

University of Southampton

**Timing studies of Seyfert
galaxies with the Rossi X-ray
Timing Explorer**

by Philip Uttley

Submitted for the degree of Doctor of Philosophy

Department of Physics and Astronomy
Faculty of Science

May 2000

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF SCIENCE
PHYSICS AND ASTRONOMY

Doctor of Philosophy

TIMING STUDIES OF SEYFERT GALAXIES WITH THE ROSSI X-RAY TIMING
EXPLORER

by Philip Uttley

Previous studies of the X-ray variability power spectra of active galactic nuclei (AGN) on time-scales of hours to days revealed striking similarities to the timing properties of black hole X-ray binary systems (XRBs) on much shorter time-scales (< 1 s). The unique capabilities of the Rossi X-ray Timing Explorer (RXTE) mission allow us to sample AGN lightcurves over a far greater range of time-scales than before, enabling us to compare the broadband power spectrum of AGN with that of XRBs. Here we present the results of an RXTE program to monitor 4 Seyfert galaxies (NGC 4051, MCG-6-30-15, NGC 5506 and NGC 5548) to measure their power spectra from 10^{-8} Hz to 10^{-2} Hz.

We use RXTE long-look observations to show that the power spectra of our sample are intrinsically non-stationary, in that their RMS variability scales with local mean flux. We show that this relation also applies to the black hole XRB Cygnus X-1 and the accreting millisecond pulsar SAX J1808.4-3658 and is highly linear, implying that it is intrinsic to the red-noise variability process which dominates the X-ray lightcurves of accreting compact objects. The scaling of RMS variability with flux occurs on all measured time-scales, posing problems for models where lightcurves are made from small building blocks, such as conventional shot-noise models. We suggest that the lightcurves are built from the top down, out of large flaring events which break up into self-similar structure on smaller scales.

Discrete sampling of lightcurves causes distortion in the observed power spectrum due to red-noise leak and aliasing effects. We develop a Monte Carlo technique, based on the ‘response method’ of Done et al. (1992), to robustly estimate the shape of the underlying broadband power spectrum. We apply this technique to our data, and data obtained from a separate program to monitor the Seyfert galaxy NGC 3516. We find that the broadband power spectra of MCG-6-30-15, NGC 5506 and NGC 3516 flatten significantly at low frequencies, and that this flattening can be well fitted by a high-frequency break model analogous to the shape of the high-frequency power spectrum in Cygnus X-1. The break frequency of NGC 3516 and the lack of a detectable break in NGC 5548 are consistent with a linear scaling of the break time-scale with their respective black hole masses (as estimated from reverberation mapping). The high break frequencies measured for MCG-6-30-15 and NGC 5506 imply that these are high state AGN, accreting at a high rate and analogous to high state XRBs.

The lightcurve of NGC 4051 is complex, showing normal activity, non-Gaussian variations (which prevent a valid measurement of the power spectrum) and two low states which lasted for 5 and 3 months respectively in early 1998 and 1999. The X-ray energy spectrum in the low state is consistent with reflection of the normal continuum off the distant molecular torus expected from unification models. We show that the X-ray and EUV continua of NGC 4051 are highly correlated. X-ray and EUV observations, together with long-term optical monitoring data, suggest that the low state corresponds to the disappearance of the inner accretion disk, or its transition to a radiatively inefficient mode. On the basis of variability time-scales and intriguing observational evidence, we argue that NGC 4051 is a ‘macro-microquasar’, the AGN analogue of the galactic microquasar GRS 1915+105.

To Mum, Dad and Granddad.

Contents

1	Introduction	1
1.1	Active Galactic Nuclei	2
1.1.1	The central engine: accretion on to a massive black hole	3
1.1.2	AGN unification	5
1.1.3	X-ray spectral properties of AGN	7
1.1.4	X-ray variability properties of AGN	10
1.2	X-ray binary systems and the XRB-AGN connection	11
1.2.1	The power spectra of XRBs	12
1.2.2	The power spectra of AGN	16
1.3	Long-term monitoring of Seyfert galaxies with <i>RXTE</i>	19
1.3.1	<i>RXTE</i>	19
1.3.2	The Seyfert monitoring program	21
1.4	Thesis overview	24

2	Power Spectral Issues	27
2.1	Measuring the power spectrum	27
2.2	Simulating lightcurves	32
2.3	The effects of sampling on power-spectral shape	33
2.3.1	Red-noise leak	36
2.3.2	Aliasing	37
2.3.3	The effects of uneven sampling	42
2.3.4	The effects of binning	43
2.4	Stationarity of the power spectrum	44
2.5	The flux dependence of AGN variability	47
2.6	The flux dependence of XRB variability	49
2.7	Summary	55
3	PSRESP: A robust method for measuring power spectral shape	57
3.1	The response method concept	58
3.2	PSRESP implementation	61
3.3	Lightcurve simulation and computing time	64
3.4	Additional caveats	67
4	The broadband power-spectral shape of Seyfert galaxies	69

4.1	Observations and data reduction	70
4.1.1	MCG-6-30-15 observations	70
4.1.2	NGC 5506 observations	72
4.1.3	NGC 5548 observations	73
4.1.4	NGC 3516 observations	76
4.1.5	Data reduction	77
4.2	A first look at the RMS variability and power spectra	79
4.3	Do the broadband power spectra really flatten?	84
4.4	Fitting models with power-spectral breaks	88
4.4.1	The knee model	89
4.4.2	The high-frequency break model	93
4.4.3	Comparison with other work	96
4.5	The dependence of break frequency on AGN luminosity and black hole mass	100
4.6	The energy dependence of RMS variability and the power spec- trum	103
4.6.1	Energy dependence of long-timescale fractional RMS	103
4.6.2	Dependence of power-spectral shape on energy	107
4.7	Summary	109

5	X-ray, EUV and optical variability of NGC 4051	111
5.1	The strange variability properties of NGC 4051	112
5.1.1	Observations and data reduction	112
5.1.2	RMS variability and evidence for non-Gaussian variations in the lightcurve	113
5.1.3	The power spectrum	117
5.2	The low states	119
5.2.1	Variability of the low state	119
5.2.2	The low state X-ray spectrum	121
5.3	Simultaneous EUV and X-ray variability	126
5.3.1	The EUV-X-ray correlation	126
5.3.2	A simple spectral model	129
5.3.3	Constraining simple Comptonisation models	133
5.4	The relationship between X-ray and optical variability	137
5.5	Summary	141
6	Conclusions	143
6.1	The physical nature of the red-noise process	144
6.2	Power-spectral scaling in AGN and XRBs	148

6.2.1	What physical time-scale does the break frequency correspond to?	151
6.2.2	How well can we measure black hole mass?	153
6.2.3	Where are the AGN QPOs?	156
6.3	Accretion states and AGN unification	157
6.3.1	Evidence for high state AGN	157
6.3.2	Is NGC 4051 a macro-microquasar?	158
6.3.3	Different states: towards an AGN Grand Unification Theory	163

List of Figures

- 1.1 The standard AGN unification model (taken from Urry & Padovani 1995). Note that the picture is not to scale, as the innermost regions of the model AGN have been expanded for clarity. . . . 6

- 1.2 Plot of the ratio of the observed *RXTE* spectrum of NGC 5506 to a power law model fitted to the 2–5,7–10 keV range (M^cHardy, Papadakis & Uttley 1998). 9

- 1.3 Comparison of power spectra of different types of XRB. The objects are: (a) an accreting millisecond pulsar, (b) an atoll source (neutron star XRB), (c) a black hole XRB and (d) a Z source (neutron star XRB), (taken from Wijnands & van der Klis 1999). 14

- 1.4 512 s binned *EXOSAT* lightcurve of NGC 5506, obtained 24–27 January 1986. 18

- 1.5 The *RXTE* spacecraft. 20

- 1.6 *RXTE* PCA 2–10 keV monitoring lightcurves of MCG-6-30-15, NGC 5506 and NGC 5548. Time units are in Truncated Julian Date minus 10000 (TJD=JD-2440000.5) 23

2.1	Lightcurves (left) and power spectra (right) simulated using the method of Timmer & König, corresponding to model power spectra of power-law slopes $\alpha = 1$ (top), $\alpha = 2$ (centre) and a broken power-law model with high-frequency slope $\alpha = 2$ breaking to $\alpha = 0$ at frequency 10^{-4} (bottom). Note that, for comparison, all three simulated lightcurves are generated using the same random number sequence.	34
2.2	The effects of red-noise leak on power-spectral shape. The grey dashed line represents the true underlying power spectrum, while the solid black line shows the measured power spectrum after distortion by the limited sampling window. Underlying power spectral slopes from top to bottom are $\alpha = 2.5, 2.0, 1.5$ and 1.0 respectively.	38
2.3	The effects of aliasing on power-spectral shape. The grey dashed line represents the true underlying power spectrum, while the solid black line shows the measured power spectrum after distortion by the effects of discrete sampling. Underlying power spectral slopes from top to bottom are $\alpha = 2.0, 1.5$ and 1.0 respectively.	41
2.4	Local flux dependence of mean σ for SAX J1808.4-3658 and Cygnus X-1. The dotted lines mark the best-fitting linear models described in the text.	52

2.5	Local flux dependence of the power spectrum of Cygnus X-1. The black line marks the power spectrum for segments of mean flux greater than the total mean of the entire observation, the dark grey line marks the mean power spectrum for segments of mean flux less than the total mean. The corresponding combined mean fluxes for all segments used to make each power spectrum are also shown in the figure.	54
4.1	Intensive monitoring (top) and long-look (512-s binned) 2–10 keV lightcurves of MCG-6-30-15.	71
4.2	Intensive monitoring (top) and long-look (512-s binned) 2–10 keV lightcurves of NGC 5506.	73
4.3	Intensive monitoring (top) and long-look (512-s binned) 2–10 keV lightcurves of NGC 5548.	75
4.4	2–10 keV lightcurves of NGC 3516, from top to bottom: total monitoring lightcurve, intensive monitoring lightcurve, May 1997 long-look lightcurve, April 1998 long-look lightcurve (long-look lightcurves binned to 512 s).	78
4.5	Raw broadband power spectra of MCG-6-30-15, NGC 5506, MGC 5548 and NGC 3516. The dashed line shows the low-frequency part of the power spectrum, made from the total monitoring lightcurves, while the dotted line shows the high-frequency part made from long-look lightcurves. Solid lines mark the medium-frequency power spectrum (made from the intensive monitoring lightcurves), VHF power spectra for MCG-6-30-15 and NGC 5506, and the power spectrum of the second long-look observation for NGC 3516.	83

4.6	Comparison of best-fitting model average power spectra with the observed power spectra for the single power-law model described in the text. Open squares mark the low-frequency model average, simple crosses mark the high-frequency model average while filled squares mark the medium-frequency model average, the VHF model average for MCG-6-30-15 and NGC 5506, and the high-frequency model average for the second long-look observation of NGC 3516. Note that the error bars represent the RMS error in the simulated power spectra used to calculate χ_{dist}^2 as described in Section 3.1.	87
4.7	Probability that a single power-law of slope α is acceptable to describe the broadband power spectrum of NGC 5548.	88
4.8	The form of the knee model power spectrum.	90
4.9	Comparison of best-fitting model average power spectra with observed power spectra for the knee model described in the text (left), and corresponding confidence contours for the parameter space searched (right). The dashed and solid confidence contours represent the 68% and 90% confidence limits respectively.	92
4.10	The form of the high-frequency break model power spectrum.	94
4.11	Comparison of best-fitting model average power spectra with observed power spectra for the high-frequency break model described in the text (left), and corresponding confidence contours for the parameter space searched (right). The dashed and solid confidence contours represent the 68% and 90% confidence limits respectively.	97

5.1	2–10 keV lightcurves of NGC 4051, from top to bottom: total monitoring lightcurve, intensive monitoring lightcurve, December 1996 512-s binned long-look lightcurve.	114
5.2	Comparison of the best-fitting high-frequency break model average power spectrum with the observed broadband power spectrum of NGC 4051.	118
5.3	Simultaneous <i>EUVE</i> (DS/S instrument) and 4–10.5 keV <i>BeppoSAX</i> (MECS instrument) lightcurves of NGC 4051 in the 1998 low state, binned to 22176 s (4 <i>EUVE</i> orbits).	120
5.4	<i>EUVE</i> lightcurve of NGC 4051 during the 1999 low–normal state transition.	122
5.5	May 9-11 <i>RXTE</i> PCA and <i>BeppoSAX</i> MECS spectra for the best-fitting multiplicative reflection model described in the text.	125
5.6	May 1996: 5544 s binned <i>EUVE</i> (scaled by factor 100) and <i>RXTE</i> (3 PCUs, 2–10 keV) lightcurves of NGC 4051 (errors are 1σ).	127
5.7	December 1996: 5544 s binned <i>EUVE</i> (scaled by factor 100) and <i>RXTE</i> (3 PCUs, 2–10 keV) lightcurves of NGC 4051 (errors are 1σ).	127
5.8	Z-transformed discrete cross-correlation function of <i>EUVE</i> and <i>RXTE</i> (4–10 keV) lightcurves, combined from separate May and December ZDCFs. The ZDCF is binned to 20 ks and errors are 1σ	128

5.9	Comparison of linear and non-linear models (dotted line and solid line respectively). For clarity the data have been averaged (and 1σ errors calculated accordingly) into bins of $0.3 \text{ counts s}^{-1}$ width. See text for details of corresponding fit parameters. . . .	132
5.10	NGC 4051 X-ray, 5100\AA continuum and $H\beta$ line monitoring lightcurves.	138
5.11	Comparison of smoothed long-term X-ray and optical continuum lightcurves (taken from Peterson et al. 1999).	139
6.1	Relative flux of C compared to the mean flux in four energy bands.	147
6.2	Comparison of $\nu P(\nu)$ power spectra of MCG-6-30-15, NGC 5506, NGC 3516, NGC 5548 and Cygnus X-1 in the low and high states.	149
6.3	The low state transition in GRS 1915+105. Infra-red and radio lightcurves are plotted in normalised luminosity units on the same axis, while the zero flux base of the X-ray observations is represented by the flat line, corresponding to Earth occultation of the source. X-ray hardness ratio is plotted at the bottom of the figure and a depiction of the cloud ejection model is printed at the top (figure taken from Mirabel & Rodriguez 1999).	161

List of Tables

- 2.1 Flux-related changes in the variance of NGC 4051, NGC 5506 and MCG-6-30-15. Mean data are shown for low and high flux segments (corresponding to segment fluxes below and above the mean for the entire observation, given by μ_{tot}). μ is the mean flux (2–10 keV, count s⁻¹) for each flux bin, n is the number of segments in the flux bin, $\overline{\sigma^2}$ is the mean variance (count² s⁻²) and σ_{frac} is the fractional RMS variability ($\sigma_{\text{frac}} = (\sigma^2/\mu^2)^{\frac{1}{2}}$). . . . 49

- 4.1 Parameters of lightcurves for use in PSRESP. T and ΔT are the lightcurve duration and sampling interval (in seconds), μ and σ_{frac} are the lightcurve mean flux (in count s⁻¹) and fractional RMS respectively and P_{noise} is the Poisson noise level expected in the power spectrum due to counting statistics (in unnormalised RMS-squared units, count² s⁻² Hz⁻¹). Notes: ^a Details given in the table are for the *EXOSAT* lightcurve, the *RXTE* lightcurve used to measure the power spectrum at the highest frequencies has $\mu = 28.1$, $P_{\text{noise}} = 79.7$. ^b Lightcurve obtained 22–26 May 1997. ^c Lightcurve obtained 13–16 April 1998. 82

- 4.2 Time resolution of simulated lightcurves used in PSRESP. 85

4.3	Results from fitting broadband power spectra of four Seyfert galaxies with a simple unbroken power-law model.	86
4.4	Results from fitting broadband power spectra of four Seyfert galaxies with a knee model, see text for details	91
4.5	Results from fitting the broadband power spectra of four Seyfert galaxies with a high-frequency break model.	96
4.6	Energy dependence of RMS variability.	104
4.7	Energy dependence of power-spectral slope	109
5.1	NGC 4051 long-term lightcurve variability parameters	115
5.2	NGC 4051 variability parameters for the Gaussian-distributed portion of the long term lightcurve	116
5.3	Best-fitting reflection model parameters for PCA, MECS and combined MECS and PCA spectra. Γ is the photon index of the incident power-law continuum, A is the incident power-law normalisation, E_K and σ_K are the line energy and actual width respectively (in keV) and F_K is the line flux in photons $\text{cm}^{-2} \text{s}^{-1}$. All errors are 90% confidence limits for 2 interesting parameters. Also shown is the χ^2 value and number of degrees of freedom for each fit, and the 2–10 keV luminosity of the incident power-law in units of $10^{42} \text{ erg s}^{-1}$ (assuming $H_0 = 50 \text{ km s}^{-1} \text{ Mpc}^{-1}$).	124

Preface

The work presented in this thesis is original work carried out by the author. At the time of writing, two papers have been published based on some of the work presented here. The work described in Section 5.2.2 is published in Uttley et al., 1999, *MNRAS*, 307, L6, while the work described in Section 5.3 is published in Uttley et al., 2000, *MNRAS*, 312, 880.

Acknowledgements

I have come to owe a lot of people a lot of thanks over the past three and a half years. In particular I would like to thank my supervisor, Ian McHardy, first and foremost for starting such a fascinating *RXTE* program and allowing me to take part in it. I would also like to thank Ian for giving me just the right amount of help and encouragement, while allowing me to develop my own ideas and choose the path my project would take. I'd like to thank my other collaborators, Iossif Papadakis who gave me plenty of assistance and encouragement and helped to introduce me to the pleasures and pitfalls of power spectral analysis, and Antonella Fruscione who provided her generous *EUVE* expertise.

I'd like to thank everyone in the Astronomy Group who made it a real pleasure to be here, it's not all work, work, work, is it? I'd very much like to thank Georg Lamer, for being a good holiday/drinking buddy on our jaunts to Bologna and Hawaii and for doing all the hard work when we were actually observing. We had some interesting work discussions too.

I want to thank all my friends outside of work, who provided much-needed perspective when things got tough and made sure that I still had a life while I was a Ph.D. student. I want to thank Andrew for being such a good friend for nearly ten years now and being a good listener and great to chat to about anything. I'd like to big up Mark for being my best drinking buddy for the past four years. Well done son, I think I owe you for that. Much kudos and respect goes to Jason for his unflinching skepticism about most things and for being interested in 'technical' stuff too. And I must thank Alex for being a great friend for so long and also for being out of my Southampton circle so I'd have someone to gossip to whenever I went home (P.S. bwaahaahaa...excellent).

Finally, and most importantly of all, I'd like to thank my parents for getting me interested in reading and learning since I can remember, and for encouraging me at every step. All this is thanks to you.

And finally, some poetry, which perhaps sums up the contents of my thesis at least as well as the abstract.

‘A trend is a trend is a trend.
But the question is: will it bend?
Will it alter its course, through some unforeseen force,
And come to a premature end?’

Alexander Cairncross

‘I know what I know, and I know it is true.
If you ask me for proof, I can talk ’til you’re blue.
I have it all here in print and in scrawl,
It’s a fact, not a theory. In fact it’s a law.
I know it’s the truth, no matter how strange.
Oh, by the way, it’s all subject to change.’

C. C. Keiser

Chapter 1

Introduction

Overview

In this chapter I shall review the current state of understanding of active galactic nuclei (AGN), including emission mechanisms, unification models and the X-ray spectral and variability properties of AGN. We then look at the similarities between the X-ray spectra of AGN and X-ray binary systems (XRBs) and introduce the power spectrum as a method of characterising the X-ray variability of XRBs which has been very well studied to date. I then review the status of power spectral measurements of variability in AGN, which until recently have been limited by the difficulties of measuring variability on long time-scales. Finally I introduce the *Rossi X-ray Timing Explorer* mission and show how we have used it to obtain long-term monitoring lightcurves of a sample of Seyfert galaxies, in order that we can measure their broadband power spectra for comparison with those of XRBs.

1.1 Active Galactic Nuclei

Roughly 10% of galaxies have powerful sources of continuum emission in their nuclei. The emission from these active galactic nuclei (AGN) is far more powerful than could be accounted for by concentrations of stars and furthermore, varies on timescales of minutes (in the X-ray band) to months (in the optical band), indicating that the emitting region is very compact. The absolute luminosities of AGN cover an extremely broad range, from $\sim 10^{41}$ erg s $^{-1}$ in the lowest luminosity Seyfert galaxy NGC 4395 (Lira et al. 1999) to $\sim 10^{47}$ erg s $^{-1}$ in the extreme accretion rate quasar PDS 456 (Reeves et al. 2000). Despite this large range of luminosities, the spectral energy distributions of AGN show a similar pattern, in that the emission covers a very broad band and the broadband energy spectrum is relatively flat (e.g. Elvis et al. 1994) so that from the infra-red to γ rays, each decade of photon frequency contributes a similar amount to the total luminosity of the AGN. The continuum luminosity of most AGN falls off dramatically at radio wavelengths, however around 10% of AGN are said to be ‘radio-loud’ and have radio luminosities roughly a factor of 1000 or more greater than radio-quiet AGN with similar optical luminosities.

Apart from the radio-loud/radio-quiet dichotomy between AGN, distinctions can be made based on other observational characteristics. The broad range in luminosity is classified by splitting AGN into high-luminosity ‘quasars’ (typically those with bolometric luminosities exceeding 10^{45} erg s $^{-1}$) which may be radio-loud or radio-quiet, while lower luminosity AGN are classed as Seyfert galaxies or radio galaxies if they are radio-quiet or radio-loud respectively. These luminosity-based distinctions are somewhat arbitrary however, since the mechanism which powers all AGN is thought to be the same, namely accretion on to a massive black hole.

1.1.1 The central engine: accretion on to a massive black hole

In order to fall inwards towards a central object, orbiting gas must somehow lose kinetic energy. The standard picture of accretion on to black holes is that conservation of angular momentum forces the gas into a flat disk structure, an ‘accretion disk’ and then viscous forces between gas particles cause the gas to transfer kinetic energy into thermal energy so that it heats up, and falls further in towards the black hole. If the gas is optically thick, it cools by losing energy in the form of radiation, so that ultimately the accretion process can liberate $\sim 10\%$ of the rest mass of accreting matter as continuum emission. The attractive feature of the model of accreting black holes as the central engine of AGN is that the size of the black hole and accompanying accretion disk simply scales linearly with black hole mass, as does the luminosity of the system for a given accretion rate (as a fraction of the maximum ‘Eddington’ rate). Therefore the huge range in luminosity of AGN can be explained as the result of differing accretion rates on to black holes of mass $\sim 10^5\text{--}10^9 M_{\odot}$.

The theory of optically thick accretion disks has been studied in detail, first by Shakura & Sunyaev (1973) who studied the properties of geometrically thin disks, with the ratio of disk scale height H to disk radius R , $H/R \ll 1$. Later, other workers studied the properties of geometrically thick disks, which are thought to be more realistic, since close to the black hole, disk temperatures are high enough that radiation pressure puffs up the disk so that $H/R \sim 1$ (Abramowicz, Calvani & Nobili 1980; Abramowicz et al. 1988). Evidence for the existence of accretion disks in AGN comes from fits to the optical and UV spectra of quasars, which show evidence of ‘big blue bumps’ towards the UV, thought to be the peak of a black body spectrum of the temperature expected for a supermassive black hole (e.g. Brunner et al. 1997). Perhaps the strongest evidence for the existence of accretion disks comes from the detection of features in the X-ray spectra of AGN that are characteristic of reflection from

optically thick matter, which I shall discuss later in this chapter.

Recently, attention has turned towards an alternative mode of accretion to the standard optically thick accretion disk, which occurs at low accretion rates (below $\alpha^2 \dot{M}_{\text{Edd}}$, where α is the viscosity parameter and \dot{M}_{Edd} is the Eddington accretion rate). These so called ‘advection-dominated accretion flows’ (ADAFs, e.g. Narayan & Yi 1995) are optically thin accretion flows where most of the energy liberated by the accretion process simply heats up the optically thin plasma (and is ultimately carried into the black hole). Therefore ADAFs are extremely radiatively inefficient, a property which makes them suitable candidates for the mechanism that powers certain accreting black hole systems, which appear to have a much lower luminosity than their estimated accretion rate would suggest. Examples of such systems are the massive black holes at the centres of elliptical galaxies (Fabian & Rees 1995) and the massive black hole in the centre of our own galaxy, Sagittarius A* (Mahadevan 1998). Whether or not ADAFs really exist, and whether they can also be applied to describe accretion in higher luminosity systems such as AGN is currently a matter of great debate. It has been suggested that ADAFs are unstable in that they have difficulty in losing angular momentum and therefore lead to powerful outflows, so that most of the accreting material is expelled from the system (Blandford & Begelman 1999). The resulting outflow may help to suppress the strong radio synchrotron emission expected from ADAFs, explaining a significant discrepancy between the predictions of simple ADAF models and observations of the nuclei of elliptical galaxies (Di Matteo et al. 2000).

Relatively weak outflows from AGN have been observed (e.g. Christopoulou et al. 1997), however the most powerful outflows are seen in the form of powerful relativistic jets which emit synchrotron radiation, significantly contributing to the radio luminosity of the AGN and thus marking the distinction between radio-loud and radio-quiet AGN. Although the jets are clearly associated with outflowing material, how they become so collimated and energetic is still unknown. The best candidate for powering such jets is by extracting energy from

a rapidly rotating black hole via magnetic fields set up by the accreting matter (the Blandford-Znajek mechanism, e.g. Meier 1999). The radio-loud/radio-quiet dichotomy may then be a manifestation of black hole spin. However, the detailed physics of such a process, and in particular how magnetic fields are maintained in an accreting system are still unknown.

1.1.2 AGN unification

We have seen that the common feature of all AGN is their central engine, namely an accreting supermassive black hole, and that the range in luminosities of AGN and the radio-loud/radio-quiet dichotomy can be explained as a consequence of, respectively the black hole mass and whether the system has a powerful jet (which may be associated with a high black hole spin). Other differences between AGN include whether they have strong broad and narrow optical emission lines (e.g. Seyfert 1 galaxies and broad line radio galaxies, BLRGs), narrow lines only (Seyfert 2 galaxies and narrow line radio galaxies, NLRGs), or an intermediate spectrum with relatively weak broad lines (Seyfert 1.5, 1.8 and 1.9), or whether the AGN is classed as a blazar and has features characteristic of strongly beamed emission from a relativistic jet, namely relatively strong X/ γ -ray emission, a flat-spectrum radio-core, high optical polarisation and rapid variability in all wavebands.

To remain true to Occam's Razor and encompass all these different types of AGN in a simple picture with the minimum number of ad hoc assumptions, we assume the AGN unification model depicted in Figure 1.1. At the centre of the AGN is the accreting black hole, which may or may not produce a powerful jet (so the AGN depicted in Figure 1.1 is radio-loud). The accretion disk may extend to several light-days in size for the largest black holes. Beyond it out to distances of light-months, rapidly moving clouds in the broad line region (BLR) reprocess the central continuum to produce broad optical




Figure 1.1: The standard AGN unification model (taken from Urry & Padovani 1995). Note that the picture is not to scale, as the innermost regions of the model AGN have been expanded for clarity.

emission lines, while at much greater distances (hundreds of light-years) low density, slow moving clouds in the narrow line region (NLR) reprocess continuum radiation into narrow emission lines. A key component of the model is a dense molecular torus at distances of light-years from the central engine. The torus is optically thick and heavily reddens or obscures emission from the BLR and central continuum source, so that the additional classification of the AGN (apart from its radio-loudness or luminosity) depends on how we view it. If we view the AGN through the torus, we see a Seyfert 2 or NLRG, whereas if we have a clear line-of-sight to the BLR we see a Seyfert 1 or BLRG. Finally, we see a Blazar if we look directly down the jet so that we see beamed emission. There is much evidence to support the standard AGN unification model, such as the detection of ionisation ‘cones’ in Seyfert galaxies indicating collimation of the continuum radiation by the torus (e.g. Tadhunter & Tsvetanov 1989) and the detection of broad optical lines in scattered polarised light in Seyfert 2 galaxies, indicating the presence of a hidden BLR (e.g. see Antonucci 1993 for details and a review of evidence for the standard unification model). However, there are a number of puzzles yet to be addressed, in particular the nature of a sub-class of Seyfert galaxies known as narrow-line Seyfert 1s (NLS1s), which I shall touch on in the next section.

1.1.3 X-ray spectral properties of AGN

AGN vary most rapidly in the X-ray band, implying that the X-rays are emitted close to the central black hole, where light-crossing time-scales are shortest, as are the variability time-scales in the accretion disk (Treves, Maraschi & Abramowicz 1988). Simple accretion disk models, which predict inner disk temperatures of $\sim 10^5$ K (i.e. peaking in the UV) cannot explain the hard X-ray emission from AGN, which requires temperatures of greater than 10^8 K. The X-ray spectra of AGN have a power-law form, with the photon flux at

energy E , $F(E) \propto E^{-\Gamma}$ (where Γ is called the ‘photon index’), which is characteristic of non-thermal emission. Non-thermal X-ray emission may be produced by synchrotron processes in a relativistic jet, as is seen in Blazars. However, the presence of a non-thermal X-ray spectrum in non-beamed sources such as Seyfert galaxies and quasars suggests a different mechanism in these sources, namely Compton up-scattering of low-energy ‘seed’ photons from a thermal distribution of electrons. The Comptonizing medium is thought to lie above the inner region of the accretion disk in some sort of tenuous hot ‘corona’ analogous to that of the Sun (e.g. Haardt, Maraschi & Ghisellini 1997). The mechanism for heating the corona is unknown, but by analogy with models for heating the Solar corona, high coronal temperatures might be achieved by magnetic reconnection in the corona driven by turbulence in the underlying accretion disk (Di Matteo 1998).

Evidence for an accretion disk is strongest in the X-ray band, since X-ray spectra show features characteristic of reprocessing in optically thick material, namely a strong fluorescent Fe K line at ~ 6 keV and a reflection ‘hump’ above 10 keV (see Figure 1.2). The most well-studied iron line, in the Seyfert 1 galaxy MCG-6-30-15 is intrinsically redshifted and shows a strong red wing (Tanaka et al. 1995), and can be well described by models where the line originates in the inner disk, close to the marginally stable orbit, where it is heavily gravitationally redshifted and Doppler broadened by the relativistic motion of the gas (e.g. Bromley, Miller & Pariev 1998). Similar evidence has been found in a number of other Seyfert galaxies (Wang, Zhou & Wang 1999, Nandra et al. 1997b), while possible reflection humps, caused by a combination of absorption at low energies and down-scattering of high energy photons have been detected in a large number of Seyfert galaxies, providing added support to the disk reflection model (Nandra & Pounds 1994). Radio-quiet quasars also show evidence for broad iron lines, although with some evidence that the accretion disks in quasars may be more heavily ionised than those in Seyferts, consistent with significantly higher accretion rates in these objects (Nandra et al. 1997c,




Figure 1.2: Plot of the ratio of the observed *RXTE* spectrum of NGC 5506 to a power law model fitted to the 2–5, 7–10 keV range (McHardy, Papadakis & Uttley 1998).

Reeves et al. 2000).

X-ray spectra of Seyfert 2 galaxies show evidence for heavy low energy absorption (absorbing column densities $> 10^{23} \text{ cm}^{-2}$) consistent with what is expected from obscuring tori, an interpretation which is further confirmed by the presence of narrow iron K lines which are not redshifted in the AGN reference frame, as we might expect if the line emission originates in reflection from the distant torus (e.g. Turner et al. 1998).

In general the X-ray spectra of radio-quiet AGN look very similar, with photon index $\Gamma \sim 1.9$ independent of X-ray luminosity. However the class of AGN known as narrow-line Seyfert 1s (NLS 1s) show significantly steeper power-law spectra ($\Gamma \sim 2.1$) and strong soft excesses at low energies which can be fitted with black body spectra of $kT \sim 150 \text{ eV}$ (Vaughan 1999). NLS 1s show significantly narrower broad permitted lines than normal Seyfert 1s, leading to the suggestion (Boller, Brandt & Fink 1996) that NLS 1s have relatively lower

central black hole masses than AGN with broader emission lines (if the line widths are associated with Keplerian orbital velocities). Low black hole masses would imply higher accretion rates in NLS 1s, which might explain their soft excesses and steeper X-ray spectra, since accretion disk temperature increases with accretion rate so that systems accreting at a high rate have hotter, more powerful disk emission (hence the blackbody soft excess) which can steepen the hard X-ray power-law continuum by Compton-cooling the corona.

1.1.4 X-ray variability properties of AGN

The X-ray lightcurves of AGN obtained by long-look observations (i.e. up to a few days duration) show scale-invariant variability with no characteristic timescales, such as periodicities (M^cHardy & Czerny 1987, Lawrence et al. 1987). Studies of the X-ray variability of reasonably large samples of radio-quiet AGN show that the X-ray variability amplitude decreases with increasing luminosity (Green, M^cHardy & Lehto 1993, Lawrence & Papadakis 1993, Nandra et al. 1997a). An anti-correlation between X-ray variability amplitude and luminosity might be expected if all AGN lightcurves are made up of pulses of the same intrinsic luminosity and variations are caused by Poisson fluctuations in the number of pulses. In this model, lightcurves of more luminous AGN are made out of more pulses and hence show relatively smaller fluctuations. Recently however, observations of luminous NLS 1s have shown strong variability which upsets the luminosity-variability trend (Turner et al. 1999). If NLS 1s have less massive black holes than normal Seyferts, their greater variability amplitudes may be associated with the smaller variability time-scales associated with lower mass black hole systems. Therefore the luminosity-variability trend exhibited by normal Seyfert galaxies may represent a similar accretion rate among normal Seyferts so that more luminous objects have more massive

black holes (and hence are less variable).

A radical interpretation of this model would suggest that *variability processes in accreting black hole systems are identical irrespective of black hole mass*. In other words, the number of varying regions or pulses stays the same but their luminosity and variability time-scales increase (possibly linearly) with black hole mass. Support for this model comes from the similarities between the variability properties of AGN and black hole X-ray binary systems, which I shall outline in the next section.

1.2 X-ray binary systems and the XRB-AGN connection

X-ray binary systems (XRBs) are the brightest X-ray sources in the sky, but are generally unremarkable at longer wavelengths. Their strong X-ray emission and extremely rapid variability (on time-scales of milliseconds) bears all the hallmarks of accretion onto a compact object. However unlike the massive central black holes in AGN, the compact objects in XRBs are much less massive neutron stars ($\sim 2 M_{\odot}$) and $\sim 10 M_{\odot}$ black holes, which accrete material from the strong stellar winds of high mass companion stars (HMXBs) or, more commonly, via Roche lobe overflow or accretion disk instability following a build-up of accreted matter from a low mass companion (LMXBs).

Despite the huge difference in size and luminosity, the X-ray spectral properties of black hole XRBs are remarkably similar to those of AGN. The classic black hole X-ray binary system Cygnus X-1, in its typical low state shows a hard power-law continuum (photon index $\Gamma \sim 1.6$) and reflection features characteristic of reprocessing of X-rays by the accretion disk (e.g. Done & Zycki 1999). In its high state (which is thought to correspond to periods of high accretion rate), the power-law continuum of Cygnus X-1 becomes steeper

($\Gamma \sim 2.3$) and a significant soft-excess appears which can be described by a spectral model corresponding to an accretion disk of temperature ~ 0.4 keV (Dotani et al. 1997), suggesting a possible analogy with the NLS 1s.

1.2.1 The power spectra of XRBs

The X-ray variability properties of XRBs have been well studied over the past two decades, owing to the rapid variability and high measured fluxes of these sources, which mean that a significant range of variability time-scales (from ms to hours) may be sampled in relatively short observations. The most common technique for characterising the variability of XRBs is to measure the power spectrum of the lightcurve, which is simply the modulus squared of the Fourier transform of the lightcurve. Power spectra show the relative contribution to the total variability of a lightcurve made by variations on different time-scales (corresponding to power-spectral frequencies). For example, a single, well-defined periodic signal (e.g. the pulse period of a pulsar) will show up as a single peak in a power spectrum at the frequency of the signal, while a ‘quasi-periodic oscillation’ (QPO) which is not so well defined (if the period changes slightly or the pulse shape is broad), leads to a broader peak. A so-called ‘white-noise’ process which corresponds to random, uncorrelated fluctuations in a lightcurve (such as that seen when the observed lightcurve of a constant source varies due to photon counting statistics, i.e. Poisson noise) leads to a perfectly flat power spectrum. A self-similar lightcurve which shows correlated variability (i.e. trends on all time-scales), such as that seen in AGN corresponds to a ‘red-noise’ process, such as a random walk or flicker noise, and shows a characteristic power-law power spectrum (power at frequency ν , $P(\nu) \propto \nu^{-\alpha}$, where α is the power-spectral slope).

The power spectra of X-ray binaries show a combination of quasi-periodic,

white-noise and red-noise processes depending on the time-scales which they are studied at. At high frequencies, neutron star XRBs show evidence of kHz QPOs which may be associated with the neutron star spin or the beat between the neutron star spin frequency and the Keplerian frequency associated with rotation of the inner edge of the accretion disk. The highest frequency QPOs shown by black hole XRBs are at 67 Hz, ~ 185 Hz and 300 Hz, found in the systems GRS 1915+105, XTE J1550-564 and GRO J1655-40 respectively (e.g. Morgan, Remillard & Greiner 1997, Remillard et al. 1999). XTE J1550-564 is an X-ray nova, while GRS 1915+105 and GRO J1655-40 are both classed as ‘microquasars’ in that they have radio jets with a morphology similar to that of radio-loud quasars and along which material has been seen to be ejected from the system (Mirabel & Rodrigues 1994). High frequency QPOs have not been seen in more ‘normal’ persistent black hole X-ray binary systems, but at lower frequencies (1–20 Hz) they are quite common (Wijnands & van der Klis 1999, and see Figure 1.3), although their origin is still unclear.

Red-noise variability contributes the most to the total variance of all XRB lightcurves (Belloni & Hasinger 1990a), which show characteristic power-law power spectra at high frequencies although at low frequencies the power spectra flatten to $\alpha = 0$ so that the lightcurves become white-noise on long time-scales. The ‘knee’ frequency where the power spectrum flattens to zero slope has been shown to scale with the frequency of the low-frequency QPO (Wijnands & van der Klis 1999) suggesting that the knees and QPOs are correlated with a single parameter, possibly the accretion rate of material in the disk.

The power-law slope above the knee frequency, α is typically ~ 1 or greater, although in many XRBs (particularly neutron star systems) the presence of QPOs confuses the situation and makes the true underlying slope difficult to quantify. In the well studied black hole XRB Cygnus X-1 however, the power spectrum above the knee (typically at ~ 0.1 Hz) is relatively clean, and in the low state shows a second, high-frequency break at ~ 2 Hz. Above the high-frequency break, the slope is ~ 2 , while at intermediate frequencies between




Figure 1.3: Comparison of power spectra of different types of XRB. The objects are: (a) an accreting millisecond pulsar, (b) an atoll source (neutron star XRB), (c) a black hole XRB and (d) a Z source (neutron star XRB), (taken from Wijnands & van der Klis 1999).

the low-frequency knee and high-frequency break, the slope is very close to 1. In the high state, Cygnus X-1 shows a higher high-frequency break at ~ 10 Hz but no evidence for a low-frequency knee, with the $\alpha \simeq 1$ power-law extending down to at least 10^{-2} Hz (Cui et al. 1997). The increase in the high-frequency break may be associated with the accretion rate increase which is thought to cause the soft high-state energy spectrum, and could possibly be associated with the inwards movement of the inner edge of the accretion disk at higher accretion rates.

Changes in the variability properties of XRBs, such as the low-state/high-state transition in Cygnus X-1 occur on time-scales of months to years. Besides the low and high states, other rarer states have been identified in black hole X-ray binaries, namely the ‘off’ state where the X-ray source becomes quiescent and the ‘very high’ state where the X-ray emission is strong but only weakly variable and shows evidence for both strong hard and soft spectral components (see Fender 1999 for a brief review).

Due to the stochastic nature of red-noise lightcurves, power spectra measured over short time periods show random fluctuations about the mean power, which average out over long time-periods if the power spectrum is *stationary*. Power spectra which change systematically in shape or amplitude over time are referred to as being *non-stationary*. Although on short time-scales XRB power spectra are classed as being stationary, all are ultimately non-stationary on time-scales of months or longer, probably due to disc instabilities in the accretion flow which are not yet understood. The microquasar GRS 1915+105 shows particularly complex and rapid non-stationary behaviour, including periods of strong flaring behaviour followed by quiescent periods which sometimes repeat quasi-periodically on time-scales of less than an hour and are associated with radio and infra-red flares which may indicate the ejection of material from the inner disk into the jet (Mirabel et al. 1998). Cygnus X-1 shows evidence of non-stationarity *within* its low state on time-scales of days to months, in that its knee and break frequencies vary, along with the slope of the power

spectrum above the high-frequency break (Belloni & Hasinger 1990). Other low state black hole XRBs show similar behaviour (Main et al. 1999).

Favoured interpretations of the red-noise power spectra of XRBs generally involve shot-noise models, where the lightcurves are built from superpositions of randomly occurring shots, whose long decay times allow for the long-term memory inherent in red-noise lightcurves. I shall discuss details of variability models in context with observational results later in this work.

1.2.2 The power spectra of AGN

Given the spectral similarities between XRBs and AGN, we might also expect to see similarities in their variability properties, although perhaps with the relevant variability time-scales scaled up by the black hole mass. An important feature of black hole systems is that their characteristic length scale (namely the Schwarzschild radius, also called gravitational radius) scales linearly with black hole mass, so that the geometry of the X-ray emitting regions of XRBs and AGN should be similar. The difference in environments between black holes in AGN and XRBs should not be important, because the X-ray emitting region is gravitationally dominated by the black hole and should be oblivious to factors such as the presence of a companion star, which might affect the outer part of the accretion disk.

Studies of the X-ray variability of AGN are difficult because of their relatively long variability time-scales (even in the X-ray band) and the relatively poor signal-to-noise of AGN lightcurves (compared to XRB lightcurves) due to their faint nature. Early power-spectral studies by Lawrence et al. (1987) and McHardy et al. (1987) showed that AGN power spectra had a red-noise (i.e. power-law) form, with no characteristic time-scales such as frequency breaks or QPOs. In fact, the QPOs in AGN lightcurves have proved particularly elusive. Until recently, the only possible QPO detection in an AGN lightcurve

was found in *EXOSAT* observations of the Seyfert galaxy NGC 5548, with a 500 s period (Papadakis & Lawrence 1993a), although due to the short duration of the lightcurves used it is difficult to assess the significance of this result. Recently, an *ASCA* observation of the Seyfert galaxy IRAS 18325-5926 showed evidence for a periodicity of 58 ksec (Iwasawa et al. 1998). Again the significance of this periodicity is difficult to assess since its likelihood was obtained using the Lomb method (Press et al. 1992), which assumes that the noise underlying the periodic signal is white noise - a situation which does not apply to the red-noise lightcurves of AGN. Furthermore, the periodic signal in the lightcurve of IRAS 18325-5926 was not detected by subsequent *RXTE* observations, so it would seem that AGN QPOs remain as elusive as ever.

Measurements of the power-spectral slope of AGN show that they are relatively steep, with $\alpha \sim 1.5-2$ (Lawrence & Papadakis 1993; Green, M^cHardy & Lehto 1993), perhaps implying that they lie above the high-frequency break if it occurs in the power spectra of AGN. However, due to the relatively short length of lightcurves obtained from long-look observations, compared to AGN variability time-scales, it is difficult to measure power spectra at low enough frequencies to determine if they really flatten. Theoretically, AGN power-spectra must flatten to $\alpha < 1$ at some point, in order to prevent the total variance of the lightcurve, equivalent to the integrated power, from diverging to infinity. The key questions are: where does this flattening occur, and do the *overall* shapes of AGN power spectra remain similar to that of black hole XRBs, even at low frequencies (i.e. does the power spectrum scale by some common factor)?

The first attempt to measure the broadband power spectrum of an AGN was made by M^cHardy (1988), who combined a power spectrum obtained from a particularly long continuous *EXOSAT* observation of NGC 5506 (see Figure 1.4) with that obtained from a crude, very long time-scale lightcurve made from sporadic observations of NGC 5506 by a variety of missions over the course of a decade. The resulting broadband power-spectrum seemed to show




Figure 1.4: 512 s binned *EXOSAT* lightcurve of NGC 5506, obtained 24–27 January 1986.

evidence of flattening, but no formal statistical significance could be placed on this result due to the uncertainties in how the measured power spectrum depended on the crude sampling pattern of the long-term lightcurve. Papadakis & McHardy (1995) applied a similar method to measure a crude broadband power spectrum of the Seyfert galaxy NGC 4151, except that they used Monte Carlo methods to show that despite the poor sampling, the flattening they detected at low frequencies was significant. However, due to the poor sampling of the long-term lightcurve, nothing more could be said about the shape of the power spectrum on long time-scales, nor could any break or knee frequencies be defined.

In order to measure a good quality broadband power spectrum of AGN with good sampling over a range of time-scales, we need a dedicated monitoring program. This has become possible in the last few years, thanks to the work of the *Rossi X-ray Timing Explorer*.

1.3 Long-term monitoring of Seyfert galaxies with *RXTE*

The launch of the *Rossi X-ray Timing Explorer* (*RXTE*) mission, on December 30th 1995, offered the first real opportunity to monitor AGN over long time-scales with good sampling, enabling power-spectral studies of AGN to extend to lower frequencies than ever before and allowing the prospect of searching for flattening in AGN power spectra and comparing AGN power spectra with those of XRBs over a broad range of time-scales.

1.3.1 *RXTE*

RXTE is specifically designed to carry out timing studies of variable X-ray sources on time-scales ranging from microseconds to years. It carries three scientific instruments (see Figure 1.5), the Proportional Counter Array (PCA), the High Energy X-ray Timing Experiment (HEXTE) and the All-sky Monitor (ASM).

The PCA consists of 5 Xenon-filled proportional counter units (PCUs, numbered 0–4), sensitive in the energy range 2–60 keV, with moderate energy resolution ($< 18\%$ at 6 keV) and 1 microsecond timing resolution. The total collecting area of the 5 PCUs is 6500 cm^2 , allowing large count rates to be measured, even for relatively faint sources (e.g. $\sim 10 \text{ count s}^{-1}$ for an AGN of flux $\sim 10^{-11} \text{ erg cm}^{-2} \text{ s}^{-1}$), however since the PCA is not an imaging instrument (spatial resolution is determined by a collimator with 1 degree FWHM), background count rates are high so that sensitivity is limited to fluxes $\sim \text{few} \times 10^{-12} \text{ erg cm}^{-2} \text{ s}^{-1}$. Since launch, discharge problems have meant that one or both of two of the 5 PCUs (3 and 4) are often switched off. This problem has extended to a third PCU (PCU 1) since March 1999. Despite these problems, adequate count rates can still be measured for many AGN. The




Figure 1.5: The *RXTE* spacecraft.

background contribution to the observed count rate is $\sim 12 \text{ count s}^{-1}$ in the 2–10 keV band (top Xenon layer only) for 3 PCUs, and since no imaging or offset-pointing information is available, the background needs to be modelled based on satellite engineering and housekeeping data. The latest background models are more than adequate for this task.

The HEXTE covers the energy range 15–250 keV and has energy resolution $\sim 15\%$ at 60 keV and time resolution of 8 microseconds. It consists of two non-imaging clusters of scintillation counters, with a 1 degree FWHM field of view, which rock on and off-source every 16–128 s so that the background is continually measured by the off-source pointing cluster (i.e. background modelling is not required). The collecting area of each cluster is only 800 cm^2 , so that only $\sim 1 \text{ count s}^{-1}$ is measured from even the brightest AGN. Despite the accurate background subtraction, the background rate (50 count s^{-1} per cluster) is sufficiently high that HEXTE data are only marginally useful for AGN studies, and then only for spectral studies with large exposure time. Therefore

I shall not use any HEXTE data in this work.

The ASM covers an energy range of 2–10 keV with very limited energy resolution, and consists of three wide-angle shadow cameras equipped with proportional counters, with a total collecting area of 90 cm². The shadow cameras sweep $\sim 80\%$ of the sky every 90 minutes, so that fluxes from individual bright sources can be integrated over time to yield a sensitivity of $\sim 10^{-10}$ erg cm⁻² s⁻¹. Unfortunately, this makes the ASM only just sensitive to a few bright AGN and certainly incapable of long-term monitoring of AGN.

One of the purposes of *RXTE* has been to use its exceptionally high time resolution capability to study the timing properties of XRB systems on very short time-scales, using relatively short (~ 10 ksec) observations. However, a second revolutionary feature of *RXTE* is its rapid pointing capability, i.e. within just a few minutes *RXTE* can slew to point at a new target. No other X-ray astronomy mission to date has had this capability, which allows *RXTE* to efficiently observe sources for short ‘snapshots’ (~ 1 ksec duration) without wasting a lot of valuable mission time in slewing to targets. For this reason, *RXTE* is ideal for carrying out monitoring observations of targets at regular (and sometimes very frequent) intervals over long periods of time, allowing the possibility of measuring reliable broadband power spectra of AGN for the very first time.

1.3.2 The Seyfert monitoring program

Much of the work presented in this thesis is based on a long-term program to monitor Seyfert galaxies with *RXTE* over a range of time-scales, initiated in the first observing cycle (AO1, Proposing Investigator: Ian M^cHardy) and continued in subsequent cycles up to the present. The initial sample of Seyfert galaxies consists of four targets. In order of ascending luminosity

they are: NGC 4051, MCG-6-30-15, NGC 5506 and NGC 5548, with typical luminosities of $\sim 5 \times 10^{41}$ erg s $^{-1}$, $\sim 10^{43}$ erg s $^{-1}$, $\sim 10^{43}$ erg s $^{-1}$ and $\sim 5 \times 10^{43}$ erg s $^{-1}$ respectively. Thus, a large range of luminosities is covered and presumably a large range of black hole mass, allowing variability as a function of luminosity to be studied. All the targets are bright (2–10 keV flux $\sim \text{few} \times 10^{-11}$ erg cm $^{-2}$ s $^{-1}$) and have been well studied previously, so that the results obtained from long time-scale monitoring can be put into context with existing knowledge of these sources. Furthermore, the four Seyfert galaxies selected for study are known to be variable to some extent, so that time is not wasted on studying the X-ray variability of constant sources. NGC 4051, MCG-6-30-15 and NGC 5548 are all Seyfert 1 objects (although NGC 4051 is also classed as NLS 1), while NGC 5506 is classed as Seyfert 2, although this classification may not be due to obscuration by a torus but rather the fact that this AGN is seen through an edge-on disk host galaxy.

The purpose of the program is to measure the broadband power spectra of the sample, to see if the power spectra flatten on long time-scales as we might expect, and to compare the power-spectral shapes with that of XRBs, in the hope of determining whether characteristic time-scales such as power-spectral breaks or knees exist and whether they scale with black hole mass. In order to achieve these goals, we must measure the power spectrum as well as we can across as broad a frequency range as possible. In order to sample variations on the shortest time-scales, we can use long-look observations, some of which were obtained as part of other programs (as in the case of MCG-6-30-15 and NGC 5548) while we obtained long-looks of NGC 4051 and NGC 5506 as part of our monitoring program. To sample variability on longer time-scales we must monitor our targets at regular intervals, using ~ 1 ksec snapshots to obtain a good measure of flux at each sampled point. However, it would not be efficient to define a single shortest sampling interval (e.g. one day) and use it to sample variability on time-scales much longer than that interval, so instead we sample the variability on intermediate and long time-scales using




Figure 1.6: *RXTE* PCA 2–10 keV monitoring lightcurves of MCG-6-30-15, NGC 5506 and NGC 5548. Time units are in Truncated Julian Date minus 10000 (TJD=JD-2440000.5)

several different observing schemes. Variability on intermediate time-scales is measured by monitoring our targets at twice-daily intervals for two weeks and then daily intervals for a month. On longer time-scales, we monitor the targets at weekly intervals (for the first cycle) and fortnightly thereafter.

The resulting long-term lightcurves (of almost three years duration) that we have obtained for three of our targets are shown in Figure 1.6, to demonstrate the sampling pattern and also show that our targets do in fact seem to vary on all time-scales. In addition to these data, I shall also use data from a separate *RXTE* monitoring program of the Seyfert 1 galaxy NGC 3516 (P.I. Rick Edelson), obtained from the *RXTE* archive, which is detailed later in this work, along with further details of the monitoring observations and more detailed lightcurves showing the long-looks and intensive monitoring periods.

1.4 Thesis overview

The main aim of this thesis is to investigate the broadband power-spectral shape of the sample of Seyfert galaxies which we have monitored with *RXTE*. In Chapter 2, I will introduce some of the basic techniques used to determine the power spectrum of a lightcurve, together with the standards and definitions which I shall use throughout this work. I will introduce the method I shall use to simulate lightcurves from a given power-spectral model and use simulations to demonstrate the effects of sampling on the power spectrum, which leads to undesirable effects such as red-noise leak and aliasing. I will end by discussing the concept of stationarity in more detail and presenting the first important new result of this work, namely that the power spectra of AGN and XRBs are intrinsically non-stationary, in that the power-spectral amplitude scales with the square of local mean flux.

In Chapter 3, I will discuss ‘PSRESP’, a method which I have developed from the ‘response’ method of Done et al. (1990), for testing different power-spectral

models against observed power spectra which are distorted by sampling effects and estimating reliable confidences on model parameters using Monte Carlo techniques.

In Chapter 4, I will apply the PSRESP technique to determine whether the broadband power spectra of our sample of Seyfert galaxies actually flatten on long time-scales. I will test knee and high-frequency break models for the power-spectral shape against the observed power spectra, and discuss the implications of the derived characteristic frequencies for scalings with black hole mass and luminosity. I will also discuss the energy dependence of fractional RMS variability and power-spectral shape.

In Chapter 5, I will discuss the unusual variability properties of NGC 4051, which entered a prolonged low state twice during the course of our monitoring campaign, and also shows evidence of unusual non-Gaussian variability in its lightcurve. I show that the X-ray energy spectrum measured during the low state is consistent with that expected if the central X-ray source has switched off, leaving only the spectrum of X-ray reflection from a distant molecular torus. I discuss the relationship between the X-ray and EUV bands on the basis of results obtained from simultaneous EUV and X-ray observations of NGC 4051, and show how we can use constraints on lags between the two bands to constrain simple Comptonisation models for the production of the X-ray continuum. I will also present and interpret some key results from a comparison of the long-term X-ray monitoring lightcurve of NGC 4051 with optical monitoring data obtained by the ‘AGN Watch’ team.

Finally, in Chapter 6 I will speculate on some of the key results obtained in this work. I will first discuss the implications of the discovery of an intrinsic flux dependence of variability described in Chapter 2, which casts doubt on shot-noise models where the basic building blocks of a lightcurve are small shots. I will propose an alternative model, where the dominant structures are large flaring regions, which break into fractal structure so that variability on all time-scales is produced and remains correlated with overall flux. I will dis-

cuss the implications of the power-spectral scalings found in Chapter 4 for the physical time-scales which break-frequencies might correspond to. I will also discuss prospects for using X-ray variability to measure the black hole mass. I will discuss the implications of the possibility that the different states seen in black hole XRBs may have analogs in AGN, and argue that the unusual variability properties of NGC 4051 suggest that it is in a state analogous to that seen in the microquasar GRS 1915+105, i.e. NGC 4051 may be a ‘macro-microquasar’. Throughout the concluding chapter I will discuss possibilities for future work using forthcoming data.

Chapter 2

Power Spectral Issues

Overview

In this chapter, I shall review how the power spectrum is measured and define the conventions which will be used to measure power spectra throughout this work. We will learn how to simulate lightcurves from power-spectral models and use qualitative arguments and simulations to investigate the distorting effects of sampling on the power spectrum. Finally, we will investigate the problem of comparing power spectra when lightcurves are non-stationary and I shall present the first major observational result of this work, that the RMS variability of AGN and the X-ray binary systems Cygnus X-1 and SAX J1808.4-3658 scales linearly with X-ray flux.

2.1 Measuring the power spectrum

Prior to measuring the power spectrum of a lightcurve, it is conventional to subtract the mean lightcurve flux from each data point, in order to remove the zero-frequency power from the power spectrum. The next step in measuring the power spectrum of a lightcurve is to take the Fourier transform of the

mean-subtracted lightcurve. However, the classical definition of the Fourier transform of a function $f(t)$ requires that $f(t)$ be continuous. In the case of real data, the time series of length N data points, $f(t_i)$ is not continuous, it is discretely sampled, and may be unevenly sampled. Therefore, throughout this work I shall use the method of determining the power spectrum applicable to discretely (and possibly unevenly) sampled data, given by Deeming (1975), who defines a discrete Fourier transform at a frequency ν :

$$F_N(\nu) = \sum_{i=1}^N f(t_i) \exp^{i2\pi\nu t_i},$$

so that the modulus squared of the Fourier transform is given by:

$$|F_N(\nu)|^2 = \left(\sum_{i=1}^N f(t_i) \cos(2\pi\nu t_i) \right)^2 + \left(\sum_{i=1}^N f(t_i) \sin(2\pi\nu t_i) \right)^2.$$

Note that the frequencies sampled by the discrete Fourier transform occur at evenly spaced intervals, $\nu_{\min}, 2\nu_{\min}, 3\nu_{\min}, \dots, \nu_{\text{Nyq}}$, where ν_{\min} is equal to T^{-1} (where T is the total duration of the lightcurve, i.e. $T = t_N - t_1$) and the Nyquist frequency $\nu_{\text{Nyq}} = (2T/N)^{-1}$.

Finally, we obtain the power $P(\nu)$ by applying a normalisation to $|F_N(\nu)|^2$:

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

By applying this normalisation, we obtain the desirable outcome that integrating the power spectrum over a given frequency range, ν_1 to ν_2 yields the contribution to the variance of the lightcurve due to variations on time-scales of ν_2^{-1} to ν_1^{-1} . Thus the total variance of the lightcurve is given by the integral of the power spectrum across all measured frequencies, ν_{\min} to ν_{Nyq} (e.g. van der Klis 1997). The units of power under this ‘root mean squared, squared’ (RMS squared) normalisation are $\text{count}^2 \text{ s}^{-2} \text{ Hz}^{-1}$ for power spectra made from lightcurves where lightcurve flux is given as a photon count rate.

Another power spectral normalisation that is sometimes used in X-ray variability studies of XRBs is the Leahy normalisation, $2/n$ where n is the total number of counts measured in the lightcurve (Leahy et al. 1983). The advantage of the Leahy normalisation is that the flat Poisson-noise level in the

power spectrum which is contributed by photon counting statistics has a value of 2 and follows a χ^2 distribution with 2 degrees of freedom, allowing an easy calculation of the significance of, e.g. peaks in the power spectrum due to periodic signals measured from pulsars. For red-noise lightcurves like those seen in AGN and XRBs, periodic signals must be measured against the background of the intrinsic broad-band noise of the source itself, rather than the Poisson level, which somewhat negates the useful aspect of this normalisation. For this reason and the ease of interpreting the power spectral amplitude in terms of lightcurve variance, most work on XRBs in the past decade has made use of the RMS squared normalisation, which I shall also use throughout this work (prior to normalising the power spectra by the mean squared flux, see Section 2.4).

Note that the Poisson level in this normalisation is simply equal to twice the total average count rate of the lightcurve. For a satellite like *RXTE* which measures relatively large background count rates compared to the source fluxes observed from AGN, it is important to remember that the noise level is equal to twice the total average count rate of the lightcurve *before* background subtraction has been carried out.

The red-noise lightcurves of AGN and XRBs are inherently stochastic, such that a single lightcurve is unique and unpredictable yet at the same time it is constrained to be drawn from the distribution of possible lightcurves whose average properties are described by e.g., the power spectrum. Therefore, the power spectrum of a single continuous lightcurve (of say 100 ksec duration) is *not* the same as the average power spectrum of the X-ray source (produced from, 1000 lightcurves of 100 ksec duration, for example).

In fact for a stochastic process such as white or red noise, the measured values of the power spectrum $P_{\text{real}}(\nu)$, of a realisation of the process are randomly distributed about the mean power spectrum describing the process $P_{\text{proc}}(\nu)$, following a χ^2 distribution with two degrees of freedom and standard deviation $P_{\text{proc}}(\nu)$ (Timmer & König 1995). Thus the measured power spectrum

fluctuates wildly and must be binned up in order to estimate the mean power (equivalent to $P_{\text{proc}}(\nu)$) at a given frequency. One way to achieve this, which is commonly used where lightcurves oversample the range of timescales of interest (which is nearly always the case in XRB observations), is to measure the power spectrum for a large number of equal-length segments of the lightcurve, and then average $P_{\text{real}}(\nu)$ across all segments. The uncertainty in the resulting mean power at each frequency is then given by the standard error:

$$\Delta\bar{P}(\nu) = \sqrt{\frac{1}{n(n-1)} \sum_{i=1}^n (P_i(\nu) - \bar{P}(\nu))^2},$$

where n is the number of power spectra (i.e. lightcurve segments) averaged over. In order for $\bar{P}(\nu)$ and $\Delta\bar{P}(\nu)$ to be good estimates of the mean power and error on the mean power, especially in the low-frequency case of red-noise variability where the power diverges rapidly, n needs to be large, so that the distribution of points about $\bar{P}(\nu)$ is approximately Gaussian. Thus the number of lightcurve segments for a given lightcurve duration needs to be large (preferably 100 or greater), which reduces the duration of each segment and hence the frequency range sampled by the power spectrum.

An alternative approach to estimating the mean power and its error is to smooth the measured power spectrum by binning up measurements of power at adjacent frequencies. The wider the frequency bins used, the more points are averaged and hence a better mean power is obtained. The power-law shape power spectra of red-noise lightcurves are suitably binned in logarithmically spaced frequency bins. However, binning up the power spectrum in frequency suffers from the drawback that at the lowest frequencies, the frequency bins must be wide in order to average a large number of power spectral measurements, causing power-law power spectra to become distorted. Therefore binning in frequency effectively limits the frequency coverage of a power spectrum to high frequencies which are well sampled, in the same way that binning up power spectra from smaller lightcurve segments does.

One way to alleviate the problems of estimating accurate mean powers and

their errors, especially at low frequencies where small-number statistics may lead to large biases, is to bin up the logarithm of power, as opposed to the power. This method was suggested by Papadakis & Lawrence (1993), who showed that the distribution of measured red-noise power in a frequency bin in logarithmic power space tends towards a Gaussian distribution with fewer points binned up than are required for points in linear power space. This method has the added advantage that biases due to fluctuations in power at low frequencies are naturally minimised by binning the logarithm of power. Therefore, throughout this work I shall bin up power in logarithmic space, using frequency binning or binning up power spectra made from lightcurve segments, or a combination of both. It is important to note that the average of the logarithms of power calculated in this way is not the same as the logarithm of the average power, which is the number we want and which we obtain by adding 0.253 to the average logarithmic power (Papadakis & Lawrence 1993). This correction is assumed throughout this work.

We can reasonably assume that we can obtain a good estimate of the power and its error by binning up power spectra with at least 10 points per bin. With a 100 ksec lightcurve, this means that we are restricted to looking at frequencies $> 10^{-4}$ Hz. As we are interested in measuring the power spectra of AGN over as broad a time-scale as possible, it may be necessary to take a different approach to estimating the uncertainty in the power at low frequencies, and hence placing confidence limits on any power-spectral models we may wish to fit. An additional, more serious problem which we face is that sampling that is not continuous (like our Seyfert monitoring data) leads to distortion in the power spectrum, which is not simple to predict analytically and which conventional methods of estimating power spectral errors does not take account of. Therefore, the approach I shall take will be to use Monte Carlo methods to make robust estimates of the power spectral shape of AGN across as broad a range of frequencies as possible, taking into account the distorting effects of sampling. I shall describe the method I will use in detail in the next chapter,

however in the remainder of this chapter, I shall consider some of the other issues relevant to estimating the shape of the broadband power spectrum. Let us first look at how we can simulate lightcurves from a given power spectral shape.

2.2 Simulating lightcurves

Throughout this work, I shall be using a method to simulate lightcurves suggested by Timmer & König (1995). Prior to their work, methods of simulating lightcurves based on a given power spectral shape were unsatisfactory. For example, it is well known that shot-noise models can reproduce red-noise power spectra of varying shapes (e.g. Lehto, 1989), so that lightcurves could be made directly from the specified shot-noise model using the appropriate shot parameters for the desired power-spectral shape. This approach is limited however, as it cannot be used to simulate lightcurves which are not red-noise or which have complicated shapes with more than three different power-spectral slopes. A more desirable method for simulating lightcurves should be free of any assumptions about the physical process which underlies the power-spectral shape. Done et al. (1990) suggested constructing a lightcurve $f(t)$ from a model power spectrum $P(\nu)$ by summing sine waves of frequency ν with randomised phases, $\phi(\nu)$ and amplitudes determined by the amplitude of the power spectrum at that frequency:

$$f(t) = \sum_{\nu} \sqrt{P(\nu)} \cos(2\pi\nu t - \phi(\nu)).$$

This procedure is flawed however, since it chooses a deterministic amplitude for each frequency (given by the underlying model power $P(\nu)$), unlike real stochastic lightcurves whose variance at a given frequency is randomly distributed about the mean for the underlying process (see previous section). The approach suggested by Timmer & König relies on this property of stochastic

lightcurves.

Consider the measured power spectrum, $P_{\text{real}}(\nu)$ of a realisation of some process with power spectrum $P_{\text{proc}}(\nu)$. As described in Section 2.1, $P_{\text{real}}(\nu)$ at different frequencies is randomly distributed about $P_{\text{proc}}(\nu)$, following a χ^2 distribution of standard deviation $P_{\text{proc}}(\nu)$. The Fourier transform $F_{\text{real}}(\nu)$ corresponding to $P_{\text{real}}(\nu)$ is a complex Gaussian random variable, with real and imaginary parts randomly distributed about zero with standard deviation $\sqrt{\frac{1}{2}P_{\text{proc}}(\nu)}$ (Timmer & König 1995). Therefore, an evenly sampled lightcurve of N data points and sampling interval ΔT can be generated by drawing two random numbers at each frequency $\nu_{\text{min}}, 2\nu_{\text{min}}, 3\nu_{\text{min}} \dots \nu_{\text{Nyq}}$ (where $\nu_{\text{min}} = N \Delta T$, $\nu_{\text{Nyq}} = (2 \Delta T)^{-1}$) and allocating them to the real and imaginary components of $F_{\text{real}}(\nu)$, and then taking its inverse Fourier transform. In this way, simply by inverting the procedure used to calculate a power spectrum from a lightcurve, it is possible to generate stochastic lightcurves from any given power spectral shape. Furthermore, since lightcurves are generated using the inverse Fourier transform, it is possible to generate lightcurves of length $N = 2^m$ (where m is an integer) rapidly from a specified power spectrum using the Fast Fourier Transform (e.g. Press et al. 1992). Some examples of lightcurves simulated with this technique, and the resulting (unbinned) power spectra are shown in Figure 2.1.

2.3 The effects of sampling on power-spectral shape

The simulated lightcurves shown in Figure 2.1 represent an ideal case, where the lightcurves are evenly sampled and the power spectra of the underlying processes are limited only to the frequency range between ν_{min} and ν_{Nyq} . In reality, the power spectrum of a real underlying process does not fall to zero above the Nyquist frequency and below the minimum frequency sampled.




Figure 2.1: Lightcurves (left) and power spectra (right) simulated using the method of Timmer & König, corresponding to model power spectra of power-law slopes $\alpha = 1$ (top), $\alpha = 2$ (centre) and a broken power-law model with high-frequency slope $\alpha = 2$ breaking to $\alpha = 0$ at frequency 10^{-4} (bottom). Note that, for comparison, all three simulated lightcurves are generated using the same random number sequence.

Variations due to power on longer and shorter time-scales than those sampled cannot be disentangled from the power spectrum, leading to distortions in the power spectral shape, which depend on the sampling pattern and duration of the lightcurve.

Consider a ‘real’ underlying lightcurve $r(t)$, on which we impose a sampling pattern $w(t)$ so that $w(t) = 1$ when we sample $r(t)$ and zero otherwise. The resulting observed lightcurve, $f(t)$ is given by:

$$f(t) = r(t) w(t).$$

According to the Convolution theorem of Fourier transforms (e.g. Press et al. 1992), the Fourier transform of two functions multiplied together in the time domain is equal to the convolution of the Fourier transforms of those functions in the frequency domain, so:

$$F(\nu) = R(\nu) * W(\nu),$$

where $R(\nu)$, $F(\nu)$ are the Fourier transforms of the underlying lightcurve and the observed lightcurve respectively and $W(\nu)$ is the Fourier transform of the sampling pattern, also known as the ‘window function’. The ideal power spectrum which we would hope to measure is given by $|R(\nu)|^2$ but what we actually measure is a power spectrum distorted by the window function, i.e. $|F(\nu)|^2$. The effect of the convolution is to smear power across the power spectrum so it is not possible to disentangle the true power at a given frequency from power which was transferred there from other frequencies. To do so requires a knowledge of the power spectral shape outside the frequency range sampled by $w(t)$, i.e. the measured power-spectral shape is model dependent.

The simplest examples of power-spectral distortion occur when a lightcurve is evenly sampled, because the pattern of distortion due to variations corresponding to frequencies above ν_{Nyq} or below ν_{min} is straightforward and easy to see. Variations on time-scales greater than the duration of the lightcurve contribute most to power at the lowest frequencies sampled. This effect is called ‘red-noise

leak'. Variations on timescales shorter than those sampled contribute an approximately constant amount across the entire power spectrum, leading to a flattening of the red-noise power spectrum at high frequencies. This effect is called 'aliasing'. I shall now describe these effects in more detail, and use simulated lightcurves to show how they distort the power spectrum and how that distortion is dependent on the true underlying power spectrum.

2.3.1 Red-noise leak

Consider a red-noise lightcurve, which we observe for a duration T . If the power-spectral slope does not flatten significantly at frequencies much lower than T^{-1} , the lightcurve will display significant long-term variations which will show up as approximately slow rising or falling trends in our observation. This trend contributes to the variance of the lightcurve we observe, and so must contribute to the power spectrum we measure. Now consider the form of the contribution a single trend makes to the power spectrum of an evenly sampled lightcurve. The variance of a section of the lightcurve of length T , is equal to the integrated power in the power spectrum, down to the frequency $\nu_{\min} = T^{-1}$. If we smooth out all variations other than the slow rising or falling trend, we can intuitively see that the variance due to the slow trend increases significantly as T increases, i.e. the integrated power increases strongly as ν decreases. Therefore the effect of red-noise leak is to contribute additional power to all frequencies, with much more power contributed at low frequencies than at high frequencies. In fact, red-noise leak contributes power with a power-law form and slope $\alpha = 2$. The amount of red-noise leak relative to the true power spectrum in the frequency range being observed is dependent on the amount of integrated power below ν_{\min} relative to the integrated power above ν_{\min} , so that steeper red-noise power spectra show stronger distortion due to red-noise leak.

I show the effects of red-noise leak in Figure 2.2, which shows the average power-spectral shape generated for 1000 simulated continuously sampled 100 ksec lightcurves (with 100 s time binning), and for comparison, the true underlying power spectra. The effects of red-noise leak on power spectral slopes of 1 or 1.5 are practically insignificant. The measured slope for $\alpha = 1$ remains unchanged, while the measured slope for $\alpha = 1.5$ is 1.54. For $\alpha = 2$, the slope remains unchanged but the power-spectral normalisation increases significantly, as expected from red-noise leak which contributes a slope of 2. For a particularly steep underlying slope, $\alpha = 2.5$ the normalisation of the measured power spectrum increases and its slope is significantly flattened to $\alpha = 2.14$. Note that the simulated 100 ksec observations were embedded in simulated lightcurves of 64 times longer duration, so that there was power only out to frequencies $\sim 1.6 \times 10^{-7}$ Hz. If the underlying power spectra were to remain significantly steep ($\alpha > 1$) below this frequency, the effects of red-noise leak would be greater than shown here.

2.3.2 Aliasing

Suppose that we sample a lightcurve discretely, at time intervals ΔT , and measure its power spectrum. We can only measure the power spectrum up to the Nyquist frequency $\nu_{\text{Nyq}} = (2\Delta T)^{-1}$, but if the underlying power spectrum contains significant power at higher frequencies, this will be reflected in the variance of the lightcurve and hence the measured power spectrum. The extra power at frequencies greater than ν_{Nyq} is transferred to lower frequencies in the power spectrum.

The way in which this power is redistributed or ‘aliased’ can be seen as follows. Consider an ‘ideal’ continuously sampled lightcurve which we use to make a power spectrum undistorted by any aliasing. We can restrict the frequency range of this ideal power spectrum to be the same as the ‘real-life’ one we




Figure 2.2: The effects of red-noise leak on power-spectral shape. The grey dashed line represents the true underlying power spectrum, while the solid black line shows the measured power spectrum after distortion by the limited sampling window. Underlying power spectral slopes from top to bottom are $\alpha = 2.5$, 2.0 , 1.5 and 1.0 respectively.

measure by smoothing the ideal lightcurve to remove variations on timescales less than ΔT , i.e. we make a continuous lightcurve with time bins of width ΔT . Therefore, with the same sampling pattern as the real lightcurve, *we can recover the ideal power spectrum if each sampled data point t_i represents the mean flux of the underlying ideal lightcurve in a bin of width ΔT centred at t_i .* For a discretely sampled lightcurve the above condition will be satisfied if there is no power in the underlying lightcurve at frequencies greater than ν_{Nyq} . However, in reality it is highly likely that there will be power at frequencies greater than ν_{Nyq} . This means that the flux measured at discretely sampled data points will not be the same as the mean flux for bins of width ΔT centred on those points. There will be an intrinsic error for each point, the magnitude of which depends on how much power there is above ν_{Nyq} . The effect on the measured power spectrum is then analogous to the contribution of Poisson noise to the power spectrum - a constant level of power is introduced which serves to flatten the power spectrum at high frequencies. However, unlike the Poisson noise contribution, this aliased component of power is dependent on the total amount of power above the Nyquist frequency, i.e. it is dependent on the power-spectral shape.

The effects of aliasing on power spectra of different slopes is shown in Figure 2.3. The average aliased power spectra were constructed from 1000 simulated lightcurves, each consisting of 200 data points sampled with a time interval $\Delta T = 100$ ksec (with each data point corresponding to a 1 ksec observation) . The power spectrum was cut off above 10^{-8} Hz, so that there is no significant red-noise leak contribution to the power spectra. The 1 ksec time bin size for each data point t_{bin} , limits the aliased power to be due only to power at frequencies $\nu < (2t_{\text{bin}})^{-1}$ (i.e. $\nu < 5 \times 10^{-4}$ Hz), since power above this frequency is removed because shorter-timescale variations are smoothed by the binning. The slopes measured by naively fitting a power-law to the observed average power spectra were 1.78, 1.2 and 0.54 for underlying slopes $\alpha = 2.0, 1.5$ and 1.0 respectively. Note that the flattening of the power-spectral

slope is more pronounced for lower underlying slopes, as expected. In fact, the constant additive component due to aliasing can be well approximated by a constant level such that its integral over the frequency range sampled by the observations is equal to the integral of the power above the Nyquist frequency (so that the total lightcurve variance measured by the power spectrum is conserved), in other words the constant power level P_C for the underlying power spectrum $P(\nu)$ is given by:

$$P_C = \frac{1}{\nu_{\text{Nyq}} - \nu_{\text{min}}} \int_{\nu_{\text{Nyq}}}^{(2t_{\text{bin}})^{-1}} P(\nu) d\nu.$$

Note however, that this approximation cannot account for stochastic variations in the amount of aliasing which apply when power spectra are made from relatively small numbers of data points (roughly a few hundred data points or less).

Binning up a discretely sampled lightcurve can reduce the effects of aliasing, since averaging fluxes measured at several points in a time bin yields a better estimate of the true mean flux in that time bin. Unfortunately, binning points into larger time bins increases the minimum time-scale sampled and hence reduces the frequency range sampled. Therefore it is necessary to over-sample a lightcurve, i.e. to sample on timescales shorter than those required, so that aliasing may be reduced by binning. To test the effect of binning, I increased the sampling rate of the $\alpha = 1$ lightcurves described above by factors of 2 and 4 (i.e. 50 ksec and 25 ksec sampling interval respectively), and made average power spectra after binning them to the 100 ksec sampling interval. The measured slopes were 0.74 and 0.88 respectively, so that even higher sampling rates would be required to reduce the intrinsic flattening in the slope to less than 0.05.




Figure 2.3: The effects of aliasing on power-spectral shape. The grey dashed line represents the true underlying power spectrum, while the solid black line shows the measured power spectrum after distortion by the effects of discrete sampling. Underlying power spectral slopes from top to bottom are $\alpha = 2.0$, 1.5 and 1.0 respectively.

2.3.3 The effects of uneven sampling

The considerations of the effects of red-noise leak and aliasing given above were restricted to the ideal case of evenly sampled lightcurves because the form that power spectral distortion takes is easiest to understand in this limit. For real monitoring programs, scheduling constraints make it difficult to obtain perfectly evenly sampled lightcurves. In the case of our Seyfert monitoring program, lightcurves are approximately evenly sampled in that the typical error from the requested sampling interval is $\sim 20\%$. The actual intervals between observations are neither systematically larger nor smaller than the requested sampling interval, so that the average sampling interval is close to that which was requested. Clearly the distortion due to red-noise leak, which is dependent on the total duration of the monitoring lightcurve, is not significantly affected by these deviations from even sampling, provided they are not very large, i.e. much less than the lightcurve duration itself. The main difference in the power spectral distortion from the case of even sampling is due to the effects of aliasing.

As described in the previous section, the effect of aliasing can be thought of as being due to the deviation in the sampled flux from the true mean flux in the bin of width equal to the sampling interval, which the sampled data point represents. The bin width in this case is the average sampling interval, i.e. it is the same as that in the evenly sampled case, however the sampled observations do not come from the centre of each bin, but are randomly scattered around the centre of each bin. First consider a white noise lightcurve which is unevenly sampled. Because there are no long-term trends comparable to the sampling interval, the average deviation of the sampled flux from the mean flux of the bin it represents is independent of where the observed data point is measured in the bin, so the effect of aliasing is not changed from the even sampling case. However, as we consider steeper power-spectral slopes, long-term trends appear in the lightcurve which mean that the average deviation in the

sampled flux from the mean flux of the bin is dependent on its position in the bin. On average, slow trends within a bin will mean that sampled data points which are further from the centre of the bin will contribute more to aliasing than those closer to the centre, i.e. there will be a greater amount of aliasing than in the evenly sampled case. Note that this additional distorting effect of uneven sampling is dependent on the power spectral shape, since the effect will be larger when long-timescale trends are larger, i.e. the power spectrum is steeper.

A different form of uneven sampling comes from lightcurves consisting of continuous segments of data which are broken up by gaps comparable to or larger than the segment length. If the gaps are roughly periodic, perhaps due to Earth-occultation effects, power is introduced at frequencies corresponding to the period of the gaps. Lightcurves which are very poorly sampled can show a mixture of all the effects described above, including red-noise leak. It is possible to determine the average power spectral shape for a given power-spectral model and window function analytically, but this is not trivial and becomes almost impossible for complex power spectral shapes (I. Papadakis, private communication). Additionally, although an analytical determination of the distorted power spectrum can give the average power spectrum from an infinite number of realisations of the given power-spectral model, convolved with the given window function, it does not give any estimate of the uncertainties on the measured power spectrum due to its stochastic nature. Therefore the simplest way to estimate a distorted power spectrum and its uncertainty is to use Monte-Carlo techniques, which are described in the next chapter.

2.3.4 The effects of binning

Finally, we should note that binning up a continuous lightcurve into time bins of width t_{bin} has the effect of steepening the power spectrum by a frequency

dependent factor $|B(\nu)|^2$ (van der Klis 1989), given by

$$|B(\nu)|^2 = \left(\frac{\sin \pi \nu t_{\text{bin}}}{\pi \nu t_{\text{bin}}} \right)^2.$$

This is due to the fact that binning has the effect of smoothing lightcurves on timescales similar to the binning timescale.

In the case of the AGN power spectra presented in this work, binning effects are not significant because either the data is not continuously sampled (i.e. in the case of monitoring lightcurves) or because the power spectrum at high frequencies is dominated by Poisson noise (in the case of long-look data).

2.4 Stationarity of the power spectrum

When we measure the power spectrum of a lightcurve, the amplitude of the power spectrum is dependent on the measured variance of the lightcurve. Thus, if two X-ray sources show the same pattern of X-ray variability (i.e. they show the same amplitude of variability relative to their mean flux), but they have differing mean fluxes then the fainter X-ray source will show a lower amplitude power spectrum than the brighter one. It is desirable then to normalise the X-ray lightcurves of different X-ray sources by their mean fluxes in order that we can directly compare the amplitudes of their power spectra. Equivalently, we can divide the power spectrum of each lightcurve by the square of the mean flux of the lightcurve. Under the RMS squared normalisation we are using, this means that the integrated power or variance σ^2 becomes the fractional variance,

$$\sigma_{\text{frac}}^2 = \frac{\sigma^2}{\mu^2}$$

where μ^2 is the mean flux, so that the square root of the integrated power, after normalising by the squared mean flux, is equal to the contribution to the fractional RMS made by variations on timescales corresponding to the frequency range integrated over.

Normalising the power spectrum by the squared mean flux of the lightcurve, although highly desirable, raises an important question: *what value of the mean should we use?* At first sight, the answer may seem simple, i.e. use the measured lightcurve mean. This approach would be suitable for a white-noise lightcurve, where the mean measured over a given time interval at a time T_1 is the same as that measured at a different time T_2 and is independent of the duration of the time interval. However, for a red-noise lightcurve, long term trends, due to power at frequencies corresponding to time-scales greater than the observation duration, lead to different means being measured at different times - the lightcurve is non-stationary.

This problem is most acute if the power spectrum is particularly steep. However, as discussed in Section 1.2.2, the power spectra of AGN must eventually flatten towards low frequencies so the lightcurve effectively becomes stationary on long time-scales. If we then measure the mean of the lightcurve on as long a time-scale as possible (i.e. a few years in the case of our monitoring lightcurves), that should be a good representation of the ‘true’ mean of the lightcurve. We can then normalise the power spectra by the square of this mean, rather than the local mean measured for a given observation.

As discussed in Section 1.3.2, in order to study the power spectrum of AGN over a broad frequency range without requiring excessive amounts of observing time, it is necessary to collect monitoring data at a number of different sampling rates. For the targets we are studying, we have obtained long exposures to study short time-scale variability, and snapshot observations twice-daily, daily or every two weeks to study medium to long time-scales. The different sets of observations are made at different times and for different durations, so they have different means. Therefore it seems right that in order to make a broadband power spectrum including separate power spectra made for each sampling pattern, we should normalise each separate power spectrum by the mean of the lightcurve measured over the longest duration, i.e. the lightcurve made of snapshots obtained every two weeks. However, we have made an

unspoken assumption throughout this discussion: *that the power spectrum is itself stationary and does not change in shape or amplitude.*

In fact, the assumption that the power spectrum is stationary, at least on relatively short time-scales, underpins all work on power spectra of XRB systems and AGN. If the power spectrum of an X-ray source were not stationary, but varied in an unpredictable way, then the power spectra obtained at any given time would not be a good representation of the average variability properties of that X-ray source. Power spectra of X-ray binaries are known to be non-stationary on relatively long time-scales, with Cygnus X-1 showing variations in the position of its low-frequency and high-frequency breaks and high-frequency power-spectral slope on time-scales of hours, although the slope and normalisation of the power spectrum between the breaks does appear to be stationary (Belloni & Hasinger 1990). The galactic microquasar GRS 1915+105 shows wild variations in its variability properties on timescales of minutes (Mirabel et al. 1998). It is thought that this non-stationarity in XRB power spectra is due to changes in the structure, inner radius or the variability mode of the accretion disk, since the time-scales for power-spectral variability are comparable to viscous time-scales in a disk around a $\sim 10 M_{\odot}$ black hole. In AGN the time-scales for similar variations must be scaled up by the black hole mass, so that variations in the shape of AGN power spectra will occur on time-scales of the order of a decade or more - longer than we sample here.

It would seem that we need not worry about long-term changes in power-spectral shape analogous to those seen in XRB systems, however we must consider the possibility that the power spectrum of AGN is intrinsically non-stationary on the time-scales which we sample here. Previous work by Green (1993) using two *ROSAT* lightcurves of NGC 4051 obtained a year apart, showed that the RMS variability seemed to be linearly correlated with the mean flux of each lightcurve, implying that the power spectrum is indeed non-stationary, in that its amplitude scales with lightcurve mean flux. However, as

we shall see in chapter 5, NGC 4051 shows evidence of unusual non-stationarity in its power spectrum on time-scales of months, unlike the other targets we have monitored. Therefore we should check that NGC 4051 is not a special case, by looking at how the variance (and hence power-spectral amplitude) of all our targets scales with flux on short timescales.

2.5 The flux dependence of AGN variability

We would like to know whether the variance, and hence RMS variability and power-spectral amplitude of AGN is dependent on flux on the short time-scales which are well sampled by *RXTE* long-look observations. To this end, I first made 16 s-binned 2–10 keV lightcurves for the long-look observations of NGC 4051, NGC 5506 and MCG-6-30-15, which all vary significantly on short time-scales. Details of these observations and data reduction can be found in Sections 4.1 and 5.1.1. The *RXTE* lightcurves are naturally split into continuous segments of ~ 2500 s duration, corresponding to the time between Earth occultations of the target AGN. These continuous segments provide a natural duration to measure the power spectrum over and thus estimate the integrated power of the lightcurve segment for comparison with its mean flux. I have chosen to use integrated power as a measure of variance rather than simply measuring the variance of each segment in the time domain, because variance is dependent on the duration of each segment, which varies between ~ 2000 and 3000 s, whereas the integrated power in a specified frequency range is independent of segment duration. I have chosen to integrate the power spectrum of each segment in the frequency range 5×10^{-4} – 5×10^{-3} Hz, since above 5×10^{-3} Hz the power is dominated by the Poisson level.

Due to the intrinsic stochastic nature of the lightcurve, as described in Section 2.1, the integrated power in different segments varies naturally about some mean which is the integrated power for the average power spectrum

which describes the underlying process. Therefore there is significant scatter in the integrated power, irrespective of mean flux, and it is necessary to bin up the measured integrated powers of each segment according to the segment mean flux. To this end I binned up the variances of each segment according to whether the segment mean flux lay below or above the mean flux of the entire lightcurve. The standard error in the resulting mean integrated powers (henceforth mean variances $\overline{\sigma^2}$) was calculated in the standard manner using the scatter of segment variances (σ_i^2) about the mean:

$$\Delta\overline{\sigma^2} = \sqrt{\frac{1}{n(n-1)} \sum_i^{i=n} (\sigma_i^2 - \overline{\sigma^2})^2},$$

where n is the number of segments used to calculate the mean variance in that flux bin.

The mean variance incorporates the variance due to Poisson noise, so the true mean variances of the source in each flux bin may be obtained from the measured mean variances by subtracting off the average Poisson noise level in the power spectrum integrated over the 5×10^{-4} – 5×10^{-3} Hz frequency range. The resulting values of noise-subtracted mean variance for the low and high flux bins for all three AGN are shown in Table 2.1, along with the fractional RMS variability measured in each flux bin. These data clearly show that the variance of the X-ray lightcurves of all three AGN is dependent on X-ray flux, and furthermore the relationship between the RMS variability and the mean segment flux seems to be linear, such that the fractional RMS variability remains approximately constant, despite significant flux changes (e.g. \sim factor 3 for NGC 4051). Therefore, the X-ray variability power spectrum of AGN measured on short time-scales, is not stationary, since the amplitude of the power spectrum scales with the square of local mean flux. However, *the power spectrum normalised by the local squared mean flux is effectively stationary, since the fractional RMS variability is independent of local mean flux.*

Having seen that the amplitude of the power spectrum measured on short time-scales scales with the square of the local mean flux, we would like to

	μ_{tot}	Low flux				High flux			
		μ	n	$\overline{\sigma^2}$	σ_{frac}	μ	n	$\overline{\sigma^2}$	σ_{frac}
NGC 4051	3.4	1.9	17	0.047 ± 0.024	11.4 %	5.3	16	0.32 ± 0.07	10.7 %
NGC 5506	28.1	25.5	20	0.21 ± 0.05	1.8 %	31.2	19	0.41 ± 0.11	2.1 %
MCG-6-30-15	12.2	10.3	73	0.28 ± 0.05	5.1 %	14.1	70	0.56 ± 0.07	5.3 %

Table 2.1: Flux-related changes in the variance of NGC 4051, NGC 5506 and MCG-6-30-15. Mean data are shown for low and high flux segments (corresponding to segment fluxes below and above the mean for the entire observation, given by μ_{tot}). μ is the mean flux (2–10 keV, count s⁻¹) for each flux bin, n is the number of segments in the flux bin, $\overline{\sigma^2}$ is the mean variance (count² s⁻²) and σ_{frac} is the fractional RMS variability ($\sigma_{\text{frac}} = (\overline{\sigma^2}/\mu^2)^{\frac{1}{2}}$).

know if this same pattern extends to variations on longer timescales, i.e. does the average amplitude of a power spectrum measured in any duration ΔT scale with the square of the mean flux measured in ΔT , while the power-spectral shape remains the same? Unfortunately we do not yet have sufficient monitoring data on long time-scales to test this for AGN, however we can look at the X-ray lightcurves of X-ray binary systems to see if a similar pattern emerges, investigate its dependence on frequency and then by analogy extend these results to the long time-scale variability of AGN.

2.6 The flux dependence of XRB variability

Power-spectral studies of X-ray binary systems measure the power spectrum of segments of the XRB lightcurves (~ 1 ksec or more) which are at least ten times longer than the maximum time-scales sampled by the power spectrum. The power spectra obtained in this way are very well defined, with high S/N , but since XRB power spectra flatten below frequencies of ~ 0.1 –1 Hz, the

mean flux does not vary significantly on long time-scales so that any power-spectral dependence on flux would be very difficult to detect in these studies. Furthermore, because it is common practice to normalise XRB power spectra by the squared mean flux of the lightcurves used to make them, any linear dependence of power spectral amplitude on flux would be missed completely. In order to search for the kind of non-stationarity we see in AGN power spectra at high frequencies, it is necessary to examine the flux dependence of XRB power spectra that are measured on short time-scales, so that a significant range of mean segment fluxes can be sampled. With this in mind, I obtained public archival data for *RXTE* observations of Cygnus X-1 (in the low state), observed 23 October 1996 (Nowak et al. 1999) and the recently discovered accreting millisecond pulsar, SAX J1808.4-3658 (Wijnands & van der Klis 1998), observed 18 April 1998, of duration 18 ksec and 23 ksec respectively. Cygnus X-1 is an obvious choice for study, as it is the most well-known black hole X-ray binary. The accreting millisecond pulsar SAX J1808.4-3658 is a good example of an accreting system with a red-noise power spectrum which is certain to contain a neutron star, so any similarity in the variations of the power-spectra of these objects will provide strong evidence that the power-spectral variability is intrinsic to the accretion process, independent of whether the central object has a hard surface or event horizon.

Using PCA binned and event mode data from the observations of Cygnus X-1 and SAX J1808.4-3658 respectively, I made 16 ms binned lightcurves for each source in the 0–13.1 keV energy range (it was not possible to make lightcurves in the 2–10 keV range due to the limited spectral resolution of the binned mode data). The mean fluxes in this energy range were 3377 count s⁻¹ and 409 count s⁻¹ for Cygnus X-1 and SAX J1808.4-3658 respectively. I then split each lightcurve into 10 s segments, and determined the noise-subtracted integrated power in each segment over the 0.1–10 Hz frequency range. Taking the square root of the integrated power yields the RMS variability σ due to variations in the 0.1–10 Hz frequency range. Due to the good data quality

and large number of 10 s segments measured in each lightcurve, it is possible to bin up σ as a function of flux using relatively narrow flux bins to examine the exact form of the σ -flux relationship over a broad range of flux. Using a minimum number of 30 measurements of per flux bin to ensure an accurate estimate of the standard error in the mean σ , the resulting σ -flux relation for both sources is plotted in Figure 2.4.

The dependence of mean σ on flux is remarkably linear in both Cygnus X-1 and SAX J1808.4-3658, implying that the type of non-stationarity seen on short timescales in AGN lightcurves is also present (and much more significantly detected) in X-ray binary lightcurves, independent of whether the central object is a black hole or neutron star. The flux dependence of power-spectral amplitude must be intrinsic to the accretion process, and operates over at least 6 orders of magnitude in central object mass.

Visual inspection of Figure 2.4 shows that the σ -flux trend in both sources does not pass through the origin, i.e. there is a constant offset in flux and possibly also σ . To test the goodness of fit of a linear model, including a constant offset from the origin, I fitted a function of the form $\sigma = k(F - C)$, where F is the flux and k and C are constants. This linear model provides a good fit to the data, yielding χ^2 values of 28.7 (for 22 degrees of freedom) for Cygnus X-1 and 17.2 (20 degrees of freedom) for SAX J1808.4-3658. The best fitting model parameters were $k = 0.326 \pm 0.017$, $C = 850 \pm 130 \text{ count s}^{-1}$ for Cygnus X-1 and $k = 0.305 \pm 0.026$, $C = 81 \pm 30 \text{ count s}^{-1}$ for SAX J1808.4-3658 (uncertainties are at the 90% confidence level for 2 interesting parameters). Note that the gradient of the σ - F trend k , is equivalent to the fractional RMS variability of the variable- σ component of the lightcurve. The constant C represents a second component to the lightcurve, which does not follow the linear σ - F trend but may still contribute to the total value of σ . It is not possible to disentangle the mean flux level of this second, constant- σ component and the contribution it makes to the total σ , however the value of C represents the flux of this component in the limit where it does not vary.




Figure 2.4: Local flux dependence of mean σ for SAX J1808.4-3658 and Cygnus X-1. The dotted lines mark the best-fitting linear models described in the text.

The X-ray power spectra of accreting compact objects are clearly not stationary, and depend in a very specific manner on the local X-ray flux. But how does this non-stationarity depend on the power-spectral frequency? Does only the amplitude of the power-spectrum change, or does its shape change also? To investigate the dependence of power-spectral shape on flux for Cygnus X-1, I made two flux-binned power spectra, by binning up the power spectra made from 100 s segments according to whether the mean X-ray flux of the segment lay above or below the mean flux of the entire observation. The resulting power spectra are shown in Figure 2.5. Although the flux variations in Cygnus X-1 on time-scales of ~ 100 s are not large, the quality of the power spectra made in this way is easily sufficient to see the difference made by a relatively small change in flux. Clearly the power spectral shape remains the same, while its normalisation scales with squared mean flux over at least 3 decades of frequency.

Having shown that the normalisation of the power spectrum of the black hole X-ray binary Cygnus X-1 scales with the *local* squared mean flux across all measured frequencies (after taking account of the constant- σ component to the lightcurve) we can assume that the power spectra of AGN show the same behaviour. Therefore we should not normalise power spectra by the mean of the lightcurve on long time-scales - we should instead normalise by the local mean, i.e. the mean of the lightcurve used to make that particular section of the broadband power spectrum. One outstanding question is whether AGN show evidence for a constant- σ component to their lightcurves. If σ is directly proportional to the local mean flux of the variable- σ component of the lightcurve (i.e. any second component to the lightcurve is constant and does not contribute to σ), then the constant flux component C is given by:

$$C = \mu_1 - \frac{\mu_2 - \mu_1}{\frac{\sigma_2}{\sigma_1} - 1},$$

where μ_1 and μ_2 are the low and high mean flux levels respectively, with RMS variabilities σ_1 and σ_2 . Using the 68% confidence upper limit to σ_2 and the lower limit to σ_1 (inferred from the uncertainties in σ^2 given in Table 2.1) we

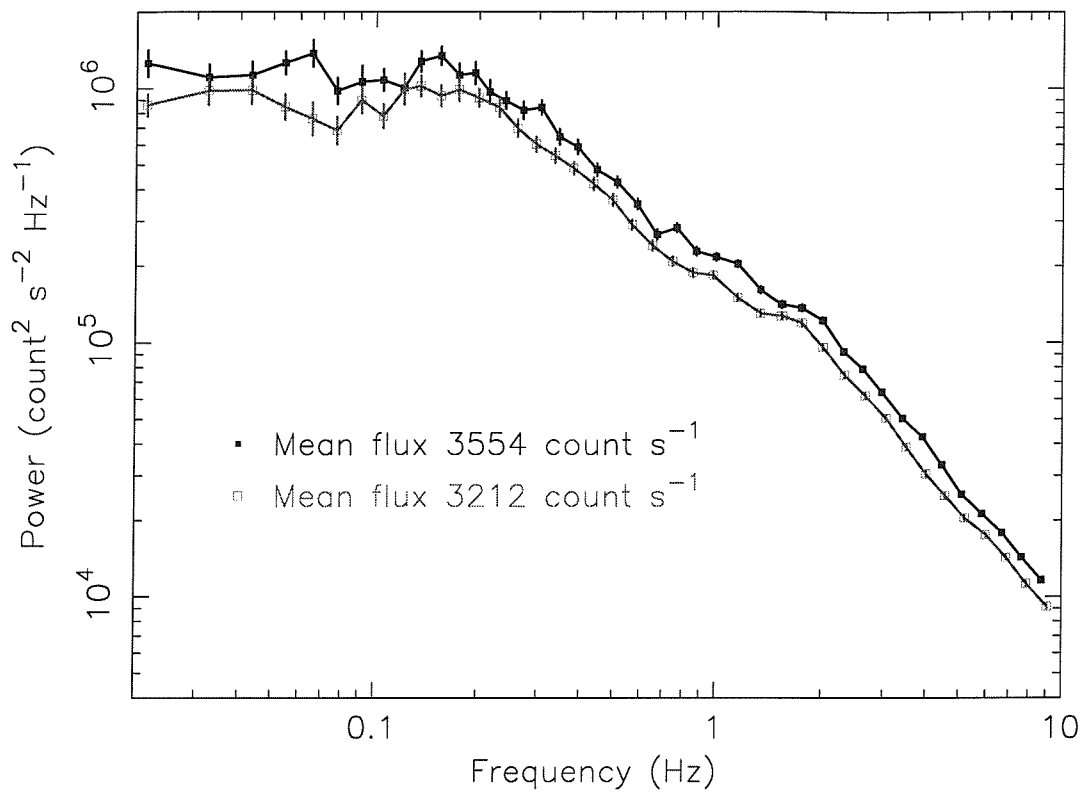


Figure 2.5: Local flux dependence of the power spectrum of Cygnus X-1. The black line marks the power spectrum for segments of mean flux greater than the total mean of the entire observation, the dark grey line marks the mean power spectrum for segments of mean flux less than the total mean. The corresponding combined mean fluxes for all segments used to make each power spectrum are also shown in the figure.

can set a 90% confidence upper limit on C of 0.82 count s⁻¹, 18.2 count s⁻¹ and 4.5 count s⁻¹ for NGC 4051, NGC 5506 and MCG-6-30-15 respectively. No lower limits can be set. It is interesting to compare these upper limits as a fraction of total mean lightcurve flux with the values of C measured for the XRB systems in the same limit where C represents a constant component to the lightcurve. The fractional values of C , relative to mean lightcurve flux, for Cygnus X-1 and SAX J1808.4-3658 are 0.25 and 0.2 respectively, compared to upper limits of 0.24, 0.65 and 0.37. The AGN lightcurves are consistent with having a second component like those of X-ray binaries, but the low upper limit on C for NGC 4051 implies that this source at least may differ from the X-ray binaries I have considered, in having a smaller (or non-existent) constant- σ component to its lightcurve.

2.7 Summary

In this chapter, we have seen how we measure the power spectrum from observed lightcurves, and discussed the conventions that will be used to measure power spectra throughout this work, namely that we shall use the RMS squared normalisation and bin the measured power spectra logarithmically. We have seen how we can simulate lightcurves and how we can use these simulations to estimate the distorting effects of the lightcurve sampling pattern on the observed power spectrum, such that it is not a fair representation of the true underlying power spectrum. The distorting effects of sampling are dependent on the underlying model power spectrum, so that in general, it is necessary to use simulated lightcurves to estimate the distortion on the power spectrum and hence compare the observed power spectrum with what we expect from a given model power spectrum. In the next chapter, we shall investigate this concept further and develop a technique, PSRESP, which can make robust estimates of

the probability that a given model is capable of describing the observed power spectrum.

We have also seen that, quite unexpectedly, the power-spectra of lightcurves of three AGN and the X-ray binaries Cygnus X-1 and SAX J1808.4-3658 are intrinsically non-stationary in that the mean RMS variability is linearly dependent on flux. The power spectrum of Cygnus X-1 shows that this linear dependence results from the power-spectral shape staying the same while the power-spectral amplitude scales with the square of local mean flux. I will discuss the implications of this remarkable dependence of power-spectral amplitude on local mean flux in the concluding chapter of this work. For now, it is important that we take with us the lesson that, despite our preconceptions, we must always normalise power spectra by the local mean flux in order to compare them.

Chapter 3

PSRESP: A robust method for measuring power spectral shape

Overview

This chapter describes the method we shall be using to estimate the model-dependent, underlying power-spectral shape of lightcurves whose directly-measured power spectra are distorted by sampling effects. I shall first describe the philosophy underlying the basic technique, called the ‘response method’, and then show how this technique can be applied using Monte Carlo methods to estimate reliable probabilities that a given underlying power-spectral model is acceptable for describing the data. The specific implementation of this method, which I have called PSRESP is then given along with caveats for its use.

In Chapter 2, we saw how the intrinsic stochastic variability of red-noise lightcurves leads to uncertainties in the measured power-spectral shape. Furthermore, sampling the lightcurves at discrete intervals for limited durations distorts the measured power spectral-shape away from that of the true power spectrum. These effects are dependent on the underlying power-spectral shape

and the sampling pattern of the lightcurve and may only be neglected if a lightcurve is continuously sampled (so that aliasing effects can be ignored) for a time duration much longer than the corresponding minimum frequency being sampled by the power spectrum (so that red-noise leak is minimised and estimated uncertainties are reliable). Unfortunately, due to observing constraints and the long time-scales which AGN vary on, neither of these criteria are satisfied by our (or any other) AGN monitoring campaign. It is therefore necessary to develop a robust technique for estimating the true power-spectral shape (and its uncertainty) of unevenly sampled lightcurves, in order that we can confidently study the broadband power-spectral shape of AGN and compare our findings with the power spectra characteristic of X-ray binary systems. The method I have developed is called PSRESP, short for Power Spectral RESPONSE, and is based on the ‘response method’ described by Done et al. (1992).

3.1 The response method concept

As described in Section 2.3, the power spectral distortion caused by a lightcurve sampling pattern can be understood as being due to a convolution of the Fourier transform of the true underlying process with that of the sampling window (the ‘window function’). The response method (Done et al. 1992) is based on an analogy with the problem of spectral response encountered in X-ray spectroscopy. In the process of measurement of an X-ray energy spectrum, the true underlying X-ray spectrum is distorted by the limited energy resolution and instrumental effects of the X-ray telescope. The effect is to convolve the true X-ray spectrum with a ‘spectral response function’ to yield the distorted spectrum which is actually measured. The convolution smears the energy spectrum so that the measured X-ray spectrum is model dependent, hence models are fitted to the data by convolving an underlying model ‘guess’ with the best estimate of the telescope response function and then comparing

the resultant model measured spectrum with the real data. The errors in the data are primarily due to simple counting statistics, so they are not distorted by the response function (since they are not intrinsic to the underlying spectrum), and may thus be used to estimate a χ^2 value for the model to determine a robust probability that the model is a good fit to the data. In this way, the best-fitting model spectrum can be found by χ^2 minimisation while stepping through a given range of model parameters. The veracity of this technique has been shown by Lampton, Margon & Bowyer (1976) and more recently, Yaqoob (1998).

In the power spectral case, we have a measured power spectrum which is distorted away from the underlying power spectrum by the effects of sampling (the window function is analogous to the response function). Therefore it would seem that we need to determine the best-fitting model power spectrum by stepping through a range of power-spectral parameters to generate model power spectra, somehow convolving them with the window function to make distorted model power spectra, and then estimating a chi-squared value of the distorted model power spectrum by comparing it with the measured power spectrum.

Unfortunately however, this is where the analogy with X-ray spectroscopy ends. The power spectrum, $P_{\text{real}}(\nu)$ of the real lightcurve which underlies the lightcurve which we sample is assumed to be a stochastic realisation of a process defined by the model power spectrum $P_{\text{mod}}(\nu)$ which we are seeking to fit (see Section 2.1). The power of any realisation of the model process is distributed randomly (with standard deviation $P_{\text{mod}}(\nu)$, Timmer & König 1995) about the model power spectrum, contributing to an intrinsic uncertainty in $P_{\text{mod}}(\nu)$. Hence the distribution of power which is otherwise used to estimate the uncertainty on $P_{\text{mod}}(\nu)$ is itself distorted, so that the method of error estimation described in Section 2.1 cannot be applied to power spectra distorted by sampling. To deal with this problem, Done et al. (1992) and later Green, McHardy & Done (1999) determined the distorted power spectrum for a given

underlying model by averaging a large number of power spectra made from simulated lightcurves with the same sampling as the data, using the RMS error on the model average power at each frequency (calculated from the spread of simulated powers at that frequency) as an estimate of the error in the power. This RMS error is then used to calculate the χ^2 for the model compared with the data. We cannot estimate errors for the observed power spectrum, because it is distorted by sampling effects and hence is model dependent. However, since the definition of the χ^2 parameter is based on the variance of the model population from which the data is drawn, the approach of assigning an error to the model is technically more correct than assigning an error to the measured data.

Using this technique, Done et al., Green, McHardy & Done and also recently Leighly (1999) determined χ^2 for various power-spectral model fits to long-look observations of AGN, and hence estimated their best-fitting power-spectral parameters. However, all these investigations suffer from the problem that the χ^2 parameter estimated from comparing distorted model average power spectra with measured power spectra, χ_{dist}^2 is not the same as the χ^2 which follows from the well-known statistical distribution. This is because the ensemble of simulated powers used to define the RMS error on the model average power is not normally distributed, due to the distorting effects of sampling. Although the χ_{dist}^2 will give an estimate of the best-fitting set of model parameters, reliable confidences cannot be placed on these parameters by assuming the usual χ^2 distribution.

It is therefore necessary to determine the distribution of χ_{dist}^2 directly from the simulated power spectra, by determining a χ_{dist}^2 for each simulated power spectrum. Then, the percentile of this distribution which the observed value of χ_{dist}^2 lies in gives the probability that the observed data can be described as a realisation of the given model parameters, i.e. the null hypothesis is satisfied. For example, if only 1% of simulated power spectra have a greater χ_{dist}^2 than the observed power spectrum, there must only be a 1% probability that the

observed data is a realisation of the model parameters, i.e. the given model parameters can be ruled out at 99% confidence. This technique of using Monte Carlo simulation of synthetic data sets to set confidence limits on estimated model parameters is described in more detail by Press et al. (1992, section 15.6). I have applied this technique to the determination of reliable confidence estimates for power spectral models fitted to AGN monitoring data in the PSRESP software, which I describe in more detail below.

3.2 PSRESP implementation

As described in Section 1.3.2, our AGN monitoring program maximises the frequency coverage of the power spectrum at a minimum cost in observing time, by observing our targets using several observing schemes to measure variability on different time-scales. Long-look observations, obtained as part of our program or from the *RXTE* public archive, are used to measure power at the highest frequencies ($> 10^{-5}$ Hz), while daily and twice-daily observations measure the medium-frequency power spectrum (10^{-6} – 10^{-5} Hz) and observations every two weeks extend the power spectrum to low frequencies (10^{-8} – 10^{-6} Hz). Assuming that the power spectral shape is stationary on time-scales of a few years (as we would expect if time-scales for changes in power-spectral shape scale with black hole mass), and given the discovery that power spectral normalisation scales with local mean flux (as reported in Chapter 2), then we can obtain a good measure of the broadband power spectrum of one of our targets by fitting the power spectra measured from each observing scheme (normalised by the local squared mean) with a single power-spectral model.

For targets which vary significantly on time-scales of less than ~ 1 ksec, it is desirable to make very-high-frequency (VHF) power spectra which are undistorted by aliasing, by using continuous data segments from long-look observa-

tions to make separate power spectra and then averaging them (so these VHF power spectra cover the frequency range $> 4 \times 10^{-4}$ Hz). The standard errors on these power spectra can be determined using the conventional method described in Section 2.1. The number of continuous segments and hence individual power spectra used to make the VHF power spectra will be large (~ 40 or greater) and hence their errors will be reliable. We can then use the VHF power spectra to tie down the power-spectral shape at short timescales, without needing to simulate lightcurves with high time resolution (16 s), which would be very time consuming. Instead, the VHF model average power spectrum can be the undistorted power-spectral model, calculated directly from the model parameters (however, see Section 3.3 for how to deal with red-noise leak).

The procedure for running PSRESP is then as follows: lightcurves are made for each observing scheme and are input together with details of their local mean fluxes, Poisson noise levels, bin sizes (1000 s for monitoring snapshots) and the lightcurve bin size to be used for making power spectra (e.g. 86400 s for the daily observing scheme) spectra. Also, details of how to bin up the resulting power spectra are included. Next, if there is significant power at high frequencies, the VHF power spectrum (which has already been made) is input, along with its Poisson noise level and corresponding local mean flux for normalisation purposes.

The basic method used by PSRESP is then summarised as follows.

1. First read each observed lightcurve, $x_i(t)$, where i denotes the lightcurve and $i = 1$ to n , the total number of lightcurves used to measure the shape of the power spectrum. Measure its sampling pattern, $s_i(t)$ and determine the power spectrum $P_{\text{obs},i}(\nu)$ normalised by the square of the mean flux of $x_i(t)$, binning the power spectrum up as desired. Also read in any previously measured, VHF power spectrum, $P_{\text{VHF}}(\nu)$.
2. Step through the range of power-spectral model parameters to be fitted.

For each set of model parameters, which define an underlying model power spectral shape, $P_{\text{mod}}(\nu)$ carry out the following:

- (a) For each $x_i(t)$ carry out the following:
 - i. Simulate N lightcurves which are realisations of $P_{\text{mod}}(\nu)$ and have sampling pattern $s_i(t)$.
 - ii. Make power spectra, $P_{\text{sim},j}(\nu)$ (where $j = 1, N$) from the simulated lightcurves, binned in the same way as $P_{\text{obs},i}(\nu)$.
 - iii. Average the $P_{\text{sim},j}(\nu)$ to get the model average distorted power spectrum, for the lightcurve, $\overline{P_{\text{sim}}}(\nu)$ and use the spread in $P_{\text{sim},j}(\nu)$ in each frequency bin to estimate the RMS error $\Delta\overline{P_{\text{sim}}}(\nu)$.
- (b) Define a statistic, which we call χ_{dist}^2 , which is calculated from the model average and observed power spectra of each lightcurve:

$$\chi_{\text{dist}}^2 = \sum_{\nu=\nu_{\text{min}}}^{\nu_{\text{max}}} \frac{(\overline{P_{\text{sim}}}(\nu) - P_{\text{obs}}(\nu))^2}{\Delta\overline{P_{\text{sim}}}(\nu)^2},$$

where ν_{min} and ν_{max} are respectively the minimum and maximum frequencies measured by $P_{\text{obs}}(\nu)$. Sum the χ_{dist}^2 determined for the power spectrum of each input lightcurve (combining with the ordinary χ^2 measured from the VHF power spectrum) to determine the total χ_{dist}^2 for the model relative to the observation. Find the correct normalisation of the model by renormalising all the $\overline{P_{\text{sim}}}(\nu)$ by the same factor k , varying k until the total χ_{sim}^2 is minimised.

- (c) Determine a reliable probability that the observed broadband power spectrum could be a realisation of the model power spectrum as follows:
 - i. For each observed power spectrum (except the VHF power spectrum), randomly select a simulated power spectrum $P_{\text{sim},j}(\nu)$ from those used to make the model average power spectrum, and measure its χ_{dist}^2 relative to the model, by substituting $P_{\text{sim},j}(\nu)$ for the observed power spectrum $P_{\text{obs},i}(\nu)$. Simulate a

VHF power spectrum by randomly drawing the power at each frequency from a Gaussian distribution of mean equal to the model power at that frequency and standard deviation equal to the estimated error at that frequency. Calculate the χ^2 of the simulated VHF power spectrum and sum with the χ_{dist}^2 measured for the randomly drawn simulated power spectra to find a total simulated χ_{dist}^2 .

- ii. Repeat the above step M times, to determine the distribution of total χ_{dist}^2 for M realisations of the model power spectrum, and count the number of realisations for which χ_{sim}^2 exceeds that measured for the observed power spectra, $M_{\chi_{\text{dist}}^2 > \chi_{\text{dist,obs}}^2}$. The probability that the data can be described as a realisation of the model is then given by $M/M_{\chi_{\text{dist}}^2 > \chi_{\text{dist,obs}}^2}$.
3. Use the probabilities obtained for each $P_{\text{mod}}(\nu)$ to set confidence limits on the model parameters, and plot best-fitting model average power spectra, confidence contours etc. as desired.

3.3 Lightcurve simulation and computing time

The main drawback with using the response method is that it is computationally demanding, requiring that large numbers of lightcurves are simulated and power spectra calculated from them. Even though PSRESP uses the inverse Fast Fourier Transform method of Timmer & König to generate lightcurves, the sheer length of lightcurves required to create model average power spectra to low frequencies means that we must consider ways to minimise the computing time needed to run PSRESP, while maintaining its reliability and accuracy. The first point to consider is the time resolution of monitoring lightcurves, which, in order to be totally realistic, should be equal to the duration of the

snapshot observations, ΔT_{obs} , i.e. 1 ksec. In practice, this constraint may be significantly relaxed, since for red-noise lightcurves, most of the contribution to aliasing and its uncertainty is due to variations at frequencies close to the Nyquist frequency, because there is not much power at frequencies much greater than this. For our monitoring data this means that we can account for most of the model average power spectrum and its uncertainty by making lightcurves of resolution ΔT_{sim} , which is ~ 10 times less than the sampling time-scale ΔT_{samp} of the observation scheme under consideration. We can estimate the small amount of extra power due to variations on time-scales less than this resolution by integrating the power of the specified model between the frequencies $(2\Delta T_{\text{sim}})^{-1}$ and $(2\Delta T_{\text{obs}})^{-1}$, which contributes a small additive component to the power spectrum analogous to the Poisson level, as described in Section 2.3. Note that we cannot simply simulate lightcurves of resolution ΔT_{samp} and then estimate the entire contribution due to aliasing in this way, since fluctuations in the underlying power spectrum at frequencies close to ν_{Nyq} contribute significantly to the uncertainty in the power at the highest frequencies, which we must determine accurately in order to estimate reliable confidences in the models. However, since we have increased the simulated lightcurve bin size from 1 ksec to ~ 10 ksec and ~ 100 ksec for daily and fortnightly monitoring schemes respectively, we have achieved a substantial decrease in computing time without compromising the accuracy of the method.

Now let us consider the durations of simulated lightcurves used by PSRESP. In order to take account of the distorting effects of red-noise leak due to power at frequencies less than the minimum frequency observed, simulated lightcurves must be significantly longer than the observation length, T (by at least a factor 10). However, we can minimise the cost of making long lightcurves by obtaining all our simulated lightcurves from a single, very long lightcurve. Thus, if we wish to simulate 1000 lightcurves of duration $10T$, we can instead simulate one single lightcurve of duration 10^3T and split it into 1000 smaller

lightcurves which we use as our sample of simulated lightcurves. Note that although any long timescale trends in the total simulated lightcurve will contribute the same amount of red-noise leak to all the simulated power spectra, the bulk of red-noise leak (and the power-spectral uncertainty it introduces) will be due to variations on shorter timescales, comparable to the observation duration. Hence the red-noise leak contribution in a model average power spectrum made from a 1000 segments of a lightcurve of duration 10^3T will be roughly equivalent to that measured from 1000 independently simulated lightcurves.

Finally, we must consider the effects of red-noise leak on the VHF power spectrum. Under the implementation of PSRESP described above, the VHF power spectrum made from continuous data segments (which does not suffer from the effects of aliasing) is used as a good representation of the underlying measured power spectrum and its uncertainties. However, if we simply compare the VHF power spectrum with the underlying undistorted model shape, we ignore the effects of red-noise leak which could be significant in distorting the shape of the VHF power spectrum, especially if the underlying power spectrum does not flatten significantly until far below the minimum frequency sampled by the VHF power spectrum. Therefore, we need to take account of the effects of red-noise leak on the model power-spectral shape at high-frequencies. It is not necessary to deal with aliasing effects since the lightcurve segments used to make the VHF power spectrum are continuous. The simulated lightcurves are made to be at least 64 times longer than ν_{\min}^{-1} , where ν_{\min} is the minimum frequency sampled by the VHF power spectrum, with resolution Δt smaller than $\frac{1}{2}\nu_{\max}^{-1}$ where ν_{\max} is the maximum frequency which contains significant power above the noise level (typically around 3×10^{-3} Hz) and is chosen so that the ratio of ν_{\max} to ν_{\min} is a power of 2. Power spectra of the lightcurves sampled to have duration ν_{\min}^{-1} may then be made using the Fast Fourier Transform, which allows a VHF model average power spectrum for 1000 simulated lightcurves to be determined very rapidly. The VHF model average power spectrum is

then used in place of the underlying model power spectrum (although the differences are generally only small since red-noise leak is not pronounced in our targets), while the errors determined from the observed VHF power are used as errors on the model (since these will not be greatly affected by red-noise leak).

3.4 Additional caveats

Although the technique presented here is currently the best method of reliably constraining the power-spectral shape of unevenly and discretely sampled lightcurves over a broad range of frequencies, it is important to note the assumptions which underly the PSRESP method as it now stands.

As I have described earlier, the measurement of a single broadband power spectrum from separate power spectra made from different observing schemes at different times relies on the assumption that the power spectral shape remains stationary on time-scales comparable to the duration of the monitoring programs described here. Proof of this assumption will come from the observational results themselves, in that the observed power spectra from different observing schemes can be well fitted with the same power-spectral model (as shown in the following chapter) although one of our targets, NGC 4051, shows evidence of unusual non-stationarity in its lightcurve which violates this assumption and will be described in Chapter 5.

A second assumption, which is made in the calculation of probabilities, is that the power spectra measured by different observing schemes are independent of one another. Clearly this is not strictly true if one observing scheme (e.g. 6 weeks of daily observations) is embedded in a longer one (three years of fortnightly observations). However, provided that the frequency coverage of the different schemes does not overlap significantly, then the only significantly correlated contribution to the power spectra will be due to red-noise leak effects

which are weak in any case and negligible if the power-spectra flatten towards low frequencies.

Finally, it is important to note the effects of the flux dependence of power spectral amplitude described in Section 2.4, since lightcurves simulated by the method of Timmer & König (or any other method known by the author) do not have this property. As described in Section 2.3, aliasing distorts the power spectrum by contributing a constant value to the power, whose integrated power is on average equal to the integrated power between the Nyquist frequency and $(2t_{\text{bin}})^{-1}$, where t_{bin} is the bin time of the sampled data points. Therefore the constant level due to aliasing is proportional to the integrated power. If the power spectrum was truly stationary, the contribution to this constant level from each data point would be the same, since the integrated power is independent of flux. However, because integrated power scales with the square of flux, data points which sample a higher local flux will contribute more to the aliasing than those which sample a lower flux. The net effect however, is that the mean level of aliasing contributed by all data points is proportional to the mean flux of the entire observation, so that it is equivalent to the case where the power spectrum is truly stationary. Therefore the fact that the simulated lightcurves don't show the intrinsic non-stationary behaviour shown by real lightcurves should not be important for the purpose of simulating power spectra.

Chapter 4

The broadband power-spectral shape of Seyfert galaxies

Overview

In this chapter, I apply the PSRESP method to measuring the broadband power-spectral shape of four Seyfert galaxies, including three of the targets we have monitored with *RXTE*, (MCG-6-30-15, NGC 5506 and NGC 5548) and NGC 3516 which has been intensively monitored as part of a separate program (Edelson & Nandra 1999). I test the measured power spectra with a simple power-law model to see if they flatten, and then test two simple models, a knee model and a high-frequency break model, to see how well they describe any flattening which we see. The implications of measured break frequencies for the hypothetical scaling of the power spectrum with luminosity and black hole mass are discussed. Finally, we look at the energy dependence of RMS variability and the power spectrum.

4.1 Observations and data reduction

In this section, I will describe the data used to make broadband power spectra for MCG-6-30-15, NGC 5506, NGC 5548 and NGC 3516, and the reduction method applied to the data.

4.1.1 MCG-6-30-15 observations

MCG-6-30-15 was observed by *RXTE* as part of our long-term Seyfert monitoring program, which consists of sequences of snapshots of ~ 1 ksec exposure obtained using several different sampling intervals. In accordance with the observing plan, which was intended to measure a broadband power spectrum at minimum cost in observing time, MCG-6-30-15 was observed weekly from 8 May until 21 August 1996, then twice-daily for two weeks from 23 August to 5 September 1996, then daily for ~ 4 weeks from 6 September to 29 September 1996. During AO2 and AO3, MCG-6-30-15 was observed every two weeks for the duration of each observing cycle until 2 February 1999, with the exception of a 4 week gap in observations (due to sun-angle pointing constraints on *RXTE*) each November. Monitoring continued in AO4 on a monthly basis, but for the purpose of this work, I only include data to the end of AO3, because instrumental gain changes on *RXTE* early in AO4 mean that count rates measured after that time are not comparable with those measured in previous AOs.

MCG-6-30-15 was also observed for ~ 400 ksec total exposure time from 03:31 UT on 4 August 1997 to 12:34 UT on 12 August 1997. This long-look observation was only interrupted by Earth-occultations of the source, and is the longest quasi-continuous observation of an AGN to date. I obtained the data for this observation from the *RXTE* public archive, in order to make the high-frequency and VHF power spectra of MCG-6-30-15. Spectral analysis of this




Figure 4.1: Intensive monitoring (top) and long-look (512-s binned) 2–10 keV lightcurves of MCG-6-30-15.

data is described by Lee et al. (1999), while Reynolds (2000) uses the data to study variability of the iron line. Power spectral analysis of the observation is reported by Nowak & Chiang (2000), who find evidence for a power spectral break at 10^{-4} Hz, however since Nowak & Chiang simply calculate the power spectrum and its errors directly from the observed lightcurve, without accounting for sampling effects and poorly-defined errors, no formal significance can be placed on their result. The additional data provided by our long term monitoring program combined with the PSRESP technique of power-spectral model fitting should allow us to place the first valid constraints on the power-spectral shape over a much broader frequency range.

The long-term monitoring lightcurve of MCG-6-30-15 was shown in Figure 1.6. The intensive monitoring and long-look lightcurves are shown here in Figure 4.1.

4.1.2 NGC 5506 observations

NGC 5506 was observed by *RXTE* as part of our long-term Seyfert monitoring program using the same observing pattern as used for MCG-6-30-15. Weekly monitoring lasted from 23 April to 28 July 1996 in AO1, followed by twice daily observations for two weeks and 4 weeks of subsequent daily observations, for a total intensive monitoring period from 8 August to 19 September. Fortnightly observations followed in AO2 and AO3 (with a two month gap during October and November of each year due to sun-angle constraints), until 2 February 1999. Monthly monitoring continued in AO4 but is not included for the reasons outlined earlier. The total monitoring lightcurve is shown in Figure 1.6 while the intensive monitoring section of the lightcurve is shown in more detail in Figure 4.2.

We obtained an *RXTE* ~ 100 ksec ‘long-look’ observation of NGC 5506 from 04:45 UT 20 June to 12:33 9 July 1997. Unfortunately, the observation was sporadically broken up over these 19 days, with no single continuous exposure (after accounting for Earth occultations) greater than 25 ksec (see Figure 4.2). Because the lightcurve is broken up in this way, it is next to useless for probing frequencies below 10^{-4} Hz, but since the X-ray variability of the source shows significant variability on short timescales, I shall just use the lightcurve to construct a VHF power spectrum using continuous data segments (i.e. down to 4×10^{-4} Hz). In order to sample frequencies down to 10^{-5} Hz, I shall use the 200 ksec long-look observation of NGC 5506 obtained with *EXOSAT* from 24-27 January 1986 (McHardy & Czerny 1987, and see Figure 1.4). The energy range sampled by *EXOSAT* (1–9 keV) is comparable to the 2–10 keV range I shall use here, so the power spectrum normalised by the squared mean flux should have similar shape and amplitude to the power spectrum measured by *RXTE*, if the power-spectral shape is stationary on time-scales of a decade. As an added precaution, however, I will also fit the broadband power spectrum of NGC 5506 without including the *EXOSAT* data.




Figure 4.2: Intensive monitoring (top) and long-look (512-s binned) 2–10 keV lightcurves of NGC 5506.

A study of the spectral variability of NGC 5506 using both the monitoring and long-look data is contained in Lamer, Uttley & M^cHardy (2000).

4.1.3 NGC 5548 observations

In common with the monitoring programs for MCG-6-30-15 and NGC 5506, we observed NGC 5548 with *RXTE* twice-daily for two weeks and then daily for four weeks, for the period from 26 June to 8 August 1996, although this intensive monitoring period did not occur at the end of the AO and was embedded in the period of weekly monitoring which lasted from 23 April to 26 October 1996. There were no gaps due to sun-angle constraints, so the monitoring campaign resumed with fortnightly observations on 12 November 1996,

which continued throughout AOs 2 and 3 (until 22 December 1998), but did not continue into AO4.

NGC 5548 was observed for number of continuous periods from June to August 1998 as part of a separate program to study spectral variability (see Chiang et al. 2000). I have obtained the longest of these observations from the *RXTE* public archive, of 90 ksec exposure covering the period 20 June to 23 June 1998. The observation is not continuous, but is much better sampled than the NGC 5506 observation, so it is adequate to measure the power spectrum down to a frequency of $\sim 10^{-5}$ Hz. However, since NGC 5548 shows no evidence for variability on timescales less than 1 ksec, I shall not use this long-look lightcurve to measure the VHF power spectrum and its errors directly.

The *RXTE* observations of NGC 5548 suffer from a minor complication in that the field of view also contains the BL Lac object 1E 1415.6+2557, offset 0.5° from NGC 5548, which is known to be X-ray bright (Nandra et al. 1993). Chiang et al. conducted separate pointings at this source and found that its contaminating contribution to the measured 2–10 keV PCA count rate (for 3 PCUs) of NGC 5548, after allowing for the effects of the PCA collimators, was only $\sim 2 \text{ count s}^{-1}$ (about 10% of the total measured count rate). The contaminating flux was estimated to vary by $\leq 0.8 \text{ count s}^{-1}$ in two months, an amount which would be indistinguishable from variations at the Poisson noise level of the NGC 5548 lightcurves. Hence, 1E 1415.6+2557 should not contribute significantly to the power spectrum of NGC 5548 which we measure, although, in order to take account of its contribution to the mean flux level of the NGC 5548 lightcurve, I shall henceforth subtract 2 count s^{-1} from the measured 2–10 keV mean flux level of NGC 5548 for the purposes of power-spectral normalisation.

Figure 1.6 shows the total long-term monitoring lightcurve obtained for NGC 5548, while the intensive monitoring and long-look lightcurves are shown in Figure 4.3.




Figure 4.3: Intensive monitoring (top) and long-look (512-s binned) 2–10 keV lightcurves of NGC 5548.

4.1.4 NGC 3516 observations

NGC 3516 was intensively monitored with *RXTE* in AO2 and AO3 as part of a separate program (P.I. Rick Edelson) to measure its broadband power spectrum, as reported in Edelson & Nandra (1999), who selected NGC 3516 as a good target on the basis that it is a high latitude X-ray source and can be continuously viewed without Earth occultation. The high-frequency part of the power spectrum was measured using a continuous 4.3 day observation (obtained 00:14 UT 22 May to 05:37 UT 26 May 1997), which was sandwiched in a sequence of 256 snapshots obtained every 12.8 hours (obtained 16 March to 30 July 1997), in order to sample medium frequencies. The low-frequency coverage was extended by sampling every 4.3 days for the remainder of AO2 and AO3. Edelson & Nandra report the detection of a knee at $\sim 4 \times 10^{-7}$ Hz, however they take no account of the effects of aliasing on the power spectrum and furthermore, determine errors using the spread in the measured power spectrum, rather than estimating them from simulations. Hence the resulting measurement of power spectral shape is biased towards a flattening at low frequencies (due to the fact that aliasing is greater for longer sampling intervals). It is therefore important to apply the PSRESP method to the observations of NGC 3516, in order to estimate a reliable confidence in any flattening in the power spectrum and compare the measurements of power-spectral shape obtained by PSRESP with those of Edelson & Nandra. To this end, I obtained data from the NGC 3516 monitoring program from the *RXTE* public archive, including long-term monitoring data up to 28 December 1998. I also obtained data from a second long-look observation from 08:00 UT 13 April to 16:13 UT 16 April 1998, which will serve to further tie down the power spectrum at high frequencies. All the lightcurves I will use to measure the broadband power spectrum are shown in Figure 4.4. Note that the long-look lightcurves show little evidence for variability on time-scales less than 1 ksec, so as with NGC 5548 I shall not use the long-look lightcurves to measure a VHF power

spectrum.

4.1.5 Data reduction

All the data described above was reduced using the `FTOOLS v4.2` *RXTE* data reduction tools. Lightcurves were made using PCA Standard 2 data in the 2–10 keV energy range, corresponding to PCA channels 7–28, for the top layer of the PCUs only, in order to maximise signal to noise, since the background count rate becomes significantly larger than the source count rate if the other layers are included. Only data from PCUs 0, 1 and 2 were extracted, because PCUs 3 and 4 are only switched on intermittently due to discharge problems. Standard screening criteria were applied, so that the pointing offset to the source was less than 0.02° , Earth elevation was restricted to be greater than 10° , electron contamination was required to be less than 10% and data obtained within and up to 20 minutes after the South Atlantic Anomaly maximum was excluded.

Background was estimated using the L7 faint source models for gain epoch 3 and subtracted from the observed lightcurves to leave the source lightcurves. As part of their program to monitor NGC 3516, Edelson & Nandra obtained offset pointed observations, near to NGC 3516 but in a region of sky believed to be free of bright sources, in order to test the accuracy of the L7 background model and search for long-term variations in any residuals due to poor background subtraction (which might contribute spurious power to the power spectrum). Although they measured a small, but constant positive flux offset in background-subtracted blank-sky observations ($0.87 \text{ count s}^{-1}$), long term variations in this offset were very small (RMS variability $0.39 \text{ count s}^{-1}$ in one year). This level of variability is not much greater than that expected due to Poisson noise, and is negligible compared to the variations we see in our




Figure 4.4: 2–10 keV lightcurves of NGC 3516, from top to bottom: total monitoring lightcurve, intensive monitoring lightcurve, May 1997 long-look lightcurve, April 1998 long-look lightcurve (long-look lightcurves binned to 512 s).

monitoring lightcurves on a similar timescale, hence we do not expect it to contribute significantly to the power spectrum. Comparison of simultaneous *BeppoSAX* and *RXTE* observations of NGC 4051 have shown that, in this source the L7 background model is accurate in subtracting background with negligible residuals (Uttley et al. 1999, and see Section 5.2.2). This suggests that much of the residual positive offset seen by Edelson and Nandra may be due to spatial X-ray fluctuations in the cosmic X-ray background or a single, relatively faint X-ray source.

Monitoring lightcurves were binned to the minimum bin-times to be used for the power-spectral analysis, 8640 s for MCG-6-30-15, NGC 5506 and NGC 5548, and 4608 s for NGC 3516, in order to avoid splitting single snapshots over two or more time bins. Although binning in this way results in some small inaccuracies in the observation times, the discrepancies will be smaller than the bin times used for lightcurve simulations and so will not affect the results. Long-look lightcurves were binned in 512 s bins in order to fit power spectral models down to $\sim 10^{-3}$ Hz without needing to simulate lightcurves with excessively small bins. Additional long-look lightcurves were made with 16 s binning for MCG-6-30-15 and NGC 5506 in order to measure VHF power spectra directly in the 4×10^{-4} – 10^{-2} Hz frequency range. Finally, the 30 s binned *EXOSAT* band D lightcurve for NGC 5506 was obtained from the HEASARC public archive and binned up to 512 s (this data is pre-reduced).

4.2 A first look at the RMS variability and power spectra

In order to calculate the power spectra, I first make additional, separate lightcurves corresponding to just the periods of intensive monitoring so that

for each source we have lightcurves for three observing schemes for input into PSRESP:

1. A long time-scale lightcurve incorporating *all* monitoring data, to measure the low-frequency power spectrum ($\sim 10^{-8}$ Hz– 10^{-6} Hz).
2. An intensive monitoring lightcurve, to measure a medium-frequency power spectrum ($\sim 10^{-6}$ Hz– 10^{-5} Hz).
3. A long-look lightcurve (two such lightcurves for NGC 3516) to measure the high-frequency power spectrum $\sim 10^{-5}$ Hz– 10^{-4} Hz.

Additionally, for MCG-6-30-15 and NGC 5506, I calculate the power spectrum (and reliable errors) directly for continuous segments of the 16 s binned lightcurves of ~ 2.5 ksec duration, to provide a measure of the VHF power spectrum at the highest frequencies ($\sim 4 \times 10^{-4}$ Hz– 10^{-2} Hz) which can be input directly into PSRESP.

The monitoring lightcurves consist of a mixture of sampling times (e.g. twice-daily and daily for the intensive monitoring lightcurves, while the long term lightcurves consist of twice-daily, daily, weekly and fortnightly monitoring). Therefore, in order to minimise any distortion, the power spectra are made from lightcurves binned up to the maximum sampling interval of the observing scheme under consideration (i.e. two weeks or 1209.6 ksec for the long-term monitoring lightcurves, 86.4 ksec for the intensive monitoring lightcurves, except for NGC 3516 where I bin the long-term and intensive monitoring lightcurves to 4.3 days and 12.8 hours respectively). Long-look lightcurves will be binned to 2048 s for the purposes of making power spectra, so that gaps due to earth occultation are minimised to be no more than one bin wide. The total lightcurve durations, bin widths, mean fluxes, fractional RMS variability (after subtracting the Poisson noise contribution to variance) and power-spectral Poisson noise levels (equal to twice the mean total count rate for the

lightcurve) for each lightcurve are given in Table 4.1. Note that mean flux and fractional RMS are calculated based on the quoted bin widths, i.e. the contributions to mean flux and fractional RMS from each bin are equally weighted, so that bins containing many data points (e.g. 2-week wide bins which contain daily or twice-daily observations) do not contribute more to the mean flux or variance than bins which contain a single data point.

It is particularly interesting to note that the largest fractional RMS variability measured on long time-scales is that of NGC 5548, which has the highest luminosity of the sample, contrary to what we would expect if variability simply scaled inversely with luminosity, as suggested by early studies of X-ray variability on short time-scales using *EXOSAT* data (Lawrence & Papadakis 1993; Green, McHardy & Lehto 1993). However, the fractional RMS variability of NGC 5548 measured from the long-look lightcurve is smaller than that of NGC 5506 (for a comparable duration), suggesting that the inverse correlation between luminosity and X-ray variability reported for *EXOSAT* data holds on short time-scales, but breaks down on longer time-scales. A possible reason for this is that the power spectra of lower luminosity Seyfert galaxies flatten at higher frequencies than higher luminosity objects, so that the integrated power in higher luminosity objects overtakes that of lower luminosity objects on long time-scales. To test this, we must look at the power spectra.

I made power spectra from the binned lightcurves using the method of Deeming (1975), determined the logarithm of power at each frequency and binned the power spectra in logarithmically spaced frequency bins, separated by a factor of 1.3 in frequency but with a minimum of two measured powers per bin, so that the bin spacing is larger at the lowest frequencies sampled by each power spectrum. The VHF power spectra for NGC 5506 and MCG-6-30-15 were calculated by measuring power spectra for each continuous lightcurve segment and binning them together in logarithmically spaced bins also separated by a factor of 1.3 in frequency, except that the minimum number of

	Long time-scale				
	T	ΔT	μ	σ_{frac}	P_{noise}
MCG-6-30-15	8.52×10^7	1.2096×10^6	14.0	26.5%	51.5
NGC 5506	8.52×10^7	1.2096×10^6	26.8	22.6%	76.7
NGC 5548	8.26×10^7	1.2096×10^6	13.6	30%	54.1
NGC 3516	5.60×10^7	3.6864×10^5	13.3	29.6%	49.5
	Medium time-scale				
	T	ΔT	μ	σ_{frac}	P_{noise}
MCG-6-30-15	3.22×10^6	8.64×10^4	14.7	21.7%	51.5
NGC 5506	3.61×10^6	8.64×10^4	22.4	15.1%	68.0
NGC 5548	3.58×10^6	8.64×10^4	14.8	20%	56.4
NGC 3516	1.18×10^7	4.608×10^4	12.7	28.7%	48.3
	Short time-scale				
	T	ΔT	μ	σ_{frac}	P_{noise}
MCG-6-30-15	7.23×10^5	2048	12.2	20.8%	47.7
NGC 5506 ^a	2.25×10^5	2048	6.9	12.1%	96.4
NGC 5548	1.97×10^5	2048	22.5	8.7%	72.1
NGC 3516 ^b	3.62×10^5	2048	11.5	7.2%	45.7
NGC 3516 ^c	2.88×10^5	2048	15.2	9.8%	52.9

Table 4.1: Parameters of lightcurves for use in PSRESP. T and ΔT are the lightcurve duration and sampling interval (in seconds), μ and σ_{frac} are the lightcurve mean flux (in count s^{-1}) and fractional RMS respectively and P_{noise} is the Poisson noise level expected in the power spectrum due to counting statistics (in unnormalised RMS-squared units, $\text{count}^2 \text{s}^{-2} \text{Hz}^{-1}$). Notes: ^a Details given in the table are for the *EXOSAT* lightcurve, the *RXTE* lightcurve used to measure the power spectrum at the highest frequencies has $\mu = 28.1$, $P_{\text{noise}} = 79.7$. ^b Lightcurve obtained 22–26 May 1997. ^c Lightcurve obtained 13–16 April 1998.




Figure 4.5: Raw broadband power spectra of MCG-6-30-15, NGC 5506, MGC 5548 and NGC 3516. The dashed line shows the low-frequency part of the power spectrum, made from the total monitoring lightcurves, while the dotted line shows the high-frequency part made from long-look lightcurves. Solid lines mark the medium-frequency power spectrum (made from the intensive monitoring lightcurves), VHF power spectra for MCG-6-30-15 and NGC 5506, and the power spectrum of the second long-look observation for NGC 3516.

points per bin was set to 10 so that the standard error could be estimated for each frequency bin. Each power spectrum was then normalised by the squared mean flux of the lightcurve used to make them, so that the resulting power is in fractional RMS squared units (see Section 2.1). The resulting measured broadband power spectrum for each object is shown in Figure 4.5.

Inspection of the power spectra in Figure 4.5 shows that they do flatten, however we must be cautious since these are raw, observed power spectra which are almost certainly distorted by sampling effects. We can only be certain that these power spectra do indeed flatten by using PSRESP to test models for unbroken power-law power spectra against the data.

4.3 Do the broadband power spectra really flatten?

To determine if the power-spectra flatten, I will test a simple power-law model for the underlying power spectrum, $P_{\text{mod}}(\nu)$ of the form:

$$P_{\text{mod}}(\nu) = A \left(\frac{\nu}{\nu_0} \right)^{-\alpha} + C_{\text{noise}},$$

where A is the the amplitude of the model power spectrum at a frequency ν_0 , α is the power-spectral slope and C_{noise} is a constant value which is fixed at the Poisson noise level for the lightcurve. Note that the Poisson noise level is included in the model rather than subtracted from the power spectra before model fitting, because the power spectra are binned logarithmically (so constants in linear space may not simply be subtracted). This is particularly important for the VHF power spectra of NGC 5506 and MCG-6-30-15, whose standard errors are determined in logarithmic space, and also for high-frequency power spectra in general, which are close to the Poisson noise level, since fluctuations in the power spectrum lead to some measured powers lying

	Long time-scale	Medium time-scale	Short time-scale
MCG-6-30-15	86400 s	8640 s	512 s
NGC 5506	86400 s	8640 s	512 s
NGC 5548	86400 s	8640 s	512 s
NGC 3516	46080 s	4608 s	512 s

Table 4.2: Time resolution of simulated lightcurves used in PSRESP.

below the Poisson noise level (so subtraction of this level would lead to negative measured powers).

The model is tested against the measured power spectra described in the previous section by stepping through values of α from 1.0 to 2.4 in increments of 0.1 (i.e. test the model with $\alpha = 1.0, 1.1, 1.2$ etc.). These values of α cover the range of reasonable values which could possibly be fitted to the data. Probabilities that the measured power spectra could be a realisation of the model are calculated by PSRESP, as described in Section 3.2 using $N = 1000$ simulated lightcurves to determine the distorted model average power spectrum and the distribution of χ^2_{dist} of the realisations of the model (which is determined for each set of power spectra by randomly measuring χ^2_{dist} for 10000 sets of simulated power spectra). The simulated lightcurves have time resolutions ΔT_{sim} given in Table 4.2. Additional distortion in the power spectrum due to model power at frequencies greater than $(2\Delta T_{\text{sim}})^{-1}$ is calculated directly from the model, as described in Section 3.3. Distorted model power spectra which take account of red-noise leak for the highest frequency power spectra (with directly measured errors) included in the broadband power spectra of MCG-6-30-15 and NGC 5506 are determined using the method also described in Section 3.3. The best-fitting α and corresponding probability that the measured power spectra could be a realisation of a model with a single power-law slope are given in Table 4.3. The first and second of each of these values shown

	Best-fitting α	Probability
MCG-6-30-15	1.5	0.17%
NGC 5506	1.4/1.5	9.4%/1.4%
NGC 5548	1.6	33%
NGC 3516	1.8	3.4%

Table 4.3: Results from fitting broadband power spectra of four Seyfert galaxies with a simple unbroken power-law model.

for NGC 5506 correspond to fits without or including the *EXOSAT* data respectively. The simple power-law model, without any flattening is rejected at better than 99% confidence for MCG-6-30-15, better than 90% confidence (or close to 99% confidence including the *EXOSAT* data) for NGC 5506 and better than 95% confidence for NGC 3516. Only the NGC 5548 broadband power spectrum does not reject the model at a significant confidence, although even in this case, the model fit is not particularly good. The best-fitting models are compared with the measured power spectra in Figure 4.6. It is apparent from these plots that the simple power-law model does not fit the observed power spectra of MCG-6-30-15, NGC 5506 and NGC 3516, even after allowing for the distorting effects of sampling, because the intrinsic power spectrum of each of these objects does indeed flatten towards low frequencies. There is no significant evidence for flattening at low frequencies in the power spectrum of NGC 5548. Figure 4.7 shows the fit probability measured at each fitted value of α for NGC 5548, which demonstrates how PSRESP is capable of finding well-defined probability maxima in the same way that χ^2 fitting can, using more conventional data sets. We have now established that the power-spectra of at least three of the four Seyfert galaxies studied here flatten significantly towards low frequencies. That the power spectrum of NGC 5548 does not flatten significantly suggests that the break in its power spectrum occurs at a lower




Figure 4.6: Comparison of best-fitting model average power spectra with the observed power spectra for the single power-law model described in the text. Open squares mark the low-frequency model average, simple crosses mark the high-frequency model average while filled squares mark the medium-frequency model average, the VHF model average for MCG-6-30-15 and NGC 5506, and the high-frequency model average for the second long-look observation of NGC 3516. Note that the error bars represent the RMS error in the simulated power spectra used to calculate χ^2_{dist} as described in Section 3.1.




Figure 4.7: Probability that a single power-law of slope α is acceptable to describe the broadband power spectrum of NGC 5548.

frequency than the breaks in the power spectra of the other objects, which supports the interpretation for the higher long time-scale RMS of NGC 5548 given in the preceding section. The next step is to try to fit the observed power spectra with more complex models which flatten at low frequencies, in particular, can we distinguish between models where the power spectrum flattens to $\alpha = 0$ or $\alpha = 1$, and can we constrain any break frequencies to test if the break frequency scales with luminosity?

4.4 Fitting models with power-spectral breaks

Unfortunately, due to the large amount of computing time required to run PSRESP, it is not possible to fit complex multi-parameter models to the data on the computing facilities (and in the time) available to the author. Furthermore, the quality of the data presented here does not yet warrant such fits. Therefore, I will leave more detailed model fits as a future work. For now, it is sufficient

to determine whether we can rule out any particular model for the power-spectral flattening which we see, and to try and constrain any characteristic break frequencies in any models which successfully describe the data. I shall consider two possible models for the power-spectral flattening, a ‘knee’ model where the power spectrum turns over to $\alpha = 0$ at low-frequencies, and a ‘high-frequency break’ model, where the power spectrum breaks to $\alpha = 1$ below some break frequency, analogous to the high-frequency break of Cygnus X-1.

4.4.1 The knee model

Knee models where the power-law power spectrum turns over to slope $\alpha = 0$ (i.e. white noise) below some knee frequency are predicted by simple shot-noise models. In the simplest case (e.g. Halford 1968), where all the shots have a single exponential decay time-scale, τ , the resulting power spectrum is predicted to have the form:

$$P(\nu) \propto \frac{1}{\tau^{-2} + (2\pi\nu)^2},$$

so that on frequencies corresponding to time-scales longer than the shot decay time-scale, the power spectrum flattens to $\alpha = 0$, while at higher frequencies $\alpha = 2$ (see Figure 4.8). Note that the flattening occurs relatively quickly, in less than a decade of frequency. However as described in Section 1.2.1, measured power spectra of X-ray binaries have slopes of $\alpha \sim 1$ which persist over greater than one decade in frequency before flattening and cannot be the result of the rapid flattening from $\alpha = 2$ predicted by simple shot-noise models. Furthermore, *EXOSAT* measurements of the power spectrum of AGN showed power-law slopes with α intermediate between 1 and 2 (Green, McHardy & Lehto 1993; Lawrence & Papadakis 1993). To account for these observations, more complex shot-noise models have been invoked, where the shots have a power-law distribution of decay time-scales (e.g. Lehto 1989; Lochner, Swank & Szymkowiak 1991), which can produce power-spectral slopes intermediate




Figure 4.8: The form of the knee model power spectrum.

between 0 and 2. These intermediate slopes flatten to zero slope at frequencies corresponding to the longest shot time-scales, so that their power spectra are described by:

$$P(\nu) = \frac{A}{\left(1 + \left(\frac{\nu}{\nu_{\text{knee}}}\right)^2\right)^{\alpha/2}},$$

where A is the constant amplitude of the power-spectrum at zero slope and ν_{knee} is the ‘knee frequency’ and α is now defined as the power-spectral slope above the knee frequency.

Let us now test this model against the measured broadband power spectra, to see if it can explain the flattening we see. Using the equation given above for the underlying power spectral shape (also including the constant Poisson noise level), we can fit the model in the same way as fitting a simple power law in the previous section. The free parameters to be stepped through are α , which is again incremented in steps of 0.1 between 1.0 and 2.4, and ν_{knee} which is stepped through by multiplicative factors of 2, from 10^{-8} Hz to 10^{-4} Hz, since a very broad range in frequency must be covered. Approximately 200 pairs of α and ν_{knee} must be tested (as opposed to only 15 parameters when

	Probability	α	$\nu_{\text{knee}}/10^{-6}$ Hz
MCG-6-30-15	19%	1.8 ± 0.1	5.12 (2.56,5.12)
NGC 5506 ^a	58%	$1.7 \pm_{0.3}^{0.7}$	0.64 (0.08/0.0,5.12/10.24)
NGC 5506 ^b	68%	$1.9 \pm_{0.2}^{0.3}$	2.56 (1.28/0.16,5.12/10.24)
NGC 5548	47%	$1.6 \pm_{0.1}^{0.3}$	0.02 (0.02/0.02,0.16/1.28)
NGC 3516	17%	2.1 ± 0.3	0.64 (0.32,1.28)

Table 4.4: Results from fitting broadband power spectra of four Seyfert galaxies with a knee model, see text for details

fitting the simple power law in the previous section), so to save on computing time, the number of lightcurve simulations used to estimate each model average power spectrum for each pair of parameters is reduced from 1000 to 100. The χ_{dist}^2 distribution (and hence probabilities that the data can be described as a realisation of the model) is determined by determining χ_{dist}^2 for 1000 sets of simulated power spectra. The restricted number of simulated power spectra will increase the uncertainty in the calculated probabilities, however this effect is not very significant, as contour plots of the fitted parameter space will show. The best-fitting parameters and probabilities are shown in Table 4.4. The superscripts ^a and ^b mark the NGC 5506 results excluding and including the *EXOSAT* data respectively. The errors quoted for α in Table 4.4 correspond to the values of α below which the fit probability is reduced to less than 32% (i.e. they represent 68% confidence limits), except for models where the maximum fit probability is itself less than 32%, in which case the errors represent 90% confidence limits. The numbers in brackets, which follow the best-fitting knee frequencies, correspond to the 68%/90% confidence lower limits and (after the comma), upper limits on the knee frequency (except where the maximum fit probability is less than 32%, in which case only the 90% upper and lower limits are quoted). Contour plots for each of the knee model fits (excluding the

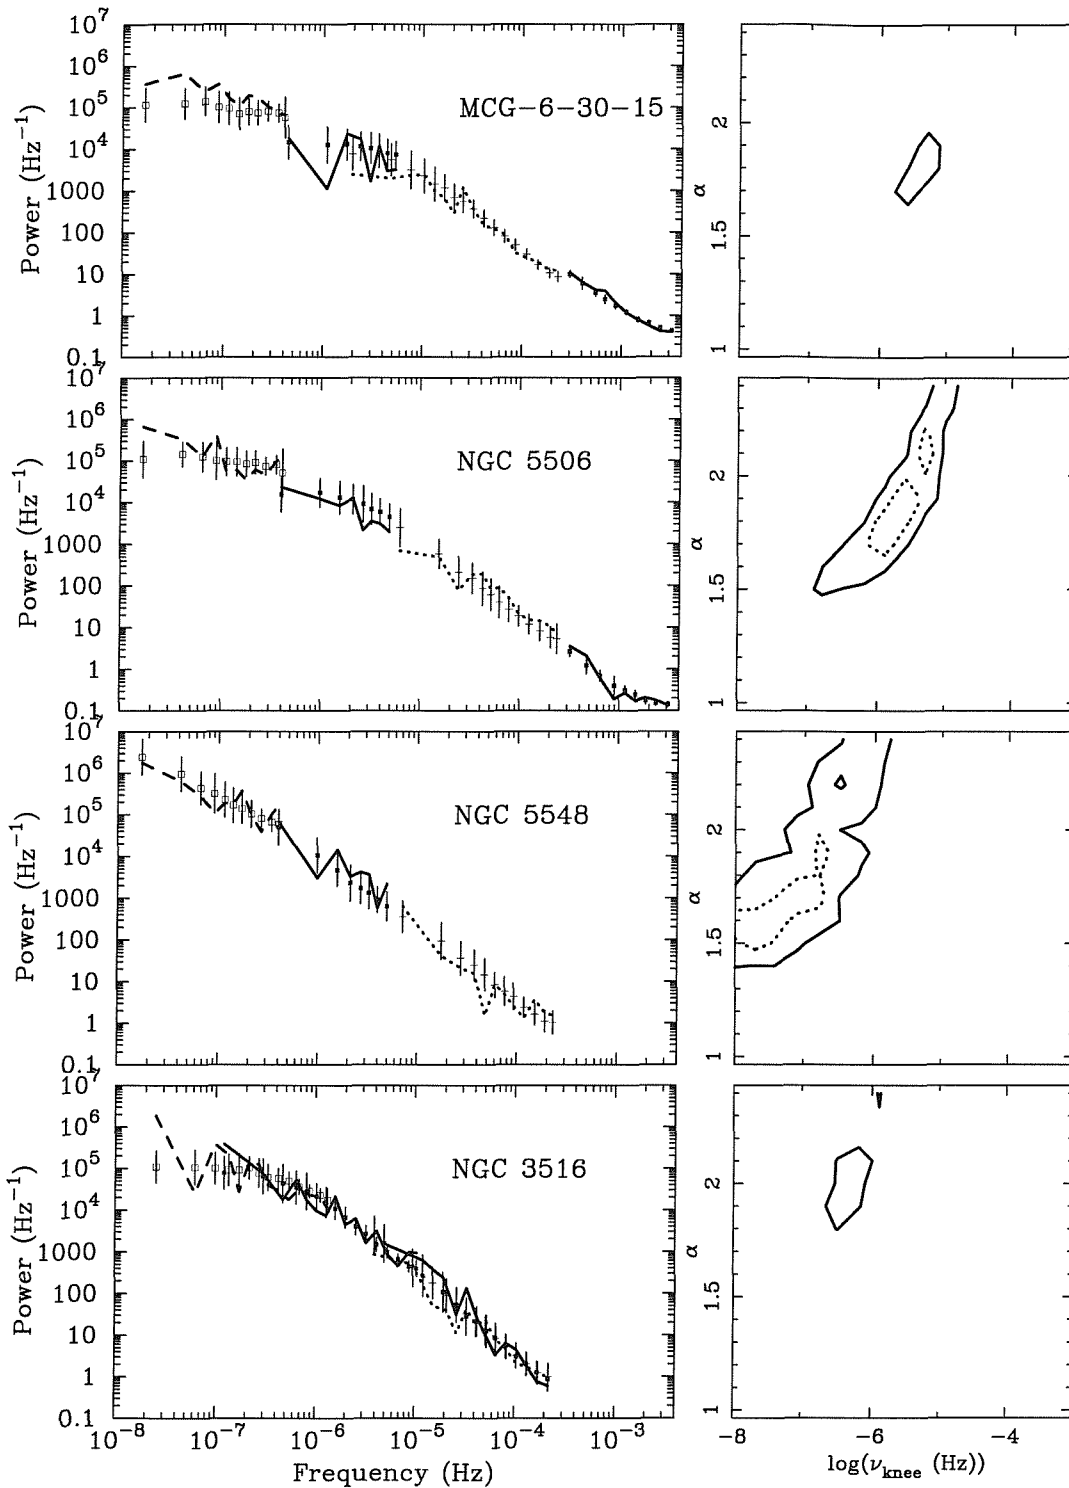


Figure 4.9: Comparison of best-fitting model average power spectra with observed power spectra for the knee model described in the text (left), and corresponding confidence contours for the parameter space searched (right). The dashed and solid confidence contours represent the 68% and 90% confidence limits respectively.

fit of the NGC 5506 broadband power spectrum which excludes the *EXOSAT* data), together with the best-fitting model average power spectra, are shown in Figure 4.9. The contour plots show that the acceptable regions are broad and well-defined, rather than consisting of many separate ‘islands’, which implies that using only 100 simulated lightcurves per measured power spectrum is sufficient to determine reliable maxima in the probability space.

As Table 4.4 shows, the fit probabilities for all objects are better than those obtained from fitting a simple power-law model to the data, although the knee model is still not able to fit the observed power spectra of MCG-6-30-15 and NGC 3516 particularly well. Note that the MCG-6-30-15 and NGC 3516 data sets are the best in the sample, since they have the longest long-look and intensive monitoring lightcurves respectively, suggesting that perhaps the knee model is not the correct model for representing the low-frequency flattening in the sample, and that NGC 5548 and NGC 5506 would also show poorer fits if better quality lightcurves were obtained for them. Therefore we shall now examine a second model, the high-frequency break model, to see if this can fit all the observed broadband power spectra adequately.

4.4.2 The high-frequency break model

The motivation for the high-frequency break model comes from the power spectrum of the black hole X-ray binary Cygnus X-1 in the low state, which shows two frequency breaks (as described in Section 1.2.1, also see Figure 2.5): a high frequency break which varies between 1–5 Hz, above which the slope α varies between 1.4 and 2.4 (and below which $\alpha \simeq 1$), and a low-frequency break varying between 0.04–0.4 Hz, below which the slope flattens to $\alpha \simeq 0$. If AGN have a similar power spectral-shape to Cygnus X-1 (albeit scaled down in frequency by some factor, possibly the black hole mass), then because the power spectral slopes of AGN lightcurves measured by *EXOSAT* (and appar-




Figure 4.10: The form of the high-frequency break model power spectrum.

ently also by *RXTE*) are significantly greater than 1, we may be measuring the analog of the high-frequency break in Cygnus X-1.

To test this possibility, we should try fitting the observed power spectra with a model of the form used to fit the high-frequency power spectrum of Cygnus X-1 (e.g. Nowak 1999):

$$P(\nu) = \begin{cases} A \left(\frac{\nu}{\nu_{\text{bk}}}\right)^{-\alpha_{\text{hi}}} & \text{if } \nu > \nu_{\text{bk}} \\ A \left(\frac{\nu}{\nu_{\text{bk}}}\right)^{-\alpha_{\text{lo}}} & \text{otherwise} \end{cases}$$

Where A is the power-spectral amplitude at the break frequency ν_{bk} , and α_{hi} and α_{lo} are the high and low-frequency slopes respectively, such that $\alpha_{\text{hi}} > \alpha_{\text{lo}}$. An example of a high-frequency break model with $\alpha_{\text{hi}} = 2$ and $\alpha_{\text{lo}} = 1$ is shown in Figure 4.10.

We do not consider the low-frequency break in this model in order to minimise the number of free parameters. This approach is valid since, if the model is correct, low-frequency breaks will occur at least a decade lower in frequency than any measured high-frequency breaks and so will not contribute as significantly to any flattening (besides which, if additional low-frequency breaks are significant they will be apparent from the residuals in any comparison of the

data with the model).

An important question is whether it is valid to fix the slope below the high-frequency break. In the case of the knee model, fixing the slope at zero was valid on physical grounds, since the model was based on shot-noise models which predict a flattening to zero slope. The high-frequency break is not based on any physical model however, since we have chosen it as a good empirical representation of the data. Hence there is no physical basis for the sharpness of the break, but since this model can adequately describe the high-frequency power-spectral shape of Cygnus X-1, it should also serve as a possible empirical representation of the power spectra of much poorer quality which we measure here. A further advantage is that if the model is a good fit to the data, any break frequencies measured can be compared directly with that of Cygnus X-1, to investigate how break frequency scales with black hole mass. Clearly it is desirable, on grounds of computation time, to restrict the number of free parameters by fixing the slope below the break, but there are also observational reasons why we might fix the slope to $\alpha_{lo} = 1$. One particularly striking aspect of all the Cygnus X-1 power spectra is that, despite the variations in the position of the high and low-frequency breaks and the slope above the high-frequency break (e.g. as shown by Belloni & Hasinger 1990), the slope of the intermediate power-spectrum, between the two breaks, is always remarkably close to 1. Furthermore, Nowak et al. (1999) show that the power spectra of Cygnus X-1 made from simultaneous lightcurves in different energy bands show an energy dependence above the high-frequency break (in that α_{hi} decreases towards higher energies) but maintain the same shape (i.e. $\alpha_{lo} = 1$) below the break. These results suggest that a power-spectral slope of 1 (or very close to 1) below the high-frequency break may, in fact, be the rule. We can determine if the power spectra we measure are at least consistent with this possibility by fitting the high-frequency break model (including the constant Poisson noise level, as before), fixing $\alpha_{lo} = 1$ and stepping through the same parameter ranges as used to fit the knee model (i.e. $\alpha_{hi} = 1.0$ -2.4 in

	Probability	α	$\nu_{\text{bk}}/10^{-6}$ Hz
MCG-6-30-15	67%	$2.0 \pm_{0.2}^{0.1}$	51.2 (25.6/12.8,102.4/102.4)
NGC 5506 ^a	90%	$2.1 \pm_{0.6}^{0.3}$	25.6 (0.4/0.0,51.2/102.4)
NGC 5506 ^b	97%	$2.4 \pm_{0.8}^0$	51.2 (1.6/0.4,51.2/102.4)
NGC 5548	73%	$2.4 \pm_{0.9}^0$	2.56 (0.0/0.0,5.12/10.24)
NGC 3516	61%	2.2 ± 0.2	2.56 (2.56/0.64,5.12/5.12)

Table 4.5: Results from fitting the broadband power spectra of four Seyfert galaxies with a high-frequency break model.

increments of 0.1, $\nu_{\text{bk}} = 10^{-8}$ – 10^{-3} Hz in multiples of 2). The resulting best-fitting parameters are reported in Table 4.5, with the results presented in the same format as for Table 4.4. The best-fitting model average power spectra and confidence contour plots are shown in Figure 4.11. The high-frequency break model fits the data better than the knee model for all the objects in the sample, and importantly, provides good fits to the measured power spectra of NGC 3516 and MCG-6-30-15, which the knee model was unable to fit well. Although the best knee model fits to the power spectra of NGC 3516 and MCG-6-30-15 were only ruled out at better than 80% confidence, the fact that the high-frequency break model fits *both* object’s power spectra well implies that it is the best model for describing the data.

4.4.3 Comparison with other work

Power spectral measurements of all the Seyfert galaxies described here have been attempted before, although with questionable statistical validity, so it is interesting to compare the reliable results obtained with PSRESP with these earlier efforts. The long-look *EXOSAT* observation of NGC 5506 was used




Figure 4.11: Comparison of best-fitting model average power spectra with observed power spectra for the high-frequency break model described in the text (left), and corresponding confidence contours for the parameter space searched (right). The dashed and solid confidence contours represent the 68% and 90% confidence limits respectively.

by McHardy (1988), together with observations from earlier missions, to provide a crude broadband power spectrum which showed evidence of flattening within the *EXOSAT* power spectrum (i.e. above 10^{-5} Hz), which continued towards lower frequencies. McHardy interpreted this flattening as being a gradual change over several decades in frequency, rather than being due to breaks in the power spectrum. No confidence limits could be placed on this apparent flattening however, as stochastic variations in the power spectrum were unaccounted for, as were the undoubtedly strong aliasing effects of the very low-frequency sampling used to make the power spectrum at low frequencies. The results of using PSRESP to fit the much better *RXTE* monitoring data, even combined with the *EXOSAT* long-look, only show that flattening is significantly detected but cannot strongly constrain the model which describes the flattening. This implies that although the flattening originally described by McHardy was undoubtedly real, any interpretation of the form of power spectral shape which describes the flattening was unwarranted by the data. Even the continuous long-duration *EXOSAT* long-look contains too few measurements at the lowest frequencies to provide a reliable measure of the power at those frequencies.

The long-look *RXTE* observation of MCG-6-30-15 was used by Nowak & Chiang (2000) to measure the power spectrum down to $\sim 10^{-6}$ Hz. Nowak & Chiang reported a power-spectral shape similar to that of Cygnus X-1, with a high-frequency break at $\simeq 10^{-4}$ Hz with slope $\alpha \simeq 2$ above the break, and a slope $\alpha \simeq 1$ below the break, flattening to $\alpha \simeq 0$ below a second break at $\simeq 10^{-5}$ Hz. The high frequency break and slope measured by Nowak & Chiang match up quite well with the slope and break measured by fitting the high-frequency break model in PSRESP. However, there is no evidence whatsoever from the monitoring data for a break to zero slope at a frequency of $\simeq 10^{-5}$ Hz. Comparison of the MCG-6-30-15 model-average power spectrum with the data in Figure 4.11 shows that although the section of the power spectrum measured from the long-look lightcurve does flatten at low frequencies this is not signif-

icant at all and is perfectly consistent with stochastic variations in the power spectrum. Again, the power spectrum is over-interpreted at low frequencies because the errors at low frequencies are not well defined.

Chiang et al. (2000) measure the power spectrum of NGC 5548 at high-frequencies ($> 10^{-5}$ Hz) using *RXTE* long-look data, and find no evidence of flattening, but report a flattening at low frequencies measured from the *RXTE* ASM lightcurve of this object. Unfortunately, as discussed in Section 1.3.1, the ASM count rates measured for sources as faint as NGC 5548 show strong spurious variations, so the ASM cannot be used to measure the power spectrum.

Edelson & Nandra (1999) report the first broadband power spectrum made from well sampled lightcurves, for NGC 3516, with the specific intention of measuring low-frequency flattening. They report a significant flattening of the power spectrum, with a break at 4.14×10^{-7} Hz obtained by fitting a knee model to the data, however they do not take account of aliasing effects on the low-frequency power spectrum or calculate valid errors on the measured power spectral points (since they bin as few as two points together in the lowest frequency bins). Fitting similar data (i.e. the same data set as used by Edelson & Nandra but also including three months of additional long-term monitoring data and the second long-look observation), we find that the knee model fits a similar break of 6.4×10^{-7} Hz, although this fit is not particularly good, and a better fit can be obtained by fitting a high-frequency break model with a frequency break at 2.56×10^{-6} Hz. The high-frequency slope measured by Edelson & Nandra is $\alpha = 1.74 \pm 0.12$, in contrast with the slope, $\alpha = 2.2 \pm 0.2$ which we measure. The reason for this discrepancy seems to be that stochastic variations in the lowest frequency power measured from the long-looks cause large uncertainties in the high-frequency slope which can only be accounted for using the response method. Additionally, using both long-look lightcurves to make power spectra leads to a more accurate ‘average’ of the high-frequency slope.

4.5 The dependence of break frequency on AGN luminosity and black hole mass

Now let us consider how the measured break frequencies (in the context of the models described here) relate to AGN luminosity, and possibly the black hole mass. Since the only well-constrained breaks are those measured for MCG-6-30-15 and NGC 3516, we shall consider these two objects and then speculate on likely break frequencies for NGC 5548 and NGC 5506.

What is most striking about the breaks measured for MCG-6-30-15 and NGC 3516 is that they are significantly different, in either model, at a level of better than 99% confidence, i.e. the 90% confidence limits of the break frequencies do not overlap. This is unexpected, because both objects have similar 2–10 keV X-ray luminosities of $L_{2-10} \sim 1.5 \times 10^{43}$ erg s⁻¹, so the fundamental parameter driving the position of the break frequency is not strictly related to the luminosity. One intriguing possibility is that although both objects have a similar luminosity, they may have different black hole masses. It is quite possible that the break frequency scales linearly with black hole mass, as would be expected if the timescale for the break corresponds to a characteristic time-scale of the accretion disk or the characteristic size scale of the system. If this is the case, then MCG-6-30-15 must have a significantly smaller black hole than NGC 3516 and must be accreting at a much higher fraction of its Eddington limit.

Let us assume that the break frequency scales linearly with black hole mass, and that the mass of the black hole in Cygnus X-1 is $10 M_{\odot}$. The high-frequency break in Cygnus X-1 in the low state is typically ~ 2 Hz, but can vary between 1–6 Hz, while the low-frequency break is typically at ~ 0.1 Hz but can vary between 0.04–0.4 Hz. If the knee model is correct, the knee frequency should correspond to the low-frequency break in Cygnus X-1 (i.e. where the power-spectral slope flattens to $\alpha = 0$). Scaling the knee frequency measured for NGC 3516 (see Table 4.4) with the typical value for Cygnus X-1

yields a black hole mass of $\simeq 1.6 \times 10^6 M_{\odot}$.

Such a low black hole mass implies super-Eddington accretion, since the Eddington luminosity $L_{\text{Edd}} \simeq 1.3 \times 10^{38} M_{\text{BH}} \text{ erg s}^{-1}$, (where M_{BH} is the black hole mass in solar masses) is approximately $2 \times 10^{44} \text{ erg s}^{-1}$, compared with a bolometric luminosity $L_{\text{bol}} \simeq 4 \times 10^{44} \text{ erg s}^{-1}$ estimated using the bolometric correction $L_{\text{bol}} \simeq 27L_{2-10}$ (Padovani & Rafanelli 1988). The best estimate of the mass from the knee model is thus inconsistent with the assumption that the source is in the low state. The upper limit to the mass according to the knee model is $\simeq 1.25 \times 10^7 M_{\odot}$, taken by scaling the 90% upper limit to the break frequency with the highest measured low-frequency break in Cygnus X-1 (at 4 Hz).

By scaling the frequency break for NGC 3516 measured by the high-frequency break model with the high-frequency break of Cygnus X-1, we estimate a black hole mass of $\sim 10^7 M_{\odot}$, with upper limit $\sim 10^8 M_{\odot}$. The black hole mass of NGC 3516 has been independently estimated using optical reverberation mapping techniques to be $\sim 2 \times 10^7 M_{\odot}$ (Wanders & Peterson 1994), similar to the best estimate from the high-frequency break model. Therefore, not only does the high-frequency break model provide a better fit to the power spectral shape of NGC 3516 than the knee model, the black hole mass it estimates is consistent with an independent measure and is consistent with sub-Eddington accretion of the order of 20% of the Eddington limit, consistent with the comparison with the power spectrum of Cygnus X-1 in the low state.

The break frequency for MCG-6-30-15 is significantly higher than that for NGC 3516, so that the best estimate for its black hole mass, assuming the high-frequency break model, is very low at $4 \times 10^5 M_{\odot}$ implying strongly super-Eddington accretion. Even the estimated upper limit to the black hole mass ($5 \times 10^6 M_{\odot}$) implies accretion at a substantial fraction (40%) of the Eddington limit, which seems inconsistent with the comparison with Cygnus X-1 in the low state. An intriguing alternative possibility is that MCG-6-30-15 is in a state analogous to the high state of Cygnus X-1, which has a characteristic

steep ($\alpha \simeq 2$) high-frequency power spectrum breaking at 10 Hz to $\alpha \simeq 1$, with no low-frequency breaks down to frequencies as low as 10^{-2} Hz (Cui et al. 1997b). Inspection of the power spectrum of MCG-6-30-15 in Figure 4.11 shows no evidence for any residuals due to flattening at low frequencies which might contradict the hypothesis that MCG-6-30-15 has a high-state power spectrum. Scaling by the high state frequency-break yields an estimated black hole mass of $2 \times 10^6 M_{\odot}$ (upper limit $\sim 8 \times 10^6 M_{\odot}$), consistent with close-to-Eddington accretion (which we might expect from a high state source) but not super-Eddington accretion.

Although the black hole mass in MCG-6-30-15 has not yet been measured by reverberation mapping, Reynolds (2000) has pointed out that the low luminosity (absolute B magnitude $\simeq -19$) of the S0 host galaxy of this AGN corresponds to a black hole mass of $\sim 10^7 M_{\odot}$ if the black hole mass-bulge luminosity relation is the same as that determined by Magorrian (1998) for a sample of normal elliptical and S0 galaxies. I note here that the black hole masses of all the Seyfert galaxies which have been reverberation mapped show smaller black hole masses by a factor of ~ 10 than the Magorrian relation suggests (Wandel 1999). This discrepancy may represent a systematic difference between Seyfert galaxies and normal galaxies which I shall not touch on here, except to note that it implies a black hole mass for MCG-6-30-15 which is significantly lower than $10^7 M_{\odot}$, consistent with the mass suggested by the high-state interpretation of the power spectrum.

The power-spectral shape of NGC 5506 is also consistent with a high state interpretation of this source, although the uncertainties on the break frequency are sufficiently large that a low-state interpretation would suffice to fit the data self-consistently. I note here however that the low luminosity of the NGC 5506 galaxy (absolute B magnitude $\simeq -20$), suggests that the central black hole mass is relatively low. Furthermore, both NGC 5506 and MCG-6-30-15 have quite steep 2–10 keV X-ray spectra compared to other Seyfert galaxies (power-law photon index $\Gamma \sim 2$, McHardy, Papadakis & Uttley 1998; Lamer, Uttley

& McHardy 2000) which is also suggestive of the high-state analogy since the hard power-law component of the X-ray spectrum of Cygnus X-1 is significantly softer in the high state than in the low state. (Cui et al. 1997a)

The frequency break in NGC 5548 is not significantly detected and so is not well defined. The black hole mass of NGC 5548 measured from reverberation mapping is $\sim 10^8 M_{\odot}$ (Wandel, Peterson & Malkan 1999; Kaspi et al. 2000), implying that the accretion rate is relatively low ($< 10\%$ Eddington), so applying the low state interpretation we can estimate that the high-frequency break occurs at $\sim 2 \times 10^{-7}$ Hz, and may become detectable with further monitoring observations.

4.6 The energy dependence of RMS variability and the power spectrum

Studies of the X-ray variability of black hole X-ray binaries suggest that the power-spectral amplitude changes as a function of the energy band of the lightcurve and that, in the case of Cygnus X-1 at least, the power-spectral slope at high frequencies is steeper at lower energies. It is therefore of interest to determine whether AGN variability is also energy-dependent.

4.6.1 Energy dependence of long-timescale fractional RMS

Lin et al. (2000) investigate the energy dependence of fractional RMS variability for a sample of 4 black hole X-ray binary systems including Cygnus X-1. To do so, they measure the fractional RMS variability of each source integrated over the 0.002–10 Hz frequency range, in three energy bands, classed as low (2–5 keV), medium (5–10 keV) and high (10–20 keV). They find a surprising variety of behaviours, with the fractional RMS increasing with energy for

	2–5 keV			5–10 keV			10–20 keV		
	μ	σ^2	σ_{frac}	μ	σ^2	σ_{frac}	μ	σ^2	σ_{frac}
MCG-6-30-15	7.44	4.65	29%	6.59	2.47	24%	2.11	0.21	22%
NGC 5506	12.78	9.32	24%	13.99	8.96	21%	4.65	0.74	18.5%
NGC 5548	8.02	5.07	28%	7.61	3.43	24%	2.56	0.35	23%
NGC 3516	6.45	5.02	35%	6.83	2.94	25%	2.44	0.27	21%

Table 4.6: Energy dependence of RMS variability.

the source 1E 1740.7-2942, decreasing with energy for the source GX 339-4, Cygnus X-1 showing both these behaviours and GRS 1758-258 showing a peak in RMS variability in the medium energy band.

In order to test whether the fractional RMS of the Seyfert galaxies in our sample is energy dependent, I made long-term monitoring lightcurves for each source in the same low, medium and high energy bands as used by Lin et al.. Then, I calculated the mean, variance and fractional RMS variability for each object using the same technique of equally weighting time bins as used to calculate the same values for the 2–10 keV lightcurves in Section 4.2 (using the long time-scale bin widths). The variance of the monitoring lightcurves is equivalent to the integrated power in the 10^{-8} – 10^{-3} Hz frequency range, so that after scaling the frequencies by the black hole mass, the frequency range is roughly equivalent to the 0.002–10 Hz range used by Lin et al. (as is evidenced by the similar fractional RMS values of 20–40%). The resulting mean count rate (μ), noise-subtracted variance (σ^2) and fractional RMS (σ_{frac}) values for each energy band are shown in Table 4.6, which shows that the fractional RMS variability of all the objects in the sample increase towards lower energies. Note that because the variations in different energy bands are correlated, the *relative* variances and fractional RMS between bands are accurate even though there are some small uncertainties in the absolute variance and fractional RMS due to the longest time-scale variations which contribute

to the variance. Additionally, the contribution of the contaminating BL Lac object, 1E 1415.6+2557 has not been accounted for in determining the RMS values for NGC 5548. However as 1E 1415.6+2557 shows a steeper energy spectrum than NGC 5548 (and so contributes relatively more to the mean flux at lower energies) and does not contribute significantly to variability, its effect will probably be to reduce the inverse correlation between fractional RMS variability and energy.

An increase in fractional RMS variability towards lower energies could be explained if the power-law X-ray continuum becomes steeper with increasing X-ray flux, so that as the flux increases by a small amount at high energies, the spectrum becomes steeper so that the relative variation is larger at lower energies. This picture is consistent with earlier reports of just this type of spectral variability in the Seyferts Mkn 766 (Leighly et al. 1996) and NGC 4051 (Guainazzi et al. 1996). Furthermore, in Lamer, Uttley & McHardy (2000), we present a spectral study of NGC 5506 using the monitoring and long-look data which confirms that the continuum slope is correlated with source flux in that it gets steeper as source flux increases, while Chiang et al. (2000) show the same behaviour in long-look observations of NGC 5548.

An alternative explanation for the decrease in fractional RMS with increasing energy could be that there is a constant component in the lightcurves which has a harder energy spectrum than the varying-flux component. This constant hard component could also contribute to an apparent steepening of the energy spectrum at high fluxes, even if the underlying continuum slope of the varying component of the lightcurve was constant. A good candidate for a constant flux component which has a hard spectrum is that due to reflection from distant cold material, perhaps the molecular torus. Lamer, Uttley & McHardy find evidence for a contribution from torus reflection in the spectral variability properties of NGC 5506, however even allowing for this constant hard component the source still shows spectral variation which must be intrinsic to the continuum itself. If we assume that the discrepancy between the fractional

RMS variability in the 2–5 keV and 5–10 keV energy bands is entirely due to the presence of a constant component in the 5–10 keV band which does not contribute to the 2–5 keV band, we can set a lower limit to the flux of the constant component, as follows.

The true underlying fractional RMS variability in the limit where there is no constant component, is given by the fractional RMS in the 2–5 keV band, $\sigma_{\text{frac},2-5}$. The mean flux in the 5–10 keV band, μ_{tot} is given by:

$$\mu_{\text{tot}} = \mu_V + C,$$

so that the measured fractional RMS $\sigma_{\text{frac},5-10} = \sigma/\mu_{\text{tot}}$ (where σ is the absolute RMS measured in the 5–10 keV band), but the true fractional RMS variability is equal to the ratio of σ to μ_V , so that $\sigma_{\text{frac},2-5} = \sigma/\mu_V$. Rearranging these equations we find that:

$$C = \mu_{\text{tot}} \left(1 - \frac{\sigma_{\text{frac},5-10}}{\sigma_{\text{frac},2-5}} \right).$$

Applying this equation to the data in Table 1 we find that MCG-6-30-15, NGC 5506, NGC 5548 and NGC 3516 require constant flux contributions of *at least* 17%, 12%, 14% and 29% of their total mean 5–10 keV fluxes respectively, compared to the $\simeq 10\%$ (determined from the XSPEC ‘href’ reflection model) expected from a reflecting torus which covers 50% of the sky seen from the source. Molecular tori are generally expected to present a lower covering fraction than this, so it seems likely that most (if not all) of the trend for fractional RMS to decrease with increasing energy is due to intrinsic spectral variability of the continuum (although in the next chapter, we will see direct evidence for a possible torus reflection component in the Seyfert galaxy NGC 4051).

The question remains as to why all the AGN in our sample show the same energy dependence of RMS variability while black hole X-ray binary systems show a variety of trends. Until a more systematic and extensive survey of spectral variability in AGN has been carried out, we cannot tell if decreasing fractional RMS with increasing energy is the rule for all radio-quiet AGN. However, it is worth noting that this behaviour is expected under certain Comptonisation

models for the production of the X-ray continuum, where flux increases are caused by increases in the flux of low energy seed photons, which then serve to cool down the Comptonising electrons and hence soften the X-ray continuum (e.g. Haardt, Maraschi & Ghisellini 1997). Perhaps X-ray binaries show similar behaviour, but this is sometimes masked by the presence of constant soft spectral components (analogous to that seen in the high state of Cygnus X-1), which reduce the fractional RMS at low energies. If these constant soft components are associated with the hot accretion disk in X-ray binary systems it is perhaps not surprising that we do not see them in systems with much more massive black holes, where the accretion disks are much cooler.

4.6.2 Dependence of power-spectral shape on energy

Nowak et al. (1999) have shown that above the high-frequency break, the power-spectral slope of Cygnus X-1 is dependent on energy, in that it is flatter in power spectra made from higher energy lightcurves. The change in α is only 0.2 between the 0–3.9 keV band and 8.2–14.1 keV, but it is significant and raises important questions as to why this change in slope occurs above the high-frequency break but not below. Lin et al. find no evidence for a change in power-spectral shape with energy in the other black hole X-ray binary candidates they study, GX 339-4, GRS 1758-258 and 1E 1740.7-2942, but I note here that the power spectra of all these objects do not show any evidence of high-frequency breaks either.

In this chapter, we have found that the flattening seen in the power spectra of Seyfert galaxies is consistent with that expected from high-frequency breaks analogous to those seen in the power spectrum of Cygnus X-1, except scaled down in frequency by the black hole mass. Therefore, we might expect to see an energy-dependence of power-spectral slope above the break frequency in

AGN power spectra. In order to test this possibility, I made lightcurves in the 2–5 keV and 7–15 keV bands using the same schemes outlined in Table 4.1, and fitted their power spectra with the high-frequency break model described in Section 4.4.2. Although the model fitted the 2–5 keV power spectra well, and returned the same fitted break frequencies as found from the power spectral fits in Section 4.4.2, the model was rejected at 90% confidence by the 7–15 keV observed power spectra of NGC 3516 and NGC 5506 (although adequate fits were obtained for MCG-6-30-15 and NGC 5548). Visual inspection of the best-fitting models in the 7–15 keV band showed that these poor fits were due to a systematically greater spread in the power spectral points than predicted by the simulated data, and not because of any residual trends indicating that the model was incorrect. It is possible that errors due to poorer background estimation at high energies (where the background flux is large relative to the source flux) could lead to additional spread in the power which is unaccounted for by the simulated data sets, however it is difficult to explain why this should only apply to NGC 3516 and NGC 5506. Alternatively, the source variability itself may deviate from the Gaussian-distributed stochastic variations expected from pure red-noise variability. Consideration of this issue is complex and will not be explored further here.

If we assume that the best-fitting power spectral shapes are not biased in any way by the unusually large spread in the power spectrum at high energies, then we can fix the high-frequency break at the best-fitting value for each object measured in Section 4.4.2 (except NGC 5548, whose break is not well defined, so we set it at 2×10^{-7} Hz, commensurate with the measured black hole mass) and measure the best-fitting slopes in each energy band. The resulting slopes are given in Table 4.7, which shows that there is a systematic trend for the slopes at higher energies to be flatter in all the objects in the sample, consistent with what is seen in the Cygnus X-1 data, although we cannot quantify any uncertainties on the slope differences. It is possible that this tendency is caused by systematic errors in the power spectrum at high energies. A more detailed

	$\alpha_{2-5\text{keV}}$	$\alpha_{7-15\text{keV}}$
MCG-6-30-15	2.0	1.7
NGC 5506	2.2	1.9
NGC 5548	1.6	1.3
NGC 3516	2.2	2.0

Table 4.7: Energy dependence of power-spectral slope

study requires better quality data or a better understanding of any systematic uncertainties in the high energy data, so I shall leave this for future work.

4.7 Summary

We have seen that the power spectra of three of the four Seyfert galaxies under study show significant evidence for flattening towards low frequencies, and that this flattening is consistent with that expected from a high-frequency break model power spectrum, analogous to the model which describes the high-frequency break in the power spectrum of Cygnus X-1, where the low frequency power spectral slope, α_{10} is fixed to 1. A knee model can successfully describe the power spectra of NGC 5506 and NGC 5548, which have the poorest data sets in the sample, but does not provide a very good fit to the power spectral shapes of NGC 3516 or MCG-6-30-15, which have the best data sets, suggesting that the high-frequency break model is the better model in that it is consistent with *all* the data.

The possibility that we are measuring the AGN analog of the high-frequency break in Cygnus X-1 and that the break-frequency scales inversely with black hole mass is reinforced by the estimated black hole masses of NGC 3516 and

MCG-6-30-15, derived under this assumption, which are consistent with independent estimates of the mass based on, respectively, optical reverberation mapping and considerations of host galaxy luminosity. The high break-frequency of MCG-6-30-15 suggests that this source may be in a high state analogous to that occasionally seen in Cygnus X-1. NGC 5506 may also be a high-state source although the break-frequency is not well defined. The power spectrum of NGC 5548 shows no strong evidence of a break, consistent with the high black hole mass of this AGN, estimated by reverberation mapping. It is important to note however, that the data are only good enough to determine whether a given model is consistent with the data or not, We cannot categorically say that the high-frequency break model is correct, since other power-spectral models, perhaps with more gradual flattening or low-frequency slopes with α other than 1, may also describe the data well.

The energy dependence of fractional RMS variability suggests that the slope of the medium-energy X-ray spectrum scales with X-ray flux in all four objects in the sample, consistent with spectral studies. The energy-dependence of power-spectral slope is confused by the fact that the power-spectral models which work well for all the objects in the sample at low energies (2–5 keV, 2–10 keV) do not provide a good fit the power spectra of NGC 5506 and NGC 3516 at high energies (7–15 keV). The poor fits are due to unusually large fluctuations in the observed power spectra, which may be caused by poor background subtraction or may be due to an intrinsic deviation from the expected red-noise behaviour in the high-energy lightcurves of these two sources. If we assume that the break-frequency is independent of energy and that there are no systematic trends in the power spectrum caused by the unexpectedly large fluctuations in the high-energy band, then the high-frequency power spectral slopes of all the objects are steeper in the 2–5 keV band (by $\Delta\alpha = 0.2$ – 0.3) than in the 7–15 keV band, consistent with what is seen from Cygnus X-1 (although uncertainties are difficult to quantify).

Chapter 5

X-ray, EUV and optical variability of NGC 4051

Overview

In this chapter we will first investigate the unusual variability properties of NGC 4051 in the X-ray and EUV band, revealing evidence of a variety of behaviours during the course of the *RXTE* monitoring campaign, including two quiescent low states lasting several months in early 1998 and 1999. I then discuss evidence of non-Gaussianity in the lightcurve and measure the broadband power spectrum. I also investigate EUV variability during the low state. The low state X-ray spectrum is studied, revealing strong evidence for reflection from a molecular torus and implying that the primary X-ray source switches off almost completely during the low state. Simultaneous EUV and X-ray variability during the normal active state is discussed and constraints on the lag between the two bands are used to constrain the size of the EUV/X-ray emitting region, if simple Comptonisation models are assumed. Finally, we present and discuss some key results from the ‘AGN Watch’ team’s optical monitoring campaign of NGC 4051, in the context of the relationship between the X-ray and optical continua. The optical, X-ray and EUV measurements together



suggest that the inner accretion disk disappears or changes to a radiatively inefficient mode in the low state.

5.1 The strange variability properties of NGC 4051

In this section, I shall first introduce the *RXTE* lightcurves we have obtained for NGC 4051 and then investigate their unusual properties in more detail.

5.1.1 Observations and data reduction

NGC 4051 was monitored with *RXTE* as part of our Seyfert monitoring program, using the same observing plan as used for NGC 5506, NGC 5548 and MCG-6-30-15 in AO1-AO3, i.e. \sim two weeks of twice-daily monitoring, four weeks of daily monitoring and weekly monitoring for the remainder of AO1, followed by fortnightly monitoring in AO2 and AO3. The intensive monitoring schemes were not consecutive with one another however, as the twice-daily monitoring period took place from 20 May to 29 May 1996, while the daily monitoring period lasted from 4 October to 31 October 1996. The weekly monitoring which took place throughout the remainder of AO1 (until the beginning of the daily monitoring period which led into AO2) began on 23 April 1996. In AO4, monitoring continued at fortnightly intervals.

During the period of twice-daily monitoring, we also observed NGC 4051 with the Deep Survey Spectrometer (DS/S) instrument on board the *Extreme Ultraviolet Explorer* (*EUVE*) satellite (which detects photons in the 124-188 eV energy range), to study simultaneous variability in the medium-energy X-ray and EUV bands.

From December 13 to December 16 1996, we observed NGC 4051 with *RXTE* for ~ 70 ksec useful exposure, which overlapped for 120 ksec duration with

a further continuous *EUVE* observation. The results of a combined analysis of these simultaneous *RXTE/EUVE* observations will be discussed in Section 5.3. In this section, I shall describe the general variability properties of the *RXTE* lightcurves, which were reduced in the standard manner described in Section 4.1.5. The resulting 2–10 keV long-term lightcurve (including data up until 1 March 1999, before the gain change of the PCA) and daily intensive monitoring lightcurves are shown in Figure 5.1, together with the December long-look lightcurve.

What is immediately apparent from the long-term monitoring lightcurve is that the X-ray variability process underlying the lightcurve is not stationary on long time-scales. The source enters two distinctive low-flux states, which are characterised (apart from their low mean fluxes $\sim 1 \text{ count s}^{-1}$) by a very low level of variability. The first low state lasted for ~ 150 days, from January 1998 to the end of May 1998, while the second began in late December 1998 and lasted until March 1999. We were fortunate enough to have a long-look observation of NGC 4051 scheduled in early May 1998, simultaneous with observations by *EUVE* and the *BeppoSAX* satellite which enabled us to confirm that the X-ray source was not significantly variable and measure a good-quality X-ray spectrum, which is described in the next section. In the remainder of this section however, I shall concentrate on the unusual variability properties of the long-term lightcurve and use the December long-look lightcurve to investigate the power spectrum.

5.1.2 RMS variability and evidence for non-Gaussian variations in the lightcurve

First, let us look at the RMS variability of the long-term monitoring lightcurve. The mean flux, μ (count s^{-1}), noise-subtracted variance σ^2 ($\text{count}^2 \text{ s}^{-2}$) and fractional RMS, σ_{frac} , calculated by equally weighting contributions to mean

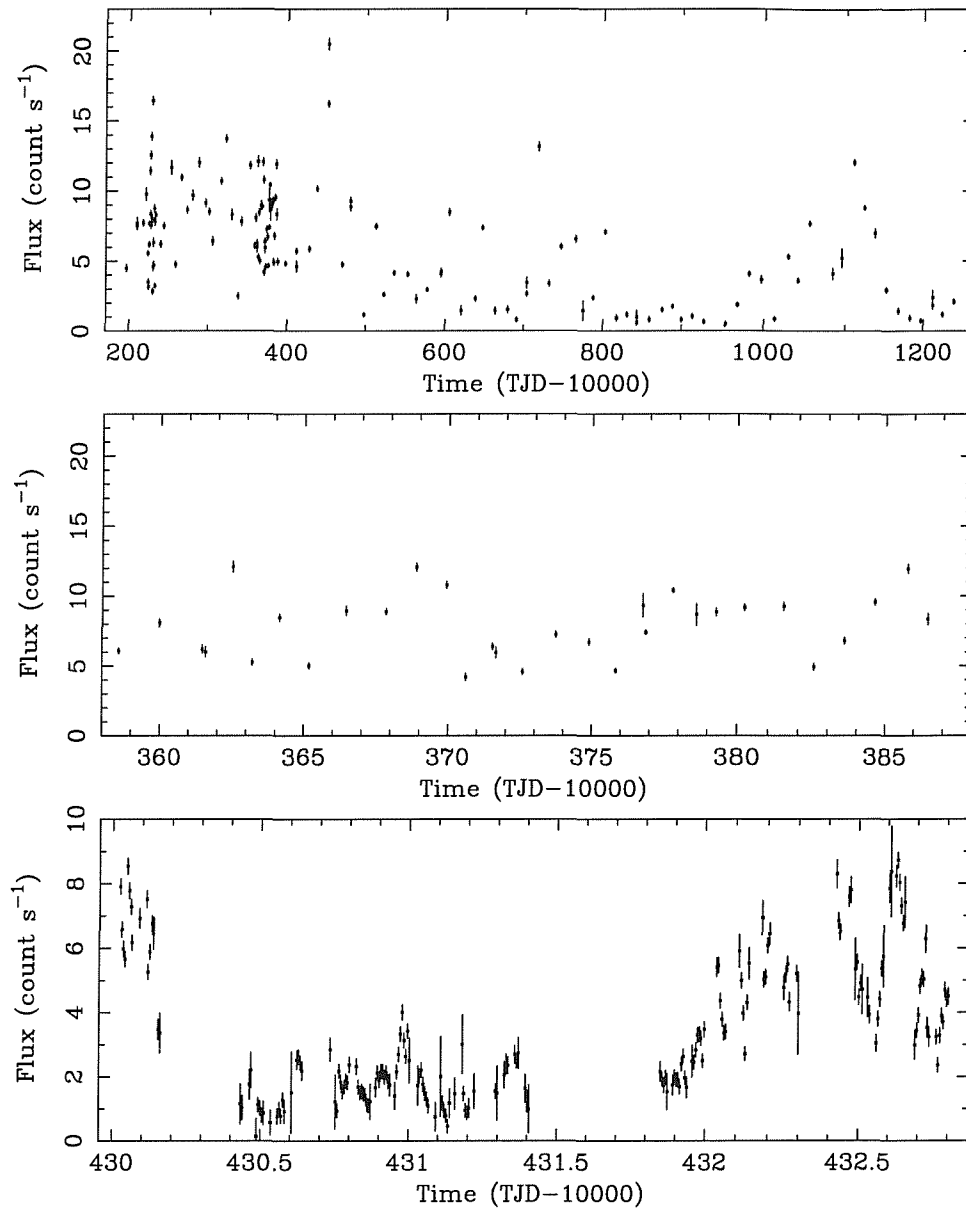


Figure 5.1: 2-10 keV lightcurves of NGC 4051, from top to bottom: total monitoring lightcurve, intensive monitoring lightcurve, December 1996 512-s binned long-look lightcurve.

Band	μ	σ^2	σ_{frac}
2–10 keV	4.86	13.91	77%
2–5 keV	2.55	5.07	88%
5–10 keV	2.31	2.21	64%
10–20 keV	0.75	0.15	52%

Table 5.1: NGC 4051 long-term lightcurve variability parameters

and variance from 2-week bins, are shown in Table 5.1 for the 2–10 keV lightcurve and each of the energy bands considered in Section 4.6.1. The fractional RMS values measured from the long-term lightcurve are exceptionally high. The reason for this is that the lightcurve mean is pulled down by the low-state periods and also by a period of generally low fluxes lasting ~ 1 year between TJD 10500 and the beginning of the first low state, so that the period of large variability observed in the first AO contributes much more to the variance than it would if it were considered in isolation. The measured RMS values are clearly distorted because the lightcurve is so strongly non-stationary.

The period from \sim TJD 10500–10800 is striking in that it seems to show evidence for non-linear variability, with the source flux typically being fairly low, punctuated by large flares, i.e. the measurements of flux are not evenly distributed about the mean. Leighly (1999) points out that a skewed distribution of flux does not immediately imply non-linearity in the lightcurve, since it is possible that the underlying lightcurve is linear but non-Gaussian. A linear non-Gaussian lightcurve might be caused by an underlying linear process which must reach a certain threshold level before triggering a corresponding flux increase in the lightcurve. Therefore a Gaussian distribution of fluxes in a lightcurve implies linearity, but a non-Gaussian distribution does not strictly imply non-linearity. A rigorous test for non-linearity has been proposed by Keenan (1985), however the Keenan test is not possible to apply here since it requires a longer, more intensively sampled time series than is available here.

Band	μ	σ^2	σ_{frac}
2–10 keV	8.12	9.95	39%
2–5 keV	4.52	3.89	44%
5–10 keV	3.60	1.42	33%
10–20 keV	1.03	0.09	29%

Table 5.2: NGC 4051 variability parameters for the Gaussian-distributed portion of the long term lightcurve

Instead, we can test the Gaussianity of the flux distribution of the lightcurve by measuring the distribution of fluxes below and above the mean flux. If the distribution is Gaussian, the distribution of points about the mean will be symmetric. We can then apply the KS test to the distributions of absolute flux deviation (i.e. the magnitude of mean flux minus measured flux) below and above the mean, to ascertain the probability that they come from the same parent distribution. Green, McHardy & Done (1999) applied this method to two *ROSAT* observations of NGC 4051 to show that the November 1991 lightcurve was significantly non-Gaussian while the November 1992 lightcurve was consistent with Gaussian variability. I have split the long-term monitoring lightcurve prior to the first low state into two equal halves, and find that the first half of the lightcurve is consistent with a Gaussian flux distribution (the distribution of fluxes above and below the mean are the same at 53% confidence) while the second half of the lightcurve (TJD 10503–10810) shows a significantly non-Gaussian distribution (the distributions of fluxes above and below the mean are not the same at 90% confidence).

If we only consider the Gaussian portion of the long-term lightcurve (i.e. up until TJD 10503), we find that the RMS values are similar to those obtained for the Seyfert galaxies described in the previous chapter (see Table 5.2). I shall discuss the variability properties of the low state in the following section. For now, let us turn to the properties of the power spectrum of NGC 4051.

5.1.3 The power spectrum

The non-Gaussian variability in the long-term lightcurve of NGC 4051 poses a problem in that the power spectrum for a non-Gaussian or non-linear process is not well defined and cannot be well fitted by any simple model, such as a power law or broken power law. Green, McHardy and Done (1999) demonstrate this by using the response method to fit models to power spectra of *ROSAT* lightcurves of NGC 4051 obtained in November 1991 and November 1992. Although the Gaussian distributed lightcurve power spectrum was well fitted by various broken power law models, the non-Gaussian lightcurve power spectrum from November 1991 was not well fitted at all by any simple power-law or broken power-law model, showing large, scattered residuals. Therefore, we are restricted to measuring the low-frequency power spectrum only from the Gaussian-distributed portion of the long-term monitoring lightcurve.

The broadband power spectrum is made from the 2–10 keV lightcurves, using the method described in Section 4.2, except that the medium time-scale lightcurve only includes the daily monitoring scheme and not the twice-daily scheme, as the two were not consecutive. The long time-scale lightcurve used to make the low-frequency power spectrum includes data only up to TJD 10503. The high-frequency power spectrum was made using the December long-look lightcurve, which was also used to measure the VHF power spectrum, in the same manner as for MCG-6-30-15 and NGC 5506. Mean fluxes for the long time-scale, medium time-scale and long-look lightcurves are 8.15, 7.73 and 3.31 count s⁻¹ respectively, while the corresponding Poisson noise levels in the power spectrum are 39.5, 38.7 and 29.1 count² s⁻² Hz⁻¹.

Fitting the measured broadband power spectrum with a high-frequency break model, we find that the best-fitting probability is only 3%, corresponding to a break frequency of 6.55×10^{-4} Hz and high-frequency slope $\alpha = 2.3$. Since the range of break frequency stepped over was large and only 100 simulated lightcurves were used for each observing scheme, this ‘optimum’ fit is consis-




Figure 5.2: Comparison of the best-fitting high-frequency break model average power spectrum with the observed broadband power spectrum of NGC 4051.

tent with noise in the fitting process. There is no preferred break frequency (including the possibility of no break at all). If we examine the best-fitting power spectrum, shown in Figure 5.2, we see that the problem stems from the high-frequency power spectrum which shows much wilder fluctuations than expected from the stochastic variations of a red-noise power spectrum. Fitting the high-frequency power spectrum alone, with a simple power-law model, we find that the best-fitting probability is only 7%, again due to the wild fluctuations in the power spectrum.

If we examine the December 1996 long-look lightcurve (see Figure 5.1) used to make the high-frequency power spectrum, we see that it looks suspiciously non-Gaussian. The uneven sampling and relatively short duration of the lightcurve do not allow us to use the KS test to determine if it is non-Gaussian, but note that the fractional RMS variability of the lightcurve is 63%, higher than that measured for the Gaussian-distributed part of the long-term monitoring lightcurve. The December 1996 long-look observation was obtained between TJD 10430 and TJD 10433, only 70 days prior to the start of the period we have

defined as the non-Gaussian part of the lightcurve. Therefore, it seems likely that the period of non-Gaussian behaviour in the lightcurve began some time prior to the December long-look observation, and only becomes apparent at a later date because of the uncertainties in the form of the underlying long-term lightcurve caused by the sparse sampling. The large scatter in the December power spectrum is caused by the non-Gaussianity, in the same way that the scatter in the November 1991 power spectrum measured by Green, McHardy & Done was associated with non-Gaussianity in the *ROSAT* lightcurve. Therefore, we cannot measure a valid broadband power spectrum for NGC 4051 because the different parts of the power spectrum are associated with different lightcurve behaviours.

5.2 The low states

We shall now investigate the variability and spectral properties of the low states of NGC 4051.

5.2.1 Variability of the low state

Inspection of Figure 5.1 shows that the long term variability (i.e. on time-scales of a month) is minimal during the first and second low states. The low level variability that can be seen is probably real, although due to the particularly low flux levels, we should be cautious about interpreting such variations, since systematic variations due to poor background subtraction might also contribute a low level of variability. During the first low state, we were fortunate enough to observe NGC 4051 with two satellites with imaging capabilities, *EUVE* and the *BeppoSAX* satellite, allowing more confident background




Figure 5.3: Simultaneous *EUVE* (DS/S instrument) and 4–10.5 keV *BeppoSAX* (MECS instrument) lightcurves of NGC 4051 in the 1998 low state, binned to 22176 s (4 *EUVE* orbits).

subtraction and constraints on any variability. *EUVE* observed NGC 4051 from 01:10 UT May 8 to 11:48 UT May 15 1998, while *BeppoSAX* observed NGC 4051 from 09:57 UT May 9 to 18:05 UT May 11 1998. Lightcurves from these observations are shown in Figure 5.3. The medium energy X-ray lightcurve shows no significant evidence of variability, nor does the EUV lightcurve on the 20 ksec time-scale at which it is binned. However, the EUV lightcurve seems to show evidence of a slow trend of increasing flux as the observation progresses. This is confirmed by splitting the entire EUV lightcurve into two bins of 324 ksec duration, which yields significantly different fluxes (to better than 99% confidence) in each half of the lightcurve. Therefore it seems that although variability on short time-scales (<day) is not significant, there is significant variability on longer time-scales, at least in the EUV lightcurve. Singh (1999) reports evidence from *ROSAT* HRI observations of significant extended soft X-ray emission in NGC 4051, which may be associated with the nuclear outflow seen in this source at optical and radio wavelengths (Christopoulou et al. 1997). Converting the measured *ROSAT* HRI count rates due to the

extended emission to EUVE DS/S count rates using the PIMMS software, gives a contribution to the *EUVE* DS/S flux of $\sim 5 \times 10^{-3} \text{count s}^{-1}$ (assuming a thermal bremsstrahlung model of temperature 1 keV). The estimated *EUVE* flux contributed by the extended emission is not particularly sensitive to the model used, in that an extreme 10 keV bremsstrahlung model yields a rate of $\sim 3 \times 10^{-3} \text{count s}^{-1}$ while smaller temperatures yield higher fluxes. The mean *EUVE* flux over the entire low state observation is $\sim 7 \times 10^{-3} \text{count s}^{-1}$, implying that a significant proportion of the low state *EUVE* flux originates from the extended emission region, so that the flux intrinsic to the central source is even lower than that measured.

During the second low state in early 1999, we observed NGC 4051 for 2 orbits each day with *EUVE* for two 5 day observing periods separated by 8 days, to coincide with optical monitoring being conducted by the ‘AGN Watch’ team. The resulting lightcurve is shown in Figure 5.4, which shows that although the source was in the low state for the first four days of the first observing period (March 8–12 1999), it had suddenly recovered to its normal state before the observation on March 13, i.e. the recovery-time from the low-state was constrained to be less than one day. Note that the EUV flux from NGC 4051 prior to its exit from the low state is greater than that seen during the low state in May 1998, which indicates that significant EUV flux from the central source remains during the low state (i.e. it has not completely switched off in the EUV band) and that the central source varies significantly in the EUV on long time-scales.

5.2.2 The low state X-ray spectrum

We now investigate the spectrum of the source in its low state, as measured by the PCA on board *RXTE* and the MECS instrument on *BeppoSAX* during

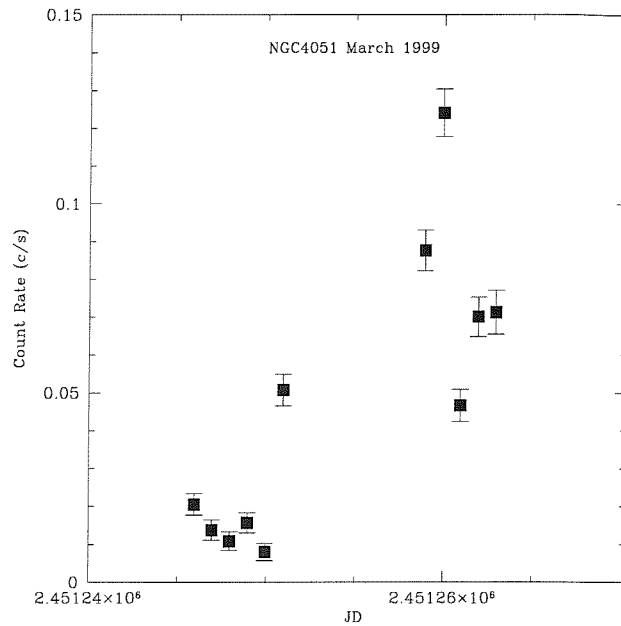


Figure 5.4: *EUVE* lightcurve of NGC 4051 during the 1999 low-normal state transition.

the long-looks of May 9-11 1998. The *BeppoSAX* data was first considered separately by Guainazzi et al. (1998), henceforth G98, who found that the spectrum at medium X-ray energies was particularly hard and showed a prominent iron line, consistent with what might be expected if the central source had switched off in medium X-rays to leave a constant hard component due to reflection of the central continuum from the distant molecular torus. Here, we shall constrain the spectral parameters of this possible torus reflection spectrum by fitting a reflection model to the *BeppoSAX* MECS spectrum together with that measured from the *RXTE* PCA. The 2-10 keV lightcurve obtained by the PCA shows no significant variability above the expected level for systematic errors in the background estimation, consistent with the observed lack of variability in the *BeppoSAX* lightcurves (G98). We therefore use the PCA and MECS spectra integrated over the whole observation.

We will not consider the data from *EUVE* and the LECS instrument on board

BeppoSAX in our fits, because these data show evidence for a separate low-energy component at energies below 4 keV, in addition to the component seen at medium energies by the PCA and MECS. As described in the preceding section, the bulk of this low-energy component is almost certainly associated with the extended soft X-ray emission discovered by Singh (1999). Since we are interested in the medium energy spectral component, we shall only consider the PCA and MECS spectra in the energy ranges 4–15 keV and 4–10.5 keV respectively. We use a PCA response matrix generated by the PCARSP v2.36 script; details of the MECS calibration and data reduction can be found in G98.

We fit the spectra in XSPEC v10.0. Simple power-law fits show a very flat spectrum so, as in G98, we shall attempt to account for this hard spectrum in terms of a reflection model. By fitting up to 15 keV we can use the simple `href` multiplicative model for reflection of a power-law spectrum off a slab of cold material. We also include a Gaussian iron line and Galactic absorption. This simple model approximates the reflection spectrum of cold material with an unknown distribution around the primary X-ray source. Like G98, we assume that the reflecting material subtends 2π steradians of sky, as seen from the source of the incident continuum. The inclination angle of the reflector to the line of sight is unknown, but since it does not significantly affect the fits, we freeze it arbitrarily at 30 degrees. We find that the best-fitting observed fraction of the illuminating power-law continuum in all our model fits where it is left free is zero (i.e. the source has switched off completely); so we fix this parameter to zero for the purpose of constraining the other model parameters. In Table 5.3, we show the resulting best-fitting parameters for separate model fits to the PCA and MECS spectra. Both sets of data are fitted reasonably well by the model and the model parameters are consistent with being the same in both the PCA and MECS spectra. The agreement between the two instruments confirms the accuracy of the PCA background model. We therefore attempt to constrain the model parameters further by fitting both the

	Γ	$A/10^{-2}$	E_K	σ_K	$F_K/10^{-5}$	$\chi^2/\text{d.o.f.}$	L_{2-10}
PCA	$2.1 \pm_{0.6}^{0.4}$	$0.6 \pm_{0.5}^{1.0}$	$6.42 \pm_{0.17}^{0.18}$	$0.29 \pm_{0.29}^{0.47}$	$2.7 \pm_{1.1}^{2.0}$	16.4/26	0.30
MECS	$2.2 \pm_{0.5}^{0.4}$	$0.9 \pm_{0.5}^{1.1}$	$6.46 \pm_{0.11}^{0.32}$	$0 \pm_0^{0.56}$	$1.60 \pm_{0.6}^{1.2}$	28.2/22	0.39
PCA/MECS	2.30 ± 0.25	$1.0 \pm_{0.3}^{0.7}$	$6.46 \pm_{0.09}^{0.16}$	$0.1 \pm_{0.1}^{0.35}$	$2.0 \pm_{0.6}^{0.8}$	49.0/53	0.38

Table 5.3: Best-fitting reflection model parameters for PCA, MECS and combined MECS and PCA spectra. Γ is the photon index of the incident power-law continuum, A is the incident power-law normalisation, E_K and σ_K are the line energy and actual width respectively (in keV) and F_K is the line flux in photons $\text{cm}^{-2} \text{s}^{-1}$. All errors are 90% confidence limits for 2 interesting parameters. Also shown is the χ^2 value and number of degrees of freedom for each fit, and the 2–10 keV luminosity of the incident power-law in units of $10^{42} \text{ erg s}^{-1}$ (assuming $H_0 = 50 \text{ km s}^{-1} \text{ Mpc}^{-1}$).

PCA and MECS spectra jointly with the same model parameters. The resulting best-fitting parameters are also shown in Table 5.3. Figure 5.5 shows the model fitted jointly to the spectra from both instruments. The inferred slope of the illuminating continuum, as obtained by the joint fit, is higher than that obtained from the individual fits to both the PCA and MECS spectra. The higher slope is due to the improved definition of the continuum flux at lower energies by the MECS data, which sets the 1 keV continuum normalisation to a higher value than that given by the PCA fit alone, combined with the greater sensitivity of the PCA at high energies which holds down the continuum at higher energies. As one might expect, given that none of the illuminating continuum is directly visible, the slope of that continuum is not well determined. However, we note that a continuum photon index of 2.3 was observed during the simultaneous *RXTE* and *EUVE* observations in May 1996 (see next section).




Figure 5.5: May 9-11 *RXTE* PCA and *BeppoSAX* MECS spectra for the best-fitting multiplicative reflection model described in the text.

The value inferred for the luminosity of the primary continuum incident on the reflector is fairly typical for NGC 4051 in its active state (e.g. Guainazzi et al., 1996). Note that the inferred value for the incident luminosity assumes the slab geometry which is inherent in the reflection model (i.e. 50% covering fraction). Since the inferred continuum luminosity is compatible with observations, the actual covering fraction must be of this order. Fixing the continuum slope of the combined-fit model to its best-fitting value, we can set a 99% confidence upper limit (for 2 interesting parameters) of 0.024 for the fraction of the illuminating primary flux which is directly observed. Combining this observation with the lack of variability from the EUV to medium X-ray bands, the simplest assumption is that the primary continuum has switched off completely.

The iron line parameters are not strongly affected by the parametrization of the underlying continuum. The iron line equivalent width is ~ 1 keV, consistent with the interpretation that the entire medium-energy spectrum originates in cold reflecting material. The line energy and width are also consistent with this interpretation.

Given that the May 1998 long-looks occur towards the end of the low state, it would seem that the reflecting material must lie some considerable distance from the central continuum, in order that we can still see reflection of the source in the active state. Therefore the distance of the reflector from the central continuum must be > 150 light-days, consistent with the interpretation that we are witnessing reflection from the molecular torus, which is assumed to lie at distances \sim light-years from the central source. The fact that the X-ray power-law continuum seems to disappear during the low states (or at least drop to a very low level) suggests that the inner disk, which presumably powers the X-ray emitting region, has disappeared, possibly becoming radiatively inefficient (e.g. an ADAF). We will consider more evidence for the disappearance of the inner disk in the light of results from optical spectral observations during the low state, later in this chapter.

5.3 Simultaneous EUV and X-ray variability

In this section, we shall consider the interband EUV and X-ray variability properties of NGC 4051 during periods of normal activity, as measured by simultaneous observations of NGC 4051 with *EUVE* and *RXTE* in May and December 1996. In particular, we shall use the data to constrain any lags between the medium X-ray and EUV band, and use these constraints to constrain Comptonisation models for the power law continuum.

5.3.1 The EUV-X-ray correlation

We show the background-subtracted *RXTE* (2–10 keV) and *EUVE* lightcurves for the May observations in Figure 5.6. For presentation purposes, we have scaled up the *EUVE* lightcurve by a factor of 100. The bin width is 5544 s,

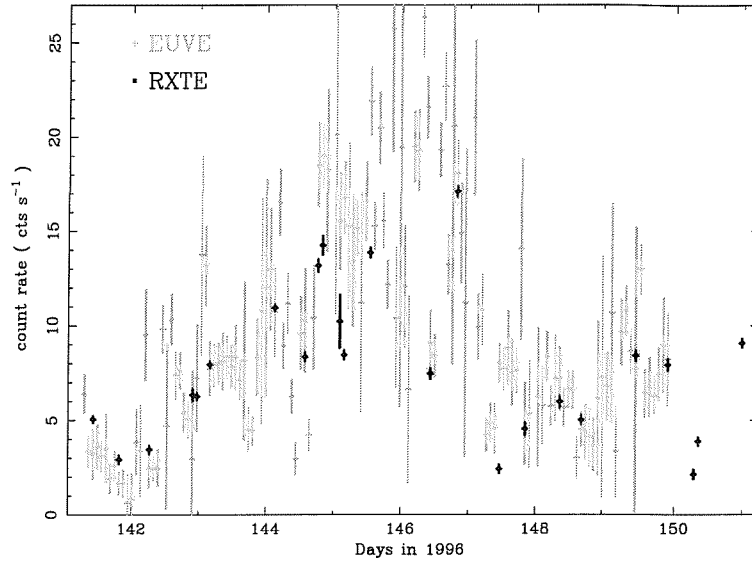


Figure 5.6: May 1996: 5544 s binned *EUVE* (scaled by factor 100) and *RXTE* (3 PCUs, 2–10 keV) lightcurves of NGC 4051 (errors are 1σ).

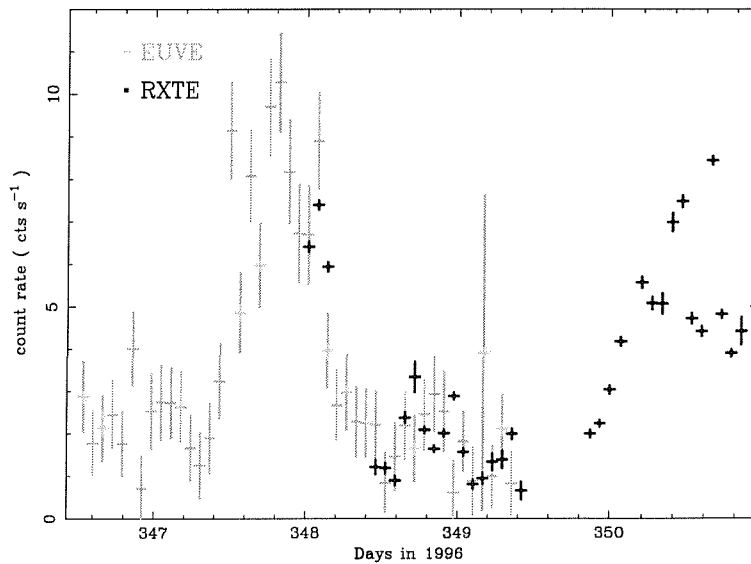


Figure 5.7: December 1996: 5544 s binned *EUVE* (scaled by factor 100) and *RXTE* (3 PCUs, 2–10 keV) lightcurves of NGC 4051 (errors are 1σ).




Figure 5.8: Z-transformed discrete cross-correlation function of *EUVE* and *RXTE* (4–10 keV) lightcurves, combined from separate May and December ZDCFs. The ZDCF is binned to 20 ks and errors are 1σ .

which is the orbital period of *EUVE*. For clarity we do not show the May 6th observations (but we note that the 2–10 keV and *EUVE* count rates for this time were ~ 8.0 cts s^{-1} and 0.12 cts s^{-1} respectively). The lightcurves for the December observations are shown in Figure 5.7, using the same bin width and the same scaling factor for the *EUVE* lightcurve including, for completeness, the data obtained by both instruments outside the times of overlap.

Visual inspection of the lightcurves shows a striking correlation between the two bands. We can quantify this correlation and search for lags between the two bands by carrying out a cross-correlation analysis of the lightcurves. In the context of this work we are particularly interested in the relation of the power-law component of the X-ray spectrum to the EUV emission. Since NGC 4051 is known to display both variable low energy X-ray absorption and an iron fluorescent emission line at ~ 6 keV (Guainazzi et al. 1996), we will exclude these components and obtain the ‘pure’ power-law contribution to the emission by extracting a lightcurve in the 4–10 keV band, excluding 5–7 keV (corresponding to PCA channels 12–14 and 21–28 in the current gain epoch

3).

We compute the cross-correlation function (CCF) of the time series using the Z-transformed Discrete Correlation Function (ZDCF) method of Alexander (1997), which is based on the DCF method of Edelson & Krolik (1988) but estimates errors more reliably. We calculate separate ZDCFs for May and December 1996 and bin up the resulting noisy ZDCFs (which have different binning) into identical 20-ks-wide bins, before combining them by adding together values corresponding to the same lag (weighting according to errors). We show the resulting combined ZDCF in Figure 5.8.

A simple visual inspection of the ZDCF indicates that the EUV band leads the X-rays by between 0 and 20 ks. Simulations of perfectly correlated lightcurves (with zero lag) which use the same sampling pattern as our data show that the peak at 20 ks is probably artificial, due to an excess of data pairs sampled at that particular lag. Therefore, we believe that the most likely lag is within the 0 ks bin. We now turn to an alternative technique for constraining the lag, which is not affected by the sampling pattern.

5.3.2 A simple spectral model

The simplest explanation for the strong correlation between the *EUVE* and *RXTE* lightcurves is that both instruments are sampling the same continuum, i.e. *the X-ray power-law extends to the EUV band*. A detailed spectral analysis of the *RXTE* data will be described in a later work, but here we can test the hypothesis that the EUV continuum is an extension of the X-ray power-law continuum by fitting a simple power-law to the 4–10 keV (excluding 5–7 keV) region of the spectrum for the brightest observation in May. We find a power law slope (photon index, Γ) of 2.3 ± 0.1 , corresponding to a 2–10 keV flux of 6.5×10^{-11} ergs cm $^{-2}$ s $^{-1}$. If we use these parameters in the PIMMS v2.3 count rate calculator, assuming a galactic absorption of 1.3×10^{20} cm $^{-2}$, we

obtain a predicted *EUVE* DS/S (Lexan/B) count rate of 0.27 ± 0.1 cts s⁻¹. This agrees well with the actual *EUVE* count rate (simultaneous with the 500 s long *RXTE* observation) of 0.24 ± 0.05 cts s⁻¹. Note that because the *EUVE* count rate can be adequately explained as the extrapolation of the X-ray power law modified by Galactic absorption, there is no requirement for a significant neutral column in the AGN host galaxy, a result consistent with previous *ROSAT* and *ASCA* observations. The lack of any significant absorbing column (in addition to Galactic) implies that the significant low energy X-ray absorption seen in previous observations of NGC 4051 (M^cHardy et al., 1995; Guainazzi et al., 1996) must be due to ionised gas (so that hydrogen and helium are virtually completely ionised). This is again consistent with previous observations, which indicate that the X-ray absorption is due to an ionised ‘warm absorber’.

Previous observations of NGC 4051 with *EXOSAT* (Papadakis & Lawrence 1995) and recently *ASCA* (Guainazzi et al. 1996), indicate that the photon index of the X-ray power-law is positively correlated with the source luminosity, in the sense that the continuum becomes softer as the source flux increases. An inspection of the scaled *EUVE* and *RXTE* lightcurves indicates that this may indeed be the case; there is a tendency for the linearly scaled *EUVE* flux to exceed the 2–10 keV count rate when the source is bright, and vice versa when the source is dim. If the *EUVE* count rate can be described as a simple function of the *RXTE* count rate, we can fit such a function to the data. The scaled lightcurves indicate that the *EUVE* count rate scales with the *RXTE* count rate in a non-linear way. We suggest a simple function of the form:

$$R_{\text{EUVE}} = A R_{\text{RXTE}}^n + C$$

where R_{EUVE} and R_{RXTE} are the predicted *EUVE* count rate and the *RXTE* (4–10 keV, excluding 5–7 keV) count rate respectively. A , n and C are constants. The error on the predicted values is given by:

$$\Delta R_{\text{EUVE}} = (R_{\text{EUVE}} - C) \frac{n \Delta R_{\text{RXTE}}}{R_{\text{RXTE}}}$$

where ΔR_{EUVE} and ΔR_{RXTE} are the errors on the *EUVE* and 4–10 keV count rates respectively.

For given parameters A , n and C we can calculate the predicted value (and error) of R_{EUVE} for each point in the *RXTE* lightcurve and compare with the actual values by calculating a χ^2 value, defined as:

$$\chi^2 = \sum^n \frac{(R_{\text{EUVE,measured}} - R_{\text{EUVE}})^2}{\Delta R_{\text{EUVE,measured}}^2 + \Delta R_{\text{EUVE}}^2}$$

Where $R_{\text{EUVE,measured}}$ and $\Delta R_{\text{EUVE,measured}}$ are the measured values of the *EUVE* count rate and associated error respectively, and the sum is over the n pairs of data points that are measured simultaneously in both bands. By stepping through a range of values of the parameters we can attempt to find the set of best-fitting model parameters.

We use the 4–10 keV band of the *RXTE* data (excluding 5–7 keV), so that we sample only the continuum component of the X-ray spectrum. Due to orbital constraints there are gaps in the *EUVE* lightcurve so that not all *RXTE* data points have corresponding simultaneous *EUVE* data points. We therefore bin the lightcurves into identical 1 ks bins (so the corresponding time for each bin is the same for both the *EUVE* and 4–10 keV lightcurves). In this way we only sample times where the *EUVE* and *RXTE* data are simultaneous to within 1 ks. This restriction is necessary because large flux changes can occur on timescales of a few ks, so the flux in both bands must be measured as simultaneously as possible to allow a good comparison between bands. We next fit the simultaneous data with our model.

We first fit a simple linear model to the data ($n = 1$). We obtain best-fit values of $A = 0.029 \pm_{0.004}^{0.005}$ and $C = -0.006 \pm_{0.007}^{0.005}$, with $\chi^2 = 35.14$ for 31 degrees of freedom. This model is acceptable, but a non-linear model (allowing n to be a free parameter) improves the fit ($\chi^2/d.o.f. = 30.00/30$), with best-fit values $A = 0.012 \pm_{0.006}^{0.013}$, $n = 1.68 \pm 0.55$ and $C = 0.006 \pm_{0.011}^{0.007}$. Negative and positive values of C (as given by the linear and non-linear model fits) correspond to constant flux components in the *RXTE* and *EUVE* bands respectively. The




Figure 5.9: Comparison of linear and non-linear models (dotted line and solid line respectively). For clarity the data have been averaged (and 1σ errors calculated accordingly) into bins of $0.3 \text{ counts s}^{-1}$ width. See text for details of corresponding fit parameters.

constant soft component, which may be largely associated with the extended emission in the host galaxy imaged by *ROSAT* (Singh 1999), contributes an EUV flux consistent with the constant value estimated by our non-linear model fit. However, the constant component from the molecular torus would lead to a negative value of C . The exact value of C , which is at present not well constrained, is determined by a trade-off between these two constant components. Using our data, we cannot formally rule out either the non-linear or linear model, however the F -test indicates that the non-linear model is better at describing the data at 63% confidence. In Figure 5.9 we show a comparison of the linear and non-linear models with the data.

The fact that we can successfully model (reduced $\chi^2 = 1.0$) the relationship between both energy bands to a simultaneous time resolution of 1 ks implies that there is no significant lag between the bands (i.e. any lag is less than

1 ks). This result is consistent with the less-than-unity value of the zero-lag peak in the lower resolution CCF. We can rule out the possibility that our model is so general that it will fit any time lag by shifting the lightcurves with respect to each other and fitting the model. If we cause the *EUVE* lightcurve to lag the *RXTE* lightcurve by 1 ks, the best-fitting non-linear model yields $\chi^2/d.o.f. = 65.78/23$. If we cause the *RXTE* lightcurve to lag the *EUVE* lightcurve by 1 ks we obtain $\chi^2/d.o.f. = 39.38/35$. If the *RXTE* lightcurve lags the *EUVE* lightcurve by 2 ks we find $\chi^2/d.o.f. = 60.78/38$. Clearly the model does not easily fit other lags, although a lag of 1 ks between the X-ray and EUV bands is allowed by the model. We might expect the model to fit short lags even if the true lag is zero, since both lightcurves are autocorrelated on short timescales (\sim ks). In any case, it remains true that effectively zero lag between the bands is adequate to describe the data.

Recent simultaneous observations of NGC 5548 with *RXTE* and *EUVE* showed evidence that the 2–20 keV X-ray variations lagged changes in the EUV by 35 ks (Chiang et al. 2000). Scaling this lag by the black hole mass, using the masses of NGC 4051 ($1.4 \times 10^6 M_{\odot}$) and NGC 5548 ($10^8 M_{\odot}$) measured by reverberation mapping (Wandel, Peterson & Malkan 1999), we predict an X-ray–EUV lag of ~ 600 s in NGC 4051, consistent with our upper limit.

5.3.3 Constraining simple Comptonisation models

We have shown that the variable X-ray and EUV fluxes almost certainly come from the same spectral component, and that the best-fitting model to describe the relationship between the fluxes in both bands is probably non-linear, implying that increases in continuum flux are accompanied by spectral slope changes. Specifically, the power-law steepens as the continuum flux increases. This behaviour was observed previously by *ASCA* (Guainazzi et al. 1996). Using PIMMS we can estimate the change in power-law slope associated with a

doubling of the *RXTE* count rate assuming $n = 1.68$. The corresponding slope increase is ~ 0.1 , a value consistent with the results obtained by Guainazzi et al. (1996).

By constraining the lag between the EUV and 4–10 keV bands to be less than 20 ks or less than 1 ks (depending on whether we take the conservative result of the ZDCF, or the more speculative result of the scaling model), we can constrain the physical size of the Comptonising region in the context of simple upscattering models. In this context, the lag between the EUV and 4–10 keV bands corresponds to the time taken to upscatter EUV photons to medium X-ray energies (regardless of the energy of the initial seed photons). We now consider a simple Comptonisation model, where the EUV photons originate at the centre of a homogeneous spherical cloud of thermal electrons, and are upscattered on their way through the cloud to produce the medium-energy X-ray photons. We use this simple model to estimate an upper limit to the size of the X-ray emitting region.

The time spent to upscatter a photon from an energy E_1 to a higher energy E_2 is:

$$t_{\text{up}} = \frac{N \lambda}{c} \quad (5.1)$$

where N is the number of scatterings required to raise the energy from E_1 to E_2 and λ is the mean free path of the photon between scatterings. The mean free path is given by:

$$\lambda = \frac{1}{n_e \sigma_T} \quad (5.2)$$

where n_e is the electron density and σ_T is the Thomson cross-section (assuming $E_1 \ll m_e c^2$). We can express n_e in terms of the optical depth of the Comptonising cloud, τ and the radius of the cloud, R :

$$n_e = \frac{\tau}{\sigma_T R}. \quad (5.3)$$

Incorporating equations 2 and 3 into equation 1, and rearranging yields:

$$R = \frac{c \tau t_{\text{up}}}{N} \quad (5.4)$$

so that substituting our maximum lag for t_{up} sets an upper limit on R . The number of collisions required to scatter the photons from energy E_1 to E_2 depends on the electron temperature, T_e (e.g. see Longair, 1992). In the non-relativistic regime, the number of scatterings is given by:

$$N = \frac{\log(E_2/E_1)}{\log(1 + (4kT_e/m_e c^2))}. \quad (5.5)$$

The electron temperature is not well constrained, but because the power-law continuum extends right across the useful PCA band (up to 15 keV) with no cut-off at the highest energies, we are justified in assuming a lower limit to the electron temperature of ~ 25 keV. OSSE observations of Seyfert galaxies show high energy cutoffs in their continua which correspond to typical electron temperatures of ~ 100 keV (~ 60 keV in the case of NGC 4151) (Zdziarski et al. 1997), consistent with our simple non-relativistic assumption.

The ratio of photon energies between the two bands is > 20 , so for an example where $kT_e = 25$ keV, $N > 16$. For $kT_e = 100$ keV, $N > 5$. For photon energies $E \ll kT_e$, the optical depth for Compton scattering in a spherical cloud can be determined from the spectral slope of the resulting Comptonised continuum (Posdnyakov, Sobol & Sunyaev 1983):

$$\tau = \left(\frac{\pi^2}{4 [(\alpha + 3/2)^2 - 9/4]} \frac{m_e c^2}{kT_e} \right)^{\frac{1}{2}} - 0.5$$

where α is the energy spectral index of the Comptonised continuum. Using the spectral index of 1.3 measured by *RXTE*, we find $\tau = 1$ for $kT_e = 100$ keV, and $\tau = 2.5$ for kT_e of 25 keV. Therefore, in the range of temperatures 25–100 keV, we expect the ratio $\tau/N < 0.2$. Using this limit in equation 4, we arrive at an upper limit for R assuming $t_{\text{up}} < 1000$ s:

$$R < 6 \times 10^{12} \text{ cm.}$$

A size of 6×10^{12} cm corresponds to ~ 20 Schwarzschild radii for a $10^6 M_{\odot}$ black hole. Thus, if the power-law continuum is produced by a single central source, the emitting region is very close to the central black hole. Alternatively, in disk corona models (e.g. see Haardt, Maraschi & Ghisellini 1997) the

power-law continuum is produced by the scattering of thermal photons from an accretion disk in a hot corona above the disk. In this case, the continuum may originate in many scattering regions spread over the disk, so the size constraint we impose may represent the size of a typical scattering region. Note that this size limit increases by a factor ~ 20 if we instead choose the more conservative 20 ks upper limit on the lag, inferred from the ZDCF.

We note here that Monte Carlo simulations of thermal Comptonisation in the relativistic regime (Skibo et al. 1995) show that if $\alpha > 1$, $\tau < 0.1$ for $k T_e > 250$ keV. In the case of very low optical depth to scattering, EUV and X-ray photons may undergo only one upscattering from the original seed photons before leaving the Comptonising cloud. In this case, upper limits to lags between the two bands yield no information regarding the size of the Comptonising region. This situation is extreme however, but in the case of NGC 4051 we cannot be completely confident in our assertion of a small Comptonising region until finite lags are measured between different energy bands, or a spectral cut-off is determined.

Finally, we note that although our constraint on the size of the emitting region in NGC 4051 is model dependent, simple causal arguments place a model-independent upper limit of 3×10^{13} cm on the separation between the EUV and X-ray emitting regions, if the lag between both bands is less than 1 ks. Combining this constraint with the rapid variability seen in both bands, we can infer that the continuum in both bands is emitted co-spatially within a region of that size (i.e., 1000 light-seconds).

5.4 The relationship between X-ray and optical variability

The relationship between variations in the optical and X-ray lightcurves of AGN can constrain models for the production of the optical continuum (or at least the part of it which is varying). Variability in the optical lightcurves of radio-quiet AGN has been studied for several decades, but no understanding of it has emerged. In radio-quiet AGN, the bulk of the optical continuum should come from the accretion disk. Variability in the optical continuum may then be caused by variations in the accretion flow of matter through the disk or perturbations in the accretion disk leading to local changes in disk temperature, and hence fluctuations in the black body flux. An alternative possibility however, is that the variations are caused by optical reprocessing of the X-ray emission, which is known to be strongly variable. The X-ray flux which is absorbed by the accretion disk heats it and so is effectively reprocessed to optical radiation, leading to fluctuations in the optical lightcurve which track the X-ray fluctuations. Simultaneous optical and X-ray monitoring of AGN is required to test this hypothesis.

Done et al. (1990a) used simultaneous optical and *Ginga* X-ray observations of NGC 4051 to search for correlations between the two bands. They found that although the X-ray flux varied greatly during the observation (by a factor of 3), the optical flux did not vary at all. Apart from the possibility that the X-ray reprocessing model of variability is wrong, one reason for the lack of optical response to large X-ray variations could be that the X-ray reprocessing region is large compared to the X-ray emitting region, so that rapid X-ray variations are smeared out and the reprocessor only responds to longer time-scale variations in the X-rays. Long-timescale X-ray and optical studies are needed to test this possibility.

In order to investigate the relationship between optical and X-ray variabil-




Figure 5.10: NGC 4051 X-ray, 5100Å continuum and $H\beta$ line monitoring lightcurves.

ity on long time-scales, we collaborated with the ‘AGN Watch’ team (Brad Peterson, Belinda Wilkes and collaborators) who monitored variations in the optical spectrum of NGC 4051 for much of the duration of our *RXTE* monitoring campaign. Besides comparing the long-term optical and X-ray lightcurves of NGC 4051, the other purpose of optical spectral monitoring was to measure lags in the response of optical emission lines to variations in the optical continuum, and so estimate a mass for the central black hole using reverberation mapping techniques (e.g. Peterson et al. 1998). NGC 4051 is classed as a narrow-line Seyfert 1 (NLS 1) and correspondingly, the mass measured by the optical monitoring program was low at $\simeq 1.1 \times 10^6 M_{\odot}$, implying an accretion rate of $\sim 20\%$ of the Eddington limit. Further details of the optical monitoring and reverberation mapping results are described in Peterson et al. (2000). Here I will present the key results found with regard to the relation between the X-ray and optical bands (also presented in Peterson et al. 2000).

The X-ray, optical continuum and $H\beta$ line lightcurves are shown in Figure 5.10.

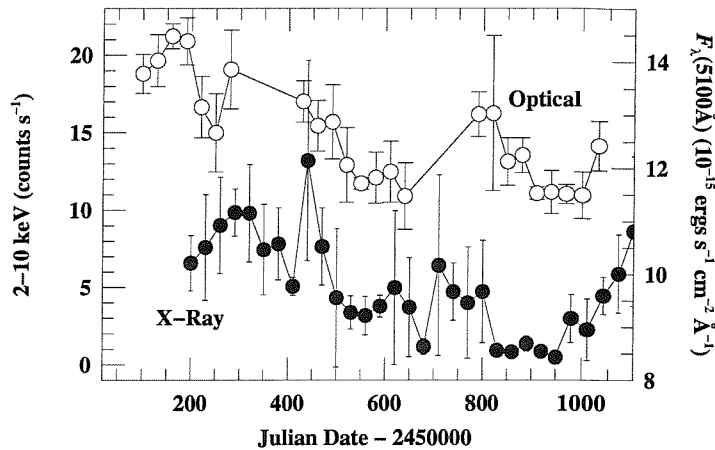


Figure 5.11: Comparison of smoothed long-term X-ray and optical continuum lightcurves (taken from Peterson et al. 1999).

The amplitude of X-ray variability is clearly much greater than any optical variability. In particular, when the X-rays virtually disappear during the low state of early 1998, the optical continuum emission remains, so clearly most of the optical emission is not driven by the X-rays. Short time-scale optical variations occur on time-scales of days but these are not correlated with short time-scale X-ray variations. However, careful examination of Figure 5.10 reveals that there does seem to be a slow downwards trend in optical continuum and line flux on time-scales of months, which corresponds to the decline in X-ray emission from a normal active state to the period of non-Gaussian activity with lower mean flux and ultimately to the low state. This correlation between the long-term optical and X-ray variability is made more clear by smoothing the X-ray and optical lightcurves with a 30-day sliding window, as shown in Figure 5.11 (Peterson et al. 2000).

We can infer from Figure 5.11 that the X-ray and optical lightcurves are cor-

related on long time-scales (i.e. months), but what does this mean for models of the optical variability? In order to smear out variability on time-scales of days any reprocessing region must be light-weeks across. The gravitational radius of a black hole of mass $10^6 M_{\odot}$, such as that inferred for NGC 4051 is only 10 light-seconds, so any reprocessor would need to be of the order of 10^5 gravitational radii in size, far larger than the accretion disk, which emits the bulk of its optical continuum within 100 gravitational radii of the black hole. Furthermore, in order to maintain a high enough brightness temperature to emit in the optical band, any reprocessor of size \sim light-weeks would need to be irradiated by extreme quasar-like luminosities. Clearly the variable optical emission is not caused by reprocessed X-rays.

An alternative explanation for the correlation between long-term X-ray and optical variations could be that both are correlated with the same underlying parameter which varies on long time-scales. An obvious choice for this parameter is the accretion rate. If the variable part of the optical continuum is associated with the inner accretion disk where disk variability time-scales are shortest, then variations in the accretion rate in the inner disk will cause variations in the optical continuum and X-ray continuum together (since the X-ray continuum is probably associated with a corona above the inner disk). There are thus two components to the optical variability of NGC 4051. A short-timescale component is associated with perturbations in the accretion disk and leads to the short-timescale optical variability, while a long-term component is associated with gradual changes in the accretion rate, leading to the long-term correlation between X-ray and optical emission.

During the low state, $H\beta$ line emission continues to vary on short time-scales. This poses a problem however, since $H\beta$ emission is driven by photons above 13.6 eV, and yet the EUV flux above 100 eV is very faint and not strongly variable during the low state. To account for this, the varying EUV flux must cut off somewhere between 13.6 eV and 100 eV. Examination of the RMS optical spectrum obtained during the low state shows that the variable part of the

broad HeII emission line (which is ionised by photons of energies above 54 eV) disappears during the low state (Peterson et al. 2000). This implies that the EUV component above 54 eV, which normally drives the HeII variations is no longer active during the low state, i.e. the variable EUV component cuts off between 13.6 eV and 54 eV. As discussed earlier, the low state itself may correspond to the disappearance of the inner disk, possibly a transition to an ADAF. This possibility is supported by the disappearance of variable EUV photons implied by the measurements of optical line variability, i.e. the inner disk is truncated within the region where disk temperatures peak in the EUV. Assuming that the inner disk is radiatively dominated, the disk temperature T of a $10^6 M_{\odot}$ black hole at a radius R (where R is in units of gravitational radii) is given approximately by $T \sim 1.5 \times 10^6 R^{-3/8}$ K (Treves, Maraschi & Abramowicz 1988), so that a cut-off temperature of between 13.6 eV and 54 eV (i.e. $1.6\text{--}6.3 \times 10^5$ K) corresponds to a disk truncation radius of between 10 and ~ 100 gravitational radii.

5.5 Summary

We have seen how the X-ray lightcurve of NGC 4051 shows unusual non-stationary behaviour, including a ‘normal’ Gaussian distributed part, an unusual non-Gaussian distributed part, associated with the relatively low-flux period prior to a quiescent, low state which lasted 5 months in early 1998 and recurred for a further three months in early 1999. The non-stationarity in the lightcurve makes it impossible to describe the broadband power spectrum with a valid power-spectral model, since the December long-look lightcurve which is used to calculate the high-frequency power spectrum seems to be contained within the non-Gaussian part of the lightcurve, and cannot be well fitted by any power-spectral model.

EUVE observations obtained during the first low state indicate that there is

a low level of variability in the low state on long time-scales, at least in the EUV band. The bulk of the EUV flux measured in the low state is probably non-variable and associated with extended emission detected by *ROSAT*. The EUV lightcurve obtained during the second low state caught the source transiting back to the normal active state, showing that this transition occurs on time-scales $< \text{day}$.

Simultaneous long-look observations of NGC 4051 in the first low state, with *BeppoSAX* and *RXTE*, show that the spectrum is particularly hard with a prominent iron $K\alpha$ line, and is well described by a model where the low-state X-ray emission originates almost entirely from reflection of the normal, active continuum off a distant (> 150 light-days) molecular torus.

Simultaneous *EUVE* and *RXTE* observations while the source was active in May and December 1996 show that the EUV emission is highly correlated with the X-rays and may be simply described as the low-energy extension of the X-ray power-law continuum. Assuming a simple spectral model, the lag may be constrained to be less than 1 ksec, which implies that the size of the X-ray/EUV emitting region is < 20 Schwarzschild radii for a $10^6 M_{\odot}$ black hole, if the X-rays are produced by Comptonisation of *EUVE* photons.

Finally, we have seen that although the optical and X-ray continua are not correlated on short time-scales (days, implying that simple X-ray reprocessing models for the optical variability are wrong), they do seem to be correlated on time-scales of months–years, which suggests that the long-term variations in both continua may be associated with long-term changes in accretion rate or the pattern of the accretion flow. This result is confirmed by the inference from optical line measurements that the continuum close to the Lyman edge remains during the low state, while it cuts off somewhere between 13.6 and 54 eV. Together with X-ray measurements, these results suggest that in the low state the inner disk disappears or enters a radiatively inefficient mode of accretion. An intriguing interpretation of the unusual variability of NGC 4051 will be presented in the next chapter.

Chapter 6

Conclusions

Overview

In this concluding chapter, I discuss some of the implications of the results presented in this thesis, commenting throughout on the potential for future work. First I raise the problems caused to shot-noise models by the flux-scaling of power-spectral amplitude uncovered in Chapter 2, and suggest a ‘fractal flares’ model to explain this result. Then I discuss the implications of the power-spectral measurements presented in Chapter 4, for the physical time-scales which the break-frequency represents and the potential of using X-ray variability measurements as a black hole mass estimator. I also consider implications for AGN QPOs. Then I discuss the evidence for different accretion states in AGN, consider the evidence for high state AGN and suggest that the unusual variability properties of NGC 4051 imply that it is the AGN analogue of the microquasar GRS 1915+105. Finally I comment on the implications of these results for AGN unification models.

6.1 The physical nature of the red-noise process

The intrinsically non-stationary behaviour we have seen in the power spectra of AGN and XRBs is particularly puzzling. The linear dependence of RMS variability on flux over a broad range of time-scales implies that this kind of non-stationarity is intrinsic to the red-noise variability process and is distinct from the changes in power-spectral shape reported on long-timescales in XRBs, which may be associated with changes in the accretion flow (e.g. Belloni & Hasinger 1990). The fact that the amplitude of the power spectrum scales with the local squared mean flux contradicts what we expect from simple shot-noise models for the variability. Shot-noise models form part of a class of ‘stationary increment processes’ which construct lightcurves from individual increments, e.g. shots, and are often used to simulate red-noise variability which is surprisingly common in Nature, being found in phenomena as diverse as fluctuations in electrical resistance in simple circuits, economic time series, and traffic density fluctuations on a motorway (Keshner 1982). The key feature of a stationary increment process is that its power spectrum is stationary (Solo 1992). In many applications of models based on stationary increment processes, stationarity of the power spectrum is a desirable feature, because the observed power spectra are indeed stationary. However, we have seen that the red-noise X-ray lightcurves of compact accreting systems are intrinsically non-stationary, even on the shortest time-scales. It would seem then that a stationary increment process such as shot-noise is not a good representation of the real physical process underlying the variability.

We should not be too hasty to rule out shot-noise processes however, because the intrinsic stationarity in the red-noise variability process is such that the AGN/XRB power spectra are stationary provided we divide by their local mean. This suggests that some kind of shot-noise model may be applicable

provided the size of the shots are somehow linearly correlated with the local mean flux. Conventional shot-noise models assume that lightcurves are built from shots with a broad distribution of decay time-scales, in order that power-spectral slopes other than $\alpha = 2$ can be produced above the knee frequency, so that the knee frequency corresponds to the decay time-scale of the longest shots (e.g. Lehto 1989; Lochner, Swank & Szymkowiak 1991). The fact that the power spectral amplitude scales in the same way with local mean flux over three orders of magnitude in frequency (as shown in Figure 2.5) implies that shots over an equivalent range of duration must be equally correlated with flux. Here we must confront a fundamental problem however, since it is the shots themselves which the lightcurve is built of and hence the flux itself is correlated with the shots. Therefore the shots must be correlated with one another *on all time-scales*, i.e. the smallest, shortest duration shots must ‘know’ what the longest duration shots are doing and vice versa.

Perhaps the simplest solution is to discard the idea that lightcurves are built ‘from the bottom up’ out of individual shots and instead suppose that the shot-like variability we see is the result of flares which cascade ‘from the top down’. The fundamental units of variability may be a few large flaring regions that are broken into a fractal structure of many levels of smaller flaring units. Under one realisation of this model, as a large flare erupts it breaks into smaller sub-flares which break into smaller sub-sub-flares, and so on, in a cascade. The number of flares at each level of the cascade varies stochastically about some mean which is independent of the luminosity of the parent flares. The smallest flares (i.e. farthest down the cascade) vary on the shortest time-scales, so that at high frequencies, variability is due to large numbers of small flares. Then, if the number of flares at a given level of a cascade obey Poisson statistics, the levels farthest down the cascade will contribute least to the variability of the lightcurve, so that the integrated power decreases towards higher frequencies, producing something like a red-noise power spectrum.

A numerical study of lightcurves resulting from ‘fractal flaring’ models of this

type is beyond the scope of this work. Presumably the fractal structure of the X-ray flares would be intrinsically related to self-similar structure of the varying regions themselves, such as a fractal structure of the magnetic field loops which might power the corona. Given the abundance of fractal structure elsewhere in Nature, we should not be surprised if it also underlies the self-similar variability of the X-ray lightcurves of compact accreting systems. Investigation of a possible physical physical basis for fractal structures in compact accreting systems will be left as a future work.

A further question resulting from the RMS vs. flux relationship discovered for Cygnus X-1 and SAX J1808.4-3658 is *what is the nature of the second, possibly constant component to the lightcurve?* Note that this component, revealed by the non-zero intercept in the linear RMS vs. flux trend, is not necessarily constant, but may contribute a relatively small constant level to the RMS variability of the lightcurve. For this reason, I shall refer to it as the ‘constant- σ ’ component of the lightcurve, as opposed to the ‘variable- σ ’ component which causes the linear RMS vs. flux trend. We can investigate the spectrum of the constant- σ component by measuring C , the value of the intercept on the x-axis, as a function of energy using lightcurves made in smaller energy bands than the 0.0-13.1 keV lightcurve used to uncover the RMS vs. flux relationship. Figure 6.1 shows the ratio of C in each energy band to the mean flux of the lightcurve in the energy band, for four energy bands (0–3.9 keV, 3.9–6.1 keV, 6.1–8.3 keV and 8.3–13.1 keV), for Cygnus X-1 and SAX J1808.4-3658. The error bars are estimated from the linear fits to the RMS vs. flux relation in each band, but given that the lightcurves in each energy band are correlated with one another, the uncertainties in the overall spectral shape are probably much smaller than the error bars imply.

Interestingly, the spectrum of the constant- σ component is very similar to the mean spectrum of the entire lightcurve. What this means is unclear, and is dependent on the assumption that the power spectral shape of the constant- σ component does not vary with energy in a manner different to the power




Figure 6.1: Relative flux of C compared to the mean flux in four energy bands.

spectrum of the total lightcurve. If we assume the simplest case, where the constant- σ component is constant and does not contribute to the variability, then the spectral shape measured here is a direct measure of the underlying constant spectrum. In this case, the similarity of the spectrum of the constant component to the mean spectrum of the entire lightcurve suggests that it is not a distant reprocessor, as in the case of the torus component in NGC 4051, but that it originates from the same source as the variable continuum. One possibility is that it represents a measurement of the non-varying part of the X-ray emitting corona. The prevalence of the shot-noise ‘building block’ paradigm and the notable variability of XRBs and AGN makes us assume that the continuum X-ray emission from accreting systems is by necessity made up of varying parts, however this need not be the case. Provided its temperature can be maintained, the corona will Comptonise seed photons from the accretion disk regardless of whether it is static or active. If the corona lies relatively close to the surface of the accretion disk, the constant emission component might represent the quieter parts of the corona, while the active regions of the corona are fairly localised in the same way that solar flares are localised to certain regions of the solar corona. On the other hand, if the constant- σ component is

not truly constant and contributes to σ , it may represent regions of the corona with a different magnetic field configuration to the regions which display fractal flaring of the kind discussed earlier.

The power spectra of Cygnus X-1 and SAX J1808.4-3658 which were used to uncover the RMS vs. flux trend have a relatively simple red-noise form. It will be particularly interesting to see how more complex XRB power spectra vary as a function of local mean flux. For example, do QPO amplitudes show the same dependence on flux as shown by the red noise, and is it possible to use the linear dependence of red-noise RMS variability on flux as a tool to disentangle the red-noise variability component in a lightcurve from other variable or constant components? The *RXTE* archive is a rich source of data for these kinds of studies.

6.2 Power-spectral scaling in AGN and XRBs

The main result of this work has been to show that the power spectra of AGN flatten towards low frequencies. Whether this flattening has the same form as that in XRBs is still not certain. The data do not yet allow us to categorically rule out simple knee models or more complex models where the power spectra might have multiple breaks or flatten gradually. However, it is encouraging that a simple high-frequency break, similar to that seen in the classic black hole XRB Cygnus X-1, can fit the data well, and yield plausible black hole masses if we assume that the black hole mass scales with the break frequency.

A useful way of comparing power spectra is to plot frequency \times power, rather than power, as a function of frequency. The $\nu P(\nu)$ plot produced in this way is analogous to the $EF(E)$ method of displaying energy spectra, in that the peak in the $\nu P(\nu)$ plot shows which time-scales most of the variability occurs on, as well as the magnitude of variability on those time-scales. In Figure 6.2 I have plotted the $\nu P(\nu)$ power spectra for the best-fitting high-frequency break




Figure 6.2: Comparison of $\nu P(\nu)$ power spectra of MCG-6-30-15, NGC 5506, NGC 3516, NGC 5548 and Cygnus X-1 in the low and high states.

models fitted to the observed power spectra of MCG-6-30-15, NGC 5506 and NGC 3516 (see Section 4.4.2). The $\nu P(\nu)$ power spectrum of NGC 5548 is also included in the figure, except that since the break is not very well defined for this object, I have assumed a break of 2×10^{-7} Hz, corresponding to the black hole mass of $10^8 M_{\odot}$ estimated from reverberation mapping. For comparison, $\nu P(\nu)$ power spectra of Cygnus X-1 are included, corresponding to typical power-spectral parameters in the low and high states (Nowak et al. 1999, Cui et al. 1997). Note that the flat peaks in the $\nu P(\nu)$ power spectra correspond to the $\alpha = 1$ part of the power spectrum, where there is equal integrated power per decade of frequency. The low-frequency drop-off in the $\nu P(\nu)$ power in the low state of Cygnus X-1 corresponds to the low-frequency break or ‘knee’ in the power spectrum, which we have not yet detected in the AGN power spectra (it could occur at frequencies $< 10^{-6}$ Hz, but the low-frequency data is not yet adequate to detect it).

It is apparent that the $\nu P(\nu)$ power spectra of AGN are similar to those in Cyg X-1, in that they have similar peak powers, but they are shifted at least 5 decades down the frequency axis. The fact that the peak powers are similar implies that the number of varying regions and the general pattern of variability is the same in AGN and XRBs, while the luminosity of the varying regions and their variability time-scales are scaled by some factor which is similar to the black hole mass. In the remainder of this chapter, we assume that this model is indeed the best model to describe the broadband power spectra of AGN, and we shall examine the specific implications of this interpretation of the data in the remainder of this section.

6.2.1 What physical time-scale does the break frequency correspond to?

The first question we might ask is: *what physical time-scale does the break frequency represent?* Physical time-scales we might consider are the light-crossing time-scale of the emitting region, t_{light} , the dynamical time-scale t_{dyn} (corresponding to the orbital time-scale of matter around the black hole), and the thermal, sound-crossing and viscous time-scales of the accretion disk t_{therm} , t_{sound} and t_{visc} respectively (e.g. Treves, Maraschi & Abramowicz 1988). The values of these various time-scales are given by:

$$t_{\text{light}} \simeq 10 (M_{\text{BH}}/10^6) R \quad \text{seconds,}$$

$$t_{\text{dyn}} \simeq 88 (M_{\text{BH}}/10^6) R^{3/2} \quad \text{seconds,}$$

$$t_{\text{therm}} \simeq t_{\text{dyn}}/\alpha \quad \text{seconds,}$$

$$t_{\text{sound}} \simeq \frac{R}{H(R)} t_{\text{dyn}} \quad \text{seconds,}$$

$$t_{\text{visc}} \simeq \left(\frac{R}{H(R)} \right)^2 (t_{\text{dyn}}/\alpha) \quad \text{seconds,}$$

where M_{BH} is the mass of the black hole in units of solar mass, R is the radius (measured from the black hole, in gravitational radii R_{G}) and t_{therm} , t_{sound} and t_{visc} assume a thin accretion disk with viscosity parameter α (Shakura & Sunyaev 1973) and scale height $H(R)$ at radius R .

The light-crossing time-scale is clearly much too small to correspond to the break frequencies measured, since it would require an X-ray emitting corona of size $> 10000 R_{\text{G}}$ which is physically unfeasible (since a huge amount of power would be required to heat and support such a large corona) and is ruled out by the short time-scale variability we measure. The break frequency might correspond to the dynamical time-scale if the X-ray emitting region extends out to $\sim 60 R_{\text{G}}$ (e.g. assuming a black hole mass of $10^7 M_{\odot}$ in NGC 3516, and a break timescale of $\sim 4 \times 10^5$ s). The thermal and viscous time-scales

depend on the viscosity parameter α which is related to the detailed physics of the disk and is unknown. However, α is probably significantly less than 1 (its maximum value) so that $t_{\text{therm}} \gg t_{\text{dyn}}$, in which case the break frequency might correspond to the thermal time-scale in the inner disk ($R \sim 10 R_G$).

The ratio of the viscous time-scale to the thermal time-scale is given by the square of the ratio of R to the disk scale-height. In a conventional thin disk, $R \gg H(R)$ so that the viscous time-scale is much longer than the thermal time-scale, which would seem to imply that the break time-scale is too short to be explained as a viscous time-scale. However, it is likely that the inner disk is dominated by radiation pressure, which acts to puff the disk up so that $R \sim H(R)$ in the innermost part of the disk. Thus the break time-scale might correspond to the viscous time-scale if the inner disk is sufficiently hot. However, since the disk temperature scales with $M_{\text{BH}}^{-1/4}$, the inner disks of more massive black holes are less dominated by radiation pressure, such that the ratio $R/H(R)$ scales with $\sim M_{\text{BH}}^{1/10}$. Therefore, if the break time-scale corresponds to the viscous time-scale it should not scale linearly with black hole mass, but scales as $\sim M_{\text{BH}}^{6/5}$. Although the extra exponent is only small, its effect over more than 5 orders of magnitude in black hole mass would mean that the black hole masses estimated in Section 4.5 are too high by a factor of 10. This would imply that NGC 3516, MCG-6-30-15 and NGC 5506 are accreting well into the super-Eddington regime and that their black hole masses are much lower than even the smallest black hole masses estimated from reverberation mapping. It therefore seems unlikely that the break time-scale corresponds to the viscous time-scale. A similar argument can be used against the possibility that the break time-scale corresponds to the sound-crossing time-scale, which is proportional to $R/H(R)$ and hence does not scale linearly with black hole mass. However, the argument against the sound-crossing time-scale is not so strong, since the deviation from linear scaling is not as great as for the viscous time-scale.

A more definitive estimate of the mass-scaling of the break time-scale can be

made by measuring the break frequencies of a large sample of AGN which have independently-estimated black hole masses (e.g. using reverberation mapping). There are now more than half a dozen reverberation-mapped AGN being monitored by *RXTE*, so these data will help to constrain the scaling somewhat better than the small sample presented here. Of particular interest is NGC 5506, which has been found to have a water megamaser (Braatz, Wilson & Henkel 1994). VLBI monitoring of the megamaser in this source should yield a very accurate mass for the central black hole, free from the assumptions and uncertainties inherent in reverberation mapping. The possibility that NGC 5506 is a high state AGN can be confirmed by such a measurement, which would provide a highly accurate measure of the mass-scaling of the frequency break, provided that an accurate break-frequency can be determined. Obtaining a better-quality broadband power spectrum for this target should be a priority.

6.2.2 How well can we measure black hole mass?

The results presented here support the hypothesis that the break frequency we measure for AGN scales linearly with M_{BH}^{-1} . The next question is, how accurately can we measure the black hole mass? Assuming that the slope below the break frequency can be fixed at or close to 1, then with the best data sets presented here, we can constrain the break frequency to within a factor 2, and this accuracy will also be achieved for other objects using forthcoming data sets. However, the main uncertainty is not the break frequency we measure, but rather what value of the break frequency in Cygnus X-1 should we compare it to? As discussed earlier, Belloni & Hasinger (1990) have shown that the position of the high-frequency break in Cygnus X-1 in the low state varies from 1–6 Hz, so it seems likely that AGN show similarly variable break frequencies (although presumably the time-scales for changes in break frequency in AGN are of the order of decades to centuries or longer). Since we do not

know whether the break frequency we measure in an AGN power spectrum is lower or higher than average, we cannot be sure which value of the Cygnus X-1 break frequency to scale it to, and this causes a further factor ~ 6 uncertainty in our estimate of black hole mass.

We can remove this intrinsic uncertainty if we can find other characteristics of the variability or spectra of XRBs (and Cygnus X-1 in particular) which are correlated with the changes in break frequency. Belloni & Hasinger (1990) found no such correlations for the high-frequency break, but did find that although the position of the low-frequency knee in Cygnus X-1 varies by a factor of 10, the knee-frequency is tightly inversely correlated with the power at the knee. Thus by measuring the power at the knee it is possible to calibrate out the variability in the knee and hence estimate an accurate scaling factor between the power spectrum of Cygnus X-1 and the AGN power spectrum (which may be used to estimate a black hole mass if the knee frequency scales in the same way as the high-frequency break). If knees exist in AGN power spectra, they must occur at relatively low frequencies so that well sampled long-duration lightcurves (e.g. 5 year duration), which are not yet available, are required to detect them. It is still possible that we may find other correlations between the high-frequency break and, for example spectral features such as iron line width or equivalent width (which we might expect if changes in the break frequency are related to physical changes in the accretion disk), so this remains an important future work.

An earlier attempt at estimating black hole mass in AGN from X-ray variability power spectra was made by Hayashida et al. (1998), who defined a ‘characteristic’ threshold frequency for AGN power spectra as being the point at which the $\nu P(\nu)$ power crossed a certain threshold (defined as $\nu P(\nu) = 2 \times 10^{-3}$ using the fractional RMS normalisation used here). By comparing the threshold frequency estimated from power spectra measured from *Ginga* long-look AGN lightcurves with that measured for Cygnus X-1 and scaling linearly with black hole mass, they estimated black hole masses which were generally much lower

than those measured by reverberation mapping, and implied super-Eddington luminosities. Besides the fact that Hayashida et al. did not take proper account of sampling effects in measuring AGN power spectra, it is easy to see why this method is unsuitable for estimating black hole masses. The threshold frequency defined in this way cannot be used as a characteristic frequency because it is strongly dependent on the measured high-frequency power-spectral slope α , in addition to being dependent on the break frequency.

Consider two AGN power spectra, with the same $\nu P(\nu)$ power at the break frequency, $\nu_{\text{bk}} P(\nu_{\text{bk}})$. The threshold $\nu P(\nu)$ power is given by:

$$\nu_{\text{th}} P(\nu_{\text{th}}) = \nu_{\text{bk}} P(\nu_{\text{bk}}) \left(\frac{\nu_{\text{th}}}{\nu_{\text{bk}}} \right)^{1-\alpha}.$$

Rearranging this to give ν_{th} we find:

$$\nu_{\text{th}} = \nu_{\text{bk}} \left(\frac{\nu_{\text{th}} P(\nu_{\text{th}})}{\nu_{\text{bk}} P(\nu_{\text{bk}})} \right)^{\frac{1}{1-\alpha}},$$

so that the ratio of threshold frequencies, $R_{\text{th}} = \nu_{\text{th},1}/\nu_{\text{th},2}$ measured for two AGN power spectra with the same ν_{bk} and $\nu_{\text{bk}} P(\nu_{\text{bk}})$ but different slopes, α_1 and α_2 is given by:

$$R_{\text{th}} = \left(\frac{\nu_{\text{th}} P(\nu_{\text{th}})}{\nu_{\text{bk}} P(\nu_{\text{bk}})} \right)^{\frac{\alpha_1 - \alpha_2}{(1-\alpha_1)(1-\alpha_2)}}.$$

So for a reasonable value of $\nu_{\text{bk}} P(\nu_{\text{bk}}) = 0.01$ and taking $\nu_{\text{th}} P(\nu_{\text{th}}) = 2 \times 10^{-3}$, if the range of values of α is 1.5–2.4 (as seen in Cygnus X-1), the variation in threshold frequency is increased by a further factor of 8, i.e. the total error in using this ‘characteristic time-scale’ can be up to factor ~ 50 .

It is not very surprising then that the method of Hayashida et al. is very poor at estimating reliable black hole masses. The fact that it tends to underestimate black hole masses (for example, Hayashida et al. estimate a black hole mass of $10^7 M_{\odot}$ for NGC 5548 versus $10^8 M_{\odot}$ measured by reverberation mapping) can be explained as a selection effect, due to the fact that more massive black holes will only have variable lightcurves on short time-scales if they have unusually flat power spectra, hence their threshold frequencies

will be higher than objects with more typical power spectra. However, the method of Hayashida et al. may prove useful if we can assume that an AGN is in the high-state and that the high-state power spectrum also scales with black hole mass, as the MCG-6-30-15 power spectrum presented here seems to indicate. The reasons for this are three-fold. First, the break frequency in the high-state power spectrum of Cygnus X-1 only varies by factor ~ 2 (from 6.5–14 Hz, Cui et al. 1997). Second, the power-spectral slope is always steep, varying between 1.9 and 2.6, so that the resulting variation in threshold frequency for AGN with the same break frequency is only a factor ~ 2 , for the threshold-to-break $\nu P(\nu)$ ratio assumed earlier. Thirdly, the data presented by Cui et al. (1997) indicate that the break frequency in the high state may be positively correlated with the power-spectral slope above the break frequency, so that variations in break-frequency or slope can be calibrated out. Taking these factors into account, it seems that we can use the method of Hayashida et al. to determine masses of AGN in the high state to within a factor ~ 2 , using only long-look data or relatively short duration monitoring lightcurves. The main uncertainty lies in whether or not an AGN is in a high state, which I shall discuss later.

6.2.3 Where are the AGN QPOs?

The scaling of characteristic time-scales of accreting black hole systems with black hole mass can provide clues as to the location of the seemingly elusive AGN QPOs. For the AGN considered in this work, of probable black hole masses $\sim 10^6 M_{\odot}$, $\sim 10^7 M_{\odot}$ and $\sim 10^8 M_{\odot}$, we expect to find the 1–20 Hz QPOs seen in black hole XRBs at frequencies of 10^{-4} – 10^{-7} Hz. Although this frequency range is sampled by the power spectra we have measured, the poorer quality of the power spectra at low frequencies, combined with the like-

likelihood that any QPOs only contribute a few per cent to the fractional RMS variability, conspire to make it improbable that we would see any evidence for them. Much better quality low-frequency power spectra are required to search for these QPOs. On the other hand, high-frequency (~ 100 Hz) QPOs such as those seen in the microquasars might be detectable at $\sim 10^{-3}$ Hz in the high-frequency power spectra of the lowest mass objects in our sample (e.g. NGC 4051, NGC 5506 and MCG-6-30-15), although we see no evidence for them. This is probably not surprising, since high frequency QPOs are apparently rare among black hole XRBs.

6.3 Accretion states and AGN unification

6.3.1 Evidence for high state AGN

One of the key results to emerge from the power-spectral study of the small sample of AGN presented in this work, is the discovery that at least one AGN in the sample, namely MCG-6-30-15, shows strong evidence for a high state power spectrum. NGC 5506 may also have a high state power spectrum, although the break frequency is not particularly well-defined so we cannot be certain of this. Further support for this hypothesis can be seen in the $\nu P(\nu)$ power spectra shown in Figure 6.2, which shows that the peak powers of MCG-6-30-15 and NGC 5506 are the lowest of the four AGN in the sample, similar to the lower peak power of the high state power spectrum of Cygnus X-1 compared to the low state. This result is robust, since the peak power is related to the RMS variability and is fairly resilient to errors in the power-spectral shape. The long-term fractional RMS variability of MCG-6-30-15 and NGC 5506 is 26.5% and 22.6% respectively (see Section 4.2), compared to 30% for NGC 5548 and 29.6% for NGC 3516, which are likely to be low state sources. For comparison,

the fractional RMS variability measured for Cygnus X-1 in the high state is $\sim 18\%$ (Cui et al. 1997) while the low state fractional RMS varies between 30%–40% (e.g. Belloni & Hasinger 1990). The greater range of RMS variability between low and high states in Cygnus X-1 compared to that seen in our sample may be due to the presence of a strong constant soft X-ray component to the high state emission of Cygnus X-1, which acts to reduce the RMS variability of the red-noise lightcurve. This component is likely to be weaker in high-state AGN if it is associated with a hot accretion disk, since disk temperatures in AGN must be lower than in XRBs.

It may seem surprising that we find evidence for high-state activity in two out of four AGN in our sample, when high states in black hole XRBs are relatively rare, but this is probably due to a selection effect. The AGN in the sample were partly selected on the basis that they were fairly bright and showed significant enough variability that there was a reasonable chance of measuring a frequency break in their power spectrum. We have seen that rapidly variable AGN probably have relatively low black hole masses, but because these systems are quite rare (as the large sample of reverberation-mapped AGN indicates), they tend to be distant and hence must be highly luminous to satisfy our brightness criterion. Therefore we preferentially select high accretion rate AGN, i.e. potential high state objects. A more sensitive monitoring instrument than *RXTE* is required to avoid this bias in the future.

6.3.2 Is NGC 4051 a macro-microquasar?

Given that AGN show evidence for at least two of the accretion states shown by black hole XRBs, it is of interest whether there are AGN which show evidence for other patterns of variability seen in black hole XRBs. NGC 4051 is particularly puzzling, since it shows evidence for several different variabil-

ity states within just three-years (see Section 5.1), including ‘normal’ Gaussian variability seen in other AGN, a period of non-Gaussian variability which lasted about a year and two quiescent low states, lasting several months each, where the X-ray source seemed to almost switch off. Earlier *ROSAT* observations of NGC 4051, in November 1991, and November 1992 showed evidence of non-Gaussian and Gaussian variations respectively (Green, M^cHardy & Done 1999).

It is interesting to note that the mean flux of the non-Gaussian lightcurve measured in November 1991 is a factor ~ 5 lower than that of the Gaussian lightcurve measured in November 1992, and similar to that measured for the non-Gaussian part of the long-term monitoring lightcurve, suggesting that the non-Gaussian behaviour is associated with periods of lower long-term mean flux. This result can be explained if the basic building blocks of the lightcurves are a few large flaring regions, which vary stochastically, as suggested by the fractal flaring model posited in Section 6.1. If the mean number of flaring regions is large, then stochastic variations about the mean will follow a Gaussian distribution, hence the resulting lightcurve will be Gaussian. However, if the number of flaring regions is decreased in proportion with the mean flux (perhaps associated with a change in accretion rate) then stochastic variations will push the lightcurve into non-Gaussianity, since the number of flaring regions cannot drop below zero but can still rise significantly above the mean.

The low state is a distinguishing feature of the long-term X-ray lightcurve of NGC 4051 and might correspond to the disappearance of the inner accretion disk. Peterson et al. (2000) suggest an analogy with black hole X-ray transient systems, in that the low state may correspond to the quiescent states seen in these sources. However, this analogy is misapplied, since black hole X-ray transients show transient behaviour on time-scales of months to years, which would correspond to $\sim 10^5$ years in NGC 4051 (assuming a black hole mass of $\sim 10^6 M_{\odot}$ inferred from reverberation mapping), not the \sim year time-scale which we see.

A more profitable approach is to assume that the time-scale for the state-transitions seen in NGC 4051 scales linearly with black hole mass, and consider whether any XRB systems show similar variability patterns on the much shorter time-scales implied by the mass scaling. Therefore, we must look for changes of state in a black hole XRB system on time-scales of only ~ 100 – 1000 s. To the author's knowledge, the only black hole XRB system which shows this behaviour is the microquasar GRS 1915+105. Amongst its many, varied behaviours, GRS 1915+105 has been observed to enter quiescent low states of duration \sim hundreds of seconds, following periods of large amplitude variability (Mirabel et al. 1998). The low states have characteristically hard X-ray energy spectra and repeat on similar or slightly longer time-scales than the duration of the low state. This pattern of repeated sequences of large-amplitude variability followed by quiescence is similar to what we observe in NGC 4051.

A typical active–low–active cycle for GRS 1915+105 is shown in Figure 6.3. Also shown are simultaneous radio and infra-red lightcurves, which show corresponding flaring events that are associated with the cycle. The radio flare lags the infra-red flare and this behaviour can be ascribed to the adiabatic expansion of plasma clouds ejected along the jet, which emit synchrotron radiation towards longer wavelengths as the plasma cloud expands and cools (also depicted in the figure). It is believed that the ejection of plasma clouds is intrinsically associated with the disappearance of the inner disk in GRS 1915+105 (Feroci et al. 1999). Optical and EUV observations suggest that the inner disk in NGC 4051 also disappears during the low state.

Could it be that NGC 4051 also shows infra-red and radio flares associated with the transition to the low state? A categorical answer to this question requires good quality ground-based monitoring observations in the infra-red and radio, both during and after low-state transitions, which we hope to obtain in the future. However, there is circumstantial evidence that NGC 4051 does flare in the infra-red and that this flaring might be associated with the transition




Figure 6.3: The low state transition in GRS 1915+105. Infra-red and radio lightcurves are plotted in normalised luminosity units on the same axis, while the zero flux base of the X-ray observations is represented by the flat line, corresponding to Earth occultation of the source. X-ray hardness ratio is plotted at the bottom of the figure and a depiction of the cloud ejection model is printed at the top (figure taken from Mirabel & Rodriguez 1999).

to the low state. Salvati et al. (1993) report infra-red monitoring observations of NGC 4051 which show that the continuum source flared during the first 6 months of 1992, increasing in flux by a factor of 2, before dimming to its original flux level. Salvati et al. were unable to explain this flare in the context of standard UV reprocessing models. Note however, that *ROSAT* observed NGC 4051 in a low-flux, non-Gaussian state in November 1991, and that in our *RXTE* monitoring lightcurve, the period of non-Gaussian variability precedes the first low state. Furthermore, sparse *ROSAT* observations of NGC 4051 in December 1991 (Komossa & Meerschweinchen 2000) show little or no evidence of variability and that the *ROSAT* flux is similar to that expected from the extended emission reported by Singh (1999). Therefore, I speculate that NGC 4051 entered a low state some time in November–December 1991 and that the infra-red flare reported by Salvati et al. later in 1992 may be associated with this transition, in the same way that infra-red flares follow the low-state transition in GRS 1915+105.

One might query the speculation above, on the grounds that microquasars show radio jet emission and as such are compared to radio-loud AGN, whereas Seyferts are not considered to be radio-loud. However, sensitive radio observations of Seyfert galaxies show that they have compact radio cores and radio jets (Kukula et al. 1999), although the radio emission is about a thousand times weaker compared to their total luminosities than that seen in radio-loud objects. If we compare the strength of radio emission with the total luminosity in GRS 1915+105, it turns out that GRS 1915+105 has a similar relative radio strength to Seyfert galaxies. Other black hole XRBs (Including Cygnus X-1) also show evidence for weak radio jets (Fender et al. 2000) and have similar relative radio strengths to Seyfert galaxies, although they do not show the plasma ejections which characterise microquasars. It seems that the term ‘microquasar’ is something of a misnomer and that in terms of their relative radio strengths at least, microquasars should in fact be called ‘micro-Seyferts’. However, in view of the attraction of the term ‘microquasar’, I shall not try to

supplant it. Hence I will suggest that NGC 4051 is a ‘macro-microquasar’.

6.3.3 Different states: towards an AGN Grand Unification Theory

We have seen strong evidence that AGN exist in at least two of the states seen in black hole XRBs, specifically the high and low states, and circumstantial evidence which suggests that even the unusual variability properties of microquasars may have an AGN analog. Therefore, I suggest that an essential missing piece in the AGN unification puzzle is the fact that AGN can exist in a number of accretion states analogous to those seen in XRBs. For example, what are the Narrow Line Seyfert 1s? NLS 1s are characterised by rapid, sometimes non-Gaussian X-ray variability and typically-soft X-ray spectra, as well as the narrow permitted lines which denote the class. Rapid variability and soft X-ray spectra could be associated with relatively low mass black holes accreting in the high state, which might account for at least some of the NLS 1s. However, we should note that NGC 4051 is also classed as an NLS 1 and that its long-term lightcurve shows evidence of prolonged periods of non-Gaussian activity, similar to what is seen in some other NLS 1s (Leighly 1999). Perhaps some of the NLS 1 population is associated with the same kind of accretion instability which we see evidence for in NGC 4051 and which might be analogous to that seen in microquasars. The X-ray spectrum of GRS 1915+105 during periods of activity between low states is particularly steep ($\Gamma \sim 2.5$) and has a strong thermal emission component (Feroci et al. 1999), so that a soft spectrum might also be characteristic of the unstable state. Some of the more luminous and extreme NLS 1s which show evidence of non-linear variability could be associated with just such a state.

Although the suggestions above are highly speculative, it remains true that if the basic patterns of accretion disk variability are the same irrespective of

central black hole mass, then many AGN are different because they occupy different states. The time-scales for changes in state are likely to be very large, of the order of a hundred thousand years or more, (except for unstable systems analogous to GRS 1915+105, which might show changes on time-scales of years to decades), but as we study large representative samples of AGN we are likely to find objects in all manner of states. It is to be hoped that future studies of AGN variability and comparison with XRBs, together with an improved theoretical understanding of accretion disks, might lead to an AGN ‘Grand Unification Theory’ which encompasses every animal in the AGN zoo.

Bibliography

- Abramowicz M. A., Calvani M., Nobili L., 1980, *ApJ*, 242, 772
- Abramowicz M. A., Czerny B., Lasota J. P., Szuszkiewicz E., 1988, *ApJ*, 332, 646
- Alexander T., 1997, in: *Astronomical Time Series*, eds. Maoz D. et al., Kluwer Academic Publishers, Netherlands, p. 163
- Antonucci R., 1993, *ARA&A*, 31, 473
- Belloni T., Hasinger G., 1990, *A&A*, 227, L33
- Belloni T., Hasinger G., 1990a, *A&A*, 230, 103
- Blandford R. D., Begelman M. C., 1999, *MNRAS*, 303, L1
- Boller T., Brandt W. N., Fink H., 1996, *A&A*, 305, 53
- Braatz J. A., Wilson A. S., Henkel C., 1994, *ApJ*, 437, L99
- Bromley B. C., Miller W. A., Pariev V. I., 1998, *Nature*, 391, 54
- Brunner H., Müller C., Friedrich P., Dörrer T., Staubert R., Riffert H., 1997, *A&A*, 326, 885
- Chiang J., Reynolds C. S., Blaes O. M., Nowak M. A., Murray N., Madejski G., Marshall H. L., Magdziarz P., 2000, *ApJ*, 528, 292
- Christopoulou P. E., Holloway A. J., Steffen W., Mundell C. G., Thean A. H. C., Goudis C. D., Meaburn J., Pedlar A., 1997, *MNRAS*, 284, 385
- Cui W., Heindl W. A., Rothschild R. E., Zhang S. N., Jahoda K., Focke W., 1997a, *ApJ*, 474, L57
- Cui W., Zhang S. N., Focke W., Swank J. H., 1997b, *ApJ*, 484, 383
- Deeming T. J., 1975, *Astrophys. & Space Sci.*, 36, 137

- Di Matteo T., 1998, *MNRAS*, 299, L15
- Di Matteo T., Quataert E., Allen S. W., Narayan R., Fabian A. C., 2000, *MNRAS*, 311, 507
- Done C., Ward M. J., Fabian A. C., Kunieda H., Tsuruta S., Lawrence A., Smith M. G., Wamsteker W., 1990, *MNRAS*, 243, 713
- Done C., Madejski G. M., Mushotzky R. F., Turner T. J., Koyama K., Kunieda H., 1992, *ApJ*, 400, 138
- Done C., Zycki P. T., 1999, *MNRAS*, 305, 457
- Dotani T, et al., 1997, *ApJ*, 485, L87
- Edelson R. A., Krolik J. H., 1988, *ApJ*, 333, 646
- Edelson R., Nandra K., 1999, *ApJ*, 514, 682
- Elvis M., et al., 1994, *ApJS*, 95, 1
- Fabian A. C., Rees M. J., 1995, *MNRAS*, 277, L55
- Fender R. P., 1999, to be published in: *Proc. ESO workshop 'Black Holes in binaries and galactic nuclei'*, eds. Kaper L., van den Heuvel E. P. J., Woudt P. A., Springer-Verlag (astro-ph/9911176)
- Fender R. P., Hendry M. A., 2000, preprint, astro-ph/0001502
- Feroci M., Matt G., Pooley G., Costa E., Tavani M., Belloni T., 1999, *A&A*, 351, 985
- Green A.R., 1993, *Ph. D. Thesis - Southampton University*
- Green A. R., McHardy I. M., Lehto H. J., 1993, *MNRAS*, 265, 664
- Green A. R., McHardy I. M., Done C., 1999, *MNRAS*, 305, 309
- Guainazzi M., Mihara T., Otani C., Matsuoka M., 1996, *PASJ*, 48, 781
- Guainazzi M., et al., 1998, *MNRAS*, 301, L1
- Haardt F., Maraschi L., Ghisellini G., 1997, *ApJ*, 476, 620
- Halford D., 1968, *Proc. IEEE*, 56, 251
- Hayashida K., Miyamoto S., Kitamoto S., Negoro H., Inoue H., 1998, *ApJ*, 500, 642
- Iwasawa K., Fabian A. C., Brandt W. N., Kunieda H., Misaki K., Terashima Y., Reynolds C. S., 1998, *MNRAS*, 295, L20

-
- Kaspi S., Smith P. S., Netzer H., Maoz D., Jannuzi B. T., Giveon U., 1999, preprint, astro-ph/9911476
- Keenan D. M., 1985, *Biometrika*, 72, 39
- Keshner M. S., 1982, *Proc. IEEE*, 70, 212
- Komossa S., Meerschweinchen J., 2000, *A&A*, 354, 411
- Kukula M. J., Ghosh T., Pedlar A., Schilizzi R. T., 1999, *ApJ*, 518, 117
- Lamer G., Uttley P., McHardy I. M., 2000, submitted to *MNRAS*
- Lampton M., Margon B., Bowyer S., 1976, *ApJ*, 208, 177
- Lawrence A., Watson M. G., Pounds K. A., Elvis M., 1987, *Nature*, 325, 694
- Lawrence A., Papadakis I., 1993, *ApJ*, 414, L85
- Leahy D. A., Darbro W., Elsner R. F., Weisskopf M. C., Kahn S., Sutherland P. G., Grindlay J. E., 1983, *ApJ*, 266, 160
- Lee J. C., Fabian A. C., Brandt W. N., Reynolds C. S., Iwasawa K., 1999, *MNRAS*, 310, 973
- Lehto H., 1989, in: *Two topics in X-ray Astronomy*, 499, eds. Hunt J., & Battick B., ESLAB symposium, ESLAB SP-296, ESTEC
- Leighly K. M., Mushotzky R. F., Yaqoob T., Kunieda H., Edelson R., 1996, *ApJ*, 469, 147
- Leighly K. M., 1999, *ApJS*, 125, 297
- Lin D., Smith I. A., Böttcher M., Liang E. P., 2000, *ApJ*, 531, 963
- Lira P., Lawrence A., O'Brien P., Johnson R. A., Terlevich R., Bannister N., 1999, *MNRAS*, 305, 109
- Lochner J. C., Swank J. H., Szymkowiak A. E., 1991, *ApJ*, 376, 295
- Longair M. S., 1992, *High Energy Astrophysics Vol. 1*, Cambridge University Press, Cambridge
- Magorrian J., et al., 1998, *AJ*, 115, 2285
- Main D. S., Smith D. M., Heindl W. A., Swank J., Leventhal M., Mirabel I. F., Rodriguez L. F., 1999, *ApJ*, 525, 901
- Mahadevan R., 1998, *Nature*, 394, 651
- McHardy I., Czerny B., 1987, *Nature*, 325, 696

- McHardy I., 1988, *Mem. Soc. Astron. Ital.*, 59, 239
- McHardy I. M., Green A. R., Done C., Puchnarewicz E. M., Mason K. O., Branduardi-Raymont G., Jones M. H., 1995, *MNRAS*, 273, 549
- McHardy I. M., Papadakis I. E., Uttley P., 1998, in *Proc. Symp. The active X-ray sky: first results from BeppoSAX and RXTE*, eds Scarsi L., Bradt H., Giommi P., Fiore F., *Nucl. Phys. B (Proc. Suppl.)*, 69/1-3, 509
- Meier D. L., 1999, *ApJ*, 522, 753
- Mirabel I. F., Rodriguez L. F., 1994, *Nature*, 371, 46
- Mirabel I. F., Dhawan V., Chaty S., Rodriguez L. F., Marti J., Robinson C. R., Swank J., Geballe T., 1998, *A&A*, 330, L9
- Mirabel I. F., Rodriguez L. F., 1999, *ARA&A*, 37, 409
- Morgan E. H., Remillard R. A., Greiner J., 1997, *ApJ*, 482, 993
- Nandra K., et al., 1993, *MNRAS*, 260, 504
- Nandra K., Pounds K. A., 1994, *MNRAS*, 268, 405
- Nandra K., George I. M., Mushotzky R. F., Turner T. J., Yaqoob T., 1997a, *ApJ*, 476, 70
- Nandra K., George I. M., Mushotzky R. F., Turner T. J., Yaqoob T., 1997b, *ApJ*, 477, 602
- Nandra K., George I. M., Mushotzky R. F., Turner T. J., Yaqoob T., 1997c, *ApJ*, 488, L91
- Narayan R., Yi I., 1995, *ApJ*, 452, 710
- Nowak M. A., Chiang J., 2000, *ApJ*, 531, L13
- Nowak M. A., Vaughan B. A., Wilms J., Dove J. B., Begelman M. C., 1999, *ApJ*, 510, 874
- Padovani P., Rafanelli P., 1988, *A&A*, 205, 53
- Papadakis I. E., Lawrence A., 1993, *MNRAS*, 261, 612
- Papadakis I. E., Lawrence A., 1993a, *Nature*, 361, 233
- Papadakis I. E., Lawrence A., 1995, *MNRAS*, 272, 161
- Papadakis I. E., McHardy I. M., 1995, *MNRAS*, 273, 923
- Peterson B. M., 1998, *Adv. Space Res.*, 21, 57

-
- Peterson B. M., et al., 2000, to appear in *ApJ*
- Posdnyakov L. A., Sobol I. M., Sunyaev R. A., 1983, *Astrophys. Space Phys. Rev.*, 2, 263
- Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992, *Numerical Recipes*, 2nd edn. Cambridge Univ. Press, Cambridge
- Reeves J. N., O'Brien P. T., Vaughan S., Law-Green D., Ward M., Simpson C., Pounds K. A., Edelson R., 2000, *MNRAS*, 312, L17
- Remillard R. A., Morgan E. H., McClintock J. E., Bailyn C. D., Orosz J. A., 1999, *ApJ*, 522, 397
- Reynolds C. S., 1999, preprint, astro-ph/9912001
- Salvati M., et al., 1993, *A&A*, 274, 174
- Shakura N. I., Sunyaev R. A., 1973, *A&A*, 24, 337
- Singh K. P., 1999, *MNRAS*, 309, 991
- Skibo J. G., Dermer C. D., Ramaty R., McKinley J. M., 1995, *ApJ*, 446, 86
- Solo V., 1992, *SIAM J. Appl. Math.*, 52, 270
- Tadhunter C., Tsvetanov Z., 1989, *Nature*, 341, 422
- Tanaka Y., et al., 1995, *Nature*, 375, 659
- Timmer J., König M., 1995, *A&A*, 300, 707
- Treves A., Maraschi L., Abramowicz M., 1988, *PASP*, 100, 427
- Turner T. J., George I. M., Nandra K., Mushotzky R. F., 1998, *ApJ*, 493, 91
- Turner T. J., George I. M., Nandra K., Turcan D., 1999, *ApJ*, 524, 667
- Urry C. M., Padovani P., 1995, *PASP*, 107, 803
- Uttley P., McHardy I. M., Papadakis I. E., Cagnoni I., Fruscione A., 1998, in *Proc. Symp. The active X-ray sky: first results from BeppoSAX and RXTE*, eds Scarsi L., Bradt H., Giommi P., Fiore F., *Nucl. Phys. B (Proc. Suppl.)*, 69/1-3, 490
- Uttley P., McHardy I. M., Papadakis I. E., Guainazzi M., Fruscione A., 1999, *MNRAS*, 307, L6
- Uttley P., McHardy I. M., Papadakis I. E., Cagnoni I., Fruscione A., 2000, *MNRAS*, 312, 880

- van der Klis M., 1989, in: *Timing Neutron Stars*, 27, eds. Ögelman & Van den Heuvel, NATO ASI C262, Kluwer
- van der Klis M., 1997, in: *Statistical Challenges in Modern Astronomy II*, astro-ph/9704273
- Vaughan S., Reeves J., Warwick R., Edelson R., 1999, *MNRAS*, 309, 113
- Wandel A., 1999, *ApJ*, 519, L39
- Wandel A., Peterson B. M., Malkan M. A., 1999, *ApJ*, 526, 579
- Wanders I., Horne K., 1994, *A&A*, 289, 76
- Wang J. X., Zhou Y. Y., Wang T. G., 1999, *ApJ*, 523, L129
- Wijnands R., van der Klis M., 1998, *Nature*, 394, 344
- Wijnands R., van der Klis M., 1999, *ApJ*, 514, 939
- Yaqoob T., 1998, *ApJ*, 500, 893
- Zdziarski A. A., Johnson W. N., Poutanen J., Magdziarz P., Gierlinski M., 1997, *Proc. of the 2nd INTEGRAL Workshop*, ESA SP-382, 373