

UNIVERSITY OF SOUTHAMPTON

**X-ray power spectral densities of
Active Galactic Nuclei**

Daniel Paul Summons

Submitted for the degree of Doctor of Philosophy

SCHOOL OF PHYSICS AND ASTRONOMY

SCHOOL OF ENGINEERING SCIENCE

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS

December 2007

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS

SCHOOL OF PHYSICS AND ASTRONOMY

SCHOOL OF ENGINEERING SCIENCE

Doctor of Philosophy

X-ray power spectral densities of Active Galactic Nuclei

by Daniel Paul Summons

I address the question of whether active galactic nuclei (AGN) exhibit similar X-ray temporal variability properties to black hole X-ray binaries (BHXRBs). I utilise long time-scale X-ray monitoring data to produce broadband power spectral densities (PSDs) for a variety of AGN. Unfortunately X-ray monitoring programmes of AGN often produce light curves that are irregularly sampled, which causes the observed PSD to become distorted. I therefore further developed a Monte Carlo modelling technique (PSRESP), based upon Uttley *et al.* (2002), to determine the undistorted broadband PSD and associate a robust acceptance probability to the fitted model. Using archival and proprietary *Rossi X-ray Timing Explorer (RXTE)* data, along with X-ray data from *XMM-Newton*, I apply PSRESP to determine the undistorted PSD of NGC 3783, which had previously been suggested to be analogous to a ‘hard’ state BHXRB. I show that a second break is not statistically supported, and I show that this PSD is, in fact, well fitted by a ‘soft’ state model that has only one break at higher frequencies. These results leave Arakelian 564 as the only AGN which shows a second break at low frequencies. I present a power spectral analysis of a 100 ksec *XMM-Newton* observation along with *RXTE* and *ASCA* data of the narrow line Seyfert 1 galaxy Ark 564. I demonstrate that the PSD of Ark 564 is well fit by the sum of two Lorentzian-shaped components, similar to those seen in BHXRBs. The PSD and lag spectrum of Ark 564 strongly suggests that a two-component Lorentzian interpretation of the origin of most of the variability is correct. Given the super-Eddington accretion rate of Ark 564 it is most likely analogous to a very-high state system. These results further confirm the connection between AGN and BHXRBs. I use proprietary X-ray data from our *RXTE* monitoring programme, together with other data from the *RXTE* archive and complementary data from *XMM-Newton*, *ASCA*, and *EXOSAT* to produce broadband PSDs for 32 AGN spanning the frequency range from $\sim 10^{-8}$ to $\sim 10^{-3}$ Hz. The X-ray data for each AGN are reduced and analysed in a consistent manner. In particular, to ensure consistency between different *RXTE* gain epochs, all *RXTE* monitoring data is calibrated with respect to public archival Crab PCA data. I apply PSRESP to the resulting PSDs to determine the undistorted shape. All of the 32 AGN are consistent with either an unbroken or single-bend power-law interpretation of their PSDs, with the exception of Ark 564. I thus show that, for an expanded sample, AGN do behave like scaled-up BHXRBs, with the bend timescale scaling as mass divided by accretion rate. Thus, I further endorse the result of McHardy *et al.* (2006).

ACKNOWLEDGEMENTS

Over the past three years I have accumulated a fair list of people that have, in some way, coached me towards writing this magnum opus.

First and foremost I would like to extend my eternal gratitude to Ian M^cHardy, Phil Uttley, Patricia Arévalo and Atul Bhaskar. Without their patience and encouragement this work would not have been possible. I would also like to extend my thanks to all the remaining members of the Southampton astronomy community who have always offered their unwavering support, I owe you all a pint and if you are lucky I might even stretch to some nuts, dry roasted of course!

I have made some great new friends during my time in the Astronomy Group, but I feel obligated to give special mention to Joker and Clown or their nom de voyage - Chris and OJ. They have always ensured I never took anything too seriously, which was especially helpful when writing this tome. Their never-ending bombardment of quotes, tomfoolery, and shenanigans have kept a smile on my face. This one is for Joker and Clown:

Saturn, Jupiter, Pluto... Mars, Venus... George, Chr... Chrissy. That is a list of my best friends in order... order of preference.

NOEL FIELDING (1973 –), AS THE MOON

I must also give thanks to my friends who have ensured that I never got too buried in my work and for making sure I always take a break and gain some much need perspective; how could I ever forget Thursday night games and the Wii-lympics!

I thank Mum, Ian, and Becky for their encouragement, love and support they have always shown me. I could never of got this far without them. My only regret is that Ray will never read the work that has been dedicated to him. Finally, I would like to thank Leanne Cannell. She has been my rock and my friend since my journey through the astrophysics looking glass began. Many thanks to you all!

Daniel Summons, December 2007

They've done studies you know, 60% of the time it works everytime.

PAUL RUDD (1969 –), AS BRIAN FANTANA

For Ray

CONTENTS

1	Introduction	1
1.1	Active Galactic Nuclei	2
1.1.1	Accretion onto a black hole	3
1.1.2	Unified model of AGN	5
1.1.3	X-ray emission	7
1.1.4	X-rays and temporal variability of AGN	10
1.2	Black hole X-ray binary systems	12
1.2.1	Properties of BHXRBs	13
	Quasi-periodic oscillations	14
1.3	Comparing AGN and BHXRBs	16
1.4	AGN X-ray monitoring	18
1.4.1	<i>Rossi X-ray Timing Explorer (RXTE)</i>	19
1.4.2	<i>XMM-Newton</i>	21
1.5	Summary of thesis	22
2	Measuring temporal variability	24
2.1	Variability and the power spectral density	25
2.1.1	The importance of variance	25
2.1.2	Periodogram estimates of the power spectrum	26
	Properties of the periodogram	30
2.2	Light curve simulation	32
2.3	Sampling biases in AGN power spectra	34
2.3.1	Aliasing	36
2.3.2	Red-noise leak	38
2.3.3	Added complication of uneven sampling and binning	38
2.4	Significance of red-noise leak	39
2.5	Stationarity	41
2.6	Summary	44
3	Power spectral modelling: PSRESP	46
3.1	Response method	46
3.2	Improved PSRESP implementation	49

3.3	Technical overview of the major improvements to PSRESP	55
3.3.1	The Poisson noise level and simulated error bars	56
3.3.2	Distortions, skewed distributions and acceptance probability determination	59
3.4	Computation, optimisation and Iridis2	67
3.4.1	FFT	67
3.4.2	Aliasing and red-noise leak	68
3.4.3	Optimisation	69
3.4.4	Iridis2	74
3.5	Assumptions and limitations of PSRESP	74
3.6	Overview	76
4	The broadband power spectrum of NGC 3783 and Ark 564	78
4.1	Timing analysis and accretion state determination of NGC 3783	79
4.1.1	Observations and data reduction	81
4.1.2	Combining <i>RXTE</i> and <i>XMM-Newton</i> data	84
4.1.3	PSRESP and the power spectrum of NGC 3783	85
4.1.4	Fitting models to the power spectrum of NGC 3783	87
	Unbroken power-law PSD model	87
	Single-bend power-law PSD model	90
	Double-bend power-law PSD model	92
	Single-bend power-law PSD model with a Lorentzian component	96
4.1.5	Discussion and Conclusions	99
4.2	The Narrow Line Seyfert 1 Ark 564: power spectral analysis	100
4.2.1	Observations and data reduction	101
	<i>RXTE</i>	101
	<i>ASCA</i>	101
	<i>XMM-Newton</i>	102
4.2.2	The power spectrum of Ark 564 and model fitting	103
	Double-bend power-law PSD model	104
	Double-Lorentzian PSD model	106
4.2.3	Two-component Lorentzian PSD model and lag spectra	112
4.2.4	Discussion and Conclusions	114
5	X-ray power spectra	117
5.1	Observations and data reduction	118

5.1.1	<i>RXTE</i> Data Reduction	126
	Calibrating <i>RXTE</i> AGN data	126
5.1.2	<i>XMM-Newton</i> Data Reduction	127
5.1.3	The Light Curves	129
5.2	Power spectral modelling	129
5.2.1	Combining X-ray data	130
5.2.2	Monte Carlo simulations: PSRESP	130
5.3	Results	131
5.3.1	Unbroken power-law PSD model	131
5.3.2	Single-bend power-law PSD model	132
5.3.3	Double-bend power-law PSD model	149
5.3.4	Single-bend power-law PSD model with a fixed low-slope	150
5.4	Correlation between bend time-scale, black hole mass and bolometric luminosity	152
5.4.1	Correlation of bend time-scale with the $H\beta$ optical emission linewidth	164
5.5	Discussion and Conclusions	166
6	Conclusions	168
6.1	The future of PSRESP and possible optimisations	169
6.1.1	Statistics	169
6.1.2	Non-linear light curves and blazar data	171
6.1.3	Optical PSDs	172
6.1.4	Optimisations	173
6.2	Cycle 12 of <i>RXTE</i> : NGC 3783	173
6.3	Is Ark 564 unique?	174
6.4	Survey results	175
6.5	Intermediate mass black holes	177
6.6	Accretion state connections	178
6.7	Physical interpretation of the bend-frequency	179
6.8	A brief introduction to future missions in high-energy astrophysics	182
6.8.1	<i>Swift</i>	182
6.8.2	<i>Astrosat</i>	183
6.8.3	<i>The Gamma-ray Large Area Space Telescope (GLAST)</i>	183
6.8.4	X-Ray Evolving Universe Spectrometer Mission (<i>XEUS</i>)	184
6.9	Final remarks	185

APPENDICES

A	<i>RXTE</i> monitoring X-ray light curves	187
B	<i>RXTE</i> medium/short time-scale X-ray light curves	193
C	<i>XMM-Newton</i>, <i>ASCA</i> and <i>EXOSAT</i> X-ray light curves	198
	Bibliography	204

LIST OF TABLES

3.1	Comparison of χ^2_{pseudo} distribution to χ^2 for <i>RXTE</i> data	60
3.2	Comparison of χ^2_{pseudo} distribution to χ^2 for <i>XMM-Newton</i> data	62
3.3	Comparison of PSRESP versions	65
3.4	Comparison between parameters found by grid search and downhill simplex	71
4.1	Summary of the <i>RXTE</i> and <i>XMM-Newton</i> light curves used in the analysis of NGC 3783	83
4.2	Unbroken, single and double-bend power-law fits to the PSD of NGC 3783	91
4.3	Fit result of a single-bend power-law with Lorentzian component model to the PSD of NGC 3783	98
4.4	Best-fitting model parameters to Ark 564	111
5.1	AGN survey list	120
5.1	AGN survey list	121
5.1	AGN survey list	122
5.1	AGN survey list	123
5.1	AGN survey list	124
5.1	AGN survey list	125
5.2	Unbroken power-law model PSD fit results	137
5.3	Single-bend power-law model PSD fit results	147
5.3	Single-bend power-law model PSD fit results	148
5.4	Double-bend power-law model PSD fit results	151

5.5	Single-bend power-law with fixed low-frequency slope model PSD fit results	159
5.6	Black hole masses and accretion rates	160

LIST OF FIGURES

1.1	Artist's impression of an accretion disk and torus	4
1.2	Unified model of AGN	6
1.3	Comparison between an optical and X-ray image of NGC 1365	9
1.4	Example of a short time-scale AGN X-ray light curve	11
1.5	Artist's impression of an XRB	12
1.6	Power spectra of black holes in various emission states	15
1.7	Hardness-intensity diagram	16
1.8	PSDs of Cyg X-1 and AGN	18
1.9	The <i>RXTE</i> spacecraft	20
1.10	The <i>XMM-Newton</i> spacecraft	22
2.1	Illustration of the Poisson noise level	29
2.2	Simulated periodograms from an underlying PSD model	32
2.3	Light curves generated from the simulated periodograms in Fig. 2.2	33
2.4	Example light curve with regular sampling pattern	35
2.5	Illustration of aliasing	37
2.6	Red-noise leak in measured periodograms	42
3.1	An example <i>RXTE</i> light curve with typical sampling patterns	50
3.2	A typical broadband periodogram	52
3.3	Ensemble of simulated periodograms	53
3.4	Two simulated power spectra illustrating the addition of the Poisson noise level in the original version of PSRESP	57
3.5	The high-frequency periodogram using the improved Poisson noise calculation in PSRESP	58
3.6	Typical χ^2_{total} distributions for <i>RXTE</i> data constructed from χ^2_{pseudo} and the new χ^2 statistic	60
3.7	Typical χ^2_{total} distributions for <i>XMM-Newton</i> data constructed from χ^2_{pseudo} and the new χ^2 statistic	62

3.8	Comparison of rejection probability levels produced by different versions of PSRESP	66
3.9	A typical probability surface encountered in AGN fitting	73
3.10	Comparison between a linear and exponentiated light curve	75
4.1	<i>RXTE</i> X-ray light curves of NGC 3783 in flux units	82
4.2	<i>XMM-Newton</i> X-ray light curves of NGC 3783 in flux units	85
4.3	Periodogram of NGC 3783	88
4.4	The best fit unbroken power-law PSD of NGC 3783	89
4.5	The best fit single-bend power-law PSD of NGC 3783	90
4.6	NGC 3783 rejection probability levels for the single-bend model fit .	93
4.7	The best fit double-bend power-law PSD of NGC 3783	94
4.8	NGC 3783 rejection probability levels for NGC 3783 for the double-bend model fit	95
4.9	The best fit single-bend power-law with a Lorentzian component PSD model of NGC 3783	97
4.10	<i>RXTE</i> long and medium- term light curve of Ark 564 in the 2-10 keV energy range	102
4.11	Ark 564 <i>ASCA</i> light curves	103
4.12	Ark 564 <i>XMM-Newton</i> PN light curves	104
4.13	Power spectrum of Ark 564 from <i>RXTE</i> , <i>ASCA</i> and <i>XMM-Newton</i> data in the 2–10 keV <i>RXTE</i> energy band.	105
4.14	Power spectrum of Ark 564 from <i>RXTE</i> , <i>ASCA</i> and <i>XMM-Newton</i> data in the 0.2-2 <i>ASCA</i> energy band.	106
4.15	Ark 564 rejection probability levels for a double-bend power-law PSD model	107
4.16	Two-component Lorentzian PSD model fit to Ark 564	109
4.17	Ark 564 rejection probability levels for the high and low peak-frequency components in the two-component Lorentzian model fit .	110
4.18	Ark 564 rejection probability levels for the low and high-frequency Lorentzian normalisations in the two-component Lorentzian model fit	113
4.19	Combined plot of the Ark 564 lag spectrum and two-component Lorentzian PSD model	115
5.1	<i>RXTE</i> 2–10 keV light curve of the Crab	128
5.2	Unbroken power-law PSD model fits to all targets	133
5.2	Unbroken power-law PSD model fits to all targets	134
5.2	Unbroken power-law PSD model fits to all targets	135

5.2	Unbroken power-law PSD model fits to all targets	136
5.3	Single-bend power-law PSD model fits to all targets	139
5.3	Single-bend power-law PSD model fits to all targets	140
5.3	Single-bend power-law PSD model fits to all targets	141
5.3	Single-bend power-law PSD model fits to all targets	142
5.4	Single-bend power-law rejection probability levels for all targets . .	143
5.4	Single-bend power-law rejection probability levels for all targets . .	144
5.4	Single-bend power-law rejection probability levels for all targets . .	145
5.4	Single-bend power-law rejection probability levels for all targets . .	146
5.5	Double-bend power-law PSD model fits	150
5.6	Single-bend power-law PSD model with fixed low-frequency slope for selected targets	153
5.5	Single-bend power-law PSD model with fixed low-frequency slope for selected targets	154
5.5	Single-bend power-law PSD model with fixed low-frequency slope for selected targets	155
5.6	Single-bend power-law PSD model with fixed low slope rejection probability levels	156
5.6	Single-bend power-law PSD model with fixed low slope rejection probability levels	157
5.6	Single-bend power-law PSD model with fixed low slope rejection probability levels	158
5.7	Correlation between black hole mass and bolometric indices for 20 objects	161
5.8	Predicted break time-scale versus observed break time-scale	163
5.9	Observed bend time-scale versus FWHM of $H\beta$ emission line for 18 objects	165
6.1	Predicted power spectrum of NGC 3783 with new <i>RXTE</i> data	174
6.2	Power spectrum of Ton S 180	176
6.3	Radio observations of NGC 4051	180
A.1	<i>RXTE</i> long time-scales light curves	187
B.1	<i>RXTE</i> medium/short time-scales light curves	193
C.1	Short time-scales light curves	198
C.1	Short time-scales light curves	199
C.1	Short time-scales light curves	200
C.1	Short time-scales light curves	201

C.1 Short time-scales light curves	202
C.1 Short time-scales light curves	203

It's not whether you get knocked down, it's whether you get up.

VINCE LOMBARDI (1913 – 1970)

A hero is an ordinary individual who finds the strength to persevere and endure in spite of overwhelming obstacles.

CHRISTOPHER REEVE (1952 – 2004)

Nothing shocks me. I'm a scientist.

HARRISON FORD (1942 -), AS INDIANA JONES

1

Introduction

In this chapter I will review the prevailing understanding of some of the Universe's most intriguing objects - the active galactic nuclei (AGN). I will begin by outlining the AGN observationally and note some of their more prominent characteristics. I will then describe the energy generation mechanism, namely accretion, and tie together the different observed characteristics of AGN into a single unified framework. This work is based on measuring the temporal X-ray variability of AGN, so I will describe the importance of X-rays and how temporal variability can probe some of the most extreme environments in the Universe: the inner accretion disk around a black hole. AGN are not the only objects that are thought to be powered by accreting black holes, (Galactic) black hole X-ray binary systems or BHXRBs are also thought to contain accreting black holes, albeit smaller in mass. I will briefly describe the BHXRBs and discuss how BHXRB properties such as quasi-periodic oscillations and emission states might be translated to their more massive brethren. I will conclude this chapter with an introduction of the instruments and satellites used to gather the X-ray data used herein.

1.1 Active Galactic Nuclei

In the proverbial zoo of astronomical phenomena, the active galaxy truly is one of the most sensational objects in the extragalactic Universe. Detected over a broad range of redshifts and representing between 20 and 40 per cent of the galaxy population¹ (Miller *et al.*, 2003), the active galaxy is noteworthy by its extraordinary luminosity manufactured in a minute region in space at the core of these massive objects. The shortest time-scale variations in X-ray flux, on the order of minutes, indicates that these cores or active galactic nuclei (AGN) are constrained to volumes less than 10^{-3} pc^3 and represent a miniscule fraction of the host galaxy. At the core of these active galaxies, the AGN can emit tremendous amounts of energy over a broad range of frequencies, from radio wavelengths through to X-ray and gamma-rays, and exhibit a flat spectral energy distribution over several decades in frequency (Elvis *et al.*, 1994). Most AGN exhibit a marked reduction in luminosity at radio wavelengths with only 10 per cent of AGN being classed as *radio-loud* and the remaining subset called *radio-quiet*. AGN are also observed to have an extreme range of luminosities, from the exceptionally bright quasar PDS 456 (Reeves *et al.*, 2000) at $\sim 10^{47} \text{ ergs s}^{-1}$ through to the low-luminosity AGN NGC 4395 (Lira *et al.*, 1999) at $\sim 10^{41} \text{ ergs s}^{-1}$. Aharonian *et al.* (2005) recently suggested that AGN may be the source of some ultra-high energy GeV and TeV emission, and in November 2007 the Pierre Auger Collaboration published a result in *Science* claiming to correlate the origin of these cosmic rays to AGN². Throughout their lifetime, AGN will be one of the most prolific emitters of broadband radiation in the known Universe. The AGN menagerie is vast and complicated, with AGN categorised and classed according to several observed characteristics. The first of these characteristics were introduced above: radio-loud and radio-quiet, which as the nomenclature suggests represents the classes of AGN distinguished by their radio luminosities. Fanaroff and Riley (1974) further developed the distinction between radio galaxies by subdividing them into *FR1* and *FR2* type based upon radio luminosity and morphology. High-luminosity AGN are often classified as *quasars* and are either radio-loud or radio-quiet, which are often denoted by the acronyms *QSRS* (quasi-stellar radio source) or *QSO* (quasi-stellar object), respectively. *Blazar* is the term used to denote a breed of AGN where the variations in the optical band are marked by atypically rapid and large amplitude variability. It is currently thought that blazar emission

¹Calculated from over 10000 galaxies between $0.05 < z < 0.095$.

²Correlation of the Highest-Energy Cosmic Rays with Nearby Extragalactic Objects, The Pierre Auger Collaboration (9 November 2007), *Science* 318 (5852), 938. [DOI: 10.1126/science.1151124].

could be beamed due to observations looking down the jet structure. Seyfert galaxies are characterised by relatively weak radio emission and are subdivided into two types (type 1 and type 2) based on whether their spectra contain both narrow and broad optical emission lines (type 1), or only narrow lines (type 2). AGN classified as a Seyfert type galaxy are generally found to have lower luminosities than some of the other classes. The final category I will mention is *Low-Ionisation Nuclear Emission Region (LINERs)*, which are similar to Seyfert galaxies but exhibit line emission from weakly ionised atoms, whereas the emission from strong ionised emission lines are weak (Heckman, 1980). The distinction between the AGN classifications is defined observationally, since the underlying mechanism thought to power these active galaxies is the same, namely accretion onto a supermassive black hole (SMBH).

1.1.1 Accretion onto a black hole

The prodigious emission of AGN is more than can be accounted for by nuclear fusion, as observed in stellar reactions. It is commonly thought that the only practical method that could potentially generate energy with a great enough efficiency is accretion onto a compact object. The process of nuclear fusion is only 0.7 per cent efficient at converting mass into energy, which if radiated is far too small to explain AGN observations. Salpeter (1964) propose emission is driven by accretion onto a SMBH, with a mass between $10^6 - 10^9 M_{\odot}$, which has a much larger efficiency than nuclear fusion of ~ 10 per cent³. The accretion disk is a structure of gas in orbital motion around a central massive object. The diffuse matter rotates and settles in the equatorial plane, and the interplay of magnetic and viscous forces promotes instabilities within the disk, which transport angular momentum outwards throughout the disk. As a result, the matter spirals inward towards the central black hole and the potential gravitational energy released is radiated away. Fig. 1.1 is an artist's impression of an accretion disk surrounded by an obscuring torus (see Section 1.1.2)

The observed energy spectra of quasars often display a peak at ~ 0.01 keV (optical to UV energies), which is frequently referred to as the *Big Blue Bump*. The observed peak is consistent with models of optically thick, geometrically thin accretion disks, that is where the ratio between the scale height, H , and disk radius, R , is much less than unity (Shakura and Syunyaev, 1973). It is common for accretion disk

³The range of efficiencies extend up to approximately 42 per cent if the black hole is extreme Kerr.




Figure 1.1: Artist's impression of an accretion disk and surrounding torus. The cold dusty torus is clearly seen, with an inner accretion disk feeding the central supermassive black hole. The outflow represents a possible relativistic jet structure. Credit: NASA/CXC/M.Weiss; <http://chandra.harvard.edu/resources/illustrations/agn/>

models to correctly predict the existence of the big blue bump in UV (e.g. Laor 1990); however, data that spans a broad energy range becomes less well fit by thin accretion disk models and these models often erroneously predict the flux at soft X-ray or infrared energies. Even more complicated properties have been included in accretion disk modelling to better replicate observed data; for example, radiation pressure close to the central black hole could inflate the disk so that $H/R \sim 1$. However, it is not clear whether these geometrically thick disks are a more faithful representation of real accretion disks.

It has been shown that a key assumption of optically thick accretion disk theory is not necessarily justifiable; namely, that the energy dissipated at a given radius is radiated near to this radius (Paczynsky and Wiita, 1980; Rees *et al.*, 1982; Abramowicz *et al.*, 1988; Narayan and Yi, 1994). The so-called *advection-dominated accretion flow* (ADAF) has become an alternative model of accretion to the previously mentioned optically thick accretion disk, when the accretion rate is low. When the cooling time is long compared to the viscous time-scale (see Section 6.7), the energy is advected inwards with the accretion flow rather than being radiated. The radiative efficiency of an ADAF can be significantly smaller than the radiative efficiency

predicted from optically thick accretion disk models. Black hole systems where the observed luminosity is less than predicted for optically thick disks can be explained if an ADAF model is applied; for example, the SMBH at the centre of our own galaxy, Sagittarius A*, can be explained in terms of an ADAF (Mahadevan, 1998). Moreover, radiatively inefficient accretion disk models are often invoked to explain the differences between black hole X-ray binaries in the *low/hard* and *high/soft* state. The high/soft state is often thought of as having a radiatively efficient accretion disk, whereas the low/hard state is often envisaged as having a radiatively inefficient accretion disk; I will describe these accretion states later in this chapter. Another use of radiative inefficient accretion is to explain the lack of strong emission from the SMBHs at the centre of massive elliptical galaxies in clusters, where high luminosities and accretion rates might be predicted using an optically thick disk model (Fabian and Rees, 1995).

1.1.2 Unified model of AGN

The central regions of AGN remain spatially unresolved making direct imaging of features impossible. It was once thought that AGN constituted a diverse population of objects; however, Antonucci and Miller (1985) observed broad lines in the spectrum of NGC 1068, a Seyfert 2 galaxy, when measured using polarised light. The spectrum of NGC 1068 consequently resembles that of a Seyfert 1 galaxy; thus, this discovery hints that perhaps the nucleus of NGC 1068 is heavily obscured so that direct observation is impossible, but in scattered polarised light the characteristics of a Seyfert 1 galaxy are clearly seen. The concept of viewing angle dependence was further developed by Antonucci (1993) into a unification formulation of active galaxies. The unified model of AGN consolidates all the observational descriptions of AGN into a single framework. The unified model assumed throughout this work is presented in Fig. 1.2. The unified model of AGN consists of the following features: a central SMBH and accretion disk, constituting the primary source of emission; a relativistic jet can also be observed in some sources, thought to emerge from a region close to the central black hole; high-velocity clouds in Keplerian orbits constitute the broad-line region (BLR) i.e. emission lines exhibit broad wings due to Doppler broadening. Much further out are lower density and lower velocity clouds where the narrow permitted and forbidden lines are measured and constitute the narrow-line region (NLR). The BLR region is encompassed by a large high-column density material, the so-called molecular torus. Almost any observed AGN can fit into the unified paradigm, at least phenomenologically. If the AGN is radio-quiet




Figure 1.2: Diagram of the standard AGN unification paradigm (adapted from Urry and Padovani 1995). Most of the common types of observed AGN can be represented on the above figure. The majority of the analysis in this work will focus on AGN observed from the ‘radio-quiet type 1’ position i.e. Seyfert 1 galaxies. Seyfert 2 galaxies are heavily obscured by a relatively high column density and are thought to be orientated edge-on to the observer i.e. observed through the torus. The arrows pertaining to the types of galaxy represent the viewing direction to the central SMBH.

(without a relativistic jet) and the BLR is visible, a standard Seyfert 1 galaxy would be observed; if the BLR is not visible and the AGN is heavily obscured, the viewing angle is edge-on to the torus and would be representative of a Seyfert 2 galaxy. In the radio-loud case, the AGNs are referred to as a broad-line radio galaxy (BLRG) or as a narrow-line radio galaxy (NLRG) based on whether the BLR is visible, respectively. Blazar type objects exhibit dramatic variability and are an orientation effect of viewing the AGN along the relativistic jet axis of radio-loud sources and is probably a result of relativistic beaming⁴.

One class of objects, at first inspection, does not fit nicely into the paradigm presented in Fig. 1.2, and maybe provides an indication that orientation does not fully determine the class of AGN. The objects in question form a class called Narrow

⁴Relativistic beaming is the process by which the apparent luminosity of a relativistic jet is modified due to special relativity.

Line Seyfert 1 (NLS1) galaxies. Despite the neat framework provided by the unified model, it is clear that the viewing angle does not account for all variations observed between different types of AGN and most notably the unified model does not explain the presence of relativistic jets. There are expected to be intrinsic differences between the AGN systems; in particular, parameters such as black hole mass, spin and accretion rate might be expected to generate differences between the classes, and could be responsible for the NLS1 objects. NLS1 galaxies, as their nomenclature suggest, are optically similar to standard Seyfert 1 galaxies, but the broad component of their permitted optical Balmer emission lines are narrow in comparison to other Seyfert 1 galaxies ($\text{FWHM H}\beta < 2000 \text{ km s}^{-1}$); moreover, NLS1s tend to exhibit strong FeII emission. The emission lines in NLS1 galaxies can be interpreted as broad lines from ‘high’-velocity clouds near smaller mass black holes, so consequently exhibit lower Keplerian velocities. Although, the emission lines could also be explained by a bigger BLR due to the high accretion rate often observed in NLS1s i.e. increased ionisation to a further distance from the accretion disk. NLS1 galaxies are excellent cases for X-ray variability studies, since they exhibit rapid X-ray variability on time-scales less than one day.

1.1.3 X-ray emission

X-ray emission from active galaxies is an extensive theme and I will not undertake an exhaustive review of it here. See Blumenthal and Gould (1970) and Rybicki and Lightman (1986) for a review of radiation processes.

AGN are observed over a broad range of energies and are often bright up to 100 keV or more, well above the prediction of the optical/UV model of peak emission predicted by accretion disk theory. In particular, the simple models of accretion predict inner disk temperatures of $\sim 10^5 \text{ K}$, which cannot reproduce the ‘hard’ X-ray emission⁵. Hard X-ray emission from AGN requires temperatures of at least 10^9 K , which cannot be accounted for with pure accretion disk thermal emission. The X-ray energy spectrum of AGN exhibits a power-law form, which could be explained by the presence of a relativistic jet by synchrotron processes; however, the energy spectrum is still observed in radio-quiet AGN with no apparent jet feature. Consequently there must be a different mechanism that produces the observed X-ray energy spectrum. It is thought that the source of X-rays is inverse Compton scattering of soft optical/UV *seed* photons through an optically thin *corona* of hot electrons

⁵Hard is often used to denote 2-10 keV and soft 0.3–3.5 keV. In general hard and soft are broadly used to indicate the high or low energy X-rays.

(e.g. Sunyaev and Titarchuk 1980); however, the exact location and characteristics of the corona are not clear, but it is thought to behave in a way akin to the Sun's corona (Haardt and Maraschi, 1993; Haardt *et al.*, 1997). The radiation physics of corona models is well understood, but there is significant uncertainty concerning the geometry and origin of the X-ray emitting corona. Following the strategy of coronal heating in the Sun, it is thought that the corona above accretion disks is heated to X-ray emitting temperatures ($\sim 10^9$ K) by magnetic reconnection (di Matteo, 1998). Corona models provide popular methods of X-ray emission in active galaxies, but there is still great uncertainty in these models and there is still much work to be done to constrain model parameters. Fig. 1.3 illustrates the differences between X-ray and optical observations. Despite the better spatial resolution of optical telescopes, X-ray observations provide better insight into the central engine of the AGN since X-rays are thought to be emitted from the hot region surrounding the SMBH.

X-ray energy spectra of Seyfert galaxies often exhibit a bump in their continuum peaking between 20–50 keV, thought to be caused by reflection off optically thick gas (Lightman and White, 1988). The reflection physics applies if the reflecting material is taken to be a 'cool' optically thick accretion disk or dusty torus (also optically thick) at large radii (Krolik *et al.*, 1994). Low-energy photons would be absorbed and high-energy photons (~ 100 keV) would suffer multiple reflections and be scattered to lower energies, since photons ~ 100 keV or greater lose a significant amount of energy in a single scattering event; thus, eliminating the low and high-energy extremes of the reflection spectrum, leaving only the 20–50 keV hump observed (Nandra and Pounds, 1994). Another prominent characteristic feature of X-ray energy spectrum is the strong fluorescent Fe K line. The line energy of 6.4 keV, for neutral Fe, rises slightly with increasing ionisation until the atom is completely ionised. The line energy is 6.7 keV and 6.9 keV in He-like Fe and H-like Fe, respectively (Nandra, 2006). Many Fe K lines are observed to be several tens of thousands of kms^{-1} broad and heavily redshifted (Tanaka *et al.*, 1995; Mushotzky *et al.*, 1995), suggesting that the Fe K line originates close to the innermost stable orbit deep inside the gravitational potential of the black hole. The observed Fe K line profiles are often interpreted as being gravitationally redshifted and Doppler broadened by orbital motion, and may provide an excellent diagnostic tool for probing the innermost accretion disk in a region of strong gravity (Miniutti *et al.*, 2003; Miniutti and Fabian, 2004). However, not all Fe K lines exhibit a broad red wing and are sometimes observed to be narrower than 1000 km s^{-1} , implying that the source of the Fe K line might be far from the central black hole (Weaver *et al.*, 1996).

Further to the previous section on AGN unification, Fe K lines also support that interpretation. Seyfert 2 galaxies display strong soft X-ray absorption with column densities often greater than 10^{23}cm^{-2} - consistent with an observation edge-on to the obscuring torus; moreover Seyfert 2 galaxies are noted by their lack of broad Fe K lines, compatible with production of the emission line in the surrounding torus and obscuration of the central source.

In high accretion rate AGN there is a contribution to the X-ray energy spectrum below ~ 1 keV called the *soft excess*. The peak energy of the soft excess is too high to represent thermal emission from a standard accretion disk model, and super-Eddington accretion disk models (e.g. Tanaka *et al.* 2005) cannot explain why the soft excess is always seen at the same energy in AGN X-ray energy spectra (Gierliński and Done, 2004). The unusually constant nature of the soft excess com-




Figure 1.3: The image above illustrates the difference between optical and X-ray observations. The optical image of NGC 1365, from the European Space Observatory’s Very Large Telescope, shows a typical active galaxy. The X-ray image, from *Chandra*, shows the X-ray bright central region. The emission is thought to be due to an accretion disk close to the central supermassive black hole. Credit: NASA/CXC/CfA/INAF/Risaliti Optical: ESO/VLT; <http://chandra.harvard.edu/photo/2007/ngc1365/>

ponent suggests that its origin might be based in atomic physics. Much work has been done on the nature of the soft excess with both reflection and absorption based models describing the observations well. Current X-ray observations in 0.3–10 keV cannot distinguish between the predictions of reflection and absorption models (Chevallier *et al.*, 2006; Sobolewska and Done, 2007), and RMS spectra are well described by either model (Ponti *et al.*, 2006; Gierliński and Done, 2004).

1.1.4 X-rays and temporal variability of AGN

X-ray light curves of AGN exhibit significant variability over a broad range of time-scales and it is one of their defining characteristics. AGN are variable in a seemingly undetermined and erratic fashion, and display X-ray light curves composed of aperiodic fluctuations. Variability of the above type is often referred to as *noise*, that is the variability is a result of a stochastic process so the intrinsic source signal is itself random and the output of a noise process. The stochastic equations governing the underlying noise process are assumed to be constant and each observation, no matter how different it looks, is merely a single representation of that process. The stochastic nature of AGN X-ray data means that interpretation of single realisations should be treated with caution unless rigorous statistical analysis is undertaken. Care should be taken not to confuse noise in this context with measurement errors such as Poisson noise. Temporal X-ray variability of AGN is a prime way to explore the innermost region around a SMBH, with high-frequency variations implying that the emitting region is less than one light day across. Fig. 1.4 shows a typical AGN short time-scale X-ray light curve as observed with *XMM-Newton* see Section 1.4.2).

A common way to quantify the observed variability structure is the *power spectral density* (PSD). The PSD represents the amount of variability as a function of Fourier frequency. The aforementioned aperiodic X-ray variability of AGN is further classified as *red-noise*, which is characterised by correlated variability (i.e. observed fluxes are not independent), such as flicker noise or a random walk, and the PSD exhibits a power-law dependence on frequency. White-noise in contrast does not exhibit correlated variability, which means that the PSD of white-noise is flat i.e. the signal contains equal power at each Fourier frequency (e.g. Poisson noise). The tools of variability studies will be covered in detail in the next chapter.

Temporal variability studies of AGN are complicated by the long time-scale trends often observed in X-ray light curves and the relatively poor signal-to-noise due to low count rates. Analysis of *EXOSAT* X-ray light curves of NGC 5506 by McHardy




Figure 1.4: 100s binned *XMM-Newton* light curve of MCG-6-30-15. The light curve is clearly variable on a variety of time-scales.

(1988) showed that AGN PSDs are well represented by a power-law over a broad range of frequencies ($P(\nu) = \nu^{-\alpha}$, where $P(\nu)$ is the power at a Fourier frequency ν), where the power spectral slopes are typically $\alpha = 1 - 2$. Moreover, the PSD of NGC 5506 did exhibit some evidence of the power-law model flattening at low frequencies, but no formal significance could be determined due to the poor quality of data. AGN PSDs must flatten in some way and cannot continue to exhibit steep power-law slopes to low frequencies, since the integrated power would diverge to infinity. Papadakis and McHardy (1995) used *EXOSAT* observations of NGC 4151 and statistically determined, using Monte Carlo methods, that the PSD does indeed *break* or flatten at a characteristic frequency. Edelson and Nandra (1999) used *Rossi X-ray Timing Explorer* (*RXTE*, see Section 1.4.1; Swank 1998) data to measure a PSD of NGC 3516 and determined a clear break in the observed power-law. Power spectral analysis of other Seyfert 1 galaxies with long-term *RXTE* monitoring campaigns have demonstrated that most AGN PSDs are often well fit by a power-law model that breaks at a characteristic frequency (Pounds *et al.*, 2001; Uttley *et al.*, 2002; Markowitz *et al.*, 2003; McHardy *et al.*, 2004, 2005; Uttley and McHardy, 2005).




Figure 1.5: Artist's impression of a low mass X-ray binary system. The stellar mass black hole accretes matter from the low mass donor star and an accretion disk is formed via Roche lobe overflow. Credit: http://www.jca.umbc.edu/press/binary_artist_impression.jpg

1.2 Black hole X-ray binary systems

Black holes are not only thought to be present in the centre of massive active galaxies⁶, but much less massive ($\sim 10 M_{\odot}$) black holes are thought to exist in Galactic X-ray binary systems (XRBs). XRBs can be extremely bright sources of X-rays, and their rapid variability properties are indicative of accretion onto a compact object. Many XRBs harbour black hole candidates (black hole X-ray binaries, BHXRBS) accreting matter from the donor companion stars. The black holes in the aforesaid systems can accrete matter via Roche lobe⁷ overflow from a low mass companion star (LMXB) or through strong stellar winds from a high mass companion star (HMXB). HMXBs are greater than $5 M_{\odot}$ and can also accrete matter via Roche lobe overflow. Note that the nomenclature of low and high mass refer to the mass of the companion star, not the compact object. Fig. 1.5 is an artist's impression of an LMXB accreting via Roche lobe overflow.

It should be clear that there are similarities between qualitative descriptions of AGN

⁶There is evidence to suggest that supermassive black holes might lurk in the centre of all galaxies (e.g. Kormendy and Richstone 1995; Schödel *et al.* 2002) and they may play an important role in galaxy evolution (Silk and Rees, 1998).

⁷The Roche lobe is a region around a star in a binary system within which the stellar material is gravitationally bound to the star.

and XRB accretion, with the only difference appearing to be the source of accreted matter. Perhaps it is still surprising to learn that despite an enormous difference in mass and X-ray luminosity the X-ray power spectra of BHXRBS (in particular the BHXRBS Cygnus X-1, often abbreviated Cyg X-1) and AGN are remarkably similar when scaled by mass and accretion rate (McHardy *et al.* 2004; Uttley and McHardy 2005; McHardy *et al.* 2006; and Section 1.3). I will thoroughly examine the BHXRBS-AGN connection in Chapter 4 and Chapter 5.

1.2.1 Properties of BHXRBS

X-ray binaries present an excellent opportunity to study the variability properties of accreting compact objects. Rapid X-ray variability and high X-ray flux allows XRBs to be studied with relative ease through short observations with X-ray telescopes, and *RXTE* is ideal for studying these objects. A full review of BHXRBS variability properties is far beyond the scope of this work, thus in this section I will briefly review some of the interesting properties of BHXRBS, with a focus on the temporal variability properties.

One of the earliest BHXRBS to be discovered was Cyg X-1, which I will often refer to throughout this work. Cyg X-1 is unusual amongst many of the known black hole XRB candidates in that it is persistently bright in X-rays. Cyg X-1 is classed as a HMXB, since it accretes from a stellar wind of a massive O/B type star and is one of three known *persistent* sources (McClintock and Remillard, 2004). Many of the ~ 20 established BHXRBS are LMXBs and exhibit nova-like behaviour, that is they are outbursting sources with periods of dramatically increased X-ray flux and are the so-called *transient* BHXRBS. One interesting transient source that is often referred to in this work is GRS 1915+105, and it is often classed as curious amongst the transient population in that it has remained bright for more than 15 years since its last outburst. Many transients will fade into *quiescence* after their outburst, that is the source enters a very-low flux rate state and may escape X-ray detection with the current generation of X-ray telescopes.

The power spectra of BHXRBS are often described in terms of two different types of variability. The first type, as previously mentioned for AGN PSDs, is aperiodic red-noise variability, which is represented by a power-law dependence on frequency and can be band limited - that is the PSD is cut-off at low and high frequencies. *Quasi-periodic oscillations* (QPOs) are the second type of variability observed and are marked by power concentrated at a much more limited range of time-scales.

Cyg X-1 has a well studied power spectrum and in the so-called high-flux state

or high/soft state, the power spectrum displays a power-law form that breaks at $\sim 10\text{--}20$ Hz from a slope of 1 at low frequencies to ~ 2 at high frequencies (Reig *et al.*, 2002; McHardy *et al.*, 2004). In the low/hard state the power spectrum exhibits two breaks in the observed power-law. The PSD breaks at ~ 0.2 Hz to a slope of 1 and breaks higher in frequency to a slope of 2 at ~ 3 Hz (Belloni and Hasinger, 1990). The low/hard state of Cyg X-1 is in fact better described by multiple Lorentzian components rather than a double-broken power-law (Nowak *et al.*, 1999; Pottschmidt *et al.*, 2003), but this will be considered later in this work. Cyg X-1 spends most of its time in the low/hard state and transitions between low/hard and high/soft on time-scales of \sim months; although, in general transitions can be much more rapid. Many BHXRBs (including Cyg X-1) exhibit other transition states, such as the *very-high* (or intermediate) state where the energy spectrum displays a mixture of low/hard and high/soft components, and the aforementioned quiescent state.

Quasi-periodic oscillations

QPOs are a common type of variability found in XRB power spectra and come in several flavours. The low-frequency QPO (LF QPO) is defined in the range ~ 50 mHz to ~ 30 Hz, whereas the high-frequency QPO (HF QPO) range is defined as above ~ 100 Hz. The physical interpretation of QPOs is still unclear, but the HF QPOs are suspected to be associated with the inner accretion disk and its characteristic frequency is perhaps a function of black hole mass (McClintock and Remillard, 2004). Both neutron star and black hole XRBs exhibit QPOs in the above frequency ranges; however, QPOs in known BHXRBs are not known to be faster than 100–450 Hz (detected in the very-high state). Thus the detection of kHz QPOs would rule out a black hole interpretation of the compact object (Van der Klis, 2004), since only neutron star XRBs display QPOs in the kHz domain (Strohmayer *et al.*, 1996),

Fig. 1.6 illustrates typical power spectra observed from the different BHXRB emission states, represented in P_ν and $\nu \times P_\nu$ space. The important states for consideration with this work are the high/soft state (HS) and the low/hard state (LS); it should be clear from Fig. 1.6 that both the HS and LS state can be roughly approximated by a single-break and a double-break power-law, respectively. It should also be clear that strong QPOs are observed in the transition states, namely the intermediate (IMS) and very-high state (VHS). Although, the intermediate and very-high states are thought to represent the same emission state. See Lewin and van der Klis (2006)

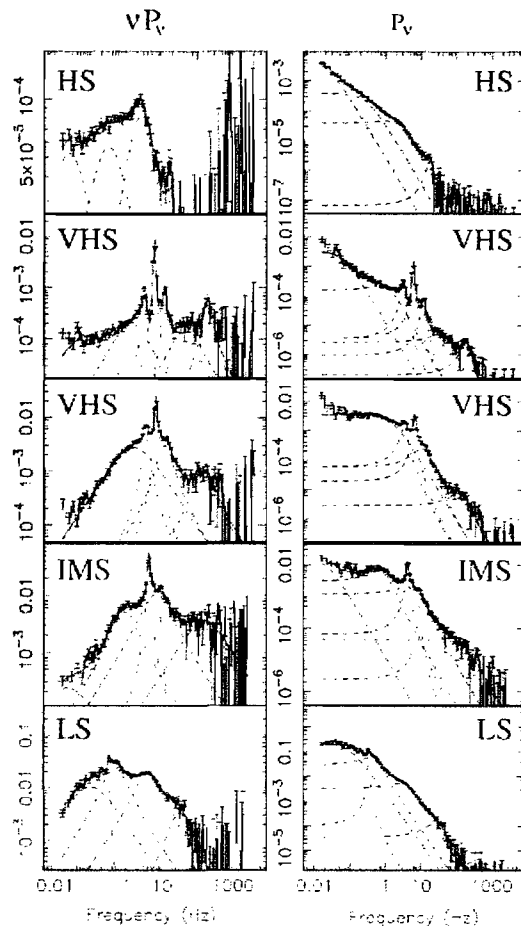


Figure 1.6: Power spectra of black holes in various emission states in P_ν and $\nu \times P_\nu$ space. Top to bottom, the plots represent the high/soft state, very-high state (with band limited noise; hard high luminosity intermediate state), very-high state (without band limited noise; soft high-luminosity intermediate state), intermediate state (a hard lower luminosity example), and the low/hard state. Credit: Van der Klis (2004).

and Klein-Wolt and van der Klis (2007) for a complete overview of BHXRBS. The black hole emission states are shown on the spectral hardness-intensity diagram (see Fig. 1.7), with the main states of interest denoted in bold face. Transient BHXRBS in outburst trace out the bold path through various emission states. The dashed jet-line indicates the transition from radio-quiet to radio-loud states, with the HS being radio-quiet (Fender *et al.*, 2004; Körding *et al.*, 2006b). As noted by Klein-Wolt and van der Klis (2007), during a hard flare in the HS the radio loudness increases significantly.

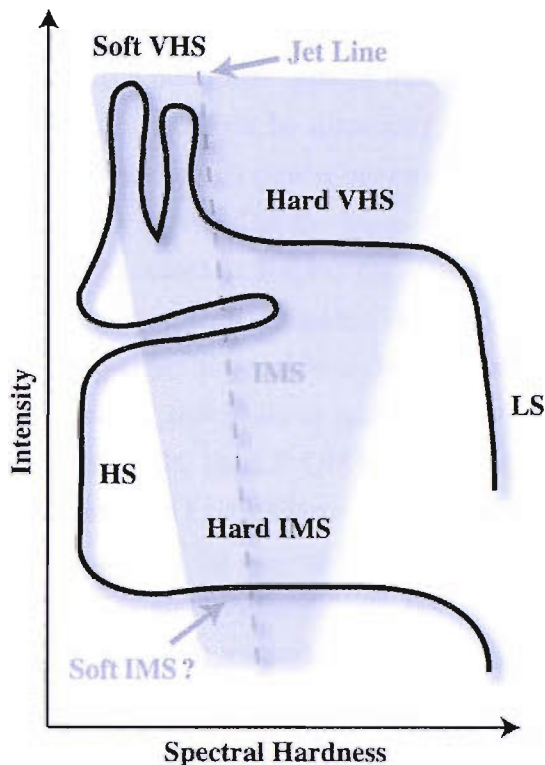


Figure 1.7: The hardness-intensity diagram for black hole states. The track shown illustrates a typical path followed by BHXRBs in outburst. Persistent BHXRBs such as Cyg X-1 can also follow parts of the track (Belloni *et al.*, 2005). Credit: Klein-Wolt and van der Klis (2007).

1.3 Comparing AGN and BHXRBs

For a number of years AGN have been compared to their BHXRB brethren, which are better understood phenomenologically and this comparison is ultimately based on the conjecture that accretion onto a stellar mass black hole ($\sim 10 M_{\odot}$) should be the same as accretion onto a SMBH ($\sim 10^8 M_{\odot}$). In the last few years there has been substantial evidence to support this hypothesis (Edelson and Nandra, 1999; Uttley *et al.*, 2002; Markowitz *et al.*, 2003; McHardy *et al.*, 2004, 2005; Uttley and McHardy, 2005; McHardy *et al.*, 2006; Summons *et al.*, 2007). However despite the obvious similarities between AGN and BHXRB, there have been no robust QPO detections in AGN PSDs⁸ (see Vaughan and Uttley 2006). There is no known reason, theoretical or otherwise, that would prohibit AGN QPOs. The two types of QPO detected in BHXRBs might be expected to be found in AGN power spectra, albeit at a somewhat lower frequency. HF QPOs are rare in BHXRBs, with less than 1/3 of sources exhibiting QPOs of this type (Remillard *et al.*, 2003). The ma-

⁸Previous QPO detections by Papadakis and Lawrence (1993b) (NGC 5548) and Iwasawa *et al.* (1998) (IRAS 18325-5926) have been ruled out.

majority of AGN PSDs presented in the literature to date are ‘soft’⁹, and given that HF QPOs seem to be confined to a particular type of the very-high state a detection in the current sample of AGN PSDs might be surprising. It is also worth mentioning that QPO detection in broadband AGN power spectra would be acutely dependent on the number of QPO cycles sampled in the time-series. Even though LF QPOs are stronger than HF QPOs, to detect a LF QPO the observed light curve must cover more than 100 QPO cycles and be well sampled (Vaughan and Uttley, 2005). Most AGN X-ray monitoring campaigns are unable to dedicate enough regular on-source time to fulfil these requirements; moreover as with the previous argument concerning emission states and HF QPOs, the LF QPO is also not detected in the HS of BHXRBs and so the detection of a LF QPO in a known ‘soft’ state AGN would also be surprising. It is unfortunate that the well observed Seyfert galaxies that form the majority of known AGN PSDs are consistent with HS emission. As suggested by Jester (2005) and Nagar *et al.* (2005), radio-loud low-luminosity AGN would make the best candidates for LS AGN, and Vaughan and Uttley (2005) also suggest that these types of AGN would make the best candidates if LF QPOs were to be actively sought out.

A final remark regarding the observed X-ray count rates of AGN and BHXRBs. It might appear as if all the timing properties of AGN are poorer than those of BHXRBs, since BHXRBs are brighter and their time-scales more easily probed. However, analysis of very-high frequency AGN PSDs may in fact prove superior to that of BHXRB PSD analysis. Consider, a Seyfert galaxy is typically $\sim 10^6$ times more massive than an BHXRB, thus the time-scales of variability in the Seyfert galaxy are expected to be $\sim 10^{-6}$ times slower. The nearest bright BHXRB is a factor of $\sim 10^4$ pc closer than the nearest bright Seyfert galaxy. Consequently the flux of a typical Seyfert galaxy is $\sim 10^{-2}$ to 10^{-3} times that of a typical BHXRB; however given that variations occur on much longer time-scales in the AGN, a Seyfert galaxy would provide $\sim 10^2$ to 10^4 times as much flux per unit light crossing time than in BHXRBs. A consequence of the above result is that AGN PSDs are often far better determined than BHXRB PSDs (in relative terms) at the highest frequencies. The importance of a thorough and consistent power spectral analysis of AGN data cannot be understated. Thus far, AGN timing analysis has been inconsistent and biased towards to brighter radio-quiet AGN. In order to provide a proper comparison between AGN and BHXRBs, analysis of aforesaid objects must be consistent and rigorous. In this work I will analyse a sample of AGN of varying masses, accretion

⁹AGN data that is well fit by a power-law with a single characteristic break-frequency are often referred to as ‘soft’.

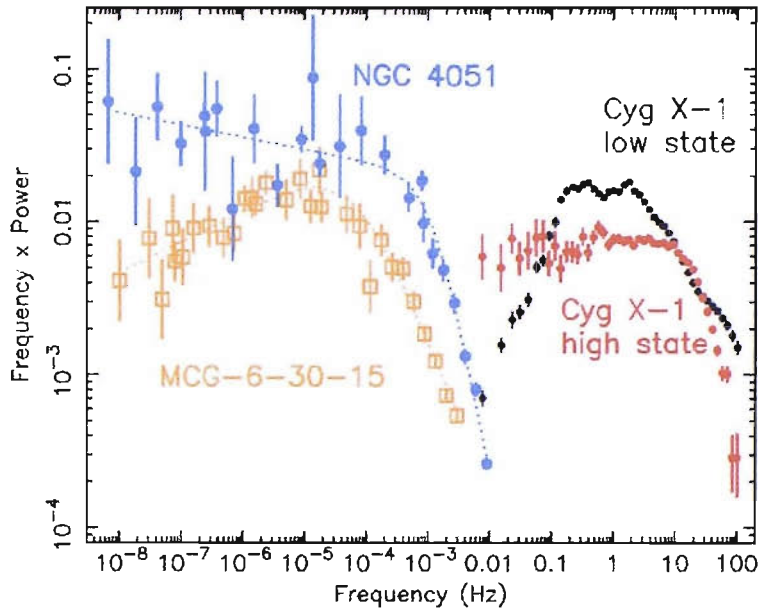


Figure 1.8: Here I show the PSDs of two AGN (NGC 4051 and MCG-6-30-15) and the PSDs of Cyg X-1 in the LS and HS. The PSDs of NGC 4051 and MCG-6-30-15 look like PSD of Cyg X-1 in the HS scaled by mass and accretion rate (this scaling is shown in Chapter 5). Credit: Ian McHardy.

rates and radio-loudness both rigorously and consistently to determine if there does indeed exist a similarity between AGN and BHXRBs. Similarities have already been suggested by McHardy *et al.* (2006) and in Fig. 1.8 I show a comparison between the PSDs of two AGN (NGC 4051 and MCG-6-30-15) and the PSDs of Cyg X-1 in the LS and HS. It is clear that the PSDs of NGC 4051 and MCG-6-30-15 are suggestively close to the PSD of Cyg X-1 in the HS. It is important to determine if the relationship between BHXRBs and AGN is real, which can be resolved by thorough unbiased modelling of AGN PSDs.

1.4 AGN X-ray monitoring

Long-term X-ray monitoring of AGN poses an interesting logistical problem. The purpose of AGN monitoring campaigns is to sample all time-scales of variability with the minimum amount of on-source time. Regular observations (once or twice a week) would be required to constrain the long time-scale variability properties of AGN, which could last tens of years or more. Fortunately, on 30 December 1995 the *Rossi X-ray Timing Explorer (RXTE)* was launched and *RXTE* would, for the first time in X-ray astronomy, provide a realistic chance of producing high quality long-term X-ray light curves. Almost four years later the *X-ray Multi-Mirror Mission* observatory (*XMM* and later called *XMM-Newton*) was launched. *XMM-Newton*

is a tremendously sensitive X-ray telescope capable of continuously measuring the short time-scale variability of AGN. As I will describe in the forthcoming section, *RXTE* and *XMM-Newton* work well together to measure and constrain the broad-band power spectrum of active galaxies.

1.4.1 Rossi X-ray Timing Explorer (*RXTE*)

There are three scientific instruments aboard *RXTE* specifically designed to carry out timing studies of X-ray variable sources on time-scales of \sim microseconds to years. *RXTE* has revolutionised X-ray timing due to its remarkable timing resolution, but it is the rapid pointing capability of *RXTE* that has made this work possible. *RXTE* is in a low-Earth orbit (period \sim 90 minutes) and is able to slew between X-ray sources extremely rapidly allowing multiple pointings to be taken at regular intervals - the exact property required to construct long-term variability light curves of AGN.

Fig. 1.9 shows a schematic of the *RXTE* spacecraft along with the three timing instruments: the All-Sky Monitor (ASM), the High-Energy X-ray Timing Experiment (HEXTE) and most importantly for the work herein the Proportional Counter Array (PCA, Zhang *et al.*, 1993).

The PCA aboard *RXTE* carries out microsecond time resolution X-ray astronomy, sensitive in the energy range 2–60 keV. The PCA consists of five proportional counter units (PCUs, labelled from 0 to 4) collecting data simultaneously. The total collecting area of the PCA is 6.5 m^2 , with the field of view set by the FWHM of the collimator or \sim 1 degree. *RXTE* is not an imaging satellite, thus an observation merely consists of ‘counting’ photons that pass through the collimators. The true power of the PCA is due to high event rate throughput, whereas imaging satellites are often limited as timing instruments by telemetry constraints and saturation of pixels (pile-up) for bright sources. The typical count rates per PCU for a moderately bright Seyfert 1 is \sim 2–3 cts s^{-1} , but can be as high as \sim 20–30 cts s^{-1} for very bright sources. The *RXTE* PCA is well understood and well described (e.g. Swank 1998; Jahoda *et al.* 2006).

As of the time of writing this thesis, four of the five PCUs in the PCA have suffered repeated discharge problems effectively limiting the amount of good-time. Only PCU 2 is operating at 100 per cent good-time i.e. it does not currently suffer discharge problems. PCU 0 and PCU 1 have both suffered a pinhole leak in their propane layer, consistent with a micrometeorite hit on 12 May 2000 and 25 December 2006, respectively. The effect of the missing propane layer is to dramatically




Figure 1.9: A schematic of the *Rossi X-ray Timing Explorer (RXTE)* spacecraft. The main scientific instrument used through this work, the PCA, is clearly labelled. Credit: http://mamacass.ucsd.edu/hexite/pictures/xte_spacecraft_line.gif

increase the background for faint sources, such as AGN, and for this reason I do not use PCUs 0 and 1 after the aforementioned dates of micrometeorite impact. Despite these problems with the PCA, AGN monitoring can still be undertaken and good count rates obtained. Owing to the fact that *RXTE* is not an imaging satellite, the background has to be simulated based on satellite information e.g. time since *South Atlantic Anomaly* (SAA) passage and satellite positional information. In the twelve years since the launch of *RXTE* the background models have become very sophisticated and blank sky pointings show they are more than adequate to correctly model the background.

Both of the remaining two scientific instruments, HEXTE and the ASM, are not used in this work. HEXTE and the ASM produce very low count rates even for the very brightest AGN due to their vastly reduced effective area in comparison to the PCA. The HEXTE covers an impressive energy range 15-250 keV and may be useful for spectral studies with a long exposure. The ASM covers 2-10 keV,

but merely lacks the sensitivity to detect all but the brightest AGN. For the above reason, only the PCA is used for AGN timing studies found in this work.

1.4.2 *XMM-Newton*

XMM-Newton has three main scientific instruments designed to undertake a variety of tasks. There are three X-ray CCD cameras comprising the *European Photon Imaging Camera* (EPIC). Two of the three EPIC instruments are *metal oxide semiconductor* CCDs (MOS) (Turner *et al.*, 2001). The MOS cameras are behind the X-ray telescopes that are equipped with *Reflection Grating Spectrometers* (RGS) and gratings direct ~ 50 per cent of the incident flux to each MOS camera and RGS instrument. The final X-ray telescope is unobstructed and has the PN camera at its focus and, as the name suggests, uses pn¹⁰ CCDs (Strüder *et al.*, 2001). The second main instrument, as I have already introduced, is the RGS. The third instrument aboard *XMM-Newton* is the *Optical Monitor* (OM) (Mason *et al.*, 2001), which is co-aligned with the X-ray telescope to provide optical and UV coverage. A schematic of *XMM-Newton* is presented in Fig. 1.10. *XMM-Newton* covers an energy range of 0.1 – 12 keV with the best response at softer energies. The effective area of *XMM-Newton* is 4.3 m² at 1.5 keV and 1.8 m² at 8 keV. *XMM-Newton* was placed in a 48-hour elliptical orbit around the Earth. An *XMM-Newton* light curve was previously shown in Fig. 1.4.

The sensitive optics and highly elliptical orbit of *XMM-Newton* allow well-sampled high resolution timing analysis of faint active galaxies to be undertaken. While *RXTE* is formally a better timing satellite, the highest time resolution of *XMM-Newton* is more than sufficient to monitor the fastest time-scales of variability often observed in AGN, but the real strength of *XMM-Newton* is the amount of continuous on-source time available. *RXTE* will be able to monitor the high-frequency variability of AGN with greater precision, but given its low-Earth orbit the amount of on-source time is often very limited and largely forbids extended continuous observations of order one day. However, *XMM-Newton* is able to continuously measure the X-ray variability of AGN for ~ 36 hours. In this regard the short to medium time-scales as measured by *XMM-Newton* are unparalleled by *RXTE*. Given the superior monitoring capabilities of *RXTE* and the excellent short time-scale continuous measurement proficiency of *XMM-Newton*, a monitoring campaign that utilises the strengths of both these satellites will provide the best broadband power spectral analysis currently available at X-ray energies.

¹⁰pn is a type of semiconductor




Figure 1.10: A schematic of the *XMM-Newton* spacecraft. Credit: <http://www.astro.phys.ethz.ch/staff/guedel/private/xmm/pictures1.html>

1.5 Summary of thesis

The aim of the work presented herein is to provide a comprehensive study of AGN power spectra. I use archival and proprietary *RXTE* data spanning 12 years, along with shorter time-scale observations from *XMM-Newton* to provide robust estimates of broadband AGN power spectral shapes for a large sample of sources.

In the next chapter (Chapter 2), I will further progress some of the terms introduced in this chapter such as the power spectrum and describe them in a more formal framework. I establish several other definitions and techniques that are inherently associated with time series analysis. I describe the method by which I simulate red-noise light curves, which will be of great importance in later chapters, and use the simulated data to illustrate the disagreeable effects of red-noise leak and aliasing. The effect of red-noise leak, in particular, is examined in detail due to its significance in later chapters. Chapter 2 will be concluded by a discussion of the assumption of stationarity.

Chapter 3 will focus on our Monte Carlo software PSRESP, which is a technique based on the ‘response’ method of Done *et al.* (1992). I have further developed PSRESP from its original implementation (Uttley *et al.*, 2002) and introduce several

new features which represent the first computational result of this work. The most important new features in PSRESP are the ability to simulate the high-frequency power spectrum and correctly determine the error bars in the presence of Poisson noise, but I also implement a more robust acceptance probability determination in the presence of extreme sampling distortions. I review the new implementation of PSRESP by testing the rejection probability levels in the presence of extreme and minimal red-noise leak. I conclude Chapter 3 by discussing the possible optimisation techniques, but also the limitations of the current implementation.

In Chapter 4, I apply PSRESP to determine the power spectral shape of two AGN. I examine the power spectrum of NGC 3783 by testing it against a variety of underlying PSD models. Markowitz *et al.* (2003) provide tentative evidence that the power spectrum of NGC 3783 contains a second lower frequency break in its power spectrum, which is possibly the tell-tale sign of LS emission. I re-evaluate the reliability of a second lower frequency break in the power spectrum of NGC 3783 by using an improved PSRESP implementation and proprietary *RXTE* data. I also discuss the possibility of a QPO feature in NGC 3783. The second half of Chapter 4 is dedicated to Ark 564, which is known to exhibit a power spectrum with a low and high-frequency break. I not only fit the broadband power spectrum of Ark 564 in a soft and hard energy range with a double-bend power-law, but I also fit the power spectra with a two-component Lorentzian model - a *first* in AGN timing analysis. I conclude Chapter 4 by discussing the evidence for a two-component model and the possible emission state of this unique object.

In Chapter 5, I apply the improved PSRESP technique to 32 AGN to provide a comprehensive survey of AGN power spectra. All data in Chapter 5 are gathered, reduced and analysed within an entirely consistent framework. The results represent the first power spectral survey of this type and are used to reanalyse the BHXRB-AGN correspondence reported in McHardy *et al.* (2006).

In the final chapter (Chapter 6) I review the results reported in this work and discuss the possible avenues of research this work could follow, as well as the future of PSRESP. I conclude Chapter 6 by describing the implications of this work with regard to future missions and AGN emission states in particular.

In the beginning the Universe was created.
This has made a lot of people very angry
and been widely regarded as a bad move.

DOUGLAS ADAMS (1952 - 2001), THE HITCH HIKER'S
GUIDE TO THE GALAXY

2

Measuring temporal variability

In this chapter I will describe how the temporal X-ray variability amplitude of active galactic nuclei is quantified and studied, with particular attention on the power spectral density (PSD). The ‘red-noise’ light curves typical of AGN show apparently random and supposedly unpredictable aperiodic variability, which is often referred to as *noise* since it is the result of a stochastic process. I will describe various statistical quantities and techniques often employed to quantify the variability amplitude of AGN red-noise light curves, namely the fractional RMS amplitude and the power spectral density (PSD), respectively. The PSD describes how the power of a time series is distributed with frequency. I demonstrate how to reliably simulate light curves, given an underlying power-spectral model, and determine the biasing effect of a sampling pattern on the power spectrum. The effect of red-noise leak and aliasing will be covered in detail and I shall present the first result of this work, the effect of red-noise leak on the measured PSD and the consequence for hypothesis testing. I will finally describe key assumptions in generating a broadband PSD, namely the stationarity property of AGN PSDs.

2.1 Variability and the power spectral density

Throughout this work I have utilised various statistical tools with which to determine quantitative information about the highly variable temporal structures often encoded in the variability pattern of AGN time series data. Accurate determination of the aforementioned temporal structure is crucial in making robust comparisons to physical models and in hypotheses testing. Here I will briefly comment on the various techniques and methodologies often encountered throughout this work.

It should be noted that due to the confusions of power spectrum and periodogram in the literature I will be following the conventions of Priestley (1981). I will use power spectrum to describe the true underlying process and periodogram to represent a realisation of the process. The periodogram generated from a single stochastic realisation would look different to the periodogram obtained from averaging together a large finite number of realisations or observations.

2.1.1 The importance of variance

Consider a discretely sampled time series, x_i , with measurement errors, $\sigma_{err,i}$, such as Poisson noise in the case of X-ray photon statistics. The uncertainties in the light curve count rate would contribute additional scatter or variance (Vaughan *et al.*, 2003) and overestimate the intrinsic source variance. The intrinsic source variance can be estimated by measuring the ‘excess variance’, which is given by the sample variance minus the contribution to the variance expected from the Poisson noise, that is the noise or fluctuations due to Poisson statistics. The sample variance, S^2 , mean-square-error, $\bar{\sigma}_{err}^2$, and excess variance, σ_{XS}^2 , are given by

$$S^2 = \frac{1}{N-1} \sum_{i=1}^N (x(t_i) - \mu)^2 \quad (2.1)$$

$$\bar{\sigma}_{err}^2 = \frac{1}{N} \sum_{i=1}^N \sigma_{err,i}^2 \quad (2.2)$$

$$\sigma_{XS}^2 = S^2 - \bar{\sigma}_{err}^2 \quad (2.3)$$

The quantities $\bar{\sigma}_{err}^2$ and μ (average count rate) are time averaged over the entire signal. The square root of Eqn.2.3 is also known as the ‘root-mean-square (RMS) amplitude’ and is sometimes written as σ_{RMS} . Normalising the RMS amplitude by the mean count rate of the source allows the intrinsic source variability to be compared to other sources with a different brightness. If two X-ray sources exhibit

the same variability structure, but have different mean flux values then the lower flux source will display a reduced RMS amplitude. Normalising the X-ray light curves by their mean X-ray flux allows direct comparison between their variability amplitudes¹. The resultant quantity is called the fractional RMS amplitude (Edelson *et al.*, 1990; Rodriguez-Pascual *et al.*, 1997)

$$F_{\text{var}} = \sqrt{\frac{S^2 - \bar{\sigma}_{\text{err}}^2}{\mu^2}} \quad (2.4)$$

The fractional RMS amplitude is often preferred over the normalised excess variance, $\sigma_{\text{XS}}^2/\mu^2$, since F_{var} is a linear statistic and can be represented in percentage terms, even though they communicate identical information. Another statistical test that will be often encountered in this work is the χ^2 test. The χ^2 statistic will be used as a goodness-of-fit throughout this work and is a comparison between the observed data distribution, D_i , and the theoretical data distribution, M_i , with uncertainties specified by σ_i

$$\chi^2 = \sum_{i=1}^N \left(\frac{D_i - M_i}{\sigma_i} \right)^2 \quad (2.5)$$

The number of observed data points minus the number of parameters in the theoretical data distribution is known as the degrees of freedom, ν . Given the degrees of freedom, the reduced χ^2 : $\chi^2_{\nu} = \chi^2/\nu$ can be determined. An acceptable reduced χ^2 should be approximately equal to unity. Use of the χ^2 statistic will be explored in greater detail in Section 3.2.

2.1.2 Periodogram estimates of the power spectrum

From a light curve or time series the periodogram can be evaluated by taking the modulus squared of the Fourier transform of the mean-subtracted light curve. Removal of the mean light curve count rate ensures the periodogram does not contain a zero-frequency power component, that is, the power contribution due to a constant offset in the light curve. The Fourier transform is formally a mathematical process which uses integral calculus on a continuously varying function; however, real time series data is discretely sampled and finite in length. That is, the resulting observation will be a time series of data length T , sampled at N discrete times, t_j . The discrete form of the Fourier transform (Deeming, 1975; Oppenheim and

¹This is mathematically equivalent to normalising the periodogram by the mean flux squared

Schafer, 1975; Brillinger and Krishnaiah, 1983) uses summations over finite limits rather than the infinite limits associated with the integral form of the Fourier transform. Moreover, the discrete Fourier transform (DFT) can be applied to data of both even and uneven sampling with equal merit, and is defined for data with arbitrary spacing. The discrete Fourier transform, $F_N(\nu)$, as given by Deeming (1975), is as follows

$$F_N(\nu) = \sum_{j=1}^N f(t_j) e^{i2\pi\nu t_j} \quad (2.6)$$

The discrete Fourier transform is generated over a range of $N/2$ discrete frequencies $\{F_N(\nu_1), F_N(\nu_2), \dots\}$, and occurs at evenly spaced intervals of frequency from ν_{\min} to ν_{Nyquist} . The lowest obtainable frequency, ν_{\min} , is determined by the length of observation i.e. $1/T$. The highest obtainable frequency or ‘Nyquist frequency’, ν_{Nyquist} , is derived from the reciprocal of twice the sampling interval of the light curve i.e. $1/2\Delta T$. Here I note a fundamental result of information theory: in order to recover all the Fourier components of a periodic waveform, it is necessary to use a sampling interval at least twice the highest waveform frequency in the signal. This is the so-called ‘Nyquist-Shannon’ sampling theorem (Shannon, 1949). The discrete Fourier transform, given in Eqn.2.6, involves the imaginary unit, i , so it is convenient to use Euler’s formula and write the discrete Fourier transform in terms of trigonometric functions for computation. The squared modulus of the discrete Fourier transform is thus given by

$$|F_N(\nu)|^2 = \left| \sum_{j=1}^N f(t_j) e^{i2\pi\nu t_j} \right|^2 \quad (2.7)$$

$$= \left(\sum_{j=1}^N f(t_j) \cos(2\pi\nu t_j) \right)^2 + \left(\sum_{j=1}^N f(t_j) \sin(2\pi\nu t_j) \right)^2 \quad (2.8)$$

The form of periodogram used throughout this work is determined by applying a normalisation of $2T/\mu^2 N^2$, of units $(\text{RMS}/\text{mean})^2 \text{ Hz}^{-1}$. Thus the formula used to generate the periodograms seen throughout this work is given by

$$P(\nu) = \frac{2T}{\mu^2 N^2} |F_N(\nu)|^2 \quad (2.9)$$

The resultant periodogram is normalised by the square of the mean flux of the original light curve, μ^2 , so that periodogram amplitudes can be directly compared.

The so-called ‘fractional root mean square’ (fractional RMS-squared) normalisation is often very practical since the integrated periodogram over Fourier frequency is equal to the fractional variance, (σ^2/μ^2) , of the light curve (van der Klis, 1997) i.e.

$$\int_{\nu_{\min}}^{\nu_{\text{Nyquist}}} P(\nu) d\nu = \frac{\sigma^2}{\mu^2} \quad (2.10)$$

The aforementioned fractional RMS-squared normalisation is a practical choice of normalisation for BHXRB and AGN data; however, other choices of normalisation exist and are often very useful.

Rather than normalising the periodogram by μ^2 , the periodogram can be given in absolute units

$$\text{Normalisation} = \frac{2T}{N^2}, \quad (2.11)$$

so that the integrated periodogram returns the total variance in absolute units. This normalisation is seldom used in AGN timing analysis, since the resultant power spectrum would depend on the source flux, and it is often convenient to compare AGN variability independent of flux.

The Leahy normalisation (Leahy *et al.*, 1983), is given by

$$\text{Normalisation} = \frac{2}{n}, \quad (2.12)$$

where n is the total counts measured in the light curve. The Leahy normalisation is useful when searching for a periodic signal in white-noise, since the expected Poisson noise level is merely 2 in the periodogram. A pulsar pulse period should be easily identifiable against the expected Poisson noise level when using the Leahy normalisation. The significant advantage of using the Leahy normalisation is nullified when examining AGN periodogram data, since the light curves of AGN are red-noise dominated. Thus any test of a strong periodic signal or otherwise must be made against the red-noise dominated continuum in the periodogram of the source rather than the Poisson noise level (Vaughan, 2005).

Given the red-noise nature of AGN light curves and the convenience of relating integrated power of the periodogram with light curve variance, I will be using the fractional RMS-squared normalisation throughout this work. Within the framework of the fractional RMS-squared normalisation the effect of Poisson noise is to add an approximately constant amount of *mean* power to the periodogram at all frequencies. In the absence of distortions (detector or otherwise) the predicted power

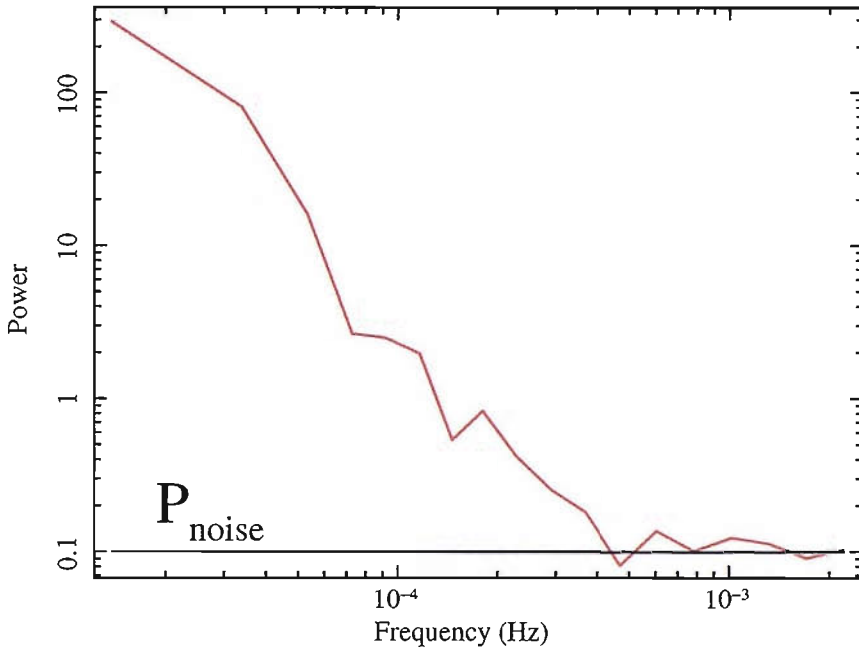


Figure 2.1: The Poisson noise component is observed as a nearly constant power contribution that flattens the power-law slope to 0 at high frequencies. The solid red line represents the observed periodogram effected by the additional Poisson noise component. The solid black line represents the approximate power in the Poisson noise level.

contribution of the Poisson noise level is

$$P_{\text{noise}} = \frac{2(\overline{\text{cts}} + \overline{\text{bgd}})}{\overline{\text{cts}}^2}, \quad (2.13)$$

where $\overline{\text{cts}}$ and $\overline{\text{bgd}}$ are the average source and background count rates, respectively. In the limit where the background drops to zero and the observation is contiguous then $P_{\text{noise}} = 2/\overline{\text{cts}}$ - this is often the case for *XMM-Newton* observations of X-ray bright AGN. The factor of 2 originates from requiring the correct variance when integrating over only the positive frequencies in the observed periodogram. If the effect of aliasing is included for binned light curve, the above equation becomes

$$P_{\text{noise}} = \frac{2T(\overline{\text{cts}} + \overline{\text{bgd}})}{N\Delta T_{\text{bin}}\overline{\text{cts}}^2}, \quad (2.14)$$

where ΔT_{bin} is the time bin width. The factor of $T/N\Delta T_{\text{bin}}$ accounts for the effect of aliasing (see Section 2.3.1) in the light curve if it contained gaps. Fig 2.1 illustrates the effect of a Poisson noise level on a periodogram at high frequencies where the signal to noise is often reduced, this feature will be particularly relevant when using *XMM-Newton* observations.

Properties of the periodogram

The X-ray light curves of AGN and BHXRBs are stochastic with no two being exactly alike. That is, a single observed light curve is apparently erratic and unpredictable but is in fact a realisation from a distribution of possible light curves. In a noise process, the observed periodogram values are randomly distributed about the true power spectrum describing the underlying noise process, following a χ^2_2 distribution² (van der Klis, 1989) with the standard deviation equal to the power. The periodogram is however inconsistent when employed for real data, that is, as the number of points in the light curve increases, the scatter in the measured periodogram does not decrease (Jenkins and Watts, 1969). The periodogram is required to be averaged or binned in some way, so that as the number of points per frequency bin increase, the observed scatter in the periodogram decreases - as is required of a consistent estimator of the PSD. The requirement of a binned periodogram, in order to make both consistent and reliable estimates of power spectrum, holds true not only for linear stochastic processes but also non-linear stochastic processes.

In order to obtain consistent periodogram estimates for BHXRB observations, the necessary time-scales can be sufficiently oversampled so that the total light curve can be partitioned into several shorter light curve segments of the same length. Calculating the periodogram of each light curve segment yields a distribution of powers at each Fourier frequency, from which a mean power and standard deviation can be determined and the ensemble of power spectrum estimates binned to form a periodogram, and the distribution of mean power estimates will approach a Gaussian distribution (by the central limit theorem). The method of determining the periodogram by oversampling³ requires the total light curve, prior to partitioning, to contain many cycles of the lowest useful time-scale, which is often true in BHXRB data. In AGN data the longest time-scales may not be sampled by even the longest of *RXTE* monitoring campaigns, and as a result there will be insufficient cycles with which to produce a reliable periodogram at the lowest time-scales.

Papadakis and Lawrence (1993a) suggest an alternative method for periodogram binning which mitigates the issues associated with low-frequency mean power determination, where the *bias* due to a low number of cycles can dominate. Papadakis and Lawrence (1993a) propose that by binning the periodogram of a red-noise process in logarithmic power space a Gaussian distribution of binned power is obtained

²A χ^2 distribution following 2 degrees of freedom.

³Oversampling in this context is taken to mean multiple measurements or samples of the longest useful time-scale in the time-domain.

significantly faster than the aforesaid method utilising light curve segments in linear power space. Logarithmic binning of power inherently reduces the distortion of low-frequency power variations. Furthermore by binning the periodogram logarithmically in frequency, a greater number of data points are used with which to determine a more robust power spectrum estimate. However despite the obvious advantage of binning the periodogram in logarithmic frequency-power space, the lowest measured frequencies in the periodogram will be combined with neighbouring points with which to produce an accurate power-spectral estimate, effectively distorting and band-limiting the overall periodogram at these frequencies. This last point can be understood in terms of the simulated periodograms exhibited in Fig. 2.2; the lowest frequencies of the raw periodogram only have a few points from which to determine a PSD estimate, and if they are binned with neighbouring points then the overall binned periodogram will not extend to as low frequencies and band-limit the observation. Care must be taken when attempting to logarithmically bin the periodogram, since the average value of the logarithm of the periodogram is *not* the same as taking the average of logarithmic power. The difference originates from the shape of χ^2_2 distribution in logarithmic space and is a constant bias, which can be removed by adding 0.25068 (Vaughan, 2005) to the binned periodogram value. Determination of the periodogram using Eqn.2.9 is a relatively straightforward task; however, astronomical observations are subject to a plethora of constraints, are rarely continuous, and produce irregular sampling patterns. Much of the data encountered in this work is not evenly sampled and is, at best, continuous for short periods of time. The irregularity of the sampling pattern seen in most observations of AGN can have a dramatic effect and produces significant biases or distortions on the resulting periodogram. The distortions on the periodogram are merely a function of sampling pattern and if these distortions are not determined and understood, the periodogram will be a poor estimator of the power spectrum. Determining a robust estimate of the power spectrum with reliable error estimates, that takes into account the aforementioned distortions due to sampling is crucial when attempting to establish the power spectral shape. In Chapter 3 I will further discuss how a robust estimate of the power spectrum is determined using our Monte Carlo method. Before I describe how our Monte Carlo method determines the periodogram, I shall discuss in greater detail how we use simulated data to probe the problems that arise due to sampling when measuring the periodogram.




Figure 2.2: Periodograms simulated using the method of Timmer and König (1995). The parameter α corresponds to the power-law slope, and ν_B corresponds to the frequency at which the power-law breaks. Plots (a) and (b) were generated from unbroken power-law models of $\alpha = 1$ and $\alpha = 2$, respectively. In plot (c) the power-law has a slope of $\alpha = 2$, which breaks to a slope of zero below the break-frequency, $\nu_B = 10^{-4}$ Hz.

2.2 Light curve simulation

Power spectrum estimates of a noise process will inherently exhibit a large scatter of 100 per cent about the true underlying power spectrum, and the real and imaginary parts of the discrete Fourier transform are normally distributed for stochastic processes. These features of a stochastic process can be exploited to produce simulated light curves with the same statistical properties as the observed data. The method described by Timmer and König (1995) is used throughout this work to simulate red-noise light curves.

The method of Timmer and König (1995) produces simulated light curves in the following way. Choose an underlying power spectral model and normalisation (e.g. a power-law slope for X-ray AGN data), and at each Fourier frequency draw *two* random deviates scattered according to a Gaussian distribution; multiply these numbers by the square root of half the power spectral model value at that Fourier frequency, and repeat over the required frequency domain. The resultant two values at each Fourier frequency represent the real and imaginary parts of the Fourier transform of the simulated data. In order to produce a real-valued time series, the Fourier com-




Figure 2.3: The periodograms in Fig. 2.2 produce the light curves shown above. The choice of power spectral parameters were chosen to exhibit how the power at different Fourier frequencies manifest in the resultant light curve. Plot (a) exhibits weak trends, whereas plot (c) displays strong trends on time-scales of $\sim 10,000$ s, which correspond to the input break frequency. Plot (b) was generated from a steep unbroken power-law to low frequencies and as a result shows a strong trend on long time-scales. Note that an arbitrary constant were added to the simulated light curves.

ponents at negative frequencies are merely the complex conjugate of the positive frequencies. The inverse Fourier transform of the simulated data yields a real light curve ($\Im = 0$ i.e. the imaginary component is zero) drawn from a distribution of possible light curves, given an underlying power-spectral model. Simulating a light curve in this way produces periodogram values that follow a χ_2^2 distribution about the underlying power spectrum, as required for a noise process. Given the freedom to choose the number of points in the resulting time series, it is often advantageous to simulate light curves of length $N = 2^m$ (where $m \in \mathbb{Z}^+$ i.e. where m is a positive integer), so that the fast Fourier transform (FFT) (e.g. Press *et al.* 1992) may be used. An example of three possible power-spectral models are shown in Fig. 2.2, and the light curves generated from these models are shown in Fig. 2.3. Each simulation was initiated with the same seed for direct comparison.

Previous attempts to simulate AGN light curves have been deemed inadequate. A method suggested by Done *et al.* (1992) constructs a light curve from a power spectral model by summation of sine waves with a random frequency distribution; however, the amplitudes are determined, which is unlike a real noise process.

2.3 Sampling biases in AGN power spectra

X-ray observations of AGN are subject to various observational constraints, as already mentioned in Chapter 1, which produce biases or distortions in the measured periodogram. Simulated light curves, like those in Fig. 2.3, represent the quintessential light curves i.e. they are evenly sampled and the corresponding periodograms (see Fig. 2.2) are undistorted by a sampling pattern. X-ray observations of AGN are subject to various observational constraints, which largely forbid light curves of the aforementioned quality. It is not possible to monitor an AGN to sufficiently measure the longest time-scales of variability, which could be several centuries or even millennia; moreover, it is completely prohibited to continuously measure the variability at twice the highest waveform frequency for long periods of time (which is *approximately* known to be of the order of minutes), as is required by the Nyquist theorem to negate aliased power. Thus, unlike the simulated light curves and periodograms, the measured periodogram from a real observation does not necessarily vanish above the Nyquist frequency nor below the lowest sampled frequency. The variability power in the underlying process that is not sampled due to insufficient temporal resolution or monitoring extent cannot be extricated from the measured periodogram, which leads to an inherently biased power spectral shape. The real underlying power spectrum of the observed process differs from the measured periodogram in a fashion that is dependent on the underlying power spectral shape and specific sampling pattern. This statement can be considered mathematically as

$$f(t) = l(t) \cdot s(t)$$

where $l(t)$ and $f(t)$ represent the true unsampled light curve and the observed light curve, respectively, and $s(t)$ is the sampling pattern function. The sampling pattern function is defined as

$$s(t) = \begin{cases} +1 & : \text{ when sampling light curve } l(t) \\ 0 & : \text{ otherwise} \end{cases}$$

As previously considered for real observations AGN cannot be continuously observed (e.g. $s(t) \neq 1 \forall t$), so the corresponding sampling pattern can be somewhat more complex. Fig. 2.4 is an example of a simple sampling function and the corresponding light curve. The convolution theorem states that the Fourier transform of two functions multiplied in the time domain is equal to the convolution of the




Figure 2.4: An example light curve with a simple sampling pattern. The top light curve, $f(t)$, has several gaps which correspond to zeros in the sampling pattern, $s(t)$, which is when the target is not observed and $l(t)$ was not sampled.

Fourier transform of these functions in the frequency domain, so

$$F(\nu) = L(\nu) * S(\nu).$$

The functions $F(\nu)$, $L(\nu)$ and $S(\nu)$ represent the Fourier transform of the observed light curve, the underlying light curve and the sampling pattern (or ‘window function’), respectively. Thus the resultant periodogram of the observed light curve, $|F(\nu)|^2$ is distorted due to the convolution of $L(\nu)$ and $S(\nu)$.

The convolution of the window function with the underlying periodogram produces a measured periodogram with power redistributed throughout the measured frequency range. It is unworkable to unravel the true power at a given Fourier frequency from the measured power, due to redistribution of power from other Fourier frequencies *outside* the time-scales sampled by the light curve. These biases can be described in the frequency domain as an interference between frequencies. The amount of power redistributed throughout the periodogram is dependent on the PSD shape. This transfer of power, essentially due to the window function, is a major cause of bias in AGN periodograms. For this work it is convenient to split this bias into two and describe the bias as separate distortions.

Variability power on time-scales longer than the duration of the observed light curve contribute additional power across the entire periodogram. This particular type of

distortion is called ‘red-noise leak’. ‘Aliasing’ is the false translation of power from time-scales shorter than those measured about the Nyquist frequency, and contributes an almost constant amount of power across the periodogram in logarithmic space - but it is also acutely dependent on the window function. I will now discuss some of the technical details of red-noise leak and aliasing, and use simulated light curves to exhibit the effect of these biases on the periodogram.

2.3.1 Aliasing

When a time series is not sampled at twice the highest waveform frequency, additional power is contributed to the measured periodogram in the form of aliasing. Given a sampling interval ΔT , if the true power spectrum contains power at frequencies above the Nyquist, $\nu_{\text{Nyquist}} = 1/(2\Delta T)$, then this power is translated back into the measured periodogram (van der Klis, 1989). The only way to eliminate the effect of aliasing is to use a low-pass filter and filter the frequencies above the highest frequency sampled, but this is obviously far from ideal and not often possible.

To determine how power is redistributed throughout the periodogram, consider the effect of discretely sampling a continuous time series as opposed to smoothing the original continuous time series. Smoothing the underlying time series to remove variability on time-scales less than ΔT is akin to producing a binned continuous time series of time bin width ΔT . If there is *no* variability power on time-scales less than ΔT the mean flux of a bin (width ΔT), centred at the same discrete time as a discretely sampled light curve with sampling interval ΔT , will be identical to the mean flux determined from the discretely sampled light curve; however, in general this cannot be guaranteed and there will be variability on time-scales shorter than those sampled i.e. *RXTE* monitoring campaigns of AGN. If there are frequencies higher than the Nyquist frequency that contain power, then the mean flux in a given time bin of the smoothed underlying function (width ΔT) will not be the same as the mean flux in a time bin calculated from discretely sampled data. There will be an error on the mean flux evaluated at discretely sampled data points, the size of which depends upon the power spectral shape and amplitude above the Nyquist frequency. This is the origin of aliasing in the periodogram.

Fig. 2.5 illustrates the typical effect of aliasing on power spectral shape. The black dashed line represents the true underlying power spectrum, while the solid black line represents the measured average power spectrum after distortion by the effects of discrete sampling at the Nyquist frequency. It is clear that power beyond the Nyquist frequency is reflected back (red dashed line) distorting the measured power.

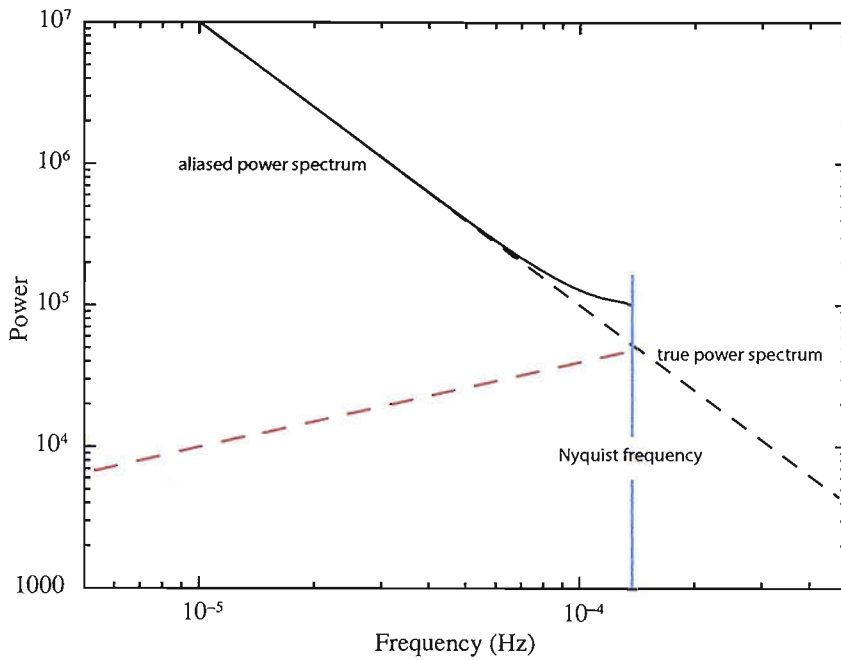


Figure 2.5: Power outside of the observed power spectrum is reflected or ‘aliased’ back about the Nyquist frequency. The solid line is the aliased power spectrum and the dashed lines represent the true underlying power spectrum (black) and reflected power respectively (red).

It should also be clear that aliasing effects the highest frequencies much more dramatically than the lower frequencies and causes the observed high-frequency “turn-up”.

Aliased power, as seen in Fig. 2.5, can be well approximated by a constant power component independent of Fourier frequency, and we assume that the power is redistributed evenly throughout the measured frequency domain. The additional power component due to aliasing can be evaluated by determining the integrated power above the Nyquist frequency - provided the PSD shape is known. The integrated power is dependent on the power spectral shape and it is important to realise that this approximation to the power content above the Nyquist frequency is a determination of the *average* aliased power, since the individual realisations will contribute different amounts of aliased power.

The cause of aliasing can be described as an incorrect determination of mean flux over a time interval of a discretely sampled light curve. It follows that incorporating more points into a time interval affords a more robust estimate of the true mean flux, and thus reduces the amount of aliasing; however, increasing the sampling interval to determine a more accurate mean flux has the unfortunate effect of reducing the Fourier frequencies sampled by the periodogram (in particular at high frequencies), since the Nyquist frequency is reduced by a broader temporal resolution.

2.3.2 Red-noise leak

Real observations are, of course, finite in length and the effect of a finite observation must be considered on the resultant periodogram. Consider a finite observation of length T , if there is significant power at frequencies below $1/T$ the light curve will exhibit slow long time-scale trends (e.g. see panel (b) of Fig. 2.3). Partial cycles of long time-scale trends contribute additional variance to the observed light curve, which cannot be differentiated from the smaller amplitude trends on time-scales that are sampled by the periodogram. The effect of these long time-scale trends is to transfer power across all frequencies in the measured periodogram with most power contributed to the lowest measured frequencies. The amount of power ‘leaked’ into the periodogram is dependent upon the power spectral shape and amplitude. Red-noise leak contributes power with a power-law form of slope $\alpha = 2$, which flattens the measured periodogram. The amount of red-noise leak grows with the amount of integrated power below the lowest measured frequency, ν_{\min} , and consequently, periodograms with steeper power-law slopes exhibit a more pronounced distortion than those with shallower slopes. In fact, red-noise leak is almost negligible when the underlying power-spectral slope is approximately less than $\alpha = 1.5$. At slopes $\alpha > 2$ the additional power due to red-noise leak can be significant. Red-noise leak will be considered in more detail later in this chapter.

2.3.3 Added complication of uneven sampling and binning

Aliasing and red-noise leak are significant problems when dealing with discretely sampled data and I have described their effects on the measured periodogram. So far I have restricted discussion to evenly sampled data; however, real X-ray monitoring observations are not necessarily regular and can be sampled in a somewhat irregular pattern. Many of the Seyfert galaxies in our monitoring campaign have fairly evenly sampled light curves, within an error of ~ 20 per cent. The sampling intervals are not systematically correlated, so the total average sampling interval over an entire campaign is approximately the interval that was requested. This is not always the case however, several AGN observations in the *RXTE* archive do not have such well sampled observing campaigns and can be both irregularly sampled and contain large gaps. Thus red-noise leak and aliasing must be considered for irregularly sampled data.

For a given time bin, the observations do not come from the centre of the time bin, but are randomly scattered about the centre. The deviation of measured flux from

the true mean flux is dependent upon the location of the measured flux within the time bin, and trends within the bin will result in sampled points being further away from the centre - supplementing the effect of aliasing. As before, long time-scale trends with the light curve can cause this effect to be large as these trends affect the distribution of power within the time bin. Red-noise leak complements this effect by further correlating the frequency bins and further increasing the correlation for steeper power spectral slopes. This is extremely problematic since it artificially reduces the point-to-point scatter between bins, the effect of this will be examined in greater detail in Section 2.4.

Other than intrinsic scatter in the sampling interval, many AGN monitoring campaigns have observation gaps that constitute a significant fraction of the overall light curve. In these cases it is best to break the light curve up into segments, since the gaps would contribute significant low-frequency power. Gaps in the light curve due to Earth-occultation, say, are periodic and introduce power at the frequency corresponding to the period of the gaps. Most AGN periodograms, in general, will exhibit all these features to varying degrees of severity. Another, more minor, effect that needs to be considered is the effect due to binning continuous light curves. Binning a continuous light curve steepens the periodogram, since the effect of binning is to smooth the light curve on time-scales alike to the binning time-scale (van der Klis, 1989). However, AGN periodograms at high frequencies (often produced from *XMM-Newton* data) are dominated by the Poisson noise level and are therefore not significantly affected by this phenomenon. While it might be possible (for well-behaved sources) to calculate the periodogram and remove the effect of these aforementioned biases analytically, it would be highly non-trivial and an average rather than a stochastic realisation would be calculated; moreover, the stochastic nature of AGN data would make it impossible for analytical techniques to produce error estimates on power spectral model parameters. Given the complexity involved in evaluating the underlying power spectral model parameters analytically, it is significantly easier to calculate the periodogram and corresponding parameter estimates (with errors) using Monte Carlo methods.

2.4 Significance of red-noise leak

In Section 2.3.1 I discussed the effect of aliasing, and it was immediately clear (see Fig. 2.5) that aliasing could be a serious cause of periodogram distortion (e.g. it could lead to the false identification of a PSD break if a higher frequency PSD is

used in combination); however, it was not clear if the distortion due to red-noise leak, as discussed in Section 2.3.2, could be a significant contribution to overall periodogram distortion.

Red-noise leak could potentially be a more serious effect than aliasing at high frequencies, since the highest periodogram frequencies will often be covered by *XMM-Newton* observations, which are continuous and do not suffer the effects of aliasing. Moreover, the steep power spectral slopes often encountered at high frequencies would only exacerbate the effect of red-noise leak. It is important to determine if there are any statistical implications of correlated frequency bins within the periodogram, which would adversely affect hypothesis testing.

I followed the prescription set out in Section 2.2 using the method of Timmer and König (1995) to simulate several light curves according to an underlying model. The final simulated light curves had a duration of ~ 135 ks (approximately the same length as a typical *XMM-Newton* orbit) and were continuous so aliasing does not affect the resultant periodograms. The light curves were simulated to be ~ 10 times longer than the final light curves so there exists power at lower frequencies than those eventually sampled. AGN periodograms, as described in Chapter 5, often exhibit power-law forms (e.g. Green *et al.* 1993) that break or bend at a characteristic frequency (e.g. Edelson and Nandra 1999; Uttley *et al.* 2002), so the first set of light curves were simulated according to an unbroken power-law form of slope⁴ $\alpha = 2.7$. In this case the steep power-law slope would extend to low frequencies outside the domain of sampled frequencies, it is expected that red-noise leak affects the resultant periodogram by leaking power across all frequencies into the measured frequency range. The power-spectral shape is expected to flatten due to red-noise leak, but by comparing the periodograms to a case where red-noise leak is minimal it can be determined whether there is a strong correlation between the frequency bins due to the addition of a constant power component. A second set of simulations using the same random seeds, but with a slightly different model were produced. In the second case light curves were simulated with a power-law slope of $\alpha = 2.7$ that broke to a power-law slope of $\alpha = 0$ at $\nu = 1 \times 10^{-4}$ Hz. In this second case the periodograms are expected to be only marginally effected by red-noise leak, since the power content of the low-frequency power spectrum is small. For clarity only three periodograms are plotted in Fig. 2.6. The break in the second plot in Fig. 2.6 is clearly visible, but the important information is the scatter between individual realisations. In Fig. 2.6 the underlying power-law extends to low

⁴ $\alpha = 2.7$ was arbitrarily chosen and is representative of the typical high-frequency power-law slope of the PSDs found in this work.

frequencies and thus there is a significant amount of red-noise leak, which appears to be reflected in the reduction of point-to-point scatter in the periodograms - but only at high frequencies.

A reduction of this nature is expected if the bins are strongly correlated, an increase in the overall scatter between the realisations is also expected. Indeed, both these effects are observed in the top plot in Fig. 2.6. The bottom plot in Fig. 2.6 is not expected to exhibit these effects in such a dramatic fashion, since the effects due to red-noise leak would be significantly reduced. It appears if even for fairly modest slopes, red-noise leak can have a fairly dramatic effect on the resultant periodogram. The correlation between the frequency bins due to red-noise leak (and aliasing to a certain extent) could have grave implications for statistical analysis if not properly taken into account. The χ^2 statistic which is often used as a measure of goodness-of-fit assumes statistically independent bins. A naive χ^2 test to determine the goodness-of-fit of an underlying model would be skewed by the correlated bins at high frequencies, whereas the lower frequency bins would remain relatively unaffected by this problem since the slopes are often less than 1.5. The problem of statistically dependent bins is non-trivial and will be discussed in detail in Chapter 3 when we overview the Monte Carlo method for AGN power-spectral estimation.

2.5 Stationarity

To produce a broadband periodogram that spans several decades in Fourier frequency, several different periodograms need to be combined to measure different time-scales e.g. the best AGN periodograms may have 11 or more years of *RXTE* observations and span approximately six decades in Fourier frequency. However in order to combine periodograms of different time-scales it must be assumed that the true underlying power spectrum is not a function of time i.e. the power spectrum is *stationary*. Non-stationarity would cause significant problems in the measured periodogram, since it could manifest as an artificial change in the power spectral shape and be misinterpreted as a characteristic feature. The question of stationarity can only be answered observationally.

BHXRBS exhibit dramatic changes in the true underlying power spectrum (i.e. non-stationarity) on long time-scales. In particular, the change between black hole spectral states is associated with a change in power spectrum (e.g. Cui *et al.* 1997; Belloni *et al.* 2005). These changes in spectral state occur on a variety of time-scales from months to days or even minutes in some transients; however, the changes in




Figure 2.6: Top plot: Three periodograms affected by red-noise leak. The underlying power spectrum is of a power-law form and extends down to low frequencies with a slope of $\alpha = 2.7$. It is important to notice the scatter between periodograms. Bottom plot: Here we plot the same three periodograms as seen above (i.e. using the same random seeds), but using a different underlying model. The model used here breaks from a slope of $\alpha = 2.7$ to a slope of $\alpha = 0$ at $\nu = 1 \times 10^{-4}$ Hz. The overall spread of realisations at high frequencies is reduced in comparison to the top plot, which is to be expected given the reduction in red-noise leak.

the power spectral shape and amplitude of BHXRBs are significantly less dramatic within any given spectral state than the changes observed over longer time-scales (Belloni and Hasinger, 1990; Cui *et al.*, 1997). If the aperiodic X-ray variability seen in black hole systems scales approximately linearly with mass, we might expect to observe similar spectral changes (state transitions) and non-stationarity in AGN power spectra. However if we scale the typical stationary time-scale observed in BHXRBs by the difference in mass between an AGN and BHXRB ($\sim 10^6 M_{\odot}$), we would expect to find that AGN stationary time-scales might last for decades or centuries i.e. much longer than our longest monitoring campaigns. The scaling argument for stationary AGN power spectra seems fair and is supported by the lack of evidence to the contrary given that an 11 year *RXTE* light curve of NGC 4051 and MCG-6-30-15 exhibit no evidence for a change in power spectral shape. It is probably sufficient to assume that the targets mentioned in this work do not exhibit changes in their power spectrum on the time-scales we have been monitoring. In order to generate a broadband periodogram over several decades in frequency, it is required to observe the target with several different sampling intervals. Where possible we have obtained long time-scale light curves spanning months to years, medium time-scale light curves spanning days to months and the continuous *XMM-Newton* observations spanning minutes to hours. These different observations can be combined to form a broadband periodogram using the fractional RMS-squared normalisation i.e. removal of the mean flux dependence. Constructing a broadband periodogram in this fashion removes the requirement to obtain a prohibitive amount of data. A few objects have *RXTE* ‘long-look’ data available, which is a prolonged observation that may span several thousand seconds. For a few fortunate sources these long-looks may be *quasi*-continuous with a few minor gaps, but the observing constraints of a low-Earth orbit generally forbid any sufficiently extended continuous monitoring.

In Section 2.1.2 I introduced the concept of fractional RM-squared normalisation so that different periodogram amplitudes can be compared. However even though normalised AGN power spectra are expected to be stationary, the same cannot be said for AGN light curves. AGN light curves are red-noise light curves, and are inherently non-stationary on the time-scales currently probed with *RXTE*. The mean flux when calculated at different times and over different durations can vary. Time-dependence in the mean flux of a red-noise light curve could be problematic for the fractional-RMS squared normalisation (see Eqn.2.10), since it requires a constant mean value. The problem with a dynamic mean flux is deciding which mean flux value to use, and the problem is only exacerbated by steeper slopes where there

exists a strong trend throughout the light curve. We expect AGN power spectra to flatten to slopes less than unity at low frequencies, so we might expect the light curve to be stationary on long time-scales. If we assume this is the case, a mean flux measurement on the longest time-scales (\sim years) should yield a result that is a good approximation to the true value. However Uttley and McHardy (2001) discuss the flux dependence of AGN and BHXRB variability and discovered that there exists a linear RMS-flux relation - a linear relationship between RMS variability and the mean segment flux, such that the fractional RMS variability remains constant despite a change in the mean segment flux. They also found that regardless of the non-stationarity of the AGN light curves, the periodogram normalised by the *local*⁵ mean flux is approximately stationary, since there is no dependence on local mean flux in the fractional RMS variability. Thus in order to compare AGN periodogram estimates from different times and over different time-scales, we must normalise the periodograms by the local mean flux.

2.6 Summary

In this chapter, I described the various conventions and terms that will be encountered throughout this work. I have considered some of the variability properties of red-noise light curves typically seen in AGN X-ray observations, and discussed how to measure the periodogram and its relationship to the underlying power spectrum. Various PSD normalisations have been considered, but I described the merits of using the fractional RMS-squared normalisation and I use this normalisation throughout this work. Various properties of the periodogram have been reviewed and in particular the logarithmic binning of the measured periodogram is described. Light curve simulation was also described as it is a useful tool when determining the effect of the window function upon the underlying power spectrum, since the simulated data can be treated in an identical way to the observed data. It is essential to simulate ‘realisations’ of the real data based on an underlying model so that the window function distortions can be determined and a true representation of the underlying power spectrum estimated, correct determination of these distortions is dependent on the power-spectral model. The quantitative effects of distortion on the measured periodogram are detailed, namely red-noise leak, aliasing and the effect of binning. I investigated the effect of correlating frequency bins with red-noise leak by considering two cases: periodogram realisations where there exists significant power at

⁵Flux of individual light curve used to measure PSD rather than any longer-term mean

low frequencies and similar realisations where there exists very little power at low frequencies. I determined that the effect of red-noise leak is to strongly correlate frequency bins, and this effect cannot be ignored when undertaking statistical model fitting. The concept of stationarity and combining multiple periodogram measures concluded this chapter.

In the next chapter I will discuss PSRESP: a power-spectral response method and how it can be used to estimate underlying power spectrum parameters from AGN observations in the presence of significant bias. I will also develop further some of the concepts introduced in this chapter to modify and improve the current Monte Carlo method.

Computers are useless.
They can only give you answers.

PABLO PICASSO

3

Power spectral modelling: PSRESP

In this chapter I will describe how to determine the best fit underlying model power spectrum parameters from an observed periodogram distorted by the window function. The principal technique involved in determining the model parameters that describe the observed periodogram is called the ‘power spectral response method’. I will briefly describe the motivation behind the power spectral response method and show how this technique can be applied to AGN observations using Monte Carlo simulations. Monte Carlo simulations provide a means by which the model parameters that best fit the observed periodogram can be determined for a noise process in the presence of sampling biases; moreover, periodogram error estimation is naturally included. I will review the method of Uttley *et al.* (2002), discuss its caveats and describe the specific improvements I have made to our Monte Carlo software, PSRESP. I will conclude this chapter by discussing the computational requirements and possible optimisations of PSRESP and its current limitations.

3.1 Response method

In order to determine the parameters that best describe the observed periodogram, it is practical to base our analysis on the response method of Done *et al.* (1992), which

is described by analogy with X-ray spectroscopy. X-ray detectors have a finite energy range and a limited spectral energy resolution, the effect of which causes the measured X-ray energy spectrum to distort, since the true X-ray energy spectrum is convolved with the ‘spectral response function’. Underlying models that are theorised to fit the true X-ray energy spectrum are convolved with the most reliable estimate of instrument response and compared to the real data. Using standard χ^2 statistics, the best fit underlying model parameters can be determined, and an acceptance probability that the model is good is associated to the fit (Lampton *et al.*, 1976; Yaqoob, 1998).

The finite energy bandpass and limited energy resolution in X-ray spectroscopy is analogous to observation duration and sampling interval in X-ray timing. Power spectral distortions can be described as the convolution of the true Fourier transform of the observed light curve with the spectral window function (see Eqn.2.3). The spectral window function in X-ray timing is analogous to the instrumental response function in X-ray spectroscopy. Thus it would be sensible to assume that a similar method for determining the true X-ray energy spectrum exists for power spectral modelling. The method requires that an underlying power spectral model be convolved with the spectral window function in Fourier frequency space to produce periodograms distorted in an identical fashion to the real observations; although, we initially work in the time domain and do not directly compute the convolution. The underlying model can be explored and χ^2 measured to determine the best fit model parameters by comparing it to the real observed periodogram; this concept forms the foundation of the response method and PSRESP (our Monte Carlo simulation software).

There is a crucial difference between determining the true X-ray energy spectrum and the true power spectrum: inherent uncertainty in the power spectral shape caused by the stochastic nature of AGN light curves. There is no such analogy in X-ray spectroscopy for the stochastic nature of red-noise light curves. Each observed light curve and corresponding periodogram is merely a single stochastic realisation of a noise process drawn from a distribution of possible outcomes given the true underlying power spectrum that best describes that process. The best way to emulate the stochastic nature of AGN light curves is through random number generators in Monte Carlo algorithms.

A reliable and robust way to determine the true power spectrum in the presence of sampling distortions, while taking into account the stochastic nature of red-noise data, is to simulate pseudo-observations based upon an underlying model (Done *et al.*, 1992; Green *et al.*, 1999). The basic concept involves simulating a large

number of light curves, according to an underlying model, that are sampled in an identical way to the real data and consequently produce periodograms that contain the same built-in distortions as the real periodogram. Each periodogram will span the same Fourier frequencies, so for each frequency bin there will be a large number of periodograms forming a distribution of power estimates. The power average (also called the distorted model average PSD) and RMS spread of the power distribution can be calculated and compared to the real periodogram by evaluating the χ^2 statistic (see Eqn.2.5). The power distribution at each frequency represents the power distribution expected from the underlying noise process given the input model parameters. The fundamental question is: is the observed periodogram consistent with the power distribution at each frequency? The model can be changed or parameters searched until the parameters that minimise the χ^2 statistic are found. It should be noted that assigning errors to the model rather than the observed data is formally correct within the framework of the χ^2 statistic, since the χ^2 is based on the variance of the model population from which the observation is drawn.

Here we remark on a crucial element of determining the best fit model parameters. The results of χ^2 minimisation has a formal interpretation only if the measured errors are normally distributed. The χ^2 statistic determined by comparing the distorted power spectrum model average to the real periodogram is *not* necessarily the same as a true χ^2 statistic. The measured χ^2 statistic can nonetheless still be used to determine the best-fitting model parameters, but it must not be used to formally evaluate the acceptance probability or even estimate confidence on model parameters. Throughout the rest of this work whenever a measured χ^2 value is used in hypothesis testing that does not necessarily conform to a standard χ^2 distribution, it will be labelled as χ_{pseudo}^2 .

In order to produce robust model acceptance probabilities for hypothesis testing, the distribution of χ_{pseudo}^2 values measured from each simulated periodogram compared to the distorted model average PSD must be determined. The null-hypothesis that *the observed data is consistent with being a realisation of the underlying model* can be estimated by comparing the observed χ_{pseudo}^2 statistic between the real data and distorted model average PSD to the distribution of χ_{pseudo}^2 values calculated between the simulated realisations to the distorted model average PSD. The probability percentile defines the acceptance probability. If 10,000 periodograms were simulated and only 500 of these had a χ_{pseudo}^2 value greater than the χ_{pseudo}^2 value measured between the observed periodogram and the distorted model average PSD, there would be an approximate 5 per cent probability that the real data is a stochastic realisation of the underlying power spectral parameters, and these particular parameters could

be rejected with 95 per cent confidence; moreover, if a sufficient number of parameters are tested, rejection probability levels of the parameter space can be constructed to determine errors on the best fit parameters. The errors on the best fit parameters are often taken to be the width of the rejection probability level i.e. in the contour space of parameter values, the 90 per cent error estimate is taken to be the width of the 90 per cent rejection probability level. The x per cent rejection probability level denotes the point at which the acceptance probability for a given combination of parameters drops to less than $100 - x$, e.g. for parameter values outside the 90 per cent rejection probability level, the acceptance probability for those parameters drops to less than 10 per cent. The method by which these rejection probability contours are formed is based on the technique of Press *et al.* (1992).

Done *et al.* (1992) was first to use the response technique when measuring the periodogram of *Ginga* data and it was later implemented by Green *et al.* (1999) for *ROSAT* observations. However, the observations were relatively short in duration and the technique could only be applied to a single light curve. The restrictions on AGN observations dictate that a broadband periodogram can only be constructed from several smaller segments of different durations and temporal resolutions; thus, multiple periodograms must be modelled simultaneously. A more sophisticated Monte Carlo based method that is able to constrain the broadband periodogram was developed by Uttley *et al.* (2002), the details of which I will describe below. I will also examine the caveats of the original implementation and discuss the extensive improvements I have made to our software, PSRESP.

3.2 Improved PSRESP implementation

Here I will describe the Uttley *et al.* (2002) implementation of the Monte Carlo software, PSRESP, and further discuss the improvements I have made.

Our AGN monitoring campaigns are often constructed in such a way as to minimise the amount of telescope time required, but to provide comprehensive coverage of the Fourier frequency space in the resulting periodogram. A typical AGN sampling pattern can be seen in Fig. 3.1. Each of the observed data points are constructed from at least a single ~ 1 ksec snapshot, and the region of significantly more intensive sampling is easily identifiable at the start of the light curve. A broadband periodogram is constructed as follows. The intensively sampled region (medium time-scales) is extracted as a separate light curve with a typical sampling interval of up to a few times daily. *RXTE* medium time-scale light curves may be a few months




Figure 3.1: A typical *RXTE* light curve that exhibits the type of sampling patterns often observed in AGN data. The intensive or medium sampling interval (\sim multiple times daily) is immediately noticeable to the start of the light curve. The data prior to the gap is used as a long time-scale observation with an average sampling of approximately 7 days. Finally the entire light curve is used as a single very long-time scale observation, binned up on time-scales approximately similar to the gap as to reduce the spurious power introduced by the gap.

in duration and are designed to cover 10^{-6} to 10^{-5} Hz in Fourier frequency space. Some AGN observations also have *RXTE* long-look data (very-high resolution, \sim 16 sec, quasi-continuous monitoring) for covering Fourier frequencies greater than 10^{-5} Hz. *RXTE* long-looks are very useful, but it is often more convenient to use *XMM-Newton* to cover the fastest time-scales, since *RXTE* long looks are rarely as long or as regularly sampled. The long time-scales are determined by the duration of the total campaign. Fig. 3.1 exhibits a long time-scale sampling pattern (before the gap), which can be used to constrain frequencies below 10^{-6} Hz. The long time-scale region would be binned to the average sampling interval (typically \sim 4–14 days), and cover the time-scales that the more intensive region cannot probe. In the case where a gap in the light curve is not excessive (i.e. not a significant fraction of the overall light curve), as in Fig 3.1, binning the entire light curve up on time-scales approximately equal to the gap size yield a periodogram that can probe the lowest possible time-scales. If enough data exists past the gap then another periodogram can also be calculated. Normalised by the fractional RMS-squared normalisation, the component periodograms from each segment can be combined

to produce a broadband periodogram covering many decades in Fourier frequency. Combining the periodograms for each light curve region given in Fig. 3.1, along with a continuous light curve obtained from *XMM-Newton* yields the broadband periodogram seen in Fig. 3.2. The three *RXTE* periodograms (very-long, long, and medium time-scales given by red, blue, and yellow, respectively) are taken from the sampling pattern in Fig. 3.1 will be affected by all the previously mentioned spectral window effects; although the *XMM-Newton* periodogram (shown in purple) will not be affected by aliasing, it flattens to zero slope at high frequencies due to the effect of Poisson noise.

A new feature of PSRESP is the ability to frequency-limit periodograms so that very-long input light curves inclusive of gaps can be used. Given that the very-long time-scale light curve is largely the same data as the long time-scale light curve, we must band-limit the very-long time-scale periodogram (red in Fig. 3.2) so that it *does not* cover the same time-scales already covered by the long-time scale periodogram (blue in Fig. 3.2). Ensuring the overall periodogram is band-limited at the intersection between component periodograms means that the data will never be included more than once at a given time-scale. Note that aliased power will have been redistributed throughout the entire periodogram prior to frequency-limiting, but the contribution due to aliased power from frequencies outside the frequency-limit is negligible.

The basic PSRESP implementation is presented and discussed in Uttley *et al.* (2002), but here I will describe and the new implementation of PSRESP. Technical details of the improvements, along with a discussion of why the improvement is warranted can be found in Section 3.3.

Once the various sampling patterns have been identified and several component light curves created, they can be input in to PSRESP as segments to produce the broadband periodogram. In order for PSRESP to produce power spectral model parameter estimates, information pertaining to light curve statistics must also be provided. PSRESP also requires details of the component light curve mean flux, and three light curve binning parameters (see Section 3.4): simulated temporal resolution, resolution of the input light curves, and the intrinsic snapshot resolution (~ 1 ks for *RXTE*). Parameters that describe how to bin up the resultant periodogram are also included, along with several other parameters for new supplementary functionality. The Poisson noise level is directly calculated from the input light curves (see Section 3.3.1).

An indicative structural overview of PSRESP is as follows. Each observed light curve that is to be used in the overall broadband periodogram is read into PSRESP.




Figure 3.2: The broadband periodogram of Fig. 3.1 and additional *XMM-Newton* data, producing 4 separate segments. This periodogram is constructed from 3 different *RXTE* sampling patterns: very-long, long and medium time-scales - red, blue and yellow respectively. The purple periodogram is from a continuous *XMM-Newton* observation, notice the flattening of the high-frequency periodogram to zero slope. i.e. the Poisson noise level. The long time-scale periodogram (blue) is band-limited so that it does not infringe on the frequencies measured by the very-long periodogram (red), since both periodograms have been made using largely identical data.

For each light curve the mean count rate is subtracted, which is equivalent to removing the zero-frequency power component in the resulting periodogram. The periodogram is calculated for each light curve, $P_{\text{obs},i}(\nu)$, and normalised by fractional RMS-squared normalisation. The periodogram is then binned as determined by the input binning parameters, and band-limited if required. The input light curve count rate errors are used to determine the Poisson noise level for each segment (see Section 3.3.1). For a given power spectral model, underlying model parameters are selected from a parameter space of possible combinations and input into PSRESP for testing (see Section 3.4). The PSD parameters define the power spectral model, and the following algorithm is performed.

- For a given set of model parameters, simulate N light curve realisations from the underlying power spectral model. Sample each of these light curves according to the observed light curve, i .
- Evaluate the periodogram of each simulated light curve and logarithmically bin the simulated periodograms identically to the observed periodograms,




Figure 3.3: This plot illustrates the simulated ensemble of periodograms given an underlying power spectral model. From the distribution of periodograms, an average power and RMS spread is determined at each Fourier frequency. The observed periodogram is checked against the periodogram distribution at each Fourier frequency to determine the chance that it could be a realisation of the underlying power spectral model. The apparent discontinuity between the periodogram segments at $\sim 10^{-6}$ Hz is a direct result of aliasing.

producing i segments in the broadband periodogram.

- Average the ensemble of simulated periodograms at each Fourier frequency to obtain the distorted model average power spectrum and determine the RMS error from the distribution at each Fourier frequency. This is done for each segment. See Fig. 3.3 for an example plot of the ensemble of simulated periodograms.
- Calculate the χ_{pseudo}^2 statistic between the model average power spectrum, \bar{P}_{pseudo} , and the observed periodogram, P_{obs} , for each component light curve,

$$\chi_{\text{pseudo}}^2 = \sum_{\nu_{\min}}^{\nu_{\max}} \frac{(\bar{P}_{\text{pseudo}}(\nu) - P_{\text{obs}}(\nu))^2}{\Delta \bar{P}_{\text{pseudo}}(\nu)^2} \quad (3.1)$$

The frequency limits span the Fourier frequency space as determined by P_{obs} . $\Delta \bar{P}_{\text{pseudo}}(\nu)$ is the RMS error on the mean power spectrum.

- Transform the χ_{pseudo}^2 statistic into a real χ^2 value by determining the probability percentiles for each simulated periodogram in each segment, this step

will be explained in detail in Section 3.3.2. Construct χ_{total}^2 by summation of χ^2 over the input periodograms i.e.

$$\chi_{\text{total}}^2 = \sum_i \chi_i^2 \quad (3.2)$$

where i denotes the segment, which for Fig. 3.2 would be 4.

- Renormalise the segments simultaneously to minimise the χ^2 statistic between the observed periodogram and the distorted model average PSD. I will explain technical difficulties of minimising the χ^2 statistic in Section 3.3.2.

In order to provide an estimate of the acceptance probability, without making any assumptions of the statistics, we use the raw distributions constructed in the Monte Carlo simulations. We test the null-hypothesis and determine the probability that the observed broadband periodogram is a realisation of the underlying model as follows.

- For each segment in the broadband periodogram, randomly select a simulated periodogram used to make the distorted model average PSD. Evaluate the χ_{pseudo}^2 statistic between the realisation and the distorted model average PSD, convert to χ^2 for each segment and combine the values to form χ_{total}^2 and minimise this value by cycling through renormalisation values.
- Repeat the above process M times to form a distribution of χ_{total}^2 for a given choice of model parameters. The probability (acceptance probability) that the observed data can be described as a realisation of the model power spectrum is determined by the number of points in the distribution that *exceed* the observed χ_{total}^2 value.
- The acceptance probabilities can be used to set confidence limits on the input model parameters and plot rejection¹ probability levels, which can be used to determine errors.

A χ_{total}^2 distribution that is constructed from every possible combination of simulated periodogram segments would be preferred; however, the number of possible combinations is equal to N^W , where W is the number of periodogram segments in the broadband power spectrum and N is the number of simulated periodograms. The number of combinations becomes prohibitory for even a modest number of

¹ $(1 - P)$ probability, where P is the acceptance probability

segments and simulated periodograms (e.g. for low, medium and high frequency power spectral segments constructed from a modest 300 light curves, the number of possible combinations would be 27 million), thus we choose a smaller sample, of M , randomly selected periodograms from all power spectral segments. We restrict $M > 1000$ so that we obtain a reliable and robust estimate from the distribution of χ_{total}^2 , with an acceptance probability resolution at least better than 0.1 per cent.

3.3 Technical overview of the major improvements to PSRESP

I have made significant modifications to the Uttley *et al.* (2002) version of PSRESP and here I will detail these improvements.

The first improvement to PSRESP is, to first order, an increase in computational performance and a more robust error determination on the high-frequency periodogram. Due to computational limitations at the time, the version of PSRESP implemented by Uttley *et al.* (2002) does not attempt to simulate the high-frequency part of the broadband periodogram, so it would use a directly measured periodogram that is not affected by aliasing to constrain the shortest time-scales and determine the error bars from the RMS spread of the logarithmically binned data. In order to provide a consistent way of evaluating the broadband periodogram over all time-scales, I have modified the software to simulate up to the highest useful frequencies, which is typically $\sim 10^{-3}$ Hz for AGN but can be as high as 0.1 Hz in extreme cases. I have modified our software to correctly evaluate the error bars on the simulated data in the presence of dominant Poisson noise; moreover, the Poisson noise component in the power spectrum is directly calculated from the input light curves. Even though aliasing in continuous light curves is thought to be negligible, this modification allows aliasing, red-noise leak and the effect of binning to be taken into account consistently throughout all frequencies in the periodogram.

Perhaps the most significant modification to PSRESP concerns acceptance probability determination for underlying model PSD parameters. A consequence of simulating the high-frequency periodogram consistently with the broadband periodogram is that the goodness-of-fit measurement (measured in the same way as a χ^2) needs to be reconsidered. The high-frequency periodogram segment can be significantly more distorted by red-noise leak and as a consequence has fewer statistically independent bins. The assumption that all segments contribute equally to the overall goodness-of-fit no longer holds true. I implement a method to implicitly re-weight

each periodogram segment so that the total goodness-of-fit is not biased towards any segment. The modification to acceptance probability determination not only allows more robust probabilities to be established, but more accurate calculation of the errors on the best fit model parameters.

3.3.1 The Poisson noise level and simulated error bars

Enabling PSRESP to simulate the highest frequencies in the broadband periodogram dictates the simulated errors be refined in the presence of Poisson noise. Previously, Poisson noise was treated as an additive component that was included when fitting real data to the distorted model average PSD. However the frequencies most dramatically affected by Poisson noise were usually determined by a directly measured periodogram, and the errors were essentially evaluated based upon the logarithmic binning explain in Section 2.1.2, which do not take into account the distortion on the periodogram due to red-noise leak. Simulations of high-frequency data in PSRESP caused the simulated error bars to be adversely affected by the Poisson noise level; this problem is illustrated in Fig. 3.4. It is clear that the simulated error bars at high-frequency shown in Fig. 3.4 as green data do not correctly represent the scatter in the observed periodogram (solid black line in Fig. 3.4). The simulations do not include Poisson noise and are dominated by red-noise leak at high-frequencies, which reduces the point-to-point scatter between the periodogram bins. The error bars on the green periodogram appear to grow at high-frequency (due to red-noise leak), but the power spectrum should be dominated by the Poisson noise level in logarithmic power space - the error bars have merely been translated. The errors bars should decrease with increasing frequency since as the number of cycles increase, the power is better determined. In order to produce accurate error bars on the high-frequency data, the effect of the Poisson noise level on the simulated error bars must be included correctly i.e. the absolute error bars must be recalculated rather than relative error rescaled in the presence of Poisson noise.

It is clear that the point-to-point scatter of the real data in the Poisson noise region is minimal, but the red-noise leak dominated errors in the simulations are much too large to represent this scatter. In this regime, the Poisson noise level should dominate the errors, not red-noise leak. The errors were corrected by recalculating them in the presence of Poisson noise i.e. the new spread was determined by taking the error bar extrema and computing the absolute error with the additional power of the Poisson noise level. The new error bars are completely dominated by the power in the Poisson noise level with the simulated spread making a small contribution to




Figure 3.4: The solid black line is the observed periodogram and the red dashed line is the underlying model. The blue simulated power spectra do not contain the additional Poisson noise component, the green simulated power spectra are modified by the Poisson noise level addition. However, it is clear that the simulated error bars are not altered by the Poisson noise level in the original version of PSRESP and are overestimated. The scatter in the presence of Poisson noise is not representative of the true scatter, so the high-frequency error bars are not accurate.

the signal at high-frequencies. The simulated spread in the presence of noise will be too small for the χ^2 fitting to produce reliable acceptance probabilities, since the observed periodogram contains the intrinsic scatter due to the Poisson noise level determination. I include this additional scatter due to the Poisson noise in the simulated error bars by considering the errors of the input light curves. The errors are used to generate 1000 realisations of the observed error “light curve” by taking a Gaussian scatter about the normalised errors that can be sampled and binned in an identical fashion to the observed data. The periodograms for all 1000 noise realisations are calculated and binned to form a distorted average noise level periodogram for each light curve segment. The resultant periodogram which denotes the Poisson noise level contains the same sampling distortions as the real data; moreover, it varies with frequency - the Poisson noise level is not constant in this version of PSRESP.

The error on the mean Poisson noise level (determined from the simulations) is calculated analytically from $P_{\text{Noise}}/\sqrt{N}$, where N here is the number of point in the frequency bin, and is combined with the simulated RMS spread of the realisations by linear addition in Fourier space (rather than in quadrature) so to preserve the




Figure 3.5: The solid black line is the observed high-frequency periodogram, the red dashed line is the underlying model and the green points denote the simulated model average plus error bars. The recalculation of the RMS spread in the presence of Poisson noise, along with the inclusion of the intrinsic scatter of the Poisson noise level provides a more robust error calculation at high-frequencies. The errors decrease, which is consistent with expectation.

χ^2 distribution of power. The new method of determining the errors on the high-frequency data can be seen in Fig. 3.5, where the errors now reduce with frequency and are representative of the observed point-to-point scatter; moreover, the errors include components due to simulated RMS and intrinsic Poisson noise level determination.

It is worth summarising the advantage of this method of determining the periodogram. Through the improvement I have made to the Poisson noise level calculation, PSRESP is now able to simulate high-frequency data and reliably calculate the correct error bars in the presence of red-noise leak and aliasing. The broadband periodogram can be determined in an entirely self-consistent manor, including effects due to sampling and binning on the Poisson noise level, which is no longer assumed to be constant.

This new feature of PSRESP also provides a means by which correlated light curves can be properly dealt with. Consider when *XMM-Newton* observations are extended over multiple consecutive orbits, the light curve will be comprised of multiple continuously sampled segments separated by orbital gaps. The beneficial feature of *XMM-Newton* data is the continuous sampling, so it would be advantageous to calculate the periodogram of the separate segments in the multiple-orbit light curve

and bin the periodograms together to provide a more robust PSD estimate at the high-frequencies (this was described for BHXR data in Section 2.1.2 i.e. over-sampling). The added number of periodogram estimates per bin at the lowest frequencies would slightly extend the periodogram measurement to lower frequencies, thus marginally increasing the frequency span of the *XMM-Newton* data. It would not be rewarding to bin-up the *XMM-Newton* data on the gap time-scale so that multiple orbit data can probe lower frequencies, since those frequencies would probably be covered by *RXTE* monitoring. Multiple consecutive orbits of *XMM-Newton* data would be correlated due to the same long time-scale trend pervading through the light curve, thus the simulated data at high-frequencies (for multiple consecutive orbits) is simulated sequentially as one long light curve so that the simulated light curves for each orbit contain the same long time-scale trend, similar to the observed data i.e. the component periodograms of the binned periodogram are not treated as independent.

3.3.2 Distortions, skewed distributions and acceptance probability determination

The most significant change to PSRESP concerns the determination of the acceptance probability. The Uttley *et al.* (2002) version of PSRESP made an implicit assumption concerning the nature of the χ^2_{pseudo} statistic. It was assumed that each χ^2_{pseudo} value from each segment contributes to χ^2_{total} as if each frequency bin was statistically independent. The high-frequency periodogram segment distribution is treated as contributing ν independent bins; however, I have previously shown that frequency bins in the high-frequency periodogram are often strongly correlated (see Section 2.4). As a result the overall χ^2_{total} distribution contains a contribution of χ^2_{pseudo} values from the high-frequency segment that are falsely weighted as having ν independent bins - when the actual number of independent bins might in fact be much lower than ν . Correlated bins, if not properly dealt with, could have a dramatic affect on the acceptance probability and reduce its reliability.

To determine whether or not PSRESP was subject to this flaw, I undertook the following analysis. I exported χ^2_{pseudo} distributions for each periodogram segment and inspected the shape of the distributions to see what effect they would have on the overall distribution. The distributions are not necessarily expected to behave like a true χ^2 distribution, since each periodogram segment is subject to sampling biases. The χ^2_{pseudo} distribution, constructed from 3000 random selections, for typical *RXTE* and *XMM-Newton* data is shown in the top plots of Fig. 3.6 and Fig. 3.7.

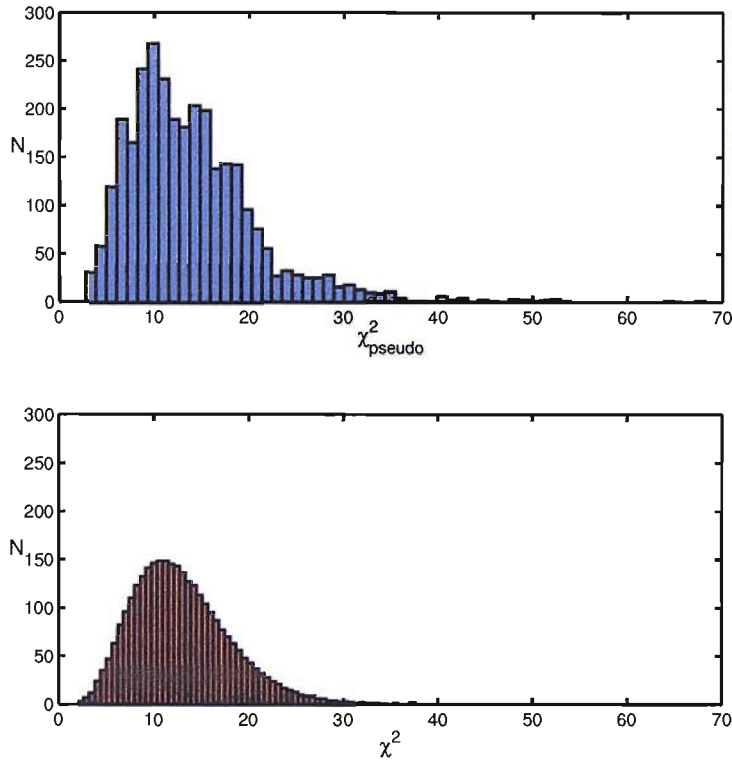


Figure 3.6: The top panel represents the distribution of measured χ^2_{pseudo} values (blue) for typical *RXTE* data. It is apparent that there are more values in the extrema of the distribution, with a slight deviation away from χ^2 . The bottom panel represents the transformed distribution of χ^2 values (red) generated from the top distribution on the same scale. This distribution no longer contains the long (distorting) tail of artificially large χ^2_{pseudo} values. Three important statistical properties that define the above distributions can be found in Table 3.1.

Property	Predicted value	χ^2_{pseudo} distribution	χ^2 distribution
μ	13.00	13.91	13.01
σ	5.10	7.05	5.10
γ_1	0.78	1.73	0.79

Table 3.1: Typical χ^2_{pseudo} distribution properties of *RXTE* data to 2 decimal places compared to a real χ^2 distribution (predicted value) and the transformed χ^2 value.

Note that Fig. 3.6 and Fig. 3.7 represent component distributions used to generate the overall distribution χ^2_{total} .

Inspection of the *RXTE* χ^2_{pseudo} distribution data (see top plot of Fig. 3.6) suggests that the distribution is slightly skewed with the mean, standard deviation and skewness of the distribution not too dissimilar from a true χ^2 distribution, these values are tabulated in Table 3.1. For reference Table 3.1 also lists the predicted statistical

properties of a real χ^2 distribution with the degrees of freedom of the measured χ_{pseudo}^2 distribution, df . Given df the mean, μ , standard deviation, σ , and skewness, γ_1 , of a χ^2 distribution can be calculated as follows

$$\mu = df \quad (3.3)$$

$$\sigma = \sqrt{2df} \quad (3.4)$$

$$\gamma_1 = \frac{2^{1.5}}{\sqrt{df}}. \quad (3.5)$$

It is clear that the measured statistical properties (presented in Table 3.1) do not deviate from the theoretical values significantly and since red-noise leak is not a major source of bias in *RXTE* AGN data², the deviation is more likely due to the effects of binning and aliasing. However, the high-frequency data (*XMM-Newton* data) segment χ_{pseudo}^2 distribution (see Fig. 3.7) indicates that the distribution deviates away from an ideal χ^2 distribution significantly. The high-frequency segment χ_{pseudo}^2 distribution is noticeably skewed and has a significant tail extending to large χ_{pseudo}^2 values i.e. some periodogram realisations have exceedingly poor (or good) fits to the distorted model average PSD. A comparison between the statistical properties of the measured χ_{pseudo}^2 distribution for typical *XMM-Newton* data and the theoretical values predicted from a χ^2 distribution with the same number of degrees of freedom are given in Table 3.2. It is obvious from both Fig. 3.7 and Table 3.2 that the χ_{pseudo}^2 distribution for typical *XMM-Newton* data deviates significantly away from a χ^2 distribution. Thus I conclude that the acceptance probability is likely to be seriously affected if high-frequency *XMM-Newton* data is included with no account taken of the χ_{pseudo}^2 distribution distortion. The high-frequency data in question is continuously sampled and will exhibit very little (negligible) periodogram distortion due to aliasing or binning effects; thus, we can conclude that the distortion to the high-frequency segment data χ_{pseudo}^2 distribution is almost entirely caused by red-noise leak. As briefly described in Section 2.4, one effect of red-noise leak is to cause a strong linear correlation between frequency bins. The extended tail in Fig. 3.7 is a direct consequence of the steep power spectral slope and subsequent correlating effect of red-noise leak on the simulated high-frequency periodogram. As outlined, in Section 3.2 the overall distribution used to determine the acceptance probability of the real data is constructed by combining a randomly selected χ_{pseudo}^2 value from each periodogram segment, minimising the sum over the segments, and

²PSD slopes of *RXTE* AGN monitoring data are often found to be ~ 1 .




Figure 3.7: The top panel represents the distribution of measured χ^2_{pseudo} values (green) for typical *XMM-Newton* data. It is obvious that *XMM-Newton* distribution is significantly affected by red-noise leak, the extensive tail is particularly problematic. The plot has in fact been truncated, since the tail of the distribution continues out as far as $\chi^2_{\text{pseudo}} \sim 145$. The bottom panel represents the new χ^2 statistic (yellow) with the axes rescaled, calculated from the top distribution. The new χ^2 distribution does not have extreme χ^2 values in the tails of the distribution and no longer contributes artificially high (or low) χ^2 values to the overall χ^2_{total} distribution. The high χ^2 values would artificially increase the acceptance probabilities for poor fits.

Property	Predicted value	χ^2_{pseudo} distribution	χ^2 distribution
μ	18.00	15.37	18.02
σ	6.00	19.53	6.02
γ_1	0.66	2.99	0.69

Table 3.2: Typical χ^2_{pseudo} distribution properties of *XMM-Newton* data to 2 decimal places compared to a real χ^2 distribution (predicted value) and the transformed χ^2 value.

repeating typically 1000-3000 times to construct an overall distribution spanning all segments. However given that each segment χ_{pseudo}^2 is drawn from a different distribution with different distortions, there is no reason to assume that the combination of χ_{pseudo}^2 values will contribute equal weight to the final distribution of χ_{total}^2 , compare Fig. 3.6 and Fig. 3.7. Given the strong correlation between bins in the simulated high-frequency PSD, the final distribution, with which the observed periodogram is tested, is indeed distorted by this data and contains the characteristic extended tail first observed in Fig. 3.7, since the distributions were implicitly assumed to have χ^2 properties. The bias is caused by the correlated high-frequency bins contributing as if all N data points were linearly independent, which is not the case - in the most extreme manifestation the entire high-frequency periodogram would only contribute to the distribution as *one* independent point! Due to the abundance of very poor fits to the data (i.e. the tail of the distribution), which results from correlated bins, some model fits will produce higher acceptance probabilities than they should ordinarily generate. That is, there are disproportionately many points with χ_{pseudo}^2 values worse than the measured value, which enhance the acceptance probability determination. These artificially high acceptance probabilities will also cause incorrect error estimates on the best fit model parameters, since the rejection probability levels are less well constrained. Moreover, the acceptance probability is dependent upon the amount of red-noise leak and as a consequence the original version of PSRESP is predisposed to favour steeper high-frequency slopes.

I have modified PSRESP to correct for the dependence of acceptance probabilities on red-noise leak, since well-defined acceptance probabilities should not depend on the amount of distortion. Rather than determine a weighting factor for each segment distribution in the final distribution, it is easier to construct a standardised distribution based on probability percentiles. The following new steps are implemented into PSRESP:

- Randomly select a realisation from each periodogram segment to form an “observation”, calculate the χ_{total}^2 statistic against the distorted model average and minimise by cycling through model normalisations.
- When the minimised χ_{total}^2 is determined, record the χ_{pseudo}^2 value contributed by each segment periodogram to the minimised χ_{total}^2 .
- Repeat the above steps M times to construct the minimised χ_{total}^2 distribution over all segments. All segments have be fitted simultaneously, but the individual contributions to χ_{total}^2 are also stored.

Each segment in the broadband periodogram has an associated χ_{pseudo}^2 distribution which is ordered to determine the probability percentile of each realisation within the distribution. For each realisation in each segment, the corresponding percentile is used to determine a χ^2 value i.e. using the standard χ^2 distribution for df , where df is equal to the number of bins. Thus, for any minimised χ_{pseudo}^2 in any segment there exists a corresponding χ^2 value. The χ^2 statistic *cannot* be directly evaluated, it can only be determined from its correspondence to the percentile of the observed χ_{pseudo}^2 statistic i.e. there exists a mapping between χ_{pseudo}^2 and χ^2 . The following steps are performed to fit data using the χ^2 goodness-of-fit statistic:

- Select distorted model average PSD normalisation. Randomly select a realisation from each periodogram segment to form an “observation” and measure χ_{pseudo}^2 . Using the χ_{pseudo}^2 and χ^2 distributions from each segment, linearly interpolate to find the corresponding χ^2 value. Sum χ^2 values over each segment to produce a χ_{total}^2 , which has crucially been produced from χ^2 and not χ_{pseudo}^2 .
- Record χ_{total}^2 value, choose a new normalisation and repeat the above step. Locate the minimum χ_{total}^2 value and record it as part of the overall minimised distribution over all segments. The normalisation that minimises χ^2 will probably not minimise the χ_{pseudo}^2 value.
- Repeat the above steps M times to construct the minimised χ_{total}^2 distribution based on minimised χ^2 values³. All segments have been fitted and minimised simultaneously according to the new goodness-of-fit statistic i.e. the overall distribution is not biased towards any segment.
- Use the observed data and compare the χ_{total}^2 value (using above steps) against the overall distribution of minimised values in the previous step. The resultant percentile defines the acceptance probability as previously done.

The crucial difference in the new algorithm is that the combination of χ^2 from each segment, to produce the χ_{total}^2 , no longer contains a contribution due exceptionally large χ_{pseudo}^2 values in the tail of the high-frequency data distributions. In the bottom plots of Fig. 3.6 and Fig. 3.7 I show the corresponding χ^2 distributions for comparison to the χ_{pseudo}^2 distribution for identical data. It is clear that the new distributions

³The M combinations are not necessarily the same of the previous M combinations, this process is random

Version	ν_B (Hz)	α	Acceptance
High-frequency errors only	$7.6^{+18.0}_{-6.1} \times 10^{-6}$	$2.6^{+1.2}_{-0.4}$	0.45
Complete	$3.4^{+8.0}_{-1.9} \times 10^{-6}$	$2.6^{+1.2}_{-0.4}$	0.40

Table 3.3: Here I show the best fit parameters and acceptance probability for the periodogram presented in Fig. 3.2 for comparison between the two versions of PSRESP described in the text. Errors determined by the 90 per cent rejection probability levels.

in Fig. 3.6 and Fig. 3.7 do not possess the extended tail (which is a function of red-noise leak) to artificially enhance the acceptance probabilities. The statistics of the transformed χ^2 distributions for both *RXTE* and *XMM-Newton* data are presented in Table 3.1 and Table 3.2 for comparison. Since the distributions have been standardised, the resultant acceptance probabilities are no longer a function of the sampling distortions; moreover, the rejection probability levels and corresponding errors on model parameters will be correct.

The new method was tested against the original method of determining acceptance probabilities in PSRESP (Uttley *et al.*, 2002) to determine the effect on the software. The two versions of the code tested both use the improvement to high-frequency data error determination described in Section 3.3.1, but only one code uses the improvement to the acceptance probability calculation described in this Section. I took the periodogram data of Fig. 3.2 to test the new software. As previously described, the above change to PSRESP should yield more constrained rejection probability levels and a lower acceptance probability *if* the observed data is expected to be strongly affected by red-noise leak i.e. the acceptance probability should no longer be a function of the amount of distortion.

I fitted the periodogram of NGC 3783 (observation, data, and subsequent analysis will be differed until the next chapter) by a single-bend power-law model (with a fixed low slope of 0.8) using both the Uttley *et al.* (2002) version of PSRESP and the version presented in this work. The best fit parameters are presented in Table 3.3 and rejection probability levels for the bend-frequency and high-frequency slope are presented in Fig. 3.8.

While the new computation takes longer it does provide improvement over the version of PSRESP without the new acceptance probability algorithm. For simulations where red-noise leak is expected to have a large effect on the result, the high-frequency errors only version of PSRESP produces an acceptance probability ($P = 0.45$) which contains a contribution due to the amount of red-noise leak. The

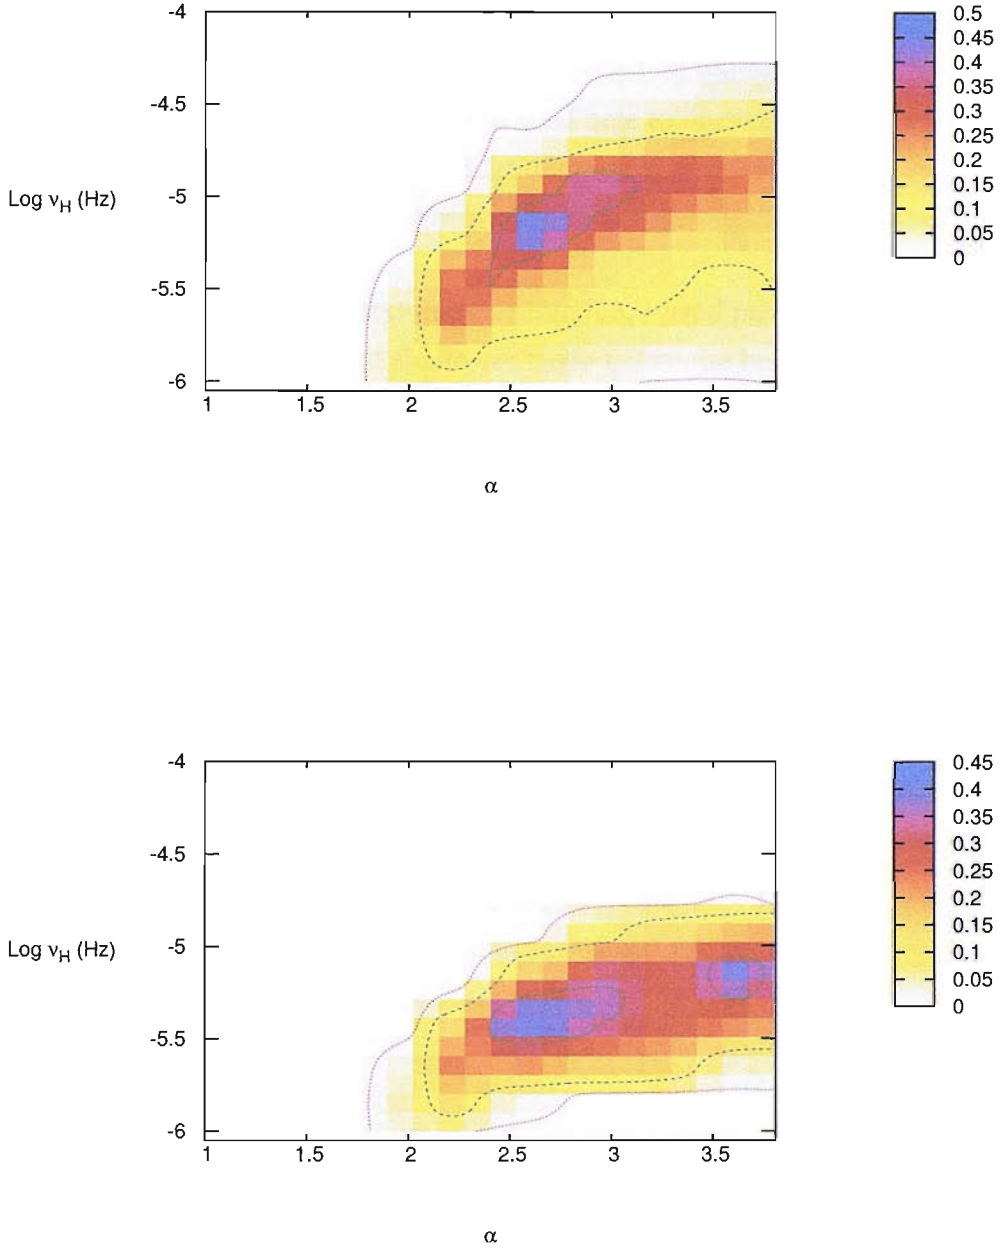


Figure 3.8: The top plot displays the rejection probability levels for the PSRESP version without the new acceptance probability calculation. The bottom plot displays the rejection probability levels for the complete version of PSRESP presented in this work. It is clear that the rejection probability levels for the bottom plot are more tightly constrained. See main text for details.

acceptance probability obtained from the complete version of PSRESP is smaller ($P = 0.40$ - by an amount dependent upon the amount of red-noise leak), which is to be expected given that we wish to measure the goodness-of-fit only. The rejection probability levels are significantly better constrained by the new software (see Fig. 3.8). It is also encouraging to find that the best-fitting parameters only change by two grid positions in bend-frequency and are consistent within error bounds. The broadband periodogram of NGC 3783 was used here to illustrate the difference between two versions of PSRESP and will be thoroughly and exhaustively examined in the next chapter; I leave further power spectral analysis of the periodogram until then.

As another test of the new PSRESP implementation, I used periodogram data with a very-high break frequency and shallower slope below the break. Red-noise leak would not be expected to cause significant distortions due to the lack of long time-scale trends in the data; thus, the difference between power spectral fitting results for both of the aforementioned versions of PSRESP is expected to be marginal. The difference between the versions is indeed marginal for data weakly effected by red-noise leak, thereby suggesting the new version of PSRESP is working as intended.

3.4 Computation, optimisation and Iridis2

PSRESP provides a robust and accurate way to determine AGN power spectra; however, due to the stochastic nature of AGN light curves, the power spectral response method employed in PSRESP is computationally gruelling. The requirement that a great many periodograms must be simulated based upon an underlying power-spectral model necessitates computational optimisations.

3.4.1 FFT

The first obvious optimisation involves utilising the fast Fourier transform (FFT) (Press *et al.*, 1992) instead of the discrete Fourier transform (DFT) where possible. The fast Fourier transform can produce mathematically identical results to the discrete Fourier transform. However, the fast Fourier transform is subject to a fairly strict constraint i.e. the size of the data set to be transformed. To achieve the maximum benefit from a fast Fourier transform the initial data set must contain 2^N points where $N \in \mathbb{Z}^+$, since it reduces the number of computations for N points from $\sim 2N^2$ to $\sim 2N \log_2 N$, where \log_2 is the base-2 logarithm. A fast Fourier transform is already adopted in the simulation of light curves, since the Timmer and König

(1995) prescription naturally uses the inverse fast Fourier transform as the length of simulated light curve can be arbitrarily adjusted. The observed *RXTE* data and sampled simulated light curves need to be evaluated using a discrete Fourier transform, since the 2^N requirement can be tricky to adhere to when dealing with real *RXTE* data. However, as discussed in Section 3.3, a new feature of PSRESP is the ability to simulate the high-frequency periodogram (usually covered by *XMM-Newton* data), which is limited at the highest frequencies by Poisson noise. *XMM-Newton* data is continuous and is limited at shortest time-scales by Poisson noise, so there is freedom to choose the sampling interval when reducing *XMM-Newton* data and generating the X-ray light curves i.e. *XMM-Newton* light curves can be rebinned without being affected by aliasing at high frequencies. Provided there is sufficient data, a prudent choice of sampling interval with which to make the *XMM-Newton* light curves will produce an observed light curve that has 2^N points - suitable for use with an FFT. The Nyquist frequency will translate depending upon the choice of sampling interval, but the part of the periodogram containing source power will not be affected if the sampling interval is chosen carefully enough. To optimise up computation within PSRESP I have implemented an optional FFT feature to exploit the aforementioned property of *XMM-Newton* data; however, no such feature can be implemented for the *RXTE* data. Given that *XMM-Newton* data can often be much longer than *RXTE* data (in terms of number of data points) this feature represents a significant working-speed increase within PSRESP, without changing the resultant periodogram.

3.4.2 Aliasing and red-noise leak

RXTE observations used throughout this work (e.g Fig. 3.1) are constructed from single snapshot observations of ~ 1 ks. To thoroughly constrain the effect of aliasing with the observed light curves the simulated light curves should be simulated with 1ks temporal resolution, but this requirement is computationally prohibitive. However, the largest contribution of aliased power comes from the frequencies close to the Nyquist, just outside the measured frequency range. At frequencies much greater than the Nyquist the power contribution is minimal, compared to those near the Nyquist, and could in fact be approximated by an average power contribution. Therefore it is sufficient (to first order) to simulate light curves with a temporal resolution at least 0.1 times the observed sampling interval. In order to determine the second-order effect of aliased power from frequencies far away from the Nyquist frequency an analytical approximation is used. We use the underlying PSD model

to determine the integrated power between the highest simulated resolution and the snapshot (1ks) resolution. The integrated power is an additive component to the distorted model average PSD. The total amount of aliased power cannot be determined in this way, since aliased power is considerably influenced by the stochastic realisations near the Nyquist, and thus would be poorly approximated by the average power approximation.

I have already described the effect of red-noise leak on the measured periodogram. To mimic these effects in a simulated light curve it is necessary to simulate these light curves to 10 times the duration of the observed data. Rather than simulating multiple, N , light curves at 10 times the duration of the observed light curve, it is computationally favourable to simulate a single light curve of length NT , where T is the observed light curve duration. The simulated light curve is divided up into N light curves of the duration T . Variations on time-scales comparable to the observed light curve contribute the most red-noise leak (due to the break in the power spectrum time-scales much longer than T contribute very little), so simulating a single light curve to be N times the length of the observed light curve is approximately indistinguishable from simulating N light curves to the same length as the observed light curve. Simulating N light curves with a minimum of 10 times the duration of the observed light curve is computationally intensive and produces a lot of unused data.

3.4.3 Optimisation

It is also important to consider the method by which the model parameters will be optimised. The most reliable and most computationally intensive method is the standard grid search, which requires a list of parameter combinations spanning the parameter space. The grid search is costly as a minimisation routine since it determines the best fit parameters by brute force, locating the minimum in the input parameter space by searching every grid point. Moreover, the grid search is not sensitive to features in the parameter space on a scale less than the separation between the grid points. Another problem with the grid search method is it becomes computationally prohibitive (within the context of PSRESP) very rapidly when we wish to explore a larger parameter space or add more dimensions e.g 3 dimensions and 10 values for each dimension require 1000 function evaluations to minimise, with each function evaluation taking between 10 and 100s depending on the input data. However, the grid search is an extremely effective method to determine the errors on the best fit parameters i.e. rejection probability levels are constructed and

errors determined in the standard fashion, as previously described.

I investigated whether PSRESP would benefit from an optimisation routine such as downhill simplex (Nelder and Mead, 1965) or simulated annealing (Kirkpatrick *et al.*, 1983; Cerny, 1985). The downhill simplex algorithm is a fast algorithm that approximately finds a *locally* optimal solution (best fit model parameters) to observed PSD using only function evaluations, not derivatives. The downhill simplex method involves defining a geometrical object consisting of $N + 1$ vertices in N dimensions; for example, in two dimensions the simplex is a triangle and in three dimensions the simplex is a tetrahedron. The vertices forming the simplex are also required to be nondegenerate so that the simplex encloses a finite N dimensional volume. The downhill simplex method evaluates the objective function at each vertex and essentially moves the vertex with the largest function evaluation through the opposite face of the simplex of the lowest evaluation. The simplex is also permitted to contract and expand in all directions. The iteration is terminated when the objective function changes by an amount smaller than some predefined tolerance. The downhill simplex method is often used as a means to quickly determine a solution rather than rigorously locate the true globally optimised solution; moreover, the true power of the downhill simplex method becomes apparent for high-dimensional parameter spaces.

To test the downhill simplex method I fitted the periodogram data of Ark 564 using both the grid search and downhill simplex method with PSRESP as the objective function (description of the data will be deferred until Section 4.2 when Ark 564 is discussed in detail). The initial simplex was defined as follows

$$\begin{pmatrix} 5.0 \times 10^{-4} & 1.0 \times 10^{-6} & 2.0 & 1.0 \\ 1.0 \times 10^{-2} & 1.0 \times 10^{-6} & 2.0 & 1.0 \\ 5.0 \times 10^{-4} & 5.0 \times 10^{-6} & 2.0 & 1.0 \\ 5.0 \times 10^{-4} & 1.0 \times 10^{-6} & 2.5 & 1.0 \\ 5.0 \times 10^{-4} & 1.0 \times 10^{-6} & 2.0 & 1.5 \end{pmatrix}$$

where the entries define the geometrical shape in four dimensions. Column 1 through 4 indicates the coordinates for each vertex of the starting simplex and denote the high-frequency break, low-frequency break, high-frequency slope, and the intermediate-frequency slope, respectively. Table 3.4 presents the best fit double-bend power-law PSD model parameters found by the downhill simplex method (Nelder and Mead, 1965), and the standard grid search. The underlying model is defined in four dimensions.

Method	v_H (Hz)	v_L (Hz)	α_H	α_L	Acceptance probability	Steps	Runtime (seconds)
Downhill Simplex	2.5×10^{-3}	9.7×10^{-7}	4.2	1.2	0.62	277	4608
Grid search results	2.4×10^{-3}	7.5×10^{-7}	4.2	1.2	0.75	14784	~410754

Table 3.4: Comparison between the best fit parameters of a grid search and the downhill simplex for a 4 parameter model. It is clear that the downhill simplex method can determine best fit parameters consistent with more rigorous method of a grid search. However, the downhill simplex does not produce any diagnostic output to verify that the optimised solution is not local, unlike the grid search. Without the grid search results it would be impossible to tell if the downhill simplex optimised parameters are a global solution without reinitialising the simplex from new starting positions.

Despite the initial appeal of the downhill simplex method due to the runtime required, there are significant complications that prevent the downhill simplex method from being the first choice optimisation routine. In particular:

- No diagnostic output to thoroughly determine if the optimised solution is local. The downhill simplex method would need to be initialised from several different positions in the parameter space to determine if the original optimised solution is a global solution. Only when several minimisation routines agree on an optimised solution can a globally optimised solution be *assumed*.
- The downhill simplex results presented in Table 3.4 are close to those provided by the grid search, but I had prior knowledge of the solution from the grid search and several simplex restarts were required to find it. Prior knowledge of the best fit solution will not generally be available.
- The downhill simplex method does not provide complete rejection probability levels, which means error estimation of the optimised parameter is not possible i.e. it does not probe the entire parameter space.

The first two points of the above problems can be very troublesome for less well behaved objective functions i.e. when the probability space associated with observed data is pathological; for example, the simplex restarts used to find the solution in Table 3.4 did not regularly find the same solution. Further tests using a different initial simplex and with models of more parameters suggest that the resultant optimised solution is acutely dependent on the initial values; thus, the reliability of the downhill simplex method can never be guaranteed, since the smoothness of the underlying topology is unknown. To illustrate this important point Fig. 3.9 is the probability surface for two parameters from a simpler single-bend power-law PSD model for a different AGN (Fairall 9; analysis is deferred to Chapter 5), even for a more simple model the probability surface is suitably irregular to cause serious problems for any minimisation routine liable to get trapped by a locally optimised solution. The downhill simplex method is an fast optimisation routine, but the test results seem to imply that it could provide useful indicative optimisation information only, and it is not used to produce final optimisation results.

I also considered the the simulated annealing algorithm which is based on the downhill simplex and takes longer to complete, but it is less prone to determine a local locally optimised solutions. Like the downhill simplex algorithm, simulated annealing can be faster than the grid search, but it does not provide any means to determine




Figure 3.9: A typical probability surface obtained in fitting AGN data by a simple single-bend power-law PSD model. It is clear that even for simple models the resultant topology of the parameter space to be searched is sufficiently complex to baffle any optimisation routine that is prone to find locally optimised solutions. The parameter space shown here is from a grid search fit of Fairall 9 data (see Chapter 5). The parameters ν and α_H are the bend-frequency and high-frequency slope of a single-bend power-law PSD model, respectively.

error bounds on the optimised parameters. Moreover, both simulated annealing and downhill simplex methods are designed for problems where the objective function is defined in a higher-dimensional space (~ 20 dimensions) i.e. where grid searches would be virtually impossible. The only other optimisation method considered was the Levenberg-Marquardt algorithm (Levenberg, K., 1944; Marquardt, D., 1963) for minimising a general non-linear function. The Levenberg-Marquardt algorithm can provide information corresponding to the errors on the measured parameters, but it requires more than one function evaluation at each tested grid point (for derivative determination), thereby increasing the runtime for PSRESP since function evaluations are expensive.

The diagnostic output (e.g. complete rejection probability levels) from a grid search were deemed too important to sacrifice for an increase in speed; moreover, *Iridis2* (covered in the next section) provides a means to reduce the runtime of a grid search quite significantly. The future of PSRESP and optimisation will be covered in detail in Chapter 6.

3.4.4 Iridis2

For many of the AGN found within this work, a simple unbroken or single-bend power-law PSD suffices when fitting their PSD; however, more complicated models (e.g. double-bend power-law PSD model) with more dimensions very quickly becomes computationally prohibitive for modest parameter spaces. In Chapter 4 when I describe the PSDs of NGC 3783 and Ark 564 with more complicated PSD models, I utilise the Iridis2 computer cluster at the University of Southampton to break up the required parameter space into several smaller computations. Iridis2 contains over a 1000 processor-cores, some of which I used to determine some of the PSDs presented in this work. Determination of AGN PSDs (with best fit model parameters) can be a time consuming process if the input parameter space is large; however, splitting-up the parameter space in N smaller spaces along one dimension and processing these N parameter spaces independently greatly reduces the time required (by a factor of the number of points in the parallelised dimension) to generate an AGN PSD for complicated models. The results from the N PSRESP simulations are combined to form a final result identical to one that would be obtained from a full unparallelised fit over the entire parameter space. In the case of the more complicated models such as those model used in the analysis of Ark 564 (see Chapter 4) the full simulations of a desktop machine would have taken approximately 20 days to complete, whereas they only took approximately 20 hours to complete on Iridis2 by testing the input parameter space in segments.

3.5 Assumptions and limitations of PSRESP

The power spectral response method is a well-used and robust technique for PSD estimation; however, it is important to recognise the inherent limitations of PSRESP as well as the underlying assumptions that underpin our Monte Carlo method.

An assumption that has not been explicitly mentioned thus far is the assumed linearity of AGN light curve data. The light curve simulation technique of Timmer and König (1995) produces linear light curves with zero mean. However, according to Uttley *et al.* (2005) the RMS-flux relation (mentioned briefly in Section 2.4) found in the X-ray light curves of BHXRBs and AGN imply that the X-ray light curves themselves are inherently non-linear. Moreover, the non-linearity has exponential form, provided the RMS-flux relation applies to all the observed time-scales.

Fig. 3.10 represents the exponentiated version (bottom) of the linear light curve (top) i.e. each count value is replaced by an exponentiated value. Their forms are




Figure 3.10: The two light curves presented here have been generated with the same seed, except the bottom figure has been exponentiated. That is, the flux in the linear (top) light curve has been exponentiated to produce the bottom non-linear plot. It is clear that the statistical properties of the exponentiated light curve have changed i.e. the mean is no longer zero and there is a greater fluctuation in flux.

fundamentally and statistically different since exponentiation is a non-linear process on the light curve. It would be preferential to use non-linear light curves since they display a closer resemblance to real X-ray light curve (Uttley *et al.*, 2005); however, it is not practical to exponentiate the simulated linear light curves since the operation of exponentiation changes the *shape* of the underlying PSD in a non-linear way. In particular the high-frequency PSD changes in a way that is impossible to determine analytically i.e. the underlying PSD model parameters would change in a non-linear and unpredictable way if the linear light curves are exponentiated. For the majority of AGN observations the linear light curve simulation has proved sufficient, since the light curves often appear to be broadly linear. However, this limitation should be noted since it not immediately clear whether this approximation extends to blazar-like sources where the observed X-ray light curves are highly non-linear (see Chapter 5 and 3C 279 for an example).

PSRESP can produce and fit broadband periodograms from several input light curves covering complementary time-scales taken at different epochs. The ability to pro-

duce a single broadband periodogram spanning several decades in frequency relies on the assumption that AGN PSDs are stationary on the time-scales probed by the long time-scale monitoring campaigns of *RXTE*. Verification of this assumption lies in observational evidence, in that two AGN (NGC 4051 and MCG-6-30-15) monitored for over 11 years do not exhibit any deviation away from a stationary PSD thus far. As we shall see in Chapter 5 NGC 4051 and MCG-6-30-15 are very well fit by a single-bend power-law over several decades in frequency to the longest measured time-scales. This implies the longest time-scale observations are still consistent with some of the shortest, in that they are well fit by the same PSD model parameters thereby implying stationarity. These *RXTE* light curves are 11 years in duration, and it is still possible that stationarity is violated on longer time-scales. Non-linearity will be considered again in Chapter 6 when I discuss the future of PSRESP.

3.6 Overview

In this chapter I have described the basic principle by which we determine the underlying power spectrum from a periodogram distorted by a window function - namely the power spectral response method. I have introduced how the response method is implemented within the framework of Monte Carlo simulations to provide a robust estimation of the underlying model parameters. The standard method of PSRESP and the improvements I have made to the original work of Uttley *et al.* (2002) were also described. The improvements include a modification of the very-high frequency ($> 10^5$ Hz) PSD determination to be self-consistently evaluated with the broadband of the PSD. Moreover, I introduce a new Poisson noise calculation to provide a more robust PSD estimate and error determination. Not only can PSRESP simulate and determine the very-high frequency PSD, but it now contains a feature whereby it can combine several observed periodograms into a single PSD estimate. This feature is particularly useful if there exist several *XMM-Newton* observations separated only by orbital gaps. The segments are simulated and binned together to provide a more robust PSD estimate. This feature can also deal with correlated or uncorrelated light curves. Perhaps the most important improvement to PSRESP is the more robust determination of the model acceptance probabilities in the presence of sampling biases such as red-noise leak and aliasing. The goodness-of-fit statistic, χ^2_{pseudo} , is transformed to a χ^2 to remove the dependence on the distortions within the acceptance probability is a major improvement, even if it does require signif-

ificant additional computation. The improved acceptance probability method was tested using existing AGN data, and performed as intended. The computational requirements of PSRESP were considered, and possible optimisation techniques such as the downhill simplex method were also briefly tested and reviewed. Optimisation of PSRESP is highly non-trivial and is discussed again in Chapter 6 when I consider the future of PSRESP.

The great tragedy of Science - the slaying of a beautiful hypothesis by an ugly fact.

THOMAS H. HUXLEY (1825 – 1895)

4

The broadband power spectrum of NGC 3783 and Ark 564

In this chapter, I apply the improved PSRESP technique to determine the broadband power spectrum of two AGN: the Seyfert 1 galaxy NGC 3783 and the narrow line Seyfert 1 galaxy Ark 564. Previous observations with *RXTE* have suggested that the power spectral density of NGC 3783 flattens to a near zero slope at low frequencies, in a similar manner to that of black hole X-ray binary systems in the low/hard state. The low radio flux emitted by the object NGC 3783, however, is inconsistent with a low/hard state interpretation. The accretion rate of NGC 3783 (~ 7 per cent of the Eddington rate) is similar to that of other active galactic nuclei with ‘soft’-state PSDs and higher than that at which the BHXRB Cyg X-1, with which AGN are often compared, changes between low/hard and high/soft states (~ 2 per cent of the Eddington rate). If NGC 3783 really does have a ‘hard’-state PSD, it would be quite unusual and would indicate that AGN and BHXRBs are not quite as similar as we currently believe. Here I present an improved X-ray PSD of NGC 3783, spanning from $\sim 10^{-8}$ to $\sim 10^{-3}$ Hz, based on considerably extended (5.5 yr) *RXTE* observations combined with two orbits of continuous observation by *XMM-Newton*. I show that this PSD is, in fact, well fitted by a ‘soft’ state model which has only one break at high frequencies. Although a ‘hard’-state model can also fit the data,

the improvement in fit by adding a second break at low frequency is not significant. Thus NGC 3783 is not unusual. These results leave Ark 564 as the only AGN which shows a second break at low frequencies. The well-studied power spectrum of Ark 564, unlike the other AGN, is clearly band-limited and I show that it can be well described by the sum of two Lorentzian-shaped components.

4.1 Timing analysis and accretion state determination of NGC 3783

The PSDs of AGN can typically be represented by red-noise type power laws (i.e. $P(\nu)$, the power at frequency ν , $\propto \nu^{-\alpha}$ where $\alpha \sim 1$) with a bend or break (to $\alpha \leq 2$) at a characteristic frequency. The time-scale, corresponding to the characteristic bend-frequency, scales approximately linearly with black hole mass from AGN to BHXRBs (McHardy, 1988; Edelson and Nandra, 1999; Uttley *et al.*, 2002, 2005; Markowitz *et al.*, 2003; McHardy *et al.*, 2004, 2005), albeit with some scatter. However, the scatter is entirely accounted for by variations in accretion rate, allowing scaling between AGN and BHXRBs on time-scales from years to ms (McHardy *et al.*, 2006).

BHXRBs are observed in a number of distinct X-ray spectral states which also have distinct X-ray timing properties. Two common states are the low/hard (hereafter ‘hard’) and high/soft (hereafter ‘soft’). In the hard state, the energy-spectrum is dominated by a highly variable power-law component and the PSDs are well fitted by multiple broad Lorentzians. For use in AGN, where signal-to-noise is lower than in BHXRBs, this PSD shape can be approximated by a double-bend power-law with slopes $\alpha = 0, 1$ and 2 , from low to high frequency, where the high- and low-frequency bends correspond to the strongest peaks in the Lorentzian parameterisation. The breaks are typically separated by only one to two decades in frequency. In the soft state, the energy spectrum is dominated by an approximately constant thermal disk component which extends into the X-ray band in BHXRBs but which in AGN is shifted down to the optical/UV band. Therefore, a meaningful comparison between the PSDs of soft state BHXRB and AGN can only be made in cases where the BHXRB power-law emission is strong enough to show significant variability. Such BHXRBs are rare and the best example is Cyg X-1 which shows a ‘ $1/f$ ’ PSD over many decades of frequencies (Reig *et al.*, 2002). The soft state is distinguished from the hard state by having only one, high frequency, break in this power-law, from slope 1 to 2.

It has been suggested that this pure simple broken or cut-off power-law PSD shape is unique to the soft state of Cyg X-1, which is a persistent source. However in transient BHXRBs with similar X-ray spectra, the power-law PSD component may be seen in combination with broad Lorentzian features (Done and Gierliński, 2005). Axelsson *et al.* (2006) also note that a mixed power-law plus Lorentzian PSD is also present in Cyg X-1 in lower luminosity, harder spectral states, but as the luminosity rises the Lorentzian features weaken and the power-law PSD component strengthens until, in the softest state, it completely dominates. Since the softest spectral states of transient BHXRBs are dominated by constant disk emission we cannot determine whether they show a similar PSD shape to Cyg X-1.

However, a direct comparison of transient BHXRBs and Cyg X-1 is complex, since the transients show much larger luminosity changes, and complex hysteresis effects in spectral hardness versus luminosity (e.g. Homan *et al.* 2001, Belloni *et al.* 2005) which are not seen in Cyg X-1. Therefore it is not clear that one can compare timing properties between Cyg X-1 and transient BHXRBs simply as a function of observed X-ray spectrum.

None the less, it is still interesting that the X-ray spectrum of Cyg X-1 never becomes totally disk-dominated, and always contains a relatively strong variable component whose PSD resembles that of X-ray bright AGN. If variability originates, at least partly, in the disk, so power spectral shape is related to the disk structure, that structure might be severely disrupted during outbursts, thereby suggesting a possible difference between the persistent Cyg X-1 and the transient BHXRB sources. The similarities between the PSDs of Cyg X-1 and AGN may also be related to the possible similarities in accretion flows between AGN and Cyg X-1 noted by Done and Gierliński (2005).

To date, NGC 3783 and the Narrow Line Seyfert 1 Galaxy (NLS1) Ark 564 are the only AGN with suggested second, low-frequency breaks in their PSDs (i.e. similar to low/hard BHXRBs) and are both commonly referred to as being unusual (e.g. Done and Gierliński 2005). The power spectral evidence for a second break in the case of Ark 564 is very strong (Pounds *et al.* 2001; Papadakis *et al.* 2002; Markowitz *et al.* 2003; McHardy *et al.* 2007). Of all the AGN with good timing data, Ark 564 shows the highest accretion-rate (possibly super-Eddington) so it would not be surprising if it were in an unusual state, e.g. the ‘very high’ state where the PSD, in BHXRBs, also displays two distinct breaks. The accretion rate of NGC 3783, on the other hand, is similar to those of AGN with proven soft-state PSDs (e.g. NGC 3227, NGC 4051 McHardy *et al.* 2004, MCG-6-30-15 McHardy *et al.* 2005), and in particular it is radio quiet (e.g. Reynolds 1997). In the hard state, BHXRBs are strong

radio sources whereas in the soft state the radio emission is quenched (Corbel *et al.*, 2000; Fender, 2001; Körding *et al.*, 2006b). I also note that NGC 3783 has a more moderate accretion rate (~ 7 per cent) than Ark 564, and more similar to the other AGN mentioned above. Cyg X-1 changes from the hard to the soft state at around 2 per cent of the Eddington accretion rate (i.e. $\dot{m}_E = 0.02$) (Pottschmidt *et al.* 2003; Wilms *et al.* 2006; Axelsson *et al.* 2006). These two facts do not lie easily with a hard state identification of NGC 3783. Thus it would be surprising, and might indicate that our current ideas regarding the scaling between AGN and BHXRBs are not entirely correct, if NGC 3783 were proven to have a hard state PSD. It is therefore important to determine whether NGC 3783 does have a second, low frequency, break in its PSD or not.

Timing analysis of NGC 3783 data, until 11 April 2002, is reported by Markowitz *et al.* (2003), who found evidence for a break in the 2 – 10 keV PSD of NGC 3783 at 4×10^{-6} Hz and found provisional evidence for a second lower-frequency break at $\sim 2 \times 10^{-7}$ Hz. Specifically, Markowitz *et al.* (2003) rejected the possibility that the PSD is described by a single-break power-law with low-frequency slope 1, similar to other AGN, at the 98 per cent confidence level. In this section I re-investigate the evidence for the second break in the PSD of NGC 3783, using new long-term monitoring data that covers the frequency range where the break appears to be. By including additional *RXTE* archival data spanning several years, along with short time-scale observations by *XMM-Newton*, I will demonstrate that the improved PSD is perfectly compatible with a single-bend power-law, consistent with the behaviour of the other moderately-accreting Seyferts.

4.1.1 Observations and data reduction

From 1999 to 2006, NGC 3783 has been the target of various monitoring campaigns with *RXTE*. These campaigns have consisted of short, ~ 1 ks duration, observations with the proportional counter array (PCA, Zhang *et al.* 1993). The overall X-ray light curve of NGC 3783 can be separated out into three regions of different sampling patterns that can be used to construct a broadband power spectrum. NGC 3783 was observed every ~ 4 days from 02 January 1999 until 11 April 2002 for ~ 1200 days. There was a ~ 20 -day period from 20 February 2001 to 12 March 2001 where NGC 3783 was observed every ~ 3.2 hours. I obtained this data from the public *RXTE* archive and combined it with data obtained from our own monitoring campaign that began on 28 February 2004 until 13 September 2006 with a sampling interval of 2 days. As the gap in the entire *RXTE* light curves is a significant frac-




Figure 4.1: *RXTE* long-term (top) and intensive (bottom) sampling light curve of NGC 3783 in the 2–10 keV band. The bottom light curve is measured in days since the start of the intensive sampling (20 February 2001) period.

tion of the overall duration, I will include data from each monitoring campaign as separate light curves in the fits.

A summary of the X-ray light curves used in our timing analysis of NGC 3783 are presented in Table 4.1.

The *RXTE* data described above was reduced using *FTOOLS* v6.0.2. The light curves were made using PCA STANDARD-2 data in the 2–10 keV energy range. I use data from the top layers of PCUs 0 and 2 up to 12 May 2000 and only top layer PCU 2 data from observations after this date. I only use the top layer of the PCUs to maximise the signal-to-noise ratio, since the background count rate increases substantially if subsequent layers are included. Data were not collected from PCUs 1, 3 and 4 since they often suffer repeated breakdown during on-source time. Data were selected according to the standard ‘good-time’ criteria, i.e. data were rejected if they were gathered within 10° from the Earth’s limb, within 30 minutes of the South Atlantic Anomaly (SAA), if the pointing offset was greater than 0.02° , or if the electron contamination was greater than 0.1.

The PCA has no capability to take simultaneous background measurements during on-source time, background data were estimated by using the L7 model for faint

Light curve	Sampling interval	Observation length	Date Range [MJD]
<i>RXTE</i> Long-term 1	~4.36 days	1194.6 days	51180.5–52375.1
<i>RXTE</i> Long-term 2	~2.1 days	928.3 days	53063.4–53991.6
<i>RXTE</i> Intense monitoring	~3.2 hours	19.9 days	51960.1–51980.1
<i>XMM-Newton</i> observations (2 orbits)	200-s	3.2 days	52260.8–52264.0

Table 4.1: Summary of the *RXTE* and *XMM-Newton* light curves used in the analysis of NGC 3783, including their sampling frequency and date range in modified Julian day.

sources using PCABACKEST v3.0 to simulate the effect of a diffuse X-ray background, particle induced background, SAA transitions and induced radioactivity. The response matrices for each PCA observation were calculated using PCARSP v10.1. The final 2 – 10 keV fluxes were calculated using XSPEC v12.2.1 by fitting a power-law with galactic absorption to the PHA data.

NGC 3783 was observed by *XMM-Newton* during revolutions 371 and 372, between 2001 December 17 and 2001 December 21. Temporal analysis of these data were first presented by Markowitz (2005) who discusses the coherence, frequency-dependent phase lags, and variation of high frequency PSD slope with energy. Here I use these data to constrain the high frequency part of the overall long and short timescale PSD. I used data from the European Photon Imaging Cameras (EPIC) PN and MOS2 instruments, which were operated in imaging mode. MOS1 was operated in Fast Uncompressed Mode and I do not use those data here. The PN camera was operated in Small Window mode, using the medium filter. Source photons were extracted from a circular region of 40 arcsec radius and the background was selected from a source-free region of equal area on the same chip. I selected single and double events, with quality flag=0. The MOS2 camera was operated in the Full Window mode, using the medium filter. I extracted source and background photons using the same procedure as for the PN data and selected single, double, triple and quadruple events. These data showed no serious pile-up when tested with the XMM-SAS task *epatplot*.

I constructed light curves, for each detector and orbit, in the 0.2–2, 2–10 and 4–10 keV energy bands. I filled in the ~ 5 ks gap in the middle of orbit 371 light curves, and some other much smaller gaps, by interpolation and added Poisson noise. The resulting PN and MOS2 continuous light curves were then combined to produce the final light curves for each orbit. The combined, background subtracted, average count rates in the 0.2–10 keV band were 11.8 c/s for orbit 371 and 15.8 c/s for orbit 372, and the 0.2–2 keV combined light curve is shown in Fig. 4.2. Poisson noise dominates the PSD on timescales shorter than 1000s, so the light curves were binned into 200s bins.

4.1.2 Combining *RXTE* and *XMM-Newton* data

To determine the PSD over the largest possible frequency range I combine the *RXTE* and *XMM-Newton* data. In BHXRBs the break-frequency and slope of the PSD below the break appear to be independent of the chosen energy band (Cui *et al.*, 1997; Churazov *et al.*, 2001; Nowak *et al.*, 1999; Revnivtsev *et al.*, 2000; McHardy




Figure 4.2: *XMM-Newton* X-ray light curve of NGC 3783 in the 0.2-10 keV band. The light curve is measured in days since 17 December 2001 and the gap is due to orbital constraints i.e. the plot above represents two orbits of *XMM-Newton* data.

et al., 2004). On the other hand, the PSD normalisation and the slope above the break are often energy-dependent (Markowitz, 2005). Therefore, when combining data from different instruments, it is preferable to use similar energy ranges. The *RXTE* data are in the 2–10 keV band and, for NGC 3783, that band has a median photon energy of 5.7 keV. The *XMM-Newton* band with the same median photon energy is 4.1–10 keV. However the count rate in that *XMM-Newton* band is low (2 c/s) so I only detect significant source power above the Poisson noise level at frequencies below 10^{-4} Hz. To probe higher frequencies I can use the 0.2–2 keV *XMM-Newton* data (8.8 c/s) but I must re-scale its PSD normalisation to that of the 4–10 keV PSD. I determined the scaling correction by producing PSDs in both energy bands and fitting the same bending power-law model to the noise-subtracted data. On the assumption that the PSD shape below the high frequency break is energy-independent, the combined *RXTE* 2–10 keV and *XMM-Newton* 0.2–2 keV PSD will then have the shape of the 0.2–2 keV PSD.

4.1.3 PSRESP and the power spectrum of NGC 3783

Prior to constructing and fitting the power spectrum of NGC 3783 with PSRESP, the component sampling pattern must be extracted from the *RXTE* and *XMM-Newton* data to produce separate light curves corresponding to the time-scale they were designed to measure. To produce a PSD that is sufficient to cover a broad range of frequencies, PSRESP will, in general, require three separate input light curves:

1. A long time-scale light curve. Usually constructed from an extended *RXTE* monitoring campaign to measure the low-frequency section of the PSD. The

long time-scale *RXTE* data is binned so that the resultant PSD estimates cover $\sim 10^{-9}\text{Hz}$ to $\sim 10^{-6}\text{Hz}$ in frequency.

2. A medium time-scale light curve. The medium time-scale light curves are usually extracted from the intensive sampling pattern found in most AGN monitoring campaigns (e.g. the 20 day observation of NGC 3783). The intensive sampling data measures the medium frequencies in the broadband PSD ($\sim 10^{-6}\text{Hz}$ to $\sim 10^{-5}\text{Hz}$).
3. A short time-scale light curve. *XMM-Newton* light curves make ideal short time-scale light curves, since they are continuous (negligible aliasing) and cover $\sim 10^{-5}\text{Hz}$ up to the frequencies dominated by Poisson noise. Moreover, *XMM-Newton* light curves cover a similar energy range to *RXTE*.

The above list of light curve inputs are the ideal and typical of many AGN observations, but it is not uncommon for AGN monitoring campaigns to be split-up into more component light curves - depending upon the sampling pattern. NGC 3783, for example, requires four input light curves: two long time-scale light curves from either side of the gap, the intensive sampling light curve and the *XMM-Newton* data¹. Some of the other AGN observations encountered in Chapter 5 will require more sophisticated partitioning.

The PSD estimates for NGC 3783 are binned with width of 1.3ν , where ν is the starting frequency, by taking the average of the log of power (Papadakis and Lawrence, 1993a). I require a minimum of 4 PSD estimates per bin. For each set of chosen underlying-PSD model parameters, I simulate red-noise light curves, as described by Timmer and König (1995). The *RXTE* light curves are simulated with time resolutions of 10.5 h, 5.0 h and 18 m for the first and second long time-scale and the medium time-scale light curves respectively. The simulated resolution, which is 10 times shorter than the typical sampling intervals of the real observations, given in column 2 of Table 4.1, is to take into account the effect of aliasing. These simulated light curves were resampled and binned to match the real NGC 3783 observations. *XMM-Newton* light curves were simulated with 200-s resolution, as at shorter time-scales the underlying variability power is negligible compared to the Poisson noise, so aliasing does not play a role. The Poisson noise level was not subtracted from the observed PSD but was added to the simulated PSDs. The simulated

¹The *XMM-Newton* data is technically two light curves separated by an orbital gap; however, since the X-ray flux is correlated in both light curves they are treated as a single input with their resultant periodograms binned together.

model average PSD is evaluated from this ensemble of PSD realisations, and the errors are assigned from the RMS spread of the realisations within a frequency bin. Each input light curve is normalised by its square mean flux, so that the resultant power spectrum is in fractional RMS squared units (see Chapter 2). The broadband periodogram of NGC 3783 is shown in Fig. 4.3. The different segments (and colours) represent the different input light curves: the blue, black and green lines represent the periodogram of the *RXTE* long time-scale observations and intensive monitoring, respectively. The red line is the periodogram of the two orbits of *XMM-Newton* data binned together. The bottom figure of Fig. 4.3 is the same periodogram plotted in $\nu \times P_\nu$ space. Throughout the rest of this work all power spectra will be plotted in $\nu \times P_\nu$ space units, since it is easier to visually identify power spectral features even though the two plots in Fig. 4.3 convey identical information.

Visual inspection of Fig. 4.3 suggests that the periodogram of NGC 3783 does indeed flatten at low frequencies, which is characteristic of many AGN periodograms. However, the periodogram in Fig. 4.3 is raw and untreated, so the sampling distortions must be accounted for before final conclusions regarding the power spectrum of NGC 3783 are made. In the next section I will use PSRESP to fit a variety of models to NGC 3783.

4.1.4 Fitting models to the power spectrum of NGC 3783

Here I present the Monte Carlo analysis for several PSD model fits in an attempt to quantify the underlying model shape that best describes the PSD of NGC 3783, and associate an acceptance probability with each model. I initially test a simple unbroken power-law PSD model. I next fit a power-law with a single-bend in the PSD, and then a model incorporating a double-bend in the power-law. I also fit a single-bend power-law with a Lorentzian component to the observed periodogram of NGC 3783.

Unbroken power-law PSD model

The unbroken power-law PSD model is the simplest model I will use to fit the underlying power spectrum. Any flattening of the measured PSD at low frequencies would be instantly recognisable as a deviation away from a pure power-law slope. Before more complicated models are tested I must attempt to rule out the unbroken power-law model.

I fitted a simple power-law model to the data of the form




Figure 4.3: The top figure is the periodogram of NGC 3783, where the black, blue and green data represents the *RXTE* long time-scale observations and intensive observation, respectively. The red line is the periodogram of the two orbits of *XMM-Newton* data binned together. The bottom figure is the same periodogram plotted in $\nu \times P_\nu$ space i.e. power density. Note that the rise in power at the highest frequencies is due to the photon Poisson noise.

$$P(\nu) = A \left(\frac{\nu}{\nu_0} \right)^{-\alpha}, \quad (4.1)$$

where A is the normalisation at a frequency ν_0 and α is the power-law slope. The model was tested against the broadband periodogram of NGC 3783 by cycling through a range of possible power-law slopes, α , from 1.5 to 2.5 in increments of 0.05. Values of α were chosen to cover a range of possible power-law slopes that could be fitted to the data. Nine hundred simulated PSDs were used to determine the distorted model average PSD. Nine thousand randomly selected PSD segments were chosen to determine the χ^2 distribution of the realisations of the model. The best-fitting power-law slope, α , and corresponding probability is given in Table 4.2.




Figure 4.4: The top figure is the best-fitting unbroken power-law PSD of the combined *RXTE* and *XMM-Newton* NGC 3783 data plotted in RMS^2/Hz . The Poisson noise level at high frequencies is clearly visible as a flattening to zero slope above $\sim 10^{-3}\text{Hz}$. The solid black, blue, green, red lines represent the observed PSD data as first exhibited in Fig. 4.3. The distorted model average PSD is described by the symboled points with associated error bars. The open squares (cyan) and open circles (magenta) denote the distorted model average for the *RXTE* long time-scale data. The open triangles (orange) and filled circles (purple) represent the *RXTE* medium and *XMM-Newton* short time-scale data respectively. The dashed line is the underlying unbroken power-law model described in the text.

The bottom figure is the best-fitting unfolded unbroken power-law PSD for NGC 3783 in $\nu \times P_\nu$ space. In an unfolded PSD, the simulated power spectrum data points are associated with the observed data points with the distortion effects removed.




Figure 4.5: The best fit single-bend power-law PSD of NGC 3783 plotted in $\nu \times P_\nu$ space. The observed PSD points are as described in Fig. 4.4.

The unbroken power-law fit to NGC 3783 is shown in Fig. 4.4. The unfolded $\nu \times P_\nu$ space plots are an insightful way to inspect the best-fitting model in terms of the relative positions between the model that has been distorted by sampling effects and the observed data. The errors in an unfolded PSD can appear to be large. This is to be expected since the Poisson noise level is removed and the unfolded PSD merely reflects the source signal strength, which can be weak compared to the power in the Poisson noise level at high frequencies. It is clear from Fig. 4.4 that an unbroken power-law is not a good description of the observed data, since the observed power spectrum deviates at $\sim 10^{-6}$ Hz to a more shallow slope.

Single-bend power-law PSD model

Here I fit a single-bend power-law to the data. The power spectrum of the BHXRB Cygnus X-1 in the high/soft spectral state is best described by a power-law model that bends at a characteristic frequency from $\alpha \sim 2$ to $\alpha \sim 1$. The single-bend power-law PSD model has previously provided a good fit to the PSDs of the AGN NGC 4051 and MCG-6-30-15 (McHardy *et al.*, 2004, 2005):

$$P(\nu) = \frac{A \nu^{-\alpha_L}}{1 + (\nu/\nu_B)^{\alpha_H - \alpha_L}} \quad (4.2)$$

where α_L and α_H are the power-law slopes below and above the bend-frequency, ν . The single-bend power-law PSD model could have no physical significance and merely appears to be a good empirical representation of the observed data. Three

Model	Normalisation (Hz ⁻¹)	α_H	α_I	α_L	ν_H (Hz)	ν_L (Hz)	Acceptance
Unbroken	3.7×10^1	2.1_{-*}^{+*}	NA	NA	NA	NA	< 0.01
Single-bend	1.5×10^{-4}	$2.6_{-0.6}^{+1.0}$	NA	$0.8_{-0.8}^{+0.5}$	$6.2_{-5.6}^{+40.6} \times 10^{-6}$	NA	0.44
Double-bend	1.0×10^2	$3.2_{-1.2}^{+*}$	1.3_{-*}^{+*}	0.0	$2.6_{-*}^{+*} \times 10^{-5}$	$1.7_{-*}^{+*} \times 10^{-7}$	0.63

Table 4.2: Best-fitting model parameters for the unbroken, single and double-bend power-law PSD model to the combined *RXTE* and *XMM-Newton* PSD of NGC 3783. The errors are calculated from the 90 per cent rejection probability levels. The bend-frequency for the single-bend model, ν_B , is denoted here as ν_H . An asterisk indicates that the limit is unconstrained.

hundred simulated PSDs were used to determine the distorted model average PSD. Three thousand randomly selected PSD segments were chosen to determine the χ^2 distribution of the realisations of the model. The best fit unfolded single-bend power-law PSD for NGC 3783 is presented in Fig. 4.5 and the best-fitting model parameters are reported in Table 4.2. The best fit bend-frequency is $\nu_B = 6.2_{-5.6}^{+40.6} \times 10^{-6}$ Hz, the high-frequency slope is $\alpha_H = 2.6_{-0.6}^{+*}$, and the low-frequency slope is $\alpha_L = 0.8_{-*}^{+0.5}$. The errors are 90 per cent rejection probability levels and represent the point at which the likelihood of acceptance drops to less than 10 per cent, an asterisk indicates that the limit is unconstrained. The rejection probability levels for the main interesting parameters are plotted in Fig. 4.6. A good likelihood of acceptance is obtained (P=44 per cent) for the single-bend power-law PSD model. The single-bend power model represents a significant improvement; thus, we rule out an unbroken power-law model. The best-fitting single-bend frequency obtained here is consistent with the value found by Markowitz *et al.* (2003). Note that the result of Markowitz *et al.* (2003) is based upon *RXTE* long-term data up to 11 April 2002 (52311.3 MJD).

Double-bend power-law PSD model

The single-bend power-law power spectrum model would allow best comparison to the high/soft state of Cygnus X-1; however, to provide the best comparison to the low/hard state of Cygnus X-1 we consider a more complex double-bend power-law PSD model to see if the goodness-of-fit improves significantly. Although it is customary to describe the low/hard state of BHXRBS using multiple-Lorentzian components (e.g. Pottschmidt *et al.* 2003), AGN data is sufficiently poor in quality that a double-bend model will suffice as an approximation to the low/hard state (Belloni and Hasinger, 1990). The double-bend power-law model is given by

$$P(\nu) = \frac{A \nu^{-\alpha_L}}{\left[1 + \left(\frac{\nu}{\nu_L}\right)^{\alpha_I - \alpha_L}\right] \left[1 + \left(\frac{\nu}{\nu_H}\right)^{\alpha_H - \alpha_I}\right]}, \quad (4.3)$$

where α_I is the intermediate-frequency slope and ν_L and ν_H are the low and high bend-frequencies respectively. I fixed the low-frequency slope to zero, to avoid making the simulations computationally prohibitive, and because a low-frequency slope of zero would allow the best qualitative comparison to the low state of Cygnus X-1 (Nowak *et al.*, 1999).

Testing a large multidimensional parameter space in PSRESP is an extremely in-

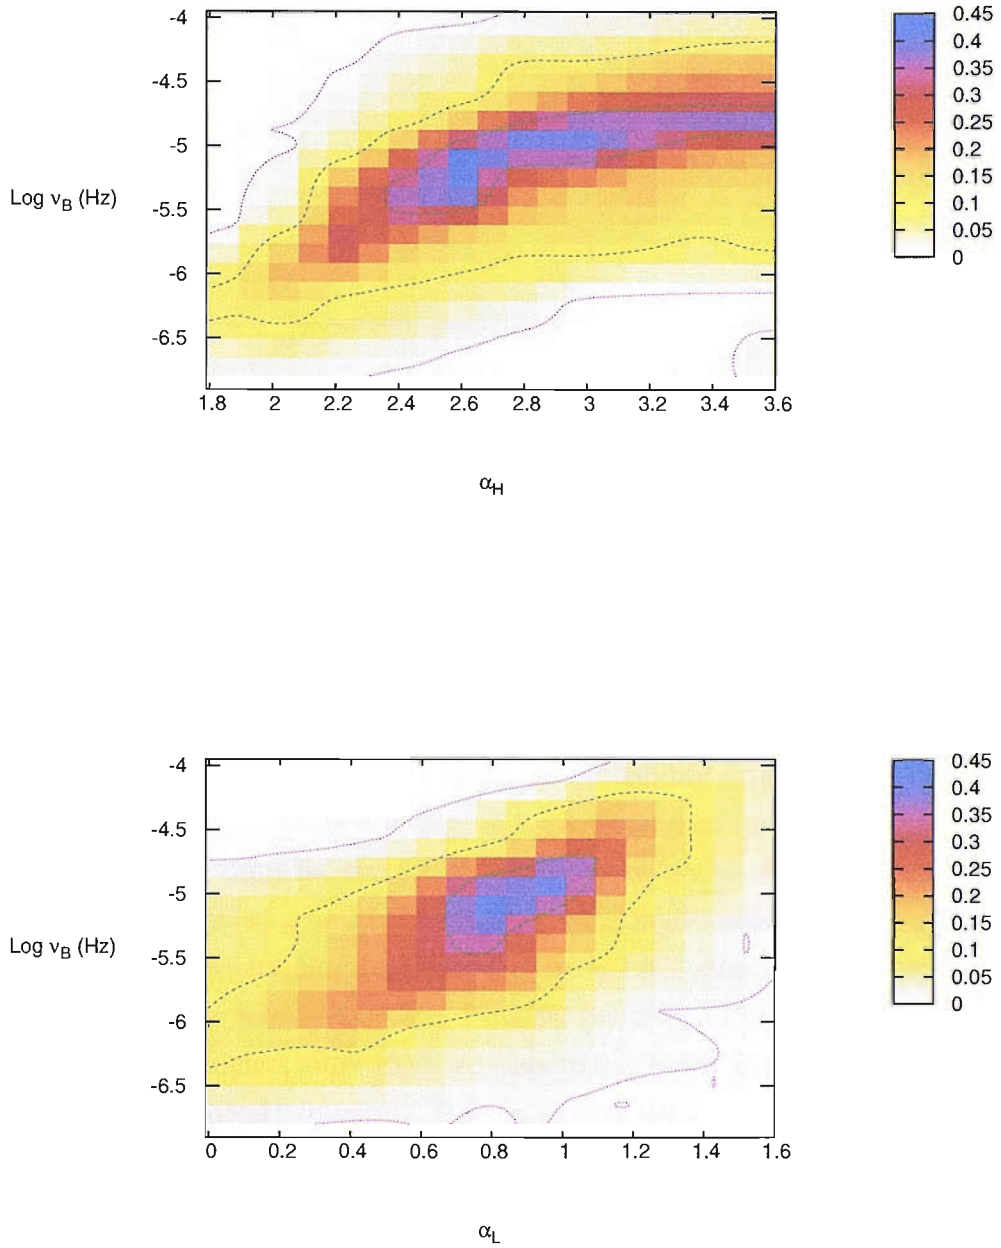


Figure 4.6: Rejection probability levels for NGC 3783 showing the errors on the best-fitting single-bend power-law model. The top contour represents the rejection probability levels for the high-frequency slope and bend-frequency taken from the projection of the three-dimensional probability space. The bottom plots represents a similar plot of low-frequency slope and bend-frequency. From the inside-out the coloured lines represent the 68, 90, and 99 per cent rejection probability levels. A color-bar represents the acceptance probability, with darker colours denoting a high acceptance probability and the best fit parameters.




Figure 4.7: The best fit double-bend power-law PSD of NGC 3783 plotted in $\nu \times P_\nu$ space. The observed PSD points are as described in Fig. 4.4.

tensive problem computationally (see Section 3.4). Three hundred simulated PSDs were used to determine the distorted model average PSD. Three thousand randomly selected PSD segments were chosen to determine the χ^2 distribution of the realisations of the model.

Fig. 4.7 presents the best-fitting double-bend power-law model, with the best-fitting parameters given in Table 4.2. The rejection probability levels for the main interesting parameters are plotted in Fig. 4.8. A good likelihood of acceptance is obtained (P=64 per cent) for the double-bend power-law model. The best-fitting high bend-frequency is $\nu_H = 2.6_{-*}^{+*} \times 10^{-5}$ Hz, the high-frequency slope is $\alpha_H = 3.2_{-1.2}^{+*}$, the intermediate-frequency slope is $\alpha_I = 1.3_{-*}^{+*}$, the low-frequency bend is $\nu_L = 1.7_{-*}^{+*} \times 10^{-7}$ Hz. As before, I use 90 per cent rejection probability levels. The added parameters allow extra freedom to find better fit probabilities for any given set of double-bend parameters. For this reason, the contour levels cover larger ranges in the parameter space and therefore, most of the 90 per cent contours in our double-bend fit remain unbounded over the fitted parameter space. The best fit low-frequency bend obtained here is consistent with the value found by Markowitz *et al.* (2003).

The best-fitting low-frequency bend is found close to the lowest frequencies probed by the data and, as seen in Fig. 4.8, it is essentially unbounded down to the lowest measurable frequency at the 68 per cent rejection probability level i.e. there is a great deal of freedom associated with the double-bend power-law PSD model fit to NGC 3783. These facts suggest that the second, low-frequency, bend might not be




Figure 4.8: Rejection probability levels for the high- and low-frequency bend as determined by the double-bend power-law PSD model. From the inside-out the coloured lines represent the 68 and 90 per cent rejection probability levels. A color-bar represents the acceptance probability, with darker colours denoting a high acceptance probability.

required by the data and that the improvement in the fit might be only due to the increased complexity of the model fitted.

The likelihood of acceptance is better in the double-bend power-law model than in the single-bend power-law model, 64 versus 44 per cent respectively, but this is to be expected due to additional free parameters available to the fitting algorithm. Given that a better fit is forecast, the increase in acceptance probability must be tested to determine whether or not the improvement is significant. In order to determine the significance of this improvement I must establish how the acceptance probability is expected to improve when a more complicated double-bend power-law model is fitted to a single-bend power-law data. Using the best-fitting single-bend PSD parameters, I generated 300 realisations of the *RXTE* and *XMM-Newton* lightcurves. Each broadband realisation was then fitted with the best-fitting double-bend parameters, exactly as was done with the real data, and the distribution of their fit probabilities was constructed. The distribution of acceptance probabilities represents the improvement expected when fitting a realisation of the best fit single-bend power-law model parameters with the best-fitting double-bend power-law model parameters,

both sets of parameters can be found in Table 4.2. I found that 121 out of the 300 single-bend simulations have a higher fit probability than the real data, when fitted with the double-bend model. Therefore, I conclude that the improvement in fit probability is no more than may be expected from fitting a model which is more complicated than required by the data: the double-bend model does not represent a significant improvement.

Single-bend power-law PSD model with a Lorentzian component

I finally consider whether the observed PSD might be best-described by adding a Lorentzian component, such as are commonly used to describe broadband noise components in BHXRBs (e.g. Nowak, 2000), to the single-bend power-law. I am motivated to consider this possibility because the PSD of the intense-sampling *RXTE* light curve is not very well described by either the single or double-bend power law model. Visual inspection of this light curve, shown in Fig. 4.1, suggests that the variability is strongly concentrated on time-scales of around a day, or equivalently, frequencies around 10^{-5} Hz, which is confirmed by the peak seen in the corresponding section of the PSD, and the drop in the same PSD at lower frequencies ($\sim 10^{-6}$ Hz). The long-term monitoring PSDs, however, do not show a dip at 10^{-6} Hz, creating a large discrepancy in the PSD measurements at this frequency. A strongly peaked component in the underlying PSD, at $\sim 10^{-5}$ Hz, could produce the observed features. Such a component would appear as a peak in a PSD that covered frequencies above and below its peak-frequency, but would be insufficiently sampled by the long-term monitoring campaigns; thus, its power would be aliased into the highest frequencies of the longer time-scale data, making them rise above the underlying model level and causing the apparent disparity.

The Lorentzian profile is described by

$$P_{\text{Lor}}(\nu) = \frac{AQ\nu_c}{\nu_c^2 + 4Q^2(\nu_c - \nu)^2}.$$

Centroid frequency ν_c is related to the peak-frequency ν_p by $\nu_p = \nu_c \sqrt{1 + 1/4Q^2}$ and the quality factor Q is equal to ν_c divided by the full width at half maximum of the Lorentzian. The variable A parameterizes the relative contribution of the power-law and Lorentzian components to the total RMS.

The single-bend power-law with Lorentzian component model was fitted in a similar fashion to the double-bend power-law model due to computational constraints. Moreover, the additional parameters of the single-bend power-law with Lorentzian




Figure 4.9: The best-fitting single-bend power-law with a Lorentzian component. The fit was done using the entire data set but here I only show the Lorentzian region. The observed PSD points are as described in Fig. 4.4. The Lorentzian feature in the model can reproduce qualitatively the spurious power at the high-frequency end of the long-term monitoring data and the turn down effect observed in the intensive-sampling data.

component model provide a great deal of freedom within the measured parameter space. As a result the rejection probability levels and corresponding parameter error bounds are not very well constrained. Three hundred simulated PSDs were used to determine the distorted model average PSD. Three thousand randomly selected PSD segments were chosen to determine the χ^2 distribution of the realisations of the model.

Fitting a Lorentzian component in addition to a single-bend power-law provides a good fit ($P=52$ per cent). The best-fitting Lorentzian contributes 20 per cent of the variance in the frequency range probed and its best-fitting parameters are quoted in Table 4.3. Fig. 4.9 shows the observed PSD compared with the best-fitting single-bend power-law model plus a Lorentzian component. The Lorentzian feature in the model can reproduce qualitatively the spurious power at the high frequency end of the long-term monitoring data and the turn down effect observed in the intensive-sampling data.

The additional Lorentzian component could, physically, represent a quasi-periodic oscillation (QPO) in the observed PSD. As mentioned in Chapter 1 no robust AGN QPOs have been detected due to the quality of AGN monitoring data (Vaughan and Uttley, 2005); although, their presence might be expected given their prevalence in BHXRB data. To determine the significance of the Lorentzian component fit I repeated the procedure used in determining the significance of the double-bend model.

ν_p (Hz)	ν_B (Hz)	Q	A	α_L	α_H	Acceptance
$4.8_{-0.8}^{+*} \times 10^{-6}$	$1.1_{-0.4}^{+0.6} \times 10^{-5}$	$5.1_{-3.6}^{+*}$	$0.9_{-0.7}^{+*}$	1.0_{-*}^{+*}	2.6_{-*}^{+*}	0.52

Table 4.3: Best fit single-bend power-law with Lorentzian component model parameters to the combined *RXTE* and *XMM-Newton* PSD of NGC 3783, where ν_p is the Lorentzian peak frequency, Q is its quality factor, ν_B is the power-law bend frequency and α_L and α_H are the power-law slopes below and above the bend, respectively. The errors are calculated from the 68 per cent confidence intervals, and an asterisk indicates that the limit is unconstrained.

I found that 222 of the 300 single-bend simulated PSDs have a higher fit probability than the data, when fitted with the single-bend power-law plus Lorentzian model. This result indicates that the increase in fit probability could be due to the added complexity of the model, and that the improvement in the fit over a simple bending power-law is not significant.

4.1.5 Discussion and Conclusions

I have combined our own new *RXTE* monitoring data with archival *RXTE* and *XMM-Newton* observations to construct a high-quality PSD of NGC 3783 spanning five decades in frequency.

I find that a ‘soft’ state model, with a single bend at 6.2×10^{-6} Hz, similar to that found earlier by Markowitz *et al.* (2003), a power-law of slope approximately 0.8 extending over almost three decades in frequency below the bend, and slope above the bend of approximately 2.6 is a good fit to the data. I also find that a ‘hard’ state model, with a double bend, fits the data, as does a model with a single bend plus an additional Lorentzian component. However the improvement in fit is marginal and, given the additional free parameters, is not significant. Thus I conclude that a simple ‘soft’ state model provides the most likely explanation of the data.

Assuming a mass of $3 \times 10^7 M_{\odot}$ for NGC 3783 (Peterson *et al.*, 2004), and an accretion rate of 7 per cent of the Eddington limit (Uttley and McHardy 2005, based on Woo and Urry 2002), then NGC 3783 is still in good agreement with the scaling of PSD break timescale as $\sim M/\dot{m}_E$ between AGN and BHXRBs found by McHardy *et al.* (2006).

Our new fits, show that the PSD of NGC 3783 is perfectly consistent with a single-bend power-law with low-frequency slope of 1, in contrast with the earlier result of Markowitz *et al.* (2003), who found that a similar model was rejected tentatively at ~ 98 per cent confidence. The difference can be understood in terms of the improved long-term data. Our new *RXTE* monitoring observations occur every 2 days, compared to 4 days previously, thereby increasing the long-term *RXTE* data set by a factor 2.6 and, in particular, providing overlap at high frequencies with the *RXTE* intensive monitoring data. The drop in long-timescale variability power, evident in the older long-term monitoring data is not reproduced by the new long-term monitoring data, showing that this drop could be just a statistical fluctuation. In addition, the very high frequencies are better constrained by the 2 orbits of *XMM-Newton* data than by the earlier *Chandra* data used by Markowitz *et al.* (2003).

The classification of the PSD as being ‘soft’ state means that NGC 3783 is no longer

considered unusual amongst AGN. The fact that this AGN is radio-quiet strongly supports the analogy with BHXRBs in the soft state. Also the accretion rate of NGC 3783 ($\dot{m}_E = 0.07$) (Uttley and McHardy 2005, based on Woo and Urry 2002) is similar to that of other AGN with soft-state PSDs (e.g. NGC 3227 Uttley and McHardy 2005, NGC 4051 McHardy *et al.* 2004, MCG-6-30-15 McHardy *et al.* 2005). This accretion rate is above the rate at which the persistent BHXRB Cyg X-1 transits between hard and soft states in either direction and at which other BHXRBs transit from the soft to hard state ($\dot{m}_E = 0.02$) (Maccarone *et al.*, 2003; Maccarone, 2003). I note that other transient BHXRBs in outburst, where the variable power-law emission in the soft state PSD is weak, can remain in the hard state to much higher accretion rates (~ 2 –50 per cent, Homan and Belloni 2005) but it is not clear whether we should expect similar PSD shapes to AGN for such outbursting sources. Thus NGC 3783 remains compatible with other moderately accreting AGN in being analogous to Cyg X-1 in the soft state. It is, of course, possible that the transition rate might not be independent of mass. Observations do not yet greatly constrain the transition rate as a function of mass but the absence of large deviations from the so-called ‘fundamental’ plane of radio luminosity, X-ray luminosity and black hole mass (Merloni *et al.*, 2003; Falcke *et al.*, 2004) argues against a large spread in the transition accretion-rates (e.g. see Körding *et al.* 2006a). In the case of Seyfert galaxy NGC 3227, the accretion rate is ~ 1 –2 per cent and a ‘soft’ state PSD is measured (Uttley and McHardy, 2005), which suggests that the transition accretion-rate in AGN should be at or below that value.

Our observations, which show that NGC 3783 does not have a highly unusual PSD, therefore confirm the growing similarities between AGN and Galactic black hole X-ray binary systems and leave Arakelian 564, which is probably a very high state object, as the only AGN showing clear double breaks (or multiple Lorentzians) in its PSD (e.g. Arévalo *et al.* 2006; McHardy *et al.* 2007).

4.2 The Narrow Line Seyfert 1 Ark 564: power spectral analysis

So far, the majority of well measured AGN with a moderately high accretion rate, exhibit a power spectrum similar to the soft state (or radio-quiet state) PSD of Cyg X-1. *RXTE* biases against low-mass low-accretion rate AGN, since they are often X-ray faint in *RXTE* and as such the AGN equivalent of a BHXRB low/hard state is largely unknown. Until recently, NGC 3783 was suggested to be a low/hard

state AGN since it was thought to exhibit two breaks in its PSD (Markowitz *et al.*, 2003), although its accretion rate is similar to that of other well observed soft state AGN PSDs. However, recent work using improved data (Section 4.1) has shown NGC 3783 to be entirely consistent with a soft state AGN PSD. Given that almost all AGN, with measured PSDs, seem to conform to a soft state interpretation it is very important to identify cases of non-soft state AGN. Such a result would suggest that the relationship between AGN and BHXRBs extends beyond the soft state.

The only case where there exists evidence of a non-soft state power spectrum is Ark 564. From observations with *RXTE* Pounds *et al.* (2001) present evidence for a low-frequency break, while Papadakis *et al.* (2002) present evidence for a second higher frequency break in the PSD using *ASCA* data. Papadakis *et al.* (2002) interpreted the combined *ASCA* and *RXTE* PSD as a ‘hard’ state power spectrum, but did note that the high-frequency break does not scale linearly with mass to the hard state of Cyg X-1. A hard state would also be surprising given the super-Eddington accretion rate of Ark 564 (Romano *et al.*, 2004) and it is much more likely that Ark 564 is in a very high state.

4.2.1 Observations and data reduction

RXTE

Ark 564 was observed by *RXTE* from 1 January 1999 to 4 March 2001. The long time-scale light curve observations were carried out approximately every 4 days, and the medium time-scale light curves (31 May 2000 to 1 July 2000) were carried out 8 times per day. Pounds *et al.* (2001) describes the observations from 1 January 1999 to 19 September 2000 in detail. Data were selected according to the standard ‘goodtime’ criteria (e.g. see previous section or McHardy *et al.* 2004; Summons *et al.* 2007). *RXTE* light curves (see Fig. 4.10) were produced in the 2–10 keV energy range.

ASCA

Ark 564 was observed with *ASCA* between 1 June 2000 and 5 July 2000. This observation suffers regular interruptions due to Earth occultations on the orbital time-scale of 5632s. Binning the data in orbit-long bins produces an evenly sampled light curve of 551 points that spans $\sim 10^{-6}$ Hz to $\sim 10^{-4}$ Hz in the broadband power spectrum.




Figure 4.10: The *RXTE* long-term (top) and medium time-scale (bottom) light curve of Ark 564 in the 2–10 keV energy range.

The SIS data was downloaded from the TARTARUS database in 0.6–2 and 2–10 keV energy bands, with background subtracted count rates of 2.7 and 0.6 cts s^{-1} respectively. The average photon energies of the 0.6–2 and 2–10 keV energy range are 1.18 and 3.2 KeV respectively.

See Papadakis *et al.* (2002); Edelson *et al.* (2002) for a further analysis of the variability properties of these data.

XMM-Newton

Ark 564 was observed by *XMM-Newton* from 19:47 5 January 2005 to 23:16 6 January 2005. The European Photon Imaging Cameras (EPIC) PN, MOS1, and MOS2 cameras were all operated in small window mode using a medium filter (Papadakis *et al.*, 2007). For the PN camera, a source extraction region was defined ($\sim 2' \times 2'$) and the background was extracted from a source-free region of equal area. Single and double events, with quality=0 were retained for the PN. For MOS1 and MOS2 cameras, source photons were extracted from a circular region of $\sim 46''$ in radius, and single, double, triple and quadruple events were used. Both MOS cameras suffered significant pile-up; thus, the central $12''$ diameter of each MOS exposure were




Figure 4.11: The ASCA SIS light curves of Ark 564 in 2-10 (top) and 0.6-2 (bottom) keV energy range in 5632s bins.

discarded. The PN did not display any signs of pile-up. The background subtracted count rate in the 0.2-10 keV PN light curve was $\sim 28 \text{ cts s}^{-1}$. There were no signs of background flaring during the observation.

As the variability properties normally depend on the energy band, *XMM-Newton* light curves from the PN data alone in the 0.5-2 and 2-8.8 keV bands were constructed. The average photon energies in these bands match those of the *ASCA* data, so the resulting periodograms should be consistent in shape and normalisation. These light curves are used throughout this analysis of Ark 564.

4.2.2 The power spectrum of Ark 564 and model fitting

To determine the PSD of Ark 564 over the largest possible frequency range we combine the *RXTE*, *ASCA*, and *XMM-Newton* data. The *RXTE* data has already been analysed by Pounds *et al.* (2001) and Markowitz *et al.* (2003) who both fit a breaking power-law model to the PSD. Markowitz *et al.* (2003) derive the following model parameters: a low-frequency break $\nu_L = 1.59_{-0.95}^{+4.73} \times 10^{-6} \text{ Hz}$, a low-frequency slope $\alpha_L = 0.05_{-2.05}^{+0.55}$, and an intermediate-frequency slope $\alpha_I = 1.20_{-0.35}^{+0.25}$. The results of Markowitz *et al.* (2003) are consistent with the earlier results of




Figure 4.12: The *XMM-Newton* PN light curves of Ark 564 in 0.5-2 (top) and 2-8.75 (bottom) keV energy ranges.

Pounds *et al.* (2001).

I use the PSRESP to estimate the underlying PSD parameter in the presence of sampling biases. Here I present the results of several PSD model fits to the observed Ark 564 data.

Double-bend power-law PSD model

The double-bend power-law PSD model should model the high- and low-frequency breaks determined by Pounds *et al.* (2001) and Papadakis *et al.* (2002) respectively; however, work by McHardy *et al.* (2004) suggested that AGN break frequencies are better fitted by a gentle-bend rather than a sharp break (e.g. Markowitz *et al.* 2003). Thus, here the combined *RXTE*, *ASCA*, and *XMM-Newton* observations are fitted by a power-law PSD model that gently bends rather than breaks. The double-bend power-law PSD model is given in Eqn.4.3. In order to provide the best comparison to the low/hard or very-high state of Cyg X-1 I fix the low-slope to zero. I utilise the 2–10 keV *RXTE* long time-scale data, 2–10 keV *ASCA* SIS data in 5632s bins, and the 2–8.8 keV *XMM-Newton* PN data in 48s bins. The *XMM-Newton* energy range is such that average photon energy matches that of the *ASCA* data. The choice of




Figure 4.13: PSD of Ark 564 from *RXTE*, *ASCA* and *XMM-Newton* data in the 2–10 keV *RXTE* energy band. The dashed black line denotes the best-fitting underlying double-bend power-law model. Data are represented by the symbols with error bars, unfolded from the distortions of the observed sampling pattern. The *RXTE* long time-scale, *RXTE* medium time-scale, *ASCA* and *XMM-Newton* data are denoted by the black open circles, green open triangles, blue open squares and the filled red circles respectively.

sampling interval for the *XMM-Newton* data is so that the light curve contains 2^{11} data points i.e. so that I can use the fast Fourier transform rather than the significantly slower discrete Fourier transform² when calculating the periodogram.

The double-bend power-law model was fitted to the data of Ark 564 using Iridis2, the University of Southampton’s Beowulf cluster, due to the large amount of computation required by these more complicated models. A very good fit is obtained to the hard energy range data (acceptance probability $P=75$ per cent - see Table 4.4 and Fig. 4.13) with $\alpha_l = 1.2^{+0.2}_{-0.1}$, $\nu_L = 7.5^{+28.1}_{-5.5} \times 10^{-7}$ Hz and $\nu_H = 2.4^{+2.3}_{-0.9} \times 10^{-3}$ Hz. These break frequencies are in good agreement with the values obtained by Papadakis *et al.* (2002) from analysis of *ASCA* data.

To make a combined PSD fit of the soft X-ray data, I use *ASCA* SIS 0.6-2 keV and *XMM-Newton* 0.5-2 keV together, but there is no such ‘soft’ band for *RXTE*. However, *ASCA* 0.6-2 and 2-10 keV data exhibit identical power spectral amplitudes, thus I assume that the correspondence between the two energy bands extends to the lower frequencies sampled by *RXTE*. Therefore I fitted the *RXTE* 2-10 keV long-term data together with the *ASCA* SIS (0.6-2 keV) and *XMM-Newton* (0.5-2 keV) to cover the same frequency range as the hard band. PSRESP calculates

²The FFT merely offers a speed increase over the DFT without loss of accuracy




Figure 4.14: PSD of Ark 564 from *RXTE*, *ASCA* and *XMM-Newton* data in the 0.2-2 keV *ASCA* energy band. The dashed black line denotes the best-fitting underlying double-bend power-law model. Data are represented by the symbols with error bars, unfolded from the distortions of the observed sampling pattern. The *RXTE* long time-scale, *RXTE* medium time-scale, *ASCA* and *XMM-Newton* data are denoted by the black open circles, green open triangles, blue open squares and the filled red circles respectively.

the amount of aliasing by the shape of the model power spectrum, but since I am combining hard and soft data extra care must be taken to ensure PSRESP correctly estimates the amount of aliasing at low frequencies i.e. hard and soft PSD model shapes may differ at high frequencies. Thus, I bin-up the *RXTE* 2-10 keV to approximately 16 days to minimise the possibility of misestimating the amount of aliasing. Both the hard and soft band PSDs will span the same frequencies. The fit is also very good (P=85 per cent - Fig. 4.14 and Table 4.4) with parameters $\alpha_l = 1.5^{+0.1}_{-0.4}$, $\nu_L = 3.6^{+5.1}_{-2.9} \times 10^{-6}$ Hz and $\nu_H = 1.8^{+1.8}_{-1.4} \times 10^{-3}$ Hz. The steeper intermediate slope reflects the lower high-frequency power in the soft band compared to the hard band. The break frequencies in both bands are consistent within 68 per cent rejection probability levels plotted in Fig. 4.15. The rejection probability levels in Fig. 4.15 provide strong evidence of a double-bend model i.e. compare to Fig. 4.8 in Section 4.1.4 where the same model provided a fit to the combined NGC 3783 data that was not statistically favoured.

Double-Lorentzian PSD model

Motivated by the strong evidence for a band limited power spectrum, I have further developed the correspondence between BHXRBS and AGN by fitting a multiple-

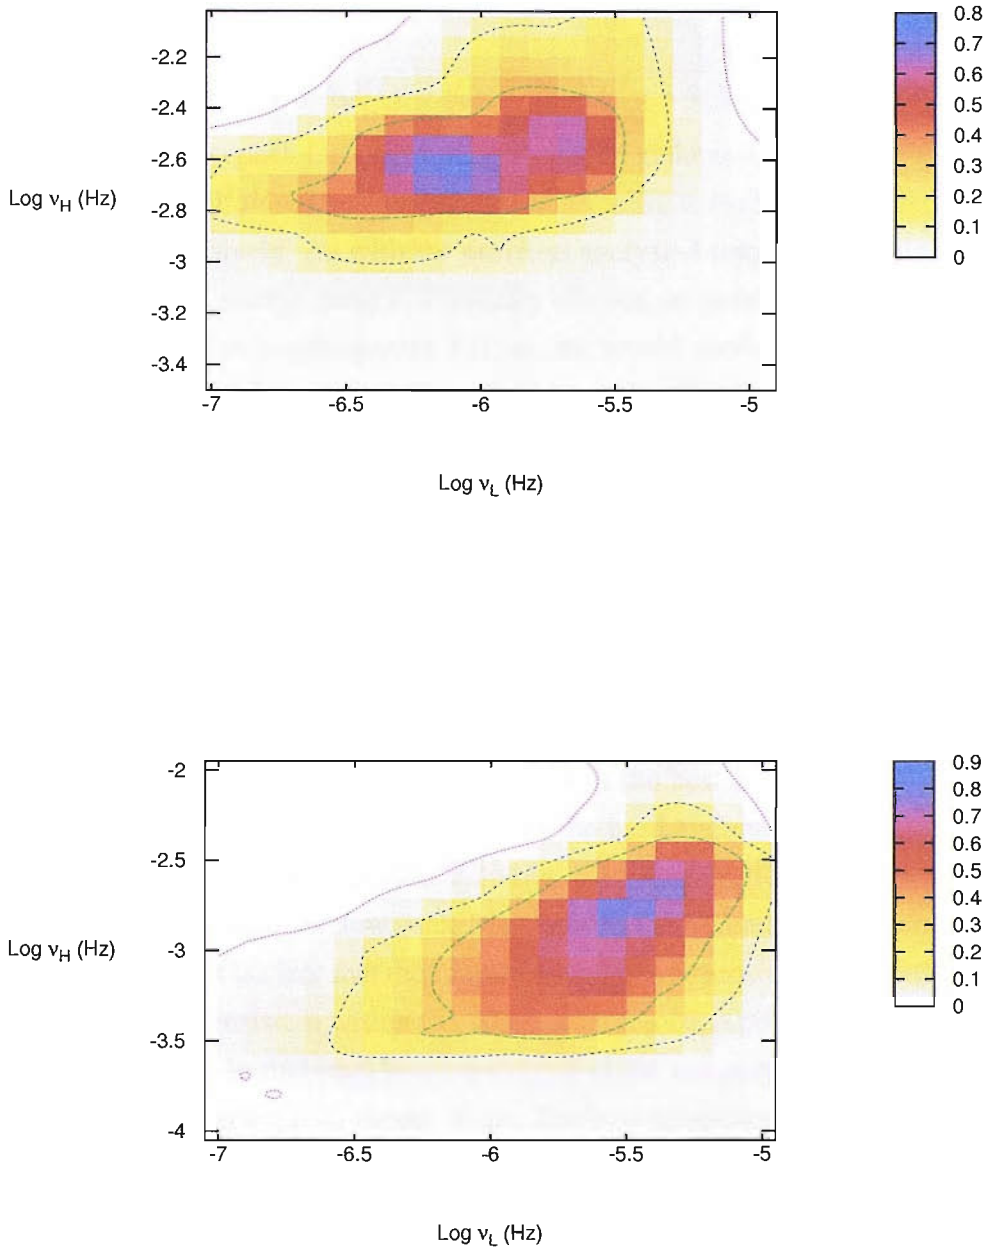


Figure 4.15: The 68, 90 and 99 per cent rejection probability levels for the high- and low-frequency bend of the double-bend power-law PSD model fit to the combined Ark 564 PSD. The top and bottom plots correspond to the hard and soft energy ranges, respectively. The rejection probability levels show conclusively that a double-bend power-law PSD model fit to Ark 564 is well constrained.

component Lorentzian PSD model, as is customary in BHXRBs analysis, to the power spectrum of Ark 564 (e.g. Pottschmidt *et al.* 2003). I fit the combined *RXTE*, *ASCA*, and *XMM-Newton* data with a model consisting of two Lorentzian components

$$P(\nu) = A_L \frac{Q_L \nu_{c,L}}{\nu_{c,L}^2 + 4Q_L^2 (\nu - \nu_{c,L})^2} + A_H \frac{Q_H \nu_{c,H}}{\nu_{c,H}^2 + 4Q_H^2 (\nu - \nu_{c,H})^2} \quad (4.4)$$

where the Lorentzian peak frequency ν_{peak} is related to the centroid frequency ν_c by $\nu_{\text{peak}} = \nu_c(1 + 1/4Q^2)^{1/2}$. The subscripts H and L refer to high- and low-frequency components, respectively. As with the previous analysis I fitted the combined PSD in the hard and soft energy ranges. I initially allowed all parameters to vary; however, I found that the low-frequency Lorentzian would converge to values of the quality factor ($Q = \nu_c/\Delta\nu_{\text{FWHM}}$) where the Lorentzian shape is made increasingly wide and is insensitive to the value of Q . A simulation with so many free parameters is computationally intensive, even on Iridis2, so in order to provide a more thorough search of the other parameters I therefore fixed the quality factor of the low-frequency Lorentzian at $Q_L = 0.01$ (with the 90 and 99 per cent rejection probability levels of $Q = 0.2$ and $Q = 0.5$, respectively) in both the hard and soft energy bands.

The best-fitting two-component Lorentzian PSD model for the hard and soft energy ranges are shown in Fig. 4.16. respectively, with the best fit parameters given in Table 4.4. The high-frequency Lorentzian in the hard energy range PSD exhibits increased variability power in 2-10 keV. In Fig. 4.17 the rejection probability levels are displayed for the high and low peak-frequency of the Lorentzian components for each energy range. It is clear that the Lorentzian peak frequencies from the different energy ranges are consistent within the 90 per cent rejection probability levels.

The PSDs presented in this section are unfolded of the sampling distortions, which depend on the underlying PSD model shape. The two-component Lorentzian model has a strong power component at 10^{-5} Hz which will cause significant aliasing in the long term *RXTE* monitoring data, whereas the double-bend power-law PSD model predicts less aliasing. The unfolded PSDs (presented) are unfolded of the sampling distortions (including aliasing) predicted by the underlying model, as a result the unfolded PSDs in Fig. 4.16 sit below the same part in the double-bend power-law model (Fig. 4.13 and Fig. 4.14), since the model did not contain so much power at high frequencies as the two-component Lorentzian model. If the two-component Lorentzian PSD model is correct, the double-bend power-law model would misestimate the amount of aliasing in the long time-scale *RXTE* data and

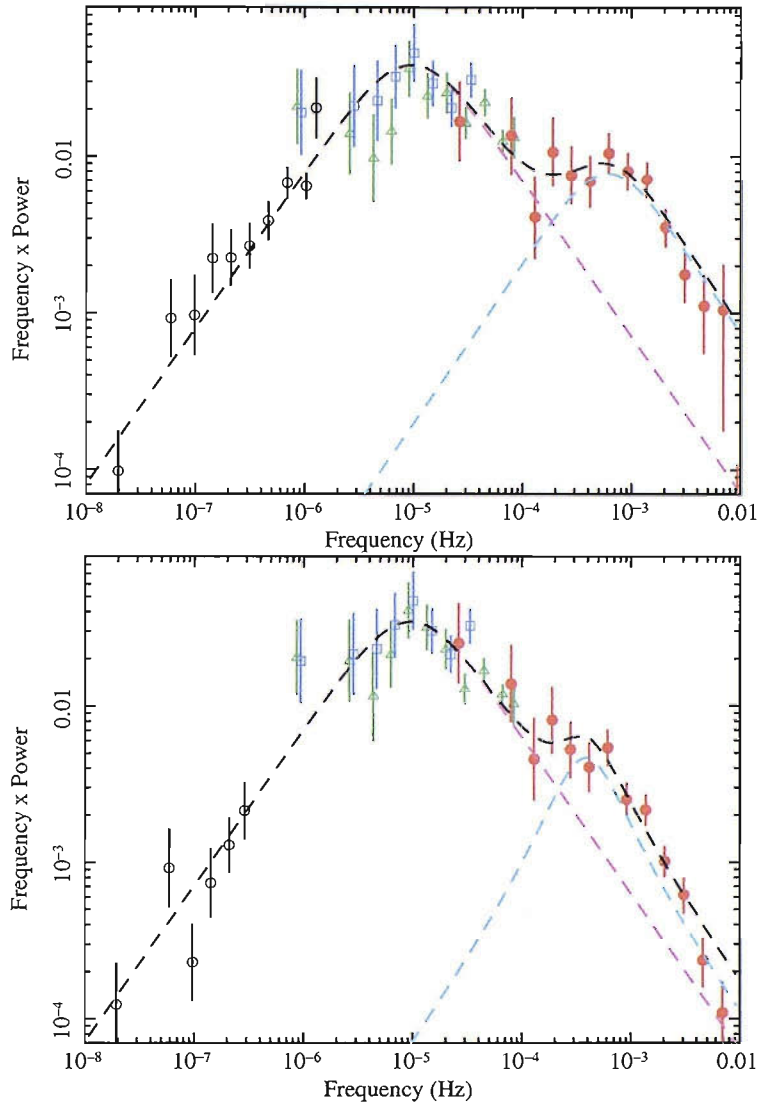


Figure 4.16: The top plot is the combined PSD of Ark 564 from *RXTE*, *ASCA*, and *XMM-Newton* in the hard energy range (2-10 keV *RXTE*). The bottom plot is the combined PSD of Ark 564 from *RXTE*, *ASCA*, and *XMM-Newton* in the soft energy range i.e. *RXTE* 2-10 keV, *ASCA* and *XMM-Newton* 0.6-2 keV. The symbols denote the four different data sets, with cyan and magenta dashed lines denoting the underlying Lorentzian components. The dashed black line denotes the sum of the Lorentzian components and represents the underlying model.

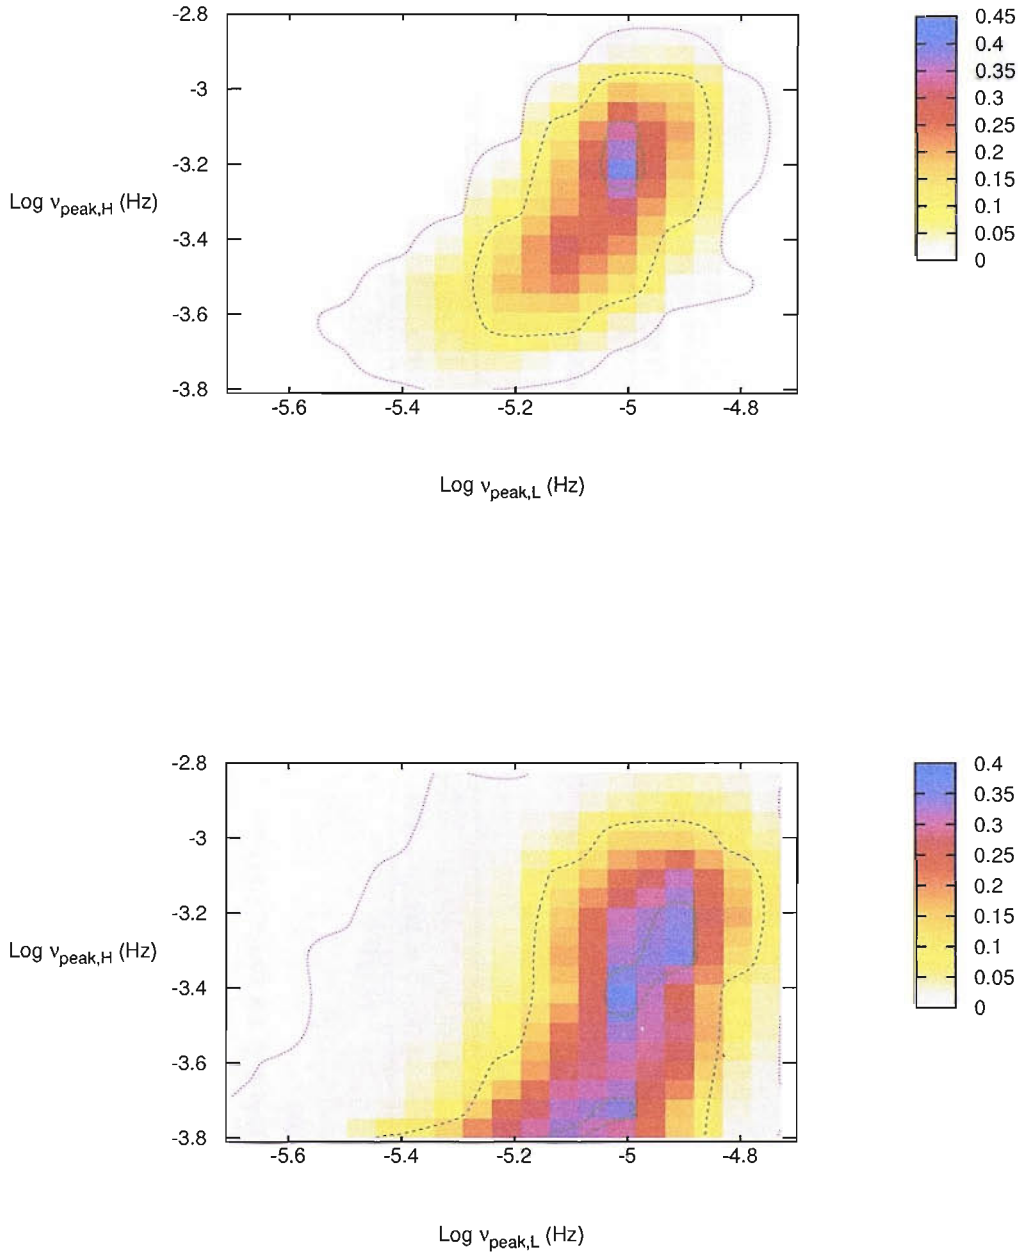


Figure 4.17: The 68, 90 and 99 per cent rejection probability levels for high and low Lorentzian peak frequencies of the two-component Lorentzian model fit (see Fig. 4.16) to the combined PSD of Ark 564. The top and bottom plots represents the hard and soft energy range rejection probability levels, respectively.

Table 4.4: Fit to the PSD using combined *XMM-Newton*, *ASCA* and *RXTE* data
Doubly-bending power-law

Energy Band	ν_L [Hz]	ν_H [Hz]	α_L	α_I	α_H	fit probability
0.6–2.0	$3.6^{+5.1}_{-2.9} \times 10^{-6}$	$1.8^{+1.8}_{-1.4} \times 10^{-3}$	0.0	$1.5^{+0.1}_{-0.4}$	3.4^{+*}_{-1}	0.85
2.0–8.8	$7.5^{+28.1}_{-5.5} \times 10^{-7}$	$2.4^{+2.3}_{-0.9} \times 10^{-3}$	0.0	$1.2^{+0.2}_{-0.1}$	$4.2^{+*}_{-1.8}$	0.75

Two-Lorentzian model

Energy Band	low- ν Lorentzian		high- ν Lorentzian		A_L/A_H	fit probability
	$\nu_{\text{peak,L}}$ [Hz]	Q_L	$\nu_{\text{peak,H}}$ [Hz]	Q_H		
0.6–2.0	$9.5^{+2.4}_{-1.9} \times 10^{-6}$	0.01	$3.9^{+2.2}_{-2.3} \times 10^{-4}$	$0.5^{+0.5}_{-0.38}$	$17.6^{+290.4}_{-11.9}$	0.37
2.0–8.8	$9.5^{+2.0}_{-1.5} \times 10^{-6}$	0.01	$6.1^{+1.5}_{-1.3} \times 10^{-4}$	$0.125^{+0.38}_{-*}$	$6.15^{+4.7}_{-1.0}$	0.40

* The errors are large and are not very well defined.

All errors are 68% confidence (1σ).

affect where the low-frequency bend is measured. Therefore the low-frequency bend in Fig. 4.13 and Fig. 4.14 is lower than in Fig. 4.16.

The high and low-frequency Lorentzian normalisations are not fitted separately in the two-component Lorentzian PSD model. The high and low-frequency Lorentzian normalisations are highly correlated with each other and I therefore fit the ratio of the normalisations A_L/A_H as a free parameter. It is clear that the rejection probability levels of $\nu_{\text{peak,H}}$ versus A_L/A_H are not coincident at the 90 per cent limit, as shown in Fig. 4.18. Thus, the high-frequency Lorentzian in the hard energy range is always relatively stronger than in the soft energy range when compared at the same frequency. The high-frequency Lorentzian component is slightly different in the hard and soft energy ranges, this is consistent with the values of the fractional RMS of 15 ± 1 per cent and 10 ± 1 per cent, respectively. The low-frequency Lorentzian RMS is consistent with being the same in both energy bands, 35 ± 1 per cent and 33 ± 1 per cent for hard and soft respectively.

4.2.3 Two-component Lorentzian PSD model and lag spectra

The power spectral analysis of Ark 564 up to this point was directly undertaken by myself using PSRESP and contributed to McHardy *et al.* (2007). This subsection on lag spectra was led by other coauthors of McHardy *et al.* (2007), but it is summarised here for completeness.

A powerful timing analysis tool is the lag spectrum. If observed light curves are separated into different energy ranges (e.g. 0.2–2 keV and 2–10 keV) the variations in one energy band can often be seen to lag behind the other; for example, in BXRBS the higher energy variations are generally seen to be lagging the lower energy variations in what we refer to as a ‘positive’ lag. Moreover, the time lags between energy bands vary as a function of the Fourier frequency of the component, with longer time-scales being associated with longer lags (e.g. see Nowak *et al.* 1999).

Previous work by Arévalo *et al.* (2006) on Ark 564 show that the time-lag between the 0.5–2 and 2–8.8 keV *XMM-Newton* bands is approximately 800s at low frequencies and approximately 0 for high frequencies. Arévalo *et al.* (2006) also show that at 10^{-4} Hz there is a sharp reduction in positive time lag values, which parallels the lag spectrum observed in Cyg X-1 in the very-high spectral state. Nowak *et al.* (1999) and Nowak (2000) have shown that the transition from constant lag to smaller constant lag often occurs at the intersection of the Lorentzian components. The coincidence of steps in the lag spectrum and transition between Lorentzian

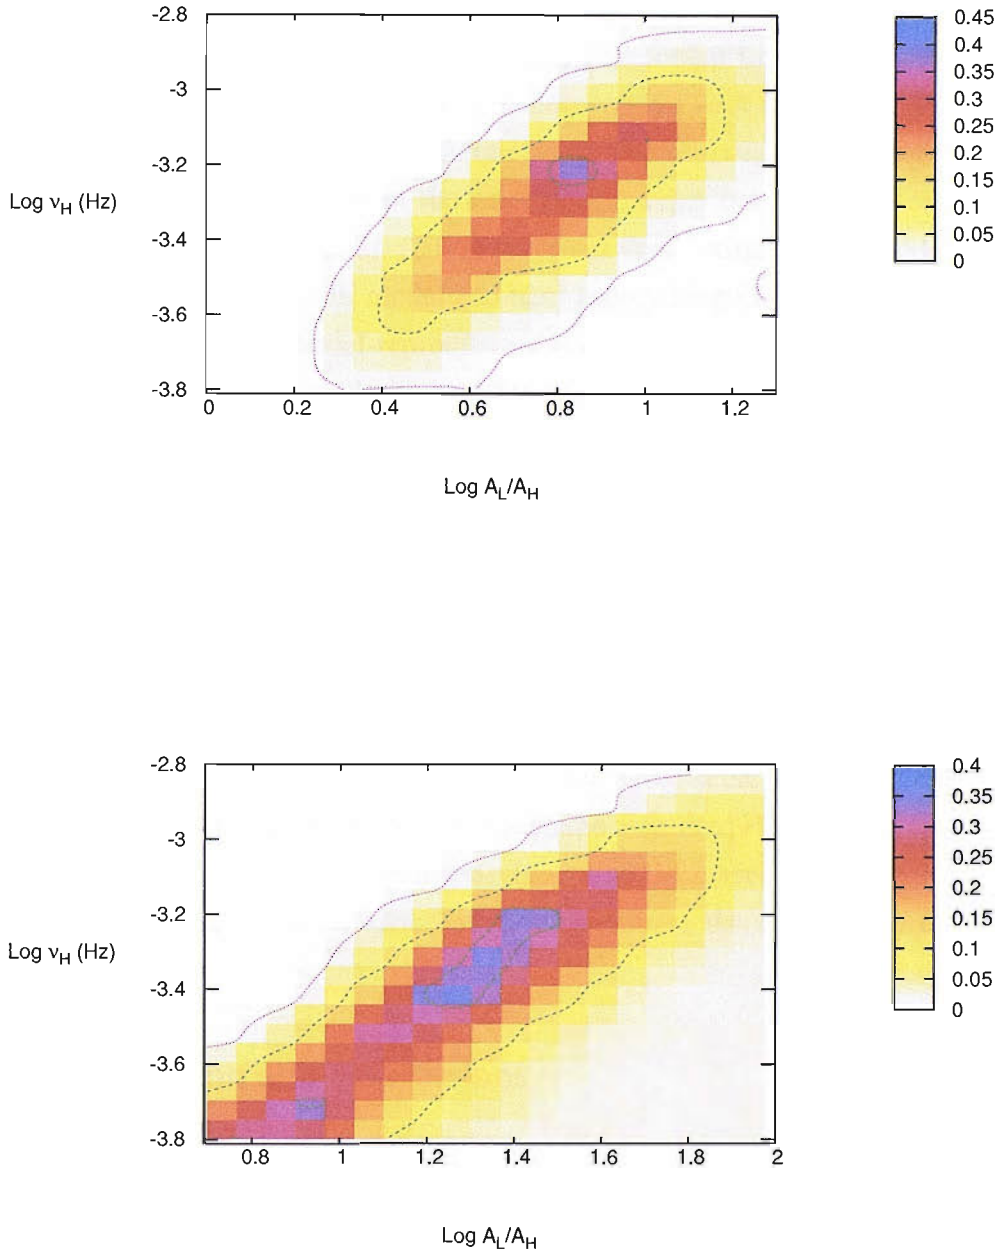


Figure 4.18: The 68, 90 and 99 per cent rejection probability levels for the high- and low-frequency Lorentzian normalisations, A_L/A_H , in the two-component Lorentzian model fit to Ark 564. The top and bottom plots represents the hard and soft energy range rejection probability levels, respectively.

components suggests that each constant lag is associated with a variability component. Therefore, the lag spectrum of Ark 564 should exhibit direct correspondence to the two-component Lorentzian model, i.e. we would expect the Lorentzian components to overlap exactly where the lag spectrum steps down.

This hypothesis was tested by directly fitting the observed lag spectrum by a model corresponding to the lags expected from a two-component Lorentzian model; the two Lorentzians exhibit different variability components with a single distinct time lag, which is constant for all Fourier frequencies. The observed time lag at a given Fourier frequency is a function of the amount of overlap between the Lorentzian components, see McHardy *et al.* (2007) for a full description of this model. The observed lag spectrum, best fit model and hard energy range PSD is shown in Fig. 4.19. The lag spectrum model reproduces the observed lag spectrum shape very well (reduced $\chi^2 = 1.35$), with the best-fitting values 638 ± 54 s (soft leads hard) and -11.0 ± 4.3 s (hard leads soft) for the low- and high-frequency lag, respectively. Recall from the best-fitting two-component Lorentzian model parameters that the low-frequency components are identical in both the hard and soft energy ranges, but the high-frequency components had a different quality factor and peak frequency. A better lag spectrum fit, with lags consistent to the previous fit, is obtained (reduced $\chi^2 = 1.13$) if the high-frequency components are forced to be the same in both energy ranges, with $Q_h = 0.5$ and $\nu_{peak,h} = 6.1 \times 10^{-4}$ Hz. The high-frequency components in both energy ranges were within and agree at the 68 per cent rejection probability level. The correspondence between the measured PSD and lag spectrum reinforces the two-component Lorentzian model as the correct description of the underlying power spectrum, which further substantiates the relationship between AGN and BHXRBs.

4.2.4 Discussion and Conclusions

Combining long time-scale archival *RXTE* monitoring data, a 35 day observation by *ASCA*, and a 100 ks observation by *XMM-Newton*, I have fitted the broadband power spectrum of Ark 564 spanning over five decades in frequency. I endorse the existence of two bend time-scales in the PSD previously reported by Pounds *et al.* (2001) and Papadakis *et al.* (2002). Given the results of the previous section on NGC 3783, Ark 564 is confirmed as the only observed AGN with a well measured PSD to exhibit band-limited variability.

I find that the observed PSD in the soft 0.5-2 keV and hard 2.0-8.8 keV energy range is well fit by a double-bend power-law PSD model, often representative of an




Figure 4.19: The top plot is the lag of soft (0.5-2 keV) to hard (2.0-8.8 keV) energy range as a function of Fourier frequency, evaluated using *ASCA* and *XMM-Newton* data. The solid line represents the model lag spectrum expected for a two-component Lorentzian PSD with model parameters determined by fitting the PSD, and forcing the Lorentzian components of the hard and soft energy ranges to be equal. Each Lorentzian is expected to have a constant lag, with the constant lag values given in the text. The bottom plot is the hard energy range PSD fitted by a two-component Lorentzian model, shown in Fig. 4.16. Notice how the lags are approximately constant in the domains where only one Lorentzian component dominates, but also how the lags change rapidly in the region where the Lorentzian overlap is maximal.

approximation to the low/hard or very-high state. The PSD is clearly band-limited with a drop in variability power at low and high frequencies. I find that a two-component Lorentzian PSD model is also a good fit to the soft and hard energy range PSD data, with consistent best-fit parameters between the bands. However, the peak frequency of the high-frequency Lorentzian component is measured to be a factor of four lower than the high-frequency bend measured with the double-bend power-law PSD model. This is common in fitting Lorentzian components, which inherently tend towards the centre of a band-limited PSD. The Lorentzian PSD model

fits exhibit increased variability power at high frequencies in the hard band, which is described by Pottschmidt *et al.* (2003) for BHXRBs as a strengthening of the high-frequency Lorentzian. The low-frequency Lorentzian component remains approximately constant between the energy ranges.

The observed lag spectrum exhibits a step-like structure when calculated between the aforementioned hard and soft energy ranges using *ASCA* and *XMM-Newton* data (Arévalo *et al.*, 2006). The time lags are approximately constant at low frequencies and reduce sharply at 10^{-4} Hz. The drop in the lag spectrum occurs exactly at the transition between the two Lorentzian components in the PSD, if each component has an associated constant time lag. This characteristic step-down feature and correspondence with observed PSD is often observed in BHXRBs (Pottschmidt *et al.*, 2003), I therefore consider the observed lag spectrum of Ark 564 as a strong indication that the two-component Lorentzian PSD model is correct.

The implication of a multiple Lorentzian component model PSD and lag spectrum is that Ark 564 should perhaps be thought of as a very-high soft or intermediate state, rather than a soft state. The low/hard state interpretation is ruled out due to the very high (super-Eddington) accretion state of Ark 564. BHXRBs in the very-high soft state (or intermediate state) are often very luminous radio sources (e.g. GRS1915+105 Miller-Jones *et al.* 2005), while BHXRBs in the soft state are very radio quiet and often undetected in this energy band. AGN with a similar mass to Ark 564 and a well-defined soft state interpretation of their PSD are radio-quiet (NGC 4051 and MCG-6-30-15), whereas Ark 564 is significantly more radio-loud. VLBI observations of Ark 564 (Lal *et al.*, 2004; Schmitt *et al.*, 2001) place its radio luminosity at 600 times that of NGC 4051 (Christopoulou *et al.*, 1997) for redshift-derived distances. The evidence supports a very-high soft or intermediate state interpretation of the observed Ark 564 data; however, the lack of a radio detection for soft state BHXRBs forbids a more complete comparison.

Mistakes are the portals of discovery.

JAMES JOYCE (1884 – 1941)

5

X-ray power spectra

RXTE has observed many AGN since its launch in 1995 and produced the best long time-scale monitoring X-ray light curves currently available. The aperiodic X-ray variability of AGN and Galactic black hole X-ray binaries (BHXRBs), as observed with *RXTE*, is best characterised as a red-noise process, with a power spectral density (PSD) of power-law form [i.e. $P(\nu)$, the power at frequency ν , $\propto \nu^{-\alpha}$ where $\alpha \sim 1$] with a break or bend (to $\alpha \sim 2$) at a characteristic frequency ν_B , or time-scale T_B . This characteristic time-scale scales approximately linearly with mass from AGN to BHXRBs (McHardy, 1988; Edelson and Nandra, 1999; Uttley *et al.*, 2002; Markowitz *et al.*, 2003; Uttley *et al.*, 2005) and inversely with accretion rate (McHardy *et al.*, 2004, 2005; Uttley and McHardy, 2005; McHardy *et al.*, 2006; K rding *et al.*, 2007).

BHXRBs are found in a number of different X-ray spectral states, which also have different PSD shapes (see Remillard and McClintock (2006) for an overview of BHXRBs states). In the low/hard state, where the energy spectrum is dominated by a highly variable power-law component, the PSD is well fitted by a multiple Lorentzian component model which, in the case of low quality AGN data, can be approximated by a double-bending power-law. In the high/soft state, where the energy spectrum is dominated by a constant thermal disk component, the PSD shape is distinguished from the low/hard state by having only one high-frequency break

in its power-law and a '1/f' PSD over many decades in frequency below the break, e.g. as seen in Cyg X-1 (Reig *et al.*, 2002). Although, the power-law component in the energy spectrum is producing the variability.

For almost all of the well observed, moderately accreting AGN, a single-bending PSD power-law model is sufficient to describe the observed X-ray variability, as we would expect if these AGN were the equivalent of Cyg X-1 'soft' states (McHardy *et al.*, 2004, 2005; Uttley and McHardy, 2005). The only exception found so far is Ark 564 whose PSD is well fitted either by double-bend power-law or by a multiple Lorentzian model, with latter being favoured because the same underlying model also explains the spectrum of the lags between the hard and soft energy bands (McHardy *et al.* 2007 or Chapter 4; see also Pounds *et al.* 2001; Papadakis *et al.* 2002; Arévalo *et al.* 2006). However there can be a problem in the comparison of published AGN PSDs in that different observers either use different analysis techniques, or different PSD models (e.g. the sharp single-break model (Markowitz *et al.*, 2003) versus the single-bending power-law model of McHardy *et al.* 2004). In this work, I present the broadband PSDs of 32 AGN, all analysed in the same way. I use public *RXTE*, *XMM-Newton*, *EXOSAT*, and *ASCA* data along with our own proprietary *RXTE* data to produce broadband PSDs which I fit with an improved Monte Carlo method (PSRESP) to determine the best fit PSD model parameters. I include improved PSDs for several previously published AGN, based upon new data, and include several previously unpublished AGN PSDs. I use the bend time-scales together with published measurements of black hole mass, bolometric luminosity, and the linewidth of the optical $H\beta$ emission line to reexamine the previously proposed correlations between these parameters (McHardy *et al.*, 2006).

5.1 Observations and data reduction

The *RXTE* archive contains a substantial amount of AGN data. AGN were only included if the observing campaign was sufficiently well sampled and the observations were of sufficient quality to produce a power spectrum. I have examined the archive and selected objects with approximately 100 or more observations. Filtering the *RXTE* archive based upon these selection criteria yield a target list of 31 objects, the majority of which are Seyfert 1 galaxies. I also select one object from the *XMM-Newton* archive (NGC 4395), since this object has a well measured power spectrum and should complement our survey of *RXTE* data (it was also included in McHardy *et al.* 2006). The final sample contains 2 QSO type objects, MR 2251-178

and PDS 456, and two low luminosity AGN (LLAGN), NGC 4395 and NGC 4258. Where periods of intensive monitoring occurred within the less frequent longer time-scale monitoring, I extracted the intensively sampled region as a separate light curve. The power spectra of the long and intensively sampled light curves then cover different time-scales, and thus make it possible to construct an overall power spectrum that spans a broad range of Fourier frequencies. The specific sampling patterns for each target are described in Table 5.1. The sampling pattern nomenclature of very-long (few years), long (months to years), intense or medium (days to months), and short (hours to days) correspond to the time-scale segments in the overall power spectrum.

Table 5.1: AGN source list

Target Name	Instrument	Time-scale	Sampling Interval	Date Range [MJD]	T_{obs}
(1)	(2)	(3)	(4)	(5)	(6)
PG 0052+251	<i>RXTE</i>	Long	4 days	53140.77 – 53432.75	291.98 days
	<i>XMM-Newton</i>	Short	300.0 seconds	53547.67 – 53547.91	0.24 days
Ton S 180	<i>ASCA SIS</i>	Intensive	~6 hours	51515.49 – 51527.72	12.23 days
	<i>XMM-Newton</i>	Short	110.0 seconds	51892.47 – 51892.83	0.36 days
Fairall 9	<i>RXTE</i>	Very Long	150 days	50390.61 – 52699.51	2308.90 days
	<i>RXTE</i>	Long	14 days	51180.56 – 52699.51	1518.95 days
	<i>RXTE</i>	Long	14 days	50390.61 – 50807.99	417.38 days
	<i>RXTE</i>	Intensive	~ 3.6 hours	52144.86 – 52178.99	34.13 days
	<i>XMM-Newton</i>	Short	435.0 seconds	51730.93 – 51731.31	0.38 days
3C 111	<i>RXTE</i>	Long	14 days	53065.41 – 53764.65	699.24days
	<i>XMM-Newton</i>	Short	230.0 seconds	51982.54 – 51983.06	0.52 days
3C 120	<i>RXTE</i>	Long	20 days	52334.92 – 53764.67	1429.75 days
	<i>RXTE</i>	Short	20 days	50812.60 – 51563.18	750.58 days
	<i>RXTE</i>	Intensive	~ 6.5 hours	52622.02 – 52630.86	65.84 days
	<i>XMM-Newton</i>	Short	490.0 seconds	52877.23 – 52878.72	1.49 days
Ark 120	<i>RXTE</i>	Very Long	35 days	50868.09 – 51662.07	793.98 days
	<i>RXTE</i>	Long	7 days	51026.21 – 51278.32	252.11 days

Table 5.1: AGN source list

Target Name	Instrument	Time-scale	Sampling Interval	Date Range [MJD]	T_{obs}
(1)	(2)	(3)	(4)	(5)	(6)
Mrk 79	<i>RXTE</i>	Long	7 days	51425.07 – 51662.07	237.00 days
	<i>XMM-Newton</i>	Short	870.0 seconds	52875.23 – 52876.53	1.3 days
	<i>RXTE</i>	Long	14 days	52720.39 – 54199.12	1478.73 days
	<i>RXTE</i>	Intensive	~ 6.5 hours	53691.42 – 53757.26	65.84 days
PG0804+761	<i>XMM-Newton</i>	Short	40.0 seconds	52025.75 – 52025.83	0.08 days
	<i>RXTE</i>	Long	4 days	51610.58 – 51971.79	361.21 days
	<i>RXTE</i>	Intensive	~ 11.9 hours	53300.26 – 53362.83	52.57 days
Mrk 110	<i>ASCA GIS</i>	Short	~0.27 hours	50756.52 – 50757.55	1.03 days
	<i>RXTE</i>	Long	14 days	53433.44 – 54199.39	765.95 days
	<i>RXTE</i>	Intensive	~ 6.8 hours	53695.48 – 53760.55	65.07 days
NGC 3227	<i>XMM-Newton</i>	Short	710.0 seconds	53324.21 – 53324.76	13.20 hours
	<i>RXTE</i>	Long	20 days	51180.45 – 53708.99	2528.54 days
	<i>RXTE</i>	Intensive	~ 7.1 hours	51636.56 – 51702.59	66.03 days
NGC 3516	<i>RXTE</i>	Long-look	1500 seconds	50405.68 – 50409.56	3.88 days
	<i>RXTE</i>	Very-long	70 days	50523.00 – 52331.58	1808.58 days
	<i>RXTE</i>	Long	14 days	50523.00 – 51593.38	1070.38 days
	<i>RXTE</i>	Intensive	4 days	50523.01 – 50659.12	136.11 days

Table 5.1: AGN source list

Target Name	Instrument	Time-scale	Sampling Interval	Date Range [MJD]	T_{obs}
(1)	(2)	(3)	(4)	(5)	(6)
NGC 3783	<i>RXTE</i>	Long-look	2000 seconds	50916.33 – 50919.68	3.35 days
	<i>RXTE</i>	Long-look	2000 seconds	50590.00 – 50594.23	4.23 days
	<i>RXTE</i>	Long	14 days	51180.53 – 52375.08	1194.54 days
	<i>RXTE</i>	Long	14 days	53063.38 – 54199.05	1135.67 days
	<i>RXTE</i>	Intensive	~ 3.5 hours	51960.14 – 51980.06	19.92 days
NGC 4051	<i>XMM-Newton</i>	Short	1000.0 seconds	52260.80 – 52264.39	3.59 days
	<i>RXTE</i>	Long	14 days	50196.49 – 54199.38	4002.89 days
	<i>RXTE</i>	Intensive	~ 6.8 hours	51665.36 – 51730.08	64.72 days
NGC 4151	<i>XMM-Newton</i>	Short	110.0 seconds	52045.44 – 52046.85	1.41 days
	<i>RXTE</i>	Long	14 days	51179.53 – 51964.62	785.09 days
	<i>RXTE</i>	Intensive	~ 3.6 hours	51870.61 – 51904.85	34.24 days
Mrk 766	<i>XMM-Newton</i>	Short	1000.0 seconds	51899.70 – 51900.08	0.39 days
	<i>RXTE</i>	Long	7 days	53065.44 – 53760.83	695.39 days
	<i>XMM-Newton</i>	Short	100.0 seconds	53513.80 – 53522.44	8.5 days
NGC 4258	<i>RXTE</i>	Long	14 days	50809.91 – 51604.58	794.67 days
	<i>RXTE</i>	Long	7 days	53436.58 – 53765.66	329.08 days
	<i>RXTE</i>	Intensive	~ 14.3 hours	53663.33 – 53765.66	102.33 days

Table 5.1: AGN source list

Target Name	Instrument	Time-scale	Sampling Interval	Date Range [MJD]	T_{obs}
(1)	(2)	(3)	(4)	(5)	(6)
	<i>XMM-Newton</i>	Short	500.0 seconds	51886.93 – 51887.20	0.27 days
NGC 4395	<i>XMM-Newton</i>	Short	21.0 seconds	52973.14 – 52974.14	1.0 days
3C 273	<i>RXTE</i>	Very Long	100 days	50122.61 – 53763.52	3639.91 days
	<i>RXTE</i>	Long	20 days	51970.87 – 53765.52	1794.65 days
	<i>RXTE</i>	Intensive	~ 11.3 hours	53358.96 – 53384.33	25.37 days
	<i>XMM-Newton</i>	Short	690.0 seconds	51710.54 – 51710.90	0.36 days
NGC 4593	<i>RXTE</i>	Long	20 days	53063.40 – 54199.45	1136.05 days
	<i>RXTE</i>	Intensive	~ 6.4 hours	53702.77 – 53766.70	63.93 days
	<i>XMM-Newton</i>	Short	550.0 seconds	52448.42 – 52449.43	1.01 days
3C 279	<i>RXTE</i>	Very Long	150 days	50104.64 – 53765.35	3660.71 days
	<i>RXTE</i>	Long	20 days	50104.64 – 51603.31	1498.67 days
	<i>RXTE</i>	Long	20 days	51970.54 – 53765.35	1794.81 days
	<i>RXTE</i>	Intensive	2 days	50427.85 – 50484.72	56.87 days
MCG 6-30-15	<i>RXTE</i>	Long	14 days	50159.77 – 54196.11	4036.34 days
	<i>RXTE</i>	Intensive	~ 6.5 hours	51622.68 – 51688.55	65.87 days
	<i>XMM-Newton</i>	Short	480.0 seconds	52123.18 – 52126.69	3.51 days
IC 4329A	<i>RXTE</i>	Long	20 days	52737.15 – 53765.09	1027.94 days

Table 5.1: AGN source list

Target Name	Instrument	Time-scale	Sampling Interval	Date Range [MJD]	T_{obs}
(1)	(2)	(3)	(4)	(5)	(6)
NGC 5506	<i>RXTE</i>	Intensive	~ 5.4 hours	52830.93 – 52864.85	33.92 days
	<i>XMM-Newton</i>	Short	330.0 seconds	52857.26 – 52858.83	1.57 days
	<i>RXTE</i>	Long	20 days	50159.82 – 52414.51	2254.69 days
	<i>RXTE</i>	Intensive	~ 6.8 hours	51620.88 – 51686.57	65.69 days
	<i>EXOSAT</i>	Intensive	2000 seconds	46454.60 – 46457.20	2.60 days
NGC 5548	<i>XMM-Newton</i>	Short	300 seconds	53200.93 – 53201.16	0.23 days
	<i>RXTE</i>	Very Long	150 days	50208.04 – 54199.19	3991.15 days
	<i>RXTE</i>	Long	14 days	50208.04 – 52749.68	2541.64 days
	<i>RXTE</i>	Long	14 days	53066.69 – 54199.19	1131.5 days
	<i>RXTE</i>	Intensive	~ 3.3 hours	52091.69 – 52125.42	33.73 days
PDS 456	<i>XMM-Newton</i>	Short	1270.0 seconds	52099.66 – 52100.77	1.11 days
	<i>RXTE</i>	Long	1 day	51963.20 – 51979.20	16.0 days
	<i>RXTE</i>	Long-look	10000 seconds	50879.71 – 50882.34	2.63 days
	<i>XMM-Newton</i>	Short	315.0 seconds	51966.41 – 51966.95	0.54 days
3C 390.3	<i>RXTE</i>	Long	4 days	51186.03 – 51964.00	777.97 days
	<i>XMM-Newton</i>	Short	400.0 seconds	53286.82 – 53287.64	0.82 days
Mrk 509	<i>RXTE</i>	Long	20 days	52726.35 – 53736.34	1009.99 days

Table 5.1: AGN source list

Target Name	Instrument	Time-scale	Sampling Interval	Date Range [MJD]	T_{obs}
(1)	(2)	(3)	(4)	(5)	(6)
	<i>XMM-Newton</i>	Short	330.0 seconds	52019.44 – 52019.96	0.52 days
NGC 7213	<i>RXTE</i>	Long	4 days	53797.13 – 54197.40	400.27 days
	<i>XMM-Newton</i>	Short	630.0 seconds	52057.93 – 52058.50	0.57 days
NGC 7314	<i>RXTE</i>	Very Long	20 days	51179.75 – 51747.86	568.11 days
	<i>RXTE</i>	Long	7 days	51257.26 – 51560.23	302.97 days
	<i>RXTE</i>	Intensive	~ 7.2 hours	51432.04 – 51438.61	6.57 days
	<i>XMM-Newton</i>	Short	330.0 seconds	52031.40 – 52031.92	0.52 days
Ark 564	<i>RXTE</i>	Long	7 days	51179.55 – 52702.85	1523.30 days
	<i>RXTE</i>	Intensive	~ 3.6 hours	51694.83 – 51726.47	31.64 days
	<i>XMM-Newton</i>	Short	48.0 seconds	53375.79 – 53376.97	1.18 days
MR 2251-178	<i>RXTE</i>	Long	7 days	53091.21 – 54196.03	1104.82 days
	<i>XMM-Newton</i>	Short	980.0 seconds	52412.88 – 52413.64	0.76 days
NGC 7469	<i>RXTE</i>	Long	20 days	52737.01 – 53773.57	1036.56 days
	<i>RXTE</i>	Intensive	~ 2.6 hours	50244.05 – 50276.00	31.95 days
	<i>XMM-Newton</i>	Short	300.0 seconds	53339.88 – 53342.97	3.09 days

5.1.1 *RXTE* Data Reduction

Typical *RXTE* AGN monitoring campaigns have consisted of many short, ~ 1 ks duration, observations with the proportional counter array (PCA, Zhang *et al.* 1993) with sampling frequencies varying from a few hours to a few days. I have analysed the archival PCA STANDARD-2 data and also our own proprietary data with FTOOLS v6.0.2 using standard extraction methods for all the AGN in Table 5.1. I use data from the top layers of all the available proportional counter units (PCUs); data were not collected from PCUs that suffered breakdown during on-source time. Moreover, data were not collected from PCUs 0 and 1 after 12 May 2000 and 25 December 2006, respectively, when the propane veto layer on the PCU was lost and hence the background in the detector increased substantially.

Data were selected according to the standard ‘goodtime’ criteria i.e. data were rejected if the target was within 10° of the Earth’s limb, within 30 minutes of the passage through the South Atlantic Anomaly (SAA), if the pointing offset was greater than 0.02° , or if the electron contamination was greater than 0.1. The PCA has no capability to take simultaneous background measurements during on-source time, so background measurements were estimated by using the L7 model for faint sources using PCABACKEST v3.0 to simulate the effect of a diffuse X-ray background, particle-induced background, SAA transitions and induced radioactivity. The response matrices for each PCA observation were calculated using PCARSP v10.1. The PCU gain settings have changed several times since launch, which would add additional variability power into the count rate light curves; thus the resultant count rates were rescaled to gain epoch 3 by calibrating with several Crab observations from the public archive. All *RXTE* count rates hereafter are normalised to 1 PCU.

RXTE light curves were generated over the 2 – 10 keV energy range for all AGN. The different time-scales (i.e. very-long, long, intensive and short) are treated as separate light curves in our fits.

Calibrating RXTE AGN data

RXTE has been in operation for over 10 years and in that time it has experienced various problems (see Section 1.4.1), which would affect the satellite’s operation. In particular the PCA has been through several gain epochs due to the breakdown of the component PCUs. Gain epochs and evolution of instrument response will contribute additional variance to measured count rates. Light curves calculated using flux

units should have the instrument response modelled out, if the response is well understood.

I use the entire archive of public *RXTE*-PCA observations of the Crab to produce a long time-scale light curve in count rate units. The Crab was chosen due to the abundance of data spanning the lifetime of *RXTE*; moreover, the Crab count rate spectrum are expected to be constant on long time-scales. It seems reasonable to assume that changes in the *RXTE* 2–10 keV Crab light curve are representative of changes in the PCA sensitivities. I undertook the following procedure to calibrate the AGN data found in this chapter.

- Produce Crab *RXTE* 2–10 keV PCA light curve in units of counts per second per PCU.
- Define gain epochs i.e. points in the light curve where count rate can change relatively dramatically. Gain epochs (documented changes in instrument response) are often associated with changes in count rate. See Fig. 5.1 for the raw count rate 2–10 keV Crab light curve; the changes in count rate due to instrument are clearly visible.
- Fit a power-law to each epoch in the Crab light curve. Eight power-laws are used to describe the observed Crab light curve.
- Each 2–10 keV AGN light curve was divided by the Crab light curve at gain epoch 3 using the aforesaid power laws to remove instrument response.

The gradual decline in count rate (see Fig. 5.1) is due to Xenon leaking into the propane layer and a broad hump in the flux light curve of the Crab produced from *RXTE* data is a mystery¹. Dead time is the time after an event that the detector system is unable to record subsequent events. In bright sources such as the Crab this is an important affect so I correct the Crab (and AGN) light curves for deadtime (Wei, 2006).

5.1.2 *XMM-Newton* Data Reduction

I used data from the European Photon Imaging Camera (EPIC) PN, and the two MOS cameras to constrain the high-frequency part of the overall power spectrum. The Science Analysis System (SAS v7) was used to reduce the raw PN and MOS data. For each AGN the single and double events were selected from a circular

¹Private communication with K. Jahoda.




Figure 5.1: *RXTE* 2–10 keV X-ray light curve of the Crab. The gain changes or large relative changes in the count rate are clearly visible. Count rates are normalised to PCU^{-1} .

region ($\sim 40''$) centred on the core of the source PSF, with quality flag=0 for PN data. The same procedure was used to extract source and background regions for the MOS camera and single, double, triple and quadruple events were selected. I also checked each target for pile-up using the *SAS* task *epatplot*, and for significant increases in the background due to increased levels of hard photons. As is customary in the cases where pile-up is not negligible, the analyses were repeated using same extraction region but removing the central $\sim 10''$ of the core of the PSF; moreover, I restrict our analysis to times when the hard bands are not affected by an increased background level. In the case of background flaring during the observation, the light curve is extracted only over those periods where flaring is not present. Data from the MOS cameras were also used if they were operated in imaging mode, but the PN dominates the resultant light curve since it has a better signal-to-noise ratio.

The *XMM-Newton* light curves for some of our sources contain a small number of minor gaps which, if left in as zero counts, introduce spurious power at high frequencies in the PSD. This effect is slight but can largely be corrected for by linearly interpolating the data from adjacent bins and adding a Gaussian scatter to simulate the counting noise. The final count rate light curves in the 0.3–10 keV and 4–10 keV energy bands were constructed from PN, and MOS data if it were available.

5.1.3 The Light Curves

The short gaps in *RXTE* light curves, as shown in Fig. C.1, are typical of many observing campaigns and are often due to orbital restrictions such as sun-angle constraints, or to a change in observing strategy. Where possible I have included all available *RXTE* data for each target. However, I have omitted data from some observing campaigns where the sampling is poor, or of limited duration, and does not help to constrain the broadband PSD.

It is not instructive to show all of the different light curves analysed in this work, but the important sampling parameters can be found in Table 5.1. I do, however, show these data in the appendices of this work.

Due to the uninterrupted nature of *XMM-Newton* light curves they are usually preferred for probing the highest PSD frequencies ($> 10^{-5}$ Hz) but occasionally, e.g. when the *XMM-Newton* observation is of limited duration, it is more useful to use *RXTE* long-look observations. These observations are broken up into many segments due to the low-Earth orbit of *RXTE* but, nonetheless, a few sources in this sample have good long-look data in comparison to relatively poor *XMM-Newton* data. NGC 3227 has a short duration *XMM-Newton* observation (~ 0.46 days), so I use the available *RXTE* long-look observation instead (~ 3.88 days). NGC 3516 has both *RXTE* long-look data and a good *XMM-Newton* observation. In our fits I use the *RXTE* long-look results due to their better quality but I do examine the *XMM-Newton* observation as a consistency check. NGC 5506 has *RXTE* long-look data, but it has a superior *EXOSAT* light curve obtained 24–27 January 1986 (McHardy and Czerny, 1987; McHardy, 1988). Assuming that AGN power spectra are stationary on time-scales of decades, the *EXOSAT* 1–9 keV band is comparable with the *RXTE* 2–10 keV band, so that the power spectrum measured from *EXOSAT* data should have a similar shape and normalisation to that of *RXTE* power spectra. For Ton S 180 and PG 0804+761 I also use archival *ASCA* GIS and SIS data, respectively, to constrain the higher frequencies. The GIS 1–10 keV light curve, and the SIS 0.5–10 keV light curve are comparable to *RXTE* 2–10 keV and *XMM-Newton* 0.3–10 keV, respectively.

5.2 Power spectral modelling

I use an improved Monte Carlo technique based upon the PSRESP method of Uttley *et al.* (2002) (see Chapter 3) to estimate the best-fitting underlying PSD model parameters in the presence of sampling distortions (see Chapter 2).

5.2.1 Combining X-ray data

To determine the observed broadband periodogram over the largest possible frequency range I combine the *RXTE* and *XMM-Newton* data, and occasionally data from other X-ray telescopes such as *ASCA* or *EXOSAT*. Power spectral slope and normalisation above the break frequency in BHXRBs are often energy dependent (e.g. Markowitz 2005). However, the slope of the PSD below the break appears to be independent of the chosen energy band (Cui *et al.*, 1997; Churazov *et al.*, 2001; Nowak *et al.*, 1999; Revnivtsev *et al.*, 2000; McHardy *et al.*, 2004). It is therefore necessary to ensure that energy ranges with similar median photon energies are used when combining data from different instruments. Throughout this work *RXTE* data are in the 2–10 keV energy band. It has previously been shown that the *RXTE* 2–10 keV is approximately the same as the *XMM-Newton* 4–10 keV, since these bands have the same median photon energy. However *XMM-Newton* has a softer response than *RXTE* and the count rate in the *XMM-Newton* 4–10 keV band is often very low so that only a marginal source variability power is detected above the Poisson noise level. To probe higher frequencies I can use 0.3–10 keV *XMM-Newton* data with significantly higher count rate, but I must rescale the resultant periodogram normalisation to match the 4–10 keV periodogram. For objects where 4–10 keV *XMM-Newton* data contained too little source power, I use the 0.3–10 keV data and determine the scaling correction by producing periodograms in both energy bands and fitting the same bending power-law model to the noise-subtracted data. On the assumption that the PSD shape below the high-frequency break is energy-independent, the combined *RXTE* 2–10 keV and *XMM-Newton* 0.3–10 keV periodogram will then have the shape of the 0.3–10 keV periodogram. The sources where the poor signal-to-noise in *XMM-Newton* 4–10 keV data adversely affect the broadband PSD are Ton S 180, 3C 120, Mrk 79, Mrk 110, NGC 4051, Mrk 766, NGC 4395, NGC 4593, and IC 4329A. In these cases I use the *XMM-Newton* 0.3–10 keV PSD.

5.2.2 Monte Carlo simulations: PSRESP

I use a Monte Carlo technique based upon the technique of Uttley *et al.* (2002) (PSRESP), to estimate the underlying PSD parameters in the presence of sampling distortions. I initially determine the observed periodogram, in parts, from the observed light curves, using the discrete Fourier transform (Deeming, 1975). The periodograms are binned in bins of width 1.5ν , where ν is the starting frequency

by taking the average of the logarithm of power (Papadakis and Lawrence, 1993a). I require a minimum of two periodogram estimates per bin. I then compare simultaneously the observed periodogram from each light curve with the model PSD determined from the average power of all the simulated periodograms. The simulated light curves have the same sampling pattern as the observed light curves to fold in the effects of sampling. I cycle the model parameters through a wide range of values to obtain the best fit model parameters for the input model. The reader is referred to Chapters 2 and 3 of this work for a full discussion of the method and sampling distortions.

5.3 Results

I present the results of several fits to a number of different PSD model shapes for each AGN in an attempt to quantify the underlying power spectral model which best describes the PSD, and associate an acceptance probability to each model. I fit all AGN with a simple unbroken power-law model and a single-bending power-law model with free low frequency slope, as well as free bend frequency and high frequency slope. Two objects (NGC 4051 and MCG-6-30-15) were also fitted with a double-bending power-law model. Nineteen objects were also fitted by a single-bending power-law model with a fixed low-frequency slope. I make no attempt to fit more exotic models that may be customary in BHXRB analysis, such as multiple Lorentzian models, since AGN data is generally not of a sufficiently high quality.

5.3.1 Unbroken power-law PSD model

To begin, I fitted a simple power-law model to the data of the form

$$P(\nu) = A \left(\frac{\nu}{\nu_0} \right)^{-\alpha},$$

where A is the normalisation at a frequency ν_0 , and α is the power-law slope. For all AGN, the model was tested by cycling through a range of ~ 20 parameter values of α in increments of 0.05. The range of values tested were chosen to produce a best-fitting value of α well bracketed by the parameter space. I made ~ 600 simulated periodograms and determined the acceptance probability from ~ 6000 combinations of simulated data.

The best-fitting models and simulated broadband PSD for all AGN are plotted in $\nu \times P_\nu$ space in Fig. 5.2. Table 5.2 summarises the best fit values of A and α for

the simple power-law along with the acceptance probability. The errors on α are taken at the 90 per cent rejection probability level. The acceptance probability is the chance that the observed data is a realisation of the underlying model.

The unfolded unbroken power-law fits to the AGN examined in this work (see Fig. 5.2) fit the observed data to varying degrees of success. The unfolded $\nu \times P_\nu$ space plots are an insightful way to inspect the the best-fitting model in terms of the relative positions between the model that has been distorted by sampling effects and the observed data. The unbroken power-law also provides an excellent way to investigate the undistorted observed PSD, without assuming a complicated model. Since these plots are unfolded of Poisson noise, many of the high-frequency PSDs appear to have somewhat large error bars (as we remove the Poisson noise component). I therefore binned the PSD at high-frequencies to clarify the high-frequency range if several points have particularly large unfolded error bars.

From our sample of 32 AGN, 14 targets (Fairall 9, 3C 111, 3C 120, Ark 120, PG 0804+761, Mrk 110, NGC 4258, NGC 4395, 3C 273, 3C 279, PDS 456, 3C 390.3, NGC 7213, and MR 2251-178) are consistent with a simple unbroken power-law model. An unbroken power-law for the remaining targets can be confidently rejected.

5.3.2 Single-bend power-law PSD model

Here I fit a single-bend power-law to the data. This model best describes the PSD of Cyg X-1 in the high/soft state, and has previously provided a good fit to the PSDs of the AGN NGC 4051, MCG-6-30-15 and NGC 3783 (McHardy *et al.*, 2004, 2005; Summons *et al.*, 2007). The single-bend power-law model is of the form

$$P(\nu) = \frac{A \nu^{-\alpha_L}}{1 + \left(\frac{\nu}{\nu_B}\right)^{\alpha_H - \alpha_L}}.$$

I tested a range of values for each free parameter: the high-frequency slope α_H , the low-frequency slope α_L , and the bend-frequency adjoining the slopes ν_B . Fig 5.3 displays the best-fitting unfolded single-bend power-law PSD model in $\nu \times P_\nu$ space. Fig. 5.4 shows the 2D projection of the 3D rejection probability levels for the main interesting parameters: high-frequency slope and bend-frequency. PSRESP demands significant computational resources; thus, in order to achieve the right balance of processing time required against result, I chose the parameter spaces so that, where possible, the rejection probability levels are largely constrained in ν_B . The high-frequency slope, α_H , is often adversely affected by Poisson noise, which

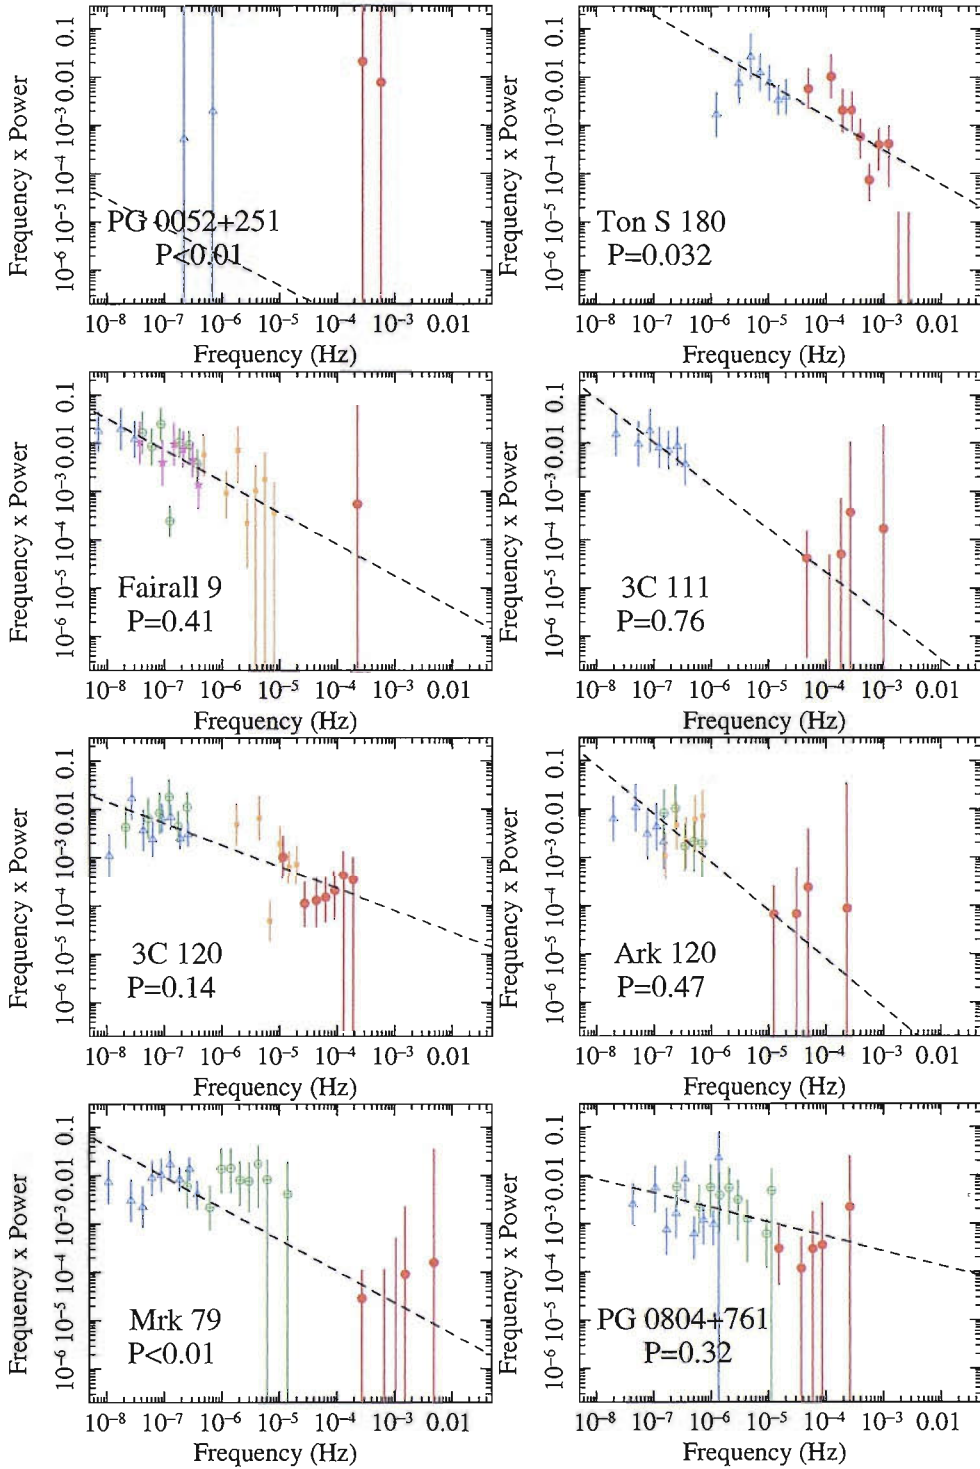


Figure 5.2: The best-fitting unfolded unbroken power-law PSD model for all targets in the selected sample, plotted in $\nu \times P_\nu$ space. The points with symbols represent the real data, unfolded from the distortions of the observed sampling pattern. A different symbol is associated with each component light curve input into the simulations. The underlying PSD model is seen as a dashed line. Noisy points at high frequencies have been binned for clarity. The acceptance probability for the unbroken power-law model is given for each fit.

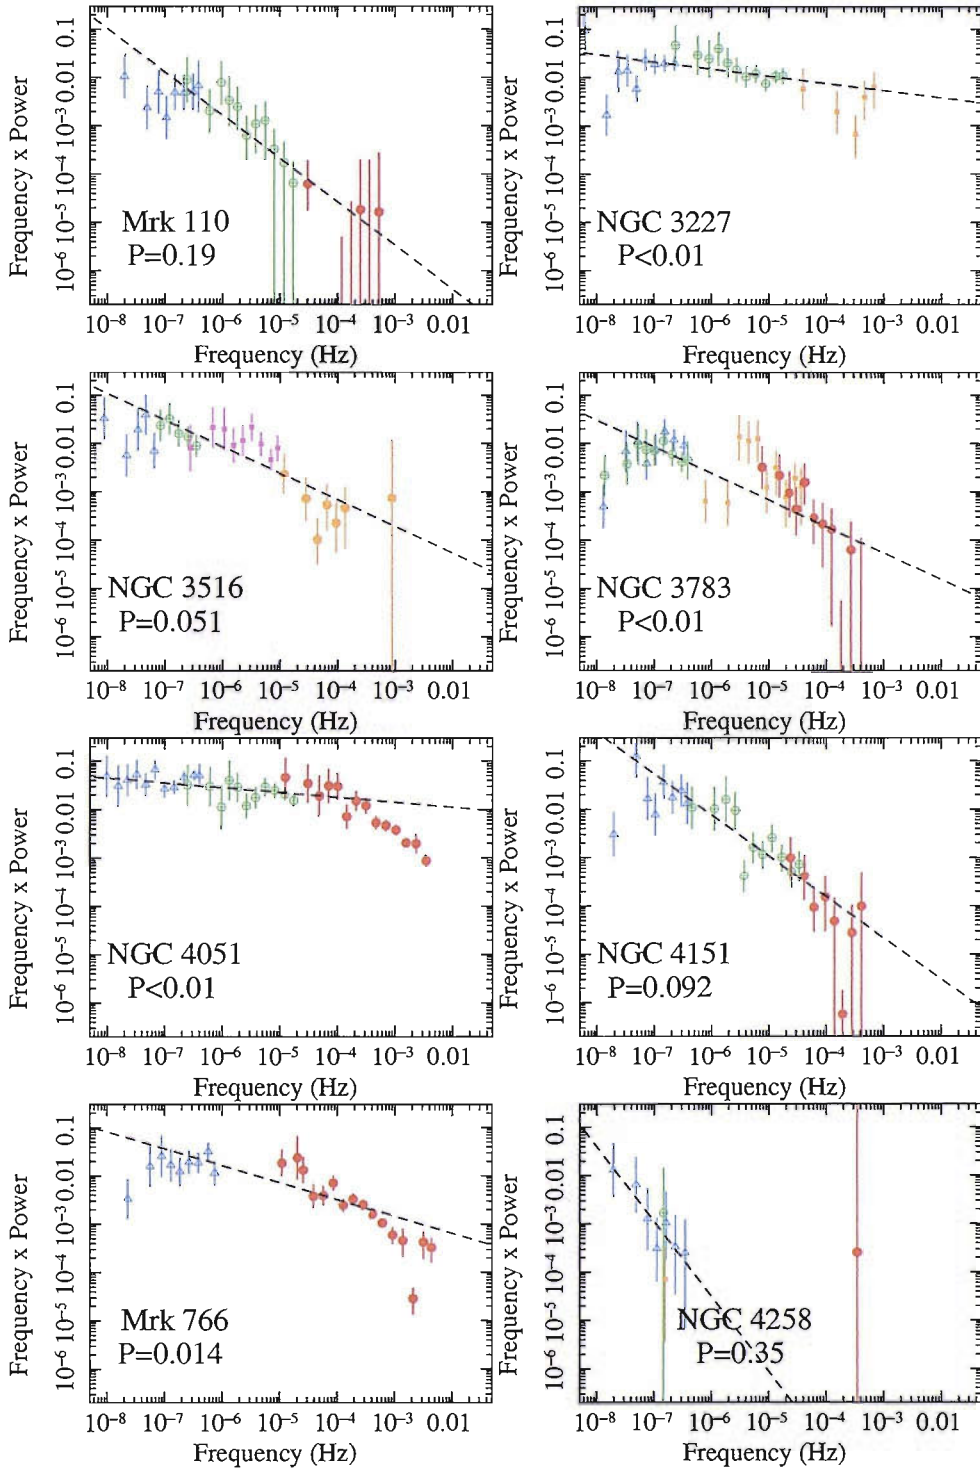


Figure 5.2: *continued.* The best-fitting unfolded unbroken power-law PSD model for all targets in the selected sample, plotted in $\nu \times P_\nu$ space. The points with symbols represent the real data, unfolded from the distortions of the observed sampling pattern. A different symbol is associated with each component light curve input into the simulations. The underlying PSD model is seen as a dashed line. Noisy points at high frequencies have been binned for clarity. The acceptance probability for the unbroken power-law model is given for each fit.

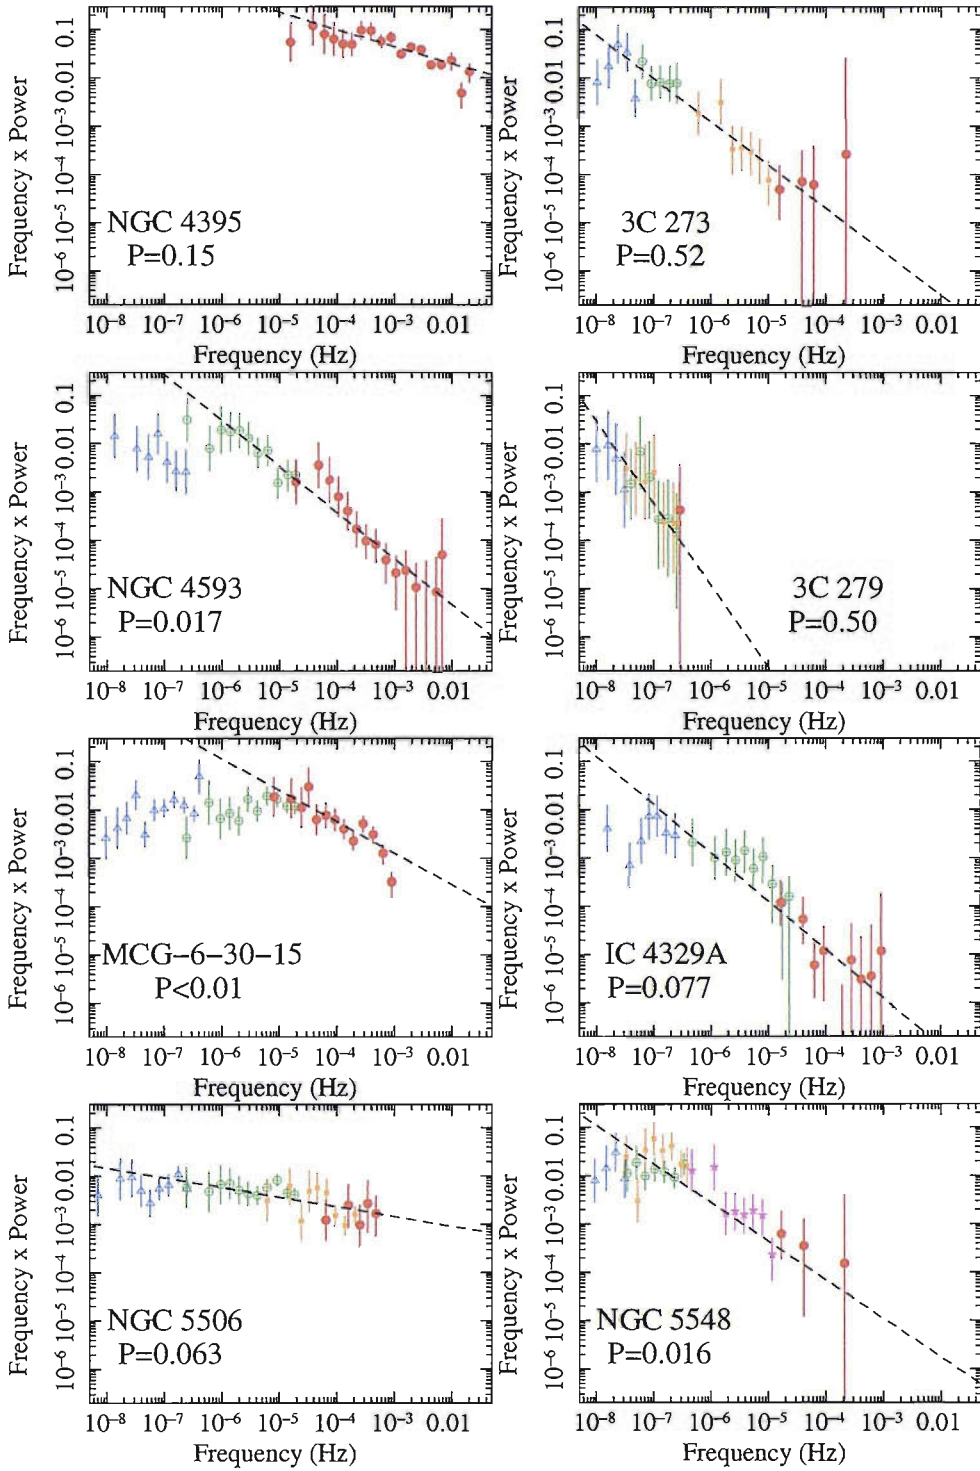


Figure 5.2: *continued.* The best-fitting unfolded unbroken power-law PSD model for all targets in the selected sample, plotted in $\nu \times P_\nu$ space. The points with symbols represent the real data, unfolded from the distortions of the observed sampling pattern. A different symbol is associated with each component light curve input into the simulations. The underlying PSD model is seen as a dashed line. Noisy points at high frequencies have been binned for clarity. The acceptance probability for the unbroken power-law model is given for each fit.

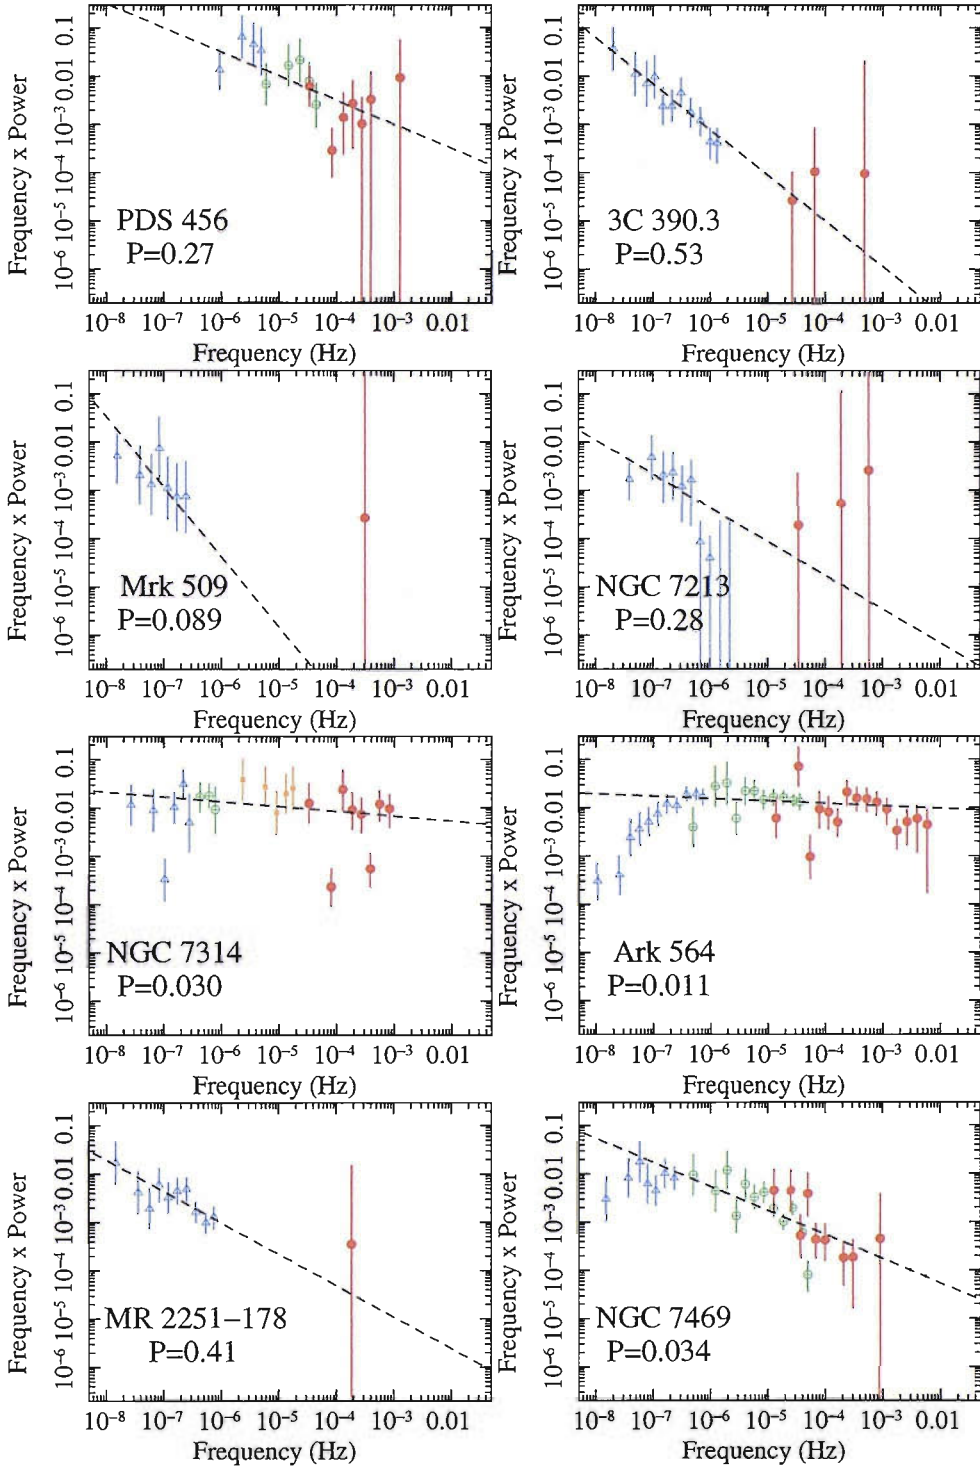


Figure 5.2: *continued.* The best-fitting unfolded unbroken power-law PSD model for all targets in the selected sample, plotted in $\nu \times P_\nu$ space. The points with symbols represent the real data, unfolded from the distortions of the observed sampling pattern. A different symbol is associated with each component light curve input into the simulations. The underlying PSD model is seen as a dashed line. Noisy points at high frequencies have been binned for clarity. The acceptance probability for the unbroken power-law model is given for each fit.

Target	Normalisation, A (Hz ⁻¹)	α	Acceptance
PG 0052+251	4.8×10^{-2}	1.60^{+*}_{+*}	< 0.01
Ton S 180	1.6×10^2	1.40^{+*}_{-*}	0.032
Fairall 9	3.6×10^1	1.65^{+*}_{-*}	0.41
3C 111	1.7×10^1	$1.90^{+*}_{-0.*}$	0.76
3C 120	6.4×10^1	1.45^{+*}_{-*}	0.14
Ark 120	7.9×10^0	$2.00^{+0.00}_{-1.00}$	0.47
Mrk 79	4.8×10^1	1.65^{+*}_{-*}	< 0.01
PG 0804+761	1.1×10^2	$1.30^{+0.3}_{-0.3}$	0.32
Mrk 110	2.1×10^1	$1.90^{+0.10}_{-0.25}$	0.19
NGC 3227	1.1×10^3	1.15^{+*}_{-*}	< 0.01
NGC 3516	2.4×10^2	1.55^{+*}_{-*}	0.051
NGC 3783	6.9×10^2	1.55^{+*}_{-*}	< 0.01
NGC 4051	2.2×10^3	1.10^{+*}_{-*}	< 0.01
NGC 4151	1.1×10^2	1.85^{+*}_{-*}	0.092
Mrk 766	5.8×10^2	1.30^{+*}_{-*}	0.014
NGC 4258	8.8×10^{-2}	2.55^{+*}_{-*}	0.35
NGC 4395	2.3×10^4	$1.35^{+0.05}_{-0.05}$	0.15
3C 273	1.6×10^1	1.90^{+*}_{-*}	0.52
NGC 4593	3.4×10^2	1.95^{+*}_{-*}	0.017
3C 279	2.4×10^{-2}	$2.7^{+0.80}_{-0.20}$	0.50
MCG 6-30-15	2.5×10^3	1.65^{+*}_{-*}	< 0.01
IC 4329A	1.3×10^1	2.00^{+*}_{-*}	0.077
NGC 5506	3.6×10^2	1.20^{+*}_{-*}	0.063
NGC 5548	4.3×10^1	1.80^{+*}_{-*}	0.016
PDS 456	1.0×10^3	$1.50^{+0.90}_{-*}$	0.27
3C 390.3	9.0×10^0	$1.95^{+0.55}_{-0.35}$	0.53
Mrk 509	1.5×10^{-1}	2.45^{+*}_{-*}	0.089
NGC 7213	8.8×10^0	1.70^{+*}_{-*}	0.28
NGC 7314	1.1×10^3	1.10^{+*}_{-*}	0.030
Ark 564	1.3×10^3	1.05^{+*}_{-*}	0.011
MR 2251-178	2.1×10^1	1.65^{+*}_{-*}	0.41
NGC 7469	1.7×10^2	1.30^{+*}_{-*}	0.034

Table 5.2: Summary of the best-fitting model parameters for the unbroken power-law model to the all targets. The errors on the power-law slope are calculated from the 90 per cent rejection probability levels. An asterisk indicates that an estimate of the limiting value was not produced.

contributes (along with aliasing and red-noise leak) to the problem of constraining rejection probability levels in this dimension. Table 5.3 summarises the best-fitting single-bend power-law model parameters. Many of the ‘noisy’ points that were binned for clarity in Fig. 5.2 are also binned in Fig. 5.3.

In six PSDs (PG 0052+251, NGC 3227, NGC 3783, NGC 4051, MCG-6-30-15, and Mrk 79) the unbroken power-law model fit is rejected with greater than 99 per cent confidence, while for a single-bend power-law PSD model they are rejected with a confidence much less than 90 per cent. In eleven other PSDs (Ton S 180, NGC 3516, NGC 4151, Mrk 766, NGC 4593, IC 4329A, NGC 5506, NGC 5548, Mrk 509, Ark 564, and NGC 7469) the unbroken power-law model fit is rejected with greater than 90 per cent confidence, while the single-bend power-law model is rejected at less than 90 per cent confidence. For the fourteen targets well fit by an unbroken power-law the suspected bend-frequency is close to the lowest measured frequency and is not well determined by the fitting.

NGC 7314 has a poor likelihood of acceptance for the unbroken power-law, but the single-bend power-law PSD model is almost as poorly fit. However, the broadband data is best fit by an unbroken power-law slope $\alpha = 1.1$. A slope this shallow is usually representative of a power-law slope *below* the characteristic bend frequency (McHardy *et al.*, 2004, 2005), it is therefore reasonable to assume that the best fit unbroken power-law slope is in fact a measurement of the low-frequency slope and that the bend is high in frequency. This conclusion is tentatively supported by a high ν_B in the single-bend fits of NGC 7314, but more data is required.

The PSD of NGC 4395 has been calculated from *XMM-Newton* data only and is well fit by an unbroken power-law and a single-bend power-law PSD model; however, as with NGC 7314 the best-fit unbroken power-law slope is fairly shallow and also implies a low-frequency slope measurement. The single-bend power-law PSD model fit to NGC 4593 provides a bend and slopes consistent with a very-high bend (highest in sample), and given the very low mass of NGC 4395 (Vaughan *et al.*, 2005; Filippenko and Ho, 2003), it can be tentatively assumed that the single-bend PSD model is preferred over the unbroken power-law PSD model. To rule out an unbroken power-law PSD model for NGC 4395, data is required below the bend to measure medium to long time-scales ($\sim 10^{-5}$ to 10^{-6} Hz). PG 0052+251 is poorly fit by the unbroken and single-bend power-law; however, PG 0052+251 represents the poorest quality data in this work.

The PSD of NGC 3516 in Fig. 5.3 was determined using *RXTE* long-look data; an independent power spectral fit using the available *XMM-Newton* data is consistent with the fit presented.

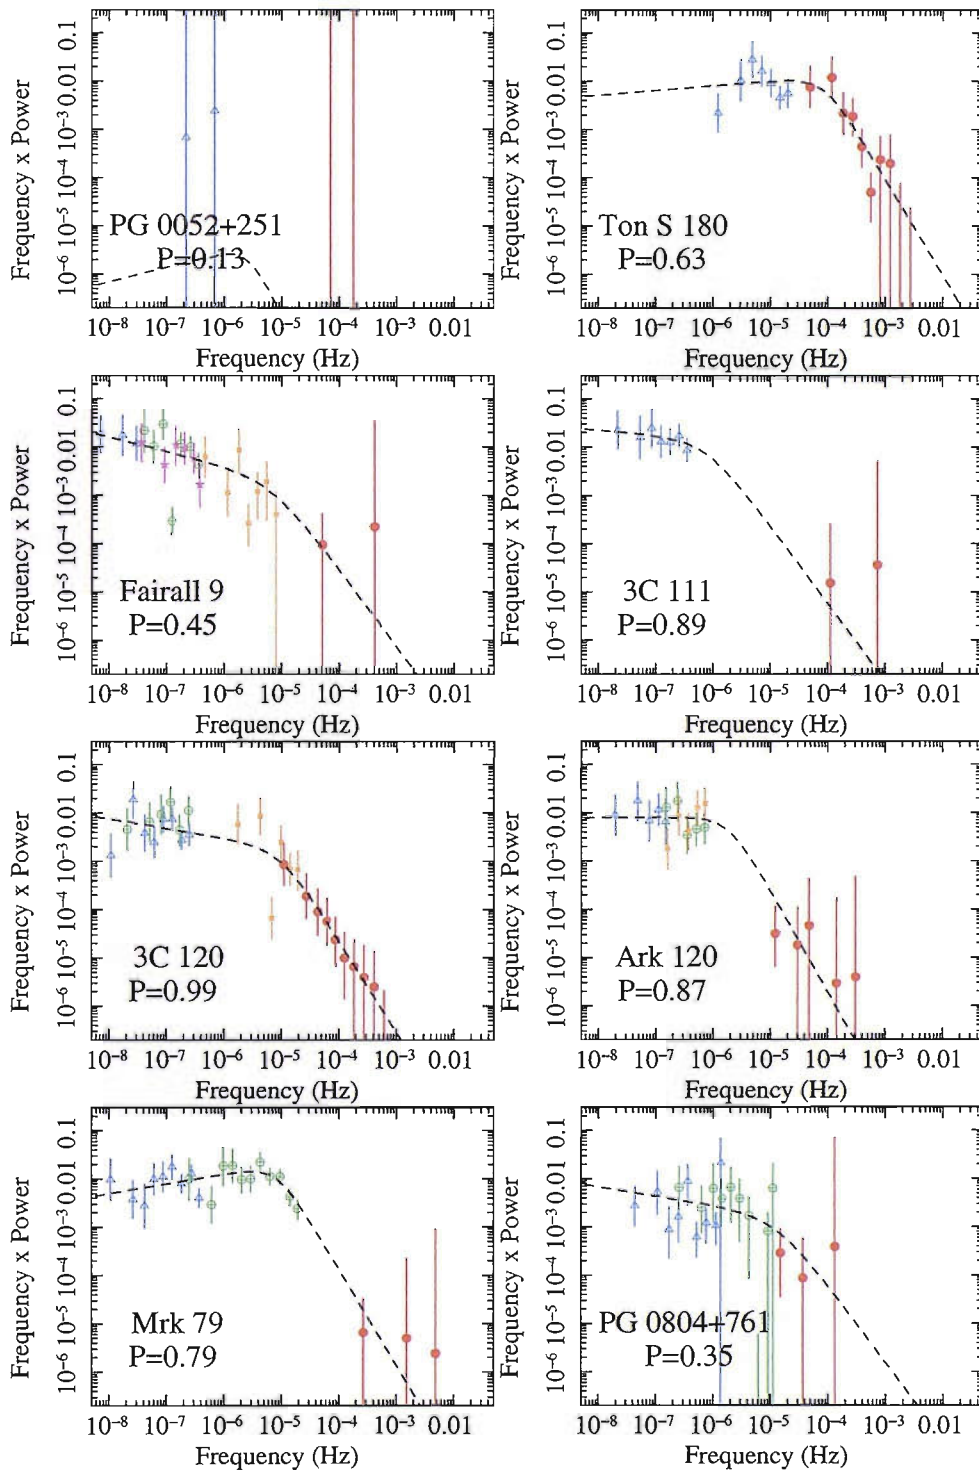


Figure 5.3: The best-fitting single-bend PSD for all targets in the selected sample, plotted in $\nu \times P_\nu$ space. See Fig.5.2 for a description of the points.

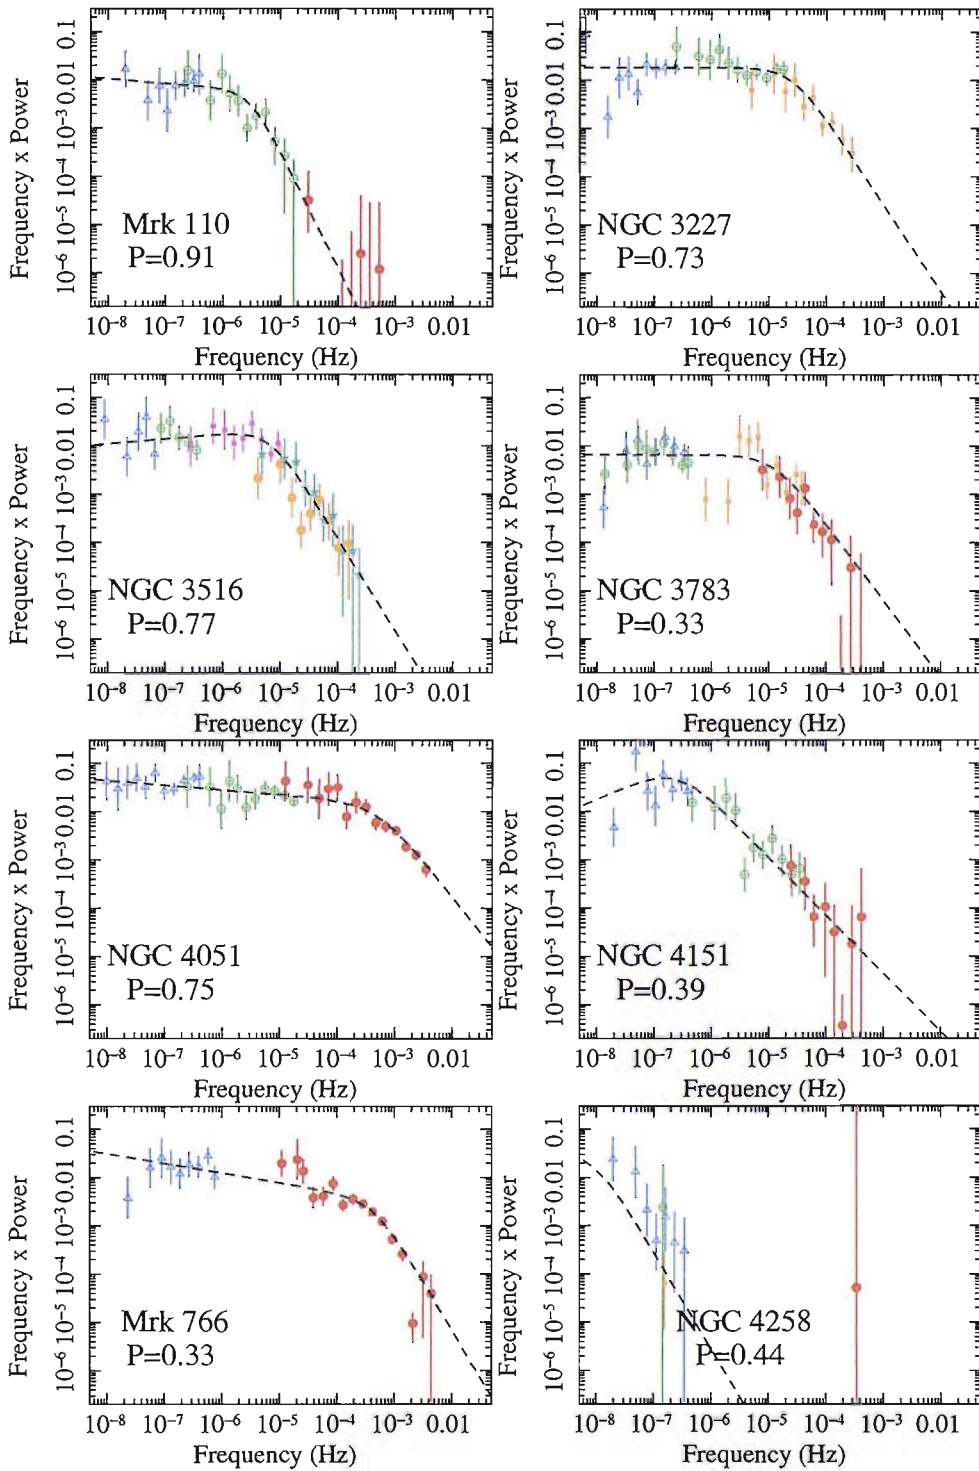


Figure 5.3: *continued.* The best-fitting single-bend PSD for all targets in the selected sample, plotted in $\nu \times P_\nu$ space. See Fig.5.2 for a description of the points.

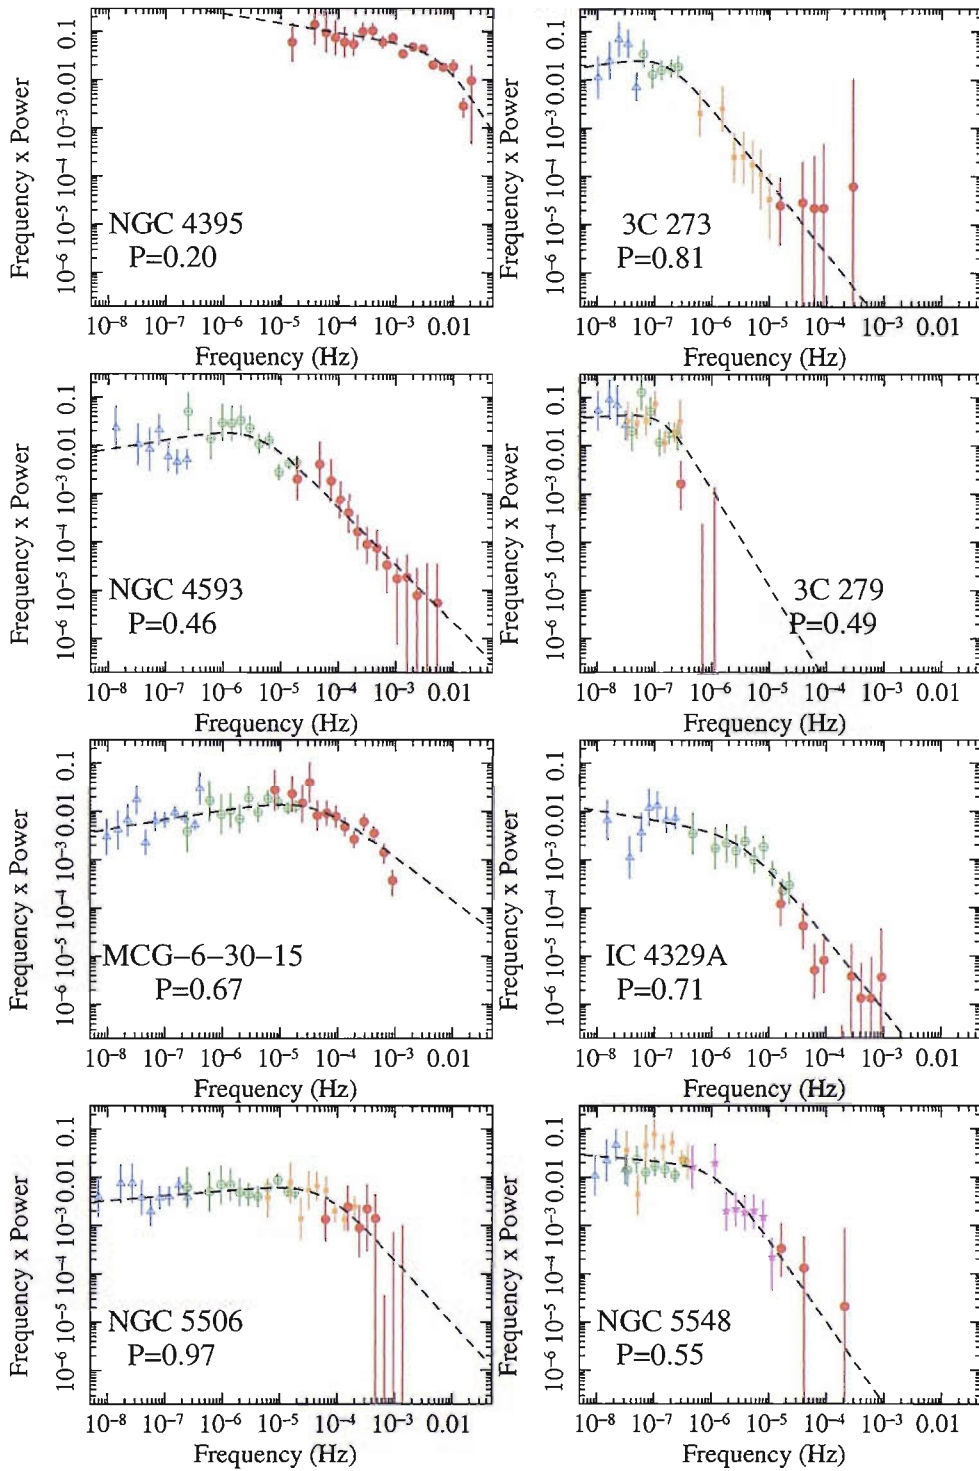


Figure 5.3: *continued.* The best-fitting single-bend PSD for all targets in the selected sample, plotted in $\nu \times P_\nu$ space. See Fig.5.2 for a description of the points.

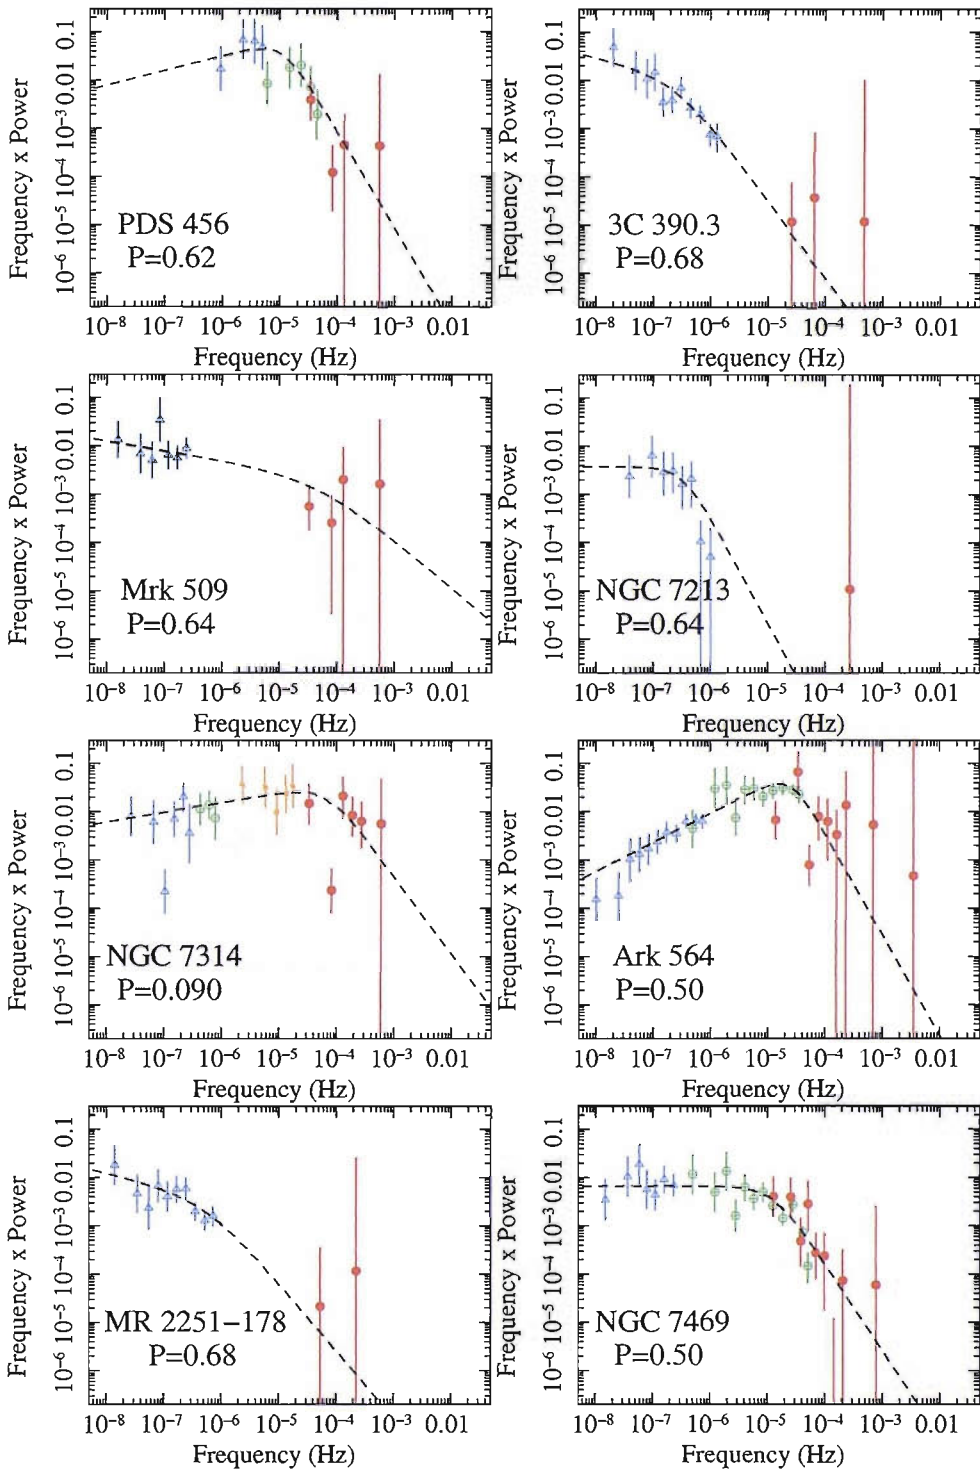


Figure 5.3: *continued.* The best-fitting single-bend PSD for all targets in the selected sample, plotted in $\nu \times P_\nu$ space. See Fig.5.2 for a description of the points.

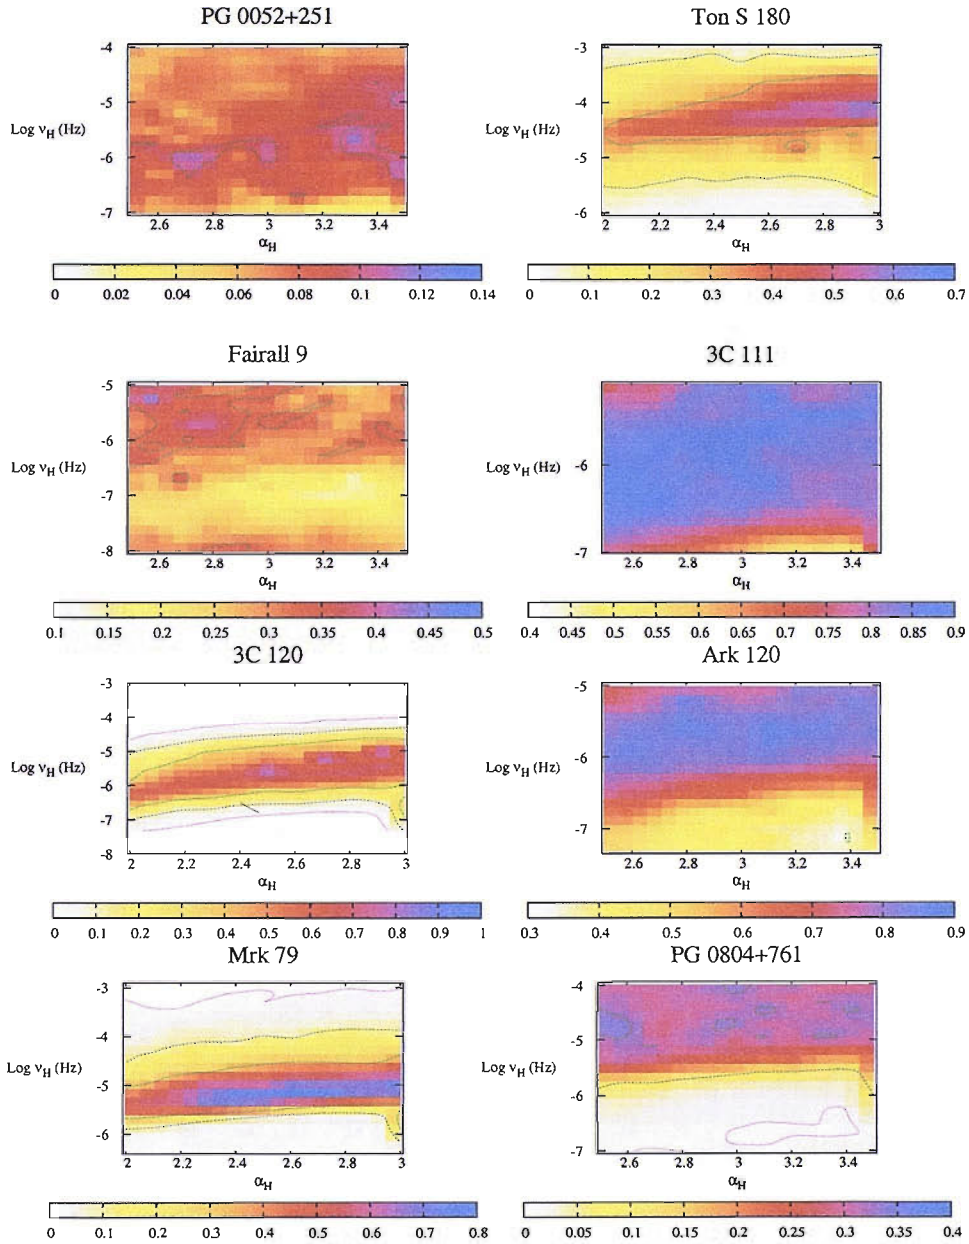


Figure 5.4: Contour plots showing the rejection probability levels on the best fit single-bend PSD model: 68, 90 and 99 per cent rejection probability levels for the bend-frequency, ν_B , and the high-frequency slope, α_H . A colour-bar represents the acceptance probability, with darker colours denoting a high acceptance probability.

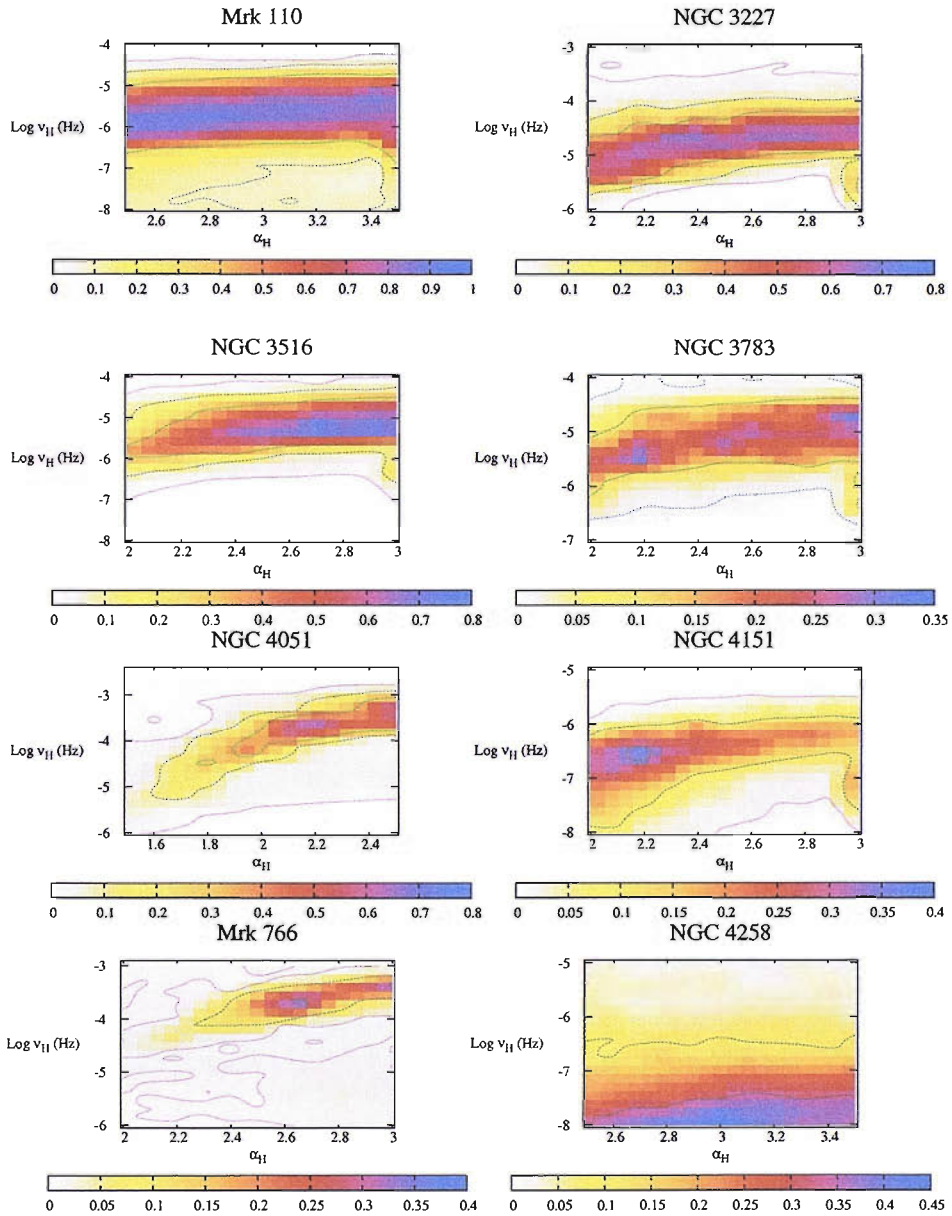


Figure 5.4: *continued.* Contour plots showing the rejection probability levels on the best fit single-bend PSD model: 68, 90 and 99 per cent rejection probability levels for the bend-frequency, ν_{B} , and the high-frequency slope, α_{H} . A colour-bar represents the acceptance probability, with darker colours denoting a high acceptance probability.

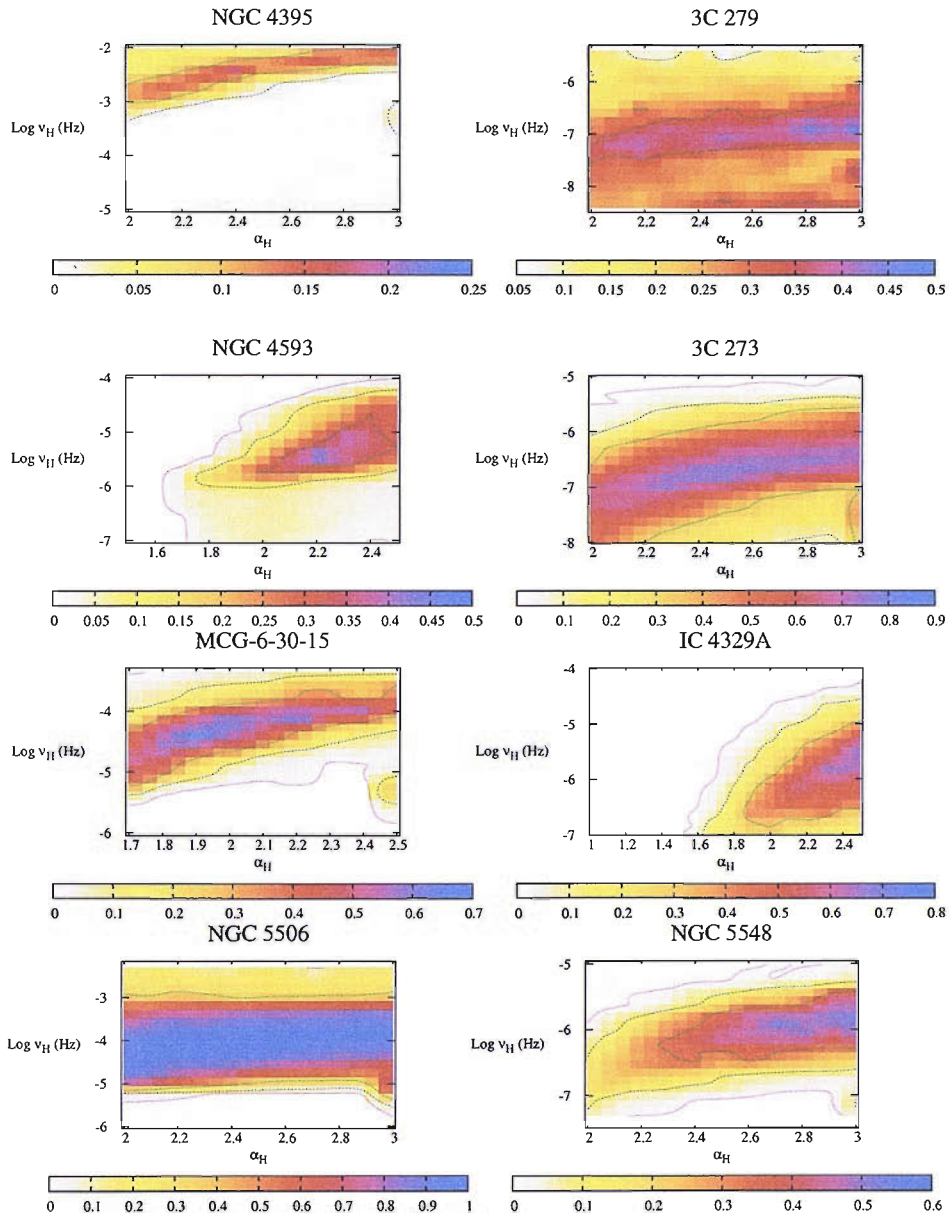


Figure 5.4: *continued.* Contour plots showing the rejection probability levels on the best fit single-bend PSD model: 68, 90 and 99 per cent rejection probability levels for the bend-frequency, ν_B , and the high-frequency slope, α_H . A colour-bar represents the acceptance probability, with darker colours denoting a high acceptance probability.

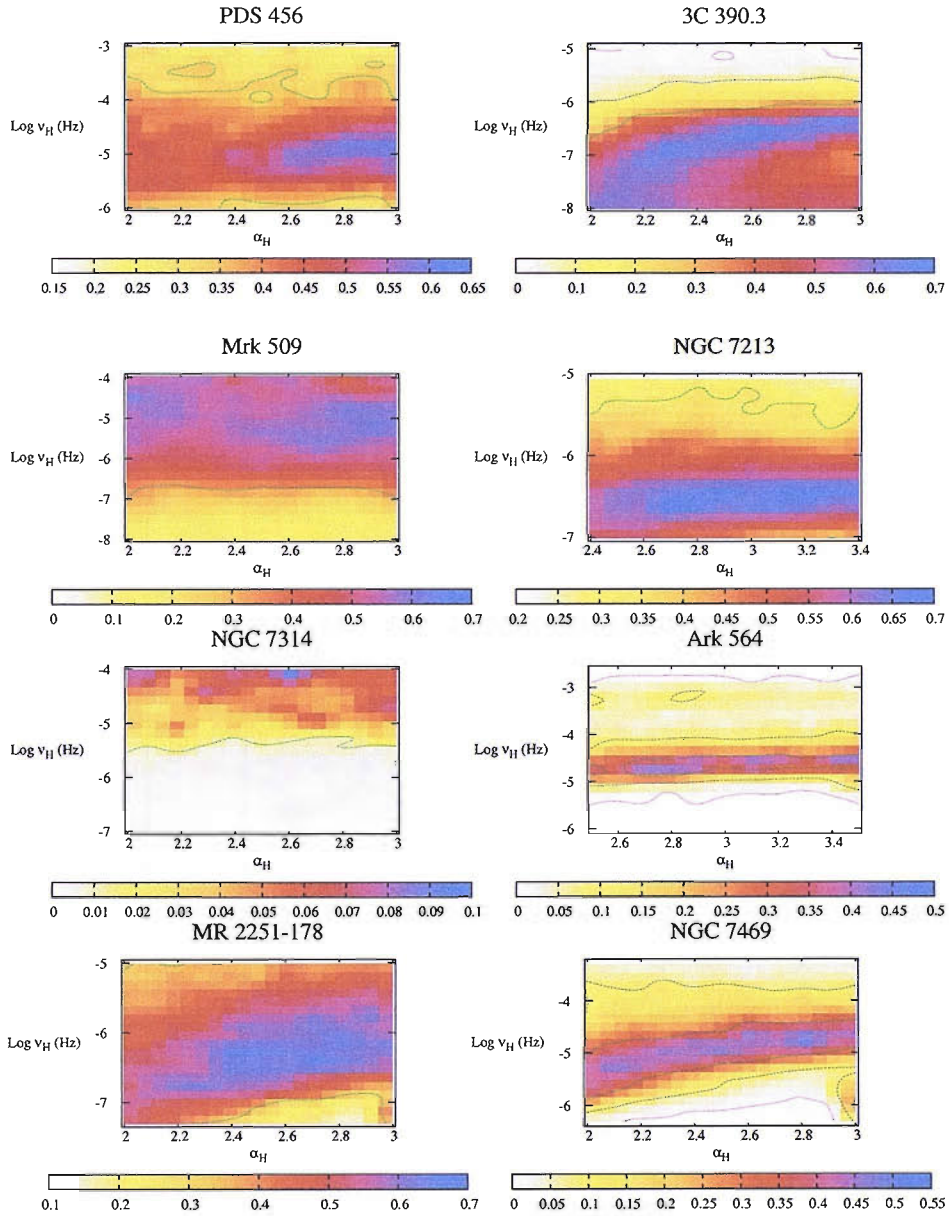


Figure 5.4: *continued.* Contour plots showing the rejection probability levels on the best fit single-bend PSD model: 68, 90 and 99 per cent rejection probability levels for the bend-frequency, ν_B , and the high-frequency slope, α_H . A colour-bar represents the acceptance probability, with darker colours denoting a high acceptance probability.

Table 5.3: Summary of the best-fitting model parameters for the single-bend power-law model to the all targets. The errors on the power-law slope are calculated from the 90 per cent rejection probability levels. An asterisk indicates that an estimate of the limiting value was not produced.

Target Name	Normalisation, A (Hz ⁻¹)	α_H	α_L	ν_B	Acceptance
PG 0052+251	1.7×10^{-4}	3.4_{-*}^{+*}	0.7_{-*}^{+*}	$2.6_{-*}^{+*} \times 10^{-6}$	0.13
Ton S 180	3.22×10^{-2}	3.0_{-*}^{+*}	0.9_{-*}^{+*}	$8.7_{-8.1}^{+20.5} \times 10^{-5}$	0.63
Fairall 9	6.3×10^{-5}	2.6_{-*}^{+*}	1.3_{-*}^{+*}	$6.6_{-*}^{+*} \times 10^{-6}$	0.45
3C 111	3.5×10^{-3}	2.6_{-*}^{+*}	1.1_{-*}^{+*}	$7.6_{-*}^{+*} \times 10^{-7}$	0.89
3C 120	1.8×10^{-4}	2.9_{-*}^{+*}	1.2_{-*}^{+*}	$9.7_{-8.5}^{+23.1} \times 10^{-6}$	0.99
Ark 120	7.9×10^{-3}	3.1_{-*}^{+*}	1.0_{-*}^{+*}	$1.9_{-*}^{+*} \times 10^{-6}$	0.87
Mrk 79	2.0×10^{-1}	$3.0_{-*}^{+0.0}$	$0.8_{-0.1}^{+*}$	$8.5_{-6.0}^{+20.3} \times 10^{-6}$	0.79
PG 0804+761	1.7×10^{-4}	2.6_{-*}^{+*}	1.2_{-*}^{+*}	$1.3_{-1.0}^{+*} \times 10^{-5}$	0.35
Mrk 110	1.7×10^{-3}	3.4_{-*}^{+*}	1.1_{-*}^{+*}	$2.9_{-1.4}^{+11.8} \times 10^{-6}$	0.91
NGC 3227	1.8×10^{-2}	2.8_{-*}^{+*}	1.0_{-*}^{+*}	$2.6_{-1.8}^{+6.1} \times 10^{-5}$	0.73
NGC 3516	6.9×10^{-2}	2.9_{-*}^{+*}	0.9_{-*}^{+*}	$6.6_{-4.6}^{+8.2} \times 10^{-6}$	0.77
NGC 3783	6.6×10^{-3}	2.6_{-*}^{+*}	1.0_{-*}^{+*}	$1.3_{-0.7}^{+1.6} \times 10^{-5}$	0.33
NGC 4051	7.0×10^{-3}	$2.5_{-0.8}^{+*}$	$1.1_{-*}^{+0.1}$	$5.1_{-2.6}^{+4.9} \times 10^{-4}$	0.75
NGC 4151	1.8×10^2	2.2_{-*}^{+*}	0.5_{-*}^{+*}	$2.6_{-2.1}^{+6.1} \times 10^{-7}$	0.39
Mrk 766	4.2×10^{-4}	$3.0_{-0.5}^{+*}$	$1.3_{-0.4}^{+*}$	$3.4_{-1.1}^{+4.2} \times 10^{-4}$	0.33
NGC 4258	7.1×10^{-4}	3.0_{-*}^{+*}	1.3_{-*}^{+*}	$1.0_{-*}^{+37.4} \times 10^{-8}$	0.44

Table 5.3: Summary of the best-fitting model parameters for the single-bend power-law model to the all targets. The errors on the power-law slope are calculated from the 90 per cent rejection probability levels. An asterisk indicates that an estimate of the limiting value was not produced.

Target Name	Normalisation, A (Hz ⁻¹)	α_H	α_L	ν_B	Acceptance
NGC 4395	1.5×10^{-2}	2.9_{-*}^{+*}	1.2_{-*}^{+*}	$6.6_{2.2}^{+3.3} \times 10^{-3}$	0.20
3C 273	8.1×10^{-2}	2.5_{-*}^{+*}	0.8_{-*}^{+*}	$1.7_{-1.6}^{+6.9} \times 10^{-7}$	0.81
NGC 4593	3.3×10^{-2}	$2.2_{-0.3}^{+*}$	0.8_{-*}^{+*}	$3.8_{-2.1}^{+4.8} \times 10^{-6}$	0.46
3C 279	2.5×10^{-1}	3.0_{-*}^{+*}	0.9_{-*}^{+*}	$1.5_{-*}^{+1.9} \times 10^{-7}$	0.49
MCG 6-30-15	1.7×10^{-1}	$1.9_{-*}^{+0.6}$	$0.8_{-*}^{+0.2}$	$3.8_{-2.1}^{+9.2} \times 10^{-5}$	0.67
IC 4329A	2.6×10^{-4}	$2.5_{-0.7}^{+*}$	1.2_{-*}^{+*}	$3.8_{-3.6}^{+9.2} \times 10^{-6}$	0.71
NGC 5506	2.1×10^{-2}	2.3_{-*}^{+*}	0.9_{-*}^{+*}	$5.8_{-0.5}^{+23.4} \times 10^{-5}$	0.97
NGC 5548	4.3×10^{-3}	2.7_{-*}^{+*}	1.1_{-*}^{+*}	$1.3_{-0.9}^{+1.6} \times 10^{-6}$	0.55
PDS 456	4.8×10^{-1}	2.9_{-*}^{+*}	0.8_{-*}^{+*}	$1.1_{-*}^{+*} \times 10^{-5}$	0.62
3C 390.3	1.1×10^{-4}	2.6_{-*}^{+*}	1.3_{-*}^{+*}	$2.6_{-*}^{+16.9} \times 10^{-7}$	0.68
Mrk 509	3.1×10^{-4}	2.0_{-*}^{+*}	1.2_{-*}^{+*}	$5.0_{-*}^{+*} \times 10^{-5}$	0.64
NGC 7213	3.7×10^{-3}	2.0_{-*}^{+*}	1.2_{-*}^{+*}	$3.4_{-*}^{+*} \times 10^{-7}$	0.64
NGC 7314	2.4×10^{-1}	2.6_{-*}^{+*}	0.8_{-*}^{+*}	$6.6_{-*}^{+*} \times 10^{-5}$	0.093
Ark 564	3.6×10^1	3.1_{-*}^{+*}	0.4_{-*}^{+*}	$2.6_{-0.9}^{+3.2} \times 10^{-5}$	0.50
MR 2251-178	4.9×10^{-5}	2.4_{-*}^{+*}	1.2_{-*}^{+*}	$5.7_{-*}^{+*} \times 10^{-7}$	0.68
NGC 7469	6.6×10^{-3}	2.8_{-*}^{+*}	1.0_{-*}^{+*}	$1.3_{-0.7}^{+1.6} \times 10^{-5}$	0.50

5.3.3 Double-bend power-law PSD model

Visual inspection of Fig. 5.3 suggests there may be a few observed PSDs that may be better fit by a more complicated double-bend PSD model. However, given the added complexity of a double-bend PSD model I would expect to find an increased likelihood of acceptance for all sources herein, so I restrict our analysis to two objects (MCG-6-30-15 and NGC 4051) where the observed data is of sufficient quality below the high-frequency bend to examine for a turn-down or second lower-frequency bend within the PSD. There are two more observed PSDs (NGC 3227, and NGC 5506) which tentatively suggest a turn-down in the PSD at low frequencies, but these objects have been recently examined for a second bend (Uttley and McHardy, 2005) and found to be entirely consistent with a single-bend power-law PSD model. NGC 3783 has also been suggested to exhibit a second bend feature, but this object was examined extensively in Chapter 4 (Summons *et al.*, 2007). Ark 564 provides the strongest evidence of a second bend in its PSD and was also thoroughly examined in Chapter 4 (McHardy *et al.*, 2007) using the same double-bend power-law model and fitting procedure. The double-bend power-law PSD model, which I fit to MCG-6-30-15 and NGC 4051 is given by

$$P(\nu) = \frac{A \nu^{-\alpha_L}}{\left[1 + \left(\frac{\nu}{\nu_L}\right)^{\alpha_I - \alpha_L}\right] \left[1 + \left(\frac{\nu}{\nu_H}\right)^{\alpha_H - \alpha_I}\right]},$$

where α_H , α_I , and α_L are the high-frequency slope, intermediate-frequency slope, and low-frequency slope, respectively. The low and high bend-frequencies, respectively, are ν_L and ν_H . In order to provide the best comparison to the low/hard state of Cyg X-1, I fix the low-slope to zero, and the intermediate-slope between the bends to 1. Although it is customary to describe the low/hard state of BHXRBs using multiple Lorentzian components (e.g. Pottschmidt *et al.* 2003), AGN data is sufficiently poor in quality that a double-bend power-law PSD model will suffice as an approximation to the low/hard state (Belloni and Hasinger, 1990; Nowak *et al.*, 1999). Analysis in Chapter 4 (reported in McHardy *et al.* 2007) on Ark 564 finds that it is best fit by a two component Lorentzian model, which is also supported by the lag spectrum (Arévalo *et al.*, 2006). Our analysis is further confirmed by the large observed residuals of the single-bend power-law PSD model fitted to the high-frequency data of Ark 564 (see Fig. 5.3).

I fitted the double-bend power-law PSD model to MCG-6-30-15 and NGC 4051, with the best-fitting model plotted in Fig. 5.5. The best fit parameters are summarised in Table 5.4. The bend frequencies are free parameters so that I might

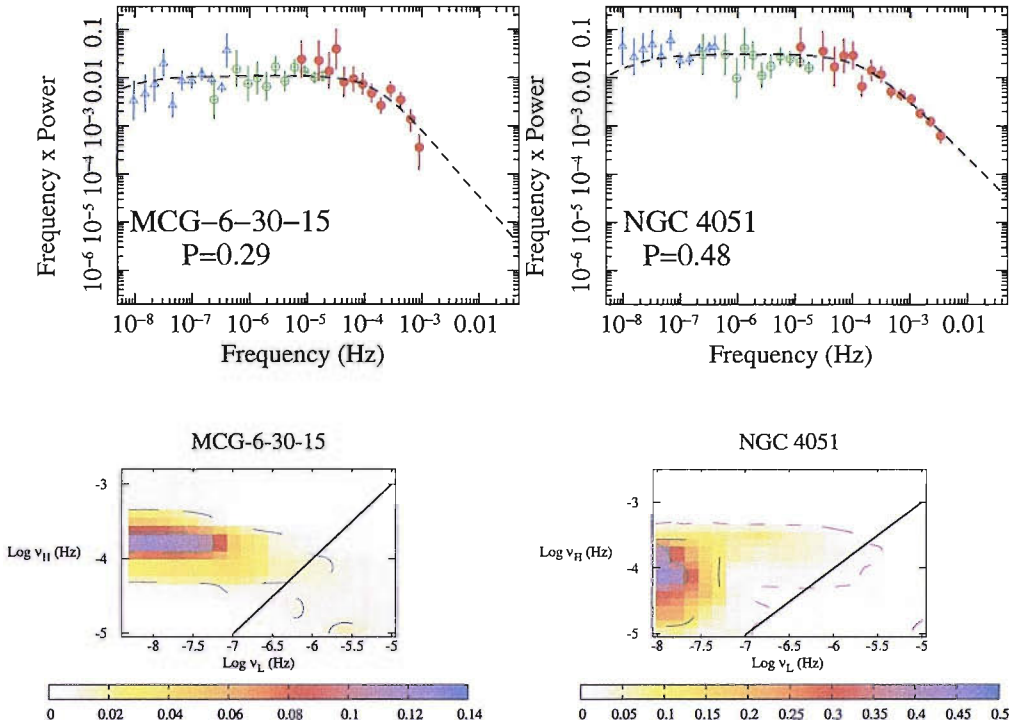


Figure 5.5: The best-fitting double-bend PSD for MCG 6-30-15 and NGC 4051, plotted in $\nu \times P_\nu$ space. See Fig.5.2 for a description of the points.

determine the ratio of the best-fitting bend frequencies, which is typically about a decade in the low/hard state of BHXRBs. The rejection probability levels for the main interesting parameters (ν_H and ν_L) are also plotted in Fig. 5.5. In Chapter 4 I obtained a very good fit (acceptance probability $P=75$ per cent) for the double-bend power-law PSD model to Ark 564 with $\alpha_I = 1.2^{+0.1}_{-0.2}$, $\nu_L = 7.5^{+25.1}_{-5.5} \times 10^{-7}$ Hz and $\nu_H = 2.4^{+2.3}_{-0.9} \times 10^{-3}$ Hz. It is clear that a double-bend power-law model is preferred for Ark 564 over the single-bend model.

The double-bend power-law PSD model fit to NGC 4051 and MCG-6-30-15 do not produce an improvement on the single-bend power-law PSD model. Moreover, the best-fitting low-frequency bend is found at the lowest permitted value. I also plot a solid diagonal line to denote the frequency ratio of 100 between breaks. The possible ratio of bend frequencies does exceed 100 at 90 per cent confidence for MCG-6-30-15 and 99 per cent confidence for NGC 4051. On this evidence, I rule out a low/hard state PSD for both objects.

5.3.4 Single-bend power-law PSD model with a fixed low-slope

Here I fit a single-bend power-law PSD model with a fixed slope to selected targets. For several objects (PG 0052+251, Ton S 180, Fairall 9, 3C 111, Ark 120,

Target	Normalisation, A (Hz ⁻¹)	α_H	ν_H (Hz)	ν_L (Hz)	Acceptance
NGC 4051	3.5×10^3	$2.2^{+0.2}_{-0.5}$	$1.6^{+1.6}_{-1.5} \times 10^{-4}$	$1.0^{+6.3}_{-*} \times 10^{-8}$	0.48
MCG-6-30-15	3.3×10^3	$2.4^{+*}_{-0.6}$	$1.6^{+1.6}_{-0.8} \times 10^{-4}$	$5.0^{+155}_{-*} \times 10^{-9}$	0.29

Table 5.4: Summary of the best-fitting model parameters for the double-bend power-law PSD model to MCG-6-30-15 and NGC 4051. The errors on the power-law slope are calculated from the 90 per cent rejection probability levels. An asterisk indicates that an estimate of the limiting value was not produced.

PG 0804+761, Mrk 110, NGC 4151, Mrk 766, NGC 4258, NGC 4395, 3C 273, 3C 279, NGC 5548, PDS 456, 3C 390.3, Mrk 509, NGC 7213, and MR 2251-178) within our sample the best fit bend-frequency is very close to the lowest measured frequency within the PSD are is not well constrained. Inspection of the measured low-frequency slopes in Table 5.3 suggest that this parameter is often determined to be ~ 1 ; thus, it is reasonable to test the less well measured PSDs against a single-bend power-law PSD model with fixed low slope at 1. A fixed low-slope may afford increased reliability of the estimated bend-frequency for objects with poorly constrained bend-frequencies.

Best fit parameters for the single-bend power-law PSD model with a fixed slope are displayed in Fig. 5.5, Fig. 5.6 and Table 5.5. It is clear that all sources tested against this fixed slope model are well fit, with all rejection probability levels (see Fig. 5.6) exhibiting increased constraint over those presented in Fig. 5.4. The rejection probability contour for Mrk 509 remains completely unbounded at the 90 per cent level for the best-fitting high-slope value.

5.4 Correlation between bend time-scale, black hole mass and bolometric luminosity

Observationally, accreting black holes come in two main flavours: stellar mass black hole systems ($\sim 10 M_{\odot}$) in BHXRBs and supermassive black holes ($\sim 10^6$ – $10^9 M_{\odot}$) in AGN. Both these systems display aperiodic X-ray variability and the existence of a mass-scaling of time-scales has previously been considered (McHardy, 1988; Edelson and Nandra, 1999; Uttley and McHardy, 2005). In McHardy *et al.* (2006) the authors show that by correcting for variations in accretion rate, the bend time-scales in the observed PSD of both AGN and BHXRBs can be physically connected, thereby implying that the accretion process must be mass-scale invariant.

McHardy *et al.* (2006) motivated by an approximate linear scaling of the bend time-scale, T_B , with mass, M_{BH} , and an inverse scaling of T_B with accretion rate, \dot{m}_E (in units of the Eddington limit), hypothesise that $\log T_B = A \log M_{BH} - B \log L_{bol} + C$. The bolometric luminosity, L_{bol} , is used rather than the accretion rate, since the majority of objects discussed here are radio-quiet i.e. \dot{m}_E is well approximated by L_{bol}/L_E . The parameters A, B, and C are determined from a grid search. McHardy *et al.* (2006) take 10 AGN with well determined bend time-scales, good mass measurements and spanning a broad range of accretion rates and find a good fit with $A = 2.10$ and $B = 0.98$. McHardy *et al.* (2006) also show that the fit can be ex-

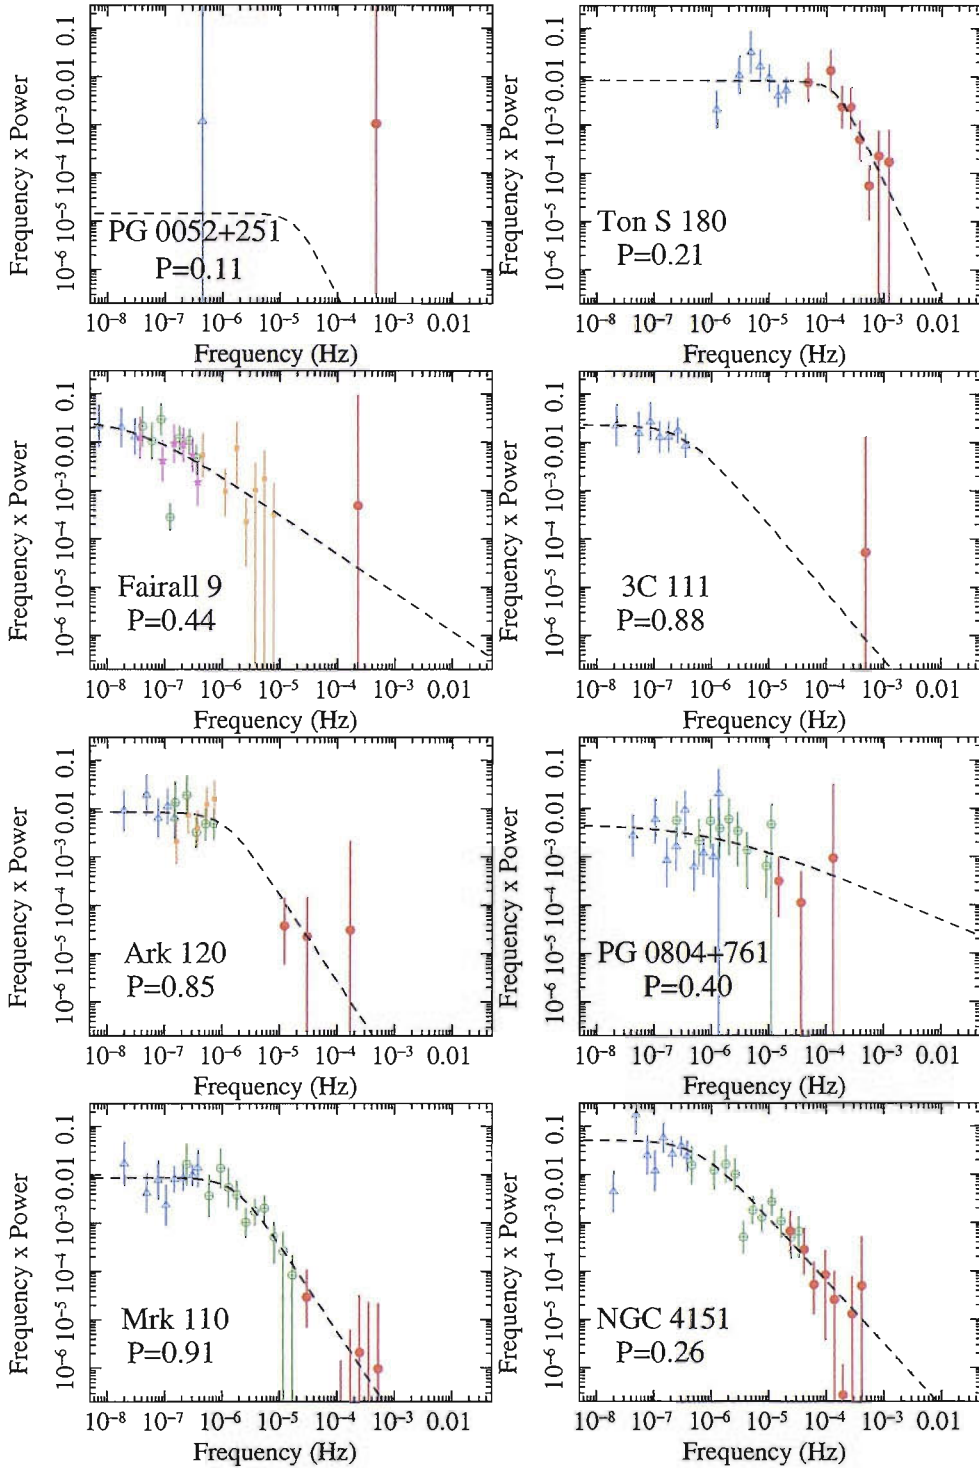


Figure 5.6: The best-fitting single-bend power-law PSD with fixed low-frequency slope for several targets in the original sample, plotted in $\nu \times P_\nu$ space. See Fig.5.2 for a description of the points.




Figure 5.5: *continued.* The best-fitting single-bend power-law PSD with fixed low-frequency slope for several targets in the original sample, plotted in $\nu \times P_\nu$ space. See Fig.5.2 for a description of the points.




Figure 5.5: *continued.* The best-fitting single-bend power-law PSD with fixed low-frequency slope for several targets in the original sample, plotted in $\nu \times P_\nu$ space. See Fig.5.2 for a description of the points.

tended to include Cyg X-1 and GRS 1915+105 in radio-quiet states i.e. the scaling relation is good over several orders of magnitude of mass.

The AGN bend time-scale data used in McHardy *et al.* (2006) were collated from the current literature, including measurements from a number of different authors, which could introduce some small additional scatter into the mass/accretion-rate correlation. Here I present an updated relation between T_B , M_{BH} and \dot{m}_E using a larger sample of AGN data; moreover, the AGN bend frequencies used hereafter have been consistently calculated and presented in this work. The bend-frequency for Ark 564 is taken from Chapter 4 (McHardy *et al.*, 2007) for the ‘hard’ energy range double-bend power-law PSD model fit.

I shortlist 21 AGN from our original list of 32 on which to examine the correlation between T_B , M_{BH} , and L_{bol} . The criterion for the shortlist was that the bend-frequency could be constrained at the 90 per cent rejection probability contour in the single-bending power-law model fits, with free low frequency slope (see Fig. 5.3, Fig. 5.4, and Table 5.3).

I fit the aforementioned list of 21 objects according to the hypothesised scaling relation and determine the best-fitting values of A , B , and C . The 90 per cent rejection probability levels which are quoted throughout this work for T_B are approximately

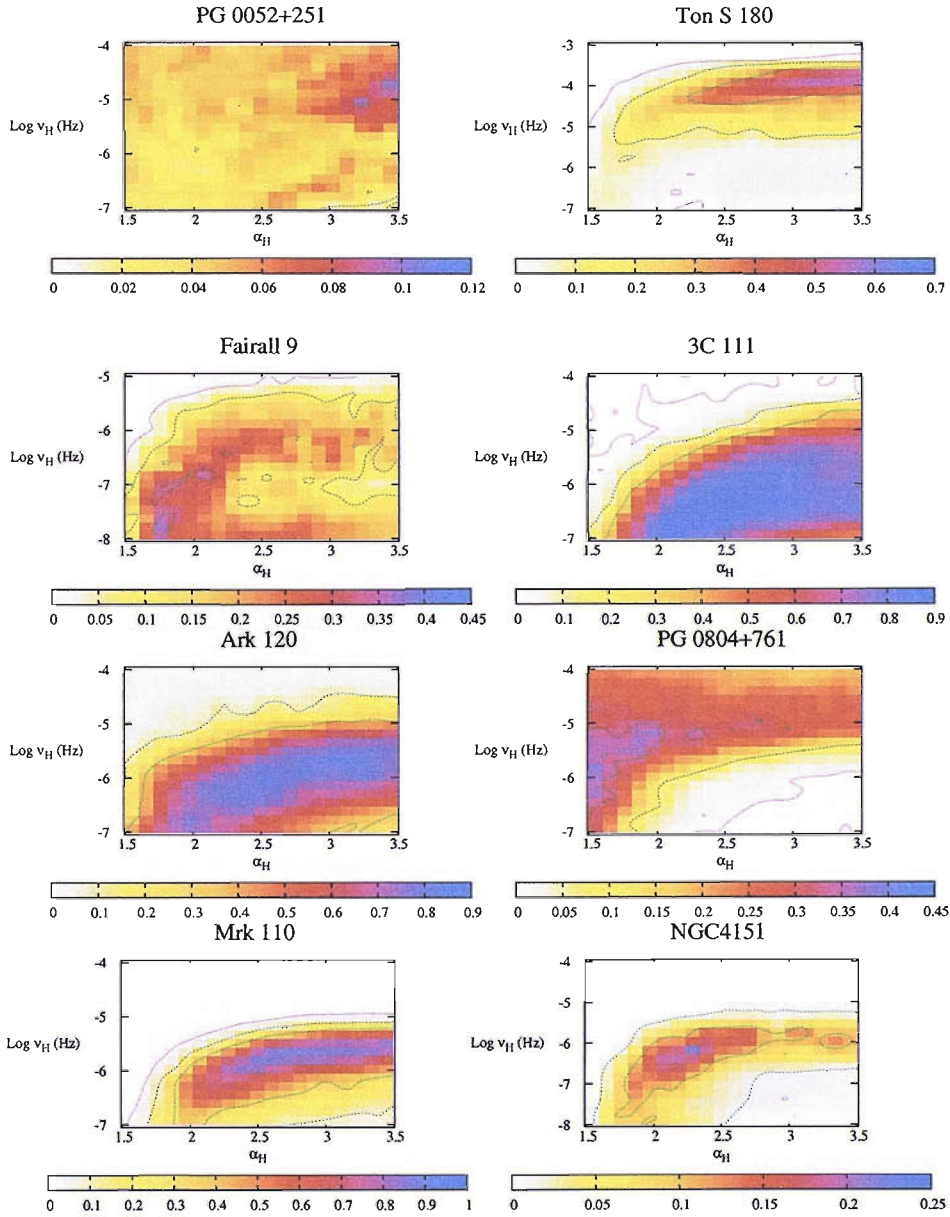


Figure 5.6: Contour plots showing the rejection probability levels on the best fit single-bend PSD model: 68,90 and 99 per cent rejection probability levels for the bend-frequency, ν_B , and the high-frequency slope, α_H . A colour-bar represents the acceptance probability, with darker colours denoting a high acceptance probability.

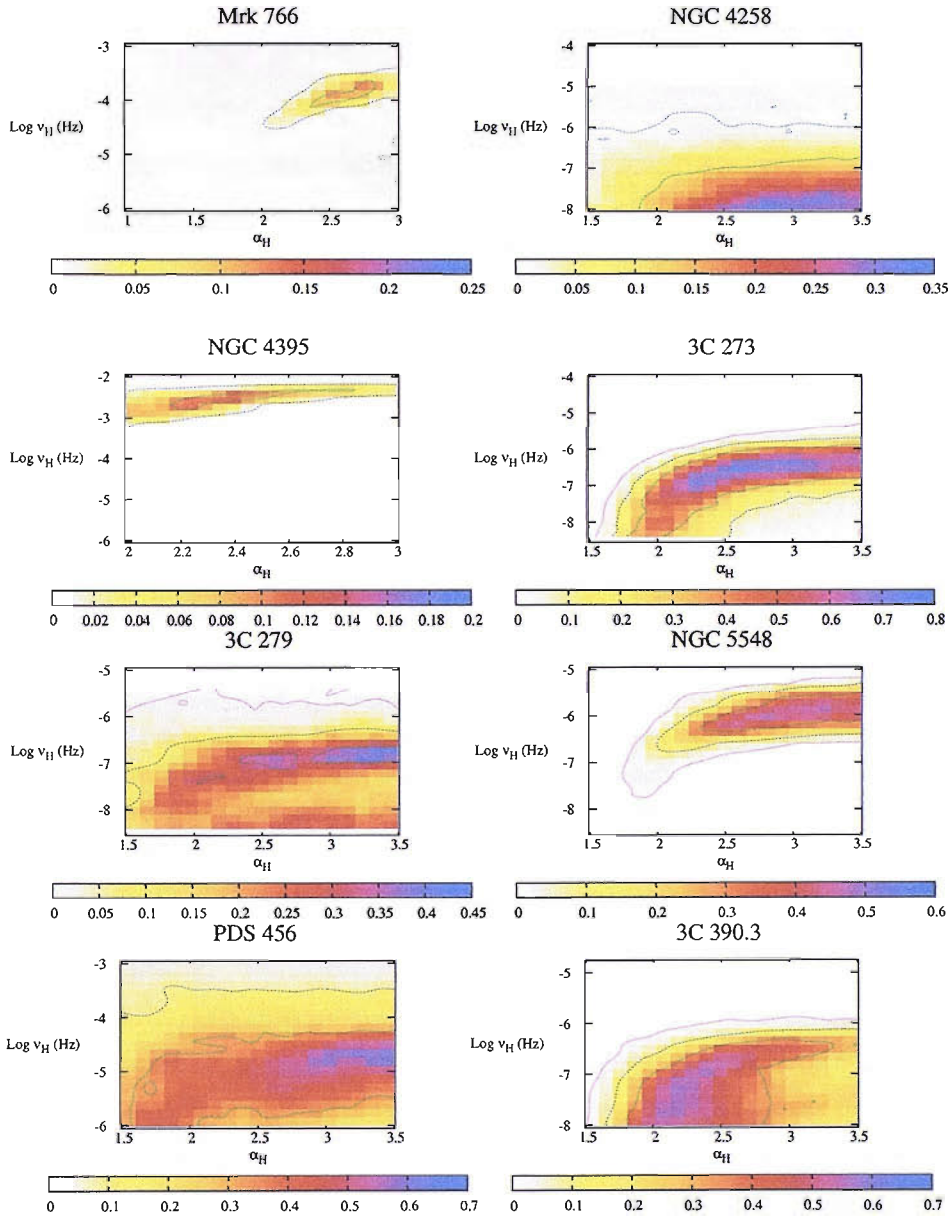


Figure 5.6: *continued.* Contour plots showing the rejection probability levels on the best fit single-bend PSD model: 68,90 and 99 per cent rejection probability levels for the bend-frequency, ν_B , and the high-frequency slope, α_H . A colour-bar represents the acceptance probability, with darker colours denoting a high acceptance probability.

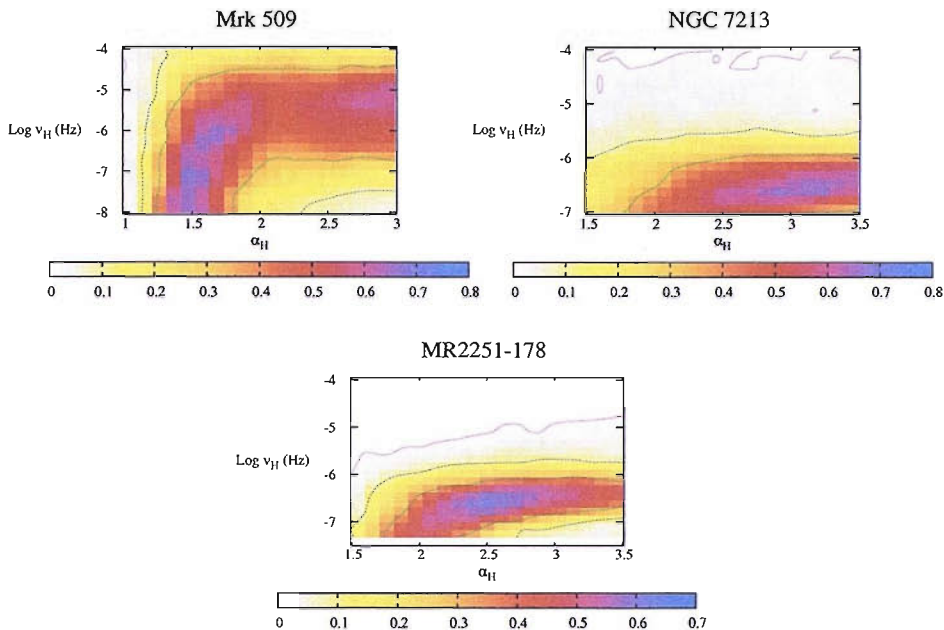


Figure 5.6: *continued.* Contour plots showing the rejection probability levels on the best fit single-bend PSD model: 68,90 and 99 per cent rejection probability levels for the bend-frequency, ν_B , and the high-frequency slope, α_H . A colour-bar represents the acceptance probability, with darker colours denoting a high acceptance probability.

Gaussian and are therefore renormalised to 68 per cent (1σ) probabilities for the purposes of fitting. Moreover since errors are often asymmetric, I take the error which lies in the direction from the predicted to observed value of T_B . Errors on the bolometric luminosity, L_{bol} , are not often provided so I assume a typical 30 per cent error which corresponds to the typical spread in estimated L_{bol} for the same AGN. The bend time-scales, T_B , have all been calculated in this work, with the exception of Ark 564 which was thoroughly examined very recently by us (McHardy *et al.*, 2007). The data for M_{BH} are, where possible, derived from reverberation mapping (Peterson *et al.*, 2004), with the bolometric luminosity data mostly taken from Woo and Urry (2002). Much of the data for M_{BH} and L_{bol} is presented in Table 5.6. The mass of Ton S 180, Mrk 766, MCG-6-30-15 are determined from stellar velocity dispersion measurements (Hao *et al.*, 2005; Botte *et al.*, 2005; McHardy *et al.*, 2005), and the mass of Ark 564 is determined from an [OIII] line width (Botte *et al.*, 2004). For NGC 5506, which is highly obscured, I use a stellar velocity dispersion measurement by Gu *et al.* (2006) (98 km s^{-1}) to determine the black hole mass rather

Target Name	Normalisation, A (Hz ⁻¹)	α_H	ν_B (Hz)	Acceptance
PG 0052+251	1.4×10^{-5}	$3.4^{+0.0}_{-0.2}$	$1.9^{+0.1}_{-1.1} \times 10^{-5}$	0.11
Ton S 180	8.4×10^{-3}	$3.5^{+*}_{-1.8}$	$1.5^{+2.2}_{-1.4} \times 10^{-4}$	0.62
Fairall 9	2.9×10^{-2}	1.8^{+*}_{-*}	$3.4^{+54.3}_{-*} \times 10^{-8}$	0.44
3C 111	2.3×10^{-2}	2.4^{+*}_{-*}	$3.4^{+288}_{-*} \times 10^{-7}$	0.88
Ark 120	8.5×10^{-3}	$2.8^{+0.7}_{-*}$	$1.1^{+18.3}_{-*} \times 10^{-6}$	0.85
PG 0804+761	4.8×10^{-3}	1.5^{+*}_{-*}	$1.1^{+*}_{-*} \times 10^{-6}$	0.40
Mrk 110	8.7×10^{-3}	$2.8^{+*}_{-1.0}$	$1.7^{+6.9}_{-1.6} \times 10^{-6}$	0.91
NGC 4151	5.2×10^2	$2.3^{+1.1}_{-0.4}$	$5.8^{+13.7}_{-3.8} \times 10^{-7}$	0.26
Mrk 766	6.0×10^{-3}	$2.9^{+0.1}_{-0.4}$	$2.9^{+1.5}_{-0.7} \times 10^{-4}$	0.14
NGC 4258	1.0×10^{-1}	$2.8^{+*}_{-1.1}$	$1.5^{+9.9}_{-*} \times 10^{-8}$	0.33
NGC 4395	7.7×10^{-2}	$2.4^{+0.5}_{-0.3}$	$2.9^{+1.5}_{-1.0} \times 10^{-3}$	0.19
3C 273	2.3×10^{-2}	2.6^{+*}_{-*}	$3.5^{+14.0}_{-*} \times 10^{-7}$	0.78
3C 279	3.9×10^{-2}	3.4^{+*}_{-*}	$2.3^{+2.9}_{-*} \times 10^{-7}$	0.43
NGC 5548	2.2×10^{-2}	$3.5^{+0.0}_{-*}$	$1.3^{+3.1}_{-0.9} \times 10^{-6}$	0.54
PDS 456	3.4×10^{-2}	3.0^{+*}_{-*}	$1.7^{+17.8}_{-*} \times 10^{-5}$	0.66
3C 390.3	2.4×10^{-2}	$2.4^{+*}_{-0.7}$	$1.1^{+4.6}_{-*} \times 10^{-7}$	0.66
Mrk 509	2.0×10^{-2}	$1.5^{+*}_{-0.4}$	$7.6^{+*}_{-*} \times 10^{-8}$	0.74
NGC 7213	3.5×10^{-3}	$3.2^{+0.2}_{-*}$	$3.4^{+35.0}_{-*} \times 10^{-7}$	0.70
MR 2251-178	7.8×10^{-3}	2.5^{+*}_{-*}	$2.5^{+16.7}_{-*} \times 10^{-7}$	0.64

Table 5.5: Summary of the best-fitting model parameters for the single-bend power-law PSD model with fixed low-frequency slope at 1 to selected targets. The errors on the power-law slope are calculated from the 90 per cent rejection probability levels. An asterisk indicates that an estimate of the limiting value was not produced.

than an alternative value from the width of the [OIII] lines. The mass estimates for NGC 4151 and NGC 3227 are determined mainly from stellar dynamics (Onken *et al.*, 2007; Davies *et al.*, 2006) and so, although consistent with reverberation estimates, should be more accurate. The short HST reverberation mapping estimate of M_{BH} for NGC 4395 is not thought to be reliable, since only one cycle of the line and continuum variations are measured in each HST observation and so these variations may not be physically related. I therefore use a mass based upon stellar velocity dispersion ($\sigma < 30 \text{ km s}^{-1}$) and low nuclear optical emission (Vaughan *et al.*, 2005; Filippenko and Ho, 2003). I use the mass of Herrnstein *et al.* (1999) for NGC 4258, which is based upon a nuclear water maser and is therefore considered very precise. The resultant fit to the 21 shortlisted AGN produces a reduced $\chi^2 = 1.9$ with $A = 2.13^{+0.33}_{-0.25}$, $B = 0.85^{+0.20}_{-0.18}$, and $C = -2.15^{+0.20}_{-0.30}$, where T_B is in days, L_{bol} is in $10^{44} \text{ ergs s}^{-1}$, and M_{BH} in $10^6 M_{\odot}$. The fitted values are in good agreement with those derived by McHardy *et al.* (2006) for their sample of 10 AGN, i.e. $A = 2.17^{+0.32}_{-0.25}$,

Target	Mass, $10^6 M_{\odot}$	Ref	$\log(L_{\text{bol}}/L_{\text{Edd}})$	Ref
Ton S 180	12.3	H05	0.474	H05
Fairall 9	255^{+56}_{-56}	P04	-1.27	W02
3C 120	$55.5^{+31.4}_{-22.5}$	P04	-0.501	W02
Ark 120	150 ± 19	P04	-1.36	W02
Mrk 79	52.4 ± 14.4	P04	-1.25	W02
PG 0804+761	693 ± 83	P04	-1.01	W02
Mrk 110	25.1 ± 6.1	P04	-0.767	W02
NGC 3227	$15.0^{+5.0}_{-8.0}$	D06	-1.40	W02
NGC 3516	42.7 ± 14.6	P04	-1.44	W02
NGC 3783	29.8 ± 5.4	P04	-1.17	W02
NGC 4051	1.91 ± 0.78	P04	-0.824	W02
NGC 4151	$45.7^{+5.7}_{-4.7}$	O07	-2.03	W02
Mrk 766	3.5	B05	-0.207	U05
NGC 4258	39.0 ± 1.0	H99	-3.52	L96
NGC 4395	$0.05^{+0.05}_{-0.04}$	V05	-2.10	L99
3C 273	886 ± 187	P04	0.306	W02
NGC 4593	$5.36^{+9.37}_{-6.95}$	P04	-0.736	W02
3C 279	269	W02	-0.427	W02
MCG 6-30-15	4.5 ± 1.5	M05	-1.21	R97
IC 4329A	$9.90^{+17.88}_{-11.88}$	P04	-0.313	W02
NGC 5506	7.60	G06	-0.430	M05
NGC 5548	67.1 ± 2.6	P04	-1.06	W02
3C 390.3	287 ± 64	P04	-1.72	W02
Mrk 509	26.0 ± 2.1	P04	-1.22	W02
Ark 564	2.6 ± 0.26	B04	0.483	R04
NGC 7213	100	N95	-3.15	S05
MR 2251-178	102 ± 20	N06	-0.496	E94
NGC 7469	12.2 ± 0.25	P04	0.0967	W02

Table 5.6: Summary of masses and accretion rates.

References used for the masses and accretion rates: H05 - Hao *et al.* (2005); W02 - Woo and Urry (2002); P04 - Peterson *et al.* (2004); D06 - Davies *et al.* (2006); O07 - Onken *et al.* (2007); B05 - Botte *et al.* (2005); U05 - Uttley and McHardy (2005); H99 - Herrnstein *et al.* (1999); L96 - Lasota *et al.* (1996); V05 - Vaughan *et al.* (2005); L99 - Lira *et al.* (1999); M05 - McHardy *et al.* (2005); R97 - Reynolds *et al.* (1997); G06 - Gu *et al.* (2006); B04 - Botte *et al.* (2004); R04 - Romano *et al.* (2004); N06 - Nandra (2006); N95 - Nelson and Whittle (1995); S05 - Starling *et al.* (2005); E94 - Elvis *et al.* (1994)




Figure 5.7: The confidence contours for the black hole mass and bolometric indices A and B . The contours represent the 68 (black), 90 (red) and 95 (green) per cent confidence contours for the dependence of bend time-scales on black hole mass and bolometric luminosity. The larger contours denote a sample of 20 reliable AGN with constrained bend time-scales, see text for sample list. The smaller highly elliptical contours represent the sample of 20 AGN with Cyg X-1 and GRS1915+105.

$$B = 0.90^{+0.30}_{-0.20} \text{ and } C = -2.42^{+0.22}_{-0.25}.$$

In order to determine whether the relation holds over the full range of masses observed in accreting black hole systems, I must include stellar mass BHXRBS. For proper comparison I should only include BHXRBS whose PSDs are well described by the same single-bending power-law PSD model that best describes the observed AGN PSDs (Axelsson *et al.*, 2005; Trudolyubov, 2001). Such sources are rare and the best examples are Cyg X-1 and GRS 1915+105 in high/soft states (radio quiet). Bolometric luminosity measurements (Wilms *et al.*, 2006) were combined with simultaneous bend time-scale measurements (Axelsson *et al.*, 2005) for Cyg X-1, and for GRS 1915+105 the bend time-scale was determined directly from observations and published fluxes (Trudolyubov, 2001) were taken with the distance taken to be 11 kpc. The BHXRBS data used herein is identical to that of McHardy *et al.* (2006), I therefore refer the reader to that paper for the full data and references.

The fit including Cyg X-1 and GRS 1915+105 is better (reduced $\chi^2 = 1.6$) with $A = 2.03^{+0.18}_{-0.14}$, $B = 0.88^{+0.17}_{-0.13}$, and $C = -2.03^{+0.18}_{-0.08}$ and is again in good agreement with

the previously derived values from McHardy *et al.* (2006) of $A = 2.10 \pm 0.15$, $B = 0.98 \pm 0.15$, and $C = 2.25 \pm 0.2$. I assumed a mass of $15 M_{\odot}$ for Cyg X-1; however assuming a mass of $8 M_{\odot}$ or $20 M_{\odot}$ changes the fit parameters only marginally. Peterson *et al.* (2004) regard their mass of NGC 7469 to be in relatively poor agreement with the virial prediction; although, they note that the mass estimate is within 2σ of the virial prediction. Thus, I remove NGC 7469 from the sample and fit the 20 remaining objects. The fit of AGN data is significantly improved (reduced $\chi^2 = 1.2$) with $A = 2.13^{+0.20}_{-0.18}$, $B = 0.98^{+0.20}_{-0.16}$, and $C = -2.3^{+0.20}_{-0.17}$ and is again in good agreement with the previously derived values from McHardy *et al.* (2006). The fit including Cyg X-1 and GRS 1915+105 is better (reduced $\chi^2 = 0.90$) with $A = 2.13^{+0.17}_{-0.15}$, $B = 0.98^{+0.15}_{-0.15}$, and $C = 2.25^{+0.15}_{-0.15}$. The fit to 20 AGN is shown in Fig. 5.7. It should be made clear that the fits for 20 and 21 AGN are consistent i.e. the confidence contours for the AGN and AGN plus BHXR data overlap at 1σ . In Fig. 5.8 I show a projection of the $T_B - M_{BH} - \dot{m}_E$ plane for all 21 AGN (shown as 20 filled circles, with NGC 7469 represent by an orange star), Cyg X-1 and GRS 1915+105. That is, I show the predicted bend time-scales $T_{\text{predicted}}$, determined by $\log T_B = 2.1 \log M_{BH} - 0.98 \log L_{\text{bol}} - 2.3$, are in good agreement with the measured bend time-scales, T_{observed} . Two of three AGN where the bend time-scale is unbounded on one side (3C 390.3 and NGC 4258) are consistent with the fit, but PG 0804+761 is an outlier. The bend time-scale in these three objects is not as well determined as the rest of the sample because it lies close to the lowest frequencies sampled by the PSD and so the slope below the bend, which is required to constrain the bend frequency, cannot be determined with a great deal of accuracy and may introduce additional scatter. However, these objects are required to constrain the lowest time-scales measured. McHardy *et al.* (2006) reported NGC 5506 as an outlier to the plot of predicted versus observed bend time-scale, although they noted that the mass derived from the width of the [OIII] lines implied an accretion rate which was too low for the obscured narrow line Seyfert classification of Nagar *et al.* (2002). However from reanalysis of its PSD and a new mass estimate based upon stellar velocity dispersion (Gu *et al.*, 2006), I now find that NGC 5506 is quite consistent with the fit.

Using our improved sample of 21 AGN, which were consistently reduced and evaluated in this work, I reinforce the result of McHardy *et al.* (2006) - that over a prodigious range of $\sim 10^8 M_{\odot}$ in mass and $\sim 10^3$ in accretion rate, accretion processes behave similarly.

In Fig. 5.8 I show a projection of the $T_B - M_{BH} - \dot{m}_E$ plane for the original 21 sources plus 7 additional sources as green open circles (Fairall 9, Ark 120, 3C 273, 3C 279,




Figure 5.8: A projection of the $T_B - M_{\text{BH}} - \dot{m}_E$ plane for 28 together with Cyg X-1 and GRS 1915+105. The predicted bend time-scales were determined using the best fit parameters for A , B , and C obtained from the best fit of 21 shortlisted objects. The BHXRB data is denoted as filled stars, with GRS 1915+105 represented by the star at the lowermost left corner. The 21 shortlisted AGN are denoted as 20 filled circles, with NGC 7469 represent by an orange star. The 7 less reliable AGN, which were not used in the fitting, are denoted as green open circles (Fairall 9, Ark 120, 3C 273, 3C 279, Mrk 509, NGC 7213, and MR 2251-178) with T_B determined by assuming a fixed low-frequency slope of 1. Even though the 7 additional AGN are not used in the fit, it is clear that they compliant with the overall correlation. I plot the 90 per cent rejection probability levels on T_{observed} .

Mrk 509, NGC 7213, and MR 2251-178) with T_B determined by an assumed low-slope ($\alpha = 1$). The best fit parameters to the 28 AGN is much worse with a reduced $\chi^2 = 1.7$ and reduced $\chi^2 = 1.6$ when Cyg X-1 and GRS1915+105 are included. Moreover, the contours for the AGN on their own and the AGN plus BHXRB do not overlap at the 90 per cent confidence level. The fits to the samples which did not include these poorly measured bend time-scales were very good. I therefore conclude that in these 7 cases there just is not enough data to allow a reliable estimate of the bend frequency or, possibly, that the low frequency slopes differ enough from 1 to affect noticeably the bend time-scale measurement. In any event, we do not include them in our best fit. Four remaining AGN of our initial sample (PG 0052+251, 3C 111, PDS 456 and NGC 7314) are not included at all, since mass data for 3C 111 and PDS 456 are not available and the PSDs of PG 0052+251, NGC 7314 are not

of sufficient quality.

5.4.1 Correlation of bend time-scale with the $H\beta$ optical emission linewidth

The width of the optical emission lines, V , broadened by Doppler motions of the gas in the broad line region (BLR) surrounding the central black hole may be expected to correlate with M_{BH} and \dot{m}_E . McHardy *et al.* (2006) determine that for a virialised BLR, we expect $V^4 \sim M_{\text{BH}} / \dot{m}_E$. Thus for objects where there exists a robust measurement of V (here defined as the FWHM of the variable, broad component, of the $H\beta$ optical emission line, best measured from the RMS optical spectrum) I plot T_{B} versus V in Fig. 5.9 for 18 objects. I primarily take the $H\beta$ FWHM linewidths detected in the RMS spectrum by Peterson *et al.* (2004) for 3C 120, Mrk 79, PG 0804+761, Mrk 110, NGC 3227, NGC 3516, NGC 3783, NGC 4051, NGC 4151, NGC 4593, IC 4329A, NGC 5548, 3C 390.3, and NGC 7469. The $H\beta$ linewidths for Ton S 180, Mrk 766, MCG-6-30-15, and Ark 564 are taken from Turner *et al.* (1999). I do not plot the three remaining AGN (NGC 4258, NGC 4395, and NGC 5506) from the shortlist sample of 21 AGN with constrained bend time-scales. NGC 4258 and NGC 5506 have obscured $H\beta$ emission lines, whereas the $H\beta$ emission line in NGC 4395 is dominated by a narrow component and the broad component is very weak and, as a consequence, poorly determined (Kraemer *et al.*, 1999).

In order to determine the form of the relationship between V and T_{B} , taking proper account of the errors on both parameters, I assume a relationship of the form $\log T_{\text{B}} = D \log V + E$, where D and E are fit parameters, and perform a χ^2 -based parameter grid search, as I did for the parameters A, B and C, earlier in this work. I take the measured errors of T_{B} and assume a 30 per cent error on the measured linewidth. Linewidth measurements between different epochs can vary as much as 50 per cent, but 10-15 per cent measurement errors are more typical; thus, 30 per cent is conservative given the archetypal systematic and statistical uncertainties. For a fit to the aforementioned 18 AGN (Ton S 180, 3C 120, Mrk 79, PG 0804+761, Mrk 110, NGC 3227, NGC 3516, NGC 3783, NGC 4051, NGC 4151, Mrk 766, NGC 4593, MCG-6-30-15, IC 4329A, NGC 5548, 3C 390.3, Ark 564, and NGC 7469) I find $D = 3.9^{+0.89}_{-0.65}$ and $E = -13.3$ (reduced $\chi^2=1.6$), which strongly supports the results of McHardy *et al.* (2006) and the predicted $T_{\text{B}}-V$ scaling relation. However the reduced χ^2 for this fit is poor, which could be a result of the outliers in Fig. 5.9, which are denoted with a crossed-circle: Mrk 110 and NGC 3227. A repeated




Figure 5.9: I show the FWHM of the $H\beta$ emission line in AGN is correlated with the observed power spectral bend time-scales produced in Chapter 5 of this work. The most notable outliers are denoted with a crossed-circle: Mrk 110 and NGC 3227. The error bars on T_B are 90 per cent rejection probability levels.

fit with Mrk 110 removed from the sample yields a formally improved fit result ($\chi^2 = 1.1$) and I find $D = 3.75^{+0.83}_{-0.70}$ and $E = -12.8$, which is also consistent with the results of McHardy *et al.* (2006). Removing both outliers from the fit (Mrk 110 and NGC 3227) results in a further improved fit ($\chi^2 = 0.70$) and I find $D = 3.90^{+0.85}_{-0.68}$ and $E = -13.2$.

It is clear from Fig. 5.9 that there are a few points that lie just off the best fit, with Mrk 110 and NGC 3227 being noticeable outliers. The reduced χ^2 for the fit presented in Fig. 5.9 using all 18 shortlisted AGN is 1.6, which is obviously affected by the outliers (Mrk 110 and NGC 3227), suggests that a further unknown source of error must be invoked in order to explain the additional scatter. This conclusion is further reinforced by the improved fit results when the outliers are removed from the fitting procedure. The error could be down to our initial choice of 30 per cent error on the measured FWHM $H\beta$ linewidth; however, a possible interpretation of the additional scatter could be AGN orientation to the observer. The FWHM measurement of $H\beta$ for Mrk 110 appears to be lower than expected given the bend time-scale (see Fig. 5.9); however, Veron-Cetty *et al.* (2007) suggest

that Mrk 110 appears almost face-on. Whereas the $H\beta$ linewidth for NGC 3227 appears to be too broad for the observed bend time-scale, but there is evidence to suggest that NGC 3227 appears almost edge-on to the observer (Lamer *et al.*, 2003). Our sample of AGNs in this work are biased towards Seyfert 1 galaxies i.e. *RXTE* is flux limited and many Seyfert 2 AGN will be heavily absorbed; moreover, the bias towards Seyfert 1 galaxies would generally restrict the distribution of AGN viewing angles within our sample (i.e. the angle between the observer and the plane of the accretion disk). If the BLR has a disk-like distribution, face-on AGN (e.g. Mrk 110) would be measured to have a reduced $H\beta$ linewidth, unlike edge-on AGN (e.g. NGC 3227) where the $H\beta$ linewidth would be almost maximal, this viewing angle dependence could be reflected in Fig. 5.9. Therefore I tentatively suggest that I have speculative evidence for a relation between viewing angle and linewidth, implying that the BLR is not isotropic.

5.5 Discussion and Conclusions

I recapitulate the following main results of our Monte Carlo analysis of 32 AGN PSDs.

I find that a simple unbroken power-law PSD model is sufficient to currently explain the broadband PSD of eight sources (MR 2251-178, PDS 456, 3C 390.3, NGC 4258, Mrk 110, PG0804+761, Ark 120, and Fairall9); however the slopes are generally measured to be fairly steep, suggesting a high-frequency slope value with a break at lower frequencies that has not been reliably measured with the current data. The fifteen remaining sources are well fit by a ‘soft’ state model, with a single bend in the power-law PSD model. Although the aforementioned eight sources are better fit by a single-bend power-law PSD model, the resultant bend-frequency is not well constrained and not statistically required. I find that a ‘hard’ state model, with a double-bend in the power-law PSD model, is not required for NGC 4051 and MCG-6-30-15; however, the results of the single-bend power-law model fit to Ark 564 do suggest, in agreement with McHardy *et al.* (2007), that a more complicated double-bend power-law model is warranted, thus Ark 564 is often referred to as a very-high state AGN. NGC 7314 is unique in our sample in that it is not well fitted by an unbroken power-law and the single-bend power-law PSD model only offers marginal improvement ($P=9$ per cent); however given the stochastic nature of AGN timing we might expect 1/24 AGN PSDs to be relatively poorly fit by the single-bend power-law PSD model even if it is the correct underlying model.

For objects where the bend-frequency is determined close to the lowest probed frequency (Ton S 180, Fairall 9, Ark 120, PG0804+761, Mrk 110, NGC 4151, Mrk 766, NGC 4258, NGC 4395, NGC 5548, PDS 456, 3C 390.3, Mrk 509, and MR 2251-178), I fitted a single-bend power-law model and assume a fixed low-slope ($\alpha = 1$). All objects fitted by this model produced better constrained rejection probability contours, which is to be expected given the reduction in model parameters, and were broadly consistent with the full single-bend power-law fits. Fairall 9, Ark 120, PDS 456, and MR 2251-178 all exhibited a notable improvement and produced moderately well-constrained contours. Mrk 509 is unusual in that the contour produced in Fig. 5.6 is completely unbounded at the 90 per cent rejection probability level.

Many of the objects presented herein have been previously analysed using similar techniques to the one employed in this work, and I find that our results are broadly consistent (e.g. see Markowitz *et al.* 2003; Uttley and McHardy 2005). Several PSDs presented here are not been previously analysed using PSRESP (Ton S 180, 3C 120, Ark 120, Mrk 79, PG 0804+761, Mrk 110, NGC 4593, IC 4329A, PDS 456, 3C 390.3, NGC 7314, NGC 7213, MR 2251-178, and NGC 7469) and some of these objects have PSDs that have not previously been published.

Using an improved PSRESP implementation to evaluate the bend time-scales of the AGN described in the work along with published black hole masses and luminosities, I calculate the relationship between these aforementioned properties. The findings presented herein support the idea that nuclear accretion in BHXRBS is similar to larger scale AGN accretion; moreover, the optical linewidth of $H\beta$ is correlated with the bend time-scale over a large range of masses. Consequently, I further endorse the findings of McHardy *et al.* (2006) - that over a broad range of masses and accretion rates the X-ray variability and optical linewidth can be explained as merely depending on M_{BH}/\dot{m}_E .

The more you know, the less you understand.

TAO TE CHING

6

Conclusions

X-ray timing studies of accreting black hole systems provide powerful insights into the nature of accretion physics and the strong gravity environment around black holes. In this concluding chapter I will further discuss some of the implications of the results presented in this work and investigate possible future avenues of research. Firstly, I will review our Monte Carlo software (PSRESP) and describe some of the possible modifications and optimisations that may need to be implemented for future work. I will then discuss the implications of the results presented in Chapters 4 and 5 with regard to the AGN-BHXRBB connection, and whether or not the relation between black hole mass, accretion rate and bend time-scale can be used to estimate the black hole mass of *ultra-luminous X-ray sources* (ULXs) i.e. do ULXs harbour intermediate mass black holes? I will also mention the possibility of determining AGN emission states (analogous to BHXRBB emission states) and locating the, thus far, missing QPOs. I will ultimately conclude this work by discussing how future high-energy missions might contribute to AGN timing analysis, and how next generation X-ray telescopes might allow an entirely new sample of AGN to be observed, namely low-mass low-luminosity AGN.

6.1 The future of PSRESP and possible optimisations

In Chapter 2 I described the basic tools used in timing analysis along with some of the caveats that must be considered in order to undertake robust analysis of an unevenly discretely sampled X-ray signal. In Chapter 3 I further developed the concepts introduced in Chapter 2 and explained how the aforesaid techniques can be implemented in Monte Carlo software to rigorously determine the power spectrum of AGN time-series data afflicted with sampling distortions. The basic implementation of our Monte Carlo software was pioneered by Uttley *et al.* (2002); however, this original implementation makes no attempt to simulate the highest frequencies often measured in AGN power spectra and also suffers significant problems when determining the acceptance probability for AGN data considerably distorted by red-noise leak. The improvements outlined in Chapter 3 namely the modification to the Poisson noise level calculation, high-frequency data simulations and acceptance probability determination, tackle many of the original limitations of PSRESP, but there are still areas where the software could be refined.

Even since the days of its original conception, PSRESP is computationally intensive, mainly due to the vast number of calculations required by the light curve simulation and Fourier transform algorithms. Fortunately if the predictions of Moore's law¹ continue, then many of today's computational limitations could be overcome in the near future with the next generation of computer hardware. For example, larger parameter spaces and more complicated models could be examined on a desktop machine; moreover, *RXTE* monitoring light curves could be simulated with the preferred temporal resolution of 1000 seconds (see Chapter 3). Since 2002, despite the various computationally intensive improvements, PSRESP has still increased in speed by a factor of ~ 3 . In a few years the current simulations that require a compute cluster, such as Iridis2, should hopefully become trivial to run on a desktop machine.

In terms of improvements that can be made to PSRESP to better analyse X-ray data, I have short-listed three main areas for development.

6.1.1 Statistics

As outlined and presented in Chapter 3 I have made several improvements to the original version of PSRESP first presented by Uttley *et al.* (2002). The majority of

¹The number of transistors that can be inexpensively placed on an integrated circuit is doubling approximately every two years.

the improvements, as I have already discussed, focus on the statistical calculations in PSRESP. I have implemented an improved method of incorporating the Poisson noise level, but I have also changed the way PSRESP determines acceptance probabilities in the presence of sampling distortions. The distortions associated with irregular sampling (aliasing and red-noise leak) introduce spurious power throughout the measured power spectrum, which produce a distorted representation of the underlying power spectral shape. The modification outlined in Chapter 3 corrects the distortion due to the redistribution of power throughout the power spectrum, which is essentially due to correlated frequency bins, but it is computationally intensive. If a test statistic can be found or developed that is not as significantly affected by correlated frequency bins, PSRESP could be both simplified and further optimised. A new choice of test statistic might negate many of the distribution problems encountered in this work, which would yield a significant computational improvement.

Another possible modification to PSRESP concerns the Poisson noise level determination. In Chapter 3 I reviewed a modification I made to PSRESP to include a Poisson noise level component that varied with frequency and it was directly determined from the observed data. The Poisson noise level is *not* simulated with the individual realisations, it is merely determined using an analytical calculation. This approach is generally good and practical; however, in sources where there exists significant source variability in the Poissonian limit i.e. for a faint AGN just within the detection limit of *RXTE*, the Poisson noise component could significantly affect the timing result. Fortunately, most of the AGN previously observed and presented in this work are well detected by *RXTE* and are not dramatically affected by Poisson noise (except at high frequencies in the power spectrum), but future sources of interest, such as low-luminosity AGN (LLAGN), could be dramatically affected. Therefore the effect of Poisson noise must be considered for weak sources, where the signal itself contains a significant contribution due to Poisson noise, if the timing results are to be considered robust. Simulation of the Poisson noise level could be included in PSRESP as a brute-force method, that is Poisson noise could be simulated directly and coupled with the realisation, but this method would be extremely computationally intensive and require an additional free parameter in the fitting routine (normalisation parameter), which would make current power spectral modelling computationally prohibitive. It is currently unclear how the effect of Poisson noise could be rigorously included without intensive simulation, but it must be considered if faint objects just above the detection threshold of a given X-ray telescope are to be reliably analysed. Nonetheless in a few years time, according to Moore's

law, computers should be powerful enough to directly simulate the Poisson noise component and determine the effect of the noise level without the entire simulation being computationally prohibitive.

6.1.2 Non-linear light curves and blazar data

There is a significant amount of blazar data in the *RXTE* archive. Blazars differ from the typical AGN found in our *RXTE* monitoring campaign, in that they display more highly variable and overtly non-linear X-ray light curves - probably a result of observing the relativistic jet feature associated with these type of objects. Moreover, future high-energy satellite missions such as *GLAST* (see Section 6.8.3) would preferentially observe the very-high (GeV/TeV) emission from these types of sources. PSRESP will be an important tool for analysing the high-energy light curves of blazar data, but work needs to be done to determine whether the linear light curve simulation routine in PSRESP is sufficient when working with such highly variable sources. This point was first considered in Section 3.5. Uttley *et al.* (2005) have shown that real X-ray light curves of AGN and BHXRBS are better described by a non-linear process, and they also show that linear light curves can be made more akin to observed light curves by the process of exponentiation. However, exponentiation changes the underlying power spectrum in a way that cannot be determined analytically.

In order to provide a full and thorough treatment of highly variable blazar data, a method of simulating non-linear data will have to be developed. The method will have to simulate data in a fast and efficient manner (consider Timmer and König 1995), but also produce non-linear light curves while retaining information pertaining to the true undistorted power spectral shape of the non-linear data.

The area where PSRESP will be truly invaluable is determining the confidence on possible periodicities in high-energy blazar light curves. Period analysis of blazar light curves is an active research field with an enormous amount of publications claiming to have found periodicities in light curves of various energy bands (e.g. Wu *et al.* 2006; Liu *et al.* 2007); however, the vast majority of these publications make crucial and fundamental mistakes regarding the stochastic nature of blazar light curves, which undermine the reliability of the result. PSRESP would be an extremely useful tool in determining the aforesaid significance and could be used to test whether a periodic signal is really present. It should be made clear that a periodic signal is possible, since *Very Large Baseline Interferometer (VLBI)* observations of relativistic jets suggest that these jets are helical in structure and might

precess (e.g. McHardy *et al.* 1990; Hong *et al.* 2004); therefore, it is possible that some periodicities might be found in jet emission. My previous statements regarding the blazar periodicities pertain to the method by which these so-called periodicities are determined, not the existence of periodicities.

GLAST will produce excellent high-energy light curves which will, of course, be examined for periodicities. In order to statistically determine whether periodicities exist Monte Carlo simulations of red-noise data must be used, which is a natural future application of PSRESP. PSRESP would require minor modification to include periodicity models, but long-term monitoring with *GLAST* should allow the current models to be used to determine any possible bend time-scales. *GLAST* data would be an exciting new application of PSRESP.

6.1.3 Optical PSDs

The process of determining X-ray power spectra has become a relatively straightforward task as many of our *RXTE* monitoring campaigns are designed to maximise PSD coverage in the most efficient manner possible. Space-based observations are generally much better suited to either extended periods of continuous monitoring or regular periods of sampling than ground-based observations. For this reason regular long-term monitoring of AGN in X-rays is often less problematic than equivalent monitoring at optical energy ranges from ground-based observatories; moreover, high-Earth orbiting X-ray satellites such as *XMM-Newton* will always provide superior uninterrupted observations than ground-based observatories.

Nonetheless, there is beginning to be enough data to produce good power spectra of AGN data in the optical energy range. Previously, NGC 5548 was the only AGN with sufficient optical flux data to produce an optical band power spectrum (Czerny *et al.*, 1999). Czerny *et al.* (1999) show that the optical power spectrum of NGC 5548 is well described by a broken power-law that breaks at $\sim 0.001 \text{ days}^{-1}$; however, Czerny *et al.* (1999) produced their result without using Monte Carlo methods and therefore do not deal with the sampling distortions. The next logical step of this work is to apply PSRESP to our current and forthcoming optical monitoring data. This project is currently underway and Elmé Breedt will be applying PSRESP to provide headway into the power spectral analysis of optical data.

6.1.4 Optimisations

I have already considered various optimisations to PSRESP, with particular attention granted to computational efficiency and the downhill simplex non-linear optimisation algorithm (see Chapter 3). It was decided that the various optimisation methods considered were not appropriate for the problems examined in this work. The standard grid search is too useful, since the resultant rejection probability levels are used to define the errors bounds of the best fit parameters. Although the decision was made to continue using the grid search, there are still possible avenues of optimisation that maintain the usefulness of the grid search. One such optimisation could be adaptive mesh refinement.

Adaptive mesh refinement is a numerical technique that resizes the grid resolution during a simulation and has historically been used in hydrodynamical simulations where grids can be enormous. In power spectral model fitting adaptive mesh refinement could be implemented to undertake the standard grid search of parameters, but in the region of a feature, such as a bend in the power-law, the grid would increase in resolution to provide a more robust estimate of the location of the feature. Perhaps most importantly, adaptive mesh refinement would still produce rejection probability levels with which to produce the errors on the input parameters.

6.2 Cycle 12 of *RXTE*: NGC 3783

In Chapter 4 I used public archive data along with our own proprietary *RXTE* data of NGC 3783 to produce a broadband power spectrum against which several different models could be tested. *XMM-Newton* data was used to constrain the highest frequencies of the power spectrum of NGC 3783. The power spectrum of NGC 3783 does exhibit suggestive features regarding a possible second bend and a speculative QPO feature in the medium time-scale *RXTE* X-ray light curve. However, both these features are not statistically required and, as previously considered, a QPO or second bend feature might be surprising given the observational properties of NGC 3783, such as its radio quietness and moderate accretion rate (see Chapter 4 for a discussion, also published in Summons *et al.* 2007). The medium time-scale *RXTE* data of NGC 3783 is poorly fit by all simple power-law models, which still highlights this data as curious.

Given the peculiar nature of the medium time-scale NGC 3783 power spectrum (between 10^{-6} and 10^{-5} Hz), we (Primary Investigator Ian M^cHardy) have been awarded a large amount of time (512 ks) to observe NGC 3783 for 128 days at




Figure 6.1: The dashed black line is the underlying single-bend power-law model with Lorentzian component, with the parameters taken from Chapter 4. The blue crossed data points represent the existing *RXTE* data used in Chapter 4 to initially test for a Lorentzian component in the single-bend power-law, and the red triangular data points represent the predicted new *RXTE* data assuming that the single-bend power-law model with Lorentzian component parameters presented in Chapter 4 are correct.

quarter day sampling to further probe the possibility of a QPO feature with *RXTE* during Cycle 12. On the assumption our initial analysis presented in Chapter 4 of the single-bend power-law model with a Lorentzian feature is correct, Fig. 6.1 exhibits the power spectrum we expect to measure with the forthcoming Cycle 12 *RXTE* data. The blue crossed data points in Fig. 6.1 is the existing NGC 3783 data (*RXTE* and *XMM-Newton*) used in Chapter 4 to initially test for a Lorentzian component and the red triangular data points are the predicted *RXTE* data. The dashed black line is the underlying model.

The additional *RXTE* data would either confirm the presence of a QPO feature or further dismiss the initial “detection” as entirely stochastic. A firm QPO detection would represent the first of its kind in AGN data and it would be a significant result in AGN timing.

6.3 Is Ark 564 unique?

In Chapter 4 I fitted the power spectrum of Ark 564 with a variety of models, which probed the band-limited nature of the power spectrum. I found that the power spec-

trum was well fitted by a double-bend power-law model consistent with a very-high state interpretation, but the most significant result was a robust two-component Lorentzian model fit. Multiple Lorentzian component models are customary in BHXRB PSD analysis and Ark 564 was the first AGN to have a similar model preferred over a power-law model. The two-component Lorentzian PSD model coupled with the lag spectrum provide strong evidence that the Lorentzian component interpretation is correct. The fit results along with the very-high accretion rate suggest that Ark 564 is a very-high state object and further strengthens the link between AGN and BHXRBs; however, an important question is whether Ark 564 is unique among the AGN population?

Ark 564 is a NLS1 and is unlike the majority of AGN observed with *RXTE* in that it has a super-Eddington accretion rate and has a relatively low black hole mass. To determine whether Ark 564 is unique more NLS1 objects must be observed, since their low-mass and high accretion rate seem to be the properties that allow the second bend in the power spectrum to be measured. Unfortunately there are not many local objects easily observable with the current generation of X-ray telescopes that fulfil this criteria. Ton S 180 is one such object with low-mass and super-Eddington accretion rate i.e. similar to Ark 564. The power spectrum of the NLS1 galaxy Ton S 180 (shown in Fig. 6.2) is not as well determined as Ark 564, but it does exhibit suggestive ‘bumps’ which *could* be fit by the same two-component Lorentzian model. A robust two-component Lorentzian model fit to the PSD of Ton S 180 would be a very important result for the AGN-BHXRB connection and prove that Ark 564 is not unique; however, such a result will have to wait until long time-scale monitoring of Ton S 180 can be undertaken to determine if the broadband power spectrum is indeed band-limited.

In Sections 6.6 and 6.8 I will revisit this point and consider what we might expect to find with the next generation of X-ray observatories and hopefully answer the question of whether Ark 564 really is unique.

6.4 Survey results

Determining AGN power spectra and corresponding bend frequencies is not an easy task. Not only do sophisticated Monte Carlo simulations need to be undertaken to correctly model the data, but light curves of sufficient duration must be obtained. The work presented in Chapter 5 represents the most thorough and consistent X-ray power spectral survey of AGN ever undertaken. Many of the light curves included




Figure 6.2: The broadband power spectrum of Ton S 180. Despite a good single-bend power-law PSD model fit, the shape of the power spectrum is suggestive of a two-component Lorentzian model; however, such a complex model cannot presently be justified and more data is required to confirm any such hypothesis. Nonetheless, given the very high accretion rate of Ton S 180, such a hypothesis remains a very valid possibility.

in this work have been produced from ~ 10 years of X-ray monitoring with *RXTE*. The main focus of this work has been the consistent and robust determination of AGN power spectra. In Chapter 5 I consistently analysed 32 AGN power spectra using our Monte Carlo software, PSRESP, and fitted a variety of PSD models to each target. As a result of the power spectral survey, a sample of bend-frequencies could be tabulated with other observable AGN quantities such as black hole mass, accretion rate and linewidth of $H\beta$.

In Chapter 5 I further confirmed the results of McHardy *et al.* (2006) - that over a range of $\sim 10^8$ in black hole mass and $\sim 10^3$ in accretion rate, accretion processes behave similarly. I endorse this result using twice the number of AGN bend frequencies than those used in McHardy *et al.* (2006). I also reexamined the proposed correlation between bend time-scale and $H\beta$ and found that these observables are indeed correlated as one would expect for a virialised BLR. However I find speculative evidence for a relation between viewing angle and linewidth, implying that the BLR is not isotropic.

In order to further improve the correlations presented in this work ($T_B - M_{BH} - \dot{m}_E$ and $T_B - H\beta$), long time-scale monitoring of AGN must continue. Long time-scale

monitoring not only allows the prospect of a second bend-frequency to be explored in the existing sample of AGN, but it will afford better determination of the high-frequency bend used in the aforementioned correlations. The sample of good bend time-scales must also be increased, which can be done by starting new observing campaigns of previously unmonitored sources. Moreover, with an increased sample of bend time-scales and $H\beta$ measurements the correlation between the two observables can be explored for further insight into the BLR and possible geometry. Other than increasing the sample of AGN with well measured bend time-scales to refine the correlations presented herein, it might become possible to determine if the proposed correlation remains true for AGN in different emission states, provided monitoring of new and existing sources continue.

6.5 Intermediate mass black holes

With an ever growing sample of AGN with well determined bend time-scales, a robust correlation between bend time-scale, black hole mass, and accretion rate would be useful when attempting to determine the black hole mass of sources that, as yet, do not have any estimates determined by conventional means.

ULXs are a bright source of X-rays that are not necessarily situated in the nucleus of a galaxy, and their luminosity generally exceeds the Eddington luminosity for neutron star or stellar mass black hole systems. One possible way to resolve the apparent super-Eddington luminosity is to suppose ULXs actually harbour intermediate-mass black holes (IMBH) (Winter *et al.*, 2006). Observations of IMBHs have never been confirmed, but there is no reason theoretically why they should not exist. Determining the mass of ULXs by conventional means (i.e. time resolved optical spectroscopy) is not possible, but X-ray timing has the potential to estimate the mass by merely evaluating their power spectra and determining the characteristic time-scales of variability (Cropper *et al.*, 2004; Strohmayer *et al.*, 2007). If the power spectra of ULXs exhibit a broken power-law form, their mass could be estimated from the work presented herein. The robust determination of an IMBH in ULXs would be an important and exciting result. However, Goad *et al.* (2006) suggest that ULXs could in fact be very-high accretion rate analogs of BHXRBs with more moderate masses ($\sim 100 M_{\odot}$).

6.6 Accretion state connections

A recurring theme of this work is the relationship between accreting black holes at all mass-scales, from the SMBHs in the centre of active galaxies through to the stellar mass black holes in X-ray binary systems. Timing studies of BHXRBs reveal that these systems can be found in a variety of accretion or emission states (see Chapter 1); however, the majority of AGN found thus far, with the exception of Ark 564, have power spectra consistent with the soft state found in BHXRBs (Cyg X-1 is the BHXRB often cited). Even though there appears to exist a link between accreting black holes at all mass-scales (see Chapter 5) there is still a lack of evidence of AGN found in other accretion states. I will now consider why the majority of AGN power spectra are ‘soft’ and where the objects in other accretion states may be lurking.

Most of the good long time-scale X-ray light curves taken with *RXTE* belong to moderately accreting Seyfert galaxies or very-high accretion rate NLS1s. The population of AGN with lower ($\ll 1$ per cent Eddington) accretion rates which are bright enough to be detected by *RXTE* are often very massive and their corresponding bend time-scales are predicted to be very long. Less massive hard state candidates with low-accretion rates are too faint to be monitored by *RXTE*. Given these facts it is perhaps not surprising that *RXTE* would preferentially observe moderately accreting ‘soft’ state or the occasional ‘very-high’ state AGN.

It is clear that several AGN, using over 10 years worth of *RXTE* data, are excellent soft state candidates. Both NGC 4051 and MCG-6-30-15 exhibit very good power spectra with no evidence of a second bend often associated with a non-soft state. The result that Ark 564 is a very-high state is also robust, so where are the low state AGN? It is most probable that most low accretion rate AGN are too faint to be observed with *RXTE*. Low state BHXRBs exhibit a radio jet feature, which does not appear in known soft states, so it is thought that the radio jet is quenched in the soft state (Fender *et al.*, 1999; Corbel *et al.*, 2000). Radio galaxies would therefore be a prime candidate for low state AGN, but the black holes in these objects are often very massive and require very long time-scale monitoring to reach the frequencies of the predicted second bend feature in the power spectrum. Low state AGN would also be good candidates to look for the elusive LF QPOs, which are not known to be prevalent in soft state BHXRBs.

The jet emission of BHXRBs could provide insight into where to look for low state AGN, but BHXRBs in the soft state might also harbour some surprises. NGC 4051 is often used to represent the definitive AGN soft state, its power spectrum exhibits

a ‘ $1/f$ ’ slope over many decades in frequency and is constructed from 11 years of *RXTE* data (e.g. McHardy *et al.* (2004); Chapter 5). However, radio observations at 5 GHz and 1.6 GHz with the *Very Large Array (VLA)* and the *VLBI* suggest that the radio-quiet Seyfert galaxy NGC 4051 does in fact produce radio emission from a relativistic jet, see Fig. 6.3. It has been suggested that radio jet emission is quenched after the transition to the soft state; moreover, there have not been any detections of even a weak radio jet after the transition to the soft state. The radio observations of a weak jet in NGC 4051 suggest that perhaps soft state BHXRBS do still contain a weak jet feature, undetected by current radio telescopes.

The next generation of radio telescopes should hopefully confirm the existence of a weak radio jet in soft state BHXRBS.

6.7 Physical interpretation of the bend-frequency

Much of this work has focused on the characteristic time-scales of variability produced by AGN. In particular, the bend frequencies found by fitting a single-bend power-law model to a variety of AGN X-ray PSDs were used to determine if there existed a correlation between the bend time-scale, black hole mass and accretion rate. However, a bend or break in the power spectrum of AGN data is an empirical result and the single-bend power-law is merely a way of characterising this observational feature. Thus the obvious question is: does the bend time-scale correspond to any physical time-scale? Assuming the variability originates in the accretion flow, the cut-off of variability power would be associated with the shortest time-scale available to the variability process. As all characteristic time-scales in the accretion disk grow with radius, the shortest possible time-scale for any process will correspond to the innermost radius. For soft state sources we expect this radius to correspond to the last stable orbit. Assuming a standard geometrically thin disk, I consider some physical time-scales of accreting black hole systems (Treves *et al.*, 1988) below.

The most obvious choice of characteristic time-scale might be the light-crossing time of the X-ray emitting region, as fluctuations on shorter time-scales will be diluted, producing a break in the PSD. However, the light-crossing time-scale is far too short to correspond to the bend frequencies measured in this work and an extraordinarily large X-ray emitting region would be required to make this a viable explanation (e.g. the light-crossing time-scale of the innermost stable orbit of a $10^7 M_{\odot}$ black hole is ~ 300 seconds).

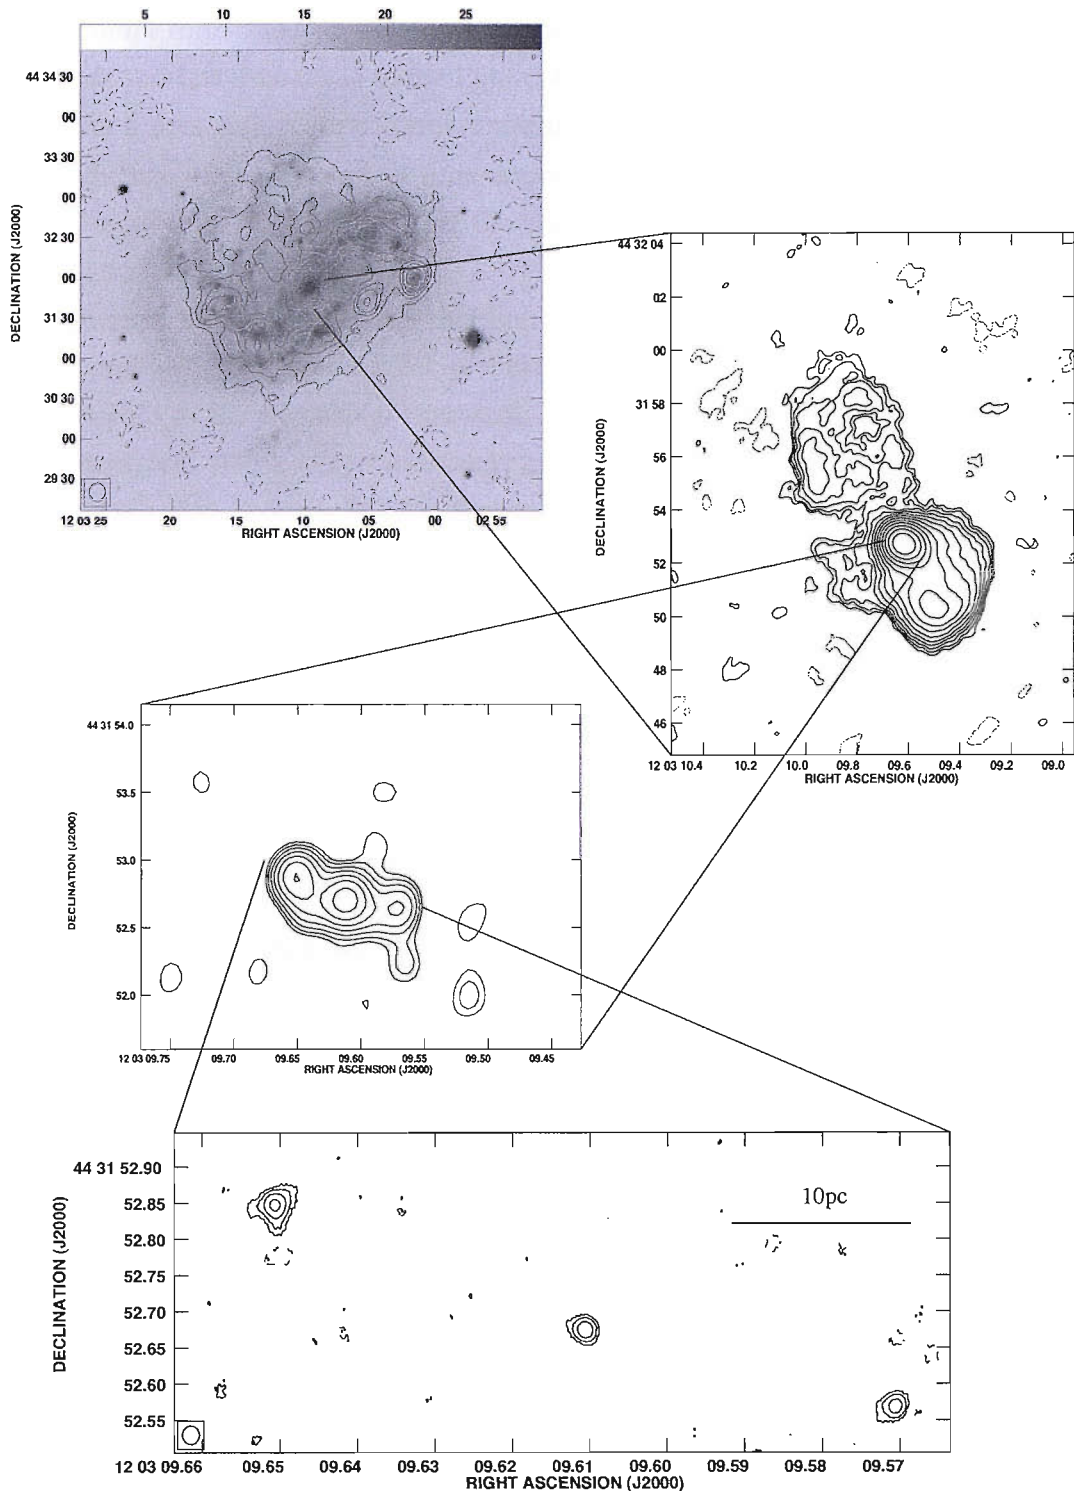


Figure 6.3: The first plot in the top left corner, shows an optical image of NGC 4051 covered by a radio contour map taken at 5 GHz with the VLA. The two subsequent images show higher resolution maps of interesting regions within the radio contours, with each map zooming in closer to the central AGN. Both these maps are taken at 5 GHz with the VLA. The final bottom image is a radio map taken with the VLBI and illustrates a central radio core with two radio hotspots either side. The highly collimated nature of the core and hotspots is a tell-tale sign of a radio jet. Source: Ian McHardy.

The orbital time-scale (also known as the Keplerian or dynamical time-scale) is the next shortest time-scale to the light-crossing time-scale. The orbital time-scale predicts a reasonably sized X-ray emitting region if the bend time-scale were a consequence of the dynamical time-scale (e.g. for a black hole mass of $10^7 M_{\odot}$ and a bend time-scale of 10^{-5} Hz the emitting region would extend out to ~ 25 gravitational radii).

The thermal time-scale denotes the time-scale for a disk to reach thermal equilibrium, provided the accretion disk is radiatively efficient. The thermal time-scale is comparable to the dynamical time-scale if the viscosity parameter ², α (Shakura and Syunyaev, 1973), is approximately maximal (unity). For smaller values of viscosity (in a thin disk model), the thermal time-scale can be much longer than the dynamical time-scale.

The slowest time-scale I will consider is the viscous time-scale, which is the time-scale on which matter diffuses through the disk under the effect of viscous forces. The viscous time-scale is not only dependent on the viscosity parameter, but also the square of the scale-height of the disk (H/R) for a standard thin disk model. The dependence of the viscous time-scale on the scale-height is important given the results of Chapter 5. In Chapter 5 I determined that the observed bend time-scale scales with mass divided by accretion rate over a wide range of masses and accretion rates. As the accretion rate increases H/R is expected to increase (e.g. as the disk becomes radiation pressure dominated) and the corresponding viscous time-scale will decrease. Most of the AGN analysed in this work are consistent with soft state interpretations, which implies that the radiatively efficient accretion disk in these sources extends to the innermost stable orbit around the central black hole. Thus the dynamical time-scale of the innermost stable orbit is not able to vary with accretion rate unless the system goes through a state transition, but the scale-height is permitted to change with accretion rate. For a population of soft state AGNs we would expect to observe a time-scale that is dependent on mass and accretion rate *if* the observed time-scale is related to the viscous time-scale and the scale-height increases with accretion rate.

Done *et al.* (2007) suggest that the viscous time-scale at the innermost stable orbit acts as a filter on high frequencies and can produce the variability properties observed. Arévalo and Uttley (2006) propose a model to explain the spectral-timing properties of the X-ray variability of accreting black holes based on fluctuating accretion disk models of Lyubarskii (1997) and Kotov *et al.* (2001). Arévalo and

²A measure of a fluid's resistance to deformation under stress.

Uttley (2006) suggest that the characteristic time-scales of variability observed in AGN is the viscous time-scale for a geometrically thick disk and is the result of an inward-propagating mass accretion fluctuations. Their model further endorses the possibility of the bend time-scale being related to the viscous time-scale and the aforesaid model can reproduce the energy-dependence of the power spectrum. Note that if the observed time-scale is the viscous time-scale then the accretion flow must be close to geometrically thick (otherwise the time-scale will be too long); moreover the accretion flow need not be an optically thick disk, an accreting optically thin flow such as a corona or ADAF would explain the X-ray variability.

6.8 A brief introduction to future missions in high-energy astrophysics

In this section I will briefly introduce some of the current and future missions in high-energy astrophysics that could potentially make important contributions to AGN timing. In the next few years it is likely that *RXTE* will no longer be a viable X-ray telescope and our X-ray monitoring of AGN will need to be re-evaluated.

6.8.1 *Swift*

The *Swift* Gamma-Ray Burst Mission (Gehrels *et al.*, 2004) was launched on 20 November 2004 and for the last few years it has been providing important information regarding the origin and nature of gamma-ray bursts (GRBs). *Swift* was primarily designed to carry out observations of GRBs using an on-board X-ray telescope (XRT) and UV/optical monitor. The XRT has an energy range of 0.2–10 keV, which is ideal for continuing our current observing campaigns, but it could also detect lower luminosity AGN than currently detected with *RXTE* due to increased sensitivity at softer energies. However, the priority of the *Swift* mission is GRBs, and as a result *Swift* is not designed to carry out AGN monitoring and any long-term monitoring campaigns with *Swift* would likely be interrupted by GRB detections. Nonetheless, long-term monitoring would be possible and would be extremely useful in constraining variability on very long time-scales.

Swift could be useful for long-term monitoring of bright AGN in the 15 keV –150 keV energy range with the Burst Alert Telescope (BAT). The BAT has a large field of view which would be ideal for monitoring bright AGN even when the GRB trigger instructs the telescope to point the XRT at a GRB afterglow. The wide field of view

of the BAT means that long time-scale monitoring of bright AGN is possible, but the energy band would be incompatible with those obtained from *RXTE* and *XMM-Newton*. *Swift* BAT light curves could be interesting, since there is little information of how AGN power spectra should behave at such high energies.

6.8.2 *Astrosat*

Astrosat (Agrawal, 2006) is an Indian led X-ray mission, which is due for launch late 2008. The primary mission objective of *Astrosat* is X-ray timing and it is essentially very similar to *RXTE*. The main instruments of interest to this work are the Large-Area Xenon-filled Proportional Counters (LAXPC) and the Soft-X-ray imaging telescope (SXT). The field of view, effective area and sensitivity of the *RXTE* PCA and *Astrosat* LAXPC are similar. The SXT uses an X-ray CCD to image the source between 0.3 – 10 keV, but the real power for AGN timing is its rapid slew and timing capabilities - the same reasons that made *RXTE* ideal for the task. The SXT would be sensitive to fainter AGN than *RXTE*, which would allow *Astrosat* to begin to probe the low-luminosity population of AGN.

Astrosat is a perfect choice to continue long time-scale monitoring of *RXTE* sources. Longer light curves (i.e. continuation of *RXTE* monitoring campaigns) would provide a two-fold benefit

- Many AGN could be examined for the possibility of a second bend feature in their power spectrum, even for the more massive AGN.
- Features in the power spectrum would be better determined, such as the bend-frequency and weak LF QPOs, since more cycles of bend-frequency (or LF QPO) would lead to a more robust detection.

The resultant light curves of *Astrosat* could also be used together with *XMM-Newton* light curves to provide broadband coverage of the power spectrum.

6.8.3 *The Gamma-ray Large Area Space Telescope (GLAST)*

In early 2008 NASA plans to launch the Gamma-ray Large Area Space Telescope (*GLAST*) (Gehrels and Michelson, 1999). *GLAST* will have a prodigious energy range of 8 keV – 300 GeV and constitute one of the most sensitive telescopes ever designed for measuring the gamma-ray sky.

Many AGN and blazars in particular are known to be powerful emitters of high-energy gamma-rays and exhibit significant variability. The relativistic jet feature in

blazar objects is also a source of very-high energy photons. The energy range of *GLAST* has been fine-tuned to detect the majority of gamma-ray emission of these active galaxies.

The Large Area Telescope (LAT) has an energy range of 20 MeV – 300 GeV and is the main instrument of the *GLAST* telescope. LAT has an effective area of $\sim 8 \text{ m}^2$ at 1–10 GeV and a wide field of view of ~ 2 steradians. The other important instrument for GRB research is the *GLAST* Burst Monitor (GBM), which has an energy range of 8 keV – 25 MeV and an enormous field of view of ~ 8 steradians, suitable for monitoring the gamma-ray sky for GRBs. *GLAST* will be placed in a low-Earth orbit to minimise the effects of a charged particle background found at high altitude orbits. At 90 minutes per orbit, *GLAST* will be able to view the entire sky in two orbits.

Given the wide field of view of *GLAST* and low-Earth orbit, many active galaxies can be regularly observed. *GLAST* light curves will be very well sampled and have excellent signal-to-noise for long time-scale monitoring. *GLAST* would provide an excellent opportunity to undertake a regular monitoring campaign of bright blazars at high energies.

6.8.4 X-Ray Evolving Universe Spectrometer Mission (*XEUS*)

The X-Ray Evolving Universe Spectrometer Mission (*XEUS*) (Parmar *et al.*, 2006) is an ESA candidate large Cosmic Vision mission with an estimated launch in 2018. The main science goals of *XEUS* are to probe clusters of galaxies, study galaxy formation and black holes at redshifts of ~ 10 , determine the nature of space-time near massive black holes in the centre of AGN, and examine the extreme matter physics of highly collapsed stars.

The energy range of *XEUS* of 0.1 keV – 15 keV is similar to that of *XMM-Newton*, but the effective area of *XEUS* is much bigger with 5 m^2 at 1 keV. *XEUS* will also be situated at Earth-Sun L2³, which is ideal for long observations up to ~ 1 Ms in duration.

It is not practical to use *XEUS* as an X-ray monitoring satellite for the objects found in this work, *XEUS* will only be used to provide long continuous look observations of faint objects. *XEUS* should be able to detect the X-ray emission of faint high redshift AGN and the more local very low accretion rate AGN. *XEUS* X-ray power

³A Lagrangian point behind the Earth relative to the Sun, where an object affected only by gravity can remain stationary i.e. the gravitational force due to the Earth and Sun are balanced by the centripetal force of the orbiting satellite.

spectra of faint obscured AGN could also be used to provide a mass estimate where more conventional methods might fail by using the $T_B - M_{\text{BH}} - \dot{m}_E$ correlation presented in Chapter 5. *XEUS* will be able to probe the high-frequency power spectrum of bright AGN and possibly detect a HF QPO in sources where a QPO might be expected e.g. Ark 564. Vaughan and Uttley (2005) have shown that a 120-ks observation with the Wide-Field Imager (WFI) camera on-board *XEUS*⁴ would be almost as good as 10 orbits or 1.2 Ms of exposure from *XMM-Newton* at detecting a possible QPO feature in the high-frequency power spectrum of NGC 4051.

Naturally, PSRESP would not require much (if any) modification to undertake a power spectral survey of *XEUS* targets. *XEUS* will provide important insights into the evasive nature of low-luminosity, low-mass AGN.

6.9 Final remarks

Physical interpretations and models of AGN power spectra will need to be further developed to probe the origin of the observed characteristic bend time-scale. Nonetheless, the work presented herein has further refined the process of determining AGN power spectra and provided a framework by which the current good sample of AGN, as observed by *RXTE*, have been reduced and analysed. I have further strengthened the relationship between AGN and BHXRBs by confirming NGC 3783 as a soft state AGN and identifying Ark 564 as the first very-high state AGN. Moreover, I have undertaken the first consistent large-scale survey of AGN power spectra and further validated the proposed correlation between $T_B - M_{\text{BH}} - \dot{m}_E$ at all mass scales.

⁴Vaughan and Uttley (2005) based their analysis on the information presented in Parmar *et al.* (2001) and estimate the WFI count rate using the publically available *XEUS* response matrices.

APPENDICES

Facts are stubborn things, but statistics are more pliable.

MARK TWAIN (1835 – 1910)



RXTE monitoring X-ray light curves

The following *RXTE* long time-scale light curves have been used.


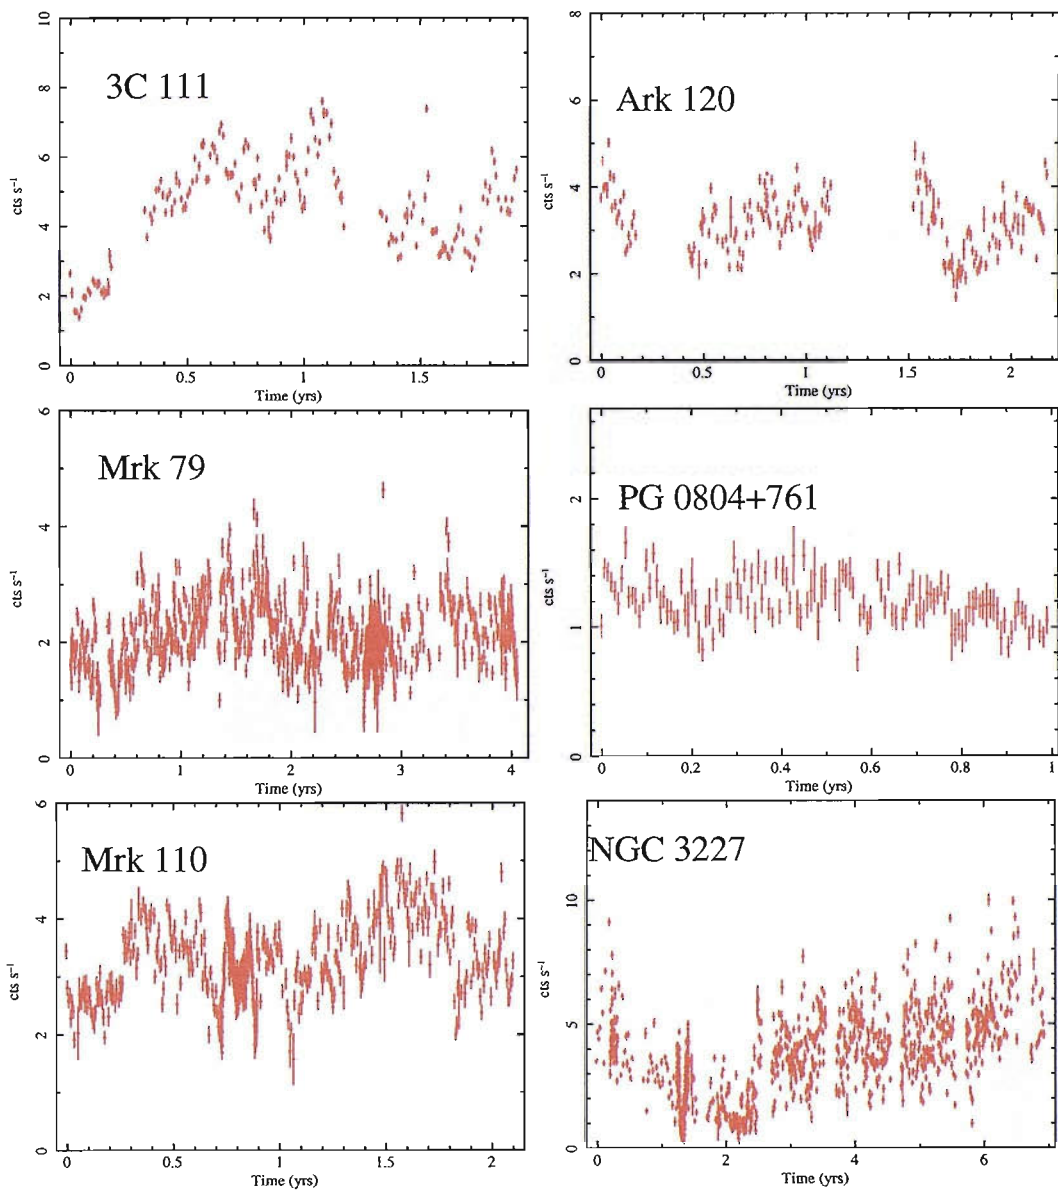






Figure A.1: The long time-scale *RXTE* 2–10 keV X-ray light curves. Count rates are normalised to PCU^{-1} .











Man will occasionally stumble over the truth,
but most of the time he will pick himself up and continue
on.

SIR WINSTON CHURCHILL (1874 – 1965)

B

RXTE medium/short time-scale X-ray light curves

The following *RXTE* medium/short time-scale light curves have been used.


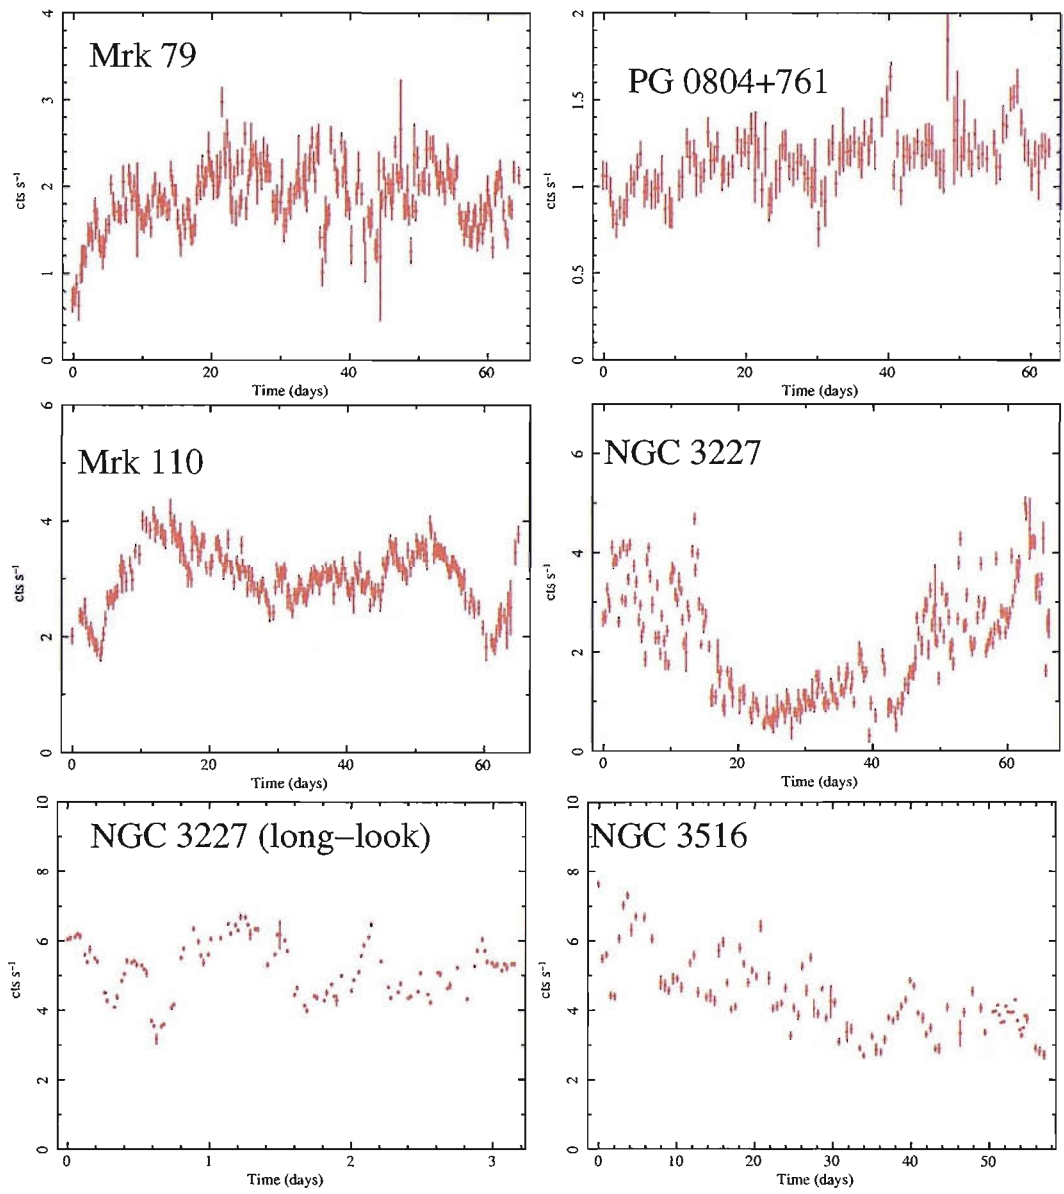





Figure B.1: The medium/short time-scale *RXTE* 2–10 keV X-ray light curves. Count rates are normalised to PCU^{-1} .









The truth is rarely pure and never simple.

OSCAR WILDE (1854 – 1900)



XMM-Newton, ASCA and EXOSAT X-ray light curves

The following short time-scale light curves have been used. The light curves are *XMM-Newton* 4–10 keV unless otherwise stated.




Figure C.1: The short time-scale X-ray light curves. The energy range of Ton S 180 is 0.3–10 keV.




Figure C.1: The short time-scale X-ray light curves. 3C 120 and Mrk 79 are in the energy range 0.3–10 keV.




Figure C.1: The short time-scale X-ray light curves. Mrk 110 and NGC 4051 are in the energy 0.3–10 keV.

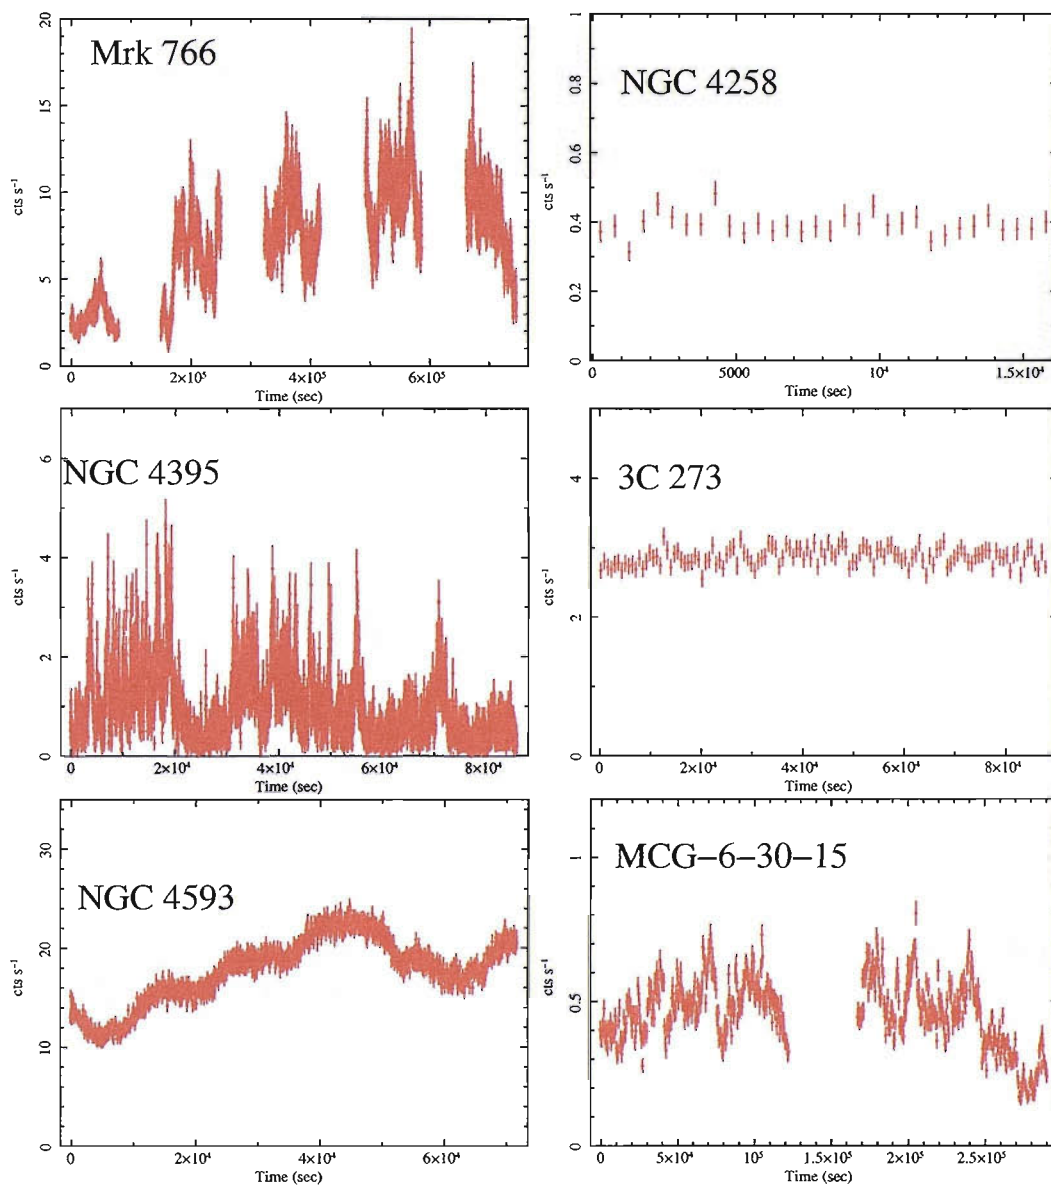


Figure C.1: The short time-scale X-ray light curves. Mrk 766, NGC 4395 and NGC 4593 are in the energy range 0.3–10 keV.

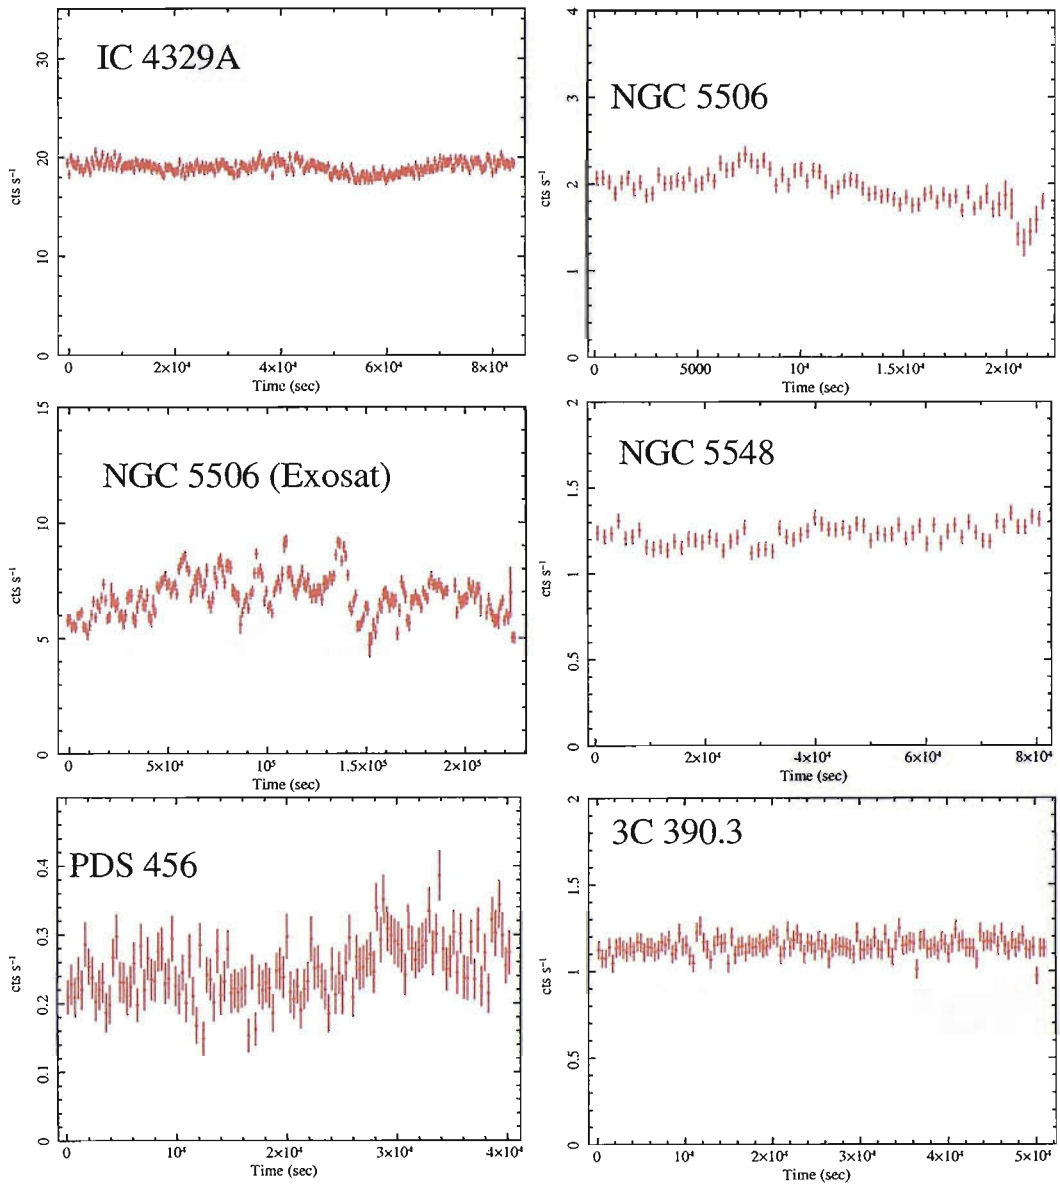


Figure C.1: The short time-scale X-ray light curves. IC 4329A is in the energy range 0.3–10 keV.




Figure C.1: The short time-scale X-ray light curves.

BIBLIOGRAPHY

- Abramowicz M. A., Czerny B., Lasota J. P., Szuszkiewicz E., 1988, *Astrophys. J.*, 332, 646
- Agrawal P. C., 2006, *Advances in Space Research*, 38, 2989
- Aharonian F., Akhperjanian A. G., Bazer-Bachi A. R., Beilicke M., Benbow W., Berge D., Bernlöhner K., Boisson C., Bolz O., Borrel V., Braun I., Breitling F., Brown A. M., Chadwick P. M., Chounet L.-M., Cornils R., Costamante L., Degrange B., Dickinson H. J., Djannati-Ataï A., O’C. Drury L., Dubus G., Emmanoulopoulos D., Espigat P., Feinstein F., Fontaine G., Fuchs Y., Funk S., Gallant Y. A., Giebels B., Gillissen S., Glicenstein J. F., Goret P., Hadjichristidis C., Hauser M., Heinzlmann G., Henri G., Hermann G., Hinton J. A., Hofmann W., Holleran M., Horns D., Jacholkowska A., de Jager O. C., Khélifi B., Komin N., Konopelko A., Latham I. J., Le Gallou R., Lemièrre A., Lemoine-Goumard M., Leroy N., Lohse T., Martin J. M., Martineau-Huynh O., Marcowith A., Master-son C., McComb T. J. L., de Naurois M., Nolan S. J., Noutsos A., Orford K. J., Osborne J. L., Ouchrif M., Panter M., Pelletier G., Pita S., Pühlhofer G., Punch M., Raubenheimer B. C., Raue M., Raux J., Rayner S. M., Reimer A., Reimer O., Ripken J., Rob L., Rolland L., Rowell G., Sahakian V., Saugé L., Schlenker S., Schlickeiser R., Schuster C., Schwanke U., Siewert M., Sol H., Spangler D., Steenkamp R., Stegmann C., Tavernet J.-P., Terrier R., Théoret C. G., Tluczykont M., Vasileiadis G., Venter C., Vincent P., Völk H. J., Wagner S. J., 2005, *Astron. Astrophys.*, 441, 465
- Antonucci R., 1993, *Ann. Rev. Astron. Astrophys.*, 31, 473
- Antonucci R. R. J., Miller J. S., 1985, *Astrophys. J.*, 297, 621
- Arévalo P., Papadakis I. E., Uttley P., McHardy I. M., Brinkmann W., 2006, *Mon. Not. R. Astron. Soc.*, 372, 401
- Arévalo P., Uttley P., 2006, *Mon. Not. R. Astron. Soc.*, 367, 801
- Axelsson M., Borgonovo L., Larsson S., 2005, *Astron. Astrophys.*, 438, 999
- Axelsson M., Borgonovo L., Larsson S., 2006, *Astron. Astrophys.*, 452, 975
- Belloni T., Hasinger G., 1990, *Astron. Astrophys.*, 227, L33
- Belloni T., Homan J., Casella P., van der Klis M., Nespoli E., Lewin W. H. G., Miller J. M., Méndez M., 2005, *Astron. Astrophys.*, 440, 207
- Blumenthal G. R., Gould R. J., 1970, *Reviews of Modern Physics*, 42, 237
- Botte V., Ciroi S., di Mille F., Rafanelli P., Romano A., 2005, *Mon. Not. R. Astron. Soc.*, 356, 789
- Botte V., Ciroi S., Rafanelli P., Di Mille F., 2004, *Astron. J.*, 127, 3168

- Brillinger D. R., Krishnaiah P. R., 1983, *Time series in the frequency domain*, Handbook of Statistics, Amsterdam: North-Holland, 1983, edited by Brillinger, D.R.; Krishnaiah, P.R.
- Cerny V., 1985, *Journal of Optimization Theory and Applications*, 45, 41
- Chevallier L., Collin S., Dumont A.-M., Czerny B., Mouchet M., Gonçalves A. C., Goosmann R., 2006, *Astron. Astrophys.*, 449, 493
- Christopoulou P. E., Holloway A. J., Steffen W., Mundell C. G., Thean A. H. C., Goudis C. D., Meaburn J., Pedlar A., 1997, *Mon. Not. R. Astron. Soc.*, 284, 385
- Churazov E., Gilfanov M., Revnivtsev M., 2001, *Mon. Not. R. Astron. Soc.*, 321, 759
- Corbel S., Fender R. P., Tzioumis A. K., Nowak M., McIntyre V., Durouchoux P., Sood R., 2000, *Astron. Astrophys.*, 359, 251
- Cropper M., Soria R., Mushotzky R. F., Wu K., Markwardt C. B., Pakull M., 2004, *Mon. Not. R. Astron. Soc.*, 349, 39
- Cui W., Heindl W. A., Rothschild R. E., Zhang S. N., Jahoda K., Focke W., 1997, *Astrophys. J. Lett.*, 474, L57+
- Czerny B., Schwarzenberg-Czerny A., Loska Z., 1999, *Mon. Not. R. Astron. Soc.*, 303, 148
- Davies R. I., Thomas J., Genzel R., Mueller Sánchez F., Tacconi L. J., Sternberg A., Eisenhauer F., Abuter R., Saglia R., Bender R., 2006, *Astrophys. J.*, 646, 754
- Deeming T. J., 1975, *ApSS*, 36, 137
- di Matteo T., 1998, *Mon. Not. R. Astron. Soc.*, 299, L15+
- Done C., Gierliński M., 2005, *Mon. Not. R. Astron. Soc.*, 364, 208
- Done C., Gierliński M., Kubota A., 2007, *Astron. Astrophys. Rev.*, pp. 3–+
- Done C., Madejski G. M., Mushotzky R. F., Turner T. J., Koyama K., Kunieda H., 1992, *Astrophys. J.*, 400, 138
- Edelson R., Nandra K., 1999, *Astrophys. J.*, 514, 682
- Edelson R., Turner T. J., Pounds K., Vaughan S., Markowitz A., Marshall H., Dobie P., Warwick R., 2002, *Astrophys. J.*, 568, 610
- Edelson R. A., Krolik J. H., Pike G. F., 1990, *Astrophys. J.*, 359, 86
- Elvis M., Wilkes B. J., McDowell J. C., Green R. F., Bechtold J., Willner S. P., Oey M. S., Polonski E., Cutri R., 1994, *Astrophys. J. Suppl. Ser.*, 95, 1
- Fabian A. C., Rees M. J., 1995, *Mon. Not. R. Astron. Soc.*, 277, L55
- Falcke H., Kording E., Markoff S., 2004, *Astron. Astrophys.*, 414, 895
- Fanaroff B. L., Riley J. M., 1974, *Mon. Not. R. Astron. Soc.*, 167, 31P
- Fender R., Corbel S., Tzioumis T., McIntyre V., Campbell-Wilson D., Nowak M., Sood R., Hunstead R., Harmon A., Durouchoux P., Heindl W., 1999, *Astrophys.*

- J. Lett., 519, L165
- Fender R. P., 2001, *Mon. Not. R. Astron. Soc.*, 322, 31
- Fender R. P., Belloni T. M., Gallo E., 2004, *Mon. Not. R. Astron. Soc.*, 355, 1105
- Filippenko A. V., Ho L. C., 2003, *Astrophys. J. Lett.*, 588, L13
- Gehrels N., Chincarini G., Giommi P., Mason K. O., Nousek J. A., Wells A. A., White N. E., Barthelmy S. D., Burrows D. N., Cominsky L. R., Hurley K. C., Marshall F. E., Mészáros P., Roming P. W. A., Angelini L., Barbier L. M., Belloni T., Campana S., Caraveo P. A., Chester M. M., Citterio O., Cline T. L., Cropper M. S., Cummings J. R., Dean A. J., Feigelson E. D., Fenimore E. E., Frail D. A., Fruchter A. S., Garmire G. P., Gendreau K., Ghisellini G., Greiner J., Hill J. E., Hunsberger S. D., Krimm H. A., Kulkarni S. R., Kumar P., Lebrun F., Lloyd-Ronning N. M., Markwardt C. B., Mattson B. J., Mushotzky R. F., Norris J. P., Osborne J., Paczynski B., Palmer D. M., Park H.-S., Parsons A. M., Paul J., Rees M. J., Reynolds C. S., Rhoads J. E., Sasseen T. P., Schaefer B. E., Short A. T., Smale A. P., Smith I. A., Stella L., Tagliaferri G., Takahashi T., Tashiro M., Townsley L. K., Tueller J., Turner M. J. L., Vietri M., Voges W., Ward M. J., Willingale R., Zerbi F. M., Zhang W. W., 2004, *Astrophys. J.*, 611, 1005
- Gehrels N., Michelson P., 1999, *Astroparticle Physics*, 11, 277
- Gierliński M., Done C., 2004, *Mon. Not. R. Astron. Soc.*, 349, L7
- Goad M. R., Roberts T. P., Reeves J. N., Uttley P., 2006, *Mon. Not. R. Astron. Soc.*, 365, 191
- Green A. R., McHardy I. M., Done C., 1999, *Mon. Not. R. Astron. Soc.*, 305, 309
- Green A. R., McHardy I. M., Lehto H. J., 1993, *Mon. Not. R. Astron. Soc.*, 265, 664
- Gu Q., Melnick J., Fernandes R. C., Kunth D., Terlevich E., Terlevich R., 2006, *Mon. Not. R. Astron. Soc.*, 366, 480
- Haardt F., Maraschi L., 1993, *Astrophys. J.*, 413, 507
- Haardt F., Maraschi L., Ghisellini G., 1997, *Astrophys. J.*, 476, 620
- Hao C. N., Xia X. Y., Mao S., Wu H., Deng Z. G., 2005, *Astrophys. J.*, 625, 78
- Heckman T. M., 1980, *Astron. Astrophys.*, 87, 152
- Herrnstein J. R., Moran J. M., Greenhill L. J., Diamond P. J., Inoue M., Nakai N., Miyoshi M., Henkel C., Riess A., 1999, *Nature*, 400, 539
- Homan J., Belloni T., 2005, *ApSS*, 300, 107
- Homan J., Wijnands R., van der Klis M., Belloni T., van Paradijs J., Klein-Wolt M., Fender R., Méndez M., 2001, *Astrophys. J. Suppl. Ser.*, 132, 377
- Hong X. Y., Jiang D. R., Gurvits L. I., Garrett M. A., Garrington S. T., Schilizzi R. T., Nan R. D., Hirabayashi H., Wang W. H., Nicolson G. D., 2004, *Astron.*

- Astrophys., 417, 887
- Iwasawa K., Fabian A. C., Brandt W. N., Kunieda H., Misaki K., Terashima Y., Reynolds C. S., 1998, *Mon. Not. R. Astron. Soc.*, 295, L20
- Jahoda K., Markwardt C. B., Radeva Y., Rots A. H., Stark M. J., Swank J. H., Strohmayer T. E., Zhang W., 2006, *Astrophys. J. Suppl. Ser.*, 163, 401
- Jenkins G. M., Watts D. G., 1969, *Spectral analysis and its applications*, Holden-Day Series in Time Series Analysis, London: Holden-Day, 1969
- Jester S., 2005, *Astrophys. J.*, 625, 667
- Kirkpatrick S., Gelatt C. D., Vecchi M. P., 1983, *Science*, 220 (4598), 671
- Klein-Wolt M., van der Klis M., 2007, *ArXiv e-prints*, 711
- Van der Klis M., 2004, "Rapid X-ray variability", Book chapter pre-print
- Körding E. G., Fender R. P., Migliari S., 2006a, *Mon. Not. R. Astron. Soc.*, 369, 1451
- Körding E. G., Jester S., Fender R., 2006b, *Mon. Not. R. Astron. Soc.*, 372, 1366
- Körding E. G., Migliari S., Fender R., Belloni T., Knigge C., McHardy I., 2007, *Mon. Not. R. Astron. Soc.*, 380, 301
- Kormendy J., Richstone D., 1995, *Ann. Rev. Astron. Astrophys.*, 33, 581
- Kotov O., Churazov E., Gilfanov M., 2001, *Mon. Not. R. Astron. Soc.*, 327, 799
- Kraemer S. B., Ho L. C., Crenshaw D. M., Shields J. C., Filippenko A. V., 1999, *Astrophys. J.*, 520, 564
- Krolik J. H., Madau P., Zycki P. T., 1994, *Astrophys. J. Lett.*, 420, L57
- Lal D. V., Shastri P., Gabuzda D. C., 2004, *Astron. Astrophys.*, 425, 99
- Lamer G., Uttley P., McHardy I. M., 2003, *Mon. Not. R. Astron. Soc.*, 342, L41
- Lampton M., Margon B., Bowyer S., 1976, *Astrophys. J.*, 208, 177
- Laor A., 1990, *Mon. Not. R. Astron. Soc.*, 246, 369
- Lasota J.-P., Abramowicz M. A., Chen X., Krolik J., Narayan R., Yi I., 1996, *Astrophys. J.*, 462, 142
- Leahy D. A., Darbro W., Elsner R. F., Weisskopf M. C., Kahn S., Sutherland P. G., Grindlay J. E., 1983, *Astrophys. J.*, 266, 160
- Levenberg, K., 1944, *Quart. Appl. Math*, 2, 164
- Lewin W. H. G., van der Klis M., 2006, *Compact stellar X-ray sources*, Compact stellar X-ray sources
- Lightman A. P., White T. R., 1988, *Astrophys. J.*, 335, 57
- Lira P., Lawrence A., O'Brien P., Johnson R. A., Terlevich R., Bannister N., 1999, *Mon. Not. R. Astron. Soc.*, 305, 109
- Liu Y., Li J., Fan J.-H., 2007, *Chinese Journal of Astronomy and Astrophysics*, 7, 380

- Lyubarskii Y. E., 1997, *Mon. Not. R. Astron. Soc.*, 292, 679
- Maccarone T. J., 2003, *Astron. Astrophys.*, 409, 697
- Maccarone T. J., Gallo E., Fender R., 2003, *Mon. Not. R. Astron. Soc.*, 345, L19
- Mahadevan R., 1998, *Nature*, 394, 651
- Markowitz A., 2005, *Astrophys. J.*, 635, 180
- Markowitz A., Edelson R., Vaughan S., Uttley P., George I. M., Griffiths R. E., Kaspi S., Lawrence A., McHardy I., Nandra K., Pounds K., Reeves J., Schurch N., Warwick R., 2003, *Astrophys. J.*, 593, 96
- Marquardt, D., 1963, *SIAM J. Appl. Math.*, 11, 431
- Mason K. O., Breeveld A., Much R., Carter M., Cordova F. A., Cropper M. S., Fordham J., Huckle H., Ho C., Kawakami H., Kennea J., Kennedy T., Mittaz J., Pandel D., Priedhorsky W. C., Sasseen T., Shirey R., Smith P., Vreux J.-M., 2001, *Astron. Astrophys.*, 365, L36
- McClintock J. E., Remillard R. A., 2004, "Black hole binaries", Book chapter pre-print
- McHardy I., 1988, *Memorie della Societa Astronomica Italiana*, 59, 239
- McHardy I., Czerny B., 1987, *Nature*, 325, 696
- McHardy I. M., Arevalo P., Uttley P., Papadakis I. E., Summons D. P., Brinkmann W., Page M. J., 2007, *ArXiv e-prints*, 709
- McHardy I. M., Gunn K. F., Uttley P., Goad M. R., 2005, *Mon. Not. R. Astron. Soc.*, 359, 1469
- McHardy I. M., Körding E., Knigge C., Uttley P., Fender R. P., 2006, *Nature*, 444, 730
- McHardy I. M., Marscher A. P., Gear W. K., Muxlow T., Lehto H. J., Abraham R. G., 1990, *Mon. Not. R. Astron. Soc.*, 246, 305
- McHardy I. M., Papadakis I. E., Uttley P., Page M. J., Mason K. O., 2004, *Mon. Not. R. Astron. Soc.*, 348, 783
- Merloni A., Heinz S., di Matteo T., 2003, *Mon. Not. R. Astron. Soc.*, 345, 1057
- Miller C. J., Nichol R. C., Gómez P. L., Hopkins A. M., Bernardi M., 2003, *Astrophys. J.*, 597, 142
- Miller-Jones J. C. A., McCormick D. G., Fender R. P., Spencer R. E., Muxlow T. W. B., Pooley G. G., 2005, *Mon. Not. R. Astron. Soc.*, 363, 867
- Miniutti G., Fabian A. C., 2004, *Mon. Not. R. Astron. Soc.*, 349, 1435
- Miniutti G., Fabian A. C., Goyder R., Lasenby A. N., 2003, *Mon. Not. R. Astron. Soc.*, 344, L22
- Mushotzky R. F., Fabian A. C., Iwasawa K., Kunieda H., Matsuoka M., Nandra K., Tanaka Y., 1995, *Mon. Not. R. Astron. Soc.*, 272, L9

- Nagar N. M., Falcke H., Wilson A. S., 2005, *Astron. Astrophys.*, 435, 521
- Nagar N. M., Oliva E., Marconi A., Maiolino R., 2002, *Astron. Astrophys.*, 391, L21
- Nandra K., 2006, *Mon. Not. R. Astron. Soc.*, 368, L62
- Nandra K., Pounds K. A., 1994, *Mon. Not. R. Astron. Soc.*, 268, 405
- Narayan R., Yi I., 1994, *Astrophys. J. Lett.*, 428, L13
- Nelder J. A., Mead R., 1965, *Computer Journal*, 7, 308
- Nelson C. H., Whittle M., 1995, *Astrophys. J. Suppl. Ser.*, 99, 67
- Nowak M. A., 2000, *Mon. Not. R. Astron. Soc.*, 318, 361
- Nowak M. A., Vaughan B. A., Wilms J., Dove J. B., Begelman M. C., 1999, *Astrophys. J.*, 510, 874
- Onken C. A., Valluri M., Peterson B. M., Pogge R. W., Bentz M. C., Ferrarese L., Vestergaard M., Crenshaw D. M., Sergeev S. G., McHardy I. M., Merritt D., Bower G. A., Heckman T. M., Wandel A., 2007, *ArXiv e-prints*, 708
- Oppenheim A. V., Schafer R. W., 1975, *Digital Signal Processing*
- Paczynsky B., Wiita P. J., 1980, *Astron. Astrophys.*, 88, 23
- Papadakis I. E., Brinkmann W., Negoro H., Gliozzi M., 2002, *Astron. Astrophys.*, 382, L1
- Papadakis I. E., Brinkmann W., Page M. J., McHardy I., Uttley P., 2007, *Astron. Astrophys.*, 461, 931
- Papadakis I. E., Lawrence A., 1993a, *Mon. Not. R. Astron. Soc.*, 261, 612
- Papadakis I. E., Lawrence A., 1993b, *Nature*, 361, 233
- Papadakis I. E., McHardy I. M., 1995, *Mon. Not. R. Astron. Soc.*, 273, 923
- Parmar A. N., Arnaud M., Barcons X., Bleeker J., Hasinger G., Kunieda H., Palumbo G., Takahashi T., Turner M., de Korte P., Willingale R., Rando N., Lyngvi A., Gondoin P., Lumb D., Bavdaz M., Verhoeve P., 2006, in *Space Telescopes and Instrumentation II: Ultraviolet to Gamma Ray*. Edited by Turner, Martin J. L.; Hasinger, Günther. *Proceedings of the SPIE, Volume 6266*, pp. 62661R (2006)., vol. 6266 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*
- Parmar A. N., Peacock T., Bavdaz M., Hasinger G., Arnaud M., Barcons X., Barret D., Blanchard A., Böhringer H., Cappi M., Comastri A., Courvoisier T., Fabian A. C., Griffiths R., Malaguti P., Mason K. O., Ohashi T., Paerels F., Piro L., Schmitt J., van der Klis M., Ward M., 2001, *X-ray Astronomy: Stellar Endpoints, AGN, and the Diffuse X-ray Background*, 599, 842
- Peterson B. M., Ferrarese L., Gilbert K. M., Kaspi S., Malkan M. A., Maoz D., Merritt D., Netzer H., Onken C. A., Pogge R. W., Vestergaard M., Wandel A.,

- 2004, *Astrophys. J.*, 613, 682
- Ponti G., Miniutti G., Cappi M., Maraschi L., Fabian A. C., Iwasawa K., 2006, *Mon. Not. R. Astron. Soc.*, 368, 903
- Pottschmidt K., Wilms J., Nowak M. A., Pooley G. G., Gleissner T., Heindl W. A., Smith D. M., Remillard R., Staubert R., 2003, *Astron. Astrophys.*, 407, 1039
- Pounds K., Edelson R., Markowitz A., Vaughan S., 2001, *Astrophys. J. Lett.*, 550, L15
- Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992, *Numerical Recipes*, 2nd edn., Cambridge University Press
- Priestley M. B., 1981, *Spectral Analysis and Time Series*, Academic, London
- Rees M. J., Begelman M. C., Blandford R. D., Phinney E. S., 1982, *Nature*, 295, 17
- Reeves J. N., O'Brien P. T., Vaughan S., Law-Green D., Ward M., Simpson C., Pounds K. A., Edelson R., 2000, *Mon. Not. R. Astron. Soc.*, 312, L17
- Reig P., Papadakis I., Kylafis N. D., 2002, *Astron. Astrophys.*, 383, 202
- Remillard R., Munro M., McClintock J. E., Orosz J., 2003, in Durouchoux P., Fuchs Y., Rodriguez J. (eds.), *New Views on MICROQUASARS, the Fourth Microquasars Workshop, Institut d'Etudes Scientifiques de Cargèse, Corsica, France, May 27 - June 1, 2002. Edited by Ph. Durouchoux, Y. Fuchs, and J. Rodriguez. Published by the Center for Space Physics: Kolkata (India), p. 49.*, pp. 49–+
- Remillard R. A., McClintock J. E., 2006, *Ann. Rev. Astron. Astrophys.*, 44, 49
- Revnivtsev M., Gilfanov M., Churazov E., 2000, *Astron. Astrophys.*, 363, 1013
- Reynolds C. S., 1997, *Mon. Not. R. Astron. Soc.*, 286, 513
- Reynolds C. S., Ward M. J., Fabian A. C., Celotti A., 1997, *Mon. Not. R. Astron. Soc.*, 291, 403
- Rodriguez-Pascual P. M., Alloin D., Clavel J., Crenshaw D. M., Horne K., Kriss G. A., Krolik J. H., Malkan M. A., Netzer H., O'Brien P. T., Peterson B. M., Reichert G. A., Wamsteker W., Alexander T., Barr P., Blandford R. D., Bregman J. N., Carone T. E., Clements S., Courvoisier T.-J., De Robertis M. M., Dietrich M., Dottori H., Edelson R. A., Filippenko A. V., Gaskell C. M., Huchra J. P., Hutchings J. B., Kollatschny W., Koratkar A. P., Korista K. T., Laor A., MacAlpine G. M., Martin P. G., Maoz D., McCollum B., Morris S. L., Perola G. C., Pogge R. W., Ptak R. L., Recondo-Gonzalez M. C., Rodriguez-Espinoza J. M., Rokaki E. L., Santos-Lleo M., Sekiguchi K., Shull J. M., Snijders M. A. J., Sparke L. S., Stirpe G. M., Stoner R. E., Sun W.-H., Wagner S. J., Wanders I., Wilkes B. J., Winge C., Zheng W., 1997, *Astrophys. J. Suppl. Ser.*, 110, 9
- Romano P., Mathur S., Turner T. J., Kraemer S. B., Crenshaw D. M., Peterson B. M., Pogge R. W., Brandt W. N., George I. M., Horne K., Kriss G. A., Netzer

- H., Shemmer O., Wamsteker W., 2004, *Astrophys. J.*, 602, 635
- Rybicki G. B., Lightman A. P., 1986, *Radiative Processes in Astrophysics*, Radiative Processes in Astrophysics, by George B. Rybicki, Alan P. Lightman, pp. 400. ISBN 0-471-82759-2. Wiley-VCH, June 1986.
- Salpeter E. E., 1964, *Astrophys. J.*, 140, 796
- Schmitt H. R., Ulvestad J. S., Antonucci R. R. J., Kinney A. L., 2001, *Astrophys. J. Suppl. Ser.*, 132, 199
- Schödel R., Ott T., Genzel R., Hofmann R., Lehnert M., Eckart A., Mouawad N., Alexander T., Reid M. J., Lenzen R., Hartung M., Lacombe F., Rouan D., Gendron E., Rousset G., Lagrange A.-M., Brandner W., Ageorges N., Lidman C., Moorwood A. F. M., Spyromilio J., Hubin N., Menten K. M., 2002, *Nature*, 419, 694
- Shakura N. I., Syunyaev R. A., 1973, *Astron. Astrophys.*, 24, 337
- Shannon C. E., 1949, *Proc. Institute of Radio Engineers*, 37 (1), 10
- Silk J., Rees M. J., 1998, *Astron. Astrophys.*, 331, L1
- Sobolewska M. A., Done C., 2007, *Mon. Not. R. Astron. Soc.*, 374, 150
- Starling R. L. C., Page M. J., Branduardi-Raymont G., Breeveld A. A., Soria R., Wu K., 2005, *Mon. Not. R. Astron. Soc.*, 356, 727
- Strohmayer T. E., Mushotzky R. F., Winter L., Soria R., Uttley P., Cropper M., 2007, *Astrophys. J.*, 660, 580
- Strohmayer T. E., Zhang W., Swank J. H., Smale A., Titarchuk L., Day C., Lee U., 1996, *Astrophys. J. Lett.*, 469, L9+
- Strüder L., Briel U., Dennerl K., Hartmann R., Kendziorra E., Meidinger N., Pfeffermann E., Reppin C., Aschenbach B., Bornemann W., Bräuninger H., Burkert W., Elender M., Freyberg M., Haberl F., Hartner G., Heuschmann F., Hippmann H., Kastelic E., Kemmer S., Kettenring G., Kink W., Krause N., Müller S., Oppitz A., Pietsch W., Popp M., Predehl P., Read A., Stephan K. H., Stötter D., Trümper J., Holl P., Kemmer J., Soltau H., Stötter R., Weber U., Weichert U., von Zanthier C., Carathanassis D., Lutz G., Richter R. H., Solc P., Böttcher H., Kuster M., Staubert R., Abbey A., Holland A., Turner M., Balasini M., Bignami G. F., La Palombara N., Villa G., Buttler W., Gianini F., Lainé R., Lumb D., Dhez P., 2001, *Astron. Astrophys.*, 365, L18
- Summons D. P., Arévalo P., McHardy I. M., Uttley P., Bhaskar A., 2007, *Mon. Not. R. Astron. Soc.*, 378, 649
- Sunyaev R. A., Titarchuk L. G., 1980, *Astron. Astrophys.*, 86, 121
- Swank J. H., 1998, *Nuclear Physics B*, 69 (1-3), 12
- Tanaka Y., Boller T., Gallo L., 2005, in Merloni A., Nayakshin S., Sunyaev R. A.

- (eds.), *Growing Black Holes: Accretion in a Cosmological Context*, pp. 290–295
- Tanaka Y., Nandra K., Fabian A. C., Inoue H., Otani C., Dotani T., Hayashida K., Iwasawa K., Kii T., Kunieda H., Makino F., Matsuoka M., 1995, *Nature*, 375, 659
- Timmer J., König M., 1995, *Astron. Astrophys.*, 300, 707
- Treves A., Maraschi L., Abramowicz M., 1988, *Publ. Astron. Soc. Pac.*, 100, 427
- Trudolyubov S. P., 2001, *Astrophys. J.*, 558, 276
- Turner M. J. L., Abbey A., Arnaud M., Balasini M., Barbera M., Belsole E., Bennie P. J., Bernard J. P., Bignami G. F., Boer M., Briel U., Butler I., Cara C., Chabaud C., Cole R., Collura A., Conte M., Cros A., Denby M., Dhez P., Di Coco G., Dowson J., Ferrando P., Ghizzardi S., Gianotti F., Goodall C. V., Gretton L., Griffiths R. G., Hainaut O., Hochedez J. F., Holland A. D., Jourdain E., Kendziorra E., Lagostina A., Laine R., La Palombara N., Lortholary M., Lumb D., Marty P., Molendi S., Pigot C., Poindron E., Pounds K. A., Reeves J. N., Reppin C., Rothenflug R., Salvétat P., Sauvageot J. L., Schmitt D., Sembay S., Short A. D. T., Spragg J., Stephen J., Strüder L., Tiengo A., Trifoglio M., Trümper J., Vercellone S., Vigroux L., Villa G., Ward M. J., Whitehead S., Zonca E., 2001, *Astron. Astrophys.*, 365, L27
- Turner T. J., George I. M., Nandra K., Turcan D., 1999, *Astrophys. J.*, 524, 667
- Urry C. M., Padovani P., 1995, *Publ. Astron. Soc. Pac.*, 107, 803
- Uttley P., McHardy I. M., 2001, *Mon. Not. R. Astron. Soc.*, 323, L26
- Uttley P., McHardy I. M., 2005, *Mon. Not. R. Astron. Soc.*, 363, 586
- Uttley P., McHardy I. M., Papadakis I. E., 2002, *Mon. Not. R. Astron. Soc.*, 332, 231
- Uttley P., McHardy I. M., Vaughan S., 2005, *Mon. Not. R. Astron. Soc.*, 359, 345
- van der Klis M., 1989, in Ogelman H., van den Heuvel E. P. J. (eds.), *Timing Neutron Stars*, vol. 262 of *NATO ASI Series C*, p. p27, Kluwer Academic Publishers, Kluwer, Dordrecht
- van der Klis M., 1997, in *Statistical Challenges in Modern Astronomy II*, pp. 321–+
- Vaughan S., 2005, *Astron. Astrophys.*, 431, 391
- Vaughan S., Edelson R., Warwick R. S., Uttley P., 2003, *Mon. Not. R. Astron. Soc.*, 345, 1271
- Vaughan S., Iwasawa K., Fabian A. C., Hayashida K., 2005, *Mon. Not. R. Astron. Soc.*, 356, 524
- Vaughan S., Uttley P., 2005, *Mon. Not. R. Astron. Soc.*, 362, 235
- Vaughan S., Uttley P., 2006, *Advances in Space Research*, 38, 1405
- Veron-Cetty M. ., Veron P., Joly M., Kollatschny W., 2007, *ArXiv e-prints*, 708

- Weaver K. A., Nousek J., Yaqoob T., Mushotzky R. F., Makino F., Otani C., 1996, *Astrophys. J.*, 458, 160
- Wei D., 2006, *X-ray Power Density Spectra of Black Hole Binaries: A New Dead-time Model for RXTE PCA*, Master's thesis, Massachusetts Institute of Technology
- Wilms J., Nowak M. A., Pottschmidt K., Pooley G. G., Fritz S., 2006, *Astron. Astrophys.*, 447, 245
- Winter L. M., Mushotzky R. F., Reynolds C. S., 2006, *Astrophys. J.*, 649, 730
- Woo J.-H., Urry C. M., 2002, *Astrophys. J.*, 579, 530
- Wu J., Zhou X., Wu X.-B., Liu F.-K., Peng B., Ma J., Wu Z., Jiang Z., Chen J., 2006, *Astron. J.*, 132, 1256
- Yaqoob T., 1998, *Astrophys. J.*, 500, 893
- Zhang W., Giles A. B., Jahoda K., Soong Y., Swank J. H., Morgan E. H., 1993, in Siegmund O. H. (ed.), *Proc. SPIE Vol. 2006, p. 324-333, EUV, X-Ray, and Gamma-Ray Instrumentation for Astronomy IV, Oswald H. Siegmund; Ed.*, pp. 324–333