, although they do require both low surface temperatures and large donor masses. However, the surface of the donor is likely to be significantly heated by any mechanisms discussed above (so  $T_4 \gtrsim 5$ ). Thus, within the limitations of the model (see above), the system is likely to be at high



**Figure 8.14:** Combinations of orbital eccentricity and donor mass that fit the observed dip occurrence rate profile. From left:  $T_4=0.1, 0.5, 1.0, 5.0, 10.0, 50.0$  and  $100.0$  respectively.

eccentricity ( $0.6 \lesssim e \lesssim 0.95$ ) for all  $M_2$  (figure 8.13). More accurate limits await the results of more detailed simulations of  $\dot{M}(t)$  for highly eccentric x-ray binaries.

### 8.6.7 Is Cir X-1 a very young system?

The results from the above toy model suggest the system is likely in a state of high eccentricity, leading to the suggestion that Cir X-1 may be in an early stage of post-supernova evolution. We urge the reader to bear in mind the limitations of the above model, however; in particular the assumption of corotation will likely not hold for a young system. In particular the balance between the competing effects of non-synchronous rotation will likely depend on orbital phase. Without a robust exploration of the effect of non-synchronous donor rotation on the accretion profile it is impossible to place limits on the discrepancy between our simulations and the true accretion profile.

The ephemeris based on the X-ray dips is a stronger indicator of the age of Cir X-1. The X-ray dip ephemeris establishes a characteristic timescale  $P/2\dot{P}$  even shorter than that suggested by the SP03 analysis, at  $\sim 1000$  years. This is an extremely short timescale for such a wide XRB, which provides further circumstantial evidence to support the suggestion that the system may still be in

a state of tidal evolution after the supernova. We examine this claim briefly here. The evolution of the orbital angular momentum of the binary

$$J_{orb} = M_1 M_2 \left( \frac{G a (1 - e^2)}{M_1 + M_2} \right)^{1/2} \quad (8.19)$$

can be straightforwardly expressed as

$$\frac{\dot{a}}{2a} = \frac{\dot{J}_{orb}}{J_{orb}} + \frac{e \dot{e}}{(1 - e^2)} - \left( \frac{\dot{M}_1}{M_1} + \frac{\dot{M}_2}{M_2} + \frac{1}{2} \frac{(\dot{M}_1 + \dot{M}_2)}{M_1 + M_2} \right) \quad (8.20)$$

As outflow is observed from this system, we follow Verbunt (1993) in parameterising the fraction kept as  $\dot{M}_1 = -\beta M_2$  and give the outflow a fraction  $\alpha$  of the specific angular momentum of the donor. Angular momentum evolution then simplifies to

$$\frac{\dot{P}}{3P} = \frac{\dot{J}_{orb}}{J_{orb}} + \frac{e \dot{e}}{(1 - e^2)} - \frac{\dot{M}_2}{M_2} \left[ 1 - \frac{\beta M_2}{M_1} - \alpha(1 - \beta) \frac{M_2}{M_1 + M_2} \right] \quad (8.21)$$

(with  $\dot{M}_2 < 0$ ). The left hand side of this relation is of order  $\sim 10^{-3} - 10^{-4} \text{yr}^{-1}$ , depending if one identifies the quadratic correction in the X-ray or radio ephemerides with orbital period change. Angular momentum loss through accretion is insignificant by comparison: even for  $\dot{M}_2 \simeq -10^{-6} M_\odot \text{yr}^{-1}$ , without giving the matter lost an extremely high angular momentum ( $\alpha \sim 50$ ) the last term on the right hand side of (8.21) is of order  $\sim 10^{-7} \text{yr}^{-1}$ . The rapid period evolution thus occurs as a reaction to a torque changing  $J_{orb}$ , with the eccentricity evolution then describing the inspiral of the binary. The long orbital period makes angular momentum loss through gravitational radiation and/or magnetic braking likely to be insufficient to produce the observed period derivative (see e.g. Zeldovich & Novikov 1999; for the magnetic braking we take the model of Verbunt & Zwaan 1981; such a model is considered fast by the standards of magnetic braking, but produces inspiral many orders of magnitude too slow for Cir X-1). Turning to tidal interactions, it has long been suggested that the approach of an initially eccentric system to circularisation and corotation leads to orbital period change in the system; however the physical timescale corresponding to such change depends on the state of the donor interior (e.g. Mazeh & Shaham 1979). However, the evolution to minimum energy is not necessarily monotonic; by considering the evolution of normal modes excited in the donor as a function of the eccentric orbit, it has been shown that for a variety of circumstances, the binary evolution may be highly complex (Mardling, 1995a,b). Resonant modes can be excited leading to quasi-stable behaviour at high eccentricity; indeed the evolution can become chaotic, with high-amplitude variations in the orbital eccentricity and period. In the simulations of Mardling (1995a,b) the chaos boundary is reached once the periastron separation shrinks below  $\sim 3 \times$  the donor radius; as the donor radius is effectively a free parameter this situation cannot be ruled out. Identification of the evolutionary state of Cir X-1 from its orbit

alone is probably not possible without a direct indication of the donor mass and radius; however indirect evidence may exist for the age of the system (see section 8.7).

If the system were in a state of tidal evolution shortly after the supernova, the system components may not yet have achieved corotation and alignment of their spin axes. The evolution to alignment is expected to give rise to high amplitude variations in mass transfer on a timescale of decades or more (Podsiadlowski 2003, priv. comm.). A look at the long-term X-ray history of Cir X-1 (e.g. SP03) shows variability by a factor of at least 10 on such a timescale. This may be further evidence that the system is evolving dynamically, though without ruling out other explanations for such long-term apparent variation, no definite claim can be made.

However, Cir X-1 is surrounded by a synchrotron radio nebula approximately  $5 \times 10$  arcmin $^2$  (Stewart et al 1993). Now that the association with SNR G321.9-0.3 has been ruled out (Mignani et al 2002), we propose that this radio nebula may in fact be the supernova remnant of Cir X-1. (When completing the paper on which this chapter is based, we became aware (Fender 2003, priv. comm.) of others exploring this possibility, but believe it has not yet been discussed in the literature.) Such a conclusion is consistent with standard SNR expansion models, and the high reddening would obscure optical emission from the nebula. With a diameter of  $\sim 8$  pc, the remnant would likely be in the Sedov-Taylor stage, which implies  $\gtrsim 4,000$  years have elapsed since the supernova (e.g. Choudhuri 1998). Comparison with observed SNRs at this stage of evolution (Hughes et al., 1998) suggests the X-ray luminosity is likely to be  $L_X \sim 1 - 50 \times 10^{35}$  erg s $^{-1}$ , or roughly 1% the “steady” output of the point source.

### 8.6.8 The External Magnetic Field of the Neutron Star

The primary objection to the young source scenario has always been that the external magnetic field of the neutron star is observed to be much weaker than expected for a young system, as suggested by the existence of Type I bursts (Tennant et al 1987) and Z-source QPOs (Shirey et al 1999), and the non-detection of X-ray pulsations (though the latter is inclination-dependent). We suggest the discrepancy may be resolved by accretion-induced screening of the core magnetic field. That low-field radio pulsars exist nearly exclusively in binaries suggests rapid external field evolution occurs as a result of the accretion process (see Battacharya & Srinivasan 1995 for a review). Detailed consideration of the accretion and screening process is beyond the scope of this work: we merely summarise the model and recent results. Matter initially channelled onto

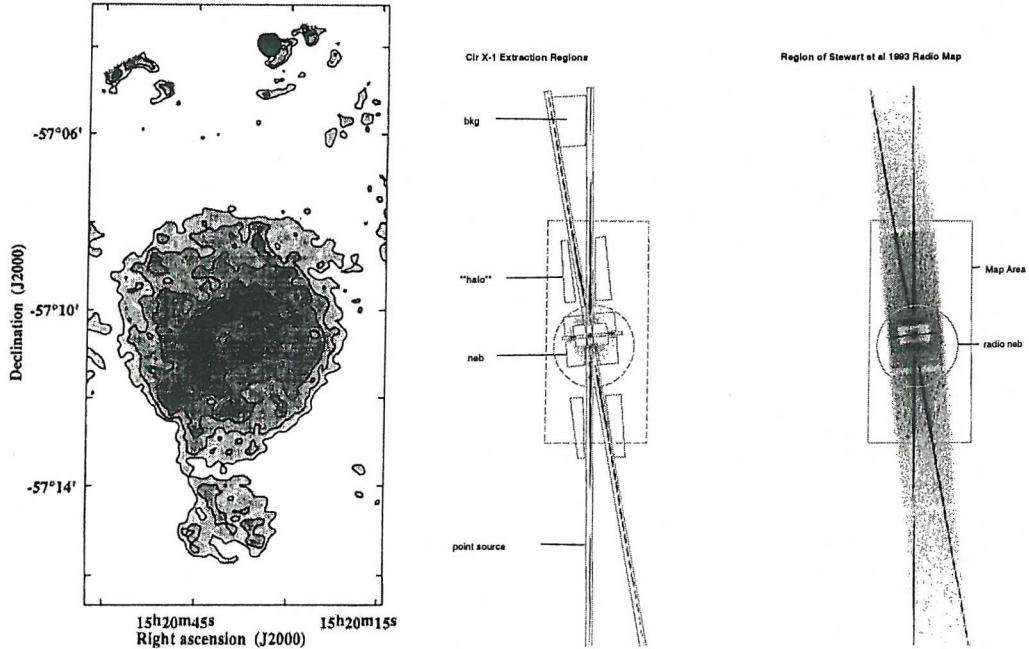
the magnetic poles spreads under the weight of the accretion column to cover the surface of the neutron star. This matter is a fully ionised plasma, and therefore is diamagnetic, leading to the screening of the internal magnetic field (Taam & van de Heuvel, 1986). Advection of the screening layer then submerges the internal magnetic field (Romani, 1990). Evolution of the external field then depends on the competition between the timescales for advection of the accreted matter and for ohmic diffusion of the internal field into the newly accreted matter (Konar & Bhattacharya, 1997). In the steady state, for  $\dot{M} \lesssim 0.1 \dot{M}_{Edd}$ , screening can occur only if the buried field is weaker than  $10^{10} - 10^{11}$  G (Cumming, Zwiebel & Bildsten 2001). However, if previously screened at higher  $\dot{M}$ , during periods of low  $\dot{M}$  the internal field will break through the screening on the timescale for ohmic diffusion across the outer crust (Cumming et al., 2001), which will be measured in decades for all sensible values of the crustal scale height. Thus one need only consider the intervals of high  $\dot{M}$  when considering the screening process. Using the RXTE/ASM count rate and source distance to estimate the mass accretion rate of Cir X-1 from the luminosity, we find instantaneous accretion rates approach  $\sim 3\dot{M}_{Edd}$  in the early stages of the RXTE/ASM lightcurve, while the “steady” accretion rate appears to reach  $\sim 0.7\dot{M}_{Edd}$  for periods of years. (We remind the reader that the type I bursts were observed during an interval when Cir X-1 was accreting at a fraction of its current rate; Saz Parkinson et al 2003, Tennant 1987). Brown & Bildsten (1998) compute the competition between advection and ohmic diffusion for polar accretion, and find that for  $\sim 0.5\dot{M}_{Edd}$  the timescale for ohmic diffusion is larger than for advection, suggesting screening may be possible on the advection timescale ( $\sim 10^3 - 10^5$  yr). Recent 2D simulations (Choudhuri & Konar, 2002) suggest that if magnetic buoyancy is neglected the interior field may be screened by 3 or 4 orders of magnitude on a timescale of  $\sim 1$  yr for a modestly stiff equation of state, however inclusion of magnetic buoyancy leads to a longer timescale of screening of  $\sim 10^5$  yr. Further work is needed to arrive at robust timescales for screening: in particular the evolution of instabilities breaking the screening (c.f. Cumming, Zwiebel & Bildsten 2001) and the extension of the model to three dimensions have yet to be explored. Without information about the likely mass transfer rate of Cir X-1 immediately after the supernova, it is difficult to judge the likely effectiveness of screening, as no information is thus available as to the  $\dot{M}$  history. We simply suggest here that if Cir X-1 has undergone many intervals of mass transfer as high as observed by RXTE/ASM, screening of the internal magnetic field by several orders of magnitude appears at present a possibility, on a timescale ranging perhaps from  $10^3 - 10^5$  years. Furthermore, we remind the reader that during enhanced accretion at periastron passage, the mass transfer onto the accretor may be spherical even with strong internal B-fields. We speculate that this may reduce the timescale necessary for screening.

## 8.7 Using Chandra to search for an associated SNR

Having raised the possibility Cir X-1 might still be in a stage of early post-SN inspiral, we now examine archival high spatial-resolution observations of Cir X-1 to determine if the associated remnant has already been observed. The Chandra observatory offers spatial resolution almost an order of magnitude finer than all other X-ray missions to date. Incoming photons are reflected off the mirror assembly into one of two instruments; the High Resolution Camera (HRC) (Murray et al., 2000) or the Advanced CCD Imaging Spectrometer (ACIS; Garmire et al., 2003). The ACIS camera may be operated in one of two modes; Imaging or Spectroscopy. The Spectroscopic mode employs a row of six CCD detectors, two of which (S1 and S3) are back-illuminated, providing better energy sensitivity than the front-illuminated chips. The High Energy Transmission Grating (HETG) (Canizares et al., 2000) consists of two reflection gratings; the High Energy (HEG; range 0.8-10.0 keV) and Medium Energy (MEG; energy range 0.4-5.0 keV) gratings are mounted with their rulings rotated 9.9° from each other, so that the 1st order diffraction patterns are spatially distinct and may be used concurrently. The zeroth-order image is also let through the grating, forming a direct image centered about the aimpoint. The presence of the grating reduces the sensitivity of the zero-order image by about a factor  $\sim 3 - 10$  (based on calculations using HEASARC W3PIMMS). The NASA/HEASARC facility has archived a high fraction of all X-ray observations taken; data from Guest Observer runs are made publicly available after an embargo period, from the HEASARC W3Browse web interface. All Chandra observations described in this chapter were obtained through this method.

Cir X-1 has been observed five times with the Chandra observatory, using the ACIS-S/HETG instrumental combination. On the first occasion, ObsID 706, the spacecraft appears to have been slightly misaligned, as the point source has missed the spatial filter, resulting in severe pileup, and was taken just after the periastron interval of dipping. ObsID 1700 was taken during an interval of dipping, while observations 1905, 1905 and 1907 were taken throughout the rest of the orbit, at phases  $\sim 0.3, 0.5 \& 0.8$  respectively.

Observation 1700 resulted in the detection of a remarkable set of strong X-ray P Cygni profiles clustered near periastron passage, corresponding to highly ionised states of Ne, Mg, Si, S and Fe (Brandt & Schulz 2000). The breadth of the lines ( $\sim 200 - 1900 \text{ km s}^{-1}$ ) and their time variability suggest identification of these lines with highly ionised outflow from the inner region of the accretion disk (Schulz & Brandt 2002). Results of observations 1905-1907 have not yet been



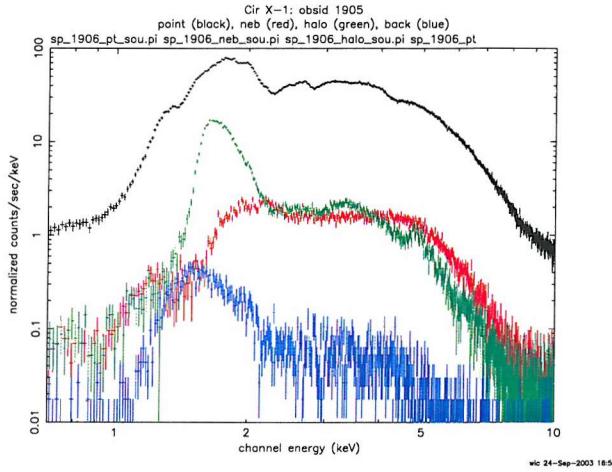
**Figure 8.15:** Regions of extraction for Chandra ACIS-S/HETG spectra of Cir X-1. **Left:** ATCA radio map of the Cir X-1 region, indicating the presence of an extended  $\sim 5 \times 10$  arcmin synchrotron nebula. From Stewart et al (1993). **Middle:** On a larger scale, the extraction regions for the spatially-resolved X-ray spectra of Cir X-1. The regions corresponding to the point source, nebula, halo and background, superimposed on a low-brightness, high-contrast image of an original ACIS-S/HETG observation (in this case 1906). **Right:** The extent of the radio nebula to the same scale, calculated from the co-ordinates given in the radio map of Stewart et al (1993).

published.

As the primary scientific goal of the observations was to follow the evolution of the P Cygni profiles throughout the orbit of Cir X-1, the optimal set-up for spectroscopy was used. However, the ACIS-S/HETG setup does transmit the zero-order image, as well as the dispersed 1st and 2nd orders of the diffraction pattern, meaning information can still be extracted from different *regions* of the source field of view. For all the analysis in what follows, data for which  $(0.3 < E < 10 \text{ keV})$  and using the provided good time intervals (GTI) during which the source was free of background contamination. Furthermore it was decided to remove readout streaks from the S4 chip, as recommended by the Chandra X-ray Center at Harvard (CXC).

### 8.7.1 Spatially-resolved spectra

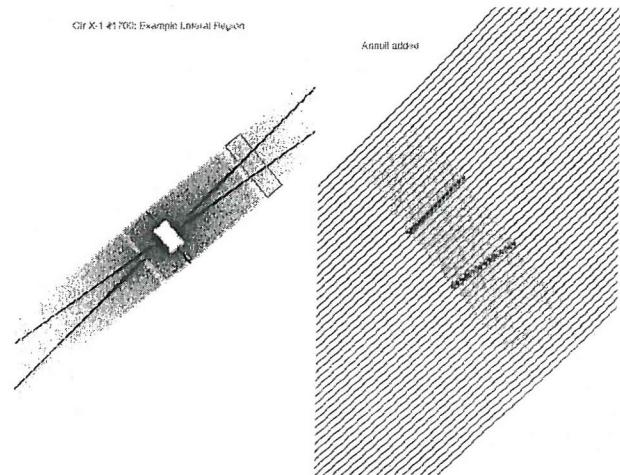
We use v3.0.1. of the *CIAO* analysis suite developed by the CXC for the Chandra mission, to extract spectra of the various regions of the field of view of Cir X-1. These regions were: the point



**Figure 8.16:** Spectra of the four extraction regions, with  $\sim 2''$  separation of the nebula/halo/background regions from the dispersion pattern. The halo region contains significant contamination from the central point source, as can be seen from the similarity in the two spectra below 2 keV.

source itself, the region corresponding to the radio nebula, a background region, and a region between the background and the nebula, which we may expect to be dominated by the X-ray “halo” from ISM scattering. These regions were selected using the Stewart et al (1993) radio map as a guide to the extent of the nebula (figure 8.15). Accurate fitting of the spectra will be deferred to future work, using the physically realistic model of Church et al (e.g. 1997); in particular to get accurate fits, points *must* be selected by spectral state to expect realistic results. The purpose of this investigation is to determine the existence of a component in the X-ray emission that does not vary with time. Under the hypothesis that *all* the X-ray emission from Cir X-1 is due to the point source and halo, spectra taken from all the regions described above should vary over time. By contrast, any extended SNR component should remain constant over the interval between the observations, which would be evident as a reduction in fractional variability.

To remove pileup, the original investigators of these observations used a spatial window to block out the zero-order image of the point source. However, spectra of the point source may still be taken by sampling the “X” - shaped dispersion pattern. Care must be taken to keep the selected nebula/halo/background regions free of contamination from the much brighter dispersed orders of the point source. The users’ guide to the HETG states that the cross-dispersion of the dispersed orders fits entirely within  $\sim 2''$  (i.e. 4 pixels) of the orders; an initial reduction taking the nebula region separate from the X-pattern by approximately this figure produced a halo spectrum with significant contamination, with a high degree of similarity to the point source below 2 keV (figure 8.16).

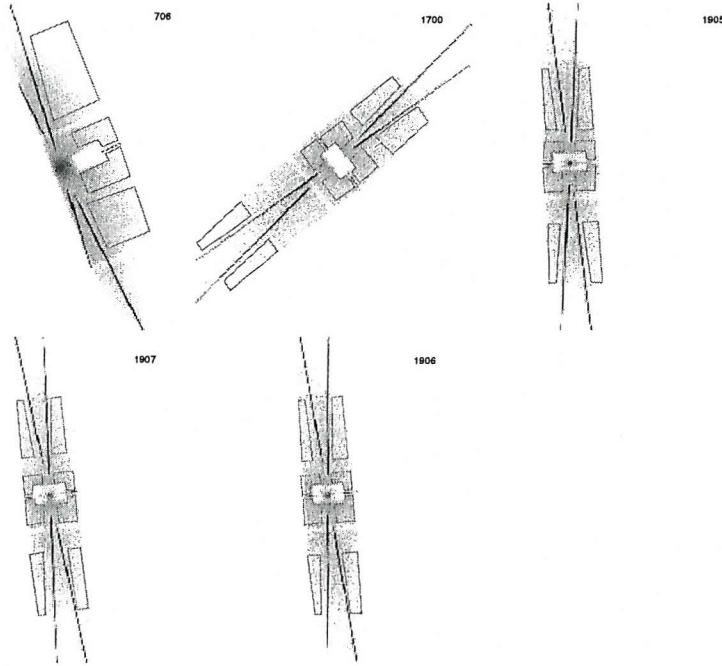


**Figure 8.17:** Illustration of the method used to obtain rectangular regions for each lateral profile. Left: a region of interest is first selected from the overall grating image. Right: Annuli are constructed with large radii and at large displacement from the point source. The result is nearly rectangular regions, parallel to one of the grating dispersion patterns. The annuli show appreciable curvature over the second arm of the “X-pattern,” so this arm is not used for lateral profile analysis.

As a result, a brief investigation into the extent of lateral contamination was undertaken. This was achieved by taking transverse cuts through the “X-pattern” of the dispersion orders and estimating the lateral distance at which the contribution from the diffracted point source is not significant. Whilst generally an extremely powerful, modular piece of software, the *CIAO* tool *dmextract* used to create radial profiles only accepts annular regions for extraction. For this first pass, profiles were generated using very large annular regions displaced from the central image so that the large thin annuli approximate rectangular extraction regions to  $\lesssim 1\%$ . This leads us to take the more conservative value of  $\sim 5''$  (10 pixels). Regions inside  $5''$  from the diffraction pattern were thus subtracted before proceeding further (c.f. figure 8.17).

The ACIS-S/HETG images of the five observations, with the extraction regions overlaid, are plotted in figure 8.18. To standardise the analysis and allow comparison, the same extraction regions were used for observations 1700, 1905, 1906 and 1907, translated and rotated as appropriate, ensuring the same amount of extraction area for these four observations. The offset of observation 706 forces the use of different extraction regions, which, combined with the expected radial drop-off of the source brightness, make selection of extraction regions comparable to that used for the other four observations impossible.

The resulting spectra are plotted in figure 8.19, immediately showing several features of interest. As expected, the point source region varies by a factor  $\sim 3$  above 3 keV and a factor  $\sim 10$  below

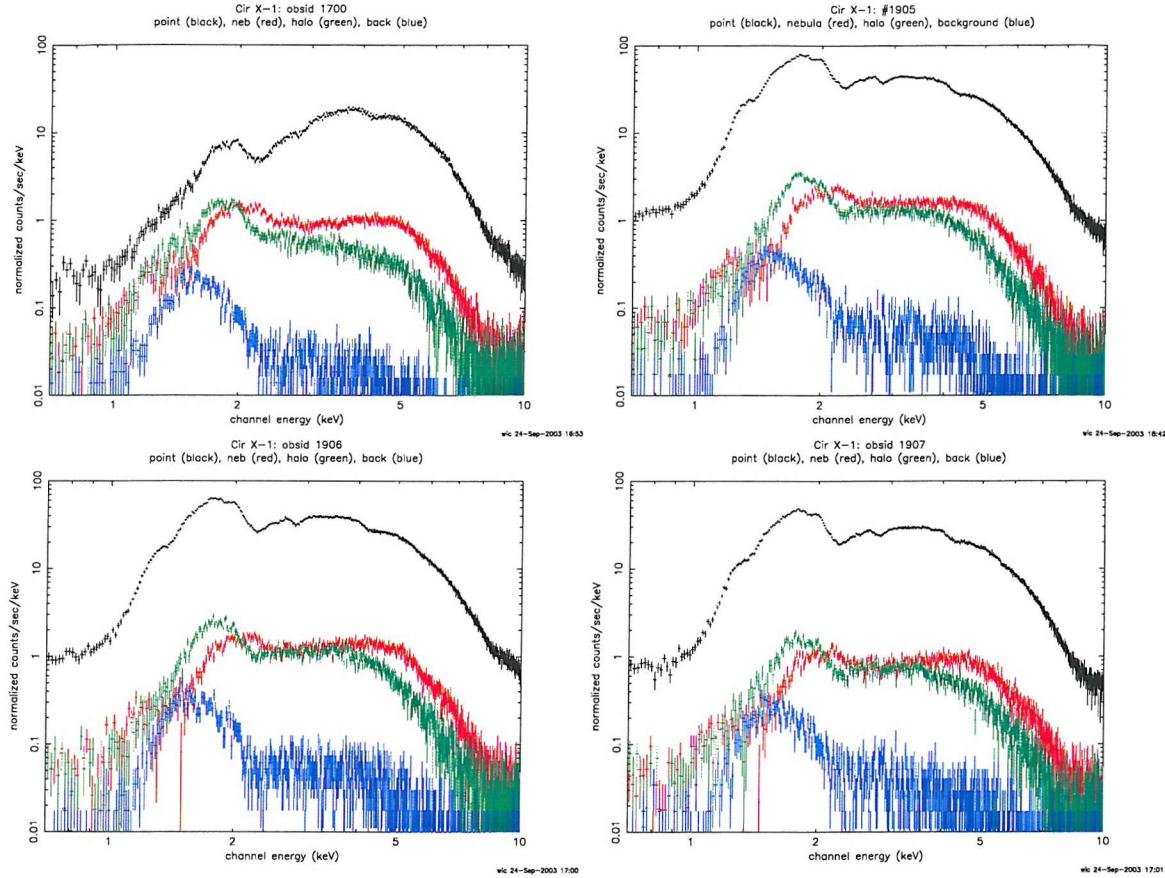


**Figure 8.18:** Regions of extraction for the five observations. The same area was used for observations 1700, 1905, 1906 and 1907, translated and rotated as appropriate, but for observation 706 the same area cannot be used.

3keV. Furthermore, as predicted from the energy dependence of  $\theta_{sc}$  and  $\Delta t$ , the spectral shape of the halo flux varies between the observations, with a factor  $\sim 4$  variation in total flux between the observations. However, the time-variation of the spectra from the regions corresponding to the nebula and halo appear consistent with each other. Thus the spectra taken over the full observation interval shows no immediate indicator of a strong SNR component.

### 8.7.2 Lateral Profiles

We now turn to spatial profiles of the emission from Cir X-1 to ascertain if an extended X-ray structure exists, and if so, if this structure is consistent with or distinct from an X-ray “halo.” As the detailed PSF at large angles of the zero-order image of the ACIS-S/HETG combination is not well known (Proposers’ Observatory Guide, cycle 6), we proceed as with the XMM observations, by empirical comparison with observations of different objects but with the same instrumental setup. For low- $nH$  we choose the IMXB Cyg X-2 (observation ID 1016), for which ROSAT fits suggest  $nH \sim 10\%$  that of Cir X-1 (Predehl & Schmitt, 1995). We also compare with GX 13+1 (observation ID 2708), the archetypal X-ray halo source, with  $nH \sim 1.3$  that of Cir X-1 (Predehl & Schmitt, 1995; Smith et al., 2002b). Also useful as a comparison is the observation of Cir X-1

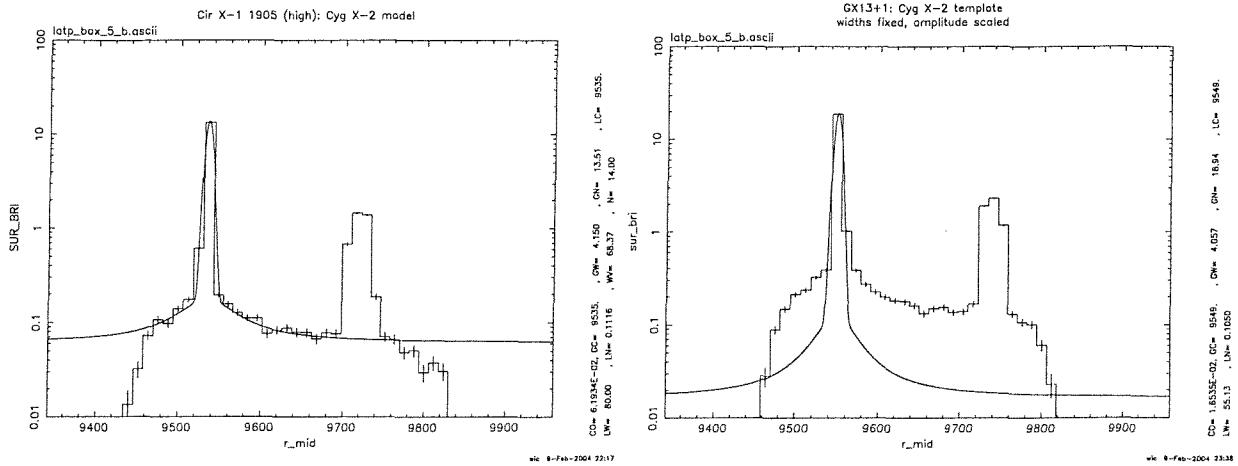


**Figure 8.19:** Point Source (black), Nebula (red), Halo (green) and Background spectra (blue), for observation 1700 (top left), 1905 (top right), 1906 (bottom left) and 1907 (bottom right). All plots are to the same scale.

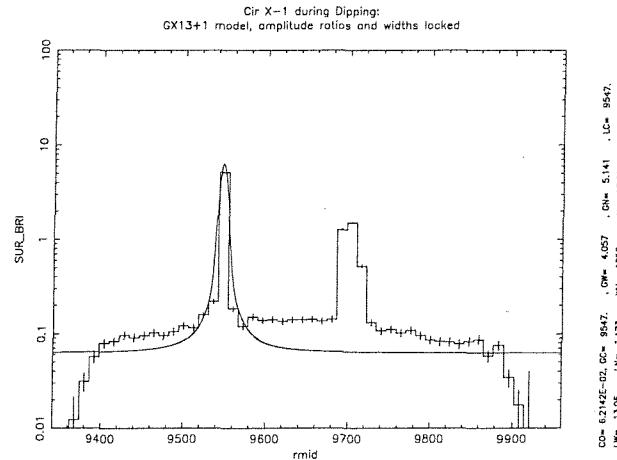
in the high-state (ID 1905), as any halo/SNR will be less significant to the diffraction pattern and point source in this case. We extract lateral profiles across the X-pattern at a distance of 4 arcmin from the point source, i.e. within the radio nebula observed by Stewart et al. (1993).

Using the large-annulus approximation described above figure 8.17 to extract lateral brightness profiles, an empirical fit was obtained to the 1-st order HEG pattern of the low- $nH$  source Cyg X-2. The exact form of the profile is a narrow gaussian profile over a broad lorentzian and a constant level. Keeping the component widths and the ratio of their normalisations constant then allows the profile to be applied to the other observations, normalising for the total count rate in each case. When the low- $nH$  (from here on “instrumental”) profile is applied to Cir X-1 in the bright state, we see that it reproduces the observed lateral brightness drop-off quite faithfully (figures 8.20 & 8.21).

However, when the instrumental profile is applied to the high- $nH$  source GX13+1, we see there is



**Figure 8.20:** Lateral profiles as a function of  $nH$ . **Left:** Cir X-1 during outburst, over the best-fit profile for the low- $nH$  source Cyg X-2. As expected, the profile is consistent with a single point source. **Right:** The same profile applied to the high- $nH$  source GX 13+1. The dispersion is clearly more broad than for the simple point-source, suggesting either increased background due to halo scattering or variation in dispersion optical path due to smearing of the point source.



**Figure 8.21:** The best-fit profile from GX13+1, scaled to the count rate of Cir X-1 during dipping. There is marginal evidence for an additional broad component.

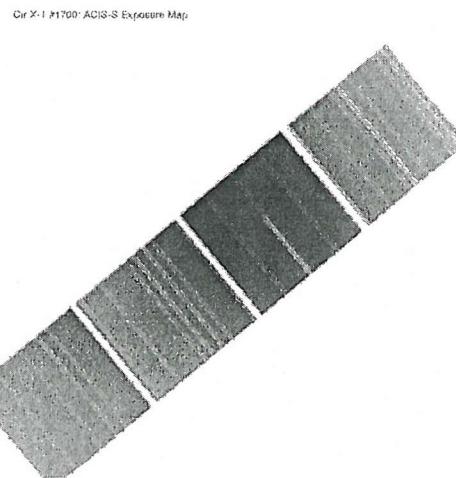
clearly a broad component present that cannot be accounted for by the fit to low- $nH$  data alone. This can be explained from two factors; (i) a contribution due to emission scattered from the halo on top of the dispersion pattern, and (ii) a broadening of the dispersion pattern itself due to effective smearing of the point source due to this halo. As expected, the same result persists for Cir X-1 itself during dipping. When the model is re-fit to GX13+1 data, we can then apply it to Cir X-1 during dipping to investigate similarities between the lateral profiles of Cir X-1 during dipping to halo-dominated diffuse emission. As before the ratio of the amplitudes of the two components are kept constant, but this time the widths of the lorentzian and gaussian are fixed as per the GX 13+1 profile. We also keep constant the ratio of the linelike components to the constant offset.

As can be seen, the resulting profile does not quite fit the data; in particular the constant offset is incorrect by  $\sim 50\%$ . In order to get the halo-dominated lateral profile to fit the dispersion of Cir X-1 during dipping, then, the constant offset must be adjusted. This implies that the lateral drop-off from the dispersion pattern, assuming *only* a halo component as per GX13+1, is not sufficient to reproduce the breadth of the lateral dispersion from Cir X-1 during dipping. However, such an interpretation is somewhat tentative, as the details of the lateral dispersion will depend on the distribution of scattering grains, which may differ substantially between the two sources. We thus must turn to radial profiles for a further test.

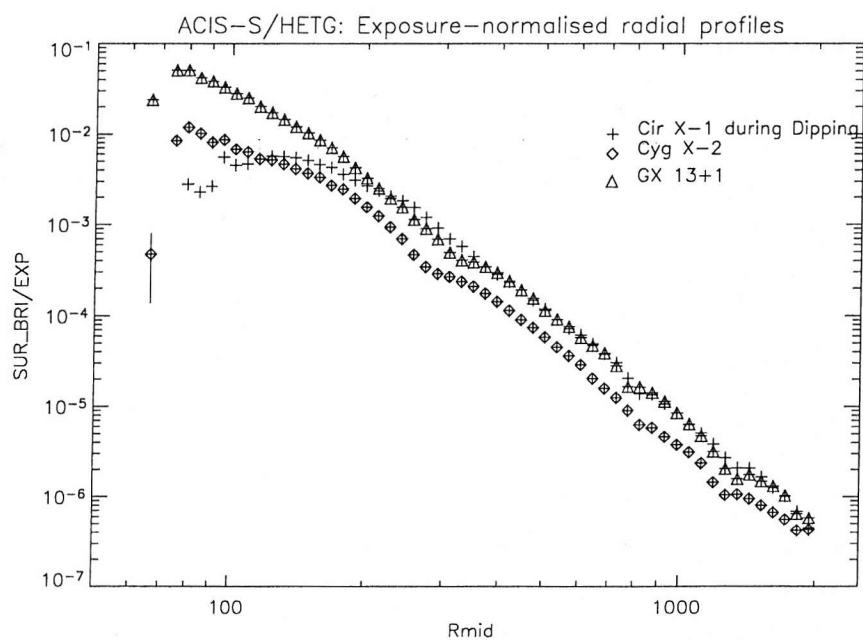
### 8.7.3 Radial Profiles

Before extracting radial profiles it is important to correct for variations in exposure across the chipset, in particular the chip gaps for which no data was taken. This allows confidence that variations along the radial profile are due to behaviour on the sky and not variations in the chipset. However, following the appropriate thread in the *CIAO* documentation introduces artifacts to the resulting radial profiles as well as producing somewhat unrealistic error figures; indeed the CXC urge caution when applying the exposure map to grating data. Instead, the raw radial profiles were used and divided manually by radial profiles over the exposure map images (c.f. figure 8.22). This procedure is not entirely free of artifacts, as can be seen from the radial profiles themselves, in which a slight deviation due to the chip gaps can be seen, but it provides an improvement over the reduction thread as is. As the PSF nominally follows a power-law (Proposers' Observatory Guide, Cycle 6), we choose log-spaced annuli to accumulate a similar number of counts in each bin and thus standardise errors somewhat. As discussed above, before proceeding computing radial profiles we remove all points within 5" of the X-pattern itself, to reduce the possibility of contamination from the point source.

The results are plotted in figure 8.23. The dramatic apparent decrease in count rate towards the centre of the Cir X-1 radial profile is due to the spatial filter applied to block the point source. Surprisingly, beyond  $\sim 2$  arcminutes from the central source ( $R_{mid} = 240$ ), the exposure-corrected radial profiles of Cir X-1, GX 13+1 and Cyg X-2, appear largely indistinguishable; except for a constant offset, all three appear to follow the same power law. Certainly within the  $\sim 5$  arcmin extent of the radio nebula observed by Stewart et al (1993) there is very little difference between Cir X-1 and GX13+1.



**Figure 8.22:** Section from the exposure map image of Cir X-1 during dipping. Clear variations both within each chip and across the chipset can be seen.



**Figure 8.23:** Radial profiles for the low- $nH$  source Cyg X-2, the archetypal X-ray halo source GX 13+1, and Cir X-1 during intervals of dipping. The radial profiles for GX 13+1 and Cir X-1 are entirely consistent with each other.

Robust comparison between the radial profiles, especially with Cyg X-2, requires normalisation by the central source flux. However, the spatial filter blocking the point source in the case of Cir X-1 makes such normalisation somewhat challenging. An attempt was made to arrive at the normalisation between the radial profiles by estimating the count rate from the readout streak. (The ACIS instrument takes data while the chip is being read out, causing photons to be clocked into the wrong row during the 3.2s frame time. For bright sources, this results in a visible streak across the position of the point source.) The ACIS-I readout streak was used to extract the point-source spectrum in the case of GX 13+1 (Smith et al., 2002b), so it was supposed that a similar method would allow determination of the central source count-rate, independently of the spatial filter (or lack thereof) at the centre of the image. Unfortunately, however, such a method does not produce results due to the different pile-up fraction between the sources. Scaling by the count rate in the X-pattern instead may be a more robust method to estimate central source flux (this will suffer less from pileup and therefore be more uniform across the sources of interest); however without selecting GTI's from each source based on spectral state, estimation of the central source flux will be unreliable due to poor spectral fits. Rigorous comparison is thus left to future work.

## 8.8 Conclusion

We have examined the periodic nature of the X-ray maxima and dips exhibited in the RXTE/ASM lightcurve of Circinus X-1. We find that the X-ray dips provide a more accurate system clock than the maxima, and thus identify them with the periastron approach of the binary components. The outbursts are interpreted as the viscous timescale response of the disk to perturbations from the varying mass transfer. We attempt to use the X-ray dips to constrain the nature of the donor and the eccentricity of the orbit, and find that, within the limitations of the model used, a wider range of parameters are in principle allowed than has previously been thought, although without further exploration of the effects of nonsynchronous donor rotation, quantitative limits are at this stage unreliable. The high rate of change of the orbital period adds to the growing body of evidence that Cir X-1 is in an extremely early stage of post-supernova evolution. If the system is indeed in an early stage of post-SN evolution, the observed weak external magnetic field must suggest strong screening of the internal magnetic field of the neutron star.

Observations to date do not yet place constraints on the brightness of any X-ray counterpart to the radio nebula surrounding Cir X-1. Chandra ACIS-S/HETG observations appear fully consistent

with diffuse emission from the ISM, but the high column density towards Cir X-1 predicts this anyway. Whether there is extended X-ray emission from a SNR has not been established. Fresh Chandra observations taken during dipping, without the grating in place, are sorely needed as under these conditions the sensitivity of ACIS-S will be dramatically enhanced. For now, no direct constraint on the age of Cir X-1 can be made from the current spatially resolved X-ray emission.

## 8.9 Update: Radio Ejections from Cir X-1

Since the majority of the work in this chapter was completed, new radio observations of Cir X-1 with ATCA appeared in print (after some two years' preparation). These observations show moving ejecta along a bipolar trajectory, with an apparent space velocity of  $\beta \gtrsim 15$ , indicating relativistic  $\Gamma_{bulk} \gtrsim 15$ . This is by far the highest space velocity of relativistic outflow from *any* stellar-mass jet system, and shows that the event horizon is *not* essential to the formation of such outflows (Fender et al., 2004). The  $\sim 5 \times 10$  arcminute radio nebula imaged by Stewart et al (1993) was shown to be concentrated into a smaller  $\sim 5 \times 3$  arcmin, which entirely fits onto a single ACIS-S chip. The X-ray counterpart to such a nebula is thus even harder to detect from archival Chandra images, as during dipping the spatial filter removes most of the flux from this region, and during observation 706 the point source is bright enough to make such a nebula undetectable. A proposal has thus been sent for Cycle 6 to use the high sensitivity of ACIS-S, with no grating, to image the circumbinary region during dipping and thus determine definitively if an X-ray counterpart to the radio nebula exists.

# Chapter 9

## Conclusions

This thesis illustrates the value of longterm X-ray lightcurves to the study of the X-ray Binaries. Chapters 4-6 and 8 have used them to constrain the properties of M15, SMC X-1, Her X-1, LMC X-4, Cyg X-2 and Cir X-1 in ways not otherwise possible, and, where applicable, the conclusions drawn have been shown to be consistent with subsequent work by other authors (see the update to Chapter 4 for an example). Chapters 5 & 6 have suggested radiation-driven warping is the most plausible mechanism to produce the longterm X-ray periodicities due to disk-warping in the bright X-ray binaries as a class, and leaves the door open for fruitful future work (see update, chapters 5 & 6), which in both cases has already started in collaboration with other groups. The work on Cir X-1 (chapters 7-9) is to some extent incomplete, as further observations have been either scheduled (AAT/IRIS2) or proposed with a high chance of success (Chandra/ACIS-S). These observations promise to place the best constraints yet on the nature of this enigmatic system. We note that no matter which interpretation turns out to be correct for Cir X-1 (i.e. young/massive or old/low-mass), its properties are clearly contradictory in a way that will challenge XRB models (see Chapter 8).

### 9.1 Future Work

The nature of long-term lightcurves is such that, as the observation interval continuously increases, the possibility is always present that the discovery that challenges the current model is always just

around the corner. This has already happened at several junctures, for example the cases of M15 (chapter 4), SMC X-1 (chapter 5) and Her X-1 (chapter 6). For each of the scientific chapters, I have suggested the obvious next step in the work in the updates at the end of each chapter. Beyond this, chapter 2 has shown there are many systems which need better characterisation to establish convincingly if their long variation is indeed due to disk precession; there is still debate, for example, over whether the long periodicity in Cyg X-2 really is due to a precessing disk-warp or some other characteristic of the mass transfer process. Of particular interest is the new predictions starting to be made of spectral evolution of sources as a function of precessional cycle when the disk is warped. There is also the issue of the ultracompacks, as the most modern theory does not predict they should warp, yet observations suggest at least one of them doesn't understand this, and does so anyway. Work on the spin period history of SMC X-1 has started, and we expect to be able to say for sure whether the warp relates to the accretion torque within a few months from now (june 2004). Turning to Cir X-1, high time-resolution RXTE/PCA time has recently been allocated to coincide with the apastron-periastron IR monitoring suggested in the update to Chapter 7. We are attempting to secure radio and INTEGRAL time to perform a truly multi-wavelength study and provide the best constraints yet on the system. Cir X-1 probably has more surprises in wait, and the Rosetta-stone nature of this system suggests such constraints may be important to our understanding of XRB behaviours.

## 9.2 Future All-Sky Monitors

As I am confident this thesis illustrates, detailed study of long-term lightcurves at high energies is extremely fruitful and indeed represents the *only* way to constrain certain properties of the X-ray Binaries. With the end of RXTE imminent (at least one time allocation committee has reduced the amount of co-ordinated time with XTE to a multiwavelength project to six months on the assumption the satellite may not last the year), a replacement X-ray monitor is sorely needed.

# Bibliography

Alcock, C., Allsman, R. A., Alves, D. R., et al. 2000, MNRAS, 316, 729

Alcock, C. & Hatchett, S. 1978, ApJ, 222, 456

Allen, C. W. 1973, *Astrophysical quantities* (London: University of London, Athlone Press, —c1973, 3rd ed.)

Alpar, M. A., Hasinger, G., Shaham, J., & Yancopoulos, S. 1992, A&A, 257, 627

Apparao, K. M. V. 2001, A&A, 371, 672

Bahcall, J. N. & Bahcall, N. A. 1972, ApJl, 178, L1+

Balbus, S. A. & Hawley, J. F. 1991, ApJ, 376, 214

Balbus, S. A., Hawley, J. F., & Stone, J. M. 1996, ApJ, 467, 76

Bałucińska-Church, M., Church, M. J., Charles, P. A., et al. 2000, MNRAS, 311, 861

Bandyopadhyay, R., Shahbaz, T., Charles, P. A., van Kerkwijk, M. H., & Naylor, T. 1997, MNRAS, 285, 718

Bandyopadhyay, R. M., Shahbaz, T., Charles, P. A., & Naylor, T. 1999, MNRAS, 306, 417

Bardeen, J. M. & Petterson, J. A. 1975, ApJl, 195, L65+

Barnard, R., Church, M. J., & Bałucińska-Church, M. 2003a, A&A, 405, 237

Barnard, R., Osborne, J. P., Kolb, U., & Borozdin, K. N. 2003b, A&A, 405, 505

Barret, D. 2001, Advances in Space Research, 28, 307

Barret, D., Motch, C., & Predehl, P. 1998, A&A, 329, 965

Basinska, E. M., Lewin, W. H. G., Sztajno, M., Cominsky, L. R., & Marshall, F. J. 1984, *ApJ*, 281, 337

Baykal, A. 2000, *MNRAS*, 313, 637

Bildsten, L., Chakrabarty, D., Chiu, J., et al. 1997, *ApJs*, 113, 367

Billington, I., Marsh, T. R., Horne, K., et al. 1996, *MNRAS*, 279, 1274

Binney, J. & Tremaine, S. 1987, *Galactic dynamics* (Princeton, NJ, Princeton University Press, 1987, 747 p.)

Blondin, J. M. & Woo, J. W. 1995, *ApJ*, 445, 889

Bode, M. F., Priedhorsky, W. C., Norwell, G. A., & Evans, A. 1985, *ApJ*, 299, 845

Boerner, G. & Meszaros, P. 1979, *A&A*, 77, 178

Bonnet-Bidaud, J. M., Motch, C., Beuermann, K., et al. 1989, *A&A*, 213, 97

Boroson, B., O'Brien, K., Horne, K., et al. 2000, *ApJ*, 545, 399

Boyle, C. B. & Walker, I. W. 1986, *MNRAS*, 222, 559

Brandt, W. N., Fabian, A. C., Dotani, T., et al. 1996, *MNRAS*, 283, 1071

Brandt, W. N. & Schulz, N. S. 2000, *ApJl*, 544, L123

Brocksopp, C., Groot, P. J., & Wilms, J. 2001, *MNRAS*, 328, 139

Brown, E. F. & Bildsten, L. 1998, *ApJ*, 496, 915

Brown, J. C. & Boyle, C. B. 1984, *A&A*, 141, 369

Buat-Ménard, V., Hameury, J.-M., & Lasota, J.-P. 2001, *A&A*, 366, 612

Butt, Y. M., Maccarone, T. J., & Prantzos, N. 2003, *ApJ*, 587, 748

Callanan, P. J., Charles, P. A., van Paradijs, J., et al. 1990, *A&A*, 240, 346

Canizares, C. R., Huenemoerder, D. P., Davis, D. S., et al. 2000, *ApJl*, 539, L41

Casares, J., Charles, P. A., & Kuulkers, E. 1998, *ApJl*, 493, L39+

Casares, J., Steeghs, D., Hynes, R. I., Charles, P. A., & O'Brien, K. 2003, *ApJ*, 590, 1041

Catura, R. C. 1983, *ApJ*, 275, 645

Chakrabarty, D., Bildsten, L., Finger, M. H., et al. 1997, *ApJl*, 481, L101+

Chakrabarty, D., Homer, L., Charles, P. A., & O'Donoghue, D. 2001, *ApJ*, 562, 985

Chakrabarty, D., Wang, Z., Juett, A. M., Lee, J. C., & Roche, P. 2002, *ApJ*, 573, 789

Charles, P. A. 1989, in *ESA SP-296: Two Topics in X-Ray Astronomy*, Volume 1: X Ray Binaries. Volume 2: AGN and the X Ray Background, 129–137

Charles, P. A., Booth, L., Densham, R. H., et al. 1983, *MNRAS*, 202, 657

Charles, P. A., Jones, D. C., & Naylor, T. 1986, *Natur*, 323, 417

Charles, P. A. & Seward, F. D. 1995, *Exploring the X-ray universe* (Cambridge, New York: Cambridge University Press, —c1995, ISBN 0521261821)

Cherepashchuk, A. M. 1982, *Soviet Astronomy Letters*, 8, 336

Chou, Y. & Grindlay, J. E. 2001, *ApJ*, 563, 934

Choudhuri, A. R., ed. 1998, *The physics of fluids and plasmas : an introduction for astrophysicists*

Choudhuri, A. R. & Konar, S. 2002, *MNRAS*, 332, 933

Christian, D. J. & Swank, J. H. 1997, *ApJS*, 109, 177

Church, M. J. & Balucińska-Church, M. 2001, *A&A*, 369, 915

Church, M. J. & Bałucińska-Church, M. 2004, *MNRAS*, 348, 955

Church, M. J., Dotani, T., Balucinska-Church, M., et al. 1997, *ApJ*, 491, 388

Clark, J. S., Miroshnichenko, A. S., Larionov, V. M., et al. 2000, *A&A*, 356, 50

Clark, J. S., Steele, I. A., Fender, R. P., & Coe, M. J. 1999, *A&A*, 348, 888

Clarkson, W. I., Charles, P. A., Coe, M. J., & Laycock, S. 2003a, *MNRAS*, 343, 1213

Clarkson, W. I., Charles, P. A., Coe, M. J., et al. 2003b, *MNRAS*, 339, 447

Clarkson, W. I., Charles, P. A., & Onyett, N. 2004, *MNRAS*, 348, 458

Coburn, W., Heindl, W. A., Rothschild, R. E., et al. 2002, *ApJ*, 580, 394

Corbet, R., Peele, A., & Smith, D. A. 1997, in *International Astronomical Union Circular*, 3–+

Corbet, R. H. D. 2003, *ApJ*, 595, 1086

Corbet, R. H. D., Laycock, S., Coe, M. J., Marshall, F. E., & Markwardt, C. B. 2004, astro-ph/0402053

Cornelisse, R., Heise, J., Kuulkers, E., Verbunt, F., & in't Zand, J. J. M. 2000, A&A, 357, L21

Cornelisse, R., in't Zand, J. J. M., Verbunt, F., et al. 2003, A&A, 405, 1033

Cowley, A. P., Schmidtke, P. C., Ebisawa, K., et al. 1991, ApJ, 381, 526

Crampton, D., Hutchings, J. B., Cowley, A. P., Schmidtke, P. C., & Thompson, I. B. 1990, ApJ, 355, 496

Cumming, A. 2003, ApJ, 595, 1077

Cumming, A., Marcy, G. W., & Butler, R. P. 1999, ApJ, 526, 890

Cumming, A., Zweibel, E., & Bildsten, L. 2001, ApJ, 557, 958

de Marchi, G. & Paresce, F. 1994, ApJ, 422, 597

—. 1995, A&A, 304, 202

—. 1996, ApJ, 467, 658

Deeter, J., Crosa, L., Gerend, D., & Boynton, P. E. 1976, ApJ, 206, 861

Demianski, M. & Ivanov, P. B. 1997, A&A, 324, 829

Di Salvo, T., Méndez, M., van der Klis, M., Ford, E., & Robba, N. R. 2001, ApJ, 546, 1107

Di Stefano, R. & Rappaport, S. 1994, ApJ, 423, 274

Dickey, J. M. & Lockman, F. J. 1990, ARA&A, 28, 215

Dotani, T., Inoue, H., Murakami, T., Nagase, F., & Tanaka, Y. 1990, Natur, 347, 534

Dubus, G., Lasota, J., Hameury, J., & Charles, P. 1999, MNRAS, 303, 139

Eggleton, P. P. 1983, ApJ, 268, 368

Eisberg, R. & Resnick, R. 1985, Quantum Physics of Atoms, Molecules, Solids, Nuclei, and Particles, 2nd edn. (New York, NY: John Wiley and Sons)

Elson, R., Hut, P., & Inagaki, S. 1987, ARA&A, 25, 565

Fabbiano, G. 1989, ARA&A, 27, 87

Fabbiano, G., Zezas, A., King, A. R., et al. 2003, *ApJl*, 584, L5

Fabian, A. C., Pringle, J. E., & Rees, M. J. 1975, *MNRAS*, 172, 15P

Fender, R., Spencer, R., Tzioumis, T., et al. 1998, *ApJl*, 506, L121

Fender, R., Wu, K., Johnston, H., et al. 2004, *nat*, 427, 222

Fender, R. P. 1998, in *ASP Conf. Ser.* 144: *IAU Colloq. 164: Radio Emission from Galactic and Extragalactic Compact Sources*, 343-+

Fender, R. P. 2001, *MNRAS*, 322, 31

Fender, R. P., Gallo, E., & Jonker, P. G. 2003, *MNRAS*, 343, L99

Feng, Y. X. & Cui, W. 2002, *ApJ*, 564, 953

Figer, D. F., McLean, I. S., & Najarro, F. 1997, *ApJ*, 486, 420

Foulkes, S., Haswell, C., & Murray, J. 2004, *MNRAS* in prep.

Frank, J., King, A., & Raine, D. J. 2002, *Accretion Power in Astrophysics: Third Edition (Accretion Power in Astrophysics: Third Edition, by Juhan Frank, Andrew King, and Derek J. Raine. Cambridge University Press, 2002, 398 pp.)*

Frank, J., King, A. R., & Lasota, J. 1992, *ApJl*, 385, L45

Frank, J., King, A. R., & Lasota, J.-P. 1987, *A&A*, 178, 137

Galloway, D. K., Cumming, A., Kuulkers, E., et al. 2004, *ApJ*, 601, 466

Galloway, D. K., Morgan, E. H., Levine, A. M., & Chakrabarty, D. 2003, in *IAU Symposium*

Gardiner, L. T., Sawa, T., & Fujimoto, M. 1994, *MNRAS*, 266, 567

Garmire, G. P., Bautz, M. W., Ford, P. G., Nousek, J. A., & Ricker, G. R. 2003, in *X-Ray and Gamma-Ray Telescopes and Instruments for Astronomy*. Edited by Joachim E. Truemper, Harvey D. Tananbaum. *Proceedings of the SPIE*, Volume 4851, pp. 28-44 (2003)., 28-44

Gerend, D. & Boynton, P. E. 1976, *ApJ*, 209, 562

Gierliński, M. & Done, C. 2002, *MNRAS*, 337, 1373

Gies, D. R., McSwain, M. V., Riddle, R. L., et al. 2002, *ApJ*, 566, 1069

Glass, I. S. 1994, *MNRAS*, 268, 742

Gorenstein, P. 1975, *ApJ*, 198, 95

Goss, W. M. & Mebold, U. 1977, *MNRAS*, 181, 255

Gottwald, M., Haberl, F., Parmar, A. N., & White, N. E. 1986, *ApJ*, 308, 213

Greenhill, L. J., Jiang, D. R., Moran, J. M., et al. 1995, *ApJ*, 440, 619

Gregory, P. C., Xu, H., Backhouse, C. J., & Reid, A. 1989, *ApJ*, 339, 1054

Greiner, J., Cuby, J. G., McCaughrean, M. J., Castro-Tirado, A. J., & Mennickent, R. E. 2001, *A&A*, 373, L37

Groot, P. J., van der Klis, M., van Paradijs, J., & et al. 1996, in *ASSL Vol. 208: IAU Colloq. 158: Cataclysmic Variables and Related Objects*, 367-+

Haberl, F., Filipović, M. D., Pietsch, W., & Kahabka, P. 2000, *A&AS*, 142, 41

Haberl, F. & Pietsch, W. 1999, *A&AS*, 139, 277

Hakala, P. J., Muhli, P., & Dubus, G. 1999, *MNRAS*, 306, 701

Hanson, M. M., Conti, P. S., & Rieke, M. J. 1996, *ApJs*, 107, 281

Hanuschik, R. W., Hummel, W., Dietle, O., & Sutorius, E. 1995, *A&A*, 300, 163

Harlaftis, E. T., Dhillon, V. S., & Castro-Tirado, A. 2001, *A&A*, 369, 210

Harmon, B. A., Wilson-Hodge, C. A., Fishman, G. J., et al. 2002, *APS Meeting Abstracts*, 17044

Hasinger, G. & van der Klis, M. 1989, *A&A*, 225, 79

Haswell, C. A., King, A. R., Murray, J. R., & Charles, P. A. 2001, *MNRAS*, 321, 475

Haswell, C. A. & Shafter, A. W. 1990, *ApJl*, 359, L47

Hatchett, S. P., Begelman, M. C., & Sarazin, C. L. 1981, *ApJ*, 247, 677

Haynes, R. F., Jauncey, D. L., Murdin, P. G., et al. 1978, *MNRAS*, 185, 661

Haynes, R. F., Lerche, I., & Wright, A. E. 1980, *A&A*, 81, 83

Heemskerk, M. H. M. & van Paradijs, J. 1989, *A&A*, 223, 154

Heindl, W. A., Coburn, W., Gruber, D. E., et al. 2001, *ApJl*, 563, L35

Hillier, D. J. & Miller, D. L. 1998, *ApJ*, 496, 407

Hoffman, J. A., Cominsky, L., & Lewin, W. H. G. 1980, *ApJl*, 240, L27

Homer, L. & Charles, P. A. 1998, *New Astronomy*, 3, 435

Homer, L., Charles, P. A., Hakala, P., et al. 2001, *MNRAS*, 322, 827

Horne, J. H. & Baliunas, S. L. 1986, *ApJ*, 302, 757

Hughes, J. P., Hayashi, I., & Koyama, K. 1998, *ApJ*, 505, 732

Hut, P., McMillan, S., Goodman, J., et al. 1992, *PASP*, 104, 981

Hutchings, J. B., Cowley, A. P., Osmer, P. S., & Crampton, D. 1977, *ApJ*, 217, 186

Hutchings, J. B. & Crampton, D. 1981, *PASP*, 93, 486

Hutchings, J. B., Crampton, D., & Cowley, A. P. 1978, *ApJ*, 225, 548

Hynes, R. I., Haswell, C. A., Cui, W., et al. 2003, *MNRAS*, 345, 292

Iaria, R., Di Salvo, T., Burderi, L., & Robba, N. R. 2001, *ApJ*, 561, 321

Ilovaisky, S. A., Auriere, M., Koch-Miramond, L., et al. 1993, *A&A*, 270, 139

in 't Zand, J. J. M., Strohmayer, T. E., & Baykal, A. 1997, *ApJl*, 479, L47+

Ioannou, Z., Naylor, T., Smale, A. P., Charles, P. A., & Mukai, K. 2002, *A&A*, 382, 130

Iping, R. C. & Petterson, J. A. 1990, *A&A*, 239, 221

Johnston, H. M., Fender, R., & Wu, K. 1999, *MNRAS*, 308, 415

Jones, D. I. & Andersson, N. 2002, *MNRAS*, 331, 203

Jonker, P. G., van der Klis, M., Wijnands, R., et al. 2000, *ApJ*, 537, 374

Joss, P. C. & Rappaport, S. A. 1984, *ARA&A*, 22, 537

Jurcsik, J., Clement, C., Geyer, E. H., & Domsa, I. 2001, *AJ*, 121, 951

Kahabka, P. 2002, *A&A*, 388, 100

Kahabka, P. & van den Heuvel, E. P. J. 1997, *ARA&A*, 35, 69

Kahn, S. M. & Grindlay, J. E. 1984, *ApJ*, 281, 826

Kaiser, C. R., Sunyaev, R., & Spruit, H. C. 2000, *A&A*, 356, 975

Kaluzienski, L. J., Holt, S. S., Boldt, E. A., & Serlemitsos, P. J. 1976, *ApJl*, 208, L71

Katz, J. I. 1973, *Nature Physical Science*, 246, 87

King, A. R., Davies, M. B., Ward, M. J., Fabbiano, G., & Elvis, M. 2001, *ApJl*, 552, L109

King, A. R., Pringle, J. E., West, R. G., & Livio, M. 2004, *MNRAS*, 348, 111

Kleinmann, S. G. & Hall, D. N. B. 1986, *ApJS*, 62, 501

Konar, S. & Bhattacharya, D. 1997, *MNRAS*, 284, 311

Kong, A. K. H., Charles, P. A., & Kuulkers, E. 1998, *New Astronomy*, 3, 301

Kong, A. K. H., Garcia, M. R., Primini, F. A., et al. 2002, *ApJ*, 577, 738

Kopal, Z. 1959, *Close binary systems* (The International Astrophysics Series, London: Chapman & Hall, 1959)

Kuiper, L., van Paradijs, J., & van der Klis, M. 1988, *A&A*, 203, 79

Kumar, P., Narayan, R., & Loeb, A. 1995, *ApJ*, 453, 480

Kunze, S., Speith, R., & Hessman, F. V. 2001, *MNRAS*, 322, 499

Kuster, M., Wilms, J., Blum, S., et al. 1999, *Astrophysical Letters Communications*, 38, 161

Kuulkers, E., van der Klis, M., & Vaughan, B. A. 1996, *A&A*, 311, 197

Kuulkers, E., Wijnands, R., Belloni, T., et al. 1998, *ApJ*, 494, 753

La Parola, V., Peres, G., Fabbiano, G., Kim, D. W., & Bocchino, F. 2001, *ApJ*, 556, 47

Lai, D. 1999, *ApJ*, 524, 1030

Lamb, D. Q., Lamb, F. K., Pines, D., & Shaham, J. 1975, *ApJl*, 198, L21+

Larionov, V., Lyuty, V. M., & Zaitseva, G. V. 2001, *A&A*, 378, 837

Leahy, D. A. 2002, *A&A*, 391, 219

Leahy, D. A., Marshall, H., & Scott, D. M. 2000, *ApJ*, 542, 446

Leahy, D. A. & Scott, D. M. 1998, *ApJl*, 503, L63+

Levine, A., Rappaport, S., Deeter, J. E., Boynton, P. E., & Nagase, F. 1993, *ApJ*, 410, 328

Levine, A. M., Bradt, H., Cui, W., et al. 1996, *ApJl*, 469, L33+

Levine, A. M., Bradt, H. V., Enevoldsen, A., et al. 2000, Bulletin of the American Astronomical Society, 32, 1260

Lewin, W. H. G., van Paradijs, J., & Taam, R. 1995, X-ray Bursts (X-ray Binaries, Cambridge Astrophysics Series, Cambridge, MA: Cambridge University Press, —c1995, edited by Lewin, Walter H.G.; Van Paradijs, Jan; Van den Heuvel, Edward P.J., p. 175)–+

Liu, Q. Z., van Paradijs, J., & van den Heuvel, E. P. J. 2000, A&AS, 147, 25

—. 2001, A&A, 368, 1021

Lochner, J. C. & Roussel-Dupre, D. 1994, ApJ, 435, 840

Lomb, N. R. 1976, Ap&SS, 39, 447

Lorimer, D. R. 2001, Living Reviews in Relativity, 4, 5

Lubow, S. H., Ogilvie, G. I., & Pringle, J. E. 2002, MNRAS, 337, 706

Lubow, S. H. & Shu, F. H. 1975, ApJ, 198, 383

Lyne, A. G., Manchester, R. N., & D'Amico, N. 1996, ApJl, 460, L41–

Lyubarskii, Y. E. 1997, MNRAS, 292, 679

Machin, G., Callanan, P. J., Charles, P. A., et al. 1990, MNRAS, 247, 205

Makishima, K., Kawai, N., Koyama, K., et al. 1984, PASJ, 36, 679

Makishima, K., Kubota, A., Mizuno, T., et al. 2000, ApJ, 535, 632

Maloney, P. R., Begelman, M. C., & Pringle, J. E. 1996, ApJ, 472, 582

Mardling, R. A. 1995a, ApJ, 450, 722

—. 1995b, ApJ, 450, 732

Margon, B. 1984, ARA&A, 22, 507

Margon, B., Deutsch, E. W., Leinhardt, Z. M., & Anderson, S. F. 1999, in International Astronomical Union Circular, 2–+

Margon, B., Lampton, M., Bowyer, S., & Cruddace, R. 1971, ApJl, 169, L23–

Marietta, E., Burrows, A., & Fryxell, B. 2000, ApJs, 128, 615

Marsh, T. R. & Steeghs, D. 2002, MNRAS, 331, L7

Marshall, N. & Millit, J. M. 1981, *Natur*, 293, 379

Martin, P. G. 1970, *MNRAS*, 149, 221

Mathis, J. S. & Lee, C.-W. 1991, *ApJ*, 376, 490

Mauche, C. W. & Gorenstein, P. 1986, *ApJ*, 302, 371

Mazeh, T. & Shaham, J. 1979, *A&A*, 77, 145

McClintock, J. E. & Remillard, R. A. 2003, To appear in "Compact Stellar X-ray Sources", eds. W.H.G. Lewin & M. van der Klis. See [astro-ph/0306213](#)

Middleditch, J. & Nelson, J. 1976, *ApJ*, 208, 567

Mignani, R. P., De Luca, A., Caraveo, P. A., & Mirabel, I. F. 2002, *A&A*, 386, 487

Mirabel, I. F. & Rodríguez, L. F. 1999, *ARA&A*, 37, 409

Mitsuda, K., Inoue, H., Koyama, K., et al. 1984, *PASJ*, 36, 741

Mitsuda, K., Inoue, H., Nakamura, N., & Tanaka, Y. 1989, *PASJ*, 41, 97

Moneti, A. 1992, *A&A*, 260, L7

Moon, D. & Eikenberry, S. S. 2001, *ApJl*, 549, L225

Motch, C., Stella, L., Janot-Pacheco, E., & Mouchet, M. 1991, *ApJ*, 369, 490

Mould, J. R. 1978, *ApJ*, 226, 923

Mukerjee, K., Agrawal, P. C., Paul, B., et al. 2001, *ApJ*, 548, 368

Muno, M. P., Fox, D. W., Morgan, E. H., & Bildsten, L. 2000, *ApJ*, 542, 1016

Muno, M. P., Remillard, R. A., & Chakrabarty, D. 2002, *ApJl*, 568, L35

Murray, J. R., Chakrabarty, D., Wynn, G. A., & Kramer, L. 2002, *MNRAS*, 335, 247

Murray, S. S., Austin, G. K., Chappell, J. H., et al. 2000, in Proc. SPIE Vol. 4012, p. 68-80, X-Ray Optics, Instruments, and Missions III, Joachim E. Truemper; Bernd Aschenbach; Eds., 68-80

Naik, S. & Paul, B. 2003, *A&A*, 401, 265

—. 2004, *A&A*, 418, 655

Narayan, R., Garcia, M. R., & McClintock, J. E. 1997, *ApJl*, 478, L79+

Natarajan, P. & Pringle, J. E. 1998, *ApJl*, 506, L97

Negueruela, I., Okazaki, A. T., Fabregat, J., et al. 2001, *A&A*, 369, 117

Negueruela, I., Reig, P., Coe, M. J., & Fabregat, J. 1998, *A&A*, 336, 251

Negueruela, I., Roche, P., Fabregat, J., & Coe, M. J. 1999, *MNRAS*, 307, 695

Nelson, R. P. & Papaloizou, J. C. B. 2000, *MNRAS*, 315, 570

Nelson, R. W., Bildsten, L., Chakrabarty, D., et al. 1997, *ApJl*, 488, L117+

Nowak, M. A. & Wilms, J. 1999, *ApJ*, 522, 476

O'Brien, K., Horne, K., Hynes, R. I., et al. 2002, *MNRAS*, 334, 426

Oegelman, H., Kahabka, P., Pietsch, W., Truemper, J., & Voges, W. 1985, *Space Science Reviews*, 40, 347

Ogilvie, G. I. 1999, *MNRAS*, 304, 557

—. 2000, *MNRAS*, 317, 607

Ogilvie, G. I. & Dubus, G. 2001, *MNRAS*, 320, 485

Okazaki, A. T. 2001, *PASJ*, 53, 119

Okazaki, A. T. & Negueruela, I. 2001, *A&A*, 377, 161

Oosterbroek, T., Parmar, A. N., Dal Fiume, D., et al. 2000, *A&A*, 353, 575

Orosz, J. A. & Kuulkers, E. 1999, *MNRAS*, 305, 132

Oudmaijer, R. D., Waters, L. B. F. M., van der Veen, W. E. C. J., & Geballe, T. R. 1995, *A&A*, 299, 69

Overbeck, J. W. 1965, *ApJ*, 141, 864

Papaloizou, J. C. B. & Pringle, J. E. 1983, *MNRAS*, 202, 1181

Parkinson, P. M. S., Tournear, D. M., Bloom, E. D., et al. 2003, *ApJ*, 595, 333

Parmar, A. N., Pietsch, W., McKechnie, S., et al. 1985, *Natur*, 313, 119

Paul, B. & Kitamoto, S. 2002, *Journal of Astrophysics and Astronomy*, 23, 33

Paul, B., Kitamoto, S., & Makino, F. 2000, *ApJ*, 528, 410

Pendleton, G. N., Paciesas, W. S., Mallozzi, R. S., et al. 1995, Nuclear Instruments and Methods in Physics Research A, 364, 567

Petterson, J. A. 1977, ApJ, 216, 827

—. 1978, ApJ, 224, 625

Phillips, S. N. & Podsiadlowski, P. 2002, MNRAS, 337, 431

Phinney, F. S. 1996, in ASP Conf. Ser. 90: The Origins, Evolution, and Destinies of Binary Stars in Clusters, 163—+

Pietsch, W., Mochejska, B. J., Misanovic, Z., et al. 2004, A&A, 413, 879

Plavec, M. & Kratochvil, P. 1964, Bulletin of the Astronomical Institutes of Czechoslovakia, 15, 165

Podsiadlowski, P. 1996, MNRAS, 279, 1104

—. 2003, astro-ph/0303660

Podsiadlowski, P. & Rappaport, S. 2000, ApJ, 529, 946

Predehl, P. & Schmitt, J. H. M. M. 1995, A&A, 293, 889

Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, Numerical recipes in FORTRAN. The art of scientific computing (Cambridge: University Press, —c1992, 2nd ed.)

Priedhorsky, W. C. & Terrell, J. 1984, ApJ, 280, 661

Pringle, J. E. 1981, ARA&A, 19, 137

—. 1992, MNRAS, 258, 811

—. 1996, MNRAS, 281, 357

Psaltis, D., Belloni, T., & van der Klis, M. 1999, ApJ, 520, 262

Ramsay, G., Zane, S., Jimenez-Garate, M. A., den Herder, J., & Hailey, C. J. 2002, MNRAS, 337, 1185

Rappaport, S., Ma, C. P., Joss, P. C., & Nelson, L. A. 1987, ApJ, 322, 842

Revnivtsev, M. & Sunyaev, R. 2003, A&A, 399, 699

Reynolds, A. P., Hilditch, R. W., Bell, S. A., & Hill, G. 1993, MNRAS, 261, 337

Ribó, M., Peracaula, M., Paredes, J. M., Núñez, J., & Otazu, X. 2001, in ESA SP-459: Exploring the Gamma-Ray Universe, 333-+

Roberts, T. P. & Warwick, R. S. 2000, MNRAS, 315, 98

Roberts, W. J. 1974, ApJ, 187, 575

Rolf, D. P. 1983, Natur, 302, 46

Romani, R. W. 1990, Natur, 347, 741

Rosenberg, A., Piotto, G., Saviane, I., & Aparicio, A. 2000, VizieR Online Data Catalog, 414, 40005

Ryter, C., Cesarsky, C. J., & Audouze, J. 1975, ApJ, 198, 103

Ryu, D. & Goodman, J. 1992, ApJ, 388, 438

Savage, B. D. & Jenkins, E. P. 1972, ApJ, 172, 491

Scargle, J. D. 1982, ApJ, 263, 835

—. 1989, ApJ, 343, 874

Schandl, S. 1996, A&A, 307, 95

Schandl, S. & Meyer, F. 1994, A&A, 289, 149

Scheuer, P. A. G. & Feiler, R. 1996, MNRAS, 282, 291

Schulz, N. S. & Brandt, W. N. 2002, ApJ, 572, 971

Schwarzenberg-Czerny, A. 1992, A&A, 260, 268

—. 1998, MNRAS, 301, 831

—. 1999, ApJ, 516, 315

Scott, D. M., Leahy, D. A., & Wilson, R. B. 2000, ApJ, 539, 392

Shafter, A. W., Szkody, P., & Thorstensen, J. R. 1986, ApJ, 308, 765

Shakura, N. I. & Sunyaev, R. A. 1973, A&A, 24, 337

Shirakawa, A. & Lai, D. 2002a, ApJ, 565, 1134

—. 2002b, ApJ, 564, 361

Shirey, R. E. 1998, Ph.D. Thesis

Shirey, R. E., Bradt, H. V., & Levine, A. M. 1999a, *ApJ*, 517, 472

Shirey, R. E., Bradt, H. V., Levine, A. M., & Morgan, E. H. 1998, *ApJ*, 506, 374

Shirey, R. E., Levine, A. M., & Bradt, H. V. 1999b, *ApJ*, 524, 1048

Sidoli, L., Parmar, A. N., Oosterbroek, T., et al. 2001, *A&A*, 368, 451

Smale, A. P. 2001, *ApJ*, 562, 957

Smale, A. P., Church, M. J., & Bałucińska-Church, M. 2001, *ApJ*, 550, 962

Smale, A. P. & Lochner, J. C. 1992, *ApJ*, 395, 582

Smith, D. A., Morgan, E. H., & Bradt, H. 1997, *ApJl*, 479, L137+

Smith, D. M., Heindl, W. A., & Swank, J. H. 2002a, *ApJl*, 578, L129

Smith, G. & Churilov, V. 2003, in *Instrument Design and Performance for Optical/Infrared Ground-based Telescopes*. Edited by Iye, Masanori; Moorwood, Alan F. M. *Proceedings of the SPIE*, Volume 4841, pp. 1337-1344 (2003)., 1337-1344

Smith, R. K., Edgar, R. J., & Shafer, R. A. 2002b, *ApJ*, 581, 562

Steeghs, D. & Casares, J. 2002, *ApJ*, 568, 273

Stellingwerf, R. F. 1978, *ApJ*, 224, 953

Stewart, R. T., Caswell, J. L., Haynes, R. F., & Nelson, G. J. 1993, *MNRAS*, 261, 593

Stewart, R. T., Nelson, G. J., Penninx, W., et al. 1991, *MNRAS*, 253, 212

Still, M., O'Brien, K., Horne, K., et al. 2001, *ApJ*, 554, 352

Strohmayer, T. & Bildsten, L. 2003, To appear in "Compact Stellar X-ray Sources", eds. W.H.G. Lewin & M. van der Klis. See [astro-ph/0301544](http://arxiv.org/abs/astro-ph/0301544)

Strohmayer, T. E., Zhang, W., Swank, J. H., et al. 1996, *ApJl*, 469, L9+

Taam, R. E. & van de Heuvel, E. P. J. 1986, *ApJ*, 305, 235

Tanaka, Y. 1997, in *Accretion Disks - New Aspects*, 1-23

Tananbaum, H., Chaisson, L. J., Forman, W., Jones, C., & Matilsky, T. A. 1976, *ApJl*, 209, L125

Tauris, T. M., Fender, R. P., van den Heuvel, E. P. J., Johnston, H. M., & Wu, K. 1999, MNRAS, 310, 1165

Tauris, T. M. & Heuvel, E. d. 2003, To appear in "Compact Stellar X-ray Sources", eds. W.H.G. Lewin & M. van der Klis. See astro-ph/0303456

Taylor, A. R. & Gregory, P. C. 1982, ApJ, 255, 210

Tennant, A. F. 1987, MNRAS, 226, 971

Tennant, A. F., Fabian, A. C., & Shafer, R. A. 1986, MNRAS, 219, 871

Tout, C. A. 2000, New Astronomy Review, 44, 37

Truemper, J., Kahabka, P., Oegelman, H., Pietsch, W., & Voges, W. 1986, ApJl, 300, L63

Ubertini, P., Bazzano, A., Cocchi, M., et al. 1999, ApJl, 514, L27

Šimon, V. 2002, A&A, 381, 151

—. 2003, A&A, 405, 199

van der Klis, M. 1994, A&A, 283, 469

van Kerkwijk, M. H., Chakrabarty, D., Pringle, J. E., & Wijers, R. A. M. J. 1998, ApJl, 499, L27+

van Kerkwijk, M. H., van Oijen, J. G. J., & van den Heuvel, E. P. J. 1989, A&A, 209, 173

van Paradijs, J., Dotani, T., Tanaka, Y., & Tsuru, T. 1990, PASJ, 42, 633

van Paradijs, J. & Kuiper, L. 1984, A&A, 138, 71

van Paradijs, J. & McClintock, J. E. 1995, Optical and Ultraviolet Observations of X-ray Binaries (X-ray Binaries, eds. W.H.G. Lewin, J. van Paradijs, and E.P.J. van den Heuvel (Cambridge: Cambridge Univ. Press), p. 58), 58-+

van Straaten, S., van der Klis, M., & Méndez, M. 2003, ApJ, 596, 1155

Vaughan, S., Fabian, A. C., & Nandra, K. 2003, MNRAS, 339, 1237

Verbunt, F. 1993, ARA&A, 31, 93

Verbunt, F. & Hut, P. 1987, in IAU Symp. 125: The Origin and Evolution of Neutron Stars, 187-+

Verbunt, F. & Zwaan, C. 1981, A&A, 100, L7

Verrecchia, F., Israel, G. L., Negueruela, I., et al. 2002, A&A, 393, 983

Vrtilek, S. D., Boroson, B., Primini, F. A., et al. 1997, Bulletin of the American Astronomical Society, 29, 1389

Vrtilek, S. D., Mihara, T., Primini, F. A., et al. 1994, ApJl, 436, L9

Vrtilek, S. D., Raymond, J. C., Boroson, B., et al. 2001, ApJl, 563, L139

—. 2003, pasp, 115, 1124

Vrtilek, S. D., Raymond, J. C., Garcia, M. R., et al. 1990, A&A, 235, 162

Wachter, S. 1997, ApJ, 485, 839

Wall, J. V. & Jenkins, C. R. 2003, Practical statistics for astronomers (Cambridge Observing Handbooks for Research Astronomers, Cambridge, New York: Cambridge University Press)

Warner, B. 1995, Cataclysmic variable stars (Cambridge Astrophysics Series, Cambridge, New York: Cambridge University Press, —c1995)

White, N. E. & Angelini, L. 2001, ApJl, 561, L101

White, N. E., Becker, R. H., Boldt, E. A., et al. 1981, ApJ, 247, 994

White, N. E. & Holt, S. S. 1982, ApJ, 257, 318

White, N. E., Stella, L., & Parmar, A. N. 1988, ApJ, 324, 363

Wijers, R. A. M. J. & Pringle, J. E. 1999, MNRAS, 308, 207

Wijnands, R., Miller, J. M., & van der Klis, M. 2002, MNRAS, 331, 60

Wijnands, R. & van der Klis, M. 1998, Natur, 394, 344

—. 1999a, ApJ, 514, 939

—. 1999b, Bulletin of the American Astronomical Society, 32, 882

Wijnands, R. A. D., Kuulkers, E., & Smale, A. P. 1996, ApJl, 473, L45+

Wijnands, R. A. D., van der Klis, M., Kuulkers, E., Asai, K., & Hasinger, G. 1997, A&A, 323, 399

Wojdowski, P., Clark, G. W., Levine, A. M., Woo, J. W., & Zhang, S. N. 1998, ApJ, 502, 253

Wojdowski, P. S., Clark, G. W., & Kallman, T. R. 2000, ApJ, 541, 963

Wolff, M. T., Wood, K. S., Ray, P. S., et al. 2000, Bulletin of the American Astronomical Society, 32, 1218

Zamanov, R. K., Martí, J., Paredes, J. M., et al. 1999, A&A, 351, 543

Zeldovich, Y. B. & Novikov, I. D. 1999, Stars and relativity (Stars and relativity / Ya. B. Zeldovich and I.D. Novikov ; translated by Eli Arlock ; edited by Kip S. Thorne and W. David Arnett. Mineola, N.Y. : Dover Publications, 1996.)

Zezas, A., Fabbiano, G., Rots, A. H., & Murray, S. S. 2002, ApJs, 142, 239

Zombeck, M. V. 1990, Handbook of space astronomy and astrophysics (Cambridge: University Press, 1990, 2nd ed.)

## Frank and Ernest



Copyright (c) 1984 by Thaves. Distributed from [www.thecomics.com](http://www.thecomics.com).

## Acknowledgements

This section is not exhaustive as there have been many people who have aided this work either through direct involvement with the analysis or background encouragement. Thanks especially to Jeno Sokoloski for an intense week of introduction to the world of astronomical observation, and to Silas Laycock for many useful discussions about time-series analysis, and to Philipp Podsiadlowski for exciting and fruitful discussions (and his patience with my misconceptions). Thanks also to my officemates Simon Shaw, Diana Hannikainen, I-Chun Shih, Julie-Ann Goodlet and Ross Collins for entertaining company and interesting discussion, and for putting up with my unpredictable hours and behavior during proposal deadlines. Nick Seymour, Sofia Araujo-Betancor and Anna Watts are thanked for invaluable help with thesis formatting (and much more besides). The students and staff at the Southampton astronomy group are thanked for providing an exciting and enjoyable environment in which to work. Thanks also to Fred Marang, Francois van Wyck, Ian Glass, Stuart Ryder and Simon Clark for help with astronomical data taking and analysis.

I also acknowledge the encouragement, support, availability and infectious enthusiasm of Phil Charles, my thesis advisor for three years of working together on some six separate research projects and at least fifteen telescope proposals. Thanks also for tolerating my somewhat unorthodox working schedule and many trips to Kansas and Africa.

At the Open University, it is a pleasure to thank Robin Barnard for many useful and enjoyable discussions and for partially “de-gimping” my understanding of X-ray spectra and BibTEX. My postdoc supervisor, Carole Haswell, is warmly thanked for her patience while I attempted to write this work at the same time as beginning a postdoc in the fast-moving Wide Angle Search for Planets.

Steve Foulkes, James Murray and Carole Haswell gave me permission to reproduce an image from

their exciting new simulations of disk-warping, and are thanked for this; I also thank Martin Still for useful discussions about wind-driven warping and permission to include information from the new work on Her X-1 and the recent ALS.

Turning away from the strictly professional, it is a pleasure to thank all the members of Cosmos Education, but most importantly the teachers and students we encountered in our travels.

Thanks also to my family, Ian, Mary and James, and grandparents Bill & Joanne, for encouragement, support and patience, and also to friends who have had to listen to me talk about binaries at great length!

Finally I especially thank Shawna, for providing inspiration, perspective, support and stability, and without whom this work could not have been completed.

"There is a theory which states that if anyone discovers just exactly what the universe is for and why we are here, that it will instantly disappear and be replaced by something even more bizarre and inexplicable.

There is another theory which states that this has already happened."

– Douglas Adams