UNIVERSITY OF SOUTHAMPTON
FACULTY OF PHYSICAL AND APPLIED SCIENCES
PHYSICS AND ASTRONOMY

Observations of X-ray binaries using the Australia Telescope Compact Array - Compact Array Broadband Backend

by

Daniel Emilio Calvelo Santos

Thesis for the degree of Doctor of Philosophy
July 9, 2012

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL AND APPLIED SCIENCES

PHYSICS AND ASTRONOMY

Doctor of Philosophy

OBSERVATIONS OF X-RAY BINARIES USING THE AUSTRALIAN TELESCOPE COMPACT ARRAY - COMPACT ARRAY BROADBAND BACKEND

by Daniel Emilio Calvelo Santos

The exploration of radio emission from the relativistic jets of X-ray binaries has become key to understanding accretion onto compact objects. Upgrades to existing radio telescope arrays have improved the ability of these instruments to detect faint sources and their unique behaviours. In this thesis I address the use of one such instrument in observing several X-ray binary systems: the Australia Telescope Compact Array - Compact Array Broadband Backend (ATCA-CABB). Firstly, a study of the adverse effects that may arise when imaging spectrally variable sources using broad bandwidths, with emphasis on ATCA-CABB. Images are produced from model data and examined to find any false structures. I then discuss complications that such features may cause when attempting to interpret real structure in images.

Secondly, the results of attempts to observe the black hole X-ray binaries, GRO J1655-40 and XTE J1550-564, in quiescence with ATCA-CABB are presented. Upper limits from the non-detection of these sources are used in exploring the lower luminosity region of the radio/X-ray correlation for hard state black hole X-ray binaries and the implications of “universal” scatter is discussed.

Thirdly, a 16 day campaign observing the radio emission of the neutron star X-ray binary Circinus X-1 over a complete orbit during an historically faint state is presented. Variability in the source is examined and corrected images are compared with archival maps in order to find any differences. I discuss my findings in terms of secular changes in jet geometry and behaviour.

Finally, the first successful millimetre detections of neutron star X-ray binaries Circinus X-1 and Scorpius X-1 are presented. Sub-arcsecond jet structure in Circinus X-1 is compared to the existing centimetre maps, providing additional support for theories proposed earlier in the thesis. The radio to mm spectrum of both sources is estimated and discussed.

This thesis concludes with a description of the broader impacts of the above projects, as well as possibilities for future research in each field.

CONTENTS

APPENDICES

LIST OF TABLES

LIST OF FIGURES

DECLARATION

I, Daniel Emilio Calvelo Santos, declare that the thesis entitled *Observations of X-ray binaries using the Australia Telescope Compact Array - Compact Array Broadband Backend* and the work presented in the thesis are both my own, and have been generated by me as a result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as:
  - Limits on the quiescent radio emission from the black hole binaries GRO J1655 – 40 and XTE J1550 – 564, 2010, MNRAS, 409, pp. 839-845
  - Radio observations of Circinus X-1 over a complete orbit at an historically faint epoch, 2012, MNRAS, 419, pp. 436-451
  - Millimetre observations of a sub-arcsecond jet from Circinus X-1, 2012, MNRAS, 419, pp. L54-L58
  - Millimetre emission from neutron star X-ray binaries, in prep.

Daniel Emilio Calvelo Santos, July 9, 2012

ACKNOWLEDGEMENTS

First of all I'd like to thank my supervisor Rob Fender for taking me on as a student, helping me through the challenges involved in an astrophysics PhD and enduring all my blank stares, occasional 'flowery' writing and one or two last minute changes of heart over the years. Secondly I'd like to thank my family for supporting me during my studies (even fiscally every now and then): Mum, for always keeping your kitchen well stocked when I visit and supplementing mine when I leave; Dad, for being very proud of what I do even though understanding it often comes with a headache; Abuela, por siempre creyendo en mi y aceptando mis decisiones, aunque podrían dejarte sin un nieto para mucho tiempo; Nick, for not rubbing in the fact that you will likely always earn more money than your big brother.

I would also like to thank my friends that have kept me sane and from becoming a social recluse. Those from my home town who have known me longest: Sean Barker and Alistair Nowell. Thanks for always challenging the existence of the things I study; thus, keeping me well versed in my counter-arguments, and generally not hurting me (much) for being a 'tax dodging' student all these years. My flatmate, James Hewitt, for surviving living with me for the last 3 years (and knowing me the 5 before that) and enduring my student lifestyle with minimal envy and pranks. My good friend, David George Cutler, and his amazing family (especially Becky Harnett and little Jessica Dax) for reminding me of some of the important things in life. To the 'Man's house' and poker lads: Rich Hextall, James Warwick, Nick Lawther, Josh Barret, Stephen Abbott and Aaron Chung: I'm pretty sure I've lost more than I've won but it's always been a good night! To my office mates who have come and gone over time: Martin Bell (whether good or bad, Australia, Thailand and Amsterdam will be trips I never forget), Liz Bartlett (for being the life of the party and mutual office rant support), Dan Plant (for sharing in the LOFAR 'experience'), Lee Townsend (for making the write up that little bit less frightening), Matthew Schurch and Mark Peacock. To my colleagues from the department who have helped make my research feel less like work: Grace Thomson, Tana Joseph, Sadie Jones, Seb Drave, Duncan Cameron, Pablo Cassatella, Gabriele Ponti, Teo Muñoz-Darias and Mickaël Coriat. Finally, to the Physics department undergrads (yes, even Ridwan) that have kept me from feeling too old by helping me enjoy some aspects of university life that little bit longer.

I've had the chance to collaborate with some excellent people over the years. Special thanks to Jess Broderick, for double checking my data reduction every now and then, helping out with observing, and for his amazing proof reading skills. To Tasso Tzioumis, for guiding me through my first steps at controlling ATCA (6 colossal dishes moving at my command - awesome) and being there so many times when I've needed help planning or covering observations. Also thanks to James Miller-Jones (particularly for delaying his own work so we could get papers out simultaneously), David Russell, Elena Gallo, Stephane Corbel, Fraser Lewis, and MAXI team members Nobuyuki Kawai and Motoki Nakajima. Finally I'd like to thank my advisers, Tom Maccarone and Malcolm Coe, for helping out with the

odd emergency, tricky astrophysical or data reduction concepts, chatting and for not being too scary during my annual reviews.

Daniel Emilio Calvelo Santos, July 9, 2012

DON'T PANIC

THE HITCHHIKER'S GUIDE TO THE GALAXY: DOUGLAS ADAMS (1952 – 2001)

Para Abuelo

The most beautiful thing we can experience is the mysterious. It is the source of all true art and all science. He to whom this emotion is a stranger, who can no longer pause to wonder and stand rapt in awe, is as good as dead: his eyes are closed.

ALBERT EINSTEIN (1879 - 1955)

"There are two kinds of people in the universe, Molly," I said. "Star Trek fans and Star Wars fans. This is shocking."
She sniffed. "This is the post-nerd-closet world, Harry. It's okay to like both."
"Blasphemy and lies," I said.

HARRY DRESDEN, GHOST STORY

So Einstein was wrong when he said, ”God does not play dice.” Consideration of black holes suggests, not only that God does play dice, but that he sometimes confuses us by throwing them where they can’t be seen.

STEPHEN HAWKING (1942 – PRESENT)

1 Introduction

1.1 A brief history of radio astronomy

On first inspection, as the lowest energy end of the electromagnetic spectrum, radio waves might inspire little excitement (I distinctly remember my undergrad self thinking along the lines of “X-ray astronomy sounds better...”); however, they have not only radically affected human society over the years but astrophysics (especially the high energy domain) as well. Given the importance of, and thus rapid spread of the use of radio waves in wireless communications during the early 20th century, it was only a matter of time before someone noticed that the Earth, be it a thunderstorm or the BBC, wasn’t the only source of radio noise in the universe. The unexpected discovery was made by Karl Guthe Jansky who, as a radio engineer, was interested in finding sources of static that could interfere with communications, and after eliminating several terrestrial signals eventually determined the remaining background hiss was coming from our own Milky Way. Though the astronomical community was slow to pick up on this revelation, the cascade of discoveries that followed once we really started looking at the radio sky have time and again redefined our understanding of the universe. These have included quasars (though their nature was confirmed with optical spectroscopy: Schmidt 1963, Oke 1963), the

detection of the first neutron stars (pulsars: Hewish & Okoye 1965, Hewish *et al.* 1968), and the Cosmic microwave background (CMB; many pigeons died to bring us this information; Penzias & Wilson 1965), the latter two of which led to a Nobel Prize.

The instruments and methods that have been employed to detect the wide variety of radio signals from these sources have evolved significantly since Karl Jansky's discovery (See Figure 1.1), often presenting great design challenges given the large wavelengths involved. Using simple principles (angular resolution $\theta \sim \lambda$/D where λ is the wavelength of observed electromagnetic radiation and D is a telescope's diameter), one can quickly grasp the difficulty involved in attempting to observe the radio sky. For example, the 14.6m wavelengths detected by Karl Jansky would require a telescope on the scale of $\sim$100km in order to create images with the same resolution as human sight ($\sim$ 1 arc-minute). Obviously, endeavours of such massive scales are beyond us, but it has not prevented us from observing as best we can, and devising innovative solutions to the resolution problem. A single parabolic dish can only become so large before weight becomes a severe issue, resulting in deformation over long periods, or even total structural failure under extreme conditions (e.g. the collapse of the Green Bank Telescope in 1988), although the Arecibo telescope avoids this by setting the reflecting surface in a sink hole (see Figure 1.1, panel 2). Even overcoming this, single aperture systems will always suffer from significant diffraction effects.

The solution to this is the use of multiple smaller antennae spread over a large area whose detected signals can be combined via interferometry to form a virtual telescope many times larger than any single dish. In such a scenario the group of dishes cover a smaller surface area (determined by the number and size of the individual dishes) compared to a filled dish and thus have a lower overall sensitivity; however, this is outweighed by the significant boost to resolution determined by the longest distance ('baseline' length) between any two antennae. This method, known as 'aperture synthesis' is described in Ryle & Hewish (1960; discussed in the following section) and has led to the development of a large number of radio telescope arrays whose antennae vary in size and arrangement, yielding different capabilities. The earliest interferometers used only two dipole antennae arrays (e.g. Ryle and Vonberg's interferometer built in 1946) with fixed positions and which relied on the rotation of the earth to allow for scanning of the sky. But as hardware improved (e.g. motorised dishes) and computers became available to manipulate and analyse received signals, interferometers with larger numbers of steerable parabolic dishes were constructed that could not only track sources across the sky, but by exploit-

Figure 1.1: Radio telescopes of the past, present and future. From the top down, first panel: Jansky's 'merry-go-round', a turntable mounted antenna with a diameter of approximately 30m that was originally located at Bell Telephone Laboratories in New Jersey, USA. Second panel: The Arecibo Observatory, the largest single filled focusing dish on Earth with a diameter of 305m situated near Arecibo city, Puerto Rico. Third panel: The Australia Telescope Compact Array (ATCA) which is made up of six 25m dishes (the sixth is fixed several kilometres from the five shown) which can have a maximum baseline of 6km, located in Narrabri, Australia. Final panel: An artist's depiction of the Square Kilometre Array (SKA) which could be made of thousands of small parabolic dishes spread over thousands of kilometres (the name itself refers to the array's proposed collecting area) to be split between Australia, New Zealand and South Africa.

ing the earth's rotation could fill out a larger synthesised aperture as the relative positions of baselines changed over 12 hours. Some of the better known examples include the Very Large Array (VLA; 27 $\times$ 25m dishes in a symmetric 3 arm pattern) in New Mexico, USA (Thompson *et al.* 1980), the Giant Metre-wave Radio Telescope (GMRT; 30 $\times$ 42m dishes, some in 3 arms with the remainder arranged randomly over 1km^2) near Pune, India (Swarup *et al.* 1991), and the Australian Telescope Compact Array (ATCA; 6 $\times$ 22m dishes on an east-west track with an optional short northern spur, picture in panel 3 of Figure 1.1) in Narrabri, Australia (Frater, Brooks & Whiteoak 1992). In the case of the VLA and ATCA, a large portion of the dishes actually lie on rails that allow them to be repositioned in order to alter the capabilities of the array. Each of these telescopes remains in use today, with both the ATCA and VLA having been recently upgraded to maintain their usefulness for years to come (see section 1.3 for details specific to ATCA). It is also possible to utilise a number of facilities spread across the globe in unison to produce images with resolutions down to milli-arcseconds, effectively taking the application of interferometry a step further (Very long Baseline Interferometry: VLBI), but requiring very precise 3-dimensional antennae positions and timing information to succeed.

Future radio telescope arrays will attempt to incorporate the best features of many that have come before while continuing to push the limits of sensitivity that can be achieved from the ground. The Low Frequency Array (LOFAR; de Vos, Gunst & Nijboer 2009) for example, is a 'phased array' currently completing commissioning which consists of many thousands (on the order of $\sim$10000) of cheap dipole antennae spread over several stations across Europe (station baselines of $\sim$ 100m, inter-station baselines from 100m to 1000s of km). Dipoles, unlike dishes, have the advantage of being omni-directional (i.e. sensitive to the entire sky) at the cost of significantly reduced sensitivity; however, their simplicity makes them far cheaper and thus can be included in very large numbers to make up for this reduction. With the aid of computers the signals from these dipoles can be modified in such a way that the array's 'beam' can be altered to target different areas of the sky without motorised detectors, and the software even allows for multiple beams which can be tiled to expand the observation area, or stacked to improve sensitivity over a smaller region. The planned Square Kilometre Array (SKA, named for its potential collecting area and featured in the bottom panel of Figure 1.1; Taylor 2008) to be constructed as seperate components in South Africa, Australia and New Zealand, will consist of tens of thousands of dipoles and thousands of small dishes effectively spanning an entire host continent (baselines $\sim$ 3000km). This will be the

most sensitive radio telescope to date, combining the advantages of dish and phased dipole arrays, as well as providing high resolutions from a single unified instrument.

Over the following sections I will briefly cover the basic mathematical principles in aperture synthesis that will aid in understanding the majority of the science covered in this thesis. I will go on to further describe ATCA and its subsequent upgrade via the Compact Array Broadband Backend (CABB; Wilson *et al.* 2011) as it was used to gather all the new radio astronomical data used in the included scientific work, as well as briefly explaining how these upgrades have improved the telescope's capabilities, but at a cost of generating new problems concerning observations of variable radio sources (discussed further in Chapter 3). This final topic is directly related to much of the material in this thesis, in that X-ray binaries are inherently variable sources and their observation has required that I tackle the problem as best I can in order to get the most out of my data, and derive useful scientific results.

1.2 Radio interferometry and aperture synthesis

Unlike the majority of telescopes, radio interferometers do not detect the sky brightness distribution directly. Instead, antennae from each baseline pair detect a radio signal which is transferred to a central receiver to be combined, after correction for geometrical path delays, into an interference pattern (hence 'interferometry'). As a result of the nature of interference patterns each baseline is more sensitive to a particular range of source sizes defined by the length of the baseline itself. Longer baselines produce patterns with narrower fringe spacings making them attuned to smaller structures in the sky (effectively a higher resolution) and less so to larger structure, with the inverse true for shorter baselines. Thus, in order to detect multiple source scales at the same time, a variety of baseline lengths is required, with the highest resolution at any particular wavelength determined by the longest baseline. The properties of the recorded interference patterns (i.e. amplitude and phase) are used to define a component of the 2-dimensional Fourier transform of the sky. The position of each baseline's component depends upon physical orientation of the baseline's two antennae, relative to the position of the target source in the night sky. This is normally described using a u,v,w co-ordinate system where w points towards the target, and u and v define the projected eastern and northern directions of the baseline position respectively, with distances typically expressed in wavelengths in order to represent everything in terms of spatial frequency (uv distance becomes

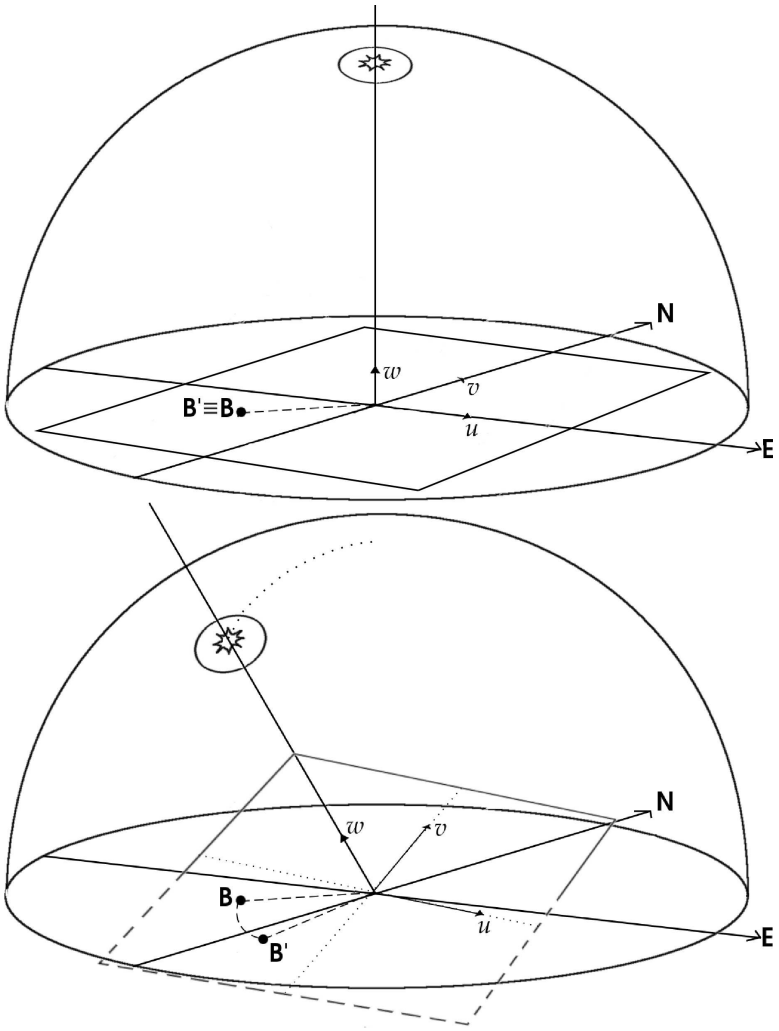

Figure 1.2: The *uv* plane. This diagram shows how baseline and source positions relate to the *uv* plane, with an example baseline represented by B, and its projection on the *uv* plane labelled B'. The top version shows the relation for a source directly overhead of the telescope (i.e. so that the *uv* plane and physical telescope plane are the same) whereas the bottom version shows how the orientation changes for an offset source (dashed edges represent regions below the physical plane).

directly proportional to angular resolution). This defines what is known as the '*uv* plane' on which the measured components of the sky's Fourier transform lie at each moment in time. The related geometry is depicted in Figure 1.2.

The amount of *uv* plane covered at any moment depends on the number of baselines, their positions, and target source declination. Additional factors also weigh in, such as instrument bandwidth which broadens the spatial frequencies encompassed by each individual baseline and thus its *uv* plane coverage (assuming no frequency averaging). Furthermore, if we then allow observations to continue over time then the projected baselines actually change as the target moves across the sky.

Interferometers exploit this to improve their *uv* coverage by allowing each baseline to trace out a track in the *uv* plane, allowing the interferometer to slowly fill in *uv* space and eventually sampling large arcs or rings of Fourier components instead of discrete points. Telescopes like the VLA with large numbers of well spaced (2-dimensions) antennae have good *uv* coverage over short periods of time (often referred to as 'snapshots'). Therefore, though improvement still occurs, long observations are less important for VLA like arrays than for east-west arrays like ATCA, where the linear arrangement (especially lower numbers of antennae) provides poor snapshot *uv* coverage. Nonetheless, with fewer antennae, east-west arrays are very cost-efficient and full synthesis (12 hour) observations can produce results comparable to larger arrays. This can be seen in Figure 1.3, where the antenna placement, snapshot and full 12 hour observations *uv* coverage (though units remain in km) are shown for the VLA and ATCA. Though the VLA will always out perform ATCA when working at the same bandwidths, you can see that a 12 hour observation vastly improves coverage for both compared to snapshot mode, but the difference is more important for ATCA.

The exploitation of these effects during long duration observations does come with drawbacks, in that one must assume that all visible sources detected during an observing run are unvarying. This is often untrue, especially in the case of X-ray binaries (ser the following Chapter for details), and can result in severe image defects when violated.

The mathematical representation of the process arises naturally from wave theory and Fourier transforms. For each detected position in the *uv* plane (i.e. each baseline) during each integrated unit of time, one measures the phase (θ) and amplitude (A) of the correlated signal, which together make up complex visibilities of the form $V(u,v) = Ae^{i\theta}$. These visibilities are the measured components of the sky's 2D Fourier transform evaluated at each *uv* co-ordinate and so are related to the true sky brightness distribution, $B(x,y)$, via equation 1.1.

$$V(u,v) = \iint_{-\infty}^{+\infty} B(x,y)e^{2\pi i(ux+vy)}\mathrm{d}x\,\mathrm{d}y \tag{1.1}$$

Since radio telescopes do not sample the entire *uv* plane, the inverse Fourier transform of the above, shown in equation 1.2 below, only provides us with an incomplete or sampled sky brightness distribution $B'(x,y)$. However, the greater the amount of *uv* space explored, the closer $B'(x,y)$ becomes to the true sky brightness distribution $B(x,y)$.

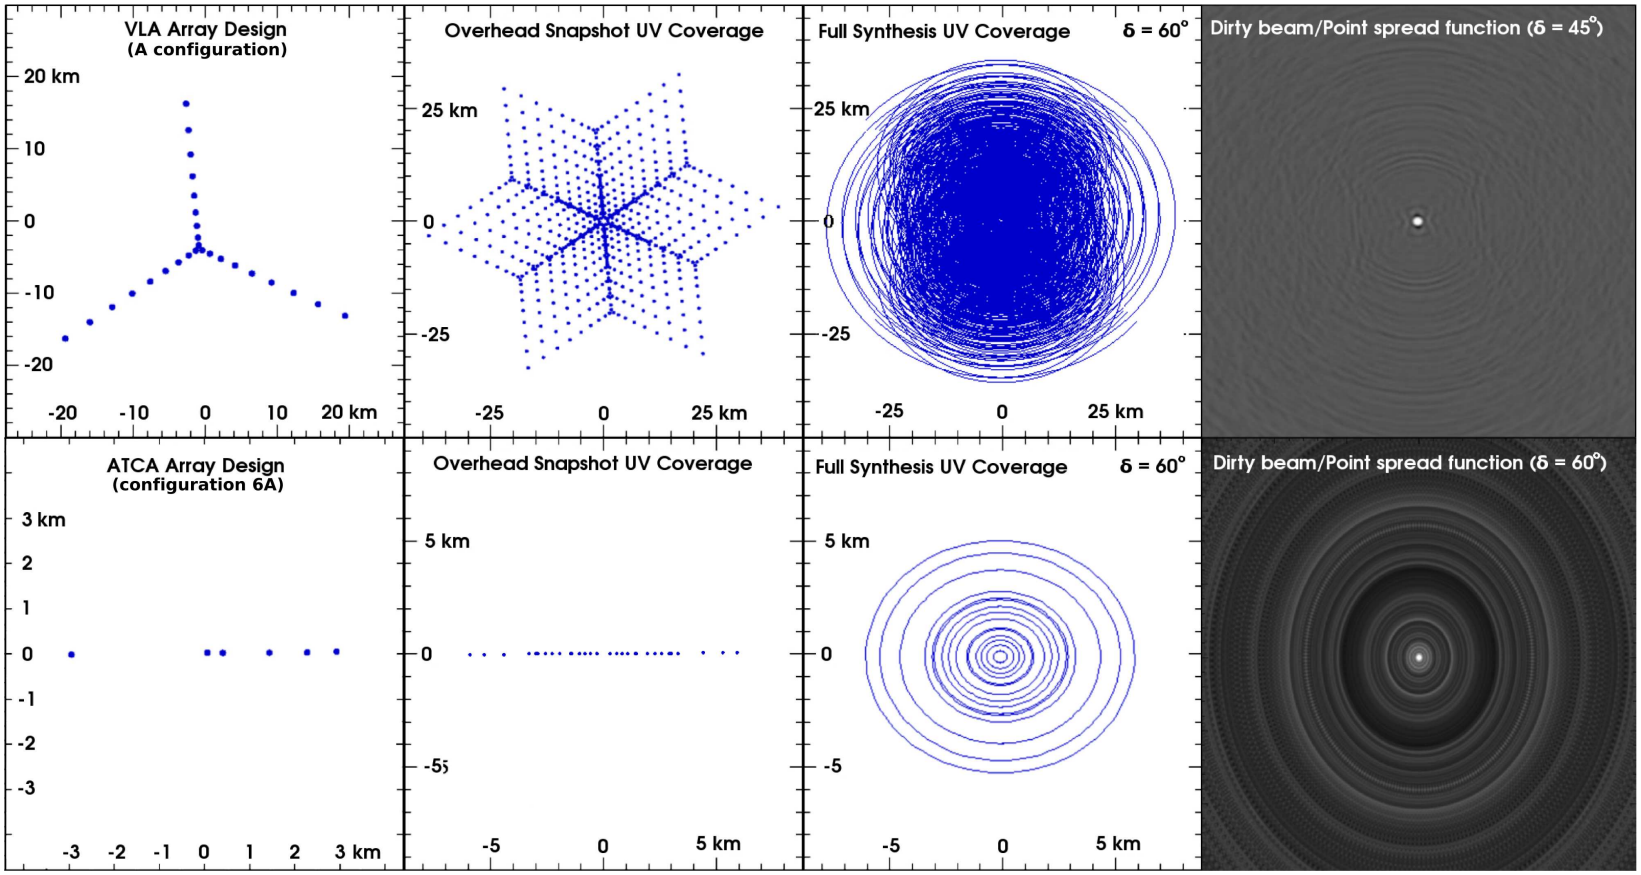

Figure 1.3: VLA and ATCA antenna configurations, example *uv* coverage and dirty beams. Top row from left to right: VLA antenna arrangement (configuration A), 'snapshot' *uv* coverage, full run (12 hour) *uv* coverage for a source at $\delta = 60^\circ$, full run dirty beam for a source at $\delta = 45^\circ$ ($\delta = 60^\circ$ unavailable but would be similar). Bottom row shows the corresponding images for ATCA antenna configuration 6A; however, note that the full run dirty beam is for a source at $\delta = 60^\circ$. Top row altered from material at `http://192.96.5.2/synthesis_school/Miod_Array_Design.pdf`, bottom row used the virtual interferometer application at `http://www.narrabri.atnf.csiro.au/astronomy/vri.html` for the first three panels, the final dirty beam was created using simulations in MIRIAD.

$$B'(x,y) = \iint_{-\infty}^{+\infty} V(u,v) e^{2\pi i(ux+vy)} \mathrm{d}u\,\mathrm{d}v \tag{1.2}$$

Coupled with the effects of varying baseline length, it is possible to manipulate the data before this Fourier inversion is carried out in an effort to enhance particular structure in the final image. While complete removal or flagging of *uv* regions is possible and effective; for example, eliminating short baselines (low *uv* values) to reduce sensitivity to large scale nebulae and structures that drown out smaller sources, this may be considered drastic as it is effectively removing useful data. An alternative is the use of weighting schemes that alter the relative contributions of different *uv* regions, and the applications of which fall between two extremes in this thesis: natural and uniform weighting. Natural weighting applies the same weighting to every *uv* measurement, whereas uniform weighting alters the weight of *uv* points based on the sampling density around them. Since actual data reduction applies this in terms of gridded area (gridding is also important for the 2D Fourier inversion, discussed later), it is useful to describe how these weighting schemes work in a grid format. The interpretation of these two schemes can become confused when discussing grids instead of individual visibilities as their roles are effectively reversed: natural weighting causes grid sectors containing larger numbers of visibilities to contribute more heavily to the imaging process and thus no longer 'levels' the contribution of the entire *uv* plane, whereas uniform weighting actually makes all grid sectors contribute equally (sector weight becomes a function of visibility density). Since most radio arrays tend to be biased towards shorter baselines and thus the detection of larger structures, uniform weighting can be of benefit in improving image resolution and revealing point-like sources by reducing the contribution from low *uv* visibilities just like the removal of shorter baseline visibilities, whereas natural weighting would be best when wanting to observe smoother large scale structure or to maximise point source sensitivity (lower noise levels). Very often, a balance must be struck between the two schemes depending on the nature of a target field and desired results.

1.2.1 Calibration and deconvolution

The initial stages of radio data reduction, prior to any imaging, requires flagging and several steps of calibration. Recorded data can be prone to radio frequency interference (RFI) as a result of terrestrial radio sources (variable levels depending on source direction and distance from the horizon) or errors from within the

instrument hardware itself. It is therefore necessary to comb the raw visibilities for such errors, via both frequency channels and time bin intervals, and flag them out to minimise adverse effects to calibration and imaging. Once flagged, multiple calibration tasks are applied in an attempt to correct data for a wide variety of effects, both observational and instrumental, including absolute flux density scale calibration, bandpass calibration, system temperature corrections, pointing accuracy, polarisation (both source and leakage corrections), primary beam correction, and phase calibration (correcting for transmission effects due to the atmosphere). Depending on the information that is required from a data set, some of these tasks are not always necessary; nonetheless, most dish based arrays that exist today (i.e. including ATCA) require observations of at least one 'calibrator source' to use as a reference for the calibrations. For example, a typical ATCA observation requires two or three (depending on the frequency used) calibrator sources: a flux calibrator that is both bright and has a stable, known absolute flux is observed at the start and end of a run; a bright, unresolved bandpass calibrator (sometimes the same source as the flux calibrator) observed at the start and end of a run (these sources are often useful for initial array phase and delay calibration as well) and finally a simple, unresolved, unvarying phase calibrator with known polarisation (easiest if it's unpolarised) that should be close in the sky to the target source. The reason for the required phase calibrator proximity is twofold; firstly, the phase calibrator must be observed regularly throughout a run to maintain the accuracy of the phase corrections and a nearby source reduces wasted slew time, and secondly, the phase calibration is based on the position of the calibrator itself and thus degrades with distance from it (i.e. most effective for sources withing the same isoplanatic patch). Each stage of calibration updates the complete correction list which can be transferred to the next calibrator source, before the final corrections can be applied to the target source data.

While this level (1^{st} generation) of calibration is often sufficient for many observations, there also exist more complex (2^{nd} generation) processes under the algorithms of SELFCAL. This uses the target source's own data to generate time-dependent bandpass, gain and/or phase corrections and can build upon the results of normal calibration routines. However, this method requires a good dynamic range (i.e. not useful for faint targets) to be effective as well as a relatively good starting model for complex source structure. Self calibration must be used with care if dealing with sources that vary significantly during an observation.

The sampled sky function (equation 1.2) can also be represented in another useful form via the convolution theorem, as shown in equation 1.3.

$$B'(x,y) = P(x,y) * B(x,y) \tag{1.3}$$

Where $P(x,y)$ defines the point source response of the array (i.e. *uv* coverage/function) which is more commonly known as a 'dirty beam'. Examples of dirty beams for VLA and ATCA observations are shown in the far right panels of Figure 1.3. As it is directly related to the available *uv* coverage of an observation, the structure of the dirty beam is also highly dependent on array configuration, source declination and observation timing. The full 12 hour VLA observation (top right panel of Figure 1.3) yields a near ideal point source with only minor ripple like distortions to the north and south which are related to the sparsely populated *uv* regions to the left and right of the full synthesis *uv* coverage plot in the panel to the left. These distortions are known as 'sidelobe' artefacts since they resemble the lobes that appear in the radiation pattern of an antenna. In the case of an ATCA full synthesis observation, we have a *uv* coverage that is more sparsely populated by distinct baseline rings, whose pattern is translated into the sidelobe rings of the dirty beam (bottom right panel of Figure 1.3). While it is evident that given a dirty beam image, a person could intuitively identify point sources in sparsely populated images by eye, difficulties begin to arise when sources become clustered and sidelobes overlap with each other and other sources.

Equation 1.3 gives us a more qualitative understanding of the final result of the 2D Fourier inversion, in that it shows us that the sampled sky distribution $B'(x,y)$ is merely the true distribution $B(x,y)$ convolved with the dirty beam $P(x,y)$, and helps us to understand the process of image cleaning known as 'deconvolution'. In this process all unsampled *uv* space is set to zero before the inverse 2D Fourier transform is carried out to produce a distorted or 'dirty' map. The inversion process itself applies a gridding and fast Fourier transform algorithm (FFT; a more efficient 2D inversion process than the original discrete Fourier transform) which divides the *u*v plane visibilities up into a grid before inversion. We can then take the resulting dirty map and, assuming that every source can be represented by a point model, identify the spread of points that when convolved with a model dirty beam account for as much structure in the map as possible. Effectively, we replace a subtracted dirty beam with a corresponding point source (defined by the array's beamsize and a flux) at each identified point.

Deconvolution is often an iterative process and the results are never perfect. A particularly busy region of the radio sky, in conjunction with a complicated dirty beam structure, can make the deconvolution process more difficult. However, there

are many different 'CLEAN' algorithms available that are designed to tackle different types of structure and data qualities that can be used alone or together to achieve better results in each situation. Two such algorithms were used in the works contained in this thesis, which I will proceed to describe in brief. The most common CLEAN methods employ variations of the Högbom algorithm (Högbom 1974, Schwartz 1978): A process which first identifies the maximum of an image and generates a component for this location with an intensity of a certain fraction of this maximum value (known as the 'gain' parameter and typically 10% by default), which is then convolved with the dirty beam and the result removed from the map. This process is then repeated, slowly removing all visible emission in a map. Several improvements to this exist including the Högbom and Clark algorithm (Clark 1980) which crops the dirty beam and includes minor and major iteration cycles that reduce the computational requirements to speed up the process.

A more advanced incarnation of the CLEAN process comes in the form of the multi-frequency synthesis or MFCLEAN algorithm (Conway, Cornwell & Wilkinson 1990, Sault & Wieringa 1994). This version attempts to account for frequency dependent variability such as sources with different spectral profiles, which become important for systems with large fractional bandwidths and observations of X-ray binaries (Chapter 2 discusses the causes of variability, whereas Chapter 3 elaborates on the effects on imaging). This is achieved by including additional terms in the dirty beam model which account for different spectral variability patterns, the most simple of which, and that which is applied in the MFCLEAN algorithm, assumes linear variability in a source spectrum:

$$B(\nu) = B(\nu_0) + \frac{\partial B}{\partial \nu}(\nu - \nu_0) \tag{1.4}$$

In terms of new beam components, the instrumental response, B_D of a point source with the above behaviour can be described using:

$$B_D(x,y) = B(\nu_0)P_0(x,y) + \nu_0 \frac{\partial B}{\partial \nu} P_1(x,y) \tag{1.5}$$

Where $P_0(x,y)$ is the standard dirty beam that would be used if one was not concerned with spectral effects (i.e. the response of an frequency invariant point source) and $P_1(x,y)$ is the 'spectral dirty beam' which defines the response of a point source whose spectrum varies with frequency by $(\nu - \nu_0)/\nu_0$. These two beam components are weighted by corresponding factors $B(\nu_0)$ and $\nu_0 \partial B/\partial \nu$, the latter of which can also be defined as $\alpha_e B(\nu_0)$ where α_e is the equivalent spectral index ($\alpha_e = \nu \partial B / B \partial \nu$). Effectively, this means that by calculating the multi-frequency

synthesised beam and therefore these weights, one estimates the flux density of a point source along with its spectral index. It is also possible to expand this method to incorporate more complicated spectral models; however, for the purposes of the work in this thesis such an exercise was never necessary.

1.3 The Australia Telescope Compact Array (ATCA)

As was briefly stated in section 1.1, the Australia Telescope Compact Array (ATCA) is a group of six 22m diameter dishes designed to observe the radio sky (an annotated satellite image of the site is shown in Figure 1.4). Prior to ATCA's construction, Australia was already well known for radio astronomy associated with single dish instruments such as the Parkes radio telescope (Bowen 1966). However, the southern hemisphere still lacked a radio dish array, and so there was an excellent opportunity for Australia to maintain a strong presence in global radio astronomy, while opening the southern radio sky to a new powerful tool. At the time of opening (1988) ATCA provided high resolving power, along with the ability to study polarisation and source spectra.

The six antennae, known as CA01 through CA06, produce 15 baselines and prior to the CABB upgrade (described in the next section) were initially capable of observing four frequency bands between 1.1 and 10.5 GHz (specifically, 1.5 GHz [20cm L band], 2.3 GHz [13cm S band], 5.0 GHz [6cm C band] and 8.6 GHz [3cm X band]) with primary beam FWHM (effectively the size of useful sky coverage) ranging between 44 and 5.1 arcminutes. Two later millimetre-wave upgrades added a further 3 bands between 15 GHz and 105 GHz (20 GHz [15mm K band], 40 GHz [7mm Q band] and 95 GHz [3mm W band]) with primary beams as small as 30 arcseconds, though the highest frequency band did not include antenna CA06. The maximum bandwidth observable at any one time was 128 MHz to be placed anywhere within the available bands; however, the system was capable of dual polarisation imaging which meant two of these 128 MHz windows were available and allowed some frequency bands to be observed simultaneously (e.g. 3 and 6 cm).

The arrangement of the antennae themselves is also not permanent, with antenna CA06 being the only fixed dish and CA01 through CA05 residing on a track that allows them to be repositioned into several useful array configurations. The track (marked in Figure 1.4) is comprised of a 3km east-west stretch which existed upon opening of the facility, as well as a northern 'spur' which was added in 1998 in order to significantly improve *uv* coverage for short observations by making 'hybrid'

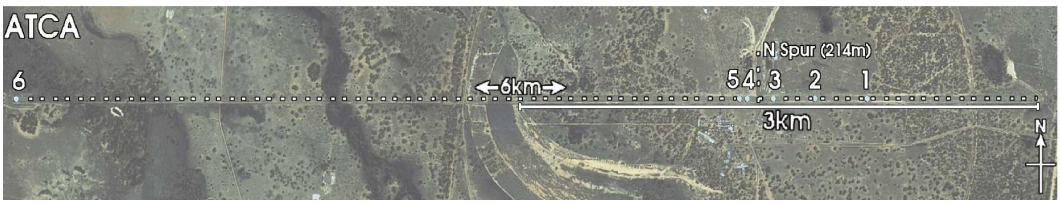

Figure 1.4: Satellite image of ATCA layout. The largest possible 6km spaced antenna position is marked by the horizontal dotted line, with the 214m northern spur marked in the vertical direction. The antennae are numbered, with antenna 6 furthest west in a fixed position. The solid horizontal line shows the extent of the 3km rail track which allows antennae 1 through 5 to change position (including into the northern spur).

telescope configurations available.

1.3.1 The Compact Array Broadband Backend (CABB)

To maintain ATCA's level of contribution to radio astrophysics it was necessary to carry out more significant hardware upgrades, beyond simply adding additional observable narrow frequency bands. While antennae design over the years has not undergone or even required radical changes, electronics, computing and communications hardware has advanced dramatically since ATCA's construction and presented an affordable avenue through which major improvements could be made to the instrument. Though improvements were also made to individual dish hardware and data transfer methods (increasing transfer rates and reducing data loss/corruption) the most significant change was a replacement of the 'correlator': the hardware that collects data signals from all antennae, frequencies and polarisations, combines, and processes them into the final output stored for analysis. It is the correlator that defines how and what information is recorded; polarisations, frequency channels, bandwidths, sampling rates and tolerance against RFI.

The new correlator allowed for an increase in ATCA's maximum bandwidth from 128 MHz to 2048 MHz, maintaining dual polarisation capability and thus giving observers two 2048 MHz windows to place at similar central frequencies to those that were available before. This is a factor of 16 improvement in bandwidth, corresponding to a factor of 4 increase in sensitivity (r.m.s. noise $\propto (B\tau)^{-0.5}$ where B is bandwidth and τ is integration time), reducing the time required to reach a desired detection threshold in images by a factor of 16. By maintaining distinct division of the band into up to 2048 frequency channels, one also drastically increases *uv* coverage (i.e. when using multi-frequency synthesis) of the array. This is illustrated in Figure 1.5 where I compare the multi-frequency synthesis *uv* coverage of

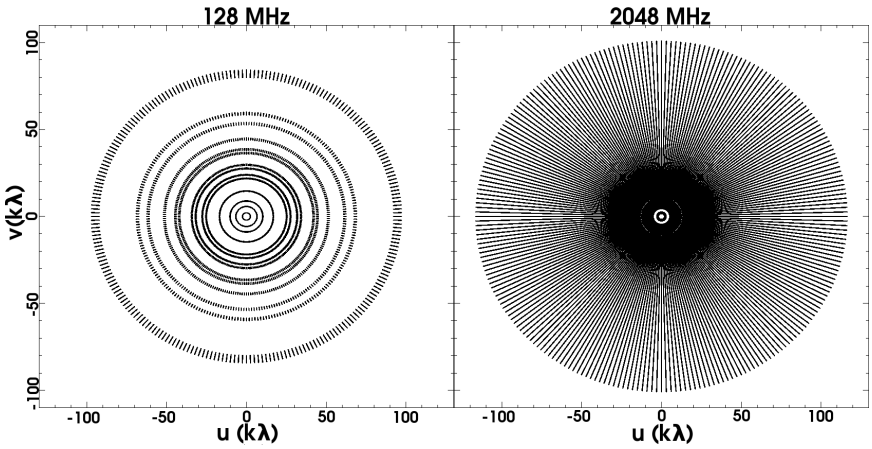

Figure 1.5: ATCA and ATCA CABB model data *uv* coverage comparison. The left hand panel shows the *uv* coverage of a model 128 MHz (the equivalent of what ATCA was capable of prior to the CABB upgrade) data set for a source at declination $\delta = 60^\circ$ with ATCA at 4.8 GHz. The right hand panel uses the same source and central frequency, but with a bandwidth of 2048 MHz (i.e. post CABB upgrade capability) to show significant change in *uv* coverage. Data models and plots were created with MIRIAD.

modelled pre-CABB ATCA data (Full synthesis/12 hours) to an equivalent following the CABB upgrade. The increase in bandwidth also improves ATCA's use in radio spectroscopy, in that far more of a spectrum is observable at any one time, reducing the need for frequency switching. This ability is even further enhanced by the use of zoom channel windows which can increase the resolution in several primary band channels by dividing them into a further 2048 channels (down to 488 Hz bins for 2048 $\times$ 1 MHz calibrator configuration) and can be placed anywhere within each original 2048 MHz band: ideal for spectral line studies requiring high velocity resolution. The new correlator also allows this detailed spectral study using all polarisation parameters. Finally, sampling resolution was also improved which provides high time resolution modes (spectra output can be increased to one every second from the default of one every ten seconds) for use in pulsar and transient studies.

ATCA's new capabilities also significantly increase the quantity of data from each observation. A standard observation (no zoom windows or increased write speeds) outputs 1.4 GB per hour, and increases by an equal amount for every zoom window added (since a zoom window represents the same number of data parameters, i.e. 2048 channels). Therefore, a full set of 16 zoom windows and the fastest write speeds would, hypothetically, increase data output rates by a factor of ~ 170. Though it might seem advantageous to gather as much data as possible during each observation, it is important to realise that the difficulty and time required to process,

calibrate and produce images from CABB data also increases with file size.

In respect to the work detailed in this thesis, this upgrade made it possible to target a variety of sources in faint states that fell below ATCA's original sensitivity levels including black holes in quiescence (Chapter 4) and neutron star X-ray binaries during historically faint epochs (Circinus X-1, chapter 5). Additionally, it also made it possible to observe neutron star XRBs at millimetre wavelengths for the first time (Chapter 6). However, this has not come without cost, in that an expansion of observable bandwidth will also increase the influence of source spectra profiles in imaging. Not only does the application of multi-frequency synthesis and MFCLEAN become essential, but spectral changes over time may have as great an effect on images as overall intensity changes, making observation of variable sources even more difficult an undertaking (discussed in Chapter 3).

1.3.2 MIRIAD

MIRIAD (Sault, Teuben & Wright 1995) is a radio interferometry data reduction suite created specifically for use with ATCA data. It provides the necessary tools to reduce ATCA and ATCA-CABB data from scratch, including loading of various data formats, viewing and manipulation (flagging) of raw visibilities, calibration, imaging, cleaning, analysis and display. This includes all the aforementioned methods and varying techniques such as multi-frequency synthesis, weighting schemes and deconvolution algorithms (CLEAN and MFCLEAN). It is capable of data simulation for a wide variety of scenarios encompassing array configuration, bandwidth, frequency channel configuration and observation intervals as well as allowing for addition and subtraction of model sources of various types (e.g. point, Gaussian). It is the dominant piece of software used in the analysis of all radio data in this thesis (Chapters 4 through 6), as well as the simulations discussed in Chapter 3. While the software has continued to be updated over time, the procedures used to discover the results discussed herein remain unchanged and the final outcome, if carried out today, should be the same.

Every time I learn something new, it pushes some old stuff
out of my brain.

HOMER SIMPSON, THE SIMPSONS

2

X-ray binaries and their electromagnetic emission

In this chapter I will briefly discuss the principal electromagnetic emission mechanisms that are believed to occur in X-ray binary (XRB) systems and their surroundings, as it will be required to understand the scientific work that will follow. I will continue to summarise several concepts related to our current understanding of XRB structure and behaviour, as well as a review of more recent discoveries that are directly associated with the scientific results in subsequent chapters.

2.1 Emission processes relevant to XRBs

Before discussing specific emission mechanics, it is important to mention the format in which radio spectra profiles are often described. For a flux density (e.g. in Janskys), S_ν which varies with frequency, ν, we define a power law relationship with spectral index α of $S_\nu \propto \nu^\alpha$.

2.1.1 Blackbody

Blackbody radiation describes the electromagnetic radiation produced by matter in thermal equilibrium with its environment, with a temperature above absolute zero, converting the thermal energy of the body into radiation via vibrations of charged particles. The spectrum's shape can be defined by the single parameter of temperature of the system, and thus gives rise to the common nomenclature of 'thermal' radiation. An optically thick astrophysical 'surface' or object that represents a thermal distribution of charged particles can be a blackbody source (at least in part), including stars (marred by absorption and emission features), thick accretion discs (actually multiple blackbodies) and the original fireball that produced the radiation that we now see as the cosmic microwave background (the most ideal natural blackbody spectrum known to exist, see top left of Figure 2.1). The power emitted by a blackbody initially increases with frequency with a spectral index of $\alpha = 2$ at low energies, where Planck's law breaks down into the simplified Rayleigh-Jeans law. As frequency increases, quantum effects become more important and curb the total emitted flux density, resulting in a peak at a frequency proportional to the body's temperature, defining the dominant 'colour' of the emitted radiation, before transitioning into a steep decline (the 'Wien tail' which follows the Wien approximation: $S_\nu \propto e^{-h\nu/kt}$) in radiated power towards higher frequencies.

2.1.2 Bremsstrahlung

Bremsstrahlung, or free-free, radiation is the result of acceleration of free electrons via deflection in the electric fields of ions, making it prevalent in hot ionised plasmas such as those found in HII regions. Emitted photon properties depend on the energy of the affected electrons, plasma ion composition and impact geometry, allowing the mechanism to produce photon energies up to and including gamma-rays though the radio regime is most relevant when discussing XRBs (and their environments). For a thermal population of electrons the calculated spectrum, ignoring any absorption effects, can effectively be described as flat ($\alpha \sim 0$, but slightly negative owing to the Gaunt factor) across all frequencies up to a high energy cut-off defined by the highest possible electron energy (related to the overall temperature of the population). At high frequencies, where photons are energetic enough to pass through the plasma unhindered (optically thin), the spectrum remains unchanged and thus flat. At low radio frequencies, self-absorption occurs preventing low energy electrons from escaping (optically thick) and causing a turnover to a thermal spectrum of α

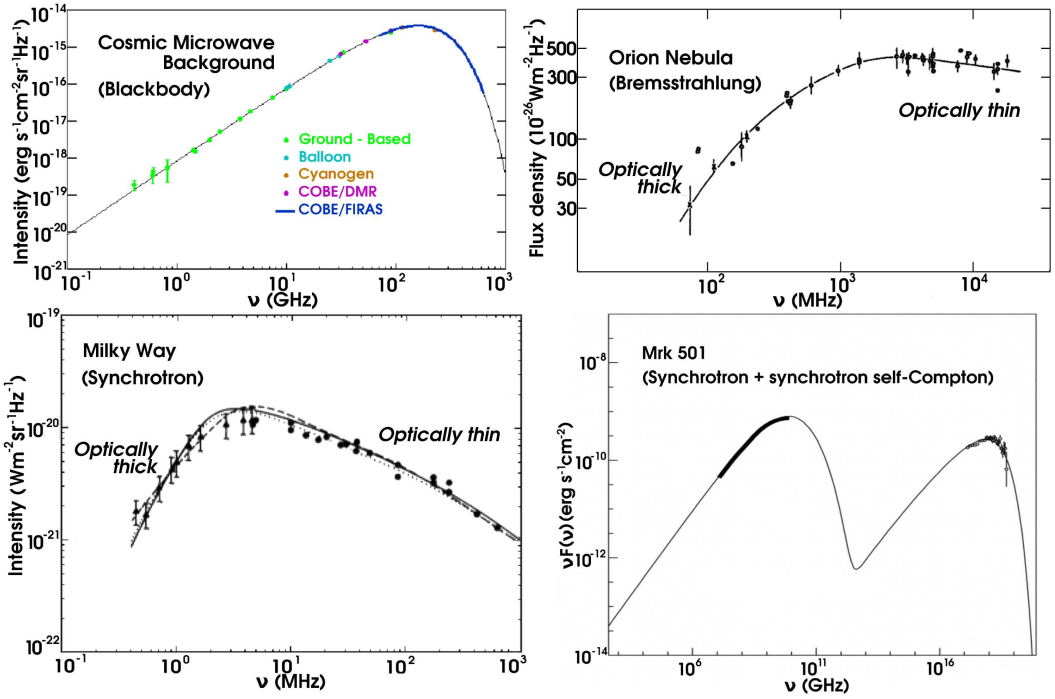

Figure 2.1: Emission process spectra. Top left: The cosmic microwave background blackbody spectrum that traces out a near perfect Planck curve (adapted from `http://arcade.gsfc.nasa.gov/cmb_spectrum.html`). Top right: Radio emission from the Orion Nebula (M42) revealing a bremsstrahlung spectrum where the turn-over due to optical thickness variation is visible (adapted from Terzian & Parrish 1970). Bottom left: Radio synchrotron emission from our Milky Way Galaxy with various spectral fits (solid, dashed and dotted lines) over-plotted (adapted from `http://gtn.sonoma.edu/images/MW_Synch.png`, originally seen in Cummings, Stone & Vogt 1973). Bottom right: X-ray to γ-ray spectrum of the active galactic nucleus Mrk 501, showing fitted synchrotron (left hump) and synchrotron self-Compton (right hump) models (adapted from Konopelko *et al.* 2003).

= +2 (i.e. Rayleigh-Jeans limit blackbody). An example spectrum (Orion nebula) is shown in Figure 2.1 (top right).

2.1.3 Synchrotron

Synchrotron radiation also results from the deflection of electrons, but via magnetic fields rather than ion charges as is the case from bremsstrahlung (hence why synchrotron is sometimes referred to as magnetobremsstrahlung). It is essentially an extension of cyclotron radiation to relativistic energies (i.e. requiring relativistic electron velocities). The path of a relativistic electron moving in a plasma containing a significant magnetic field is bent into a helix along the direction of the field lines (i.e. derived from $\bar{v}\times\bar{B}$), with the circular acceleration component giving rise

to the electromagnetic emission in the direction of the electron's motion.

The strength of the magnetic field and electron energy define the gyrofrequency of the electron's helical path, which in turn affects the frequency of the emitted radiation. Typically, astrophysical synchrotron emitting electron populations have non-thermal power law distributions of the form $dN(E) = N_0 E^{-p} dE$, where p represents the spectral index and is related to α via $\alpha = (1-p)/2$. For such populations, the resultant emission spectra have $-2 \leq \alpha \leq 0$ with values typically around $\alpha \sim -0.7$ in the optically thin regime. In the optically thick regime, one must consider the effects of synchrotron self-absorption, which for a power-law distribution yields a low-energy turn over to $\alpha = 2.5$, slightly steeper than the equivalent for a thermal population of electrons (i.e. $\alpha = 2.0$, Rayleigh-Jeans spectrum). Note that the final result is a peaked spectrum, whose maximum lies at the turning point between optically thick and thin regimes. An example spectrum is shown in Figure 2.1 (bottom left). The dependence of synchrotron radiation on field line orientation, together with magnetic fields tending to be highly ordered, yields a tell-tale side effect which can help to identify sources associated with the mechanism: synchrotron radiation can have a high level of linear polarisation ($\sim$ 70%) though the exact degree of polarisation depends on electron distribution and how ordered the magnetic field lines are.

2.1.4 Inverse Compton scattering

While not an independent form of emission (requiring a pre-existing source of photons), Compton scattering is an important mechanism in the production of unique photon populations and describes the interaction of high energy photons with matter. Often represented as an inelastic scattering process, the result is a transfer of energy from the photon to the particle, increasing the wavelength of the incident photon (the 'Compton effect'). While the particle may be an atom, the actual interacting body is of course a charged particle and most often an electron (which is expelled in the atom scenario), since their scattering capability is dependent on the Thomson cross-section: a value proportional to charge and inversely proportional to particle mass (hence hindering scattering via protons).

In the case of astrophysical studies we are more interested in the reverse of the above process, that is *inverse* Compton scattering. In this scenario, high energy electrons (again, possible for protons but less likely) interact with low energy photons and impart energy to them resulting in a decrease in wavelength (photon energy typically boosted by Γ^2, where $\Gamma = 1/(1\text{-}v^2/c^2)^{0.5}$ - the Lorentz factor). This phe-

nomenon can be used to explain the X-ray component in XRB spectra (discussed later) where low energy photons from the accretion disc are 'up-scattered' to X-ray energies by relativistic electrons in a hot corona or a jet. The effect is also of importance in the observed distortions of the cosmic microwave background by regions of high energy electrons (e.g. within hot galaxy clusters), known as the Sunyaev-Zel'dovich effect, and synchrotron self-Compton scattering (See Figure 2.1, bottom right), in which the same population of relativistic electrons that produce a synchrotron spectrum also up-scatter the resulting photons.

An important side-effect of the inverse Compton process to consider, is the 'comptonisation' of the initial spectrum. This describes a situation where the inverse-Compton process becomes a major (or indeed dominant) contributor to the energies of a population of photons, thus having a significant effect on the shape of the final spectrum. This can be a very complex effect to decompose as the final outcome depends on a large number of already complicated factors such as the initial photon energy spectrum and the properties of the scattering electron populations (e.g. density, velocity distribution). What must be kept in mind is that the effects of comptonisation can distort spectra enough so that the other emission mechanisms become difficult to distinguish. This is of particular concern when attempting to decipher the high energy region of XRB spectra (discussed further in the following section), or analysis of emission line profiles from inner accretion disc regions.

2.2 X-ray binaries

In this section I will summarise some of the basic scientific concepts related to X-ray binary studies, including 'standard model' theories which are often employed when interpreting observed behaviour. The term XRB encompasses several species of object that share the scenario of a compact, degenerate stellar corpse (black hole or neutron star) that accretes matter from a companion, or 'donor', star and in the process produce an emission spectrum with a significant contribution from X-ray photons (Figure 2.2). There are numerous sub-groups defined by additional parameters including companion type, accretion mechanics and behaviour, but for the purposes of this thesis we are most interested in low-mass X-ray binaries (LMXRBs: binaries in which the companion is of lower mass than the accreting compact object): specifically those that exhibit powerful radio jets.

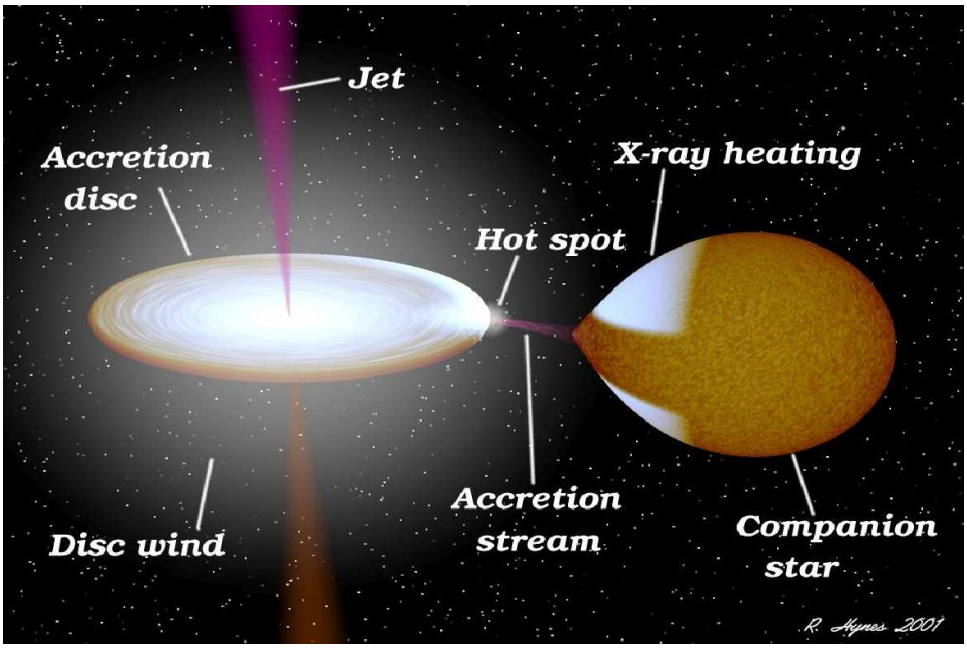

Figure 2.2: Diagram of an X-ray binary. This is an artistic impression of what an X-ray binary might look like, including many of major physical components (individually labelled) that can give rise to distinctive emission these objects are known for. Credit for the image goes to Prof. R. Hynes, Louisiana State University.

2.2.1 The compact object

Within the subdivision of LMXRBs there exists the further differentiation between systems that harbour different types of compact object. While all XRBs share a large number of characteristics due to the shared accretion mechanics, it is important to know the nature of the compact object as it may be tied to differences in the observed emission properties that can be studied to develop a better understanding of fundamental physics. The simplest way to identify the degenerate object in an XRB is to attempt to estimate the masses of the two bodies in the system via Kepler's laws, requiring that we know the orbital period, orbital separation (or alternatively, the velocity of one of the objects), and mass ratio of the two bodies. The period is often the easiest property to measure, as it can be revealed by light curve variations resulting from effects such as eclipses or companion deformation, or by spectroscopic study of emission line shifts (i.e. orbital motion Doppler effects); however, this is dependent on the inclination angle of the orbital plane (stronger effects visible in edge on systems). Orbital separation tends to be impossible to measure directly due to the distance of objects and observational resolution limitations; however, the aforementioned emission line studies also allow us to calculate the velocity of the companion along our line of sight, which together with a period

allows us to estimate the orbital separation as a function of inclination of the system (which may also be discovered as a result of eclipses or other means). Finally, the companion's spectrum and luminosity can allow us to identify its type and from that an estimate of its mass. Unfortunately, most of the above methods require one to be able to clearly observe the companion's spectrum, which is often not possible as a result of a combination of faintness and contamination from the accretion disc (which can be brighter than the donor star).

Assuming one can gain a reliable estimate for the compact object's mass, it is then possible to compare the value to the range of masses derived from the Pauli-exclusion principle for objects made up of degenerate material. The mass boundary between white dwarfs and neutron stars (the Chandrasekhar limit) defined by electron degeneracy pressure is approximately 1.4 $M_\odot$ (where $M_\odot$ is the mass of the Sun) though it is possible for a white dwarf to exceed this via rotation. Beyond this limit the white dwarf collapses under its own gravity, and in the absence of any counteractive force, results in electron-capture by protons to form neutrons and the creation of a neutron star. However, in reality a collapsing white dwarf may trigger runaway fusion, detonating as a type Ia supernova event rather than forming a neutron star. The next collapse limit is determined by the neutron degeneracy pressure (the Tolman-Oppenheimer-Volkoff limit); however, since we have yet to fully understand the required nuclear equations of state this limit is less well defined, with current estimates ranging between 1.5 and 3.0 $M_\odot$ (Arnett & Bowers 1977, Bombaci *et al.* 1995) and spin remaining an important factor. Beyond this, things get really messy, with theories predicting objects made up of exotic particles such as quarks; however, given that our understanding of equations of state degrades even further it is far more comfortable to assume that neutrons stars likely collapse into a singularity and form a black hole.

I need not focus on the specifics of white dwarf masses for the purposes of this thesis, but the range of masses where black holes and neutron stars may be indistinguishable from each other is significant, especially if there is already large uncertainty in the estimates themselves. As such, it is necessary to discuss other ways to differentiate between the two types of object. One of the more obvious differences that results in observational distinctions in the two species is the absence of an event horizon in NSXRBs, or rather, the presence of a surface on to which accreted matter can fall and interact. In the absence of the strong magnetic field required to create pulsar emission, matter accreted by a neutron star will collect and become compressed and heated under the extreme gravity of the star. The pressures and temperatures in these situations can become so high that thermonuclear burn-

ing becomes possible on the surface and triggers a massive rise in X-ray luminosity. Such events are known as Type-I X-ray bursts (reviewed in Lewin, van Paradijs & Taam 1993) and provide one of the few irrefutable indications of the presence of a neutron star; however, the absence of such bursts does not confirm the compact object to be a black hole either. In many cases, if a system whose compact object type remains in doubt displays behaviour akin to that only observed in 'confirmed' black hole XRBs, such as spectral profile similarities, familiar hardness-intensity tracks, or unique spectral components appearing during particular states or transitions, then a good case can be made for the unknown object also being a black hole. The Fourier analysis of X-ray variability can also reveal behaviours unique to each type of compact object, such as the hHz (hectohertz) and twin kHz quasi-periodic oscillations (QPOs) seen in NSXRB frequency-power diagrams (seen in 4U 1636-53 for example, Altamirano *et al.* 2008), though sometimes the differences are only revealed in particular source states. Nonetheless, these arguments are never perfect and only hold until someone discovers the same behaviour in sources from both groups.

2.2.2 The accretion disc

XRBs can accrete matter in three ways, as a result of Roche lobe overflow in which the donor star is sufficiently large and close to the compact object for the latter's gravitational field to strip matter away from the former, from the donor star's stellar winds, or a situation in which the donor spins rapidly enough to form an equatorial outflow (atmosphere of stellar disc) on which the compact object can feed (this method is exclusive to a subset of high mass XRBs known as Be XRBs). While stellar winds may remain a factor in LMXRB accretion, the dominant mechanism involved in such systems is that of Roche lobe overflow (see Figure 2.3) and is of most relevance to this thesis.

As a consequence of the orbital motion of the bodies in a binary system, the matter that falls towards the compact object via Roche lobe overflow actually has a significant amount of angular momentum imparted by the donor. This prevents the infalling matter from following a linear path, and instead approach the compact object at an angle normally significant enough to enter an elliptical orbit. The exceptions to this require long orbital periods (reducing the approach angle of the material) or a strong magnetic field associated with the compact object which alters the infalling material's trajectory (i.e. not applicable to black hole systems). The interaction of such an orbit with the companion's gravitational field leads to

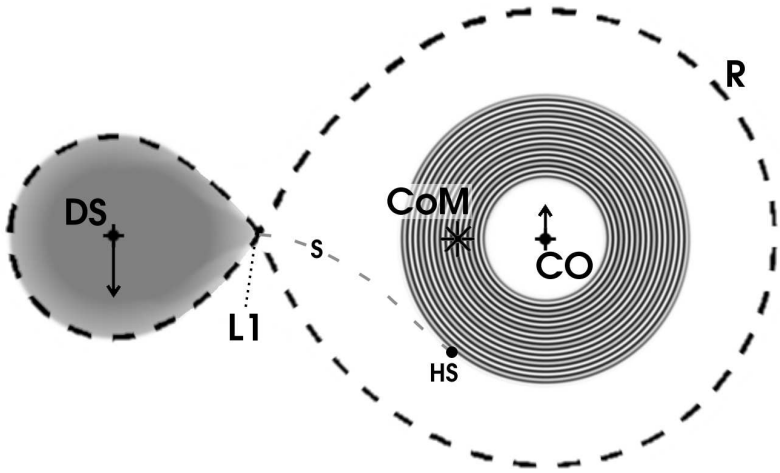

Figure 2.3: Roche lobe and accretion disc diagram. Labelled are the donor star (DS) on the left, which fills its Roche lobe (the boundaries of the lobes are defined by dashed line R) allowing matter to transfer across the inner Lagrangian point, L1, as an accretion stream following path S. As described in the main text, this matter builds up into an accretion disc, represented in this image as a series of circular orbits around the compact object (CO) on the right. Also labelled are the impact point between the stream and the disc, often referred to as the 'hot spot' (HS) and the system's centre of mass (CoM). While the Roche lobes, star positions and CoM are based on a system of mass ratio $M_{DS}/M_{CO} = 0.25$ (duplicated from figure 11 of Frank, King & Raine 1992), the depicted donor and disc sizes, stream path and relative velocity magnitudes (arrows on each star) are merely illustrative and are not intended to represent any particular system or model.

precession, which in turn results in collisions between different regions of the infalling stream. The general result of these interactions is circularisation of the orbit of accreted matter over time until it settles into the lowest energy Keplerian orbit: a ring around the compact object. The matter within this ring will continue to interact with itself and additional gas stripped from the companion via a number of possible mechanisms, such as collisions, viscous shearing, and magnetic interaction, causing the orbiting matter to heat up and radiate energy as a blackbody. The energy loss is sustained by reduction of angular momentum and thus orbital distance; however, the process is generally slow, making the drop in orbital radius a gradual effect. The path of radiating matter can be well approximated by a series of progressively smaller circular orbits around the compact object, and given a long or continuous epoch of accretion, infalling matter will fill a number of these orbits to form what is typically referred to as an 'accretion disc'. Thus, the disc serves to transport angular momentum outwards (typically, back into the binary orbit), while matter spirals in converting gravitational potential energy into kinetic and thermal energy: an incred-

ibly efficient conversion in the presence of degenerate stars (e.g. >10% of accreted matter rest mass for a typical neutron star of mass 1.5 $M_\odot$ and radius of 10km), radiating away a large fraction along the way (though black holes can incurr some additional loss due to hot matter passing the event horizon).

Understanding the structure and behaviour of matter in an accretion disc is both difficult and fundamentally important in also understanding the behaviour of phenomena associated with compact objects and X-ray binaries. Not only can the emission (blackbody and otherwise) from the hot gases within the disc dominate the observed spectrum from infrared to X-ray wavelengths, but the disc's structure and internal dynamics determine the rate at which matter actually reaches the compact object, which is in turn is believed to directly influence higher energy phenomena in the system such as the corona, jets and overall luminosity. It is no surprise then that a great deal of work has gone into determining useful models for XRB accretion discs.

To first order, XRB accretion flows may divided into two groups: radiatively efficient or radiatively inefficient (though it should be noted that other divisions exist that are, for example, based on optical depth and disc geometry). Radiatively efficient discs/flows refer to systems in which the accreted matter loses the majority of its energy (efficiency $\eta \sim 0.1$ in $L = \eta \dot{m} c^2$ for BH and NS) in the form of radiation and include the famous Shakura & Sunyaev (1973) geometrically thin, optically thick disc. This type of model describes a disc that begins relatively thicker at its outer edge, and becomes thinner as we approach the compact object, forming a near-concave cross-section (see Figure 2.4) before terminating at the inner most stable circular orbit (ISCO): the radius beyond which no stable circular orbits can form and matter becomes dominated by the radial velocity component (note that this distance is affected by both compact object mass and spin). Conversely, the matter in radiatively inefficient flows carries most of its energy all the way to the compact object and describes models such as advection dominated accretion flows (ADAFs; Ichimaru 1977; Narayan & Yi 1994). In the case of a BHXRB, this can mean a significant portion of the released gravitational potential energy actually passes the event horizon before it can be radiated away. These types of flows do not conform to the thin disc scenario since the unradiated energy is assumed to 'inflate' such regions. Some models combine the above two scenarios into a disc that transitions between radiatively efficient and inefficient disc regions, divided at a thin disc 'truncation' radius (Esin, McClintock & Narayan 1997).

The nature of angular momentum transport (disc viscosity) in these scenarios is difficult to pin down but important in explaining the rate of material flow through

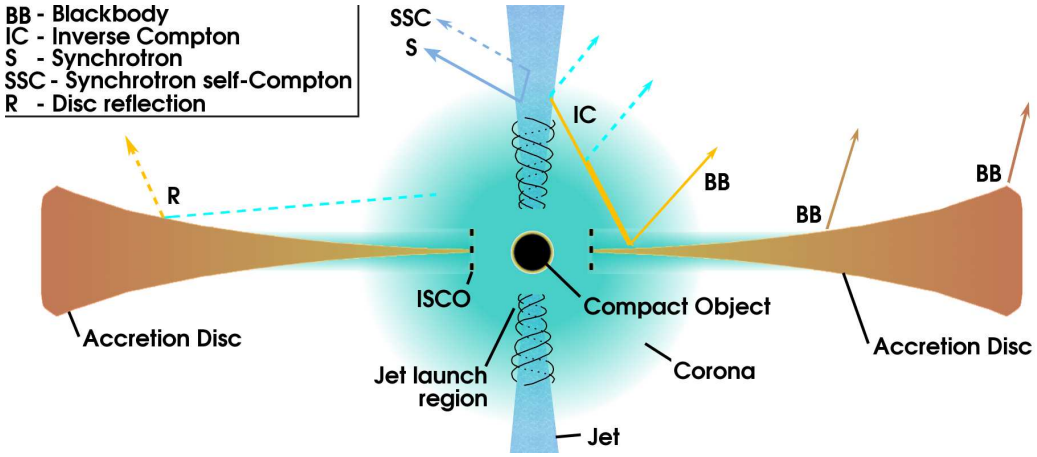

Figure 2.4: Generic accretion disc model cross-section with emission components. This figure shows a simple representation of a 'standard' model for X-ray binary structure, as well as various emission components and processes believed to contribute to XRB spectra. While I try to show most of the important contributors to an XRB spectrum, one should be aware that they do not all exist in every XRB and state. Note that the scales of various components (such as the compact object) are skewed and the shape of the disc exaggerated in order to better display important features.

the different disc radii. The Shakura-Sunyaev model, for example, employs an α parameter to account for what they regard as the two most important mechanisms: magnetic fields and turbulence. Indeed, it was later established that the magnetic field (locked into the accreted material) likely plays the greatest role in accretion disc viscosity, via magneto-rotational instabilities (Balbus & Hawley 1991). These instabilities arise when components of gas orbiting at different velocities (and thus radii) interact with each other through field lines, with closer faster components being slowed by gas in larger orbits, and vice versa, yielding the necessary transfer of angular momentum towards the outer disc. Furthermore, it now believed that large scale height magnetic fields associated with certain accretion disc models can become incredibly complex, to the point of not only affecting accretion mechanics, but driving the spectacular phenomena known as jets (see next subsection).

Crudely speaking, the basic spectrum of an optically thick disc can be represented as a sum of blackbody spectra from each Keplerian ring. However, this is not the final product of the disc's emission, as the blackbody photons can act as seeds for inverse Compton interactions (be it with jet or corona electrons), and the disc may contribute reflection components when it is illuminated by X-rays emerging from nearer the compact object (also shown in Figure 2.4).

2.2.3 Jets

The first astrophysical jets to be observed were those emerging from galactic nuclei and were quickly identified as powerful outflows of energy and matter from accreting black holes within their host galaxies (i.e. AGN such as Cygnus A: Hargrave & Ryle 1974). However, it has only recently been established that these phenomena are not unique to supermassive black holes, but in fact a natural result of accretion on all scales, with jets being detected in protostellar objects, interacting binary systems, and believed to be the cause of some types of gamma ray bursts. Study of jets is obviously important as they represent a powerful means of energy return to an accreting object's surrounding environment, as well as a plausible mechanism for angular momentum loss and thus an influence on accretion rate. On the largest scales, AGN jets are believed to play an essential role in regulating the growth of their host galaxies (McNamara & Nulsen 2007), whereas on stellar scales XRB jets produce beautiful complex jet-powered nebula which can record the history of their host system's activity. XRB jet systems have the additional benefit of allowing observation of variability on short timescales, owing to the fact that their host compact objects and accretion discs are smaller than those in their AGN counterparts (variability timescales $T_{Var} \propto M_{BH}$, e.g. McHardy 2010), which also vary but on timescales of the order of human lifetimes or more. Indeed, given that jets appear to be intimately tied to the process of accretion, it is intuitive to believe that what we learn of BHXRB jets can be scaled up to supermassive BH levels to better understand what is occurring in AGN. Nonetheless, though jets from XRBs have been actively studied for decades, they remain one of the most poorly understood features of XRB systems. I will briefly describe some of the most influential and widely used theories and models applied to various aspects of astrophysical jets, though full descriptions are beyond the scope of this thesis.

2.2.3.1 Formation

One of the most difficult questions to answer about jets is that of the nature of their formation, with questions on the subject normally followed by a great deal of 'hand waving' and a muttering of the words "magnetic fields". Though a large effort has been made to solve this problem over the years, including theoretical models, simulations and even laboratory reproductions, we are still unsure as to the exact process involved, or if indeed there is only one or several processes occurring in real systems. Initial theories involved thick accretion disc models, where the inner funnel-like region of the disc could help to collimate a wind from the compact

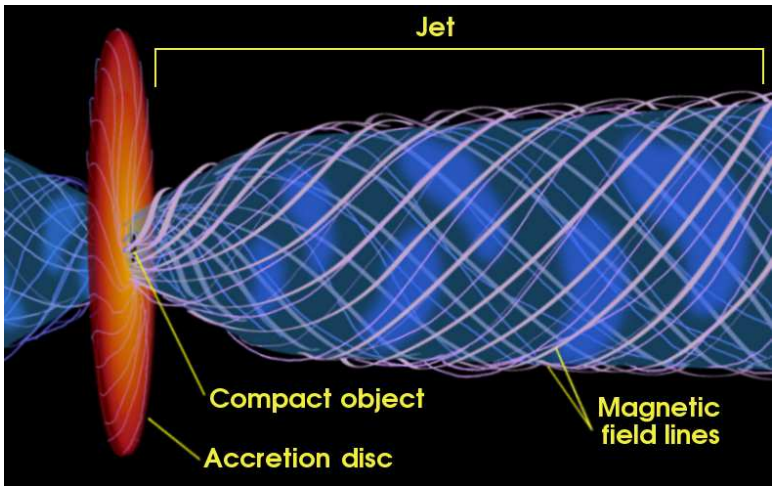

Figure 2.5: Artist's impression of Blandford-Payne magnetocentrifugal jet formation. Magnetic field lines become twisted, subsequently accelerating and collimating matter near the compact object into a relativistic jet. Modified from `http://www.nasa.gov/centers/goddard/images/content/96552main_jet_schematic.jpe`.

object (Rees *et al.* 1982); however, such models were eventually shown to be unstable (Papaloizou & Pringle 1985) though other thick disc scenarios may exist. A more attractive alternative was proposed by Blandford & Payne (1982) involving the magnetocentrifugal launching of charged particles. The basic interpretation of the model requires a disc with a large scale height poloidal magnetic field, whose field lines are frozen into the orbiting plasma and of which some are open and pitched at an angle lower than 60° from the plane of the disc. These field lines carry their own thin material, which is forced away from the disc via centrifugal acceleration as the lines are dragged around by orbital motion (the analogy of "beads on a wire" is often used); however, as distance from the disc increases, the inertia of this carried material forces the connected field lines into helical structures as their base remains connected to the rapidly rotating disc. The result is a cylindrical wall of toroidal field lines (see Figure 2.5 and panel C of Figure 2.6) whose magnetic tension acts to collimate material within it and continues to accelerate away from the disc: the beginnings of a jet. Interestingly, since thin disc fields should be of smaller scale height (bound close to the disc surface) this model would function better for thicker magnetised flows such as those presented in ADAF models (Meier 2001).

As with most models, this is an idealised situation and the original work notes several complications that can exist in real systems, including thermal effects from disc emission (affecting field line material acceleration and thus the required field-disc angle), reconnection events (field line pairs interacting and reacting elastically

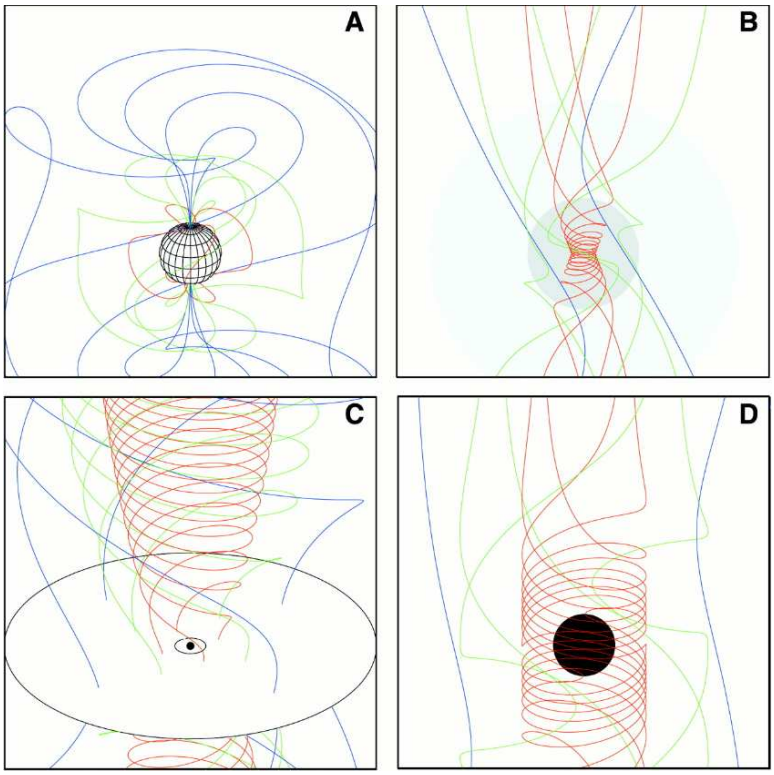

Figure 2.6: Jet formation methods via magnetic fields. A: Rotation of a neutron star's dipole magnetic field. B: Collapse of of an object drags and winds up an existing field. C: Poloidal magnetic field of a rotating accretion disc. D: Frame dragging near a black hole wraps field lines within the ergosphere. C is of most relevance to XRBs, though A and D may contribute depending on the nature of the compact object. B is shown for completeness, but applies to the collapse of stellar cores/corpses and thus possibly GRBs. From Meier, Koide & Uchida (2001).

to carry away energy) and instabilities occurring within the jet itself. Nonetheless, this work has become the foundation of our understanding of jet formation, being applied and built upon time and again to explain observations of ever more complex jet features and inspiring the use of magnetic field dynamics in explaining other jet formation scenarios (Meier, Koide & Uchida 2001; see Figure 2.6).

2.2.3.2 Composition

Another area in which there is still much debate is that of the actual composition of the matter carried by jets. Evidently, the material must be comprised of charged particles (though macroscopically neutral), but as observations of all but one jet system (on all scales) have only revealed the leptonic synchrotron emission component, we are only able to confirm the presence of electron, positrons or both. The only system in which we can confirm the presence of baryons within jets is the micro-quasar SS 433, with emission lines appearing in X-ray, optical and infrared spectra (Margon

1984). An explanation for this unique difference can be derived from SS 433's other rare jet property: a distinctly lower bulk flow velocity (0.26c) than that measured in other jets. The acceleration of heavier baryons would cost far more energy than the same process applied to a population of electrons; thus, an inverse correlation between flow speed and average particle mass should not be unexpected. However, even realising this, a small baryonic component cannot be ruled out given that related emission lines may be difficult to detect due to high ionisation parameters or extreme Doppler broadening (Mirabel *et al.* 1997).

2.2.3.3 Spectrum

Jets are one of the most well known sources of radio emission and we have had several decades to study and understand their behaviour at such wavelengths. Numerous observations have revealed several consistent features in X-ray binary radio jets: non-thermal (power-law) spectra, high brightness temperatures, and in several cases a significant level of linear polarisation. Together, these imply synchrotron processes are the source of jet emission (See section 2.1.3), specifically that of relativistic leptons within field lines following the jet stream (hence magnetocentrifugal launching fits very well, see section 2.2.3.1). An ideal, steady (an important distinction: see section 2.2.3.5) jet spectrum should thus begin with the expected α = 2.5 self-absorbed component and end with an optically thin region ($\alpha \sim$ -0.7) associated with the jets highest energy emission. However, unlike normal synchrotron emission which displays a single break between the two regimes and thus a peaked spectrum, steady jet synchrotron spectra include a near flat region (α = 0.0 - 0.3; Markoff 2010) between the two steeper slopes that can extend for the entire observable radio spectrum: a signature which is a direct result of the unique structure of jets.

As a jet travels, the leptonic matter may continue to be energised (by either the magnetic field or secondary events such as shocks: see section 2.2.3.5), but the jet also expands as internal pressure overcomes the collimation, leading to adiabatic cooling in addition to the already significant radiative cooling. The result is a gradual variation in the jet medium's synchrotron emission, with both self-absorption break frequency and lepton energies dropping (especially the number of leptons at high energies, since radiative cooling is $\propto E^2$ and adiabatic cooling is $\propto E$, where E is particle energy) as the spatial extent of the jet increases (such behaviour is well described in Blandford & Königl 1979). Indeed, it is possible to divide the jet into layers of optically thick material with specific emission properties, i.e. a 'strati-

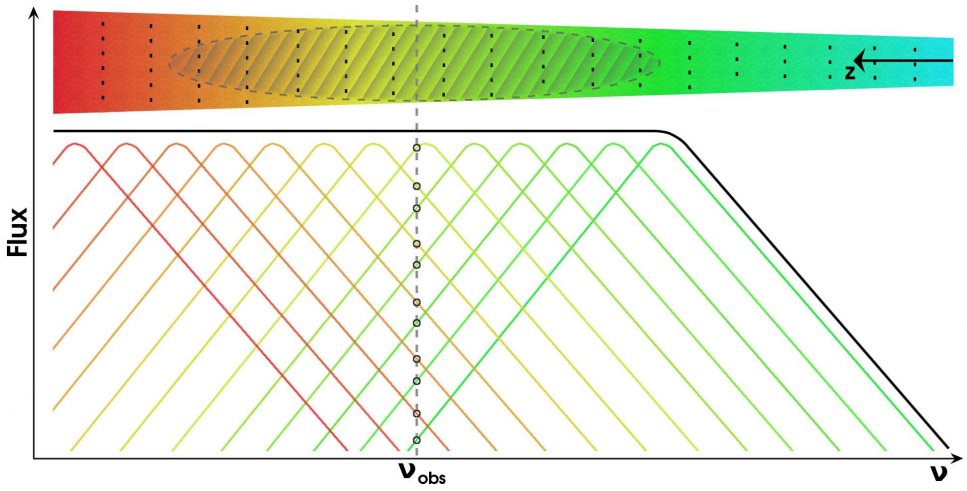

Figure 2.7: Stratified spectrum of an ideal compact jet. Each segment of the jet (base on the right, flowing in the direction of z) contributes a peaked self-absorbed synchrotron spectrum (simplified in the image) which when combined produce the typical flat spectrum displayed by jets (black line). Observations at ν_{obs} will primarily pick up the segment whose peak matches the observation frequency, as well as increasingly smaller components from segments to each side (as marked by the black circles). Thus the observed jet resembles an elliptical photosphere (marked by the grey dashed region), and is far smaller than the whole. This figure is based upon figure 1 of Markoff 2010.

fied' jet model. Each layer produces its own self-absorbed synchrotron spectrum with an associated peak frequency, but all with roughly the same spectral profile (assuming fixed equipartition between internal particle and magnetic field energies, and conservation of plasma properties along the extent of the jet) which when combined produce the tell-tale flat spectrum. Figure 2.7 illustrates this effect, as well as a possible observational effect when imaging at a single specific frequency ν_0: one only observes a smaller region of the jet centred on the layer whose spectrum peaks at ν_0 and with smaller contributions from neighbouring segments. Note that the frequency at which the jet transitions from optically thick and flat, to steep and optically thin (known as the jet break frequency) is important, as it provides information on the extent of the jet spectrum and thus its total power output, or clues about the composition and geometry of the outflow.

2.2.3.4 Velocity

It is now well known that jets from X-ray binaries reach relativistic velocities, with the aforementioned jets of SS 433 being some of the slowest at 0.26c, and 'transient' jet events yielding some of the fastest with multiple sources showing evidence for velocities $>$ 0.98c. However, in the absence of spectral lines the only means of

measuring the velocity of such jets is by capturing the motion of distinct identifiable components in the jet flow. In the case of a steady compact jet (i.e. resembling the simple models mentioned above) this is very difficult, as even a resolved steady jet will, as named, appear as a single unvarying stream. For this reason, there are yet no direct measurements of the speed of such jets, though there are some estimates based on observed secondary effects, such as jet asymmetry resulting from Doppler boosting ($>$ 0.6c for Cyg X-1, Stirling *et al.* 2001).

There are also scenarios where the jet is not steady (see the next subsection on shocks), otherwise referred to as transient jets, and are believed to reach similar or higher relativistic velocities than their steady counterparts. The advantage of observing a transient jet is that they often result in 'lumpy' flows with brighter regions of emission whose motion can be tracked and thus a jet velocity calculated (e.g. GRS 1915+105: Mirabel & Rodriguez 1994; GRO J1655-40: Tingay *et al.* 1995; XTE J1550-560: Hannikainen *et al.* 2001). While these motions implied velocities with Lorentz factors of $\Gamma \sim 2$, it was argued that the motion of such radio components could not provide an accurate upper limit for the flows with available distance measurements, and that it was actually possible to reach $\Gamma > 5$ (Fender & Kuulkers 2001, Fender 2003). Furthermore, the assumption that the motion of such objects directly reflect the properties of the underlying jet flows may be false.

At this point it is useful to review the basic formulae used to calculate the velocities of superluminal motions. The apparent angular shift (proper motion) of an object μ, in milliarcseconds per day, is related to an observed velocity parameter β (= v/c) by:

$$\beta_{obs} = (\frac{\mu}{170})D \tag{2.1}$$

where D is the distance to the source in kpc. This parameter is related to the true velocity parameter β_{re} via a known inclination to line of sight, θ, as follows:

$$\beta_{obs} = \frac{\beta_{re} sin\theta}{1 \mp \beta_{re} cos\theta} \tag{2.2}$$

where the $\mp$ represents the effects of approaching and receding components. For reference, superluminal motion (i.e. $\beta_{obs} > 1$) appears when $\beta_{re} \geq 2^{-0.5}$.

2.2.3.5 Shocks and interactions

While the events that can cause changes in jet flow properties will be handled in section 2.3, the resulting effects are simple to explain in the current context of jet

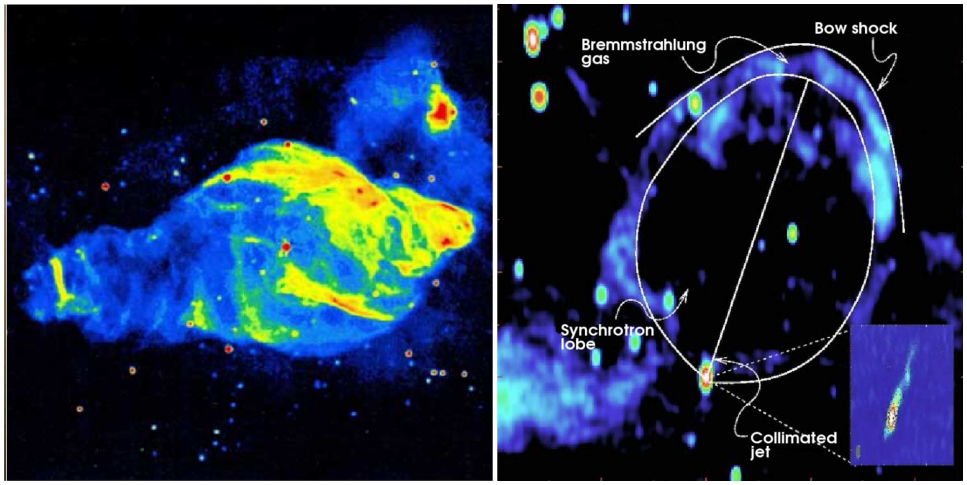

Figure 2.8: The famous jets of SS 433 and Cygnus X-1. To the left is a VLA radio (1408 MHz) image of the nebula W50 that surrounds the microquasar SS 433 (red point towards the centre) displaying spectacular arc like structures resulting from the interactions of a precessing jet with the nebula boundary (Dubner *et al.* 1998). To the right is a WSRT 1.4 GHz image of Cygnus X-1, displaying the bremsstrahlung bubble near the system: a compressed bow shock formed by the impact between the system's jet and the ISM (Gallo *et al.* 2005).

flows and fall under the title of shocks. Shocks describe the result of an interaction between two bodies of material where at least one of the components is travelling with supersonic speed, forming a shock wave where the bulk flow velocity changes abruptly: a collision which causes re-acceleration of particles in the impact region and a subsequent increase in emission. The collision itself forms two shock fronts, one moving forward and the other in the reverse direction, both of which accelerate leptons to higher energies (also into a new power law distribution) via first order Fermi acceleration (compressed magnetic fields in the shock fronts form distortions which bounce and energise leptons between them). Such collisions can be of two forms: internal or external.

Internal shocks (Kaiser, Sunyaev & Spruit 2000; Jamil, Fender & Kaiser 2010) form between two jet components moving at different speeds in the stream (often referred to as 'shells' of material). In the simplest models, relativistic momentum is conserved in inelastic collisions, and part of the released energy is believed to go towards accelerating particles in the merging shells. This often results in rebrightening of jet components downstream during events which emit faster jet flows (see section 2.3), without necessarily observing the energising stream itself (as in the case of Scorpius X-1 and Circinus X-1: Fomalont, Geldzahler & Bradshaw 2001 and Fender *et al.* 2004 respectively). These shocked shells eventually merge, and continue travelling with a new average velocity down the stream, while continuing

to expand and radiate via synchrotron processes like the original jet stream model. Since the component is compressed early on, the emission spectrum will steepen to a positive α as expected from optically thick emission, but gradually flatten again as expansion and cooling continues and the material transitions to the optically thin regime. The motion of such sites can be tracked to estimate flow velocity or their flaring behaviour correlated with that of the core emission in an effort to determine the velocity of the energising flow itself (as mentioned in section 2.2.3.4). Given a highly variable jet, it is possible for multiple shock sites to exist and for existing shocked shells to experience new collisions with fresh ejections or further downstream shells (e.g. XTE J1748-288: Hjellming *et al.* 1998).

External shocks describe collisions between a jet and a component not associated with the jet, such as the interstellar medium or supernova remnant material (in some cases that associated with the binary system's compact object). The main distinctions from internal shock models are: the larger difference in velocity between the two components (the jet is likely to be faster), and the composition of the external medium. While the former will result in a slower but intense shock front (released energy is partially dependent on the difference in shell velocities) than those seen within the jet, the latter can not only affect the merged shock's final speed, but the manner in which it radiates, since the external component is likely to contain a high concentration of atomic material to compress (i.e. bremsstrahlung can become the dominant emission mechanism, as in the case of Cygnus X-1; Gallo *et al.* 2005). Over time, the force of the jets can deform material around the impact site forming so-called 'jet-powered nebulae'; for example, the jet of Cygnus X-1 has created a large bremsstrahlung emitting bow shock with a synchrotron emitting void behind it (Figure 2.8, right panel). Such structures can be useful as calorimeters, allowing an estimation of the energy the jet has contributed to the ambient medium over its active lifetime. In the case of existing structures like supernova remnants, the jets act to deform them also, with a spectacular example in SS 433 and its surrounding nebula W50 (Dubner *et al.* 1998: Figure 2.8, left panel), which is believed to have formed as an expanding spherical nova remnant that has slowly been deformed by the jets in the east-west direction. SS 433 also serves as an excellent example of interaction sites in a nebula revealing the varying direction of the jets over time, in this case as a result of precession.

2.2.3.6 Precession

Returning to SS 433, we have the most famous example of the effects of jet precession: a situation in which the jet axis rotates periodically giving rise to helical outflow structure (individual ejections follow ballistic trajectories but in directions that vary with the jet axis), and in the case of SS 433, traces out colossal arcs as the flows impact on the interstellar medium at the boundary of a surrounding nebula (W50: see Figure 2.8, left panel). Since the cause of the precession is likely tied with the nature of jet production, we can again only speculate as to what it is; however, assuming a magnetic launch scenario, it is likely that the direction of the jet is linked to the orbital plane of the disc material whose field lines contribute to the jet's structure. Thus a scenario in which an accretion disc, or simply the deformed inner region of an accretion disc, 'wobbles' over time (i.e. disc plane changes over time) should be able to give rise to a varying jet direction, and such deformation and wobble can result from disc irradiation (Livio & Pringle 1997), loss of material via a disc wind (Schandl & Meyer 1994), misalignment of binary and disc orbits, Lense-Thirring effects due to a spinning compact object (Bardeen & Petterson 1975) or the effects of a neutron stars magnetic field (Lai 1999).

2.2.3.7 Instabilities

In the absence of precession or shocks, there are additional effects which can result in jets that show distinct distortions or deviations in their flow trajectory. Blandford and Payne already suggest that velocity shearing could cause downstream disruptions in the normal flow of a jet: an effect that is has been well studied and is known as a Kelvin-Helmholtz instability. This phenomenon can arise at the boundary between any two fluids (atmosphere layers, ocean currents, planet cloud bands) travelling side by side with different densities and flow speeds, and can be applied to jet models which include similar boundaries such as a fast core surrounded by a slower sheath (Norman *et al.* 1982). Another form of instability can arise earlier in the jet in regions where energy is dominated by poynting flux (i.e. the magnetic field is the largest energy reservoir) and yet to convert the majority of its energy into the kinetic flow. Known as a current-driven or 'kink' instabilities, smaller deformations in the jet brought about by internal magnetohydrodynamic waves can be intensified by the jet's own magnetic field, due to difference in the field strength on different sides of a forming kink (Nakamura & Meier 2004).

2.2.4 The corona

The corona describes a region of high energy plasma surrounding the compact object, possibly extending to the inner accretion disc (the cyan regions depicted in Figure 2.4). It is the corona which is identified as the source of high energy X-ray photons in several disc models, which begin as lower energy photons emitted by the disc but are upscattered (inverse-Compton) by relativistic electrons within the corona. However, it should be noted that there remains some discussion as to whether there is also some contribution to high energy X-rays from the base of the jet itself, and indeed whether the jet base and the corona may be intimately linked, with the corona material feeding the jet, or even that they are one and the same (explored further in section 2.3).

2.3 States, transitions and multi-wavelength correlations

Using standard accretion and jet models to decipher the behaviour of accretion onto compact objects is all well and good, but unfortunately XRBs are some of the most badly behaved (read: interesting) objects in the sky, often showing significant unpredictable variability, indicative of events that defy many of the assumptions used in simple modelling. It is now well known that black hole XRBs spend time in various ‘states’ defined by a number of common luminosity, spectral, X-ray timing and radio properties, and that over time systems evolve between these states, with some being triggered by periods of greatly increased luminosity called ‘outbursts’. Similar but less well understood behaviour can also be associated with neutron star XRBs. I will briefly describe some of the more popular mechanisms believed to influence these state changes, as well as the typical characteristics that mark the most well known states.

2.3.1 The Eddington limit

It is often useful to discuss relative luminosities in terms of fractions of the Eddington limited luminosity. This is the maximum luminosity possible for a spherically accreting, isotropically emitting object, and is derived by balancing the gravitational pull on matter by the object, with radiation pressure exerted by the emitted photons. The result for pure hydrogen accretion is given by

$$L_{Edd} = \frac{4\pi GMm_p c}{\sigma_T} \tag{2.3}$$

where m_p is the mass of a proton, and σ_T the Thomson cross-section, though the latter is only used assuming scattering by electrons is the dominant opacity mechanism (this can be replaced by a more complex absorption coefficient). For an astrophysical source this simplifies to 1.38×10^{38} (M/ $M_\odot$) erg s^{-1}. This can of course be violated if any of the original constraints are violated (e.g. emission beaming in jets or on-spherical accretion via an accretion disc).

Note that we can crudely relate the rate of accretion to luminosity, and thus the Eddington limit, by assuming all converted gravitational potential energy is is eventually radiated by matter falling towards the compact object. In such a scenario the luminosity L_{acc} given my matter accreted at rate $\dot{m}$ from infinity to radius R is given by

$$L_{acc} = \frac{GM\dot{m}}{R} \tag{2.4}$$

where M is the mass of the compact object.

2.3.2 Variable accretion and disc instability

As has been described, the emission of many XRB features are tied to matter and fields within the accretion disc, which in turn are affected by the rate at which material enters, passes through, and finally leaves the disc and falls onto the compact object. Therefore, it is logical to assume that variation in any one of these factors may lead to a change in observed emission characteristics.

While a change in the properties of the donor star (such as evolution over its lifetime) can alter the degree to which it fills its Roche lobe and thus rate of matter loss, such events are rare and the chance of catching a binary system at one of these turning points unlikely. Elliptical orbits can provide a stable constantly varying Roche lobe model, and thus an ideal scenario in which the compact object is better able to strip matter from the donor star at certain orbital phases, namely those closer to periastron. Such phases should therefore be punctuated by an increase in matter passing through the disc onto the compact object and a subsequent boost in luminosity, though a delay may be observed owing to the time required for the extra matter to pass through the accretion process. Indeed, this is the explanation for the flaring behaviour of some high mass X-ray binaries (e.g. GX 301-2: Layton *et al.* 1998), Be X-ray binaries (stripped from a circumstellar disc: see section

2.2.2), and the NSXRB Circinus X-1 (one of the primary sources studied in this thesis, see Chapters 5 and 6). Since the timing of these flares are related directly to orbital phase, they are periodic events which can be predicted and studied intensely in order to better understand periods of high accretion rate in XRBs. Donor star winds can also contribute to the matter accreted by the compact object, and so variability in the wind (for example in a 'clumpy' wind) can also lead to variability in emission from accretion processes. A particularly interesting example is that of Cygnus X-1, in which a line-driven wind from the companion star is slowed and subsequently accreted as it passes within the 'Strömgren zone' around the compact object: the region which can be photoionised by the X-rays (Tarasov, Brocksopp & Lyuty 2003).

The majority of XRB outbursts and state changes cannot be associated with variability in initial matter accretion rates as they are often very bright events that occur more than once and on timescales longer than a binary period, straying from the above scenarios. Indeed, it is likely that the rate of accretion from the donor for many XRBs remains stable, leading us to seek another explanation for their variable emission. The answer lies within the accretion disc itself, and the ability for the rate of matter transfer from one region of the disc to another to change abruptly as a result of instabilities. If such instabilities occur on a large scale in the disc, then the rate of matter accretion onto the compact object itself could vary dramatically. Returning to the Shakura-Sunyaev disc model, there are two primary types of instability which can cause such effects: one associated with the ionisation of hydrogen and the other with radiation pressure.

The hydrogen ionisation instability occurs at temperatures close to that required to ionise hydrogen (5×10^3 to 5×10^4 K), where the opacity of the gas increases sharply as photons begin to be absorbed in the ionisation process. The absorption heats the gas further, which produces more ionising photons, leading to a runaway heating effect (thermal instability) in that region of the disc, until most of the hydrogen is ionised. This thermal instability then triggers a viscous instability (rapid increase in mass loss from the disc region) since viscous stress is proportional to pressure in the Shakura-Sunyaev model (incorporated into the α parameter), and results in material quickly being dumped (into a smaller orbital radius) at a rate faster than it is refilled by accretion from the outside. This reduces pressure in the unstable disc region, which in turn cools the disc by reducing the heating effect, allowing the remaining hydrogen to recombine. However, once both instabilities have ceased to exist, accretion from outside is now sufficient to increase matter in the region and start the entire process over again, creating an instability cycle. Though

this describes the effect in only a single disc region, it should be obvious that the viscous instability and dumping of matter into subsequent accretion disc radii can trigger the effect in them as well, leading to a global cycle, with the disc alternating between the standard Shakura-Sunyaev model where accretion is the same at all radii, and a cooler neutral disc with decreasing accretion rate towards the inner disc. Though this type of instability was originally invoked to explain cataclysmic variable (CV) outbursts (Cannizzo, Ghosh & Wheeler 1982), the possibility that the same process could occur in XRB discs was evident, though with one significant alteration: disc irradiation. XRB outbursts show a longer exponential decay than those of CVs, which is best explained by irradiation of the outer disc by emission from the hotter inner disc and corona delaying the cooling of the outer disc (van Paradijs 1996).

The second instability arises at high mass accretion rates, in the innermost regions of a disc where material goes from being gas pressure dominated to radiation pressure dominated (i.e. $P_G = nkT$ to $P_R = \sigma T^4/3c$). In such a regime, a small increase in temperature yields a large increase in pressure, which boosts heating as in the model above (viscous stress is proportional to pressure). Again, the matter undergoes runaway heating. The Shakura-Sunyaev model predicts destabilisation of the disc at this point, with no stabilising scenario to counteract it. This can be counteracted by assuming that some energy can be carried away by advection to neighbouring disc regions and radiation (Abramowicz *et al.* 1988). However, observations show that discs that should suffer this instability (based on their Eddington ratio, L/L_{Edd}) show stable systems instead, indicating that the pressure calculations for the inner disc overestimate heating and our understanding of the inner accretion disc is incomplete. The only system which may show signs of suffering from this type of instability is GRS 1915+105, which spends a significant amount of time at near-Eddington luminosities, and thus accretion rates capable of triggering this scenario (Fender & Belloni 2004; Neilsen, Remillard & Lee 2011).

2.3.3 Black hole X-ray binary outbursts and states

Whether it is triggered by increased accretion or disc instabilities, the sudden influx of matter passing through the disc and onto the compact object increases emission over a wide range of wavelengths. These events are known as outbursts and often provide the first detections of new X-ray binaries via all-sky X-ray monitoring, due to the significant boost to X-rays from the inner disc and corona. Lightcurves for three black-hole X-ray transients are shown in Figure 2.9, illustrating how the

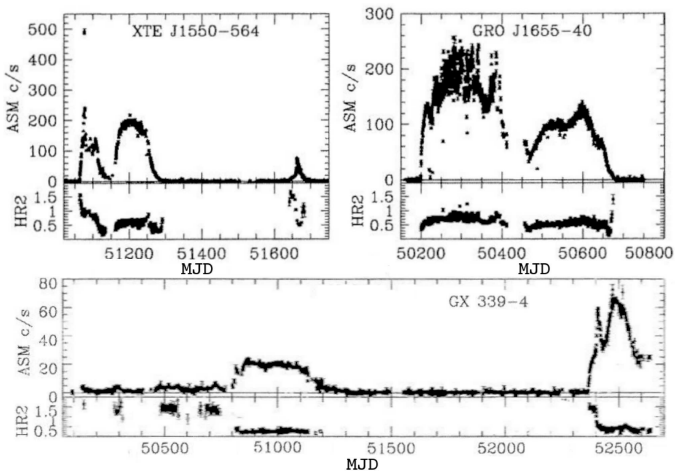

Figure 2.9: BHXRB RXTE ASM lightcurves. Each shows the unique outburst behaviour of the sources, with various profile shapes and quiescence durations. Plots were scanned and modified from McClintock & Remillard (2006), figures 4.2 and 4.3.

event profiles and duration can differ significantly, even for outbursts from the same source; however, two properties are often repeated: a sharp rise and a long term decay (continued irradiation of the disc, see section 2.3.2). In the case of non-elliptical XRBs (i.e. those with a relatively stable accretion rate), outbursts normally punctuate periods of quiescence lasting months, decades or perhaps even centuries. In this quiescent state XRBs tend to be at their faintest in the X-rays (in some cases up to 10^{-8} fainter than during outburst: A0620-00, Narayan, McClintock & Yi 1996), with a spectrum dominated by a power law component (photon index Γ = 1.5 - 2.1, defined by $N \propto E^{-\Gamma}$ where N is photon count rate in counts s^{-1} m^{-2} keV^{-1}, and E is photon energy).

When not in quiescence, XRBs are typically found in one of two brighter states: the hard and soft states. The hard state (otherwise known as the low-hard state) can effectively be considered a brighter version of the quiescent state whose name is derived for the continued dominant presence of the power-law component (Γ = 1.5 - 2.1) in the X-ray spectrum (see left hand panels of Figure 2.10), though a weak thermal component can still be detected at lower energies (peak at $\sim$ 1 keV, see GX 339-4's spectra in Figure 2.10). The soft state (also known as the high-soft or thermal dominated state) is identified by a spectrum that is dominated at lower energies ($<$ 10 keV, peak at $\sim$ 1 keV) by a thermal component, but retains a steep power law (Γ = 2 - 5) that is visible towards higher energies (see right hand panels of Figure 2.10). Note that there are also various timing features which provide additional means of differentiating between the above states. Hard state power spectra typically start flat at low frequencies, breaking between 0.01 and 1 Hz into a power-

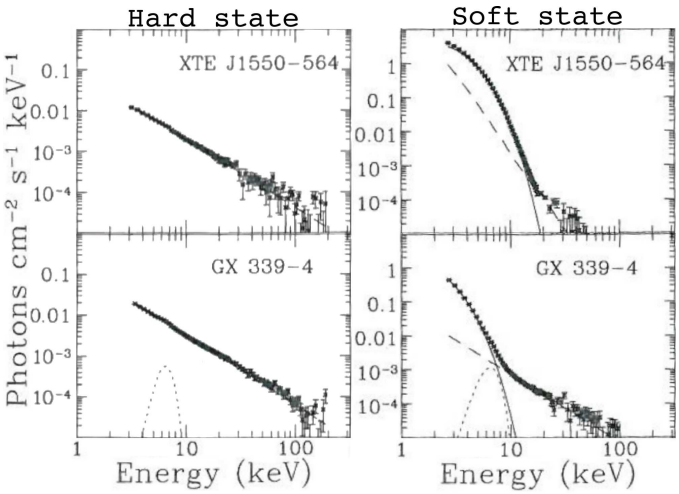

Figure 2.10: RXTE X-ray spectra of BHXRBs in the hard and soft states. The hard spectra on the left show the characteristic flat power-law component (dashed line) dominating over smaller contributions from other sources such as the accretion disc (dotted line). Conversely, the right hand spectra show a much stronger thermal component (solid line) with the power-law component dominating only at higher energies (dashed line). Plots were scanned and modified from McClintock & Remillard (2006), figures 4.8 and 4.11.

law ($\gamma \sim$ -1.5, for ν^{γ}), and again at 1 - 10 Hz into a slightly steeper component ($\gamma < 2$). Soft state power spectra are best described by a single power law ($\gamma \sim$ -1) and, typically, lower levels of integrated variability. There are also additional XRB states which correspond to transitional phases between the hard and soft states (the so called intermediate states), which can show similar behaviour to one of the above regimes (including both thermal and power-law spectral components), but at radically different X-ray luminosities and/or with unique timing features (e.g. the steep power-law or very high states) .

The behaviour of the radio jets appears to be intimately tied with what is observed from XRBs in X-rays and other wavelengths, with radio emission typically visible and persistent in the hard state (including during quiescence: e.g. A0620-00; Gallo *et al.* 2006), flares during transitions between soft and hard states (the exact point at which this occurs remains in doubt) and is much weaker or undetectable during the soft X-ray state. This behaviour is discussed further in the following section.

2.3.4 The hardness-intensity track and a unified model

Monitoring of outbursts of multiple BHXRB sources has revealed that the majority follow a very similar sequence of transitions. This sequence is illustrated very well

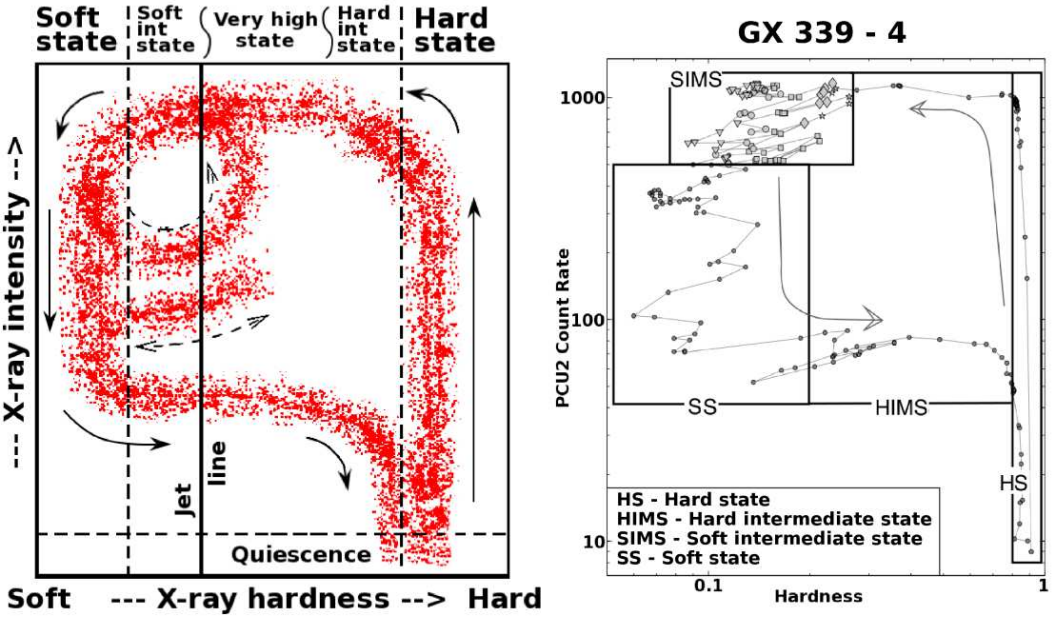

Figure 2.11: Hardness intensity diagrams for a simplified BHXRB model and GX 339-4. The left hand image shows a simple model track that BHXRBs may progress through as they transition through various spectral states (marked) during their lifetime (modified from figure 7 of Fender, Belloni & Gallo 2004). The outermost loop best represents the 'main' track for a single complete transition from quiescence (bottom right), through the hard (right side), intermediate (horizontal branches), and soft states (left side). The thick solid vertical line marks (crudely) the 'jet line', the boundary representing the activity of the jets from the source, with them inactive on the left side (soft state). The inner tracks represent deviations that some sources (GRS 1915+105 being the best example) display during their outbursts, moving in and out of the soft state and across the jet line, resulting in multiple radio flares. The right panel shows the evolution of GX 339-4 during its 2002/2003 outburst (modified from figure 18 of Belloni *et al.* 2005) with different state regions marked, and displays an excellent resemblance to the model's outer track.

by comparing the X-ray intensity of a source to its X-ray hardness (akin to the ratio of contributions from the power-law and thermal components in the spectrum) in a hardness-intensity (HI) diagram, tracing out an evolutionary track over a complete outburst (Fender, Belloni & Gallo 2004: see Figure 2.11). The primary track of the model describes the brightening of a source as it moves from quiescence into the hard state, softens (thermal spectral component becomes more dominant) during the peak of the outburst (passing through intermediate states and possibly the very high state) moving into the soft state where it declines before hardening once again (back through the intermediate states), shifting back into the hard state and finally returning to quiescence. The 2002/2003 outburst of BHXRB GX 339-4 is an example of real source behaviour that strongly resembles that which is described by

the model scenario (Belloni *et al.* 2005: right panel of Figure 2.11). An important feature of the HI diagram is the so-called 'jet-line', representing the point at which jets can turn on or off during the sources evolution. Note that the vertical jet line in Figure 2.11 represents an early, crude description of the feature, with more recent studies indicating it is more likely to be variable or tilted (Körding, Jester & Fender 2006; Klein-Wolt & van der Klis 2008; Fender Homan & Belloni 2009).

It is important to note that every outburst can display different behaviour, reaching different luminosities during the first horizontal branch (some near Eddington, others only a few percent), or transitioning back and forth between some of the states. Some sources display these types of behaviour more often or even exclusively, for example GRS 1915+105 is often found at near Eddington luminosities, and appears to bounce in and out of the soft state, with multiple radio ejection events (Klein-Wolt *et al.* 2002), suggesting it also moves across the jet line repeatedly: its behaviour could be well represented by the closed loop or spur marked on the left hand side of Figure 2.11.

Given this behaviour, it may be possible to describe the state and jet changes as a sequence of events related to the evolution of the accretion disc and its instabilities (Mirabel *et al.* 1998; Fender, Belloni & Gallo 2004). A crude scenario imagines a mostly radiatively inefficient disc being 'loaded' during the quiescent state (perhaps a truncated disc such as that of Esin, McClintock & Narayan 1997), with the system becoming brighter as the disc grows and accretion rate onto the compact object rises. As required, the spectrum in this state is dominated by the hard power-law component from a hot corona. During this phase, the jet remains relatively stable, but grows in luminosity with accretion rate. The disc contributes the thermal X-ray component, which gradually intensifies as the disc gains material. Eventually, the disc becomes hot enough for instabilities to kick in and trigger rapid migration of material to the inner disc regions. This process is fast, producing a geometrically thin, optically thick, radiatively efficient disc all the way to the ISCO (akin to the Shakura-Sunyaev disc) which in turn caused the thermal contribution to the spectrum to increase dramatically. Since Shakura-Sunyaev thin discs may find it more difficult to support the high scale height fields required to form a powerful jet (see section 2.2.3.1) the sudden propagation of the thin disc inwards could interfere with jet formation. The existing collimating field may be forced inwards, where magnetic stresses or perhaps the rotational effects of the black hole's ergosphere (see panel D of Figure 2.6) result in a significant boost in jet velocity and thus radio flaring via internal shocks, but finally collapses or is left significantly weaker as the thin disc reaches the ISCO. Eventually, enough disc material is lost from the inner regions

in the process that the system cools and fades, returning to the original quiescent scenario. While this sequence of scenarios can begin to describe what we observe, I must stress again that this is a very basic model and the changes that occur during real state transitions are likely to be far more complex. The described sequence, for example, fails to explain why we can observe both a bright optically thick disc and strong jets at the same time in some bright hard states.

2.3.5 Neutron star X-ray binary states

The notion of states is not isolated to black hole XRBs but also occurs in their neutron star counterparts, albeit with some differences. There are two classes of NSXRBs (originally classified in bright NSXRBs by Hasinger & van der Klis, 1989): Z-sources and atoll sources, named for the shape of their colour-colour diagram (CCD) tracks (i.e. comparison of hard X-ray intensity to soft X-ray intensity: see Figure 2.12). Z-sources are NSXRBs that persistently accrete at levels close to Eddington, and thus are some of the brightest stable X-ray sources in the sky. Their typical variability evolution traces a one dimensional z-shaped track, moving through the horizontal, normal, and flaring branches as luminosity increases, with radio being strongest in the horizontal branch, and weakest in the flaring branch. Atoll sources are fainter and are thus believed to accrete at lower rates ($<$ 10% Eddington), with their earliest CCD shapes showing a curved region referred to as the 'banana' branch, and a long-lived faint state in which sources were so often observed that the CCDs often displayed an isolated patch away from the banana branch, thus giving rise to the name of the group. Later CCDs showed a further branch in some atoll sources at very low luminosities that created a shape far more similar to Z sources, implying a connection between the two species (e.g. 4U 160852: Gierliński & Done 2002). Other sources also show behaviours from both sub-classes at different times (including Cir X-1: see chapters 5 and 6), with one of the best examples being observations of XTE J1701-462 which showed distinct transitions between atoll and Z-source states, attributed to variation in the source's accretion rate $\dot{m}$ (Homan *et al.* 2010).

Radio emission is typically observed in the faint harder regions (island states in atolls, horizontal and normal branches in Z-sources) of the CCDs, implying a similarity between these states and the hard and soft states of BHXRBs. In fact, the atoll NSXRB Aquila X-1 has shown an outburst HI diagram track that shows striking similarities to the BHXRB model (Tudose *et al.* 2009).

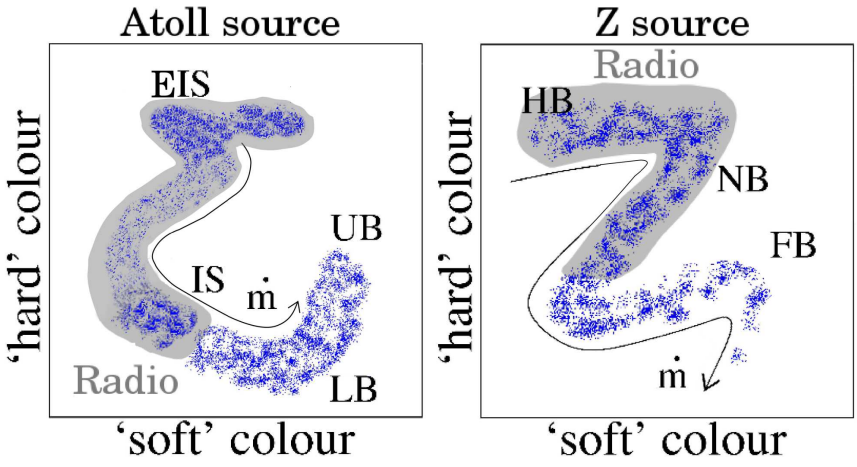

Figure 2.12: Schematic colour-colour diagrams for atoll and Z-source NSXRBs. The evolution with accretion rate $\dot{m}$ and hypothetical regions associated with radio emission (grey) are marked. Various track states and branches are labelled; extreme island state (EIS), island state (IS), lower banana branch (LB), upper banana branch (UB), horizontal branch (HB), normal branch (NB) and flaring branch (FB). Image modified from Fender (2006) figure 9.11.

2.3.6 Emission correlations

Given the similarities in the behaviour of radio and X-ray emission during the hard state rise, a comparison between the two types of emission was inevitable. A correlation between the two was first indicated in simultaneous multiwavelength studies of GX 339-4 (Hannikainen *et al.* 1998, Corbel *et al.* 2000) but was later found to exist for the majority of hard state BHXRBs (Gallo, Fender & Pooley 2003) with a relationship of $L_R \propto L_X^{0.7}$ (a value which can actually be derived from existing synchrotron jet and radiatively inefficient accretion models). This implies some common physical process in BHXRBs couples the hot X-ray emitting region near the compact object to the optically thick portion of the jet. A possible explanation for this is that the X-ray emission originates in the jet base itself, or similarly, that the jet is fed by the hot corona. A similar type of correlation was also found for NSXRBs but with a steeper slope of $L_R \propto L_X^{\sim 1.4}$ (Migliari *et al.* 2003) which is likely a result of accretion flows in neutron stars being radiatively efficient (Migliari & Fender 2006).

Recent updates to the correlation data with new observations from existing and additional sources has complicated the overall picture (See Figure 2.13). Several sources were found to lie below the existing track, and though initially easy to dismiss as outliers resulting from errors (from incorrect distance estimates for example) this population has grown to the point where it is more statistically likely that there are two distinct correlation tracks rather than one (one 'radio quiet' and

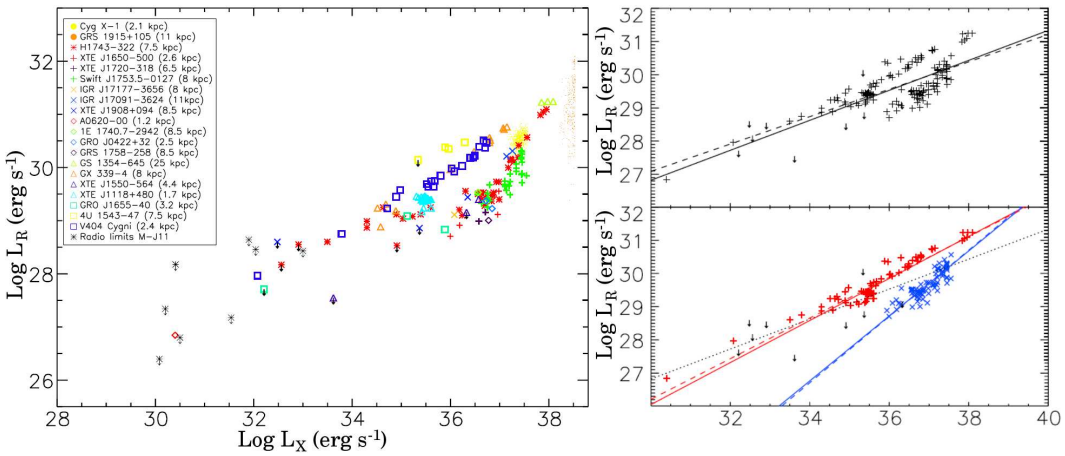

Figure 2.13: Radio/X-ray correlation of hard state BHXRBs. Figures are from figures 1 and 6 of Gallo, Miller & Fender (2012). On the left is one of the most recent collection of radio and X-ray luminosities for BHXRBs (18 simultaneous, 2 non simultaneous, plus some radio upper limits) for which a positive correlation is clearly visible. On the right are two sets of trendline fitting results, the upper right panel depicting the fit for a single population (as was originally considered in Gallo, Fender & Pooley 2003), with a slope of 0.45. The lower right panel shows the fit results for two populations, with slopes of 0.63 and 0.98 for the upper and lower tracks respectively. Note that dashed lines represent fits in which upper limit points are ignored.

the other 'radio loud': Gallo, Miller and Fender 2012). Further analysis has shown the lower track follows a steeper correlation of $L_R \propto L_X^{0.98}$ straddling the values for 'standard' black holes and neutron stars. These two correlations might at first be thought to define two distinct groups of black holes; however, the sources GRO J1655-40 and H1743-322 have been found to exist on both tracks at different times. H1743-322's correlation was studied intensely in Coriat *et al.* (2011), where it was found that the source actually exhibited a break in its correlation and that it could move from the radio loud track to a neutron star like $L_R \propto L_X^{1.4}$ track. This was interpreted as either the accretion flow changing from radiatively inefficient (radio loud) to efficient (radio quiet) during the hard state, or a different relationship between jet kinetic power and accretion rate.

Two groups independently took the radio/X-ray correlation further by testing if similar behaviour was exhibited by AGN; Merloni, Heinz & di Matteo (2003) and Falcke, Körding & Markoff (2004) They found that with the inclusion of black hole mass as an additional parameter, all weakly accreting (i.e. hard state) black holes lie on a plane (in the 3-dimensional space defined by mass, core X-ray luminosity and core radio luminosity) defined by the function $\log L_R = (0.6 \pm 0.1)\log L_X + (0.8 \pm 0.1) \log M + (7.3 \pm 4.1)$. This was dubbed the 'fundamental plane of black hole

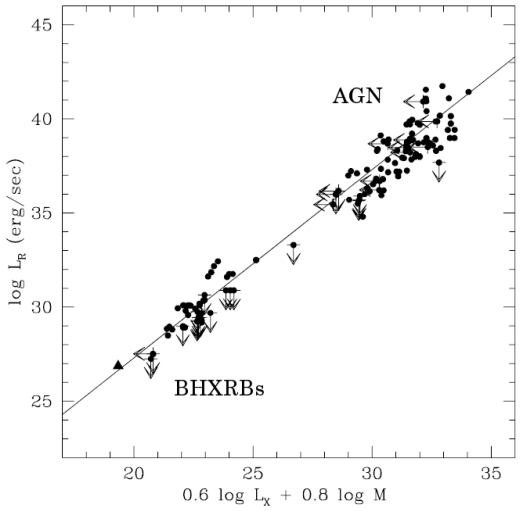

Figure 2.14: The fundamental plane of black hole activity. As seen in figure 3 of Gallo *et al.* 2006 with the addition of A0620-00 in quiescence (lower left triangle). This extended the existing slope of log L_R = (0.6 $\pm$ 0.1)log L_X + (0.8 $\pm$ 0.1) log M + (7.3 $\pm$ 4.1) (shown) from Merloni, Heinz & Di Matteo 2003.

activity' (see Figure 2.14) and strongly supports the notion of common accretion and ejection processes for black holes on all scales, and thus the scaling of XRB properties to the level of their supermassive counterparts.

2.4 Thesis overview

In the first chapter I discussed basic radio observation history and theory that will be required throughout this thesis in interpretation of data analysis methods. In this chapter I have outlined many of the basic astrophysical concepts associated with X-ray binaries and their radio emission, which will be of particular importance in understanding the results and impact of science in this thesis. In Chapter 3 I will briefly return to discussing radio observation theory, and the study of the effects of source spectral changes on upgraded telescopes such as ATCA-CABB. In chapter 4 I cover the attempts to observe the two BHXRBs, GRO J1655-40 and XTE J1550-564, in quiescence with ATCA-CABB in an effort to further explore the lower luminosity region of the hard state radio/X-ray correlation. In chapter 5 I will present the results from ATCA-CABB observations of NSXRB Circinus X-1 over a complete orbit while it was in a historically faint state, including the behaviour around the observed periastron flare, and the evolution of the system's radio structure over time. In chapter 6 I will discuss the first successful mm wavelength detections of NSXRBs Circinus X-1 and Scorpius X-1 using ATCA-CABB, discussing their ra-

dio structure, variability and spectral behaviour at these wavelengths. In the case of Circinus X-1 I will also discuss how the group of observations at cm and mm wavelengths may support the idea of recent variability in the system's jet geometry. Finally, I will conclude this thesis in chapter 7 with a discussion of the scientific impact of the enclosed studies, and future prospects for research in each field.

It's one of the great tragedies of life - something always changes.

DR. GREGORY HOUSE, HOUSE

3

The effects of source spectral changes on radio imaging

3.1 Introduction

As I mentioned in Chapter 1, the hardware upgrades installed on ATCA significantly boost its observation capabilities (see Section 1.3.1), but can also have detrimental effects to imaging of sources that undergo significant spectral changes during an observation. This is very relevant to observations of X-ray binaries and their jets since, as I described in the last chapter, structures in these sources can rapidly shift from optically thick to thin regimes and vice versa as a result of source evolution and particle (re-)acceleration. In this chapter I will first explain these possible effects in terms of observational theory and data analysis methods before moving on to the process I went through to simulate plausible observation scenarios that include variable sources. I will present attempts to image such data, discuss the resulting image artefacts that may influence interpretations of source structure, and finally summarise the implications of the results. Note that the initial goals of this investigation were to introduce me to advanced radio data reduction concepts, and to make me fully aware of the signs that could indicate the effects in question (essential for

my subsequent research), not to create a solution that could account for the error and correct images. Much of this work was carried out prior to the reclassification of some of ATCA's antenna configurations, with configuration 3.0C now also known as 6.0C.

3.1.1 Theory

As described in Section 1.2, radio images are derived from the Fourier transform of the *uv* plane (Equation 1.2) which is gradually filled by the signals from antenna pairs, with each *uv* point's position dependent on both the projection of the antennae physical positions onto the sky (for each observable location - related to the signal phase), and the observation frequency. This latter point is particularly relevant for multi-frequency observations, as it allows for a greater amount of *uv* space to be filled every moment and thus is a major benefit of upgrading telescopes to expand their observable bandwidths (such as the ATCA-CABB). However, the position of measurements on the *uv* plane (specifically, distance from origin) also influences the size of sources that can be resolved. Though this is obviously a benefit in trying to observe sources with ever increasing resolution, it also sets up a scenario where a spectrally variable source can effectively change in size in images, depending on which sets of *uv* ranges the source is most strongly detected in. I will explain this phenomenon using an extreme case which will be based on the scenarios depicted in Figure 3.1: a series of model observed source spectra and 1d representations (or cross-sections) of *uv* plane and image plane responses.

Consider first an ideal hypothetical scenario in which is is possible to observe a source's entire spectrum with an infinite range of baseline lengths (row A of Figure 3.1). In such a scenario, a true point source will appear just as such in the resultant images (i.e. the Fourier transform of an infinite, even, flat line function is a delta function at the origin). In a more realistic situation (row B of Figure 3.1), one is forced to observe a specific bandwidth of frequencies (marked by ν_{min} to ν_{max}) and with a limited number of baselines (a single baseline in the case of Figure 3.1) which causes the *uv* coverage to be truncated and imaged point sources to display the typical wave-like form prior to cleaning. A truly realistic scenario would also have to deal with multiple *uv* rings and discrete frequency sampling, but the current complexity is sufficient for the purpose of explaining the effects studied in this chapter. Next, consider an extreme scenario in which the source's spectrum resembles a step function; that is, below a certain frequency, ν_0, it cannot be detected, but behaves just like the previous scenario at frequencies higher than this (row C of Figure 3.1).

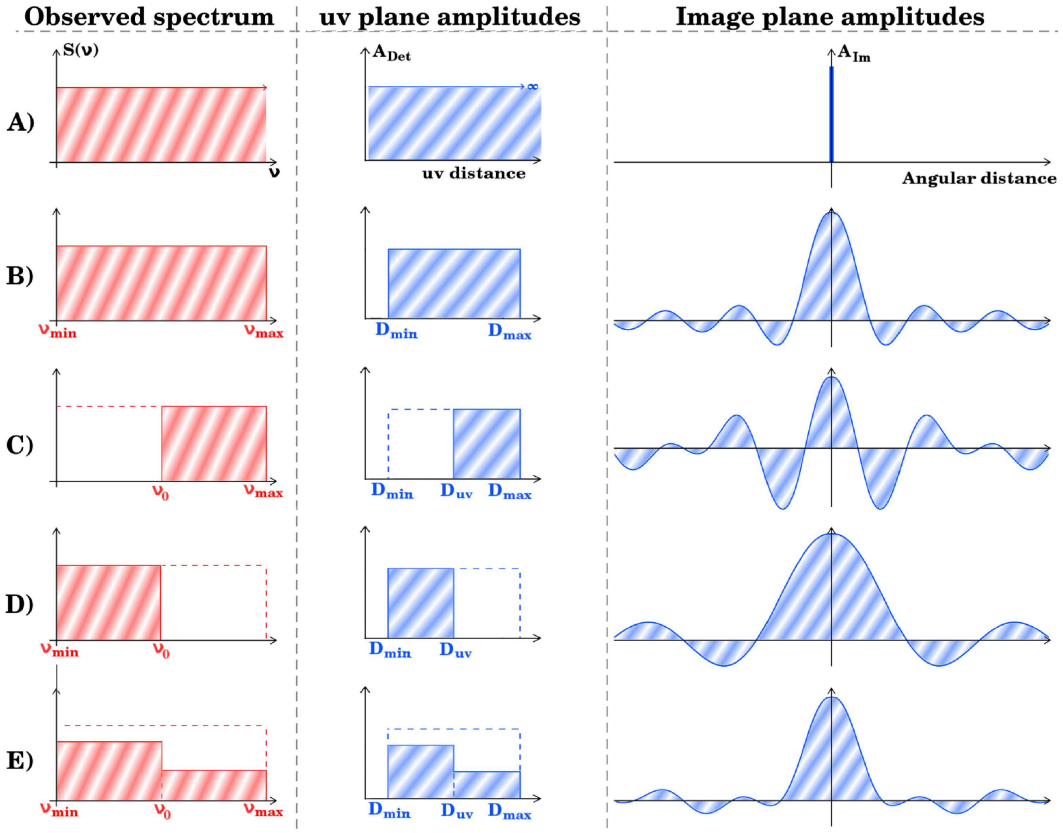

Figure 3.1: Illustration of the effects of spectral shape on *uv* plane and image plane response. The left hand plots show a model source's observed spectrum that defines each situation: A, a flat spectrum source observed by a hypothetical instrument with infinite frequency coverage; B, a flat spectrum source observed over a particular bandwidth (ν_{min} to ν_{max}); C, a source with a flat spectrum at high frequencies but a sharp spectral cut-off to 'undetectable' levels below frequency ν_0; D, a source with a flat spectrum at low frequencies but a sharp spectral cut-off to 'undetectable' levels above frequency ν_0 (i.e. the reverse of C); E, a source that varies between two spectral states (akin to C and D) resulting in a 'blended' spectrum. Middle plots show a representation of a cross section of a single baseline's (broad bandwidth) response to the source in the *uv* plane, with the horizontal axis marking the distance from the centre of the plane (assume rotational symmetry in the *uv* coverage), and the vertical axis being an arbitrary measure of the source's detected amplitude. The right hand plots show the Fourier transform of the *uv* plane response, which is equivalent to a cross-section of the image of the source prior to cleaning, or the dirty beam, with the horizontal axis being angular distance in the image plane, and the vertical an arbitrary measure of flux density. The horizontal scale of the plots remains the same in each scenario (the Fourier transforms are real calculated plots), though the vertical axis and relative scale of each column are arbitrary.

The radial distance of a *uv* measurement is proportional to wavelength; thus, the loss of low frequencies in the source's spectrum limits *uv* coverage to the largest *uv* values and higher observation resolution, or in other words, decreases the width of the primary peak in our image Fourier transform. In the reverse case, where instead the spectrum loses the higher frequencies (row D of Figure 3.1), we observe the opposite effects (as one would expect): that is, a widening of peak width and lower resolution.

The major concern now lies in attempts to clean such sources, or fit models to them in an effort to determine attributes or their resolved extent, both of which depend on use of the image plane point spread function (PSF), or 'dirty beam'. A basic dirty beam will be derived from the full normalised *uv* coverage available, and so in our scenarios will match the middle plot of situation B (remember that the dirty beam is simply a Fourier transform of the instrument's response to an ideal point source centred in the observed field: the equivalent of the plots in the right column of Figure 3.1). When used to fit sources like C and D, the beam will not match up perfectly, which aside from complicating the cleaning process, could also lead to the incorrect interpretation of a resolved source (in the case of situation C) or inaccurate measurements. Obviously, this issue should arise in any image that includes sources with different spectral shapes, even if none of them vary over time. However, as described in Section 1.2.1, multi-frequency synthesis (MFS) and clean (MFCLEAN) methods account for this scenario by including an additional term in the beam model that models the response of a source with a linearly variable spectrum (spectral dirty beam: equations 1.4 and 1.5). Though this basic function would be unable to fit some of the examples shown in Figure 3.1 perfectly, the model should be sufficient for the majority of real radio sources in which power law spectra are common (see previous chapter). The effect of a more gently varying spectrum (such as a typical power law) on the *uv* plane visibilities of a model instrument (broadband with 3 baselines) is crudely represented by Figure 3.2. Instead of distinct cut offs at specific *uv* distances, the baselines each respond in a similar manner: with a gradual rise in measured flux density towards higher frequencies/*uv* distances for a positive α spectrum source, or lower frequencies/*uv* distances for a negative α spectrum source.

Unfortunately, the spectral dirty beam (SDB) cannot account for sources whose spectra vary over time. In a situation where a source varies between situations C and D, the SDB will likely fit a model that resembles a blend of the two spectra. An example of this possibility is shown in row E of Figure 3.1, but will vary depending on the relative amount of time the source spends in each state. Such a fit will work

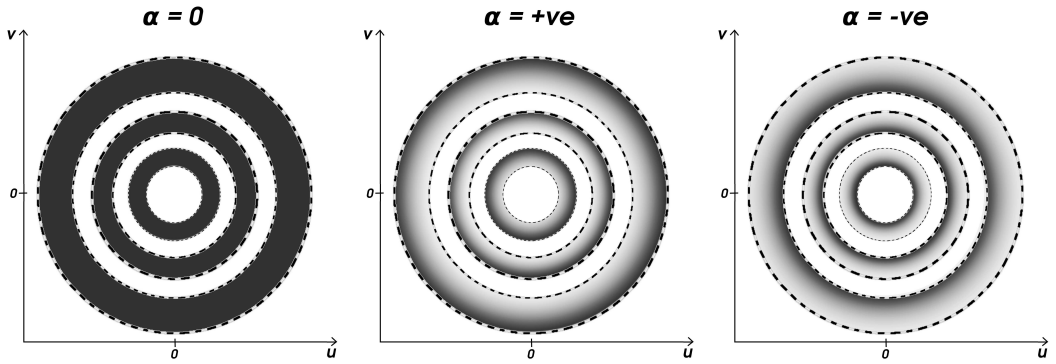

Figure 3.2: Illustration of the effects of sloped spectra on *uv* plane visibilities. These plots show the *uv* coverage of three baselines with broad bandwidths observing a single source for 12 hours (note that the plots are idealised for a source at $\delta = 90°$; i.e. circular tracks). The grey scale indicates the relative intensity of the source at each *uv* point, with a flat ($\alpha = 0$) spectrum (left) giving uniform detection, a positive ($\alpha > 0$) spectrum (middle) being brightest towards the higher frequency extreme of each baseline's coverage ring (outer edges), and a negative ($\alpha < 0$) spectrum (right) being brightest towards the lower frequency extreme of each baseline's coverage ring (inner edges). This in turn affects the dimensions of the source image, as discussed in section 3.1.1.

to some degree during periods of transition (a nearly flat spectrum as in B, which is much closer to E's model), but obviously continues to differ significantly from the two extremes. In an image, the cross section of source's shape across a particular hour angle (an axis defined by the projection of the baseline at each observation time) will approximately resemble the Fourier transform of the source's *uv* plane response at that time; thus, a variable spectrum source will appear deformed. If one attempts to subtract a point source modelled on the spectral dirty beam (symmetrical), the process may leave behind regions of over and under-subtraction around the source's location. Given that these should appear in roughly symmetrical locations about the source position, it is easy to imagine someone confusing the structure for something like a jet ejection.

It should be noted that this effect is similar to that which can arise from simple flux density variability in monochromatic observations. In such a scenario, the flux density associated with a dirty beam fitted to a variable source will not correctly match the source's flux density at all hour angles. Therefore, just as described above, there will be regions affected by over and under subtraction during the cleaning process, but as a result of incorrect amplitudes of maxima and minima rather than shifts between differing profile shapes.

3.2 Methodology

A logical path to testing the severity of these effects is to produce model data that include a source whose spectrum varies over time, and then to apply existing data reduction and analysis techniques. While MIRIAD does include a robust task capable of generating model *uv* visibilities called UVGEN, it unfortunately lacks the option of including a source with any parameter that varies over time (be it monochromatic luminosity or a complete spectrum). This can be overcome by generating multiple data sets with different source models and timing properties that can be combined in sequence using the task UVAVER which, given fine temporal and frequency binning, can closely resemble a variable source data set. The downside is that this requires a very large number of model data sets to combine; for example, a realistic representation of a 12 hour observation would require $2048 \times 4320 \approx 9$ million individual models (number of frequency channels $\times$ number of 10 second bins) that must first be combined in frequency, and then time. Suffice to say, even a reduced number of time and frequency parameters would still take a significant amount of computing time to prepare. Therefore, I prepared the python script named VASPECSIM (Appendix A) which prompts the user for input (see a prepared list in Appendix B) to produce a shell script of MIRIAD commands that can be run to produce the desired single complete *uv* visibility set. A few notes on VASPECSIM:

- Though VASPECSIM may only take a few minutes to run and prepare the script, the MIRIAD tasks required can cause the actual data generation to last hours, depending on the complexity of the models involved.
- VASPECSIM can include both point (i.e. a Gaussian restricted to the dimensions and position angle of the PSF) and Gaussian sources; however, the former is most useful in determining the severity of spectral variability effects.
- VASPECSIM can simulate spectral shapes that include both power laws and Gaussian profiles, though the former is most relevant for modelling of X-ray binaries.

Additionally, for the first tests, the following constraints were decided upon (with explanations) to maximise useful results:

- Initial parameters for UVGEN should be such that the model data sets resemble real 12 hour ATCA-CABB observations as closely as possible, but computing time is a factor and some simplification is necessary. Parameters 2-15 of

Appendix B are the same for all the included models. The main simplifications are an increase in frequency channel spacing from 1 MHz to 4 MHz, and the spectrum changes every 36 minutes (for a total of 20 spectra in a 12 hour model) rather than gradually. The latter constraint should still allow for reasonably accurate results, since even if the difference between the theoretical model SDB and the source is more significant in an early half of a time segment than it should be (due to the spectrum being stable for 36 minutes), this is offset by the opposite effect in the last half (the source is closer to the model than it would be for a gradual variation).

- Antenna array configuration 3.0C is used as it includes well spaced *uv* baseline rings ranging from 0.15 km to the ATCA maximum of 6 km. This is also an array configuration similar to those used in observing XRBs, where high resolution is important.

- Though noise can be added to images, it is better ignored in order to test the simplest effects of spectral variability.

- Initial tests need only be carried out on a single point source with a linear spectrum, as this closely represents observations of XRB 'cores' (central point source of a system) and will also match assumptions made in the calculation of the SDB.

The only parameters left to decide on are the range of spectral slopes to test and how they should vary. A reasonable range of radio emission mechanisms are covered by power laws in the range of -2 $\leq \alpha \leq$ 2, which for a 2 GHz bandwidth about the mean frequency of 5.5 GHZ (one of the standard ATCA-CABB observation frequencies) is approximately equal to linear slopes ($dS_\nu/d\nu$, no log values) ranging from $\nabla = \pm$ 3.5. Thus, to simplify the modelling process further, I can apply linear spectra in the source (set the exponent, B, during parameter input to 1: see code in Appendix A). Two types of variability should be tested, one in which observation length is shorter than variability timescales (i.e. a gradual change from one spectrum to the another), and a second in which the reverse is true (i.e. the source quickly switches to another spectrum and back again). These two scenarios should also display two different types of structural deformation in images. An important parameter to keep constant while varying the spectrum is the overall integrated flux density over the bandwidth, since, as mentioned earlier, any variability in overall luminosity will have its own effect on the source image (studied further in Section 3.4). The level of average flux density should be realistic, but given that noise is

ignored, the value I choose should not adversely affect the tests and so I can set the test source to have an average flux of 10 Jy. Once an initial set of tests is complete, I can then choose a single model to duplicate with new average flux levels in order to see how the effects scale with the parameter (Section 3.3.4). Table C.1 (Appendix C) lists the complete set of test model parameters I start with.

3.3 Results

For each model I produced an image, cleaned with MFCLEAN (200 iterations), subtracted a point source from the source location using IMFIT (subtracted point source flux densities are listed in column 5 of Table C.1) and produced a new image of the residuals. There are distinct properties that are similar across all resultant model images: firstly, the source appears identical (point like) in all pre-subtraction images so these need not be presented, though residuals do differ between sets of models and thus are included in subsequent figures; secondly, in all cases the flux of the resulting point source fit is $\approx$ 10 Jy with only 8 showing a $\approx$ 0.1% variation from this value. Minimum and maximum residual flux density measurements are listed in columns 7 and 8 of Table C.1 and will be referred to in the following subsections.

3.3.1 Stable models

In order to understand the effects of the variable models it is useful to have benchmarks of stable sources with each spectral index limit (∇ = -3.5, -2, -1, 0, 1, 2, 3.5) so that I can see the response of MFCLEAN to non-variable spectra. The residual images from such models are shown in Figure 3.3. Unsurprisingly, the most successful cleaning occurs in the model with ∇ = 0, with residual flux densities not exceeding 0.04 mJy (4×10^{-4} % of the source flux) and structure resembling diffuse sidelobe components. For all other values of ∇, sidelobe residuals are far more pronounced, appearing as ring structures around the point source, with their intensity proportional to the steepness of the source spectrum ($\sim$ 0.01% for $|\nabla|$ = 1, $\sim$ 0.02% for $|\nabla|$ = 2 and $\sim$ 0.035% for $|\nabla|$ = 3.5: See Table C.1). This already suggests MFCLEAN has more trouble modelling and subtracting sources with steeper spectral indices with the same number of iterations.

While it is likely possible to remove these completely with additional iterations of MFCLEAN, I am interested in the relative intensity of residuals in the varying models compared to these, and thus need not waste time in attempting to eliminate them entirely. In addition, it is important to realise that a trained eye should be ca-

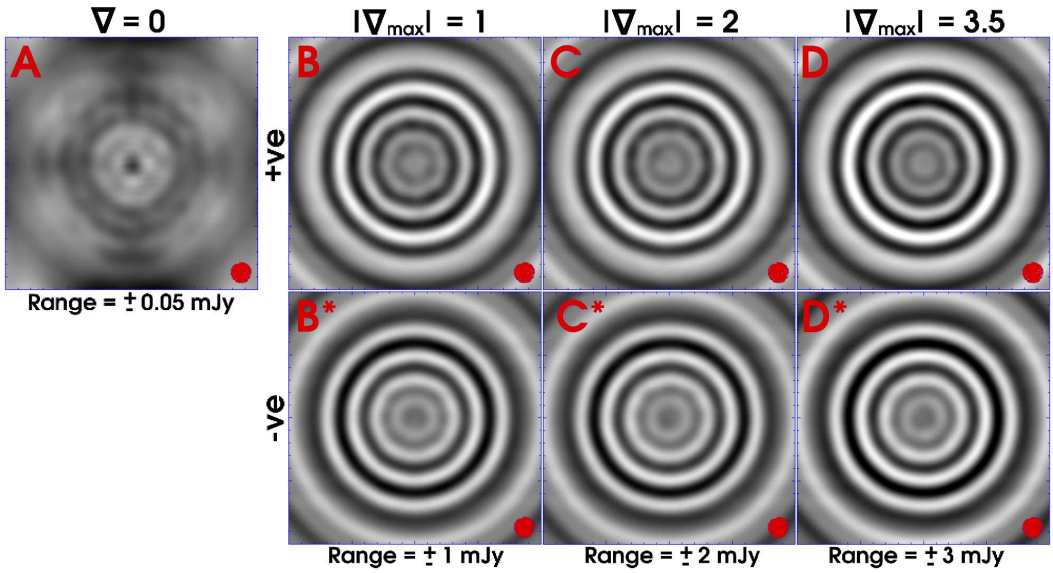

Figure 3.3: Stable model residual images. These maps show the results of attempting to subtract a fitted point source (using MIRIAD task IMFIT) from images made from generated data of a single point source (centred) whose linear spectrum ($S_\nu \propto \nu$) remains unchanged during a full 12 hour 'observation'. The ring structures are unsubtracted sidelobe residuals which the spectral dirty beam has been unable to match perfectly. Images have been cleaned with MFCLEAN for 200 iterations and flux density contrast range increases from left to right. Black corresponds to positive flux densities, white negative. See Table C.1 for image statistics (table reference matching top left labels).

pable of identifying such structures as sidelobe effects and adjust accordingly. Note that the difference in ring spacing between the upper and lower rows of Figure 3.3 resembles the effect illustrated in Figure 3.1, and the effects on the corresponding *uv* plane visibilities should be akin to that illustrated in Figure 3.2; albeit, with a larger number of baselines and different track widths (which in ATCA-CABB's case can cause some to overlap in the *uv* plane).

3.3.2 Gradual models

The residual images for the 'Gradual' models, i.e. those in which the source spectrum varies linearly over 12 hours between two values, are presented in Figure 3.4. The structure of the residuals has changed dramatically, with a number of distinct components forming instead of the ring like structures observed in Figure 3.3. The pattern formed is a natural result of the mismatch between subtraction model and what is actually in the map, and can be explained if we once again approach the image from the perspective of hour angles (HA). In these scenarios, the beginning and end of the data set (HA = -6 and +6 respectively) correspond to the north-south

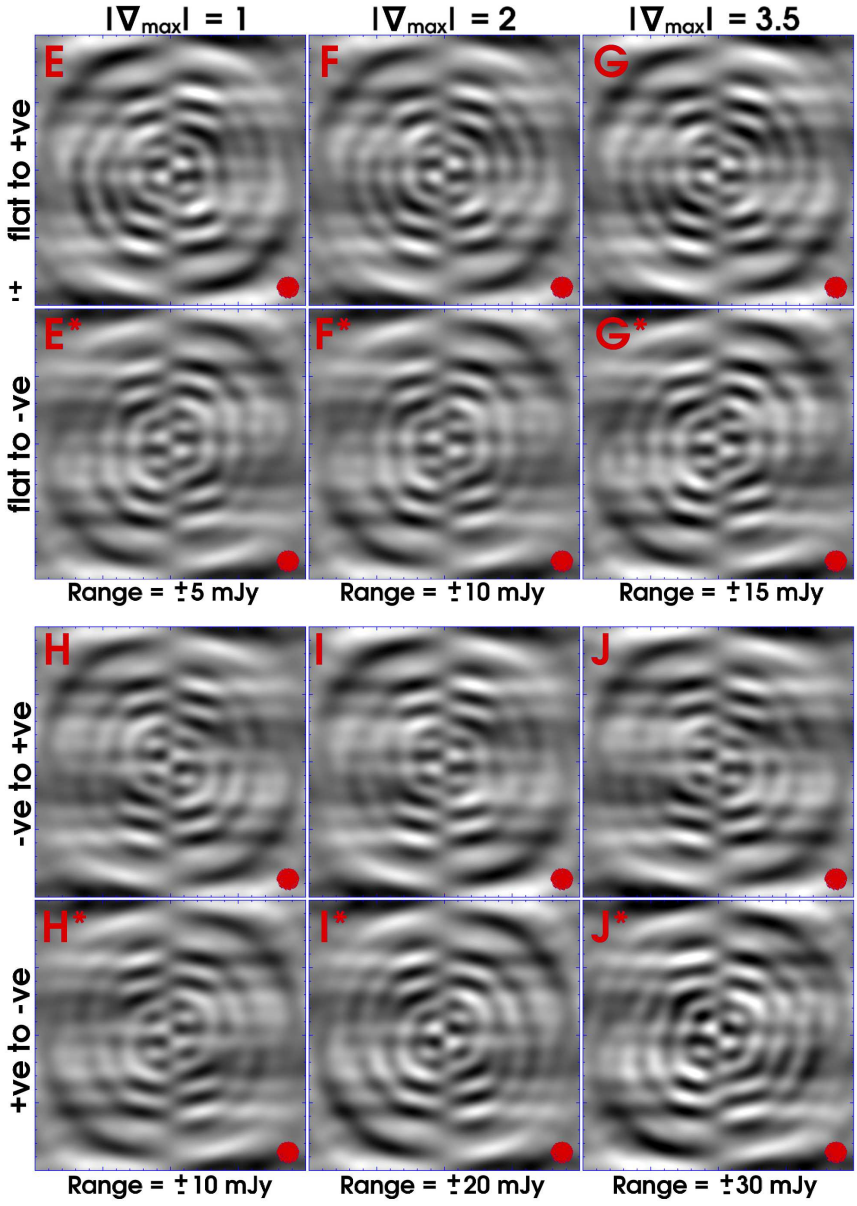

Figure 3.4: Gradual variation models residual images. The spectral index in these models varies linearly from one value to another (see left hand and top labels) over the 12 hours. The top set start with a flat spectrum, whereas the bottom set start from a non-flat spectrum. Black corresponds to positive flux densities, white negative. Images have been cleaned with MFCLEAN for 200 iterations and flux density contrast range increases from left to right. See Table C.1 for image statistics (table reference matching top left labels).

axis, with data 'filling in' along axes in an anti-clockwise direction as HA increases. Since the spectrum also changes with HA, the radii at which the sidelobes fall will gradually vary, contracting with rising ∇ (unstarred letter panels) or expanding with decreasing ∇ (starred letter panels): an effect which is most evident when comparing either side of the north-south axis (Ha =-6 one side, and +6 the other). Now, remember that the SDB (equivalent to the subtracted model) should resemble an average of the dirty beams corresponding to each individual spectral model used in a data set, and in this case should effectively be a halfway point between the two extremes of the spectral variation (since it is a linear change over the full observation). Thus, at the earliest and latest HAs (those closest to the north-south axis) the SDB matches least effectively to the source profile, and the sidelobes go unsubtracted, whereas towards the middle of observations where the SDB most closely resembles the source model (corresponding to the east-west axis in the images) the sidelobes are removed.

As expected, the brightness of these artefacts is higher in these images than in those made using stable source models, with the structures approaching 0.4% of the source's flux density (a factor of ~ 10 increase compared to stable models) in the most extreme cases (J, J*: $|\nabla| = 3.5$). The intensity increases as the difference between start and end spectral gradients widens which is logical, given that, if the SDB does represent an average, then it is roughly the same for each group of models (e.g. approximately equivalent to the SDB for $\nabla = 0$ for the cases of H, H*, I, I*, J, and J*) but the source spends time in increasingly steeper (and thus increasingly different) spectral states as $|\nabla|$ rises; therefore, the sidelobe residual regions gradually become less effectively subtracted.

It is important to note that, though these structures are brighter than those seen in stable models and no longer resemble typical sidelobes, they also do not immediately resemble physical structures observed in real sources. However, like all sidelobes, it remains possible that such artefacts arising from multiple closely clustered sources could produce brighter, distinct structures due to overlap. Nonetheless, there is an important effect which can help an observer to identify these artefacts, and that is the fact that these structures are paired with similarly intense negative regions near to them (the white regions of the maps: minima listed in column 7 of Table C.1). The SDB may fail at subtracting the positive sidelobe regions, but it also fails at eliminating the equivalent minima in the Fourier transform, or even oversubtracts/intensifies certain areas. In addition, there may be a possible pairing with similarly patterned structures on the opposite side of a source.

Table 3.1: Results of model repetitions using different source flux densities. Duplicating models J and P (10 Jy) with new source S_ν of 1 Jy and 1 mJy.

Model name	Source S_ν (Jy)	Residuals min S_ν (mJy)	Residuals max S_ν (mJy)
Gradual-3.5to3.5	10	-32.30	32.31
Gradual-3.5to3.5	1	-3.22	3.24
Gradual-3.5to3.5	1×10^{-3}	3.12×10^{-3}	3.32×10^{-3}
Burst-3.5to3.5	10	-67.26	76.66
Burst-3.5to3.5	1	-6.96	7.68
Burst-3.5to3.5	1×10^{-3}	-6.98×10^{-3}	7.60×10^{-3}

3.3.3 Burst models

The residual images for the 'Burst' models, i.e. those in which the source spectrum rapidly varies from one value to another and back during the 12 hours (spending more time at one ∇ than the other), are presented in Figure 3.5. Again we see the effects of badly subtracted sidelobes; however, they are now more concentrated towards a specific set of angles (east-west), reflecting the brief variation of the source spectrum (HA $\sim$ -2 to 2). These structures are brighter than in Gradual models by a factor of $\sim$ 2 ($\sim$ 0.2% for $|\nabla|$ = 1, $\sim$ 0.4% for $|\nabla|$ = 2 and $\sim$ 0.7% for $|\nabla|$ = 3.5). This is expected given that the source now spends more time in one spectral state than the other, shifting the SDB to resemble the former more closely, and thus resulting in a greater deviation from the source during the 'burst' period and a poorer subtraction of the sidelobes.

It is obvious then, that spectral variation artefacts are more intense when the source spends a briefer amount of time in a different spectral state to its 'normal' state. In addition, the brief switch concentrates the brightest artefacts towards a single axis which can quickly resemble the effects of jet ejections (paired structures either side of a central source). This is a far more dangerous scenario for observers, but might still be identified as false if care is taken, since the negative regions of the sidelobes also intensify.

3.3.4 Effect of source brightness

To test the intensity of artefacts resulting from sources of different flux densities, one can use some of the above models as references and vary the model spectra by a single factor, while retaining the overall spectral index (i.e. for a power law of Ax^B+C, multiply A and C by the same factor and leave the exponent B = 1). Model J had some of the most severe artefacts of the Gradual models, and P the

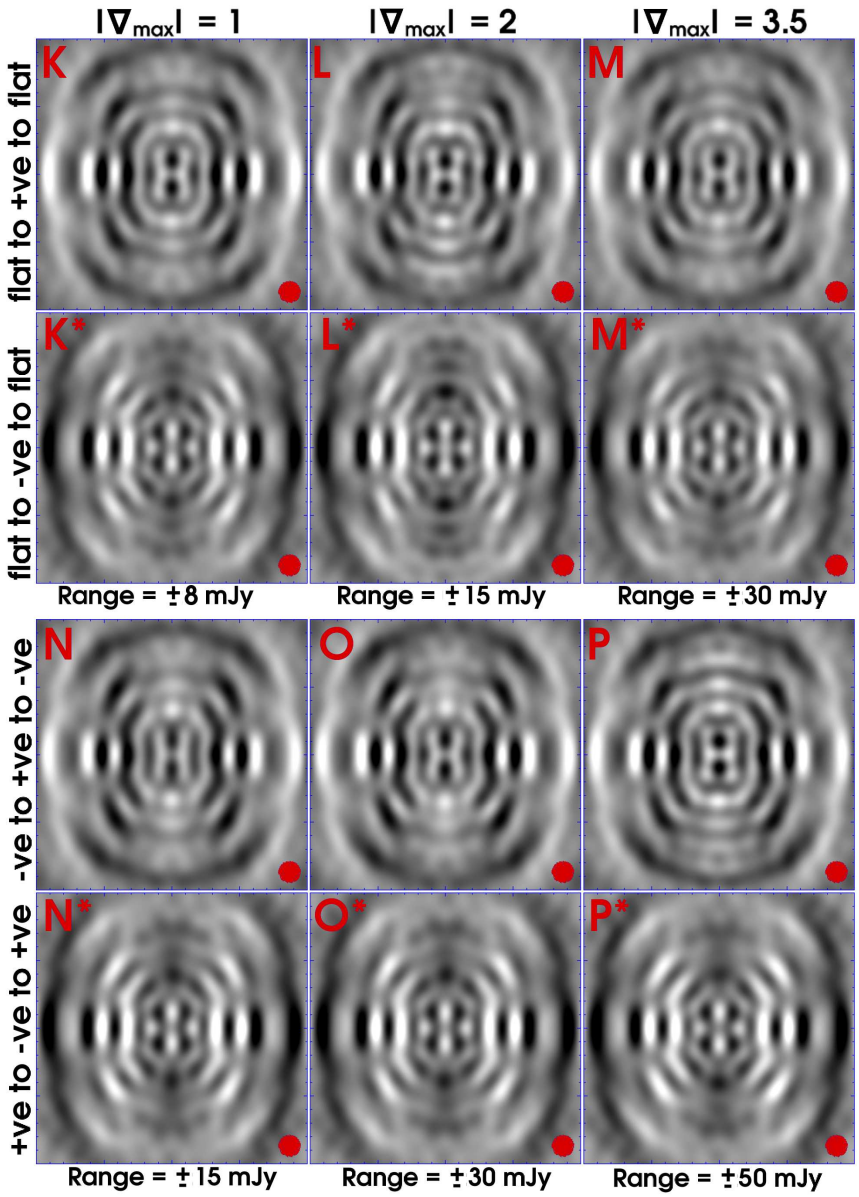

Figure 3.5: Burst variation models residual images. The spectral index in these models remains stable for the first 4.2 hours, changes sharply (over 1.2 hours) to a new value where it remains for 1.8 hours before rapidly returning to the original index (over another 1.2 hours) and stabilising for the remaining 3.6 hours. The top set starts with a flat spectrum, whereas the bottom set start from a non-flat spectrum. Black corresponds to positive flux densities, white negative. Images have been cleaned with MFCLEAN for 200 iterations and flux density contrast range increases from left to right. See Table C.1 for image statistics (table reference matching top left labels).

most severe of the Burst models, thus I chose to base the new tests on these, but with sources of 1Jy and 1mJy in place of 10 Jy. The measurements of images from these models plus the original are shown in Table 3.1. It is obvious that the intensity of the artefacts varies by the same approximate factor as the source, which is not unexpected given their origin.

3.4 A comparison with variable flux density artefacts

It is naïve to think that sources can undergo change in spectral profile without some change in flux density (see Chapter 2), and given how faint the above artefacts are in comparison to their root sources, it is important to test how they compare to the artefacts that arise from sources that vary in brightness rather than spectrum. In addition to the effect discussed in subsection 3.1.1, east-west arrays like ATCA may suffer additional complications as a result of decreased phase sensitivity along the north-west axis of the array. In such a case, the array cannot distinguish between a source that varies, and additional structure being observed along the north-south axis: the phases simply do no vary sufficiently between baseline pairs to resolve out additional sources close to each other. Thus, a source with variable flux will appear to change size much in the same way as described previously, but along axes perpendicular to the projection of the array onto the sky (i.e. perpendicular to the hour angle perspective described in the previous section).

For this case, I created models of flat spectrum sources that initially start at 10 Jy, but increase by a factor of 1.095, 1.95 and 10 over the 12 hours. The resulting residuals are shown in Figure 3.6, and their statistics in Table 3.2. The form of the artefacts is very different, showing an almost linear elongation along the east-west axis, though they still appear in with positive and negative components. However, the most important difference to note is that even in the smallest rise case, these artefacts are brighter than any of the spectral variation effects detailed in the previous section. The 1.095 $\times$ rise alone already yields artefacts with flux densities of 1% that of the subtracted source (compared to 0.77% in the most extreme case of spectral variance model P), with factors of 1.95 and 10 yielding even greater effects at 6.8% and 18.6% of source flux density, respectively.

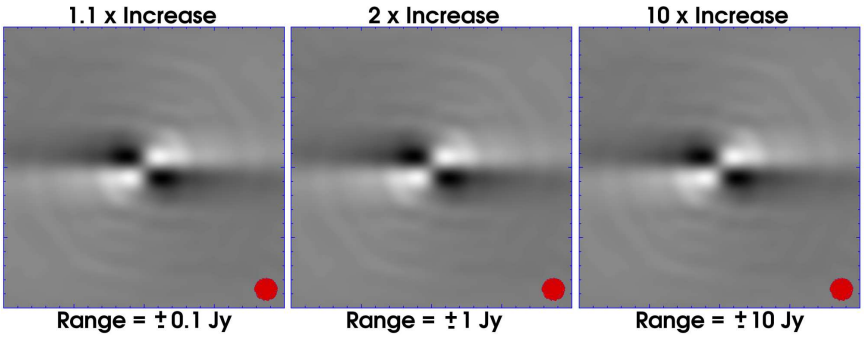

Figure 3.6: Flux density variation images. The spectral index in these models remains stable and flat but flux increases linearly over the 12 hours. Black corresponds to positive flux densities, white negative. Images have been cleaned with MFCLEAN for 200 iterations and flux density contrast range increases from left to right. See Table 3.2 for image statistics.

Table 3.2: Results of variable flux density models. Flat spectra, but linear increase in flux density over 12 hours.

S_ν increase factor ($S_{\nu end}/S_{\nu start}$)	Average Source S_ν (Jy)	Residuals min S_ν (Jy)	Residuals max S_ν (Jy)
1.095	10.47	-9.43×10^{-2}	9.98×10^{-2}
1.95	14.74	-9.46×10^{-1}	9.96×10^{-1}
10	53.09	-9.54	9.87

3.5 Conclusions

I have been able to show that spectral variance in a source during an observation can result in image artefacts, and whose intensity depend on both the behaviour of the source over time (the shorter the period over which the source deviates from 'normal' the more noticeable the effects), the difference between one spectral state and the next (the bigger the change in spectral slopes, the more intense the effects) and average source brightness (a bright source yields bright sidelobes). Having compared the relative brightness of these artefacts to their progenitor sources, I find that these effects can become visible in particularly bright sources, where flux densities approach hundreds of times the image noise (i.e. where $3\sigma \sim 1\%$ of source S_ν). However, having subsequently compared these artefacts to the well known effects of simple variation in source flux density, I find them to be of lesser concern to the latter flux density variability effects. The only regime in which spectral variability artefacts may become an equal concern, is in images of sources where flux density does not vary by more than $\sim 10\%$, and even in these circumstances the spectrum must show significant changes ($\Delta\alpha > 2$) over short periods of time. In both cases

of variability, the observed artefacts share the common trait of appearing as pairs of positive and negative components with distinct symmetry about the position of the variable source. These properties can be very useful indicators of such effects when interpreting images, and may be of use in determining whether structure is real or not.

My investigation has only scratched the surface of the problem, with a large region of parameter space not having been explored and many variables untested which could have significant impacts on the above limitations. In particular, the antenna configuration I used was intentionally selected because it included a balanced range of baseline lengths, but there are many alternatives in which baselines tend to one extreme or the other, and scenarios in which very long or short baselines might be eliminated from a data set entirely, to better observe certain types of structure. The effects will likely be of biggest concern for data sets of predominantly short baselines, since at short *uv* distances a small change in *uv* coverage has a much more significant effect on the Fourier transform's main peak width than the equivalent variation at large *uv* distances. One can crudely estimate the degree to which these effects might affect different antenna configurations by comparing the range of baseline lengths in a particular array configuration, to the range of observable frequencies in a band (which define the extent of the baselines on the *uv* plane). The latter is ≈ 1.4 for my tested band, and approaches ~ 1 for higher frequencies, the former is of the order of 100s for the 3.0 configurations which for my models gives baseline ratio/band ratio of ~ 0.01, supporting the order 1% flux density artefacts seen in the above results; however, compact configurations (such as the 0.122 configuration) can give baseline ratios of order 10 or lower (assuming one ignores antenna 6 which supplies the 6km baselines) which implies artefacts might exceed 10% flux density in such scenarios.

Be aware that the effects described in Section 3.4 are particularly important for the remainder of this thesis, as they appear in source images displayed and analysed in chapters 5 and 6.

The most exciting phrase to hear in science, the one that heralds the most discoveries, is not “Eureka!” but “That’s funny...”

ISAAC ASIMOV (1920 – 1992)

4

Observations of BHXRBs GRO J1655-40 & XTE J1550-564 in quiescence

In this chapter I present the results of radio observations of the BHXRBs GRO J1655-40 and XTE J1550-564 in quiescence, with the upgraded Australia Telescope Compact Array. The goal was to explore the sparsely populated low-luminosity region of the hard state BHXRB radio/X-ray correlation (Section 2.3.6) by exploiting the, then recent, improvements to ATCA. Though my analysis shows both sources are undetected, the radio flux density upper limits that I calculate from the images still allow me to reveal a significant amount of detail about the correlation’s low luminosity region, since the limits fall below the current track. In conjunction with quasi-simultaneous *Chandra* X-ray observations (in the case of GRO J1655-40) and Faulkes Telescope optical observations (XTE J1550-564) I found that these systems provided the first evidence of relatively ‘radio quiet’ black hole binaries at low luminosities, indicating that the scatter observed in the hard state X-ray/radio correlation at higher luminosities may also extend towards quiescent levels, and possibly even suggest that multiple tracks exist in the correlation. While both the radio and X-ray data reductions were performed by myself, the optical observation results of XTE

J1550-564 and Figure 4.2 were provided by Dr. David Russell. Additionally, my supervisor Prof. Rob Fender was kind enough to produce Figure 4.3 using existing data points in his possession, which I went on to modify for use in this project. This work was published in Calvelo *et al.* 2010a.

4.1 Introduction

As described in section 2.3, observations of accreting black holes (BHs) over the past decades have revealed the existence of spectral and temporal correlations between emission at different wavelengths, both in X-ray binaries, and active galactic nuclei. The correlations highlight the association of accretion onto a compact object with an outflowing jet, providing an important method of energy transfer from compact sources to the surrounding environment. For reviews of the topic see Belloni (2007), Done, Gierliński & Kubota (2007), and Markoff (2010).

Generally, when speaking of accretion and outflow activity in BH systems one is referring to emission in the X-ray and radio wavebands, with the radio component arising from synchrotron emission within a well collimated jet (section 2.2.3) and X-rays originating from various possible sources, including the jets themselves, the hot inner disc and a Comptonising corona (Markoff, Falcke & Fender 2001; Remillard & McClintock 2006a). BHXRB states are defined by the behaviour of the observed X-ray spectrum (section 2.3.3, also well described in Remillard & McClintock 2006b) but are also associated with significant changes in the observed radio emission. Typically, a source in the hard state exhibits steady, flat-spectrum radio emission (e.g. Fender 2001), whereas a soft state source shows little or no radio emission (Fender *et al.* 1999).

Understanding of the relationship between the X-ray and radio regimes of black hole systems has significantly improved in recent years due to a number of works which combine observations from both bands, or present the results of simultaneous (important for BHXRBs) multi-wavelength campaigns (Hannikainen *et al.* 1998; Corbel *et al.* 2003; Gallo, Fender & Pooley 2003; Falcke, Körding & Markoff 2004; Merloni, Heinz & di Matteo 2003, henceforth Ha98, Co03, GFP03, FKM04 and MHdM03 respectively). Correlations have been observed between the radio and X-ray luminosities of hard state systems (section 2.3.6). Originally noted by Ha98 in multi-wavelength observations of GX 339-4 (later confirmed in Co03), the relationship was explored further by GFP03 investigating multiple binary sources. The same correlation was discovered to hold for another BHXRB, V404 Cyg, as

well as a combination of points from additional sources (each additional source alone provided too few data points to reveal a relationship) indicating a correlation over 3 orders of X-ray magnitude. The basic characteristics of the X-ray and radio emission appeared to be independent of BH mass, and it was speculated that similarities might exist between the behaviours of stellar mass BHs and their super-massive counterparts residing within galactic nuclei. This prompted further expansion of the correlation to include BHs of all sizes by MHdM03 and FKM04 and the discovery of the *fundamental plane of black hole activity*.

As with any correlation, knowledge of the extremes can help greatly in further refining a fit, and additionally, the lower extreme of this relationship corresponds to a source regime where the dominant power output is of the form of radiatively inefficient outflows i.e. jet dominated (Fender, Gallo & Jonker 2003, Körding, Fender & Migliari 2006). Indeed, it may be that all hard state sources are jet dominated (Gallo *et al.* 2005). The possibility that the correlation between radio and X-ray extends all the way down to systems at quiescent levels should not be entirely unexpected as the quiescent state is often described as merely the hard state but at minimal accretion rates (though further differences are being uncovered; Corbel, Körding & Kaaret 2008). A quiescent binary system would likely still keep a jet, though of such low luminosity it would be difficult to detect with all but the most powerful radio telescopes (μJy level flux; GFP03).

The investigation into this regime has already met with success in the observation of A0620-00 in quiescence by Gallo *et al.* (2006). The flux measurements allowed for the expansion of the fundamental plane by a full two orders of magnitude in both radio and X-ray luminosities, as well as refining the correlation gradient; $L_R \propto {L_X}^{0.58\pm0.16}$. V404 Cyg has also been observed in 'quiescence' (Gallo, Fender & Hynes 2005) although it appears to be considerably more luminous than other quiescent binaries ($\sim$ 10 - 10^4 times brighter than A0620-00 in both radio and X-rays). At the time of my investigation's development, the next best sources for observation were the low-mass X-ray binaries GRO J1655-40 and XTE J1550-564, based on distance and relative brightness (see GFP03).

4.1.1 GRO J1655-40

GRO J1655-40 was discovered in 1994 as it went into outburst (Zhang et al. 1994) observed by the Burst and Transient Source Experiment (BATSE) aboard the Compton Gamma Ray Observatory. Radio observations revealed apparent superluminal jets (Hjellming & Rupen 1995, Tingay *et al.* 1995), only the second time such a phe-

nomena had been observed from a binary source (there are now other examples). Analysis of optical observations of the system in quiescence yielded primary and secondary masses of 5.4 ± 0.3 and 1.45 ± 0.35 $M_{\odot}$, respectively (Beer & Podsiadlowski 2002), with the primary's mass exceeding the maximum limit for a neutron star, supporting a black hole classification (see Section 2.2.1). During this initial outburst the system was observed simultaneously at X-ray and radio wavelengths (Harmon *et al.* 1995) clearly showing relativistic ejection events following X-ray flares. The system has also recently been observed in X-rays during quiescence (Pszota *et al.* 2008) where power law fits gave a flux of 6×10^{-14} erg s^{-1} cm^{-2} (0.5 - 10 keV). GRO J1655-40 remained in quiescence during these observations, making it an ideal target for exploring the lower limits of the fundamental plane.

4.1.2 XTE J1550-564

XTE J1550-564 was discovered in September 1998 (Smith 1998) with the All-Sky Monitor (ASM) aboard the Rossi X-ray Timing Explorer (RXTE), eventually reaching 6.8 Crab as measured by the RXTE (Remillard *et al.* 1998). Subsequent observations by Orosz *et al.* (2002) found the compact object's mass to be $M \sim 10.1 \pm 1.5$ $M_{\odot}$: far greater than the stable neutron star limit. Radio observations of the 1998 outburst also showed evidence of relativistic jets, found to decelerate over time (Corbel *et al.* 2002). Interestingly, the jets were also seen in X-rays and extrapolation of radio fluxes suggested that much of the X-ray emission may come from the same relativistic electron population that produces the radio component: directly revealing the transfer of kinetic energy from the jets to accelerating particles towards TeV energies. XTE J1550-564 is also in its quiescent state (Corbel, Tomsick & Kaaret 2006), observed at its faintest in X-rays to date (2×10^{32} erg s^{-1} at 0.5 - 10 keV) and with a spectrum that can be adequately fitted with a power law. Like GRO J1655-40, XTE J1550-564's distance and luminosity make it an ideal candidate for expanding the fundamental plane.

4.2 Observations & data reduction

My goal was to determine the radio luminosity of the two black hole candidates using ATCA-CABB, (quasi-)simultaneously with an estimate of the X-ray flux. For GRO J1655-40 I was able to do this directly with a near-simultaneous *Chandra* observation. In the case of XTE J1550-564 I used contemporaneous optical observations (results provided by Dr. David Russell) to indirectly estimate the X-ray flux

via the X-ray:optical relations established in Russell et al. (2006) as well as use of past X-ray flux measurements from observations of the system in quiescence.

4.2.1 GRO J1655-40

4.2.1.1 Radio

Observations of GRO J1655-40 were carried out on 2009 Jun 07 using the upgraded ATCA-CABB in 6A configuration. PKS 1934-638 was used as the primary/amplitude calibrator and 1729-37 (PMN J1733-3722) as the secondary/phase calibrator. Observations began at 07:31:25 UT (with actual source observations from 07:49:15 UT) and ended at 19:00:55 UT with time on the source being approximately 32.8 ks, giving predicted RMS noise of 6 and 8 μJy (using the ATCA-CABB Observing Characteristics Calculator at `http://www.atnf.csiro.au/observers/docs/at_sens/`) for the 5.5 and 9 GHz bands respectively (both with full 2 GHz CABB bandwidths). There was some radio frequency interference (RFI) evident in the second quarter of data at 5.5 GHz data and first quarter of 9 GHz data which was thoroughly flagged prior to image production. Inversion was straightforward and cleaning was carried out using a combination of MFCLEAN and original CLEAN subroutines. All data and image processing was carried out in MIRIAD.

4.2.1.2 X-ray

Chandra observations (using the ACIS-S detector in very faint mode) of GRO J1655-40 took place on the 2009 Jun 08 02:27:18 (UT) for $\sim$20.7 ks ($\sim$18.2 ks effective exposure) and ended at 08:11:42, 7.5 hours after ATCA-CABB observations finished, making them near-simultaneous. The analysis was carried out on the standard pipeline output level 2 data using CIAO version 4.1.2 (Fruscione et al. 2006). The data were also re-reduced manually in CIAO, with no significant difference to the pipeline output. No major background flares were detected. On source counts came to a total of 169, giving a net count rate of 9.3×10^{-3} counts s^{-1}.

4.2.2 XTE J1550-564

4.2.2.1 Radio

Observations of XTE J1550-564 took place on the 2009 Aug 04 (5.5 and 9 GHz) and 2009 Aug 06 (1.75 GHz) using the ATCA-CABB in 6D configuration. PKS 1934-638 and PKS 0823-500 were used as primary calibrators (PKS 0823-500 was

used where PKS 1934-638 data were unavailable or of bad quality) and PKS 1613-586 as the secondary calibrator. Observations on the 4th began at 02:45:00 UT (actual source observations began at 02:57:00 UT) and ended at 14:00:00 UT with approximately 33.8 ks on source. Observations on the 6th began at 02:02:30 UT, unfortunately during the run there were RFI related problems which required fixes to be made on-the-fly, resulting in observation interruptions and initial calibration errors. The observations ended at 13:59:00 UT with the final total on source time being $\sim$28.1 ks. Flagging, reduction and cleaning of the 5.5 and 9 GHz data were relatively straightforward; however the process for analysing the 1.75 GHz data was far more complicated. Approximately 75% of the data channels had to be removed because most of them displayed drops in amplitude to negligible levels. The remaining data underwent significant flagging to remove RFI spikes. The RMS in the 1.75 GHz images would therefore be significantly higher than initial predictions.

4.2.2.2 Optical

XTE J1550-564 was observed with the EM03 camera on the 2-m Faulkes Telescope South located at Siding Spring in Australia, using the SDSS i'-band filter on the nights of 2009 Aug 04, 05 and 06. Observation conditions were particularly good on the nights of the 4th and 6th; however, they varied somewhat on the 5th limiting the usefulness of resultant images. Twenty-six 200s integrations were made on the 4th and stacked (using IRAF) into groups of three to improve signal to noise. Another 12 were produced on the 5th; unfortunately, many proved unusable (see section 4.3.1). Finally, 6 more images were obtained on the 6th. The images on target were de-biased and flat-fielded using the Faulkes pipeline. Calibration was carried out using known i'-band magnitudes of stars within the field from table 1 of Sánchez-Fernández *et al.* (1999). Typical uncertainty in the magnitudes is stated to be $\sim \pm$ 0.01 magnitude.

4.3 Results

4.3.1 Flux measurements

Neither GRO J1655-40 nor XTE J1550-564 was detected in any of the ATCA-CABB radio observations. The 3σ upper limit for each band is listed in column six of Table 4.1. No extended structure previously detected in Corbel *et al.* (2002) can be seen in the 1.75 GHz images of XTE J1550-564.

Table 4.1: X-ray transient parameters, quiescent X-ray and radio flux densities. This table includes both predicted and new measurements from this work (bold values in column six). I have amended existing predictions from GFP03 using new distance estimates from Jonker & Nelemans 2004, more recent quiescent X-ray flux measurements, and expanded the list to include details of additional X-ray transients in quiescence.

Source	L_X (10^{32} erg s^{-1})	Distance (kpc)	X-ray F_ν(1 kpc) (10^{-6} Crab)	Predicted radio F_ν (μJy)	Measured "quiescent" F_ν (μJy(GHz))
A 0620-00	0.02^a-0.04^b (1,2)	1.2±0.4	1-5	13-30	51±7(8.5)
GRO J1655-40	0.2^c-3^d (1,3)	3.2±0.2	6-82	5-30	**<26(5.5)**
					<47(9)
XTE J1550-564	$\sim2^e$ (4)	5.3±2.3	∼70	∼10	**<1400(1.75)**
					<27(5.5)
					<47(9)
GRO J0422+32	0.08^d (5)	2.8±0.3	∼2	∼3	Unobserved
GS 2000+25	0.02^d (5)	2.7±0.7	∼0.5	∼1	Unobserved
.....	.....	.....	.....	.....	.....
GS 1009-45	$<0.12^d$ (6)	5.7±0.7	<3	<1	Unobserved
XTE J1118+480	$\sim0.035^c$ (7)	1.8±0.6	∼1	∼4	Unobserved
XTE J1859+226	0.14^f (8)	6.3±1.7	∼4.2	∼1	Unobserved
GS 2023+338 (V404 Cyg)	16^g (9)	$4.0^{+2.0}_{-1.2}$	∼400	∼58	350(1.4-8.4)

Energy ranges: [a] 0.4-2.4 keV;[b] 0.4-1.4 keV;[c] 0.3-7 keV;[d] 0.5-10 keV;[e] 0.5-7 keV;[f] 0.3-8 keV;[g] 1-10 keV.
References: (1) Kong *et al.* (2002); (2) Narayan, McClinktock & Yi (1996); (3) Asai *et al.* (1998); (4) Corbel, Tomsick & Kaaret (2006); (5) Garcia *et al.* (2001); (6) Hameury *et al.* (2003); (7) McClintock, Narayan & Rybicki (2004); (8) Tomsick *et al.* (2003); (9) Campana, Parmar & Stella (2001).
R.M.S. values for this work's radio limits (column 6) were measured using the CGCURS routine (mean value of multiple image regions). It is evident that the severe flagging of the 1.75 GHz observations limited the quality of the images in that band.

GRO J1655-40 is clearly detected by *Chandra*. I extracted spectra, binning the counts into rebinned channels with at least 15 photons each. I attempted to fit several model spectra to the data using XSPEC, the results of which are listed in Table 4.2. Assuming a power law fit (see Figure 4.1) and allowing N_H to vary results in a fitted value of $N_H = 2.0^{+1.0}_{-0.7} \times 10^{22}$ cm^{-2} and a corresponding unabsorbed 2-10 keV flux of $F_X = 9.04^{+2.0}_{-2.3} \times 10^{-14}$ erg cm^{-2} s^{-1}, which with the distance from Jonker & Nelemans (2004) (3.2±0.2 kpc) gives an X-ray luminosity of $L_X = 1.11^{+0.26}_{-0.30} \times 10^{32}$ erg s^{-1}. I also include power law fits with the Hydrogen column density fixed at previously calculated levels; Asai *et al.* (1998): $N_H < 0.3 \times 10^{21}$ cm^{-2}, and Kong *et al.* (2002): $N_H \sim 0.9 \times 10^{21}$ cm^{-2}. The fits yield slightly larger luminosity values of 1.2×10^{32} erg s^{-1} and 1.5×10^{32} erg s^{-1} respectively along with shallower power laws.

XTE J1550-564 is easily detected in the Faulkes optical images on all three nights, the light curves for which are shown in Figure 4.2. Little variability was detected in the source over the individual nights (the range of magnitudes is similar to the error on each magnitude); however, a noticeable drop in luminosity occurred on the third night (6th) of ~ 0.3 magnitudes. It is possible that this variation is linked to orbital modulation ($P \sim 1.5$ days, Orosz *et al.* 2002) or variation in accretion rate during quiescence. The magnitudes are consistent with the mean value obtained over 1.5 years of Faulkes Telescope monitoring (Lewis *et al.* 2008), implying the source was in quiescence at the time of observations.

Summarised optical results are listed in Table 4.3 and the light-curve for all three nights can be seen in Figure 4.2. Note that due to the variable conditions on the fifth it was only possible to obtain 3 useful magnitudes out of the total 12 images.

4.3.2 Correlations

Using the values measured in the 5.5 GHz radio band and the quasi-simultaneous X-ray/optical measurements, upper limits for both XTE J1550-564 and GRO J1655-40 can be plotted along with other current detections of hard state BHs in the L_R:L_X plane (Figure 4.3).

In the case of XTE J1550-564 I may extrapolate a predicted value for the X-ray luminosity using the relationship published in Russell et al. (2006). Within the paper it is shown that a correlation exists between the optical-infrared and X-ray band fluxes, approximated by the relation $L_{OIR} = 10^{13.1 \pm 0.6} L_X{}^{0.61 \pm 0.02}$ (see their Fig 1). Taking the mean optical flux value from the 2009 Aug 04, and subtracting the relative contribution of the secondary star in quiescence [$L_{CS} \approx 0.7 \pm 0.1(L_{OIR})$;

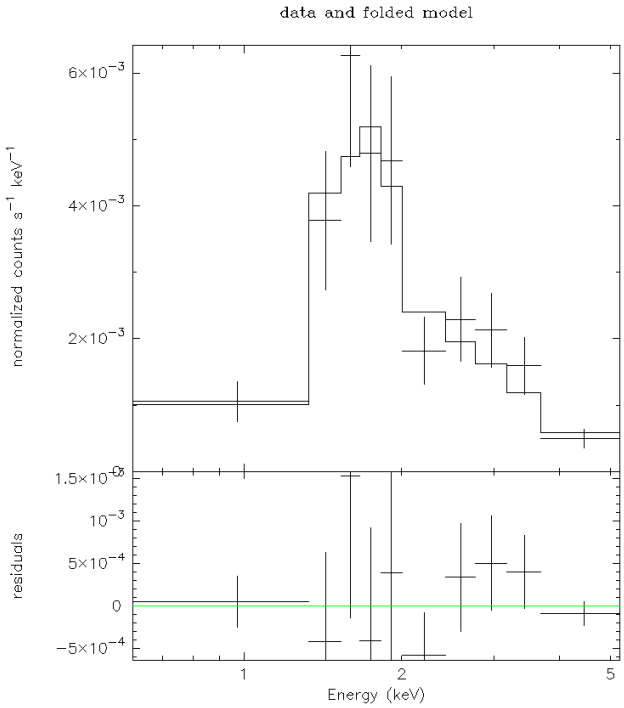

Figure 4.1: Binned *Chandra* data (0.6-8 keV) for GRO J1655-40: With fitted model (Absorption x Power law) and residuals.

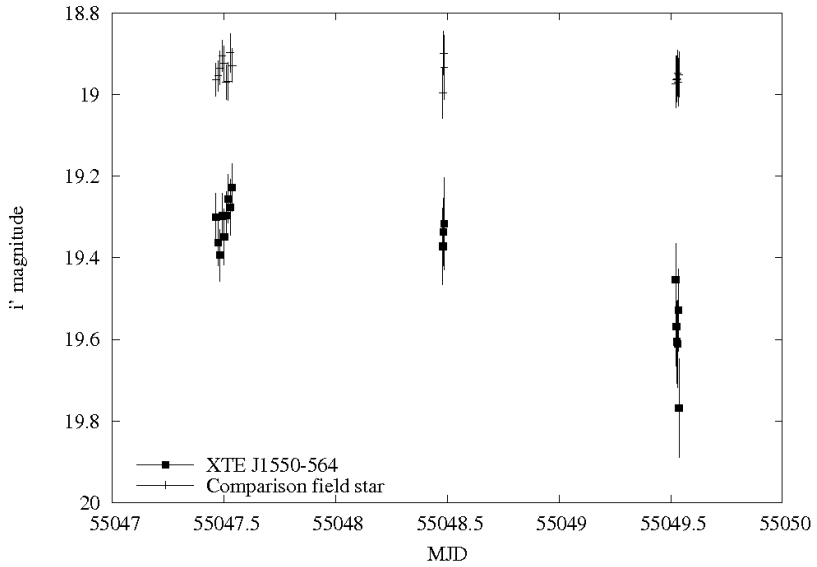

Figure 4.2: Optical light curve for XTE J1550-564: Covering the nights 2009 Aug 04-06. Included are the magnitudes of a (likely non-variable) field star slightly brighter than XTE J1550-564 (listed in Sánchez-Fernández *et al.* 1999). The XRB appears ~ 0.3 mag fainter on 2009 Aug 06 compared with 04 and 05.

Table 4.2: Best fit *Chandra* spectra parameters for GRO J1655-40. Included are the 90% confidence uncertainties and goodness of fit expressed by the reduced χ^2 parameter. Errors cannot be calculated for models whose reduced χ^2 exceed a value of 2.

Model	N_H (10^{22}cm^{-2})	Γ	kT (keV)	χ^2_r(d.o.f.)	Null hypothesis probability
Power Law	$2.0^{+1.0}_{-0.7}$	$3.1^{+1.0}_{-0.8}$	–	0.71(7)	0.67
Power Law	0.3(fixed)	0.9	–	4.5(7)	$<10^{-4}$
Power Law	0.9(fixed)	1.8	–	2.1(7)	0.4
Bremsstrahlung	$1.6^{+0.7}_{-0.5}$	–	$1.9^{+1.4}_{-0.7}$	0.63(7)	0.73
Black-body	$1.3^{+0.6}_{-0.4}$	–	$0.9^{+0.3}_{-0.2}$	0.59(7)	0.76

Table 4.3: Optical Results for XTE J1550-564.Though XTE J1550-564 remained relatively stable during each observation period, the slight drop on the night of the 6th is evident.

Date (y-m-d)	mean i' (mag)	1σ	mean error on each mag (mag)	de-red flux density (mJy)	1σ	mean flux error (mJy)
09-09-04	19.31	0.05	0.06	1.28	0.06	0.07
09-09-05	19.34	0.03	0.10	1.24	0.03	0.12
09-09-06	19.59	0.11	0.10	0.99	0.12	0.12

extrapolated from Orosz *et al.* 2002, figure 4] I can then apply the relationship to crudely estimate the 2-10 keV X-ray luminosity of 6×10^{33} erg s^{-1} at the time of observation.

Alternatively, I can combine this radio limit with recent quiescent X-ray observations; Corbel *et al.* (2006), under the assumption that the luminosity varies little during the quiescence. Corbel *et al.* (2006) supply a luminosity of 2×10^{32} erg s^{-1} (0.5-10 keV) which I convert for 2-10 keV using the mission simulator WebPIMMS (`http://heasarc.gsfc.nasa.gov/Tools/w3pimms.html`) to get a Luminosity of 9×10^{31} erg s^{-1}: significantly lower than my estimate of 6×10^{33} erg s^{-1}.

The process for GRO J1655-40 is straightforward in that the measured radio flux and X-ray luminosity can be converted and plotted directly on to the correlation graphs.

The points for both GRO J1655-40 (purple circles) and XTE J1550-564 (filled green triangles) lie at noticeably lower normalisations on the BHB plot than systems with similar X-ray luminosities ($L_X \leq 10^{34.5}$ erg s^{-1}) from the previous ensemble of measurements: up to a full order of magnitude lower radio luminosity (in the case

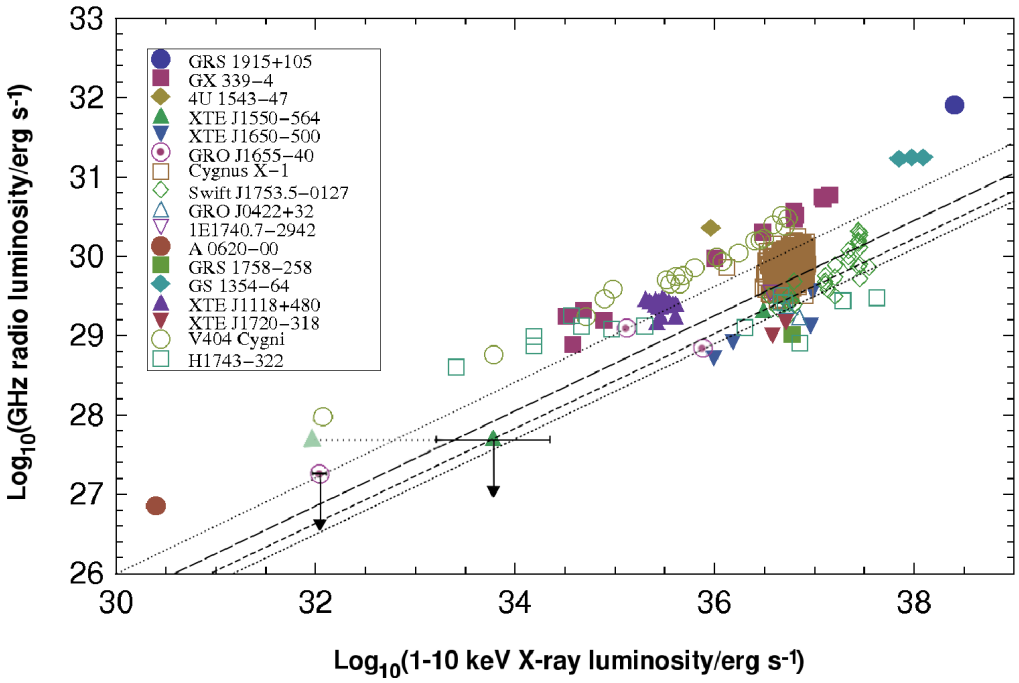

Figure 4.3: Radio/X-ray correlation plot for BHXRBs in the hard state: Additional data are taken from Fender, Gallo & Russell (2010). Radio luminosities are estimated by multiplying the 5 GHz monochromatic luminosity by the frequency (appropriate for a flat spectrum in the GHz band). Lines mark gradients of 0.6 to illustrate the possible link between my measurements and those made previously for GRO J1655-40 and XTE J1550-564. Dotted lines represent gradient extrapolations from the two previous measurements for GRO J1655-40, with the long dash line passing through an unmarked average of these two points. Finally, the medium dash line extends from the single measurement of XTE J1550-564.

of XTE J1550-564). It is interesting to note that the only measurements of XTE J1550-564 and GRO J1655-40 at higher luminosities are also towards the 'lower track', which, when taking into account these new values, could be consistent with a $L_R \propto L_X^{0.6}$ relation at a lower normalisation. The dashed lines marked on the plot illustrate this by extending the correlation gradient ($\sim$0.6) from the previous measurements for the two sources. The fainter green triangle trailing off to the left of the XTE J1550-564 estimate would be its position if I were to use the quiescent X-ray luminosity measured in Corbel *et al.* (2006). The result is a less dramatic scattering of points at lower luminosities.

4.3.3 Caveats

The best fitting power law model for my GRO J1655-40 data provides a photon index of $3.1^{+1}_{-0.8}$ which, given the large statistical errors, is broadly consistent with typical BH values ($\Gamma \approx 1.7 \pm 0.9$; Kong *et al.* 2002). Examining the results outlined

in Table 4.2 I might consider alternatives. The reduced χ^2 statistic suggests that both Bremsstrahlung and black-body disc models are better fits to the data, though the blackbody model is made unlikely by an estimated inner disc radius of < 0.2 km. Furthermore, in all models the column density value is higher than the majority of past references (only Sobczak *et al.* 2000 have such a high value: $N_H \approx 2.0 \times 10^{22}$ cm^{-2}). However, the fit becomes progressively worse as N_H is fixed lower, with the $N_H = 0.3 \times 10^{22}$ cm^{-2} being ruled out statistically.

I note that the small number of past observations of GRO J1655-40 in quiescence have yielded X-ray luminosities (all using D $= 3.2 \pm 0.2$ kpc) lower than those my power law model supplies; 3×10^{32} erg s^{-1} and 5.9×10^{31} erg s^{-1} for 0.5-10 keV in Asai *et al.* (1998) and Hameury et al (2003) respectively (my model gives 6×10^{32} erg s^{-1}) and 2.4×10^{31} erg s^{-1} for 0.3-7 keV in Kong *et al.* (2002) (my model gives 1.1×10^{33} erg s^{-1}).

Justification of the XTE J1550-564 result is more difficult since I am already dealing with a significant margin of error due to the OIR/X-ray extrapolation, and to a lesser extent the fraction of light contributed by the secondary star, which has not been measured well in i' band. In the case of the position of the XTE J1550-564 point when using the past quiescent X-ray flux, one must remain aware that measurements of quiescent flux do appear to vary over three years and the assumption of little change during quiescence may be an error: though evidence for similar luminosities from distinct quiescent periods (between outbursts) is seen.

4.4 Discussion and conclusions

Past observations have revealed that at relatively high luminosities there is considerable scatter in the 'universal' hard state radio:X-ray correlation, with a number of sources appearing to be 'radio quiet' compared to GX 339-4 and V404 Cyg (e.g. XTE J1650 – 500: Corbel *et al.* 2004, XTE J1720 – 318: Brocksopp *et al.* 2005, SWIFT J1753.5 – 0127: Cadolle-Bel *et al.* 2007). The upper limits on the quiescent radio luminosities of GRO J1655-40 and XTE J1550-564 presented here suggest that this range of normalisations for the correlation extends even to quiescent luminosities: had they followed an extrapolation of the GX 339-4 relation to A0620-00 they would certainly have been detected. Before this, the data could have been consistent with a narrow distribution of normalisations at low luminosities which would gradually broaden as the luminosity increased. In fact an explanation for exactly such a pattern of broadening distribution with luminosity) is put forward

by Soleri & Fender (2011) wherein an increasing bulk Lorentz factor with luminosity results in increasing beaming. With the inclusion of my data points, the scatter appears to exist all the way down to quiescent levels. In addition, GRO J1655-40 and XTE J1550-564 are both sources which have been observed previously to be rather 'radio quiet' in higher luminosity hard states. It is possible that these two sources may be slightly more 'radio quiet' than the majority, and as such lie along a separate parallel track.

As mentioned before, the scatter at lower luminosities is reduced if I use previous quiescent X-ray flux measurements instead of the value derived from the simultaneous optical observations. In this case XTE J1550-564 resides close to the higher luminosity track: yielding less evidence towards the 'parallel track' scenario. A scatter still exists with GRO J1655-40, but of lower magnitude in comparison to the higher luminosity region. The use of upper limits prevents one from seeing the full extent of the scatter, thus I cannot be sure if it truly decreases towards lower luminosities.

The low luminosity region of the BHB X-ray:radio correlation remains sparsely populated, as does the same region in the "fundamental plane of black hole activity". It is only with further observations of quiescent systems that we can continue to test not only the validity of GFP03 correlation, but also the possibility, as suggested above, of the correlation power law index being universal while normalisation can vary. Table 4.1 summarises the current information on radio detections and limits on black holes in quiescence. As well as tabulating my limits, and the detections of A 0620-00 and V404 Cyg, I refresh the predictions of GFP03, revising distances where appropriate, and adding new sources. Currently, the list of feasible targets at these levels is limited by the capabilities of available telescopes. However, with the completion of several instrument upgrades such as the CABB (e.g. E-VLA) and new arrays (LOFAR, ASKAP, MeerKAT etc.), the low luminosity region will become far more open to exploration.

4.5 Addendum

The above results were published prior to the release of both Coriat *et al.* (2011), in which the BHXRB H1743-322 is found to transition between different hard state radio/X-ray correlation tracks; and Gallo, Miller & Fender (2012), in which it is shown that the latest, most complete collection of hard state radio/X-ray correlation points is better fit by two populations rather than one (see section 2.3.6). As such,

it is prudent to discuss this work in light of these recent results.

From these latest projects we can only confirm two theories: that two distinct tracks exist in the hard-state radio/X-ray correlations, and sources can transition between the two; however, many questions also go unanswered. We do not know if all sources undergo such transitions, if those that do, do so in a similar region of the correlation plot as H1743-322, or even if such a source undergoes the same transition pattern repeatedly. While I am correct in postulating the existence of multiple tracks, the assumption that this track also follows the same gradient as its the radio loud counterpart is apparently incorrect. The radio quiet track of Gallo, Miller and Fender is steeper ($L_R \propto L_X^{0.98}$) which would make the non detections of XTE J1650 − 500 and GRO J1655-40 unsurprising, since if they follow the same trend they would be far too faint to detect with ATCA-CABB. At the very least, the non-detections do indicate that the lower track can extend much further than the rejoining branch marked by H1743-322, indicating that either some sources do not make a transition, or that they do so at different luminosities (lower in this case).

I repeat what I stated in the discussion: if we are to decipher the different behaviours observed in the radio/X-ray correlations of black holes, we must continue to explore the parameter space of the plot with an emphasis on faint quiescent sources. Not only are new instruments pushing the boundaries on what we can observe, but the population of known BHXRBs continues to grow, and with them, the number of available data points in the correlation. I will elaborate on such potential with recent examples in the final chapter.

I tell you: one must still have chaos within oneself, to give birth to a dancing star.

FRIEDRICH NIETZSCHE (1844 – 1900)

5

A complete orbit of Circinus X-1 at an historically faint epoch

In this chapter I will present the results from the first radio observations of a complete orbit ($\sim$ 17 days) of the neutron star X-ray binary Circinus X-1 (Cir X-1) using ATCA-CABB. These were taken while the system was in an historically faint state which, were it not for ATCA's upgrades, would have made the attempt to observe any persistent faint radio emission from the system a more difficult task. Furthermore, the enhanced sensitivity of ATCA-CABB improves the quality of short data segment imaging, allowing for the production of high detail light curves that I go on to make for Cir X-1. The observations were scheduled in such a way that the system was observed for 8 days in a quiescent 'pre-flare' state before capturing the rapid rise of a periastron passage flare and tracking its subsequent decline. I produce images from various data selections allowing me to study several aspects of the system's structure. Firstly, I tackle quasi-persistent structure in the system, its possible origins, and how it has changed since the last epochs of long term monitoring (prior to 2005), including the possible implications for the Circinus X-1's long term behaviour. I then use source modelling and subtraction techniques based on measured light curves to attempt to eliminate source variability artefacts (as discussed in Chapter 3) in the radio maps and thus reveal any non-static structure that

may arise following the flare event.

The study of variable structure in Circinus X-1 is uniquely interesting since previous observation programs have shown re-brightening of nearby structure following a flare event at the system's core (i.e. the location of the compact object) which was interpreted as re-energisation via unseen outflows with the fastest calculated jet velocities in our Galaxy ($\Gamma > 15$). All reduction and analysis of the ATCA radio data used in this chapter was carried out by myself. The Monitor of All-sky X-ray Image (MAXI) data was reduced by Motoki Nakajima, together with the team at the Tokyo Institute of Technology headed by Nobuyuki Kawai. This work was published in Calvelo *et al.* (2012a).

5.1 Introduction

Given the importance of understanding jet physics when dealing with environmental changes caused by XRBs and AGN alike, as well as its ties to accretion, it is essential that astronomers explore every possible avenue of investigation that may further our knowledge. NSXRBs share a great deal of similarities with BHXRBs, including accretion mechanics, birth scenarios and a degenerate compact object at the heart of the scene. Thus, NSXRBs represent an excellent subgroup of XRBs against which one can compare jet production that does involve a BH companion, the difference being that, it is possible to study the processes behind outflow production in the absence of complex physical attributes unique to BHs such as an ergosphere and event horizon. As a result, it might be possible to gain insight into how, and to what degree the presence of these exotic physical phenomena affect jet production, or if in fact the process is dominated by accretion mechanics alone.

The catch is that neutron star jets are often fainter in the radio than their BH counterparts, especially at lower X-ray luminosities (Section 2.3.6, see also Fender & Kuulkers 2001; Migliari & Fender 2006), and thus are often more difficult to study to the same level of detail. The emergence of upgraded radio telescopes with increased sensitivities, such as ATCA-CABB, not only allow us to study fainter BH sources but can allow us to better investigate the emission from NSXRB jets

5.2 Circinus X-1 background

Circinus X-1 (Cir X-1) is a peculiar X-ray binary (XRB) system known for its regular outbursts occurring every 16.6 days. These events are believed to be caused by

a highly eccentric orbit, wherein the periastron passage results in an increased level of accretion onto the compact object (Murdin *et al.* 1980; Nicolson, Glass & Feast 1980). The outbursts are visible at multiple wavelengths, including X-ray (Tennant, Fabian & Shafer 1986), IR (Glass 1978) and radio (Whelan *et al.* 1977). Discovered in 1971 (Margon *et al.* 1971), the system was initially classified as a black hole candidate due to X-ray spectral and timing similarities to Cygnus X-1, including millisecond variability (Toor 1977). However, a reclassification was necessary after the discovery of type I X-ray bursts (see Section 2.2.1) in 1985 which indicated a neutron star (NS) primary (Tennant, Fabian & Shafer 1986), though none had been seen again until May 2010, when the NS host was confirmed by RXTE detection of bursts (Linares *et al.* 2010). Though the system is often classified as a low mass XRB, the nature of the companion star remains under debate, with the possibility of it being a 3 - 5 $M_{\odot}$ sub-giant (Johnston, Fender & Wu 1999) or even a super-giant up to 10 $M_{\odot}$ (Jonker, Nelemans & Bassa 2007).

Cir X-1 also challenges normal NSXRB classification in that it shows X-ray behaviour reminiscent of both atoll (Oosterbroek *et al.* 1995) and Z-source classes (Shirey, Bradt & Levine 1999), as well as behaviour that defies either classification in orbital phases prior to periastron (Soleri *et al.* 2009a). This divide is likely related to variations in accretion rate, $\dot{m}$, onto the neutron star during its orbit, much like the different behaviours observed from XTE J1701-462 which have been linked to changes in $\dot{m}$ over time (Homan *et al.* 2007).

Low-frequency radio images reveal an extensive jet-powered nebula surrounding Cir X-1 (Stewart *et al.* 1993, Tudose *et al.* 2006). Early 'runaway binary' theories, which suggested the system was associated with nearby supernova remnant G321.9-0.3 (Clark, Parkinson & Caswell 1975) were supported by the existence of a tail-like structure extending from the southern end of Cir X-1's nebula. However, observations with the *Hubble Space Telescope* revealed little or no proper motion, reducing the likelihood of such an origin (Mignani *et al.* 2002).

Nearby (arcsec) ejecta, presumably previously expelled from Cir X-1, were found to brighten in the radio after flare events, suggesting re-energisation by highly relativistic invisible outflows with $\Gamma > 15$ (assuming association with the flare immediately preceding the change: Fender *et al.* 2004) and inclination angles close to the line of sight ($\theta < 5^{\circ}$). Cir X-1 was the second (after Scorpius X-1) neutron star system to show evidence of such invisible relativistic outflows. The implied velocities make the jets from Cir X-1 some of the fastest in our Galaxy. These calculations assume a distance to the source of 6.5 kpc as put forward by Stewart *et al.* (1993), but that also serves as a compromise between the more recent estimates of Jonker

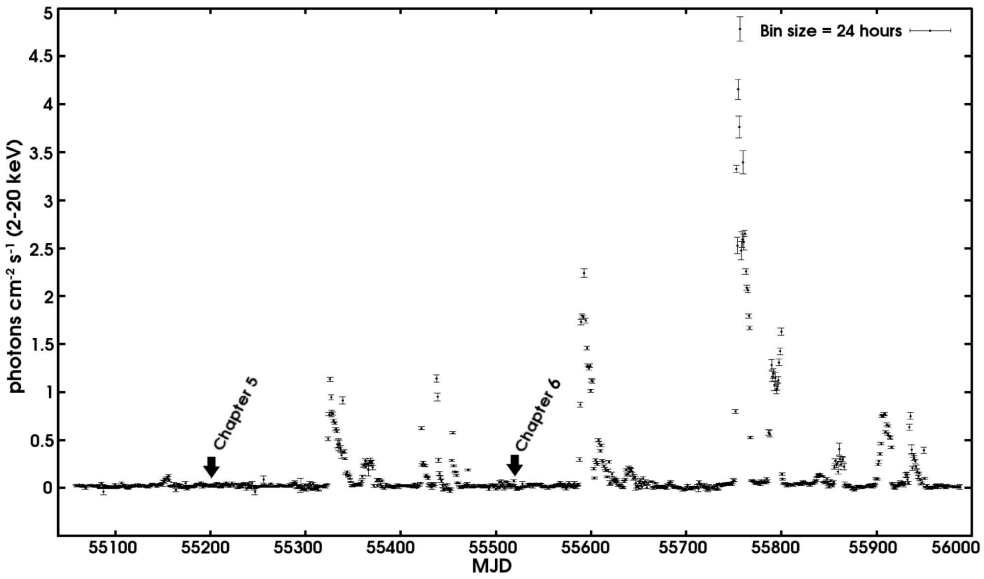

Figure 5.1: Circinus X-1 MAXI 2-20 keV X-ray light curve for the period covering 2009 Aug 15 to 2012 Mar 01. Periods of intense activity are clearly seen, separated by epochs during which Cir X-1 is fainter. The time of observations discussed in this chapter are marked by the left most arrow, as well as observations which will be discussed in Chapter 6. Data acquired from the MAXI website at `http://maxi.riken.jp/top/index.php?cid=1\&jname=J1520-571`.

& Nelemans (2004: 7.8 - 10.5 kpc) and Iaria *et al.* (2005: 4.1 kpc). The jets themselves are resolved on arc-second scales in X-rays (Heinz *et al.* 2007, Soleri *et al.* 2009b) and arc-second to arc-minutes in radio (Stewart *et al.* 1993, Tudose *et al.* 2006), with the jets appearing curved on the larger scales. X-ray shocks, believed to be caused by jet impacts, have also been observed (Sell *et al.* 2010).

Significant uncertainty exists in the orientation of the system and its jets. Though the ultra-relativistic velocities calculated by Fender *et al.* (2004) imply jet angles close to the line of sight, evidence exists that can indicate otherwise. Fitting of blue-shifted and red-shifted X-ray emission features (possibly originating in the jets) has been used to calculate lower jet speeds (0.08c) and a jet inclination near perpendicular to the line of sight (Iaria *et al.* 2008). Furthermore, there are multiple signs of an edge-on accretion disc, including X-ray P-Cygni profiles (Brandt & Schulz 2000) and X-ray dips (Shirey, Levine & Bradt 1999). However, these results need not be contradictory if one assumes that the jets are either misaligned with the orbital plane (e.g. Maccarone 2002), follow non-linear flow paths (hinted at by the curved jets observed in large scale maps of Tudose *et al.* 2006), or precess.

Tudose *et al.* (2008) report little variation in the projected jet axis ($129^\circ \pm 13^\circ$ east of north) over a decade of radio observations and imaging (1996-2006). In addition, nearly all their maps show brighter resolved structure to the south east of the

system which is interpreted as the direction of the approaching jet, with very little evidence of a receding counterpart to north west (when visible, it is significantly fainter than the structure to the south east) supporting a persistent orientation along the line of sight. Together, this would imply minimal precession occurs in the system over long periods of time. However, recent LBA observations by Miller-Jones *et al.* 2012 indicate a projected jet angle closer to east west (112° east of north) and a pair of structures evolving away from the core with approximately equal intensities, which implies the jet is aligned perpendicular to the line of sight. The fact that these latter observations map the system on milli-arcsecond scales (compared to the arcsecond scales of Tudose *et al.* 2008), means that one still cannot distinguish between precession and a curved jet.

Between 1997 and 2002 Cir X-1 had been growing steadily fainter in the X-rays, and though radio flare events were known to reach $>$ 1 Jy in the 1970s (Haynes *et al.* 1978) they too had been in decline, reaching at most 10s of mJy by the 2000s (Fender, Tzioumis & Tudose 2005). Some renewed radio activity, similar to that observed in the 70s and 80s, occurred in 2006 (Nicolson 2007) but was followed by a return to minimal activity: a historically faint period during which the observations discussed in this chapter were taken. A summary of Cir X-1's long term X-ray activity ($\sim$33 years from discovery to 2003) is given in Parkinson *et al.* (2003) with a concatenated plot shown in their figure 1. The system entered a relatively active phase in mid 2010 (a few months after this chapter's observations), with X-ray flares reaching hundreds of mCrab (Nakajima *et al.* 2010) and radio flares $>$ 0.1 Jy (Calvelo *et al.* 2010b). The system then returned to a low activity level for several months, much like that during which this chapter's observations were taken, before undergoing another period of intense flaring. MAXI (Matsuoka *et al.* 2009) provides monitoring data for a great deal of X-ray sources in the sky to the public, including Cir X-1, and I include the MAXI light curve in Figure 5.1. It appears that Cir X-1 currently spends the majority of its time in this low activity state, but does enter periods of higher activity which can last for months and during which both peak flare and inter flare X-ray flux levels (and likely radio) are significantly higher.

5.3 Observations & data reduction

The goal of the observations was to monitor Cir X-1 over a complete orbital period using ATCA-CABB, with the hope of catching the flare in significant detail, as well

as its subsequent decline. Additionally, it would be possible to see if the structure of the source varied in any way around an orbit. Moin *et al.* (2011) carried out a similar campaign on milli-arcsecond scales with eVLBI but were only able to detect the source shortly after periastron. Furthermore, with MAXI monitoring the system in X-rays, it would be possible to simultaneously compare a flare event in the two regimes.

By continuing to observe after the flare it would be possible to observe variation in ejecta luminosities caused by unseen relativistic jets, or perhaps even the jets themselves. However, given the observation time available post-flare ($\sim$ 6 days), an assumed object distance of 6.5 kpc and a maximum image resolution of 1.5", one would only be able to observe flows with minimum $\beta \approx 9.5$ (see Equation 2.1). Therefore, if hoping to observe the effects of flows related to the flare that was predicted to occur during these observations, flow velocities similar to those implied by Fender *et al.* 2004 would be required. This does not prevent the detection of structural variations at any point during the run which might arise as a result of ejections from earlier events.

5.3.1 Radio

Observations of Cir X-1 were carried out on 2009 Dec 30, 31, 2010 Jan, 01, 02, 04, 05, 06, 07, 08, 09, 10, 11, 12, 13, 14 and 15 using the ATCA-CABB in 6A configuration (minimum baseline of 337m, maximum of 5939m) at both 5.5 GHz and 9 GHz (see Table 5.1). A significant part of the observations were controlled and monitored on site by myself. PKS J0825-5010 (PKS B0823-500) was used as the primary flux and bandpass calibrator when ever possible, with PKS J1939-6342 (PKS B1934-638) used in a few cases. Phase calibration was carried out using PMN J1515-5559 (PKS B1511-55) for the first half of the observations until it was noticed that on-screen levels indicated polarisation (5-10%) of the source, at which point PMN J1524-5903 (PKS B1520-58) was made the phase calibrator, with some days including observations of both phase calibrators to make it possible to eliminate the slight polarisation effect of the first phase calibrator during analysis. Observation times varied between 8 and 12 hour runs each day (predicted RMS noise between 8 and 6 μJy at 5.5 GHz and between 10 and 8 μJy at 9 GHz), for a total of $\sim$ 135 hours on the source. All data and image processing was carried out by myself using MIRIAD (Section 1.3.2).

Table 5.1: Circinus X-1 observation log. The table lists the dates of observations, start (on source) Modified Julian Day (MJD), total on-source time, average daily radio flux densities at each frequency (5.5 GHz and 9 GHz, measured via point source fits) with the range of daily light curve values included in square brackets, along with image noise levels. Average flux densities become less useful after flare onset at the end of observations on 2010 Jan 09, due to the level of variation observed over single observation sessions (reflected in the uncertainties = daily light curve σ). Values listed for 2010 Jan 09 are measured from the data available prior to the flare.

Date (UT)	MJD start	MJD end	Total time (h)	$F_{5.5}$ [range] (mJy beam^{-1})	Noise $_{5.5}$ (μJy beam^{-1})	F_9 [range] (mJy beam^{-1})	Noise $_9$ (μJy beam^{-1})
2009 Dec 30	55195.745	55196.055	6.50	0.89 ± 0.08 [0.72:1.06]	9.7	0.58 ± 0.11 [0.39:0.77]	13.7
2009 Dec 31	55196.707	55197.056	6.91	1.16 ± 0.11 [0.88:1.42]	12.9	0.75 ± 0.10 [0.66:1.01]	24.0
2010 Jan 01	55197.756	55198.057	6.29	0.67 ± 0.15 [0.47:0.94]	14.5	0.41 ± 0.14 [0.32:0.72]	30.3
2010 Jan 02	55198.704	55199.076	5.82	0.99 ± 0.07 [0.87:1.10]	10.8	0.60 ± 0.08 [0.51:0.77]	15.5
2010 Jan 04	55200.706	55201.181	9.91	0.89 ± 0.08 [0.73:1.10]	8.6	0.55 ± 0.09 [0.39:0.73]	11.1
2010 Jan 05	55201.703	55202.162	9.43	0.77 ± 0.17 [0.37:0.91]	9.0	0.41 ± 0.18 [0.18:0.72]	11.8
2010 Jan 06	55202.703	55203.160	7.15	0.49 ± 0.12 [0.29:0.78]	11.1	0.25 ± 0.06 [0.18:0.40]	13.3
2010 Jan 07	55203.682	55204.148	9.51	0.58 ± 0.11 [0.41:0.79]	9.1	0.33 ± 0.10 [0.22:0.58]	12.1
2010 Jan 08	55204.687	55205.163	9.43	0.59 ± 0.10 [0.44:0.75]	9.1	0.30 ± 0.09 [0.17:0.50]	12.1
2010 Jan 09	55205.681	55206.157	9.41	0.93 ± 0.42 [0.42:6.67]	9.8	0.93 ± 0.65 [0.23:9.83]	12.2
2010 Jan 10	55206.683	55207.162	10.03	21.3 ± 8.81 [11.7:43.0]	38.4	16.6 ± 6.10 [6.57:29.9]	43.0
2010 Jan 11	55207.684	55208.139	9.53	9.23 ± 1.33 [7.60:11.3]	20.5	6.11 ± 1.26 [2.64:6.21]	25.4
2010 Jan 12	55208.661	55209.150	10.19	5.48 ± 0.80 [2.05:5.61]	16.5	4.02 ± 0.69 [0.96:3.65]	18.3
2010 Jan 13	55209.685	55210.153	7.46	3.72 ± 0.51 [2.48:3.96]	15.6	3.19 ± 0.21 [0.97:2.31]	20.1
2010 Jan 14	55210.672	55211.142	9.71	2.37 ± 0.21 [1.61:2.35]	10.4	2.39 ± 0.13 [0.99:1.43]	12.8
2010 Jan 15	55211.665	55212.072	8.22	1.54 ± 0.11 [1.32:1.79]	9.9	0.91 ± 0.12 [0.75:1.08]	14.7

5.4 Analysis & results

5.4.1 Flare event

Using the ephemeris detailed in Nicolson 2007 (ATEL #985) I calculated that a flare would occur near 2010 Jan 11. Observations from 2009 Dec 30 through to 2010 Jan 09 indicated Circinus X-1 was relatively stable at flux densities of $\sim$ 1 mJy at 5.5 GHz, and slightly lower at 9 GHz (for a list of average daily flux densities see Table 5.1). A sudden rise in flux density was detected towards the end of a day's observations at 2010 Jan 10 02:30 UT, continuing until scheduled time ended at $\sim$ 03:50. Upon returning to the source $\sim$ 12.5 hours later, I found Cir X-1 in decline, indicating the peak of the outburst had passed.

To better analyse the flare, I divided datasets from each day into smaller cuts ($\sim$ 30 minutes), produced images from these cuts, and used the MIRIAD command IMSTAT to measure the maximum flux density within a 10×10 arcsec^2 box around the system's previously established position (15:20:40.9 -57:10:00; based on image fitting from Tudose *et al.* 2008: tables A1 & A2); i.e. the system's 'core' from which the highest levels of radio emission are observed. This allowed me to create a detailed light curve (Figure 5.2) while maintaining an acceptable level of error on each measurement. There is the possibility that something in the selected region will outshine the core (e.g. Tudose *et al.* 2008 show some epochs where the south eastern component appears brighter than the core); however, none of my imaging showed any evidence for this and thus I could safely assume the highest measured peak flux density would only come from the core.

The flare event (shown in more detail in Figure 5.3) appears to start at MJD $\sim$ 55206.05, rising from pre-flare levels of 1 mJy to $S_{5.5} \sim 6.5$ mJy and $S_9 \sim 10$ mJy in just under 2 hours, before the day's observations finally ended. Though some variations in flux densities are visible in pre-flare day light curves (a scatter of $\sim$ 0.5 mJy visible in most) they do not follow any obvious trend. However, the light curve from the day of the flare shows a gradual rise in flux density prior to the event, with levels at the start of the run $\sim$ 0.5 mJy. There is also a brief peak in flux density at both frequencies at MJD 55205.97 (2 hours prior to the flare - more easily visible in Figure 5.3), about one hour in duration. Observations the following day began at MJD 55206.683; the initial flux densities of $S_{5.5} = 43.0 \pm 0.5$ mJy and $S_9 = 29.9 \pm 0.6$ mJy proceeded to decay rapidly (roughly power law decay: log $S_{9GHz} \approx$ -3.03(log T) - 2.12 at 9 GHz, log $S_{5.5GHz} \approx$ -3.03(log T) - 1.96, where T is days since outburst start i.e. MJD - 55206.05) before appearing

to level off at $S_{5.5} \sim 15$ mJy and $S_9 \sim 11$ mJy near MJD 55207. Subsequent days show continued decay (once again a rough power law: log $S_{9GHz} \approx$ -1.4(log T) - 1.9, log $S_{5.5GHz} \approx$ -1.6(log T) - 1.6) up to the final day of observations where flux densities remained above pre-flare levels: $S_{5.5} \approx 1.5$ mJy and $S_9 \approx 1.0$ mJy. MAXI X-ray measurements (also shown in Figure 5.2) indicate a flare occurred at, or just before, the time of the radio event, with statistical errors and the lack of a radio peak preventing a better comparison of the event in the two light curves.

There is little or no evidence of a turnover in flux densities either side of the flare, suggesting the peak of the event occurred towards the middle of the gap between observations (i.e. $\sim$ MJD 55206.5) at both frequencies. This event appears to differ from previous flares (e.g. 2000 Oct 20/21 and 2002 Dec 04/05; figure 6 of Tudose *et al.* 2008) in a variety of ways. Firstly, this flare's 'peak' (estimated to last from initial rise to the levelling off observed the next day) has a duration of almost a full day, compared to the 2000 Oct 20 flare light curve whose peak appears to last less than a third of that time. Secondly, this flare's rise is significantly steeper than many observed in the past; $\sim \times 6$ increase in 2 hours at 5.5 GHz and $\sim \times 10$ at 9 GHz, compared to the rise observed in 2002 Dec 04 where levels only doubled over half a day. However, the light curve following the 2002 Dec event does reveal a similar gradual decline post flare (the 2000 Oct 20 light curve might also show the start of a similar decay trend, but is cut short). Such comparisons are difficult to justify, as Circinus X-1 was consistently far brighter in the former epochs than that during which this chapter's observations were carried out. As a result, I can only crudely estimate the peak of the flare to have a flux density between 0.05 and 0.1 Jy at both frequencies, based on extrapolations of the initial rise and decay rates.

The lower sections of Figures 5.2 and 5.3 show the spectral index ($S_\nu \propto \nu^\alpha$) of Cir X-1's core, based on the 5.5 GHz and 9 GHz flux densities. Indices remain negative for most of the pre-flare period, with an average value of $\alpha \sim -1.0$ but with significant scatter (σ = 0.45) and statistical errors. This is, as expected, indicative of optically thin synchrotron emission (Section 2.1.3), and consistent with previous measurements of Cir X-1's spectral index (e.g. Tudose *et al.* 2008, figure 10). There does appear to be a rise to positive indices near MJD 55198, which would correspond to phase $\sim$ 0.5 (based on current ephemeris estimates). However, the day's data (2010 Jan 01/02) to which the points correspond, was particularly difficult to calibrate, making the flux densities unreliable. This effect might initially suggest an apastron radio flare (Fender 1997, Tudose *et al.* 2008), but it is more likely that the positive index is a result of errors and not a real physical effect.

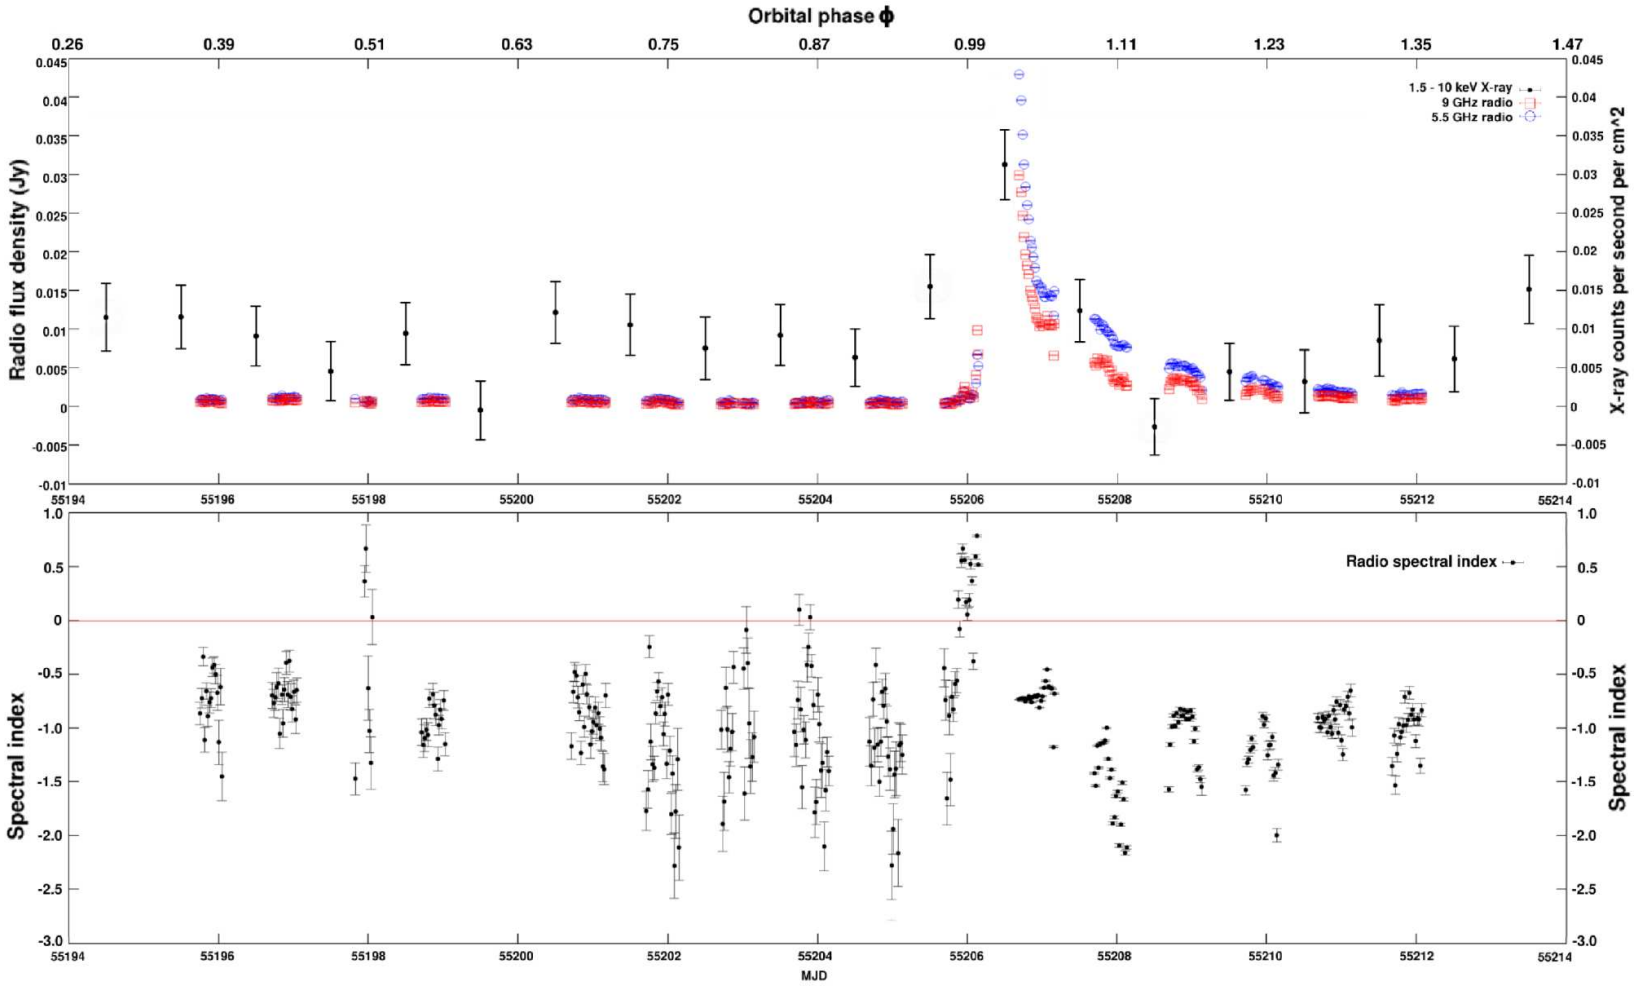

Figure 5.2: MAXI X-ray, ATCA-CABB radio (top panel) and radio spectral index (bottom panel) light curves for 2009 Dec 30 to 2010 Jan 15. Radio flux density errors (1σ) are included but are too small to distinguish clearly. The flare event is easily noticeable in all light curves, as well as the expected effects on spectral index during rise and decay.

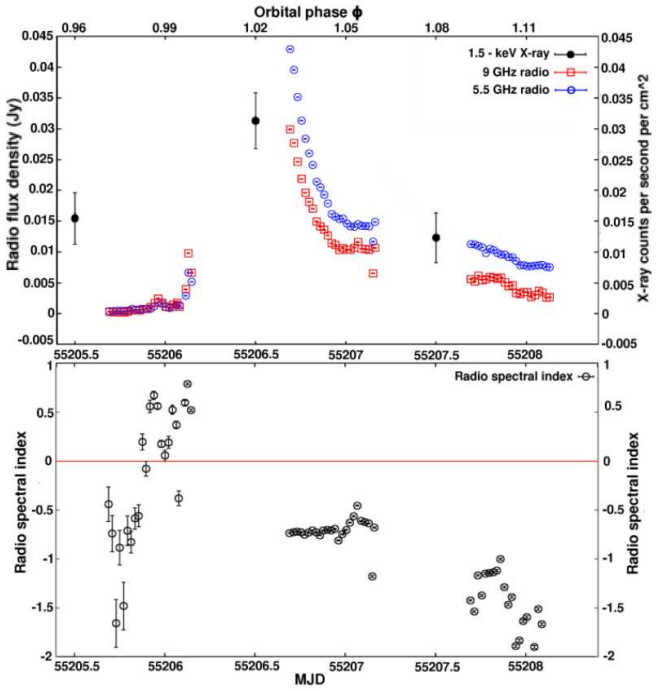

Figure 5.3: MAXI X-ray, ATCA-CABB radio (top panel) and radio spectral index (bottom panel) light curves for the flare event. Features such as the pre-flare 'hump' (MJD $\sim$ 55205.95) are more easily visible.

The spectral index begins a clear rise on the day of the flare, reaching $\alpha \sim 0.5$ twice during the days observations, firstly during times concordant with the 'bump' 2 hours prior to the flare, and then towards the end of observations with the rise of the flare itself. By the following day, during flare decline, spectral index has returned to negative values, $\alpha = -0.7 \pm 0.1$, with very little scatter. Subsequent days show increasing levels of error and scatter as values decline.

The behaviour of the light curves and radio spectrum are consistent with that expected from either internal shocks created within outflows as a result of quasi-continuous ejections (Section 2.2.3.5; Kaiser, Sunyaev and Spruit 2000) or an adiabatic expansion of ejected clouds of relativistic particles (as in van der Laan 1966), as is the case for nearly all outbursts from XRBs. If dealing with such expanding knots of emission, one could expect these objects to eventually move a sufficient distance from the core to be resolved separately, which is not observed. Such objects might still exist as part of the core emission, but if their velocities are far lower than those of the flows discussed in Fender *et al.* 2004, then there may not be sufficient time after outburst to directly observe the separation before the knots fade below detectable levels. Crude estimations of the velocities possible for this scenario can be calculated using Equations 1.1 and 1.2, together with the image synthesised beam size ($\sim$ 1.5" for a 12 hour 9 GHz CABB image) and the total

time taken for core flux densities to stabilise after the flare ($\sim$ 6 days). Taking an assumed jet inclination of $\theta \sim 5^{\circ}$, and half the beam size as a required minimum distance for detectable change in the core structure, I find that such objects could remain unresolved if moving with Lorentz factors that do not exceed $\Gamma \sim 5$. This does not eliminate the possibility of ultra-relativistic flows re-energising media further downstream, since internal shocks can fade only to reignite during additional collisions (e.g. XTE J1748-288: Hjellming *et al.* 1998).

As mentioned earlier, flux densities actually appear to level off towards the end of Jan 10 before commencing the gentler decay. There have been several examples of multi-peaked flares from Cir X-1 in the past (Thomas *et al.* 1978, Tudose *et al.* 2008) with varying numbers of peaks and timescales, but the gaps in the observations prevent me from confirming such behaviour.

There is an unfortunate effect visible in the light curve, which becomes more apparent in the spectral index plot, and that is a trend towards more negative α (and lower flux densities) at the beginning and end of each run, particularly when the source is brighter (i.e. post-flare); this is most easily seen in the observations around MJD 55209. The effect is likely a result of Cir X-1's low elevation (25° at minimum) during early and late hours of individual observation runs (hence why it becomes more evident once observations switch from $\sim$ 9 hour runs to full 12 hour runs after MJD 55200) leading to increased atmospheric opacities, or possibly, rising sensitivity to large scale structure as projected baselines decreased in length. While self calibration does appear to eliminate this turn over effect (standard calibration methods fail to do so), it also alters source flux density in such a way that makes measurements unreliable. Thus by using standard calibration routines I retain flux density accuracy, while the general trend and most important features of the light curve remain clear.

5.4.2 Imaging, modelling and subtractions

5.4.2.1 Original images

I created images for each day's individual data set, at both frequencies. These are presented in figures 5.4 (5.5 GHz) and 5.5 (9 GHz). I applied phase only self-calibration in an effort to improve image fidelity and reduce the effects of artefacts caused by the large flux density variations following flare onset. All flux density measurements on the other hand (Table 5.1), are taken from images that have not undergone self-calibration, but only the standard calibration methods detailed in the

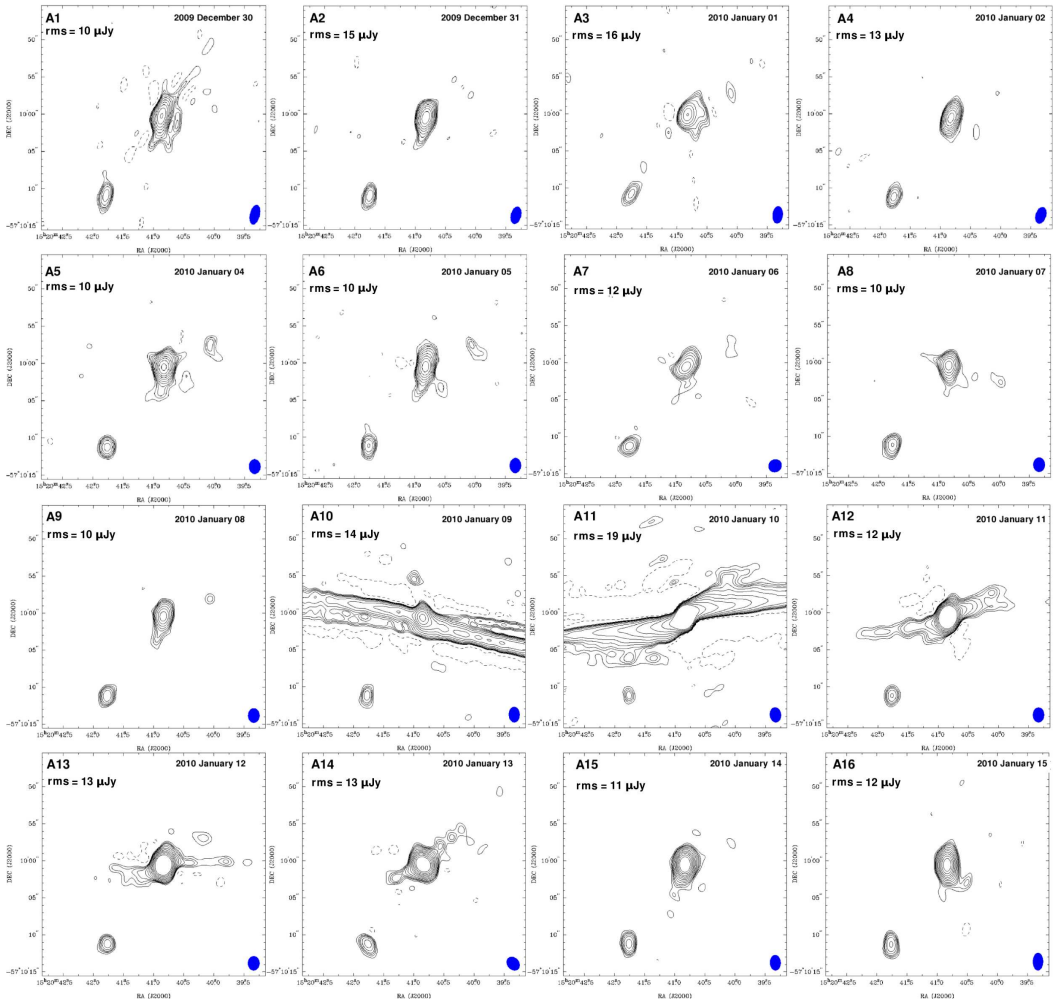

Figure 5.4: Un-altered data radio maps of Circinus X-1 at 5.5 GHz. These maps use ATCA-CABB data that has undergone both normal calibration and phase calibration routines prior to final deconvolution. Maps were cleaned using multi-frequency clean subroutines. Weighting is determined with a robust factor of 0.5 (an optimal compromise between natural and uniform weighting). Contour lines are at -2.8, 2.8, 4, 5.6, 8, 11, 16, 23, 32, 45, 64, 90$\times$rms noise of each epoch (matching the scheme used in Tudose *et al.* 2008, rms listed in the top left corner of each panel). Beam sizes are approximately 4 arcsec^2. Artefacts caused by flare variability are easily visible in images A10 through A13.

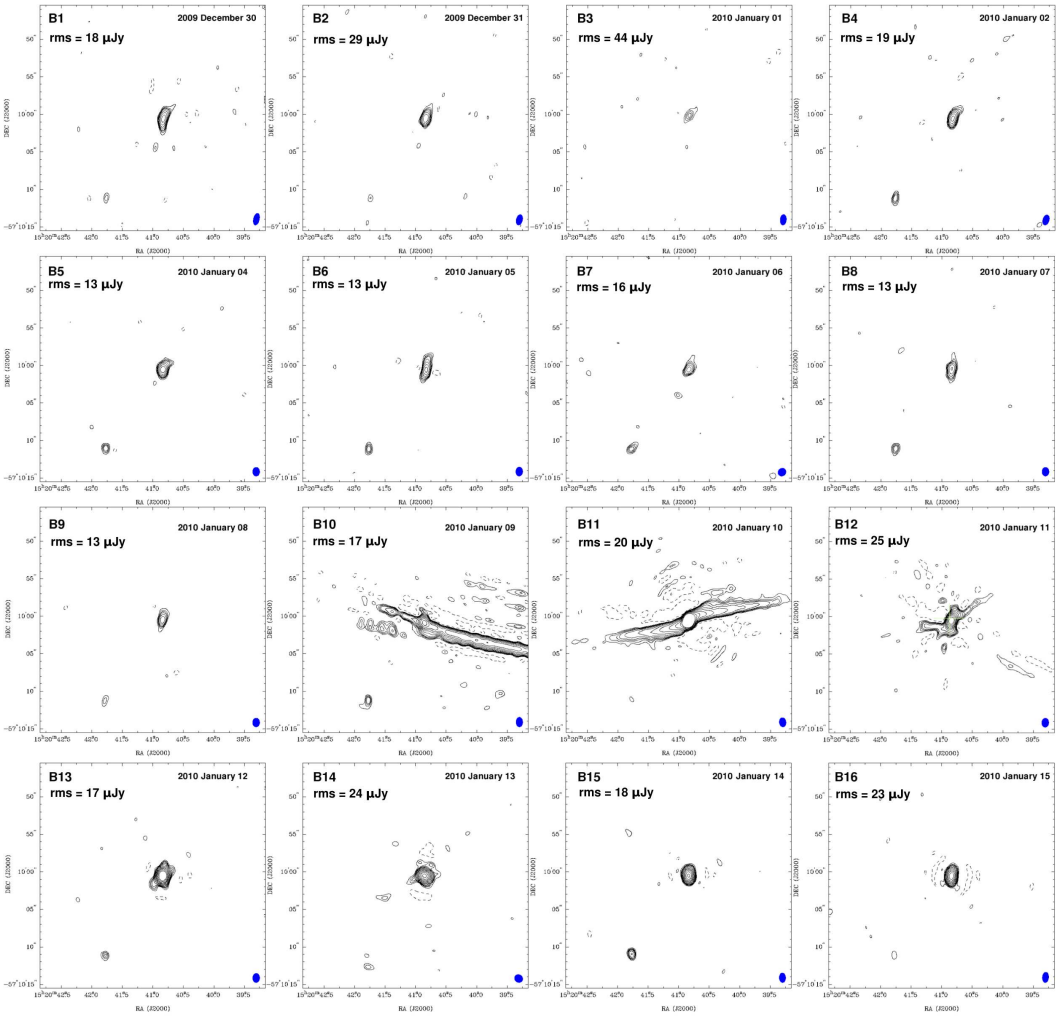

Figure 5.5: Un-altered data radio maps of Circinus X-1 at 9 GHz. These maps use ATCA-CABB data that has undergone both normal calibration and phase calibration routines prior to final deconvolution. Maps were cleaned using multi-frequency clean subroutines. Weighting is determined with a robust factor of 0.5 (an optimal compromise between natural and uniform weighting). Contour lines are at -2.8, 2.8, 4, 5.6, 8, 11, 16, 23, 32, 45, 64, 90$\times$rms noise of each epoch (matching the scheme used in Tudose *et al.* 2008, rms listed in the top left corner of each panel). Beam sizes are approximately 1.5 arcsec^2. Artefacts caused by flare variability are easily visible in images B10 and B11.

MIRIAD guide. All *uv* data with *uv* distances shorter than 6 kλ at 5.5 GHz and 10 kλ at 9 GHz were ignored in order to eliminate the more diffuse emission from the surrounding nebula (see Section 5.4.3). I will first address the 'normal' images from data sets before, and several days after the flare, in which conditions are more stable and similar from day to day, prior to discussing the images during and shortly after the event. The stability during the pre-flare days also allowed me to combine the data from the first nine observations to produce a deep 5.5 GHz map of the system's core and nearby region; this image is presented in the lower left hand panel of Figure 5.6. Figure 5.6 also presents the normal robust weighted images for the four days after the flaring episode (corresponding to images A13–A16 in Figure 5.4), as well as images in which the (variable) core has been subtracted in an attempt to measure weak variability on arcsec scales.

The system's core (centred) is well defined in all images, and in many of the 5.5 GHz examples shows a distinct south/south-east extension (particularly evident in Figure 5.6's deep map) which can be interpreted as evidence for an approaching jet, i.e. the jet component aimed along our line of sight. The position angle of this jet is similar, but appears to be slightly more southerly than that observed in previous data sets (Tudose *et al.* 2008). Some of the individual images also show a slight north/north-west protrusion, or general elongation of the core along the north-south axis which cannot be fully explained by beam shape alone, indicating the presence of the (likely) receding jet. The maps produced from 9 GHz data show similar elongation of the core, as well as extensions related to the jets. However, unlike at 5.5 GHz, the 9 GHz northern jet structure appears to be just as visible as its southern counterpart, and often more easily so.

In figure 5.7 (upper row) I compare the deep pre-flare map to a single day's observations from the pre-flare period of 2002 Dec. The difference in the structure near the core is significant, as one can clearly see a second bright component a few arc-seconds to the south west of the core in the 2002 map which does not appear in the 2009/2010 map. This component was that which varied following flare events, leading to the high jet velocity estimates by Fender *et al.* (2004), and though not persistent, structure has appeared in or very near this location multiple times in the past at various levels of intensity (see Tudose *et al.* 2008).

The lower row of Figure 5.7 shows point source subtraction results based on fits to the location of the core. It should be noted that the two far components (SE and NW) in the maps would be less likely to appear in the older map due to the lower signal to noise available (assuming similar flux densities). The differences in residual structure around the core is now more prominent. Firstly, though in both cases

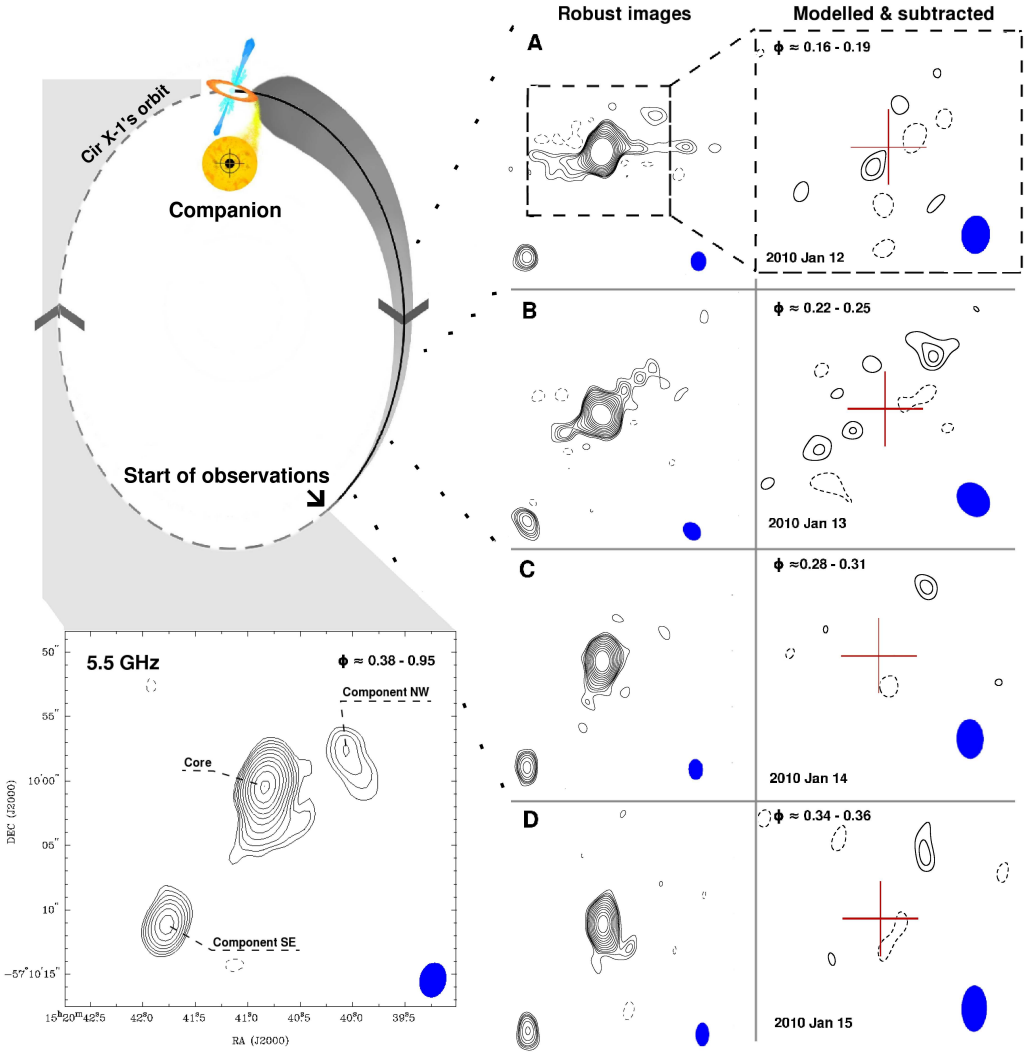

Figure 5.6: Illustration of Circinus X-1's orbital behaviour with related radio maps. My intention is to show how the observed radio structure is related to the orbital motion of the system. The upper left diagram is a representation of the system, assuming a high mass companion and an eccentric orbit. The lower left image is a deep 5.5 GHz radio map (beam size = 2.5 × 1.8 arcsec^2) of the core and nearby region of Cir X-1, combining data from the first 9 observations, i.e. those taken prior to the flare event (ignoring *uv* data with *uv* distances shorter than 6 kλ), resulting in rms noise of 3.5 μJy beam^{-1}. The images to the right show the unaltered data radio maps (left) and the final model subtracted maps (right) for the final 4 days of observations, after the flare event. Note that the right hand maps cover a smaller region around the core (as illustrated in the top set) in order to improve visible details. The core and nearby bright components referred to in section 5.4.2.1 are labelled. Contour lines in all maps are at -2.8, 2.8, 4, 5.6, 8, 11, 16, 23, 32, 45, 64, 90× r.m.s. noise of each epoch (listed in related panels from Figures 5.4 and 5.12).

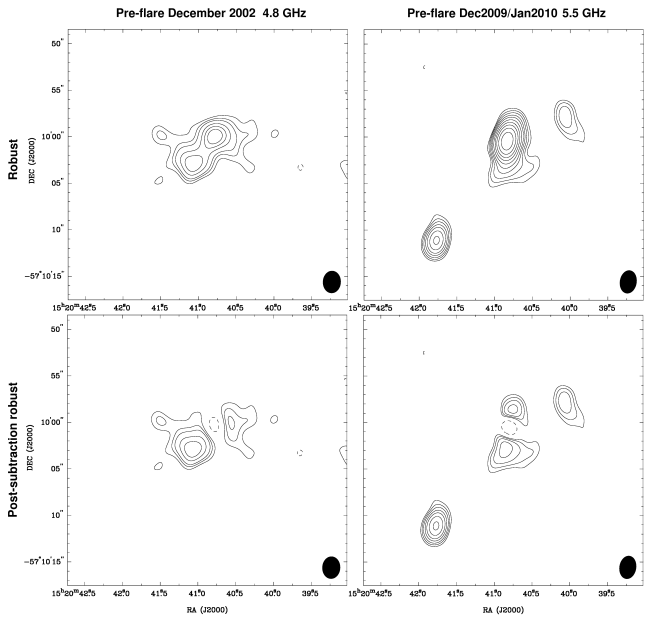

Figure 5.7: Cirxinus X-1 point source subtraction maps. The left hand column shows contours (-2.8, 2.8, 4, 5.6, 8, 11$\times$ r.m.s. noise of 55 μJy beam^{-1}, beam size = 2.4 $\times$ 1.9 arcsec^2) for data from pre-flare observations taken on 2002 Dec 02 (See Fender *et al.* 2004). The right hand column are contour images made using the nine pre-flare observations from this chapter's data-set (-2.8, 2.8, 4, 5.6, 8, 11, 16, 23, 32, 45, 64, 90$\times$ r.m.s. noise of 3.5 μJy beam^{-1},beam size = 2.5 $\times$ 1.8 arcsec^2) The upper row shows un-altered images using a weighting scheme of robust = 0.5. The bottom row show the same images but with a point source fitted and subtracted from the approximate location of the system's core. Though there are clear differences in structure near the core itself, it should be stressed that the lowest contour levels in the 2002 images are close to the flux densities measured from the peak of component SE; i.e. there is a good chance neither component SE nor NW would be visible in the older map assuming their flux densities remain constant.

one can observe residual structure on opposite sides of the core void (i.e. components likely related to the jet pair), you can clearly see the south-eastern component is significantly brighter than its north-western counterpart in the 2002 image, unlike my image which shows more similar intensity between the two (southern residual peak flux density of 79 $\pm$ 4 μJy, northern residual with 75 $\pm$ 4 μJy). Secondly, the axis along which these structures lie is different in the two images, with my map showing an axis ($\theta \approx 170 \pm 15^\circ$) far closer to north-south than 2002's near east-west orientation.

The images created from data sets during, and shortly after, the flare (panels A10 - A13 of Figure 5.4) show the artefacts caused by rapid variation in brightness of a source during an observation (linear elongation for an E-W array: see Section 1.2 and Section 3.4), and as such much of what one sees is not real, or requires very

careful interpretation. There do appear to be resolved structures along the SE-NW axis (i.e. the jet axis) that do not make up part of the general elongation of the core, and thus one may be tempted to believe they are jet related, however, given the time since flare onset and angular distance of these structures ($\sim$ 5") these can quickly be disregarded as real new components within the system. The intensity of distortions is reduced as flux densities and decay rates decline. The core at no point appears to resolve into two separate bright components as has been observed in the past (i.e. as in figure 4 of Fender *et al.* 2004, Tudose *et al.* 2008).

Other than the core, the most notable structure in the images is the separate emission component to the south east (labelled component SE in Figure 5.6's deep map), located at approximately 15:20:41.75, -57:10:11 (J2000). Upon initial inspection one could assume this to be a possible component of the approaching jet, due to its proximity to Cir X-1 and the fact its position angle would fit well with that of previously observed outflow structures. The source is detected but unresolved at both frequencies, with an average spectral index of $\alpha \sim$ -1.4. The angular distance from the core to this component is $\sim$ 13", which, based on previously estimated outflow Lorentz factors ($\Gamma > 10$, distance $\approx$ 6.5 kpc), would mean flux density variations linked to core outbursts and subsequent re-energisation via the jet will be delayed by $>$ 30 days (assuming constant flow velocity). It may still be possible to observe re-energisation of the component by flows from an older core outburst, with a large phase offset resulting from its distance. Light curves for this component at both frequencies are visible in Figure 5.8, with errors based on noise levels of each image. Measurements from 2010 Jan 9, 10 and 11 (A10–A12) are not included as low-level artefacts were noted to extend from the core and overlap with the component emission, reducing the validity of those estimates. There is at best only a hint of variation in the flux density of this component at 5.5 GHz, but some evidence at the $\sim 2\sigma$ level at 9 GHz.

A second fainter component also regularly appears north-west of the core in the 5.5 GHz images (15:20:40.1, -57:09:58 [J2000] - labelled component NW in Figure 5.6), though remaining undetected at 9 GHz. The source appears resolved in both the deep map and individual epochs, often displaying a south western extension. The angular distance ($\sim$ 7") would likely involve long delays between outbursts and subsequent reactions in the component ($\sim$ 20 days). The 5.5 GHz light curve is included at the bottom of Figure 5.8. As with component SE, there is no unambiguous evidence for variability.

In an effort to determine the origin of these emission sources, I reviewed past observations of Cir X-1. Going over images from Tudose *et al.* 2008, I found sev-

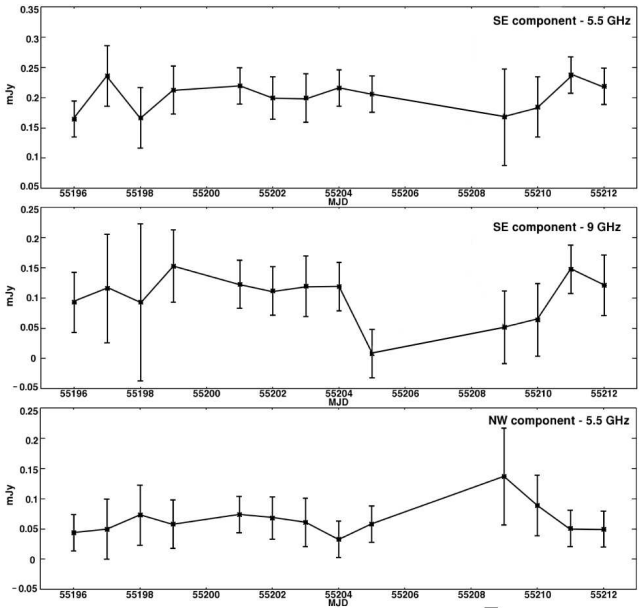

Figure 5.8: Radio light curves for Circinus X-1 components SE and NW (see Figure 5.6, bottom left). Values are from peak flux density measurements. The large gap in points is a result of unreliability of measurements from data taken during the flare and its decay.

eral with distinct compact emission near the location of component SE (a total of 5 epochs). Data sets from individual days of pre-CABB observations tend to yield images with noise levels close to the flux of these components, making detection difficult. Therefore, I created stacked data sets which combined observations of close epochs (i.e. within a maximum of 2 orbital periods difference in observation time) to produce deeper maps in an effort to track down further evidence of the components. These were: July 1996, October 1998, May 2001 and December 2002 (See Tudose *et al.* 2008 for details on these data sets). Component NW remained undetected in all revised images; however, based on image noise levels, $3\sigma > 0.05$ mJy in all cases, which compared to my flux measurements would make the component difficult to detect. Component SE was visible in nearly all stacked images with the following flux densities: Dec 2002 - 0.17 $\pm$ 0.02 mJy; May 2001 - 0.25 $\pm$ 0.03 mJy; and Oct 1998 - 0.20 $\pm$ 0.04 mJy. The exception was Jul 1996, where there was difficulty distinguishing separate components due to the abundance of structure near that portion of the image. These fluxes are not dissimilar from the range measured in this data.

Linear polarisation was detected ($> 3\ \sigma$ levels) at the position of the core on 10 Jan 2010 at 9 GHz (a day when observations of the second, unpolarised calibrator were available). No other images show any distinguishable polarisation in regions of interest and the single detection is relatively weak ($<$ 2% of core flux density: $<$

0.1 mJy), thus few strong conclusions can be garnered from it. Based on estimates quoted in Tudose *et al.* (2008) a correction of $|\Delta\theta| \leq 9^\circ$ must be applied to account for Galactic Faraday rotation. Vector orientation indicates electric fields aligned away from the jet axis, but not entirely perpendicular; similar to what was observed from the core in previous epochs (see section 7 of Tudose *et al.* 2008).

5.4.2.2 Difference images

In order to compare data sets on a day to day basis in search of structural changes associated with Cir X-1's flare, it was prudent to select a single day's data to act as a reference. This observation could then be subtracted in the *uv* plane (prior to construction of images) from each epoch, leaving only those regions of flux density which had changed in intensity. The reference data set would have to fall on a day sufficiently distant in time from a flare event, so that any subsequent variation in flux levels have diminished, both of the core and any nearby structures that may have reacted to the event. Based on the flare event captured in this chapter's observations, it is evident that even 6 days after flare onset, the core continues to remain above pre-flare flux density levels. It is best to select a day that falls beyond a half-period after a flare event as a reference, i.e. 8 days or more. In case of a flare having occurred prior to that in this data, the reference would fall on 2010 Jan 01. Unfortunately data quality must also be taken into account, as some days include gaps (or shorter observation durations) and increased levels of RFI. As mentioned in section 5.3.1, Jan 01/02 data suffered from problems making it an inadequate reference choice. Data from 2010 Jan 04 suffered from neither gaps, unexpected errors, nor high levels of RFI, and additionally covered a near full 12 hours, meaning it could not only be cropped to match hour angles of the earlier half of observations (each $<$ 9 hours long) but would also cover most of subsequent observations for effective comparisons.

Figures 5.9 and 5.10 show difference maps for all days at 5.5 GHz and 9 GHz respectively. The days prior to Jan 04 show excess flux at, or near the core; whereas the days following, but prior to the flare event, show regions of over-subtraction in the same location. This is a logical result from the trend in fluxes listed in Table 5.1 and would suggest a gradual decay in flux density over the first 10 observations, perhaps following a flare event preceding that of Jan 09. There are a pair of structures visible in the December 5.5 GHz images (panels C1 and C2) towards the north-west which should be noted, as they appear along the known jet axis and in the same locations in each of the images, then go undetected in most pre-flare maps. Images from

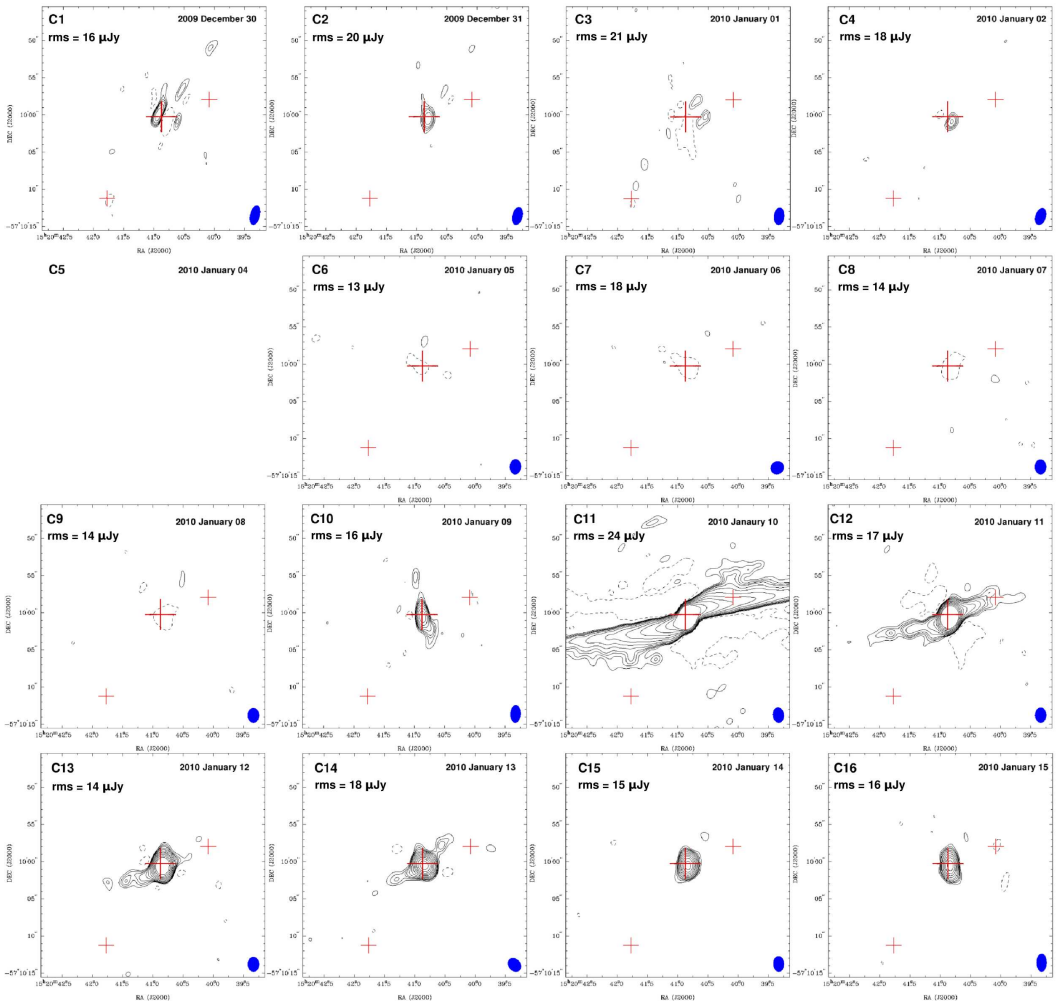

Figure 5.9: Difference radio maps of Circinus X-1 at 5.5 GHz. Each epoch has had data from 04 January 2010 subtracted prior to imaging (thus no image is available for 04 Jan itself). Weighting is determined with a robust factor of 0.5 (an optimal compromise between natural and uniform weighting). These maps use ATCA-CABB data that has undergone both normal calibration and phase calibration routines prior to final deconvolution. Maps were cleaned using multi-frequency clean subroutines. Contour lines are at -2.8, 2.8, 4, 5.6, 8, 11, 16, 23, 32, 45, 64, 90$\times$rms noise of each epoch (matching the scheme used in Tudose *et al.* 2008, rms listed in the top left corner of each panel). The larger cross marks the position of the core and the smaller crosses those of components NW and SE. Beam sizes are approximately 4 arcsec^2. Artefacts caused by flare variability remain in images C10 through C13. It should be noted that C10 only includes data prior to the initial flare rise in order to reduce severity of artefacts.

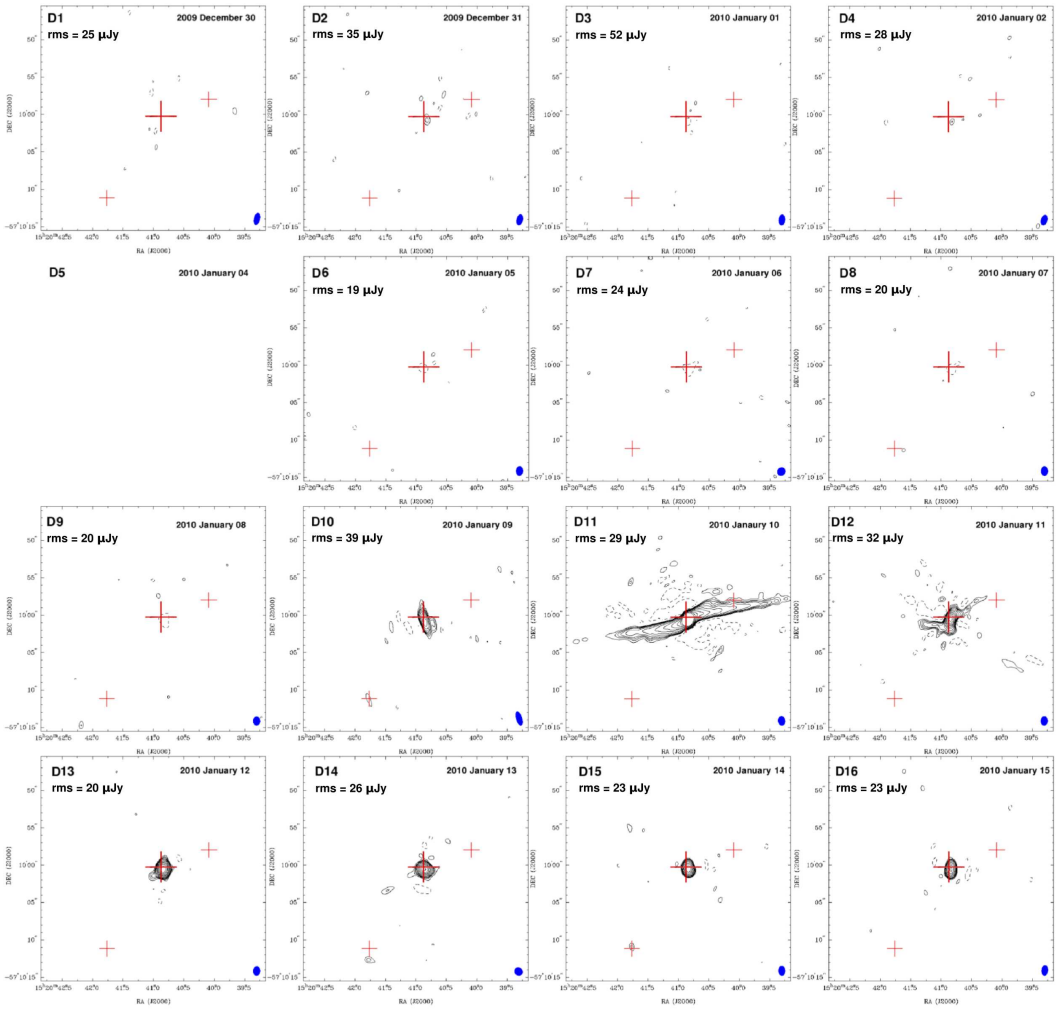

Figure 5.10: Difference radio maps of Circinus X-1 at 9 GHz. Each epoch has had data from 04 January 2010 subtracted prior to imaging (thus no image is available for 04 Jan itself). Weighting is determined with a robust factor of 0.5 (an optimal compromise between natural and uniform weighting). These maps use ATCA-CABB data that has undergone both normal calibration and phase calibration routines prior to final deconvolution. Maps were cleaned using multi-frequency clean subroutines. Contour lines are at -2.8, 2.8, 4, 5.6, 8, 11, 16, 23, 32, 45, 64, 90$\times$rms noise of each epoch (matching the scheme used in Tudose *et al.* 2008, rms listed in the top left corner of each panel). The larger cross marks the position of the core and the smaller crosses those of components NW and SE. Beam sizes are approximately 1.5 arcsec^2. Artefacts caused by flare variability remain in images D10 through D12. It should be noted that D10 only includes data prior to the initial flare rise in order to reduce severity of artefacts.

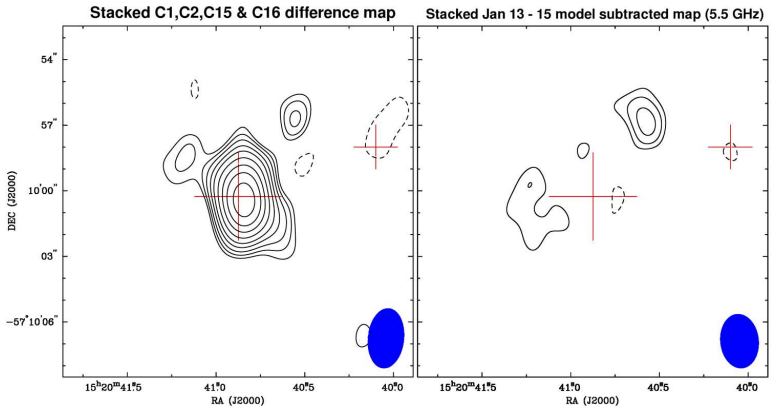

Figure 5.11: Circinus X-1 stacked difference and model subtracted image maps. The left hand image combines the difference data from 2009 Dec 30, Dec 31, 2010 Jan 14 and Jan 15 (panels C1, C2, C15 and C16 from Figure 5.9: folded orbital phase ϕ = 0.28 - 0.46, beam size = 2.75 $\times$ 1.66 arcsec^2), the right hand image combines model subtracted data from 2010 Jan 13, 14 and 15 (orbital phase ϕ = 0.22 - 0.36, beam size = 2.48 $\times$ 1.76 arcsec^2). Both images use contour levels of -2.8, 2.8, 4, 5.6, 8, 11, 16, 23, 32, 45, 64, 90 times r.m.s. of 13 μJy. The larger cross marks the core, the smaller the position of component NW. An additional (i.e. unrelated to component NW) north-western structure can be seen in both images.

the days following the flare event remain heavily affected by artefacts; however, the asymmetrical nature of some of the structures (e.g. C13) suggest that there may be some real emission towards the south-east which may persist over several days. A similar argument can be made for the structures in panel C14, and the emission component that is visible to the north-west persists in panels C15 and C16. Furthermore, its position is concordant with the structure nearest the core in panels C1 and C2. The left hand panel of Figure 5.11 shows the results of stacking data used for panels C1, C2, C15 and C16, in which one can clearly see the northern structure nearer the core. Though not shown in the image, the farther north-western structure from panels C1 and C2 was also visible, though only at a 3 σ level. The extension visible to the north-east of the core appears to be caused by a streak artefact of unknown origin running through the core perpendicular to the jet axis (further effects of this artefact can be seen to the core's south-west).

5.4.2.3 Model subtracted images

It is clear that the flare artefacts severely hinder the ability to interpret the images after Jan 09, with core variation continuing and causing elongation along an axis with an unfortunate similarity to that of the jet. Thus a second level of subtraction would be required to reveal structural variation in such images. This involves mod-

elling the flaring core, based on the variation observed from measured light curves of each observation.

The process is as follows: each day's difference data set (i.e. with the Jan 04 reference data already removed) is divided into smaller sections, approximately 5 minutes in length (>100 data segments per observation). These individual sections are then converted into images and a measurement of the maximum flux density is taken from a 30 arcsec^2 region centred on the core. It should be noted that no effort to clean the image is made prior to the measurement, as it results in no improvement to image fidelity for such small data chunks. Though these snapshot images are of very poor quality, the light curves that result from this process are very similar to that of Figure 5.2 suggesting relatively accurate representation of behaviour at the core. The flux density measurement of each 5 minute segment is then used to construct a 5 minute data model, representing a single source of the same flux density at the position (based on measurements made from the unaltered data) of Cir X-1's core. Using the MIRIAD task UVSUB, one can then create new 5 minute data segments with the models subtracted, and recombine them into a complete observation (UVCAT). This final data set can then be inverted and cleaned as normal, eliminating as effectively as possible the recorded core behaviour. A large amount of this process was scripted in python using modified components of code from VaSpecSim (see Chapter 3 and the appendices).

The results of my application of this process on the post flare images are shown in Figure 5.12. Unsurprisingly, the subtractions were least successful on the Jan 10 images owing to the rapid rates of variation soon after the flare. Many of the structures around the location of the core are likely to be remnants of flare artefacts or the result of phase errors, inferred from the symmetrical or anti-symmetrical layout of strong positive and negative counterparts (easily visible in the 5.5 GHz Jan 10 image). Though the Jan 11 images are cleaner by comparison, some care is still necessary as the day's observations still involves high flux densities and decay rates. Most visible structures can be disregarded, owing to their numbers, negative counterparts, and almost random arrangement.

The final four model subtracted images are also shown in Figure 5.6 (far right column) together with their unaltered counterparts (described in section 5.4.2.1), in an effort to summarise the behaviour that occurs within Cir X-1. Jan 12 images (row A of Fig 5.6) show some emission to the south-east (i.e. along the jet axis) of similar size and position at both frequencies, leading to jet velocity estimates with $\Gamma > 10$ (assuming a relation with the flare of Jan 9th). However, there is also a negative component at 5.5 GHz whose position would suggest these may be an

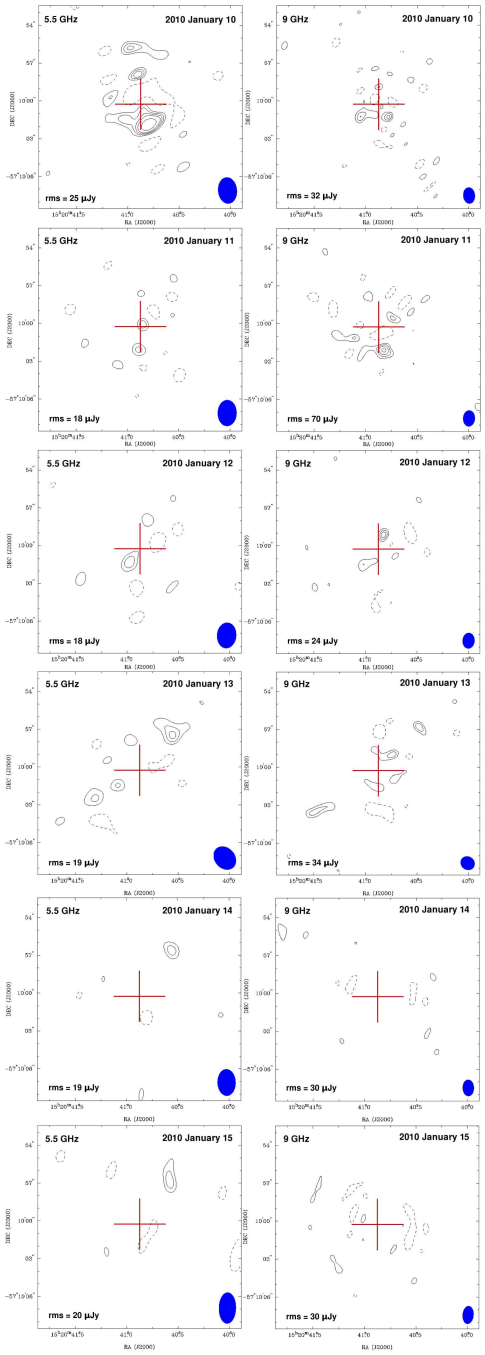

Figure 5.12: Model subtracted radio maps of Circinus X-1. These images have had both the Jan 04 data subtracted, and a model point source subtracted that is centred at the location of Cir X-1's core and whose behaviour is modelled on the core light curves. They cover the flare event and subsequent decay. These maps use ATCA-CABB data that has undergone both normal calibration and phase calibration routines prior to final deconvolution. Maps were cleaned using multi-frequency clean subroutines. Contour lines are at -2.8, 2.8, 4, 5.6, 8, 11, 16, 23, 32, 45, 64, 90$\times$rms noise of each epoch (listed on the bottom left of each panel), however the rms is difficult to measure accurately as image noise tends to increase towards Cir X-1's position, and thus values used are higher than those measured in the surrounding regions of maps. Beam sizes are approximately 4 arcsec^2.

artefact pair. One can also observe strong emission to the north at 9 GHz, but the lack of a 5.5 GHz counterpart means I cannot confirm if it is real. The images from Jan 13 onwards show the strongest evidence of real structural variation beyond the core. These were the images that showed the lowest level of artefacts caused by the flare, prior to core subtraction, and thus are the most likely to display real emission structures rather than residual artefacts.

This leaves three components of interest in the 5.5 GHz images, two of which have possible counterparts at 9 GHz that do not appear to be related to side-lobes. The component nearest the core at 5.5 GHz on Jan 13 may be related to the negative component on the opposite side of the core, and without a confirmed counterpart at 9 GHz is best ignored. The remaining two components reside on opposite sides of the core and have possible 9 GHz counterparts appearing a significant distance further from the core (1-2"), yet they remain along a very similar axis. The northern component at 15:20:40.56 -57:09:56.7 (J2000) does appear to have a negative counterpart at a similar distance to the south-east of the core, and though it is slightly smaller and fainter, may still indicate effects of flare artefacts. However, support for the reality of this northern structure is that it persists at 5.5 GHz through Jan 14 and 15 images, and the position also coincides with that of the near-core northern component observed in the Dec 30 and 31 difference maps (see section 5.4.2.2). Both positive structures are not seen again at 9 GHz, nor does any significant new structure arise over the final 2 days. The result of stacking these final three days of model subtracted data is shown in the right panel of Figure 5.11, where one can see strong evidence of the northern component ($> 5\sigma$) and weaker signs of south-eastern residual structure ($\sim 3\sigma$).

Flux density variation of the northern component over the 3 days is well within the noise levels of the individual images. The emission lies approximately 3 - 5" away from the core, and appeared $\sim$ 3 days after the peak of the flare event. Assuming it is real, then it is unlikely to be related to the observed flare event, as that would require Lorentz factors of $\Gamma > 35$. However, by considering the possibility of an earlier flare ($16.6 \times n + 4$ days delay) being the source of re-energisation, then the calculated velocities may be significantly reduced (e.g. $\sim$ 20 day delay gives Γ to be 6 - 10).

As a test of my methods I investigated whether the process of removing the generated model from the data could inadvertently create solitary structures such as those in the images. I created model data with a single point source at Cir X-1's location (much like the subtraction model) but included noise within the flux density values with a combination of sinusoidal variation and random number components.

The resulting models strongly resembled a single day's high resolution light curve of Cir X-1's core (with flux density levels set to be similar to those seen during flare decay). I proceeded to image this model data and run it through the same light curve extraction and model subtraction routines used on the real data (with varying interval sizes). The result was a near complete removal of the source each time (a weak 'core' would remain with flux density $< 1\%$ of original levels), with flux densities of residual surrounding structure in the image reaching only a few μJy (i.e. $<$ r.m.s. noise levels in standard images) and distributed in a spoke like manner rather than in clusters.

5.4.3 Large scale structure

By combining the observations from the first 9 days I was able to create a deep 5.5 GHz radio map of Circinus X-1 and its surrounding region (Figure 5.13, left), with rms levels in the majority of the map not rising above 3.5 μJy beam^{-1}. These low noise levels make it possible to see detailed filament-like structures outlining the known layout of the Cir X-1's jet powered nebula, and what appear to be faint (compared to the core) less diffuse regions within the larger nebula structure. I compare this map with contours created from the data used in Tudose *et al.* 2006 (Figure 5.13, right). The red contours are those from 1.4 GHz observations with the ATCA in a short baseline configuration (EW214: minimum baseline of 31m, maximum of 4500m) taken in September 2004, showing the outer regions of the nebula surrounding Cir X-1. The black ATCA contours (configuration 1.5A: minimum baseline of 153m, maximum of 4469m) are from August 2001 and are also at 1.4 GHz, but the observations did not include baselines as short of those from Sep 2004, and thus one observes slightly more detailed structure closer to the core. The asymmetry of the structure can be observed on both these scales: an extension of the nebula towards the south. This has been mostly attributed to projection effects as a result of the approaching south-eastern jet, and receding north-western jet. However, there are also bends in the jets (evident in the black contours) which may be attributed to precession of the jet (much like SS433, see Section 2.2.3.6) or interactions with higher-density material along the jet's path which cause the flow to deviate.

My deep map shows numerous filament structures outlining the nebula's edges, as well as multiple components along the jet axis, including the radio counterparts to the X-ray synchrotron 'caps' reported in Sell *et al.* (2010). Though they form circular outlines, the filaments are unlikely to be the residuals of side lobe structure,

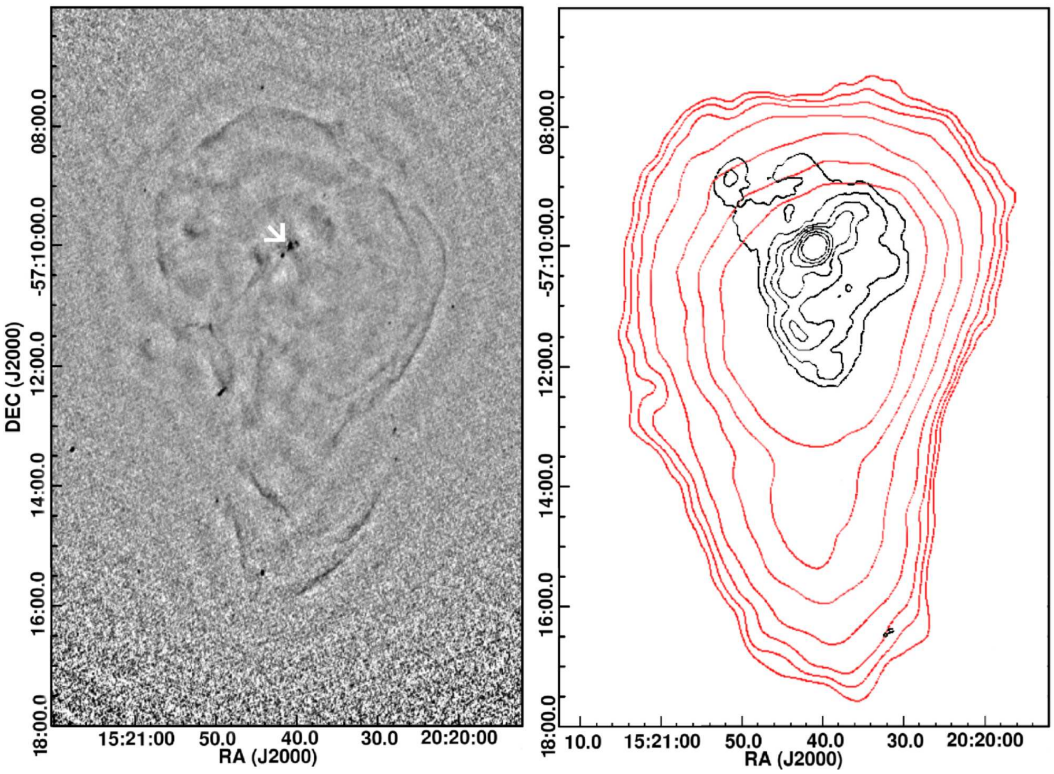

Figure 5.13: Large scale radio images of Circinus X-1's jet powered nebula. The left hand 5.5 GHz radio map (beam size = 2.3 $\times$ 1.6 arcsec^2) was produced using combined data from the first 9 days of observations (i.e. the same data as the bottom left panel of Figure 5.6, but using complete *uv* distance coverage) and shows the detailed filament like structures outlining the edges of the nebula, with Cir X-1's core marked by the white arrow. The increase in noise towards the bottom of the image is a result of primary beam correction - i.e. approaching the beam perimeter. To the right I show a contour map of the same region using 1.4 GHz data from Aug 2001 (black, contour levels of 1, 1.4, 2, 2.8, 4, 5.6 and 8 times the rms noise of 2 mJy beam^{-1}) and Sep 2004 (red, contour levels of 1, 4, 8, 16, 32, 48 and 64 times the rms noise of 0.7 mJy beam^{-1}), the difference between the two being due to array configuration (i.e. shorter baselines were available during the Sep 2004 observations). This is the same data use to produce figures 2 (middle) and 3 in Tudose *et al.* 2006, and illustrates how much of the large scale structure appears beyond the region outlined by the filaments in my maps.

based on comparison to the corresponding beam patterns; they are too large, intense and sharply defined (even prior to the application of cleaning algorithms). Much of the more diffuse structure within the confines of the filaments is affected by beam distortions, causing large scale 'lumpiness' or corrugation, though several sharper intense regions stand out (such as the caps). The majority of the outlying filaments have peak flux densities of $\sim 2\times10^{-5}$ Jy beam^{-1}, but since the 6A ATCA array configuration uses particularly long baselines (especially in comparison to those used in Tudose *et al.* 2006), it is likely much of the emission from these regions has been resolved out (the largest angular scale to which my map is sensitive is $\sim$ 40 arcsec). This is also indicated via comparison of the flux levels with those measured at the same locations at lower frequencies, yielding an extremely steep spectrum ($\alpha < -4$). It should also be noted that though I have corrected for the response of the primary beam, the distance of the filaments from the pointing centre of the observation ($\sim$ 2.5 arcmin for the 'ring' filaments, 4 - 5 arcmin for those to the south) may still add to the uncertainty in flux density measurements as a result of uncertainties in the primary beam model itself.

None of the filaments are detected at 9 GHz or in polarisation maps, and estimates for individual structures show no region has minimum brightness temperature $T_{b-min} > 0.5\times10^4$ K, meaning either synchrotron or bremsstrahlung processes can be the cause. Nonetheless, these filaments clearly define an outline similar in form to that of the nebula, including the southern extension, though smaller in comparison to the full low frequency emission (red contours). Assuming the validity of a bremsstrahlung scenario, a comparison may be made with the nebula of Cygnus X-1 where one can observe the thermal bremsstrahlung emission from ionised gas produced behind the bow shock of a jet impacting on the ISM, for which temperature estimates give T $\sim 10^4$ K (see Section 2.2.3.5 and Gallo *et al.* 2005).

It is also possible that these filaments are boundaries related to the synchrotron emission bubble formed by the outflows and being reheated by continuing jet emission pointed towards us (see Sections 2.2.3.5 and 2.2.3.6). Tudose *et al.* 2006 reported flattening of the nebula's spectrum towards its edges, in particular in the north east and south west extremes: regions in my maps where I observe longer, unbroken emission filaments. The interpretation that followed was that the sites were regions of increased interaction between the ISM and accelerated particles. The filament like morphology may simply be a result of density fluctuations in the ISM, or perhaps, if the jet does precess or the large scale structure of the outflows has changed over time, they may be regions of jet-ISM interactions much like those observed in nebula W50 (Dubner *et al.* 1998). It could easily be imagined that, if

this nebula resembles a layout much like that of SS433 and W50 (see Figure 2.8, left panel) but orientated towards us (inferred from the jet inclination: Fender *et al.* 2004), then, as suggested by the contour maps, the diffuse emission seen at lower frequencies would appear to extend beyond brighter zones reheated by the jets due to the 'middle' of the nebula being wider than the tapered ends (through which one views the system).

5.5 Discussion

Previous observations of Cir X-1 taken during its brighter past have shown most of the system's activity and structural variation occur south-east of the core, which is likely to be the direction of the approaching jet. Indeed there have been several epochs during which the emission of a distinct south-eastern component only ~ 2 - 3 arc-seconds from the core has outshone the core itself (2001 May 29 - Fender *et al.* 2004, Tudose *et al.* 2008). There has also been some evidence of the receding jet (1996 July 02, 2000 October 25); however, the receding jet never appeared without activity from its counterpart, with the latter remaining dominant in comparison.

Initial inspection of images such as the bottom left panel of Figure 5.6 would suggest the current trend is similar, but with the core being significantly diminished. There is a bright component (component SE) ~ 13 arc-seconds to the south east of the core and a fainter counterpart to the north/ north west (component NW), slightly nearer but still over 10 arc-seconds away. However the nature of component SE can be easily scrutinised. Firstly, as well as being unresolved in the radio, it was also found by Sell *et al.* (2010) to be unresolved in X-rays and coincident with a bright IR point source. Subsequently, Sell *et al.* (2010) were unable to statistically rule out the possibility that the component is in fact a background source. If the emission was indeed related to the jet output from Cir X-1, one might expect some level of long term variation in flux density (to match Cir X-1's core decline over the past decades) or position (as a previous ejection event). Archival data not only revealed it to exist in the same position during several epochs over the past decade or more, but that its flux density has not varied significantly beyond the range of values measured from this chapter's observations, and has not followed any particular trend. With this additional information in mind it becomes easier to disregard component SE as jet related emission, but it will require further observation before a firm conclusion can be made.

The northern counterpart (component NW) remains a strong candidate for jet

emission. The source is resolved in both radio and X-rays, as per Sell *et al.* 2010, and was not found to have an IR counterpart. It remains undetected in any past observation stacks (bearing in mind if at the same low flux density it would fall below the noise), though there has been evidence of intermittent structure near its location in individual epochs detailed in Tudose *et al.* (2008); specifically 2000 October 20/21 and 2001 May 27.

The strongest indication of structural variation in my maps originates in regions to the north west of the core. The structure visible in model subtracted images from Jan 13 through 15 (and possibly Dec 30, 31) lies at an angle similar to that of the jet and thus may be related to the receding outflow, though one observes comparatively little change in the south. This behaviour goes against what is expected of the system based on past observations, with the greatest level of activity coming from a brighter approaching jet located to the south. This difference could simply be due to intrinsic asymmetry of interacting material around the binary, but jet precession could be an alternative explanation. For a small angle to the line of sight, as proposed, even a small degree of precession could account for a reversal of orientation with the approaching jet appearing in the north western quadrant. It should be noted that no significant change in jet axis has been observed in radio images that have been made of the core over the last decade and a half prior to our observations (Tudose *et al.* 2008), suggesting that if any change in orientation has occurred, it was sudden and recent. Nonetheless, a near 180 degree reversal of the projection of the jet axis, though plausible, remains suspicious, and suggests that the asymmetric flow theory is more likely.

A scenario with a kinked or deviated outflow is also possible (see section 2.2.3.7). Though one already sees evidence of the south eastern outflow being curved on arcminute scales (Tudose *et al.* 2006), there may be bends earlier in the outflow that result in increased beaming effects along portions of the jet further inclined along our line of sight. Such could be the case for component SE, providing explanations for its stable position and compact structure if it were related to a persistent jet kink. Again, if dealing with a generally low jet inclination, significant deviations of the outflow direction could cause emission activity within the approaching jet to appear either side of core via projection along our line of sight.

The point source subtraction images shown in Figure 5.7 also provide evidence for precession, showing a possible axis that is almost a north-south orientation. Additionally, based on the greater degree of symmetry in residual structure around the core, my images suggest a higher inclination for the jets than that observed in the past. This latter conclusion might be seen to weaken the justification for using low

jet inclination as a possible explanation for the observed behaviour in my images. However, just as observed in the outflows of SS433, the inclination of a flow can vary as one moves further from the core, in this case as a result of either precession, flow deviation, or a combination of both.

5.6 Summary

I have presented the results from ATCA-CABB radio observations of a complete 17 day orbit of Circinus X-1 at 5.5 and 9 GHz. Successfully capturing the rise and decay of a periastron flare event from the system's core and mapping the structure of the system on each day during this period, during which it transpires Cir X-1 was in an historically faint radio state. Prior to the flare event the system was very stable at $\sim$ mJy flux densities. After the flare there is some evidence for significant spatially resolved changes in the images on arcsec scales in carefully modelled and core-subtracted images. The strongest sign of structural variation appears north west of the core, in the region previously associated with the direction of the receding jet. I interpret this unexpected behaviour as either an indication of asymmetry in the surrounding, or alternatively, of a change in outflow direction via precession or curvature of the jets. A change in jet orientation is also supported by comparisons of residual structure around the core in my maps to that of maps made from observations taken in 2002, showing an apparent 'rotation' of the arcsec-scale resolved core structure.

Deep radio maps of the area around the core show several persistent sources likely associated with the jets, and others that are counterparts to previously reported X-ray shocks. Large scale mapping of Circinus X-1's surrounding region reveals filaments outlining the known structure of the system's jet powered nebula, which, assuming a low inclination for the system's jets (as has been previously claimed), suggest the system's configuration may resemble that of SS 433 and W50, but viewed longitudinally.

Note that the results in this chapter are strongly linked to the work that will be discussed in Chapter 6, in which Circinus X-1's structure is imaged at mm wavelengths. I make a large amount of comparisons to the figures in this chapter and build upon the jet precession and deviation scenarios that I have proposed here.

Wrong, wrong, wrong, wrong, wrong, wrong, wrong, wrong. You're wrong, you're wrong, you're wrong.

DR. PERRY COX, SCRUBS

6

Millimetre observations of NSXRBs Circinus X-1 & Scorpius X-1

In this chapter I present the results of the first successful millimetre (33 GHz and 35 GHz) observations of the neutron star X-ray binary Circinus X-1 (introduced in the previous chapter), as well as centimetre (5.5 GHz and 9 GHz) and the first successful millimetre (33 GHz and 35 GHz) observations of the neutron star X-ray binary Scorpius X-1 (Sco X-1) using the Australia Telescope Compact Array. The goal of these observations was to detect a neutron star XRB for the first time at mm wavelengths, and subsequently allow probing of the radio structure at higher resolutions than had yet been reached using ATCA. Since we know nothing of NSXRBs at mm wavelengths these observations would help us to better understand the extent of the jet spectrum (see Section 2.2.3.3), and the associated energies also correspond to jet components closer to the compact object and can thus aid in deciphering the structure of outflows when compared to imaging in other radio bands.

In the case of Circinus X-1 I see strong evidence for a periastron flare during the first observation which proceeds to decline over the following four days. My analysis of spectral variability suggests that optically thin synchrotron emission dominates the unresolved core, for which I discuss possible scenarios. I proceed to present radio imaging of Cir X-1, which includes evidence of sub-arcsecond

jet structure, and variability at distances that suggest recent shock re-energisation (though I am unable to associate this with the observed flare). I go on to suggest that, if the emission is powered by an unseen outflow, then a phase delay exists between flare onset and subsequent brightening of nearby components, with flows reaching mildly relativistic velocities. Given resolved structure positions in my images, by comparing to past observations of Cir X-1 (including those of Chapter 5) I find evidence that jet direction may vary with distance from the core, or the source's precession parameters have changed. I detect Scorpius X-1 in all epochs and frequencies, with indications of rapid variability ($\sim$ hour timescales) in the mm lightcurves, as well as steady rise observed at 5.5 and 9 GHz. My images show Scorpius X-1 to be an unresolved point source at both cm and mm wavelengths. Spectral index estimates for Sco X-1 show similar behaviour as that described for Cir X-1, implying similar scenarios. All data reduction was carried out by myself, with the results of the Circinus X-1 imaging published in Calvelo *et al.* (2012b), and the spectral analysis, together with the full analysis of Sco X-1 to be included in a future publication.

6.1 Introduction

As has been discussed earlier in the thesis, neutron star X-ray binaries (NSXRBs) present a key opportunity to study radio jet behaviour in the absence of a black hole engine, and the effects this can have on jet formation. Unfortunately, jets from low luminosity NSXRBS are more difficult to study as the source group tends to be more radio quiet than their black hole counterparts, making their radio emission more challenging to detect. It should come as no surprise then that the majority of information about relativistic jets comes from studies of BH systems and that the understanding of NS jets is more incomplete by comparison.

One area in which there is a distinct lack of information is the mm emission of NSXRBs. The radio faintness of NSXRBs coupled with the increased complication of mm observations compared to longer wavelength observations (i.e. more significant phase errors due to increased refraction by atmospheric water vapour) makes it unfavourable to attempt such investigations; however, with recent advancement in radio telescope arrays and correct target selection this can eventually be overcome. In the case of BHXRB jet sources, the mm regime appears to continue the flat spectrum observed at longer wavelengths (Fender *et al.* 2000), implying the dominant contribution comes from optically thick synchrotron radiation from the compact jet

(See section 2.2.3.3). A follow up investigation showed that this mm emission became quenched during the soft state (Tigelaar *et al.* 2004): behaviour which is also seen in the radio band for many BHXRBs and is expected if the emission originates in the jets (believed to turn off or become suppressed in the soft state: Fender, Belloni & Gallo 2004). Though we cannot describe the behaviour of NSXRB mm emission as we lack actual mm detections, Migliari *et al.* (2010) show evidence that the spectrum of NSXRB 4U 0614+091 remains flat and thus consistent with emission from an optically thick compact jet between the radio and and mid-IR, at which point the spectrum breaks into a steeper power law. This value of 1.25×10^{13} Hz $< \nu_{break} < 3.71 \times 10^{13}$ Hz is interpreted as the frequency above which the synchrotron emission becomes optically thin ($\sim$ 10 times lower than break frequencies found for BHXRB GX 339-4: Corbel & Fender 2002).

Spectral index estimates for Cir X-1 (see Section 5.2 for details on the source and its history) based on 5.5 and 9 GHz observations with the Australia Telescope Compact Array Broadband Backend (ATCA-CABB) in Calvelo *et al.* (2012a: see Chapter 5) indicated the system would likely be detectable at mm wavelengths. Sco X-1 was also a viable target for mm observations (plus sharing similarities with Cir X-1) and was the second target for observations alongside Cir X-1. As both Cir X-1 and Sco X-1 are known for rapidly evolving structure tied to their jet outflows, multiple mm observations over several days were planned to provide the chance of capturing changes in structure following core variability such as moving ejecta, or downstream re-energisation events.

6.1.1 Scorpius X-1 background

Scorpius X-1 (Sco X-1) was the first X-ray source to be discovered beyond our solar system (Giaconni *et al.* 1962). An optical counterpart was discovered soon after, identifying the system as a binary star with a period of 0.787 days (derived from spectroscopic analysis by Cowley & Crampton 1975). X-ray studies showed it to be highly variable (with timing independent of the binary period) and indicated the companion was likely to be a neutron star (Davidson, Pacini, & Salpeter 1971). Radio emission was also detected from the system (Andrew & Purton 1968) and later resolved into three distinct sources in the region, one at the position of the X-ray source (thus associated with Sco X-1), and the other two at roughly symmetrical positions of about 1 arc-min from Sco X-1 (Hjellming & Wade 1971). The central radio source was found to be highly variable on timescales as short as hours, as well as displaying sparser, more intense flare events where flux densities varied more

strongly but over longer periods around the scale of day (e.g. the greatest variability reported in Hjellming & Wade 1971 was a factor of 60 over 25 hours). The variation in Sco X-1's X-ray spectrum provided some of the first evidence of differences in NSXRB behaviours that would eventually lead to the distinctions between atoll and Z source systems, with Sco X-1 falling into the latter category (Hasinger 1988; Hasinger & van der Klis 1989).

Though the position and non-thermal spectra of the two nearby radio objects suggested Sco X-1 could be a microquasar (i.e. the two objects being radio lobes), Very Large Array (VLA) studies showed that while Sco X-1 had a detectable proper motion, the 'lobes' were in fact stationary and thus unrelated background sources (Fomalont & Geldzahler 1991). Very Long Baseline Array (VLBA) studies later revealed sub-hour variability in the structure of and around Sco X-1, with ejecta moving away from the central source at relativistic velocities (0.3c: Geldzahler, Fomalont & Bradshaw 1999) establishing Sco X-1 as a radio jet source. Very Long Baseline Interferometry (VLBI) was used to study these jets in great detail, indicating variable ejecta velocities (0.32c - 0.57c) as well as evidence for re-energisation of lobe pairs by unseen jet flows travelling at $>$0.95c (Fomalont, Geldzahler & Bradshaw 2001).

6.2 Observations & data reduction

Observations were carried out on 2010 Nov 18, 19, 20, 21, 22 and 23 using the ATCA-CABB in 6A configuration (minimum baseline of 337m, maximum of 5939m), alternating between Sco X-1 and Cir X-1 each day with the first observation targeting Sco X-1 (further details listed in Table 6.1). For Cir X-1, 33 and 35 GHz observations, PKS B1253-055, PKS B1934-638, and PKS B1511-55 were used as the bandpass, flux and phase calibrators respectively. In the case of Sco X-1, 5.5 and 9 GHz observations, PKS B1934-638 was used for bandpass and gain calibration, and PKS B1555-140 for phase calibration. For Sco X-1, 33 and 35 GHz observations, PKS B1253-055, PKS B1934-638, and PKS B1622-253 were used as the bandpass, flux and phase calibrators respectively. Between 6 and 9 hours of time was spent on-source each day (predicted rms noise for average weather conditions of 26 - 34 μJy at 33 GHz, 28 - 36 μJy at 35 GHz, 12 - 15 at 5.5 Ghz and 17 - 21 at 9 GHz). All data and image processing was carried out in MIRIAD.

Table 6.1: Cir X-1 and Sco X-1 radio cm and mm observation details and measurements. For each observation the table lists the source name, date, observing frequency, start Modified Julian Day (MJD, on-source), total on-source time before and after flagging, MJD end time after flagging, measured peak flux density, and noise levels of the image from which the measurement was taken. To ease comparison between observations of the same source, I first list observations by source and then chronologically.

Source name	Date (UT)	Frequency (GHz)	MJD start	On-source time (h) [post-flag value (h)]	Post-flag MJD end	Peak S_ν [range] (mJy beam^{-1})	rms (mJy beam^{-1})
Sco X-1	2010 Nov 18	33	55518.913	1.66 [1.66]	55519.005	1.02 ± 0.31 [N/A]	0.23
Sco X-1	2010 Nov 18	35	55518.913	1.66 [1.66]	55519.005	1.24 ± 0.38 [N/A]	0.30
Sco X-1	2010 Nov 19	5.5	55519.019	5.42 [5.42]	55519.316	7.27 ± 0.79 [4.25:14.57]	0.22
Sco X-1	2010 Nov 19	9	55519.019	5.42 [5.42]	55519.316	3.11 ± 0.44 [1.26:7.52]	0.19
Sco X-1	2010 Nov 20	33	55520.896	7.45 [7.45]	55521.332	2.13 ± 0.36 [1.41:4.03]	0.19
Sco X-1	2010 Nov 20	35	55520.896	7.45 [7.45]	55521.332	1.88 ± 0.30 [1.27:3.44]	0.17
Sco X-1	2010 Nov 22	33	55522.899	6.58 [3.85]	55523.125	0.29 ± 0.06 [N/A]	0.05
Sco X-1	2010 Nov 22	35	55522.899	6.58 [3.85]	55523.125	0.34 ± 0.07 [N/A]	0.05
Cir X-1	2010 Nov 19	33	55519.841	8.86 [6.08]	55520.375	7.88 ± 0.93 [1.55:15.54]	1.58
Cir X-1	2010 Nov 19	35	55519.841	8.86 [6.08]	55520.375	7.51 ± 0.83 [1.55:14.94]	1.37
Cir X-1	2010 Nov 21	33	55521.836	8.42 [3.31]	55522.021	2.28 ± 0.21 [2.25:2.39]	0.13
Cir X-1	2010 Nov 21	35	55521.836	8.42 [3.31]	55522.021	1.38 ± 0.20 [1.12:1.78]	0.19
Cir X-1	2010 Nov 23	33	55523.804	8.14 [5.47]	55524.240	0.30 ± 0.05 [0.31:0.54]	0.07
Cir X-1	2010 Nov 23	35	55523.804	8.14 [5.47]	55524.240	0.28 ± 0.04 [0.42:0.48]	0.07

Since millimetre observations are more susceptible to atmospheric effects than longer wavelengths phase stability had to be closely monitored via the ATCA "seeing" monitor (output gives rms path length fluctuations in microns; see Middelberg *et al.* 2006). These effects increased as the runs progressed, due to rising temperature and humidity. The intensity of the rms fluctuations was so severe during the first Sco X-1 observation that a switch to cm wavelengths (where the detrimental effects of phase instability are reduced) was recommended after only 1.7 hours of mm observing, though the following mm observing epochs were deemed stable enough to continue uninterrupted. A switch to cm wavelengths was never warranted during the Cir X-1 observations (path length rms tended to remain < 700 microns), however, once analysis began it was found that large segments of each mm observation did not have sufficient phase stability to be reliable for imaging.

Due to the gradual decay in phase stability, it was necessary to define an adequate level of flagging that would improve final image fidelity. By dividing the observations into segments, and comparing the resulting maps to path length rms values VARPLT, I found that segments with low path length rms showed the expected point source at expected source locations. Higher rms segments corresponded with increasing noise towards source positions, with no single distinguishable point source visible. For Cir X-1, segments whose majority of time was spent with rms higher than 400 microns failed to produce the defined structure, which I proceeded to take as the Cir X-1 flagging limit. In the case of Sco X-1, a flagging based on a limit of 600 microns for the rms fluctuations allowed me to keep enough data to produce definable point sources in images. This limit covered the entire first and second observation runs, and almost two thirds of the final night, where rms fluctuations exceed the limit for the final 2.2 hours of observing. This limit resulted in a smaller flagged data fraction in comparison to the methods employed for Cir X-1. Post-flagging on-source time for both systems is shown in column 5 of Table 6.1.

6.3 Analysis & results

When using the cm bands, images at each frequency (5.5 GHz and 9 GHz) can show significant differences in structure, and are thus worth describing individually. However, for mm observations the relative difference in the two bands is smaller (33 GHz and 35 GHz), as are the differences in images, and so it is far more useful to combine the two mm bands during the inverse Fourier transform of calibrated *uv* data (i.e. include the visibility files of both frequencies with multi-frequency

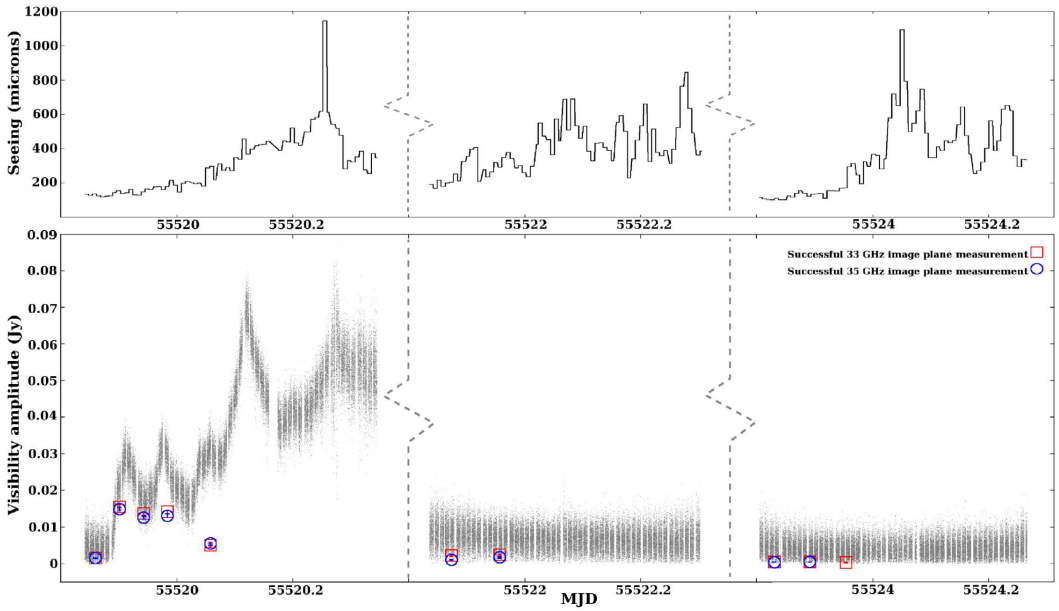

Figure 6.1: Circinus X-1 mm light curves and rms path length fluctuations. ATCA-CABB rms path length fluctuations (top), *uv* amplitude (grey points), and image plane fit plots at 33 and 35 GHz (red and blue points respectively; three observations over 5 days). Measurement errors (1σ) are included in the image plane values but are too small to appear clearly.

synthesis in MIRIAD's INVERT step) to produce maps with higher signal to noise ratios. I do however, discuss the differences in flux density between images made in each band in order to gain insight into the spectra of the two systems.

6.3.1 Circinus X-1

6.3.1.1 Light curves

Declining phase stability prevented me from producing a complete set of image plane measurements for each day's full observation. However, having produced radio maps for each day (see Section 6.3.1.3), it is apparent that Cir X-1 is the only visible source, and thus I may estimate flux density variability from *uv* data amplitude plots. However, I am still unable to distinguish whether variation arises from the 'core' of Cir X-1 (designated as a point source situated at Cir X-1's well established co-ordinates: Tudose *et al.* 2008) or surrounding structure, since the variability in amplitudes can also arise from structural variations around the source (i.e. the result radio imaging's inability to distinguish between the two possibilities, as was described in Section 3.4).

Figure 6.1 shows each observation's visibility amplitudes (combination of both frequencies), with available image plane fits at each frequency over-plotted. Note

that the visibility amplitudes are from all available data prior to flagging, and do not represent the quantity of data that was used to make images and the source measurements listed in Table 6.1. The most noticeable feature is the multi-peaked rise in the first run's amplitudes, which in comparison to the two subsequent runs, suggests flaring activity. A rise is also noticeable in the image plane fits, though it only includes a few points. The multiple peaks may be attributed to real variations from Cir X-1 (multi-peaked flares have been seen before: Thomas *et al.* 1978, Tudose *et al.* 2008), or resolved structure. In this case the radio data alone are insufficient to eliminate either possibility. My inability to produce high fidelity maps for the time periods of highest amplitude means I cannot accurately measure the peak flux density reached by this flare; however, if I assume the relative difference between mean average amplitude and core image plane values remains similar throughout the first observation, I can crudely estimate that flux densities could have reached up to 50 mJy.

The second run's amplitudes indicate a post-flare decline, although the image plane measurements show a slight rise. With this observation data set having been the most affected by phase disturbances, it is difficult to support either claim. The final observation set shows the lowest amplitudes and flux densities, with the source having likely returned to quiescent levels. The flux densities measured from the core on each individual day are shown in Table 6.1 and also show the decline in flux across the three observations, supporting much of what is observed in the plots.

6.3.1.2 Spectrum

Before I discuss the variation in Cir X-1's spectrum, it should be stressed that I am dealing with high levels of uncertainty in the calculated α values ($\pm$ 0.5 to 1: equal or greater than the magnitude of each individual measurement), which make it difficult to draw strong conclusions from the analysis. While both 33 and 35 GHz data sets show the observed flare rise and subsequent lower activity, the available measurements suggest the spectrum remains negative for the majority of all observations, as is reflected in the "complete" epoch values listed in Table 6.1, and would indicate dominance of optically thin synchrotron emission. Available light curve measurements for early in the flare rise indicate a spectrum which starts off flatter, with α=0.0 $\pm$ 1.2 and -0.7 $\pm$ 0.9 for the first and second hours, and quickly steepens to α= -1.4 $\pm$ 0.8 for the following 2 hours. One might infer the presence of optically thick structure causing flattening of the spectrum during the flare rise, as would be expected from internal jet shocks or compact ejecta comprised of rela-

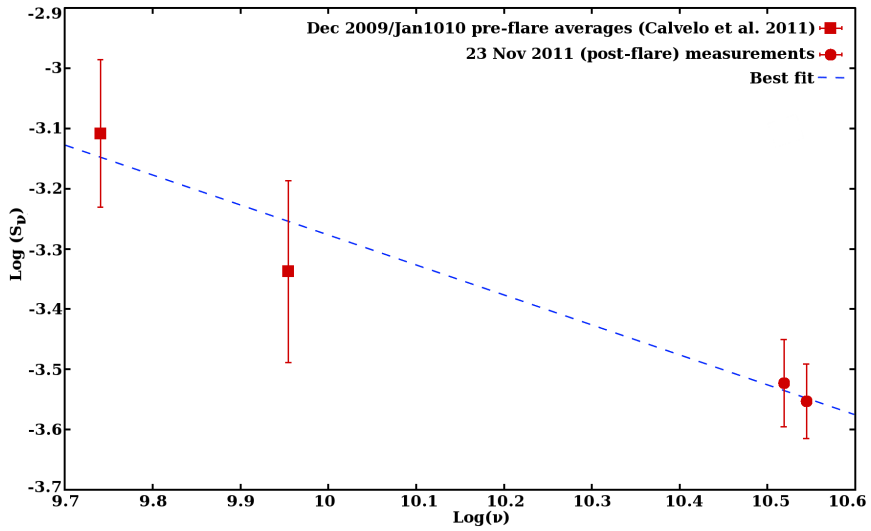

Figure 6.2: Circinus X-1 mm and cm radio spectral fits. The plotted 33 and 35 GHz points are taken from this chapter's ATCA-CABB third epoch observations and the 5.5 and 9 GHz points are averages calculated from 8 days of pre-flare ATCA-CABB observations taken in Dec 09/Jan 10 (see Chapter 6/Calvelo *et al.* 2012a for details). The fitted power law was calculated using the fit command in GNUPLOT. Since there was some variability in 5.5 and 9 GHz flux densities from day to day the sample standard deviation of the daily flux densities is used for error bar sizes. For the 33 and 35 GHz points, errors are taken directly from MIRIAD's point source fitting.

tivistic particles (see section 2.2.3.3, as well as Kaiser, Sunyaev & Spruit 2000 and van der Laan 1966). However, I must stress again the uncertainty involved in these estimates, as even with the Table 6.1 measurements I am dealing with $\Delta\alpha \sim \pm 3$. Furthermore, the values from the second epoch of Cir X-1 observations yield an unrealistically steep spectrum, indicating that point source fits have likely underestimated the 35 GHz flux density for this day (which had to be most heavily flagged as a result of high path length rms).

As mentioned in Section 6.2, at no time during our Circinus X-1 observations was it was necessary to switch to lower frequencies. While this did allow a larger amount of time observing at mm wavelengths towards better detect any structural variation in the source, it also meant the simultaneous exploration of the system's radio spectrum would be restricted. I am thus limited to use of past observations of the system at other frequencies to estimate the system's wideband spectral behaviour. As mentioned in the chapter's introduction, based on MAXI monitoring Cir X-1 currently resides in an historically faint epoch which is punctuated by brief (days to months) periods of higher activity. This chapter's mm observations took place in one of the "quiescent" intervals between these active periods, and as such it would be logical to make multi-wavelength comparisons with measurements from

observations taken during periods of similar activity.

The 5.5 and 9 GHz observations of Calvelo et al (2012a: covered in Chapter 5) include 8 days of observations prior to a flare and all taken during another faint period. Cir X-1 is detected on all 8 days, yielding mean daily flux densities of 0.78 $\pm$ 0.22 mJy at 5.5 GHz and 0.46 $\pm$ 0.16 mJy and 9 GHz, where the uncertainty arises from day to day variability. I plot these values together with my measurements from the third epoch of 33 and 35 GHz observations in Figure 6.2. I choose to use the values from the third epoch ($\sim$ 4 days following flare onset) as they are most likely to be closest to inter-flare flux density levels, though it is difficult to tell from available light curves whether the post flare decay has run its course. Based on the four available points I can make a fit with α = -0.50 $\pm$ 0.07 (fitted using the fit command in GNUPLOT). This value is shallower than the indices indicated by the separate frequency pairs of measurements which are both closer to α = -1, albeit, with higher associated uncertainties. Nonetheless, the fit agrees with the concept of optically thin synchrotron emission.

6.3.1.3 Imaging

Individual image contour plots of Cir X-1 for each day are shown in Figure 6.3. Cir X-1 is resolved in all three runs, at a position that shows no shift greater than the positional errors associated with each set of images. Jet like structure is visible along a single axis (north-west to south-east) which is similar to that seen in previous mapping of the source at cm wavelengths (images of Chapter 5, see also Tudose *et al.* 2006, 2008). Polarisation imaging was carried out for all frequency bands but no emission was detected at Cir X-1's location.

When interpreting the images, one must remain aware of the fact that I have imaged the source in both a state of rise and of decay, violating one of the major assumptions in aperture synthesis (see Chapters 1 and 3). As a result, artefacts may have been introduced into the image which can imitate the appearance of symmetrical jets (the same effects seen in the images of Cir X-1 in Chapter 5). I created simulations on the visibilities from the first observation where I assumed all emission was from a single point source at Cir X-1's location with flux density $S_\nu \approx$ average visibility amplitude (i.e. variable over the run). Images produced from the simulated data are shown in Figure 6.4 and do show signs of axial 'spokes' from the central source during peaks in amplitude; however, the intensity of these spokes beyond 0.5 arc-seconds remain below 10 per cent of the core flux density (i.e. lower than the 3σ noise levels measured in my images) and the more intense residuals

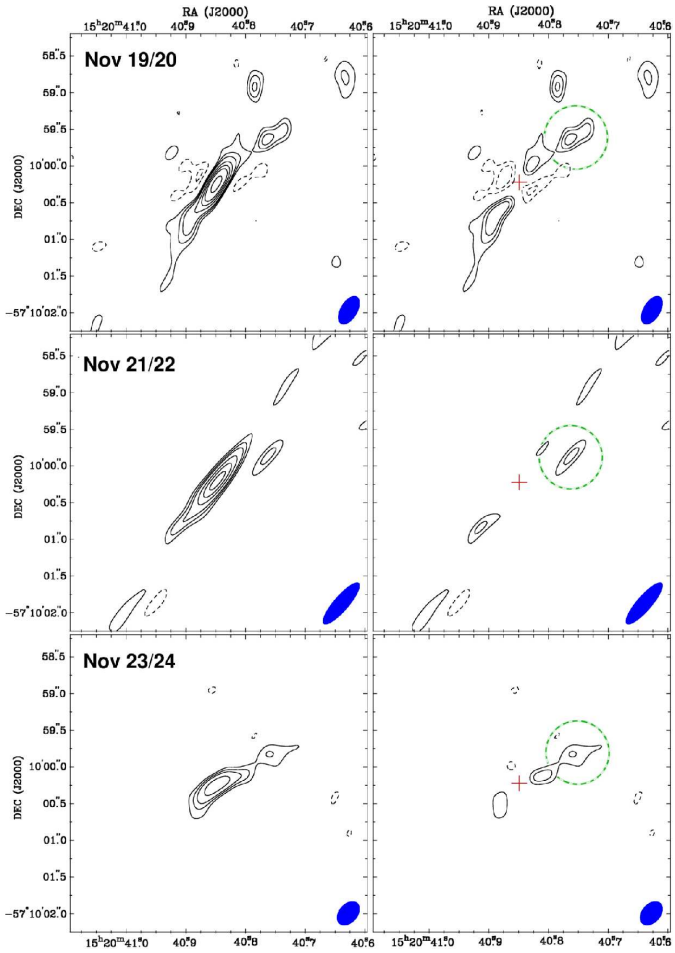

Figure 6.3: Mm contour maps of Circinus X-1. Images were created using combined 33 and 35 GHz data. Weighting uses a robust factor of 0.5 (an optimal compromise between natural and uniform weighting). Contours are set at -6, -4.5, -3, 3, 4.5, 6, 9, 12, 15, 21, 27, 36 $\times$ rms noise of each epoch. Beam sizes are 0.44$\times$0.24, 0.76$\times$0.21, and 0.39$\times$0.26 arcsec for the 19th, 21st and 23rd respectively. The right panels have had a fitted point source subtracted at the core's position (the cross). The dashed circle marks component N.

left within 0.5 arc-seconds can be identified by their symmetrical layout and pairing with nearby negative regions of similar intensity. Also of note is the fact that the position angles of the visible spoke structures do not correspond to the resolved structure visible in the real maps from the first observation epoch.

While I acknowledge that the first epoch image must be affected by the above artefacts to some degree, I remain convinced that much of the resolved structure cannot be a result of this for a number reasons. Firstly, as per the simulations, the artefacts should appear symmetrical about the point of variation, and while there is a degree of symmetry in the visible components, two distinct asymmetric features arise repeatedly in all images. These 'extensions' to the south-east and north-west of the core appear to differ in both structure and intensity, and a separate emission

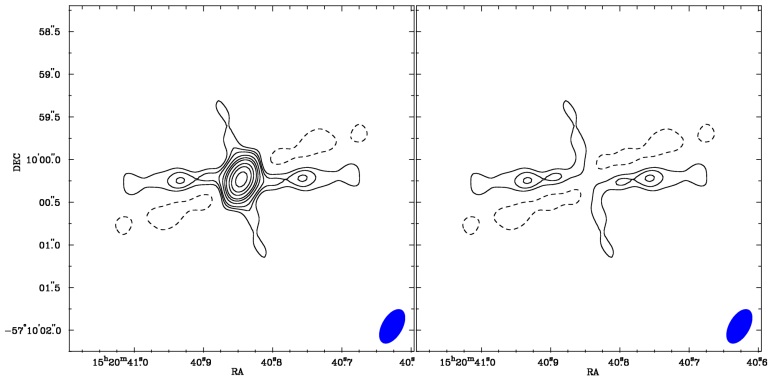

Figure 6.4: Cir X-1 core variability modelling maps. These maps have been made using model data derived from the visibility amplitudes shown in Figure 6.1 for the first observation, where it is assumed that all emission comes from a single point source at Cir X-1's location. The data uses the same data selection parameters as those used for real images (i.e. those defined by keeping rms fluctuations below 400 microns). The left panel shows the cleaned model image, whereas the right panel shows the result of subtraction of a modelled point source. Weighting uses a robust factor of 0.5 (an optimal compromise between natural and uniform weighting). Contours are set at -6, -4.5, -3, 3, 4.5, 6, 9, 12, 15, 21, 27, 36 $\times$ rms noise of 0.65 mJy. Beam size is 0.44$\times$0.24 arcsec.

component appears 0.75 ± 0.25 arc-seconds to the north-west of the core (slightly below the 'jet' axis). Secondly, the flagged data sets all share a similar starting hour angle, and though the second flagged data set is significantly shorter than the others, the position angle of the resulting beams (-30.7°, -40.9°, -41.1° for Nov 19, 21 and 23 respectively) remains similar across all data sets. With this in mind, one would expect the axis of visible artefacts to differ between the first and second runs, as each data set involves opposite types of light curve behaviour - rise and decay - and yet they remain similar. Thirdly, the final data set shows little or no evidence for continuing flux density decay, and as such should be free of the above effects. Yet, elongation of the core and the north-component continue to be seen in the final day's images.

There is a further concern in that data whose level of phase errors gradually vary could produce similar artefacts to those from source variability. If the phases vary significantly, then the position of the source can shift rapidly over an observation, which in turn causes the source to broaden in images (see Taylor, Carilli & Perley 1999, chapter 28). Though I have taken steps to eliminate afflicted data via flagging, this does not mean my images are not entirely unaffected. However, the random nature of these shifts means that over a large enough integration (i.e. data set length $>>$ phase shift time scales: true for all these data sets), artefacts should still appear symmetrical, making the first argument of the previous paragraph, coupled with multiple detections, the strongest test for validity; i.e. it is unlikely that persistent

asymmetrical structure will be an artefact.

To further examine resolved structure, I performed point source subtraction centred at the core, with the results shown in Figure 6.3 (right hand panels). The map of the 19th (first epoch) shows significant residual structure, but since it is the day of the flare, artefacts caused by source variability are more likely to occur. Signs of such effects appear in proximity to the subtraction point as roughly symmetrical negative zones along an axis perpendicular to that of positive structure. Even assuming the presence of artefacts, they cannot account for the brighter south-eastern core residual (approximately 50 per cent brighter; a difference of 0.8 mJy) nor the three other separate emission components in the north west quarter. Out of these remaining bright components, only the closest to the core (marked by the dashed circle in Figure 6.3, henceforth referred to as component N) remains in all subsequent images. The two structures farther from the core may be real and simply fade between the first and second observations (perhaps related to an earlier ejection event) but without multiple detections this is difficult to confirm.

The second map set also shows both south-eastern residuals and component N. A streak can be seen running the length of the image, parallel to the apparent jet axis, which I am unable to eliminate. This appears to be the source of several weak structures across the length of the axis, and also affects the southern extension.

The final observation continues to show component N and faint extension residuals; however, the northern extension now appears slightly offset from the 'jet' axis (further south than earlier examples) and is the more dominant of the pair of core residuals, though the difference in flux density is sufficiently low for noise to account for it. With variability likely to have reached minimum levels (as indicated by the image plane measurements), if any artefacts now exist in the map they should be a result of phase errors. This cannot be completely ruled out as a cause of the faint extensions near the core.

Figure 6.5 shows naturally weighted images from combined 21st and 23rd observations (post flare), resulting in a greater than 6σ detection of component N. Accepting component N as a real structure, it may be moving ejecta from Cir X-1 or a shock caused by re-energisation of slow material by an unseen flow (as in Fender *et al.* 2004). Point source fit flux densities are shown for the component in Table 6.2, and indicates a decline over the three observations, though I cannot determine the onset of a rise as the object is difficult to detect on shorter data intervals. The difference in flux density between the two frequencies is well within the measurement uncertainties so little of value can be said about the components spectra. LBA observations in Miller-Jones *et al.* (2012: work that was submitted in

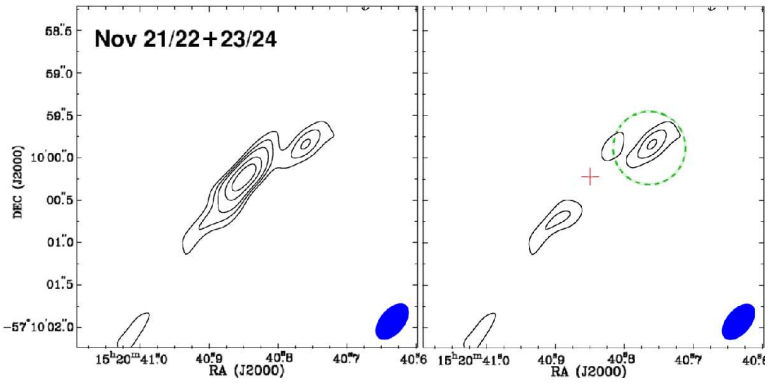

Figure 6.5: Stacked mm contour maps of Circinus X-1. Images were created using combined data from 2010 Nov 21 and 23 (33 and 35 GHz). Natural weighting is used for higher sensitivity at the cost of slight reduction in beam resolution (beam size = 0.50×0.29 arcsec). Contours are set at -6, -4.5, -3, 3, 4.5, 6, 9, 12, 15, 21, 27, 36 × rms noise of 53 μJy. Layout matches that used in Figure 6.3.

Table 6.2: Component N flux measurements. Measurements made using a point source fit, with errors taken directly from output (for corresponding image noise levels, see Table 6.1).

Date (UT)	33 GHz flux density (mJy beam^{-1})	35 GHz flux density (mJy beam^{-1})
2010 Nov 19	2.15 ± 0.48	2.19 ± 0.57
2010 Nov 21	0.53 ± 0.14	0.46 ± 0.15
2010 Nov 23	0.14 ± 0.03	0.11 ± 0.03

partnership with the publication of this work in Calvelo *et al.* 2012b), suggest the presence of objects with proper motions of 35 mas day^{-1}. Ejecta moving at such speeds in my observations would result in detectable shifts of position between the initial and final images. Analysis of component N's position in my three images shows a 0.29 ± 0.13 arc-second shift, corresponding to a proper motion of about 150 ± 70 mas d^{-1} to the south between the first and second epochs. This is the only significant change as all other shifts fall below respective fitting error limits. The shift actually brings the component closer to the core, perhaps implying a bent outflow or a sequence of different fading and brightening jet components. However, being a marginal difference and including the possibility that image errors have affected the component's position (i.e. more significant in the first day image) makes it difficult to justify further analysis.

Based on revised ephemeris estimates for the system (George Nicolson, private communication) this observation's epoch's flare should have occurred at MJD 55519.94; about an hour after the rise in the amplitude plot begins. Thus, compo-

nent N's distance from the core prohibits association with the observed flare, as it would imply a proper motion over 2000 mas d^{-1}. This argument would also be true for re-energisation via internal shocks. Such constraints, however, do not apply if the event is connected to an earlier flare. Comparison to observations of mm flares from BHXRB GRS 1915+105 (Fender & Pooley 2000) suggests that ejecta emission declines rapidly at mm wavelengths (GRS 1915+105 events have duration shorter than 1 hour), making an ejected component's possible life time of over 15 days questionable. Even if I assume a slower decay, extrapolation of component N's decay rate back towards the earlier flare would imply minimum flux densities greater than 50 mJy at launch; comparable to the peak core flare levels estimated from the amplitude plots. Shocks do not suffer such limitations, as the energising flows may travel undetected before finally brightening on interaction with slower media downstream (as is observed in other BHXRBs, e.g. XTE J1550-564: Corbel *et al.* 2002 and NSXRBs, e.g. Scorpius X-1: Fomalont, Geldzahler, & Bradshaw 2001). Thus, if I assume a single orbit's delay (16.6 days) in component N's variability, I may estimate proper motion for a flow of approximately 60 mas d^{-1}.

One may also be tempted to carry out similar calculations based on the slight shifts in position observed in the southern and northern extensions, but the possibility that they are affected by any of the artefacts I have discussed would severely limit confidence in results from such an exercise.

6.3.2 Scorpius X-1

6.3.2.1 Light curves

Figure 6.6 shows the visibility amplitudes, successful image plane measurements and path length rms for the available Sco X-1 observations. As mentioned in Section 6.2, mm observations were paused during the first epoch of observations due to decline in phase stability prompting a switch to cm wavelengths (the break resulting from the transition is clearly visible in the first epoch shortly after MJD 55519.0). While it was possible to make a measurement of the initial 1.66 hours of mm data, the short length of the observation and poor phase stability made imaging at high angular resolutions difficult. To work around this, the FWHM parameter in INVERT was used to define a Gaussian taper to apply to visibility data, effectively reducing the adverse effects of low phase stability at the cost of lowering the final resolution of the image. A concern with this method is that one may inadvertently reduce the resolution enough to cause separate nearby sources to be indistinguishable from a

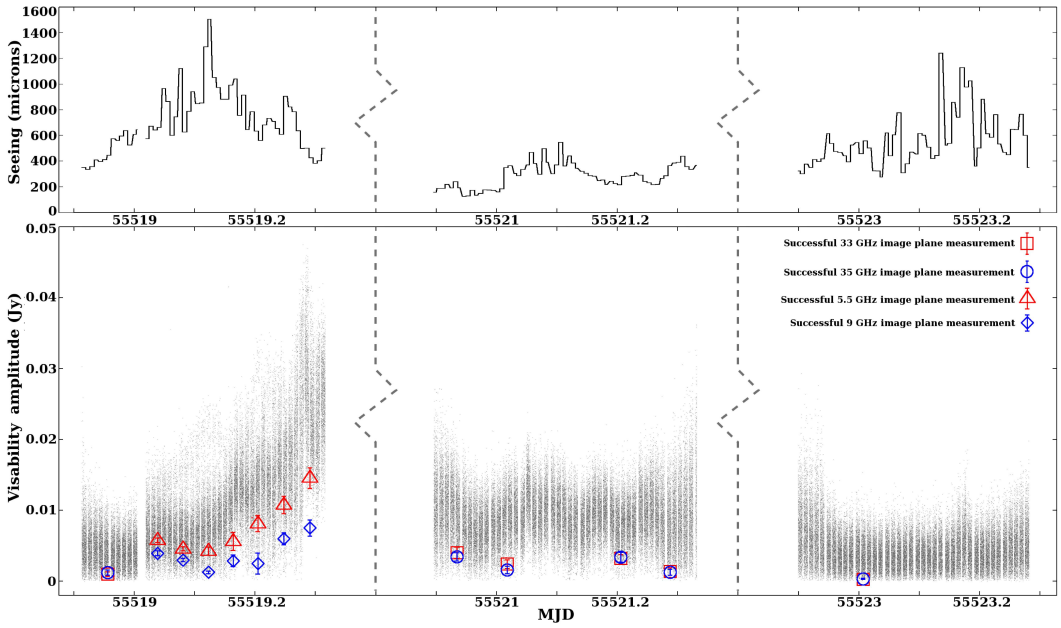

Figure 6.6: Scorpius X-1 mm light curves and rms path length fluctuations. ATCA-CABB rms path length fluctuations (top), *uv* amplitude (grey points), and image plane fit plots at 33 and 35 GHz (red and blue points respectively; three observations over 5 days). Measurement errors (1σ) are included in the image plane values but are too small to appear clearly.

single source and so tainting flux density measurements. However, based on mm images from subsequent full observations, Sco X-1 is the only visible source in the maps and so any visible flux should be attributed to it alone. Just as was the case for Cir X-1, this also means that plots of the visibility amplitudes can act as a rough estimate of Sco X-1's flux density allowing me to see variability; however, unlike Cir X-1, Sco X-1 shows no additional structure at mm wavelengths so the visible amplitudes should correspond more closely to the point like core alone. This does not apply to cm observations, with the nearest additional visible sources being the background "lobe" like sources to the north east and south west which also contribute to the visibilities. Furthermore, one must continue to take care in assuming variability in the amplitudes is a result of variability in source flux densities, as it can also be caused by resolved structure around a source.

The most notable feature in the light curve is the variability at 5.5 and 9 GHz, visible in both the visibilities and image plane measurements. During the first 2 - 3 hours visibilities and source fitting show a decline in Sco X-1's flux density, with initial measurements of 5.79 ± 0.33 mJy and 3.94 ± 0.35 mJy at 5.5 and 9 GHz respectively (MJD 55519.039), declining to minima of 4.25 ± 0.58 mJy and 3.94 ± 0.20 mJy (55519.123). This behaviour then reverses, with flux densities climbing steadily over the remaining 4 - 5 hours to reach peak values of $14.57 \pm$

1.44 mJy and 7.52 $\pm$ 1.14 mJy (MJD 55519.330). Measurements show a drop at MJD 55519.205 at 9 GHz but this is likely an error resulting from sub par image quality: source elongation along the north-south axis and subsequent decrease in phase accuracy made source fitting difficult (hence the larger uncertainty). There is no corresponding drop in the 5.5 GHz data (there was less difficulty with imaging). The cm light curve does not show any new behaviour in comparison to variability profiles of the past. A factor of $\sim$ 3.5 increase over approximately 5 hours can be seen, a similar rate seen in several past flare events (See Bradshaw, Geldzahler & Fomalont 1997), some of which have also shown multi-peak profiles which could easily result in the observed decline prior to the rise but might suggest further activity before observations began.

The second epoch of observations showed the lowest levels of path length fluctuations ($<$ 600 microns for the entire run) and thus could be used in its entirety during imaging. Both visibilities and image measurements indicate the continued presence of variability, but on shorter timescales than that observed at 5.5 and 9 GHz. Average image measurement values at both frequencies are higher than those measured during the first mm epoch but individual measurements vary by a factor of $\sim$ 3 over the observation, with the shortest time between two measurements being 2 hours (the central gap is 4.5 hours long). Source elongation in the N-S direction makes imaging of short data segments near MJD 55521.1 (towards the middle of the observation) difficult and subsequent source fitting attempts proved to be unsuccessful. Again, I see no activity here that is particularly surprising for Sco X-1 as several of the 6cm VLA observations of Bradshaw, Geldzahler & Fomalont (1997) show profiles with peak to peak spans as short as half an hour, though this is the first evidence of such rapid fluctuations at mm wavelengths.

The final mm epoch shows some regions of high path length rms, some reaching levels as high as those seen during the first epoch but more inconsistent. Path length rms was so high (above the levels seen in the second epoch overall) that I was unable to create images of sufficient quality from shorter segments of data to construct an image measurement lightcurve. I am left with a single measurement from an image created with the majority (59%) of data, essentially excluding regions towards the latter half of the observation where severely high path length rms becomes an issue (MJD 55523.125 onwards). The flux densities themselves are the lowest mm values measured from the source over the entire 3 epochs, with 0.29 $\pm$ 0.06 mJy and 0.34 $\pm$ 0.07 mJy at 33 and 35 GHz respectively. The visibilities show possible signs of gradual decline early in the observation but I am unable to confirm whether this is real without greater detail from image measurements.

6.3.2.2 Spectrum

Spectral index measurements calculated using the 33 and 35 GHz measurements plotted in Figure 6.6 suffer from large uncertainties (on a similar scale to those discussed for Cir X-1) and therefore any conclusions derived from the individual measurements would be difficult to justify. The average spectral index of the mm observations is α = -0.11 $\pm$ 0.42 (where $S_\nu \propto \nu^\alpha$). The values calculated from 5.5 and 9 GHz are better with lower fractional errors, but some care is still warranted. Ignoring the values at MJD 55519.205 due to the likely erroneous 9 GHz measurement, five of the remaining six points give indices that vary between α = -0.78 and -1.36 with an average of α = -1.11 $\pm$ 0.24 and though there is no obvious pattern, indices tend to be steeper later in the observation. The omitted sixth value is of α = -2.47 derived from the measurements of the minimum in the light curve and as an outlier may suggest that one or more of the flux densities at this point has not been measured accurately: possibly caused by the very high path length rms values at this time.

The trend towards negative indices, ignoring their uncertainty, suggests that we may be observing structure dominated by optically thin synchrotron emission in and around Sco X-1's core. This appears to be true even during the flare rise, when one might expect an inversion of the spectrum due to transition into optically thick regimes in shocks or ejecta (the same scenario discussed for Cir X-1). There is a slight flattening of the spectra early in the rise, which could be attributed to the presence of such phenomena before adiabatic expansion lowers their optical thickness and leads to a return to steeper indices; however, the aforementioned uncertainties prevent us from drawing stronger conclusions. It should be noted that an example of such spectral behaviour exists in the profile of 1982 Sco X-1 activity (see figure 2 of Bradshaw, Geldzahler & Fomalont 1997).

Figure 6.7 displays the plots of five frequency pairs of the Sco X-1 flux density measurements as well as four fitted power law for different groups of values. For 5.5 and 9 GHz I have included two pairs of measurements: one for the full data set, and a second which uses the values for the first 5.5 and 9 GHz points plotted in Figure 6.1's light curve (approximately 1 hour of observation time). For 33 and 35 GHz I have included the measurements from images of each epoch. It should be noted that though measurements are taken from image point source fits (whose uncertainties can be found in Table 6.1), due to the variability observed in both the cm and second epoch mm data sets I have opted to display conservative uncertainties calculated from the range of light curve values for each epoch.

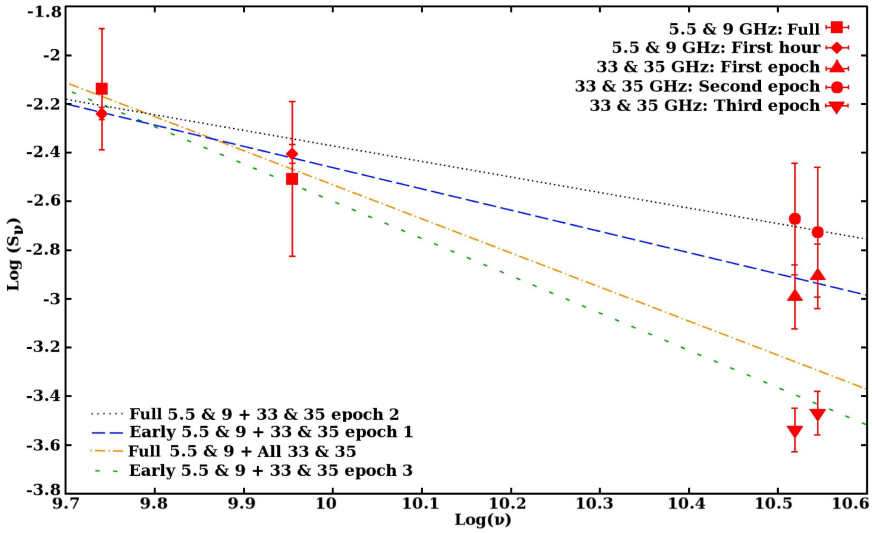

Figure 6.7: Scorpius X-1 mm and cm radio spectral fits. The plotted points represent the measurements from this chapter's ATCA-CABB observations divided by epoch, with the exception of the 5.5 and 9 GHz first hour values (diamond symbols) which are a subdivision of the full 5.5 and 9 GHz data sets. Fitted power laws were calculated using the fit command in GNUPLOT. In the case of epochs with clear variability visible in lightcurves (5.5 and 9 average, 33 & 35 second epoch) the plotted errors are the calculated sample standard deviation for viable image plane measurements. For epochs with minimal variability or no more than one image plane measurement, the uncertainty is taken directly from MIRIAD's point source fitting (5.5 & 9 first hour, 33 & 35 first epoch and 33 & 35 third epoch).

If I apply the fitting to the "full" data available (i.e. the values measured from the entirety of the cm epoch and the three sets of mm measurements) the result is a spectral index of α = -1.40 $\pm$ 0.70. Given the range of flux densities in the light curve, the uncertainty is understandably large, and while the value is not physically unrealistic for an optically thin synchrotron source, I must question its validity based on the inherent variability that exists between individual epochs. However, a steep negative spectrum is also supported when I apply a fit to the early 5.5 and 9 GHz data measurements grouped with the values from the third mm epoch, with α = -1.53 $\pm$ 0.23. The danger here lies with the corresponding observations being taken four days apart, with the significant changes in flux density occurring between. One can argue that the initial cm data may represent pre-flare flux densities, and that the third mm epoch takes place sufficiently long after the visible rise that together these values can represent "low" Sco X-1 activity levels. However, given the lack of information on variability prior to the observed rise I cannot safely make such an assumption, especially if the gentle decline early in the cm light curve is actually the decay following earlier activity as was suggested in Section 6.3.2.1. It should be noted that both these values are only slightly steeper than the indices calculated from

the later 5.5 and 9 GHz lightcurve measurements in Section 6.3.2.1, with significant overlap in uncertainties.

If focused on minimising the amount of time between multi frequency measurements, I would be inclined to apply fitting to the first mm epoch measurement together with the early component of the 5.5 and 9 GHz observations taken immediately after, yielding α = -0.87 $\pm$ 0.06. Much shallower than the previous estimates, this value is close to indices calculated for the early 5.5 and 9 GHz rise before it steepens for the remainder of the time. This is not surprising if one again considers the possibility of optically thick emission from new ejecta or shocks influencing the spectrum for a short period of time during the early flare. The shallowest spectral index results from a fit using the full 5.5 and 9 GHz measurements and the second mm epoch values, with α = -0.64 $\pm$ 0.16. While these two epochs are closer in time than the early cm and third mm epoch, they are also the two data sets that display the most significant variability during observations and thus it is difficult to distinguish whether the cause of the flatter spectrum is related to optical thickness changes or simply that the system is brighter overall during the second epoch than the first epoch.

6.3.2.3 Imaging

Figure 6.8 displays the results of 5.5 GHz and 9 GHz imaging, region scaled so that the background radio "lobes" are visible. Sco X-1 itself is visible in at both frequencies as a single point source, with traces of faint resolved structure around it. After removal of a fitted point source at Sco X-1's location I find residual structure predominantly to the south west ($\sim 5\sigma$), but also a second component on the opposite side of the subtraction point of similar intensity. When deciding whether these structures are real, one must first take into consideration the fact that the 30°$\pm$10° axis of these components differs from the known jet/ejection axis for the system (54°: Fomalont, Geldzahler & Bradshaw 2001) which has shown very little variation over long time periods, thus it is unlikely these structures are related to the system's jets. Secondly, the images may include flux density variability artefacts (see Section 3.4). The axis formed by the residuals in Sco X-1's maps corresponds to hour angles towards the end of observations, and thus the times when its flux densities are highest. As a result of these arguments I am inclined to believe that these residuals are, in fact, artefacts caused by the system's variability.

Figure 6.9 displays the results of combined 33 and 35 GHz imaging (visibility sets from both frequencies are included in the Fourier inversion during MIRIAD's

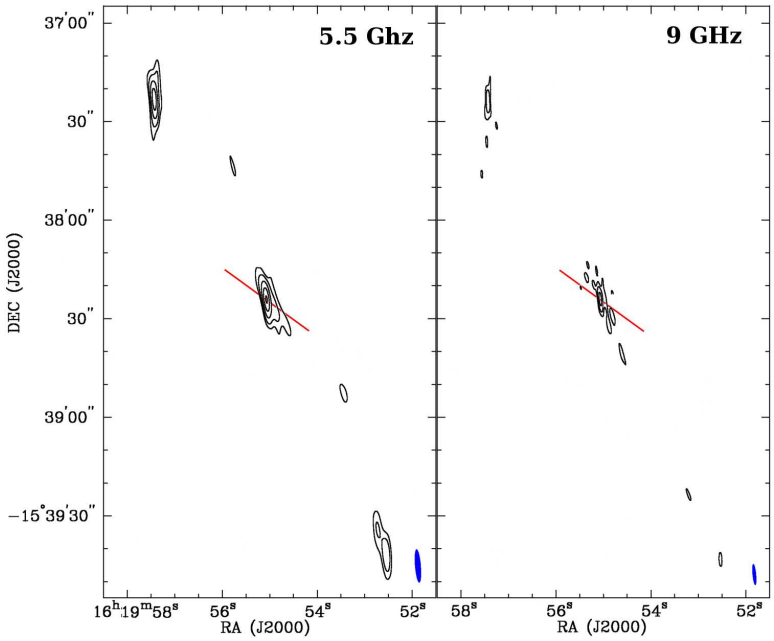

Figure 6.8: Scorpius X-1 ATCA-CABB 5.5 GHz and 9 GHz radio maps. Sco X-1 is centred but also shown are the two background radio sources (top left and bottom right) that were previously believed to be its lobes. The red line marks the known 54° jet/ejecta axis (Fomalont, Geldzahler & Bradshaw 2001). Natural weighting was used and beam size is 6.39 arcsec × 1.04 arcsec. Contours are set at -3, 3, 6, 12, 21 and 30 times the image rms of 0.22 mJy for 5 GHz and 0.19 mJy for 9 GHz.

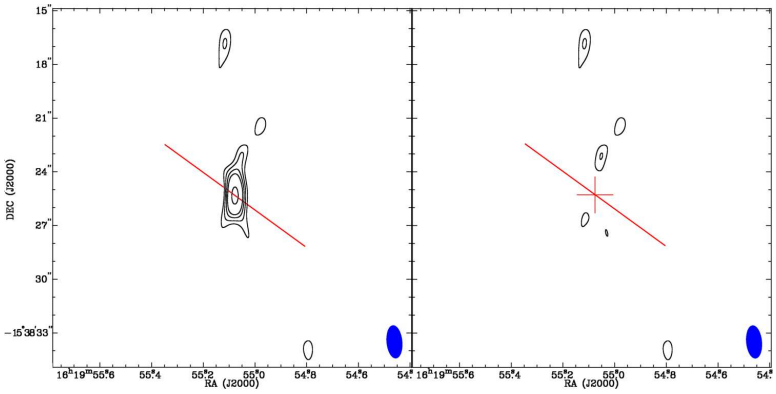

Figure 6.9: Scorpius X-1 ATCA-CABB combined 33 and 35 GHz second epoch radio maps. The left map is the original cleaned map which used a robust = 0.5 weighting scheme (robust being a compromise between natural and uniform weighting). The right map has had a fitted point source subtracted from the location marked by the cross and shows residual nearby structures. The red line marks the known 54° jet/ejecta axis (Fomalont, Geldzahler & Bradshaw 2001). Contours are set at -3, 3, 4, 5, 6 and 9 times the image rms of 0.16 mJy, and the beam size is 1.85 arcsec × 0.87 arcsec though it should be noted that a FWHM setting of 1 arcsec was enforced during the MIRIAD's INVERT step.

INVERT step) from the second epoch of mm observations. I have chosen to display the images from this epoch's observations due to their consistently low path length rms (resulting in it being the largest data set post-flagging) and thus should provide the highest image fidelity.

Once again, one can observe a point source with possible resolved structure surrounding it, revealed in the residuals left after subtraction of a fitted point source (right panel of Figure 6.9). As was found for the cm images, the axis in the mm images differs from that described in Fomalont, Geldzahler & Bradshaw (2001), making it unlikely that the structure is related to the system's jet ejections. Though I do not see the same type of flaring behaviour as the 5.5 and 9 GHz visibilities, there remains some level of variability in the epoch's data set which may contribute to the minor structures nearest the core. Furthermore, as was discussed for Cir X-1, one must remain aware of the possibility of artefacts arising from phase errors in mm observations, even though steps were taken to reduce the effects to some degree (flagging and the enforcement of the FWHM parameter described at the beginning of Section 6.3.2.1). Though not shown, the images from the other two epochs of data show no residual structures concordant with what can be seen here. Thus, having no physical reason to believe they are real, I must conclude that all the weak structures visible beyond the core can be explained as artefacts caused by phase instability, source variability or both.

Polarisation imaging was carried out for all frequency bands but no emission was detected at Sco X-1's location.

6.4 Discussion

6.4.1 Circinus X-1's varying jet structure

Having established that image artefacts predominantly affect the structures nearest to Cir X-1's core, my best jet axis estimate uses component N's position, giving an angle of $121 \pm 5^\circ$ E of N. Though this estimate falls within the range quoted in Tudose *et al.* 2008 ($129 \pm 13^\circ$ E of N) and is only marginally steeper than the estimate from milli-arcsecond structure in the Miller-Jones *et al.* (2012) LBA observations (113° E of N), it is approximately 50° shallower than the axis estimated from 5.5 GHz residual emission near the core in the images of Chapter 5. Furthermore, whereas Miller-Jones *et al.* (2012) observe symmetrical jets, my maps indicate a scenario much like that observed in Chapter 5 and Tudose *et al.* (2008) where asymmetry of near-core residuals indicates a brighter southern jet. Like in

Chapter 5, I also observe distinct activity from a region that has previously been associated with the receding jet (Tudose *et al.* 2008) in the form of component N.

The fact that the various structural orientations appear on different scales across observations taken within months of each other supports a scenario where the jet angle varies as one moves further from the system's core. Starting off at relatively high angles to the line of sight and near E-W as indicated in Miller-Jones *et al.* (2012), but steadily turning N-S and closer to the line of sight as one moves down the flow to arcsecond scales as in Chapter 5. The origin of such deviation may be precession, or quasi-static jet kinks. The unexpected appearance of structures to the north west (component N in this case) may be a result of asymmetry in the density of interacting media, or deviations in a jet close to the line of sight (as also suggested in Chapter 5).

The jet velocities reported in Miller-Jones *et al.* (2012) are of similar order to my estimates for energising flows in section 6.3.1.3 (when I assume a full orbit's phase delay), both of which fall far below the ultra-relativistic speeds calculated by Fender *et al.* (2004). However, like Miller-Jones *et al.* (2012), I too cannot completely rule out the existence of such ultra-relativistic flows. There remains the possibility that this Chapter's observations did not capture the full flare event associated with this orbital epoch (ephemeris estimates are not perfect, nor every flare event the same), altering the reference time on which my component N calculation is based.

Reconciling these combined results with historic observations such as the ultra-relativistic velocities of Fender *et al.* (2004) and lack of precession observed in Tudose *et al.* (2008) is difficult. Lacking details on sub-arcsecond structure at the time of the Tudose *et al.* (2008) observations, I cannot establish whether Cir X-1's jets have always exhibited variable inclination with distance from the core (though jet curvature in large scale images from Tudose *et al.* 2006 indicate some level of this). By grouping radio observations made after 2008 and comparing them to older results, one can observe an almost bimodal divide in Cir X-1's behaviour, with lower jet velocities, precession (indicated by the difference in jet angle observed at 5.5 GHz: Chapter 5) and multiple incidences of activity to the north-west of the core rather than the south-east being recent changes. This may indicate that, along with the decline in intensity over the last decade, the system's jet structure can also vary significantly with time.

Such a scenario draws further similarities between Cir X-1 and SS433/W50. Hydrodynamical simulations by Goodall, Fathallah & Blundell (2011) show that the structures visible in nebula W50 cannot be reproduced by extrapolation of SS433's current jet activity back through time. Instead they show that at least three distinct

epochs of jet activity, each with a different level of precession, were required to produce the nebula's layout. If the filaments visible in Chapter 5's Cir X-1 nebula maps (Figure 5.13) are indeed caused by jet-ISM interactions, then the precession angle of the jets would have had to be significantly larger than that indicated by pre-2008 radio observations. Thus, variability of Cir X-1's jet characteristics should not be entirely unexpected.

6.4.2 NSXRB radio and mm spectra behaviour

As mentioned in the Chapter's introduction, the results of Migliari *et al.* (2010) have shown that the available spectrum of at least one NSXRB (4U 0614+091) can be fitted by a flat power-law from the radio to the mid IR, implying emission originates in an optically thick jet. However, my results indicate that the radio to mm spectra of both Cir X-1 and Sco X-1 may be optically thin, suggesting that the jet break frequency (ν_{break}); i.e. the frequency at which the jet becomes transparent to its own synchrotron emission, can occur at far lower values (though, as will be discussed, care must be taken in making a direct comparison between the sources). Even taking the shallowest available power law fits for each source together with the uncertainty in the estimates, one is left with indices steeper than $\alpha <$ -0.4 for both cases: a value concordant with optically thin synchrotron emission. Nevertheless, there are several arguments that may reduce the validity of these results, or explain the existence of a steep spectrum without assuming the jet becomes optically thin early.

A major concern for both Cir X-1 and Sco X-1 is the inherent variability that occurs within the systems and the implications it may have for my estimates. Given the different disadvantages of the several possible fitting options available for the Sco X-1 measurements (See Section 6.3.2.2) I must conclude that the strongest evidence for an optically thin spectrum comes from the fitting of the first 1.66 hours of mm observations together with the first hour of 5.5 and 9 GHz (α = -0.87 $\pm$ 0.06). I choose so because Sco X-1 shows enough variability, both $\sim$ hourly and from day to day, that observation simultaneity is essential to accurately estimate the spectral index. Even using these sets of consecutive measurements, there remains the possibility that the source has varied during the interim (i.e. near the time of the observation frequency change over). In order to over estimate the steepness of the power law index, the cm measurements would have to have occurred while the system was more luminous, and though the cm light curve is in decline when observations begin, the possibility that a brief rise preceded this activity cannot be

eliminated. Even if that is the case, one might also already expect the spectrum to be flatter as a result of increased optically thick emission from regions that can appear during jet flares.

The case of variability adversely affecting my spectral index estimates is a far bigger concern in the case of Cir X-1 since in its case the data points used are separated in time by almost 10 months. Furthermore, as is evident from Chapter 5's eight 5.5 and 9 GHz daily measurements preceding the day of the flare, Cir X-1's flux density continues to vary from day to day though it is in a relatively low activity state. While this is accounted for in the uncertainties included within my calculations to some degree, it should be remembered that the 33 and 35 GHz counterparts are only the result of a single day's observations, and I cannot be sure how much variability exists at mm wavelengths during the same periods of inactivity.

If I assume that the evidence for an optically thin spectrum at these wavelengths is real, I must address why my results differ so significantly from the work of Migliari *et al.* (2010). Given the uncertainties in my measurements and the range of possible indices for optically thick and thin synchrotron radiation, I am unable to estimate ν_{break} accurately, but by assuming high optical steepness ($\alpha \sim$ -2) of the optically thin radiation I can say that $log\nu_{break}(Hz)$ must be $<$ 10.2 for Sco X-1 and $<$ 10.4 for Cir X-1. This is 500 - 2500 times lower than the ν_{break} estimated for 4U 0614+091.

One possible explanation is that this Chapter's observations not only observe the main jet, but a significant amount of nearby material associated with previous ejection events. If the compact jet is sufficiently faint, a large enough quantity of unassociated optically thin matter could dominate the spectrum at these frequencies. The frequency of Sco X-1's ejections and the limited resolution of this chapter's images (1 arcsecond and up) means that it is very possible that the central point source is not formed by a single object. In fact, the imaging done by Fomalont, Geldzahler & Bradshaw (2001) shows several examples of ejection pairs within 50 milli-arcseconds of Sco X-1's central core, the fluxes of which (at least in the case of the north eastern component) can be comparable to the core itself, with negative spectral indices. This possibility was also used to explain why the measured spectral index appeared to flatten during the flare but remained optically thin, in that the main jet's spectrum may have inverted briefly, but the overall flux remained dominated by optically thin material further down stream and muted the observed effect. One might argue that these components should eventually fade and leave the main jet as the dominant emission source, but re-energisation via unseen ultra-fast streams can extend the lifetime of visible ejecta. This argument could also be applied to Cir

X-1; however, as shown in Chapter 5 (and to a lesser degree the mm image plane measurements detailed in Section 6.3.2.1), Cir X-1's index can actually invert at the start of flares, the speed at which it does so suggesting that if such optically thin material is present it is to a lower extent than in Sco X-1, since I do not observe the same effect with similar relative variability.

If dealing with emission from the compact jet alone, the discrepancy might also be explained by inherent differences in the systems themselves. One obvious distinction between Sco X-1 and Cir X-1, and 4U 0614+091, is their NSXRB type, with 4U 0614+091 being an atoll source, Sco X-1 a Z-source, and Cir X-1 straddling the two. As mentioned in Chapter 2, it is believed that one of the main influences behind the different behaviours of the NSXRB types is accretion rate, with Z-sources having higher accretion rates than atolls. These high accretion rates likely result in a scenario parallel to that of BHXRB GRS 1915+105, wherein the source very often transitions between hard and soft states, repeatedly crossing the jet line and producing very fast ejections in a variable jet (see Section 2.3.4, also Klein-Wolt *et al.* 2002). Indeed, both Sco X-1 and Cir X-1 have shown evidence for discrete jet ejecta moving away from the core following flare events: behaviour which 4U 0614+091 lacks. Cir X-1 is of course a special case, since its overall accretion rate varies with orbital phase and so it may not reach Z-source like accretion rates at all times (likely resulting in the odd behavioural changes discussed in Section 5.2).

It is not known exactly how the jet's properties vary during such transient events, but some change in flow velocity, particle density or even jet geometry should be expected given the accretion disc structure varies with state, and its close ties to jet formation. Given the frequency of flare and ejection events in Sco X-1 and Cir X-1, it can be inferred that the properties of the jets in these systems do not remain stable for long (although this can depend on the exact ratio of matter that is drained from the disc during each flare). Therefore, we cannot predict whether we observe the jet during a particularly 'weak' epoch, during which flow rate, collimation (i.e. large opening angle), and particle density are all lower than average. Since ν_{break} should be intimately tied to the jet's properties (see Section 2.2.3.3), it will likely be lower during such 'weak jet' states. In support of this, as discussed in the previous subsection, evidence suggests that Cir X-1 may now often show a 'weaker' jet state compared to periods of monitoring prior to 2005, based on the secular changes that have been observed in X-ray and flare intensities, as well as the absence in my results, of radio structures and strong evidence of ultra-relativistic flows that were originally reported in past observations (Fender *et al.* 2004, Tudose *et al.* 2008). In contrast, 4U 0614+091 may accrete at lower rates that support a more steady jet.

It should be noted that variation in ν_{break} has been seen to occur between different observations of the same BHXRB source: GX 339-4. Corbel & Fender (2002) detected a ν_{break} for two observations of GX 339-4, but each in a slightly different position which would suggest the value can change over time in BHXRBs, possibly in relation to differences in bolometric luminosity and thus accretion rate (see section 4.1 of Migliari *et al.* 2010 for a full discussion). Given the similarities in accretion mechanics for all low-mass XRBs, regardless of the nature of the compact object, then if ν_{break} can vary in one system one would also expect the same to be possible in others.

The more extreme differences in jet behaviour between the two types of NSXRB may be unique to neutron star systems, but given the small amount of available information on these wavelength regimes for all XRBs it is too soon to draw such conclusions. Without further observations at mm wavelengths and quasi-simultaneous multi-wavelength radio campaigns (preferably during epochs that do not suffer from flaring), including attempts at improving resolution to better resolve non-core structure, it is not possible to distinguish from the above possibilities in the case of Cir X-1 and Sco X-1.

6.5 Summary

I have presented the results of centimetre (5.5 GHz and 9 GHz) observations of NSXRB Scorpius X-1 and the first successful millimetre (33 GHz and 35 GHz) observations of both Scorpius X-1 and Circinus X-1 using ATCA - CABB. I detected both systems at mm wavelengths at greater than 5σ levels, with Cir X-1 also showing sub-arcsecond jet structure. Cir X-1's light curves indicated the rise phase of a periastron flare event occurred on Nov 19, decaying over the following four days. Scorpius X-1's light curves also showed rapid mm variability, as well as a single flare event visible at 5.5 and 9 GHz.

Spectral index estimates for both systems indicate that α remains negative throughout observations, though some flattening is seen at the start of flares which is expected due to increased optical thickness of jet material. Power law fits for groups of mm and cm data points indicate that the spectra of both sources remains negative across radio and mm regimes, in contrast to the spectrum of 4U 0614+091 described by Migliari *et al.* (2010). I interpret this result as optically thin synchrotron emission dominating within the extent of the point source regions, though I cannot tell whether this originates in the main jet or high amounts of nearby opti-

cally thin material from previous ejections and shock events. In the former case, I discuss the possibility that contrasting source behaviour (atoll vs Z-source) may be the cause of significant differences in my results to the spectrum of 4U 0614+091. I also; however, put forward alternative explanations for the negative spectral indices, including the effects of variability between the quasi-simultaneous observations at different frequencies.

Initial analysis of Sco X-1 radio maps indicate possible resolved structure around a point source at Scorpius X-1's location, but I am able to effectively eliminate these as artefacts, leaving the only real structure being an unresolved point source at both cm and mm wavelengths.

Circinus X-1 radio maps showed resolved structure near the core and a north-western component whose position angle, though in agreement with the jet axes of Tudose *et al.* 2008 and Miller-Jones *et al.* (2012), indicates activity in a direction previously associated with the receding jet. Variability of the north-western component (likely the result of re-energisation by unseen outflows) cannot be physically reconciled with the observed flare, but if caused by an immediately previous periastron event, implies proper motions far lower than those observed in Fender *et al.* (2004). Given the observed differences in jet behaviour compared to results prior to 2008, I believe that this work, in conjunction with results of Chapter 5 and Miller-Jones *et al.* (2012), suggest Cir X-1's jet behaviour; including flow velocity and precession parameters, can vary on secular timescales.

For my part I know nothing with any certainty, but the sight of the stars makes me dream.

VINCENT VAN GOGH (1853-1890)

7

Conclusions & Future work

In this chapter I give a brief overview of the main results of this thesis. In addition, I will discuss the implications in a broader context, and present my thoughts on the prospects for future research that could build on what has been discovered herein.

7.1 Summary

In Chapter 3 of this thesis I presented a study of the effects various types of spectral variation in sources can have on attempts to image them in the radio with ATCA-CABB. I discovered that spectral variation in a single source does indeed produce artefacts in subsequent images, and that the severity of the effects increase proportionally with the difference in spectral shape and brightness of the progenitor source, but inversely with duration of the shift. My results implied that these effects only become important for particularly bright and variable sources, and are of secondary concern to those that arise as a result of more common variations in source flux densities; however, I suggested that the fractional severity of the effect is likely to increase as the range of baseline lengths decreases. The results of the chapter are important for observations of XRBs as they are intrinsically variable sources, and so, I apply the knowledge I gained from the investigation several times

in subsequent work covered in chapters 5 and 6.

In Chapter 4 of this thesis I presented the results of attempts to observe the BHXRBs GRO J1655-40 and XTE J1550-564 during quiescence using ATCA-CABB. The aim was to compare radio measurements with X-ray luminosity and thus add additional points to the existing radio/X-ray correlation at low luminosities. Though I was unable to successfully detect either source, the radio upper limits defined by the final images still provided valuable information when included with existing data sets. The levels at which the sources went undetected were sufficiently faint in the radio that their positions on the correlation indicated significant scatter exists at low luminosities. I went on to propose the possibility that the scatter indicates the presence of more than one correlation track in the diagram, perhaps indicative of further differentiation in the hard and quiescent states of BHXRBs, as has now been established as a fact at higher luminosities.

In Chapter 5 of this thesis I presented the results of 16 days of ATCA-CABB observations covering a complete orbit of the NSXRB jet source Circinus X-1. These observations were carried out during an epoch in which the system was particularly faint compared to its past, making the (then recent) upgrades to ATCA essential if new images of the source were to be made and reviewed for structural changes. I produced a detailed light curve for the system which clearly showed a periastron flare event rise and gentle 6 day decline, though the peak of the event went unseen. My day to day radio maps showed two quasi-persistent structures near to the source, but I could not confirm their association with Cir X-1's jets. Many of the post-flare images showed strong evidence of variability artefacts similar to those discussed in Chapter 3, which I attempted to remove before discussing the remaining structure. I also applied the knowledge gained from my work in Chapter 3 in distinguishing which components were most likely to be real physical structures. From these revised images, I found signs of additional variable structures to the north-west which I interpreted as activity resulting from jet interactions, though I once again found it more likely that they were associated with outflows from a previous flare event rather than the one present in my light curve. Their position was surprising in comparison to past images of the system (namely those of Fender *et al.* 2004 and Tudose *et al.* 2006, 2008), and I surmised that either the local environment around the system had varied since observations (yielding asymmetric interactions near the source) or the jets themselves may have shifted or now precess. I initially favoured the former scenario; however, my production of deep radio maps of Cir X-1's surrounding nebula revealed distinct filament like structures which I likened to images of SS 433/W50 and lent support to the theory of precessing jet outflows, though this

went against previous analysis of long term imaging of the system.

In Chapter 6 of this thesis I presented the results of centimetre ATCA-CABB observations of the NSXRB Scorpius X-1, as well as millimetre observations of both Scorpius X-1 and Cir X-1. The aim was to observe NSXRB emission for the first time at mm wavelengths and produce the highest resolution ATCA-CABB images of these systems to date. I successfully detected both systems, making them the first NSXRBs ever to be detected in the mm bands, including variability in the mm lightcurves. The spectral indices indicated by comparisons of the cm and mm flux densities for the two sources (using the measurements of Chapter 5 in the case of Cir X-1) suggested dominance by optically thin synchrotron emission, contrasting with predictions for compact jet emission from atoll NSXRB 4U 0614+091 being flat all the way to infra-red wavelengths in Migliari *et al.* (2010). I proceeded to provide several explanations for this effect, including the presence of nearby optically thin material unassociated with the compact jet, measurement errors, or most interestingly, that the difference arises as a result of contrasting jet behaviours in atoll and Z-source NSXRBs. The cm images of Sco X-1 showed little of interest when compared to older studies, and its mm image showed only an unresolved point source. However, the mm images of Cir X-1 showed evidence of sub-arcsecond jet structure which I proceeded to interpret using the same criteria applied in Chapter 5. I once again observed the greatest signs of activity to the north-west of the source, supporting the results of Chapter 5, and in addition the comparison of the mm images with those of the Chapter 5 revealed differences in jet orientation that could be attributed to either a twisted jet (i.e. one whose observed axis changes with distance from the source) or that the jet shifted between the two periods of observation and thus may also support the theory of precession in the system.

7.2 Broader concepts & future prospects

7.2.1 Accounting for the effects of variable radio sources

Though the work in Chapter 3 led me to conclude that the effects of spectrally variable sources will only be a concern in a small number of current radio observations, this does not mean that the effect should be ignored. Given that radio telescopes will continue to evolve in the future, the artefacts that can result from spectral variance could become a more significant concern as bandwidth and sensitivities increase. Furthermore, since the conclusions of Chapter 3 were derived from tests on ATCA-CABB, the same cannot be said for instruments with better two dimensional antenna

placement like the VLA. For such arrays (including all recently built or planned instruments; LOFAR, SKA etc.) flux density variability will not result in artefacts as severe as those seen in Chapters 5 and 6 with Circinus X-1, since phase sensitivity does not drop significantly along any one particular axis (though they are not completely unaffected). They do, however, continue to observe the effects of spectral variability, as all interferometers depend on accurate modelling of the dirty beam and its sidelobes in order to effectively clean images.

Consideration of current reduction methods for complex sources leads to a logical solution to this problem, via improvements to existing software. One can imagine an evolution of existing MFCLEAN algorithms which attempt to fit a model source spectrum that not only varies over frequency, but over time as well. The increase in calculated variables would evidently impact computing time required to find solutions, and will only be viable when either time is not an issue, or more advanced computer hardware becomes available. This method also does not circumvent the issue of degeneracy between solutions and actual real structure in the model, wherein a time variable MFCLEAN may also fit a model to a single extended, or group of sources and substitute a false point source. Care would have to be taken by the user as to when to use such algorithms, or perhaps it will be necessary to analyse smaller divisions of an observation (both in frequency and time) in order to provide a input model for the beam (much like my modelling of Cir X-1 in Chapter 5), though this will be time consuming and depend heavily on the snapshot capability of the instrument.

7.2.2 Continued exploration of radio/X-ray correlations

Thanks to the work of Gallo, Miller and Fender (2012) we now have sufficient evidence for the existence of two tracks in the hard state BHXRB radio correlation (see section 2.3.6). However, we still do not understand how these tracks behave at lower luminosities. They may completely merge at some point as a result of all sources eventually undergoing transitions like that of H1743-322, perhaps with significant scatter in the turning point of each source or all sources, or some sources may never leave a single track. As stated at the end of Chapter 4, our ability to decipher the radio/X-ray correlation(s) in hard state BHXRBs is limited mainly by the number of simultaneous observation points available to plot, which in turn, is limited by the number of BHXRBs available, and the sensitivities to which we can observed them. Our lack of information at lower luminosities is unsurprising given that sources evidently become harder to detect, but also some systems, like GRS

1915+105, currently appear to spend little or no time at quiescent luminosities, which further limits our available targets in these regions.

Our ability to detect the current list of known BHXRBs is directly related to the sensitivity of our instruments and can only improve with time, though given the surge in new upgrades and new instruments, as well as the potential power of the SKA, the sensitivity to faint sources will improve significantly over the next decade. However, the possibility of finding new BHXRB sources always exists, each with the potential to add several new points to the correlation over multiple observations. The most useful sources will obviously be those which are nearer to us, since it will be easier to observe them while they reside in their faintest states (for example, A0620-00 at 1.2 kpc). Early in 2011 a new XRB, J1357.2-0933, was detected and evidence has been found that indicates the system not only resides nearby ($\sim$ 1.5 kpc; Rau, Greiner & Filgas 2011), but likely harbours a black hole (Casares *et al.* 2011; Corral-Santanta *et al.* in prep.) making it an excellent target for correlation measurements. Furthermore, J1357.2-0933 appears to be sub-luminous in the radio band which may suggest it can reside on the 'radio quiet' branch of the relation, and when last observed was decaying in luminosity, perhaps towards quiescence. Given these facts, I happily contributed to an accepted ATCA-CABB proposal to observe the system in June, together with simultaneous *Chandra* observations, in the hope that we can detect the system in quiescence and continue to fill the low luminosity void of the correlation plot. Of course, if the radio quiet branch is as steep as implied by Gallo, Miller and Fender (2012), it will be very difficult to detect the system at low luminosities (if it remains on said branch); although, a detection of the source on the normal branch, during a transition between the two, or even as an additional upper limit (as in the cases of J1655-40 and J1550-564 in Chapter 4) will still be a significant contribution to the existing plot.

7.2.3 The twisted jets of Circinus X-1

I have found significant evidence that Cir X-1's jets have not only changed trajectory in the last few years, but show different axes depending on the distance from the source at which components are observed; however, the cause of these effects remains a mystery. As I've discussed, precession could be behind the effect, thus having significant implications for the notion that XRB jet parameters can remain persistent over secular timescales, which has already been questioned in terms of the formation of SS 433's nebula's distinct shape (Goodall, Fathallah & Blundell 2011). Of course, it is also possible for the shift to be explained by deviations of

the jet in a particularly dense medium or unstable flows (see Section 2.2.3.7). The inability to distinguish between the two scenarios is not helped by the fact that my cm and mm observations are separated by almost a year: sufficient time in which for any precession to shift the direction of the jet.

Therefore, in order to resolve this issue, I led a proposal to re-observe Cir X-1 using ATCA-CABB at a range of wavelengths; 16cm, 6cm, 3cm and 7mm, over only several days in order to minimise the effects of any existing precession. This way, if the jets are indeed twisted, the difference in jet axis would remain visible between images at different frequencies. I would also be able to compare the images made from the new data to those of the existing cm and mm observations presented in this thesis to determine if any change has occurred in the intervening time that would imply precession in the long term. The proposal was accepted and observations were carried out in December of 2011; however, at the time of writing I have yet to have sufficient time to reduce and analyse the data thoroughly. I can only present an image made from the 16cm data in Figure 7.1, making it the most detailed map of Cir X-1's jet powered nebula currently available. Once again it is possible to see the core, south-western component, X-ray shock caps and the filament like structures discussed in Chapter 5, but better defined against background noise. Furthermore, there are tantalising hints of a curved structure within the nebula, extending towards the south from the central core, which is reminiscent of the 'curved jet' structure observed in Tudose *et al.* (2006; see my Figure 5.13). A lot of work can still be done on this image as well as the remaining data from the other frequencies, but it is my hope that they can help in understanding the new behaviour exhibited by Cir X-1 in recent years.

7.2.4 The mm emission of X-ray binaries

While the inherent errors involved in mm observations make it difficult to calculate accurate spectral indices between 33 and 35 GHz, it is still possible to re-observe both Sco X-1 and Cir X-1 at cm and mm frequencies near simultaneously to confirm the presence of optically thin emission from the cores of these systems. In particular, I cannot be confident in my estimates for the spectrum of Cir X-1 in Chapter 6, given the lack of simultaneous observations at the two wavelengths. However, the new multi-frequency observations of Cir X-1 mentioned in the previous subsection can not only be used in solving the jet behaviour mystery as originally intended, but also in providing a new estimate for the source's radio to mm spectrum. Furthermore, since I have shown that some NSXRBs can be detected in the mm regime,

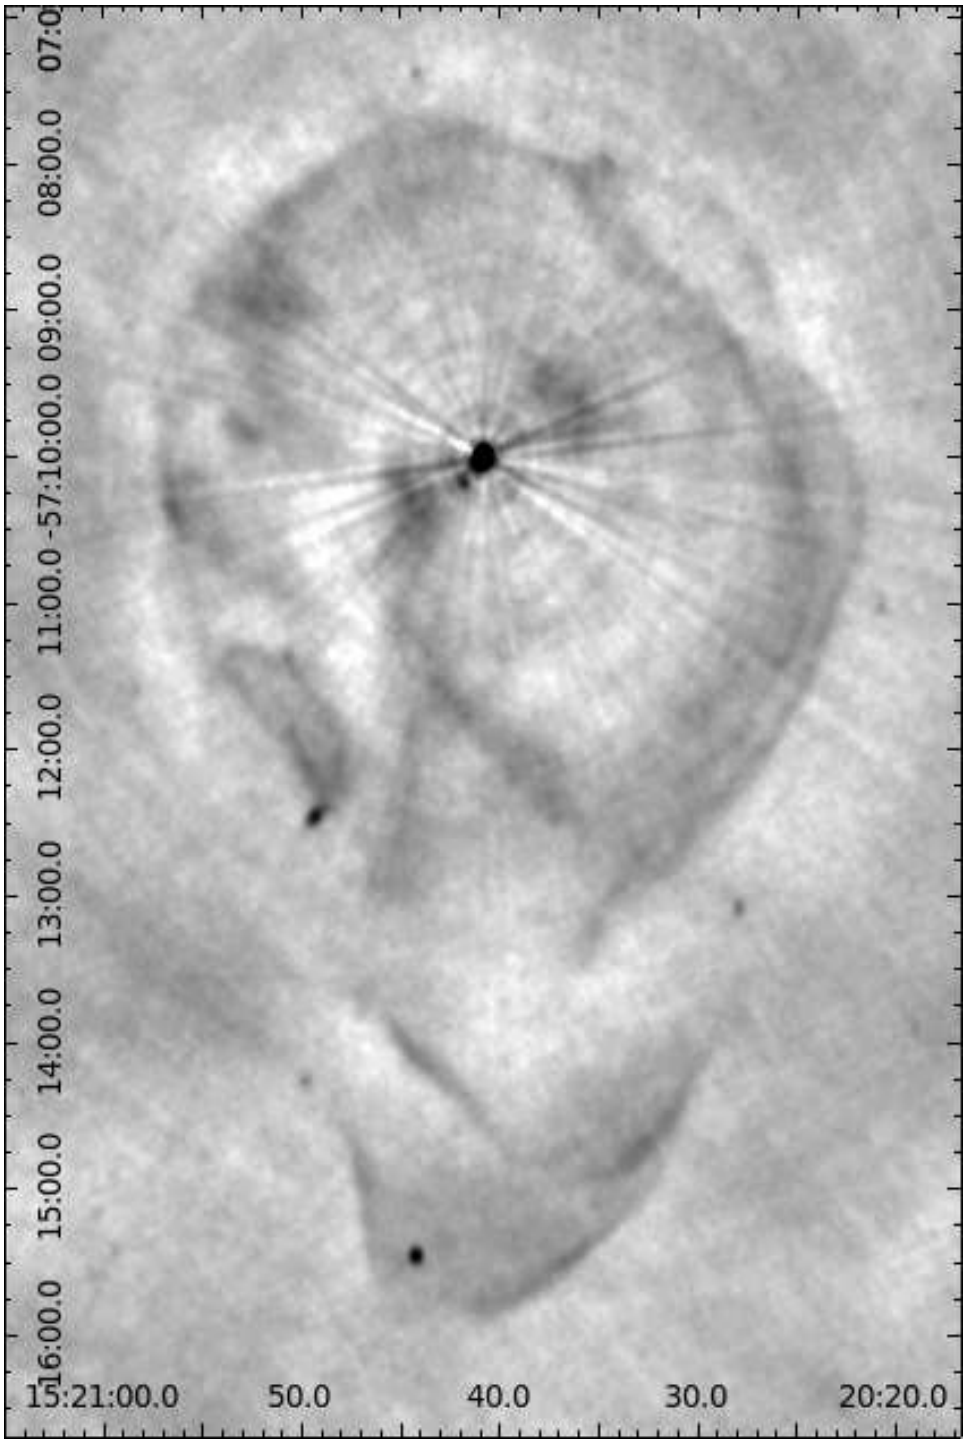

Figure 7.1: Latest 3.1 GHz (16cm) radio map of Circinus X-1's jet powered nebula. Beam size = 4.96 $\times$ 3.96 arcsec^2. Note that the reduction of the data used to make this map was relatively rushed (for example it might be possible to eliminate the streak effect radiating from Cir X-1's core with additional work), and no primary beam correction has been applied.

it is now sensible to explore the possibility of observing other NSXRBs at mm wavelengths using ATCA-CABB and other arrays wherever possible. Any such campaigns should also attempt to target both atoll and Z-source systems in order to further explore the divide in jet behaviours discussed in Chapter 6.

Based on the predicted noise of ATCA-CABB 33 and 35 GHz observations, a 12 hour run should theoretically allow the production of images with noise levels of approximately 0.03 mJy beam^{-1} (though phase errors raise this value closer to a source, as was observed in Chapter 6), meaning a 5σ detection of an NSXRB would require their mm emission be at least $\sim$ 0.15 mJy. By comparing this estimate to the measured 8.5 GHz flux densities of NSXRBs in table 1 of Migliari & Fender (2006), one can find a list of possible targets for ATCA-CABB studies. If I assume atoll sources do indeed have flat radio spectra all the way to the mid-IR (as implied by the spectrum of 4U 0614+091) then there are 4 known atoll sources which should be detectable by ATCA-CABB during some X-ray states; 4U 1728-34, Aql X-1, 4U 1608-52 and MXB 1730-335. However, if any of these sources do display optically thin spectra, they will likely fall below the detection threshold. For Z-source spectra, I should assume that the sources may have an optically thin spectrum of at least α = -1, in which case there are 5 Z-source NSXRBS which may still be detectable with ATCA-CABB: GX 17+2, GX 349+2, GX 5-1, GX 340+0 and GX 13+1, two of which have been sufficiently bright in the radio that they should remain detectable even with indices approaching α = -1.5 (GX 5-1 and GX 13+1).

Of course, if using E-VLA, then the limit for northern sources is slightly better and additional targets can be considered. Furthermore, with the introduction of the SKA in coming decades none of the currently known NSXRBs will go undetected for long.

7.3 Closing comments

Radio observation of X-ray binaries continues to revolutionise our understanding of relativistic jets and their connection to accretion mechanics. With the plethora of recent upgrades to existing arrays and new instruments appearing every few years, the range of possible jet studies is constantly expanding, and though we learn more about the physical processes behind these phenomena with every investigation, we are also faced with new mysteries as our group of targets becomes larger and their full range of behaviours easier to observe. ATCA has remained one of the primary tools in our efforts to explore XRB radio emission over the past few decades, with

the CABB upgrade giving the array a fresh lease of life and stimulating a surge in scientific output via the instrument. Though the SKA pathfinders, ASKAP and MeerKAT, and the final SKA itself will eventually overshadow it, ATCA-CABB will remain at the forefront of interferometric observations in the Southern Hemisphere for years to come, and will likely continue providing excellent astrophysical data as part of a much larger network of radio instruments beyond that.

However, one must be ever aware of the potential errors that can arise in imaging of variable sources such as XRBs with ATCA, and indeed any radio telescope to some degree. Though many tools may exist which can aid in reducing the adverse effects of such source behaviour, an element of human involvement will likely always be necessary to determine when, where and to what degree these should be used, and whether what is left is indeed real. Radio imaging will continue to help us plumb the depths of the most powerful astrophysical objects in the Universe, just as long as observers remember ‘quick and dirty’ gets them exactly that.

APPENDICES

Bazinga

DR. SHELDON COOPER, THE BIG BANG THEORY

A

VaSpecSim

Import necessary external subroutines and define input check subroutine.

```
import math
def inputtest(prompt,inputtype=str,inputlength=1,subtype=str):
    """Used to test value entry for correct type and length"""
    running=True
    while running:
        try:
            global Value
            value=input(prompt)
            if type(value)==int and inputtype==float:
                value=float(value)
            if type(value)==tuple:
                temp=range(len(value))
                for i in range(len(value)):
                    if subtype==float and type(value[i])==int:
                        temp[i]=float(value[i])
                    if subtype==float and type(value[i])==float:
                        temp[i]=float(value[i])
```

```
            value=tuple(temp)
            if type(value)==inputtype:
                if len(value)==inputlength:
                    for i in range(len(value)):
                        if type(value[i])!=subtype:
                            if subtype==str:
                                expect='text'
                            if subtype==float:
                                expect='a number'
                            if subtype==int:
                                expect='an integer'
                            print 'Value %s was not entered correctly, expected %s,
                            try again...' % (str(int(i)+1),expect)
                            running=True
                            break
                        else:
                            if type(value[i])==subtype:
                                running=False
                                Value=value
            if len(value)!=inputlength:
                print 'That is not the correct number of parameters (expected %s),
                try again...' % inputlength
        else:
            if type(value)!=inputtype:
                if inputtype==str:
                    expected='text in quotes'
                if inputtype==int:
                    expected='an integer'
                if inputtype==float:
                    expected=='a number'
                    print 'Invalid entry, expected %s, try again...' % expected
            else:
                Value=value
                running=False
    except NameError:
        print "Something wasn't in quotes, try again..."
    except SyntaxError:
```

```
        print ‘’
```

Define VaSpecSim and request information about the model.

```
def VaSpecSim():
    """Compile a script for use in miriad to build spectrally variable (over time)
    source models"""

    e=2.71828182845904523536
    print ‘’
    print ‘Time to give me some initial variables...’
    print ‘’
    print “Please note, text must be entered in quotes e.g. ‘flibble’. Numbers may be
    entered without quotes”
    print ‘’

    inputtest(‘Enter a name for this model : ’)
    name = Value
    print ‘’
    print ‘If you are unsure what value to enter for a particular parameter enter xx to
    have program set default values.’
    print ‘’

    inputtest(“Enter telescope position or name (e.g. ‘atca’) : ”)
    telescop = Value
    inputtest(‘Enter an antennae configuration (look in
    /disks/jets/blue/rpf/DATA/ATCA-antenna-locations) : ’)
    ant = Value
    inputtest(‘Enter a base unit value : ’,float)
    baseunit = Value
    inputtest(‘Enter central pointing position in sky : ’,tuple,2,float)
    radec = Value
    inputtest(‘Enter telescope lattitude : ’,float)
    lat = Value
```

Define frequency bandwidth and resolution, as well as the resolution of spectral variance (i.e. time segments).

```
inputtest('Enter minimum for frequency band (Ghz) : ',float)
vlow = Value
inputtest('Enter maximum for frequency band (Ghz) : ',float)
vhigh = Value
running = True
while running:
    inputtest('Enter step between frequency windows (Ghz). Note, this number
    must be such that it divides the bandwidth into an integer value of windows
    : ',float)
    test=(float(vhigh)-float(vlow))/float(Value)
    if int(test)==test:
        vstep=Value
        running=False
    else:
        print 'Error: not a valid step, Bandwidth/step is not an integer, please try
        again'
inputtest('Enter total time that the simulation data should cover(hours) : ',float)
time = Value
running=True
while running:
    inputtest('Enter time step (hours) i.e. this will be the maximum possible
    resolution for your spectral variance. Note, this number must be such that
    time/timestep equals an integer : ',float)
    test=float(time)/float(Value)
    if int(test)==test:
        timestep=Value
        running=False
    else:
        print 'Error: not a valid timestep, time/timestep is not an integer, please
        try again'
```

Define source parameters.

```
print ‘’
print ‘Now to define your sources...’
print ‘’

inputtest(‘Enter number of point sources to include : ’,int)
npoint = Value
inputtest(‘Enter number of diffuse sources (large scale gaussian peaks) to
include : ’,int)
ndiffuse = Value
pointcount = range(int(npoint))
diffusecount = range(int(ndiffuse))
pointralist = range(int(npoint))
pointdeclist = range(int(npoint))
diffuseralist = range(int(ndiffuse))
diffusedeclist = range(int(ndiffuse))
for i in pointcount:
    print ‘Enter RA (relative to image centre) for point source’,str(i+1),‘:’
    inputtest(‘’,float)
    pointralist[i]=Value
    print ‘Enter DEC (relative to image centre) for point source’,str(i+1),‘:’
    inputtest(‘’,float)
    pointdeclist[i]=Value
for i in diffusecount:
    print ‘Enter RA (relative to image centre) for diffuse source’,str(i+1),‘:’
    inputtest(‘’,float)
    diffuseralist[i]=Value
    print ‘Enter DEC (relative to image centre) for diffuse source’,str(i+1),‘:’
    inputtest(‘’,float)
    diffusedeclist[i]=Value

print ‘’
print ‘You must now define size and orientation of the diffuse sources.’
print ‘’

diffusebmajlist = range(int(ndiffuse))
```

```
diffusebminlist = range(int(ndiffuse))
diffusebpalist = range(int(ndiffuse))
for i in diffusecount:
    print ‘Enter FWHM for major axis of diffuse source’,str(i+1),‘:’
    inputtest(‘’,float)
    diffusebmajlist[i]=Value
    print ‘Enter FWHM for minor axis of diffuse source’,str(i+1),‘:’
    inputtest(‘’,float)
    diffusebminlist[i]=Value
    print ‘Enter position angle of major axis (measured from North to East) for
    diffuse source’,str(i+1),‘:’
    inputtest(‘’,float)
    diffusebpalist[i]=Value
```

Calculation of time and frequency segment parameters.

```
nslice = float(time)/float(timestep)
print ‘’
print ‘Your entered values for time and timestep have resulted in ’+str(int(nslice)),‘time
slices. You will be asked to enter the spectral form for each
source in each time slice.’
print ‘’
timesteps = range(int(nslice))
for i in timesteps:
    timesteps[i]=-6+i*float(timestep)
nvslice = (float(vhigh)-float(vlow))/float(vstep)
vsteps = range(int(nvslice))
for i in vsteps:
    vsteps[i]=float(vlow)+i*float(vstep)
if vsteps[0]==0:
    vsteps[0]=0.00000001
nsource = int(npoint)+int(ndiffuse)
modlist = range(len(timesteps))
for j in modlist:
    modlist[j] = range(len(vsteps))
    for k in modlist[j]:
        modlist[j][k] = range(nsource)
```

This part of the program can become rather complex and long winded if the model has significant variability resolution. What is requested from the user is the definition of the spectrum (number of power laws, gaussians and parameters), for each source, in each time segment.

```
for i in range(len(timesteps)):
    print ‘’
    print ‘For time slice between ha range ’+str(timesteps[i]),‘to ’+str(float(timesteps[i])
    +float(timestep)),‘...’
    print ‘’
    if (float(timesteps[i])+float(timestep)) < 0.001 and (float(timesteps[i])+float(timestep))
    > -0.001:
        filenameext = ‘%s.%s’ % (str(timesteps[i]),‘0.0’)
    else:
        filenameext = ‘%s.%s’ % (str(timesteps[i]),str(float(timesteps[i])+float(timestep)))

    for j in range(nsource):
        if (j+1) <= npoint:
            sourcetype = ‘point’
            sourcenum = j+1
        else:
            sourcetype = ‘diffuse’
            sourcenum = (j+1)-int(npoint)
        print ‘     For %s source %s...’ % (sourcetype,sourcenum)
        print ‘’

        print ‘Enter number of power laws (of form y=Axˆ B+C) to include
        in model’
        inputtest(‘’,int)
        npower = Value
        print ‘Enter number of gaussian forms (y=Aeˆ((-(x-b)ˆ2)/2(cˆ2))) to
        include in model’
        inputtest(‘’,int)
        ngauss = Value
        powerpar = range(int(npower))
        gausspar = range(int(ngauss))
```

```
for k in range(npower):
    inputtest('Enter Amplitude A, power B and verticle axis cross value C
    for power law '+str(k+1)+" :",tuple,3,float)
    powerpar[k] = Value
for k in range(ngauss):
    inputtest('Enter amplitude A, position B and width C for gaussian '+str(k+1)+
    " :",tuple,3,float)
    gausspar[k] = Value


fluxlist = range(len(vsteps))
for k in range(len(vsteps)):
    powercomp = 0
    gausscomp = 0
    for l in range(npower):
        x = vsteps[k]
        y = float(powerpar[l][0])*(float(x)**float(powerpar[l][1]))+float(powerpar[l][2])
        powercomp = powercomp+float(y)
    for l in range(ngauss):
        x = float(vsteps[k])
        a = float(gausspar[l][0])
        b = float(gausspar[l][1])
        c = float(gausspar[l][2])
        y = a*(e**(-((x-b)**2.0)/(2.0*(c**2.0))))
        gausscomp = gausscomp+float(y)
    fluxlist[k] = float(powercomp)+float(gausscomp)


    if sourcetype == 'point':
        modlist[i][k][j] = '%s,%s,%s' % (fluxlist[k],pointralist[int(sourcenum)
        -1],pointdeclist[int(sourcenum)-1])
    else:
        modlist[i][k][j] = '%s,%s,%s,%s,%s,%s' % (fluxlist[k],diffuseralist[int(sourcenum)
        -1],diffusedeclist[int(sourcenum)-1],diffusebmajlist[int(sourcenum)
        -1],diffusebminlist[int(sourcenum)-1],diffusebpalist[int(sourcenum)
        -1])
```

All input is complete, the program now compiles the script lines, first defining the source model text files used in UVGEN*, the actual* UVGEN *commands, and finally the* UVAVER *commands (dividing a large number of files into manageable numbers to prevent errors), cleaning up files as they become redundant.*

```
f=open('Model.%s.Script' % name,'w')
for i in range(len(timesteps)):
    for j in range(nsource):
        if (j+1) <= npoint:
            sourcetype = 'point'
            sourcenum = j+1
        else:
            sourcetype = 'diffuse'
            sourcenum = (j+1)-int(npoint)
        f=open('Model.%s.source.%s%s.dat' % (name,sourcetype,sourcenum),'a')
        f.write('Model %s %s source %s\n' % (name,sourcetype,sourcenum))
        for k in range (len(vsteps)):
            f.write('%s,%s,%s\n' % (str(i),str(vsteps[k]),str(modlist[i][k][j])))
        f.close()
for i in range(len(timesteps)):
    if (float(timesteps[i])+float(timestep)) < 0.001 and (float(timesteps[i])+float(timestep))
    > -0.001:
        filenameext = '%s.%s' % (str(timesteps[i]),'0.0')
    else:
        filenameext = '%s.%s' % (str(timesteps[i]),str(float(timesteps[i])+float(timestep)))
    f=open('Model.%s.Script' % name,'a')
    f.write('\n')
    f.close()
    for j in range(len(vsteps)):
        for k in range(nsource):
            if k==0:
                print 'echo %s > %s.%s.mod' % (str(modlist[i][j][k]),str(filenameext),
                str(vsteps[j]))
                f=open('Model.%s.Script' % name,'a')
                f.write('echo %s > %s.%s.mod\n' % (str(modlist[i][j][k]),str(filenameext),
                str(vsteps[j])))
                f.close()
```

```
            else:
                f=open('Model.%s.Script' % name,'a')
                print 'echo %s >> %s.%s.mod' % (str(modlist[i][j][k]),str(filenameext),
                str(vsteps[j]))
                f.write('echo %s >> %s.%s.mod\n' % (str(modlist[i][j][k]),str(filenameext),
                str(vsteps[j])))
                f.close()
print ''
f=open('Model.%s.Script' % name,'a')
f.write('\n')
f.close()

numuvaver = float(len(vsteps))/20
if int(numuvaver)!=float(numuvaver):
    numuvaver=int(numuvaver+1)
else:
    numuvaver=int(numuvaver)
finaloutlist=range(len(timesteps)*numuvaver)
for i in range(len(timesteps)):
    if (float(timesteps[i])+float(timestep)) < 0.001 and (float(timesteps[i])+float(timestep))
    > -0.001:
        filenameext = '%s.%s' % (str(timesteps[i]),'0.0')
    else:
        filenameext = '%s.%s' % (str(timesteps[i]),str(float(timesteps[i])+float(timestep)))
    f=open('Model.%s.Script' % name,'a')
    f.write('\n')
    f.close()
    for j in range(len(vsteps)):
        print 'uvgen ant=%s baseunit=%s telescop=%s corr=0,1,0,128 radec=%s
        ,%s lat=%s harange=%s,%s freq=%s,0.0 source=%s.%s.mod out=%s.v%s.uv'
        % (ant,baseunit,telescop,radec[0],radec[1],lat,str(timesteps[i]),str(float(timesteps[i])+
        float(timestep)),vsteps[j],str(filenameext),str(vsteps[j]),str(filenameext),str(j))
        f=open('Model.%s.Script' % name,'a')
        f.write('uvgen ant=%s baseunit=%s telescop=%s corr=0,1,0,128 radec=%s
        ,%s lat=%s harange=%s,%s freq=%s,0.0 source=%s.%s.mod out=%s.v%s.uv\n'
        % (ant,baseunit,telescop,radec[0],radec[1],lat,str(timesteps[i]),str(float(timesteps[i])+
        float(timestep)),vsteps[j],str(filenameext),str(vsteps[j]),str(filenameext),str(j)))
```

```
        f.close()
    f=open('Model.%s.Script' % name,'a')
    f.write('\nrm -rf ./%s*mod\n' % str(timesteps[i]))
    f.close()
for i in range(len(timesteps)):
    if (float(timesteps[i])+float(timestep)) < 0.001 and (float(timesteps[i])+float(timestep))
    > -0.001:
        filenameext = '%s.%s' % (str(timesteps[i]),'0.0')
    else:
        filenameext = '%s.%s' % (str(timesteps[i]),str(float(timesteps[i])+float(timestep)))
    outlist=range(len(vsteps))
    for j in range(len(vsteps)):
        outlist[j]='%s.v%s.uv' % (str(filenameext),str(j))
    for j in range(numuvaver):
        if j==(numuvaver-1):
            k=len(vsteps)-(numuvaver-1)*20
        else:
            k=20
        print 'uvaver vis=%s out=%s.a%s.uv' % (','.join(outlist[0+j*20:20+(j-1)*20+
        k]),str(filenameext),str(j+1))
        f=open('Model.%s.Script' % name,'a')
        f.write('\nuvaver vis=%s out=%s.a%s.uv\n' % (','.join(outlist[0+j*20:20+(j
        -1)*20+k]),str(filenameext),str(j+1)))
        f.close()
        print ''
        finaloutlist[i*int(numuvaver)+j]='%s.a%s.uv' % (str(filenameext),str(j+1))
    f=open('Model.%s.Script' % name,'a')
    f.write('\nrm -rf ./%s*v*.uv\n' % str(timesteps[i]))
    f.close()
print ''
numfinaluvaver = float(len(timesteps)*numuvaver)/20
if int(numfinaluvaver)!=float(numfinaluvaver):
    numfinaluvaver=int(numfinaluvaver+1)
else:
    numfinaluvaver=int(numfinaluvaver)
for k in range(numfinaluvaver):
    if k==(numfinaluvaver-1):
```

```
        l=len(timesteps)*numuvaver-(numfinaluvaver-1)*20
    else:
        l=20
    print 'uvaver vis=%s out=%s%s.uvfinal' % (',',join(finaloutlist[0+k*20:20+(k
    -1)*20+l]),name,k+1)
    f=open('Model.%s.Script' % name,'a')
    f.write('\n\nuvaver vis=%s out=%s%s.uvfinal\n' % (',',join(finaloutlist[0+k*
    20:20+(k-1)*20+l]),name,k+1))
    f.close()
print ''
print 'rm -rf ./*uv'
f=open('Model.%s.Script' % name,'a')
f.write('\n\nrm -rf ./*uv\n\n')
f.close()
```

B

Example VaSpecSim response template

‘P.5gradx2’ *-Model name*
‘atca’ *-Telescope for use in model*
‘/disks/jets/blue/rpf/DATA/ATCA-antenna-locations/3.0c.ant’ *-Location of telescope array layout file*
51.0204 *-’baseunit’ value for uvgen, i.e. Conversion factor for distance values used in the layout file.*
0,-90 *-pointing RA and DEC*
-30 *-telescope Latitude*
4.5 *-Lower frequency limit*
6.5 *-Upper frequency limit*
0.004 *-Increment between frequencies*
12 *-Observation time*
0.6 *-Increment between spectral ‘steps’*
1 *-Number of point sources*
0 *-Number of diffuse (Gaussian) sources*
0 *-Source RA.*
0 *-Source DEC. Note: if there was more than one source than there would be several more lines like this and the above, prior to the next variable.*
X,X *-If diffuse sources were included then their parameters would be entered now.*

Now, for EACH time segment (a loop until all time segments, observation time divided by increment, i.e. 36 loops in this case, have been addressed) one enters the parameters defining the spectral shape of the source:

1 *-Number of power laws for the spectral model*

1 *-Number of Gaussian profiles for the spectral model*

1,1,1 *Power law parameters.*

1,1,1 *Gaussian parameters.*

C

VaSpecSim model parameters

Table C.1: VaSpecSim model parameters. From left to right I list: Label associated with each model, the name of the model (based on variation type and index values), start and end values for spectral index ∇, pattern of variation over the 20×36 minute time segments, flux density of the final image measured via a point source fit, the dynamic range used in corresponding images shown in Figures 3.3 through 3.5, the minimum measured flux density of image residuals after point source subtraction, and finally the maximum measured flux density of image residuals after point source subtraction.

Label	Name	Start/end ∇ values	Spectral index variation	Fitted S_ν (Jy)	Image range (mJy)	Residuals min S_ν (mJy)	Residuals max S_ν (mJy)
A	Stable0	$\pm$0	none	10.00	$\pm$0.05	-0.02	0.04
B	Stable1	1:1	none	10.00	$\pm$1	-0.98	0.87
B*	Stable-1	-1:-1	none	10.00	$\pm$1	-0.78	1.01
C	Stable2	2:2	none	10.00	$\pm$1	-1.94	1.77
C*	Stable-2	-2:-2	none	10.00	$\pm$1	-1.63	2.04
D	Stable3.5	3.5:3.5	none	9.99	$\pm$3	-3.48	2.91
D*	Stable-3.5	-3.5:-3.5	none	10.01	$\pm$3	-2.83	3.43
E	Gradual0to1	0.05:1	+0.05/0.6hrs	10.00	$\pm$5	-6.06	5.04
E*	Gradual0to-1	-0.05:-1	-0.05/0.6hrs	10.00	$\pm$5	-4.47	4.85
F	Gradual0to2	0.1:2	+0.1/0.6hrs	9.99	$\pm$10	-10.23	9.71
F*	Gradual0to-2	-0.1:-2	-0.1/0.6hrs	10.00	$\pm$10	-9.19	9.96
G	Gradual0to3.5	0.175:3.5	+0.175/0.6hrs	9.99	$\pm$15	-16.37	16.26
G*	Gradual0to-3.5	-0.175:-3.5	-0.175/0.6hrs	10.00	$\pm$15	-15.90	16.87

Table C.1: (continued from last page)

Label	Name	Start/end ∇ values	Spectral index variation	Fitted S_ν (Jy)	Image range (mJy)	Residuals min S_ν (mJy)	Residuals max S_ν (mJy)
H	Gradual-1to1	-1:1	+0.1/0.6hrs skip ∇=0	10.00	±10	-9.25	9.42
H*	Gradual1to-1	1:-1	-0.1/0.6hrs skip ∇=0	10.00	±10	-9.18	9.43
I	Gradual-2to2	-2:2	+0.2/0.6hrs skip ∇=0	10.00	±20	-21.60	19.95
I*	Gradual2to-2	2:-2	-0.2/0.6hrs skip ∇=0	9.99	±20	-20.16	19.60
J	Gradual-3.5to3.5	-3.5:3.5	+0.35/0.6hrs skip ∇=0	9.99	±30	-32.30	32.31
J*	Gradual3.5to-3.5	3.5:-3.5	-0.35/0.6hrs skip ∇=0	10.00	±30	-37.61	33.62
K	Burst0to1	0:0	none for 4.2hrs +0.5/0.6hrs for 1.2hrs none for 1.8hrs -0.5/0.6hrs for 1.2hrs none for 4.2hrs	10.00	±8	-10.10	12.45

Table C.1: (continued from last page)

Label	Name	Start/end ∇ values	Spectral index variation	Fitted S_ν (Jy)	Image range (mJy)	Residuals min S_ν (mJy)	Residuals max S_ν (mJy)
K*	Burst0to-1	0:0	none for 4.2hrs -0.5/0.6hrs for 1.2hrs none for 1.8hrs +0.5/0.6hrs for 1.2hrs none for 4.2hrs	10.00	±8	-11.97	10.10
L	Burst0to2	0:0	none for 4.2hrs +1/0.6hrs for 1.2hrs none for 1.8hrs -1/0.6hrs for 1.2hrs none for 4.2hrs	9.99	±15	-18.62	20.99
L*	Burst0to-2	0:0	none for 4.2hrs -1/0.6hrs for 1.2hrs none for 1.8hrs +1/0.6hrs for 1.2hrs none for 4.2hrs	10.00	±15	-23.04	19.72

Table C.1: (continued from last page)

Label	Name	Start/end ∇ values	Spectral index variation	Fitted S_ν (Jy)	Image range (mJy)	Residuals min S_ν (mJy)	Residuals max S_ν (mJy)
M	Burst0to3.5	0:0	none for 4.2hrs +1.75/0.6hrs for 1.2hrs none for 1.8hrs -1.75/0.6hrs for 1.2hrs none for 4.2hrs	9.99	±30	-33.84	40.64
M*	Burst0to-3.5	0:0	none for 4.2hrs -1.75/0.6hrs for 1.2hrs none for 1.8hrs +1.75/0.6hrs for 1.2hrs none for 4.2hrs	10.00	±30	-41.16	36.34
N	Burst-1to1	-1:-1	none for 4.2hrs +1/0.6hrs for 1.2hrs none for 1.8hrs -1/0.6hrs for 1.2hrs none for 4.2hrs	10.00	±15	-17.94	21.15

Table C.1: (continued from last page)

Label	Name	Start/end ∇ values	Spectral index variation	Fitted S_ν (Jy)	Image range (mJy)	Residuals min S_ν (mJy)	Residuals max S_ν (mJy)
N*	Burst1to-1	1:1	none for 4.2hrs -1/0.6hrs for 1.2hrs none for 1.8hrs +1/0.6hrs for 1.2hrs none for 4.2hrs	10.00	±15	-23.35	20.38
O	Burst-2to2	-2:-2	none for 4.2hrs +2/0.6hrs for 1.2hrs none for 1.8hrs -2/0.6hrs for 1.2hrs none for 4.2hrs	10.00	±30	-39.97	45.15
O*	Burst2to-2	2:2	none for 4.2hrs -2/0.6hrs for 1.2hrs none for 1.8hrs +2/0.6hrs for 1.2hrs none for 4.2hrs	10.00	±30	-46.03	40.16

Table C.1: (continued from last page)

Label	Name	Start/end ∇ values	Spectral index variation	Fitted S_ν (Jy)	Image range (mJy)	Residuals min S_ν (mJy)	Residuals max S_ν (mJy)
P	Burst-3.5to3.5	-3.5:-3.5	none for 4.2hrs +3.5/0.6hrs for 1.2hrs none for 1.8hrs -3.5/0.6hrs for 1.2hrs none for 4.2hrs	10.00	±50	-67.26	76.66
P*	Burst3.5to-3.5	3.5:3.5	none for 4.2hrs -3.5/0.6hrs for 1.2hrs none for 1.8hrs +3.5/0.6hrs for 1.2hrs none for 4.2hrs	10.00	±50	-76.44	66.05

BIBLIOGRAPHY

Abramowicz, M.A., Czerny, B., Lasota, J.P., & Szuszkiewicz, E., 1988, ApJ, 332, 646

Altamirano, D., van der Klis, M., Méndez, M., Jonker, P.G., Klein-Wolt, M., & Lewin, W.H.G., 2008, ApJ, 685, 436

Andrew, B.H., & Purton, C.R., 1968, Nature, 218, 855

Arnett, W.D., & Bowers, R.L., 1977, ApJS, 33, 415

Asai, K., Dotani, T., Hoshi, R., Tanaka, Y., Robinson, C.R., & Terada, K., 1998, PASJ, 50, 611

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

Bardeen, J.M., & Petterson, J.A., 1975, ApJ, 195, L65

Beer, M.E., & Podsiadlowski, P., 2002, MNRAS, 331, 351

Belloni, T., Homan, J., Casella, P., van der Klis, M., Nespoli, E., Lewin, W.H.G., Miller, J.M., & Méndez, M., 2005, A&A, 440, 207

Belloni, T., 2007, MmSAI, 78, 652

Blandford, R.D., & Königl, A., 1979, ApJ, 232, 34

Blandford, R.D., & Payne, D.G., 1982, MNRAS, 199, 883

Bombaci, I., Prakash, M., Prakash, M., Ellis, P.J., Lattimer, J.M., & Brown, G.E., 1995, Nuclear Physics A, 583, 623

Bowen, E.G., 1966, The Physics Teacher, 4, 99

Bradshaw, C.F., Geldzahler, B.J., & Fomalont, E.B., 1997, ApJ, 481, 489

Brandt, W.N., & Schulz, N.S., 2000, ApJ, 544, L123

Brocksopp, C., Corbel, S., Fender, R.P., Rupen, M., Sault, R., Tingay, S.J., Hannikainen, D., & O'Brien, K., 2005, MNRAS, 356, 125

Cadolle-Bel, M., Ribó, M., Rodriguez, J., Chaty, S., Corbel, S., Goldwurm, A., Frontera, F., Farinelli, R., & 6 coauthors, 2007, ApJ, 659, 549

Calvelo, D.E., Fender, R.P., Russell, D.M., Gallo, E., Corbel, S., Tzioumis, A.K., Bell, M.E., Lewis, F., & Maccarone, T.J., 2010a, MNRAS, 409, 839

Calvelo, D.E., Fender, R., Broderick, J., Moin, A., Tingay, S., Tzioumis, A.K., & Nicolson, G., 2010b, ATel, 2699, 1

Calvelo, D.E., Fender, R.P., Tzioumis, A.K., Kawai, N., Broderick, J.W., & Bell, M.E., 2012a, MNRAS, 419, 436

Calvelo, D.E., Fender, R.P., Tzioumis, A.K., & Broderick, J.W., 2012b, MNRAS, 419, L54

Campana, S., Parmar, A.N., & Stella, L., 2001, A&A, 372, 241

Cannizzo, J.K., Ghosh, P., & Wheeler, J.C., 1982, ApJ, 260, L83

Casares, J., Torres, Manuel A.P., Negueruela, I., Gonzalez-Fernandez, C., Corral-Santana, J.M., Zurita, C., Llano, S.R., 2011, ATel, 3206, 1

Clark, D.H., Parkinson, J.H., & Caswell, J.L., 1975, Nature, 254, 674

Clark, B.G., 1980, A&A, 89, 377

Conway, J.E., Cornwell, T.J., & Wilkinson, P.N., 1990, MNRAS, 246, 490

Corbel, S., Fender, R.P., Tzioumis, A.K., Nowak, M., McIntyre, V., Durouchoux, P., & Sood, R., 2000, A&A, 359, 251

Corbel, S., Fender, R.P., Tzioumis, A.K., Tomsick, J.A., Orosz, J.A., Miller, J.M., Wijnands, R., & Kaaret, P., 2002, Science, 298, 196

Corbel, S., & Fender, R.P., 2002, ApJ, 573, L35

Corbel, S., Nowak, M.A., Fender, R.P., Tzioumis, A.K., & Markoff, S., 2003, A&A, 400, 1007

Corbel, S., Fender, R.P., Tomsick, J.A., Tzioumis, A.K., & Tingay, S., 2004, ApJ, 617, 1272

Corbel, S., Tomsick, J.A., & Kaaret, P., 2006, ApJ, 636, 971

Corbel, S., Körding, E., & Kaaret, P.,2008, MNRAS, 389, 1697

Coriat, M., Corbel, S., Prat, L., Miller-Jones, J.C.A., Cseh, D., Tzioumis, A.K., Brocksopp, C., Rodriguez, J., G.R., & 2 coauthors, 2011, MNRAS, 414, 677

Corral-Santana, J.M., *et al.* , in prep.

Cowley, A.P., & Crampton, D., 1975, ApJ, 201, L65

Cummings, A.C., Stone, E.C., & Vogt, R.E., 1973, Proceedings of the 13th International Conference on Cosmic Rays, 1, 335

Davidson, K., Pacini, F., & Salpeter, E.E., 1971, ApJ, 168, 45

Done, C., Gierliński, M., & Kubota, A., 2007, A&A Rv, 15, 1

Dubner, G.M., Holdaway, M., Goss, W.M., & Mirabel, I.F., 1998, AJ, 116, 1842

Esin, A.A., McClintock, J.E., & Narayan, R., 1997, ApJ, 489, 865

Falcke, H., Körding, E., & Markoff, S., 2004, A&A, 414, 895

Fender, R.P., 1997, AIPC, 410, 798

Fender, R.P., Corbel, S., Tzioumis, T., McIntyre, V., Campbell-Wilson, D., Nowak, M., Sood, R., Hunstead, R., & 3 coauthors, 1999, ApJ, 519, L165

Fender, R.P., Pooley, G.G., Durouchoux, P., Tilanus, R.P.J., & Brocksopp, C., 2000, MNRAS, 312, 853

Fender, R.P., & Pooley, G.G., 2000, MNRAS, 318, L1

Fender, R.P., 2001, MNRAS, 322, 31

Fender, R.P., & Kuulkers, E., 2001, MNRAS, 324, 923

Fender, R.P., 2003, MNRAS, 340, 1353

Fender, R.P., Gallo, E., & Jonker, P.G., 2003, MNRAS, 343L, 99

Fender, R.P., & Belloni, T., 2004, ARA&A, 42, 317

Fender, R.P., Belloni, T.M., & Gallo, E., 2004, MNRAS, 355, 1105

Fender, R.P., Wu, K., Johnston, H., Tzioumis, A.K., Jonker, P., Spencer, R., & van der Klis, M., 2004, Nature, 427, 222

Fender, R.P., Tzioumis, A.K., & Tudose, V., 2005, ATel, 563

Fender, R.P., 2006, In: *Compact stellar X-ray sources*, Cambridge, UK: Cambridge University Press, 2006, Cambridge Astrophysics Series, No. 39, Edited by Walter Lewin & Michael van der Klis, ISBN 978-0-521-82659-4, arXiv:0303.456

Fender, R.P., Homan, J., & Belloni, T.M., 2009, MNRAS, 396, 1370

Fender, R.P., Gallo, E., & Russell, D.M., 2010, MNRAS, 406, 1425

Fomalont, E.B., & Geldzahler, 1991, ApJ, 383, 289

Fomalont, E.B., Geldzahler, B.J., & Bradshaw, C.F., 2001, ApJ, 553, L27

Frank, J., King, A., & Raine, D., 1992, *Accretion power in astrophysics*, Second edition, Cambridge, UK: Cambridge University Press, 1992, Cambridge Astrophysics Series, No. 21, ISBN 0-521-40863-6

Frater, R.H., Brooks, J.W., & Whiteoak, J.B., 1992, Journal of Electrical and Electronics Engineering Australia, 12, 103

Fruscione, A., McDowell, J.C., Allen, G.E., Brickhouse, N.S., Burke, D.J., Davis, J.E., Durham, N., Elvis, M., & 13 coauthors, 2006, SPIE, 6270, E60

Gallo, E., Fender, R.P., & Pooley, G.G., 2003, MNRAS, 344, 60

Gallo, E., Fender, R.P., Kaiser, C., Russell, D., Morganti, R., Oosterloo, T., & Heinz, S., 2005, Nature, 436, 819

Gallo, E., Fender, R.P., & Hynes, R.I., 2005, MNRAS, 356, 1017

Gallo, E., Fender, R.P., Miller-Jones, J.C.A., Merloni, A., Jonker, P.G., Heinz, S., Maccarone, T.J., & van der Klis, M., 2006, MNRAS, 370, 1351

Gallo, E., Miller, B.P, & Fender, R.P., 2012, MNRAS, temporary version, 2849

Garcia, M.R., McClintock, J.E., Narayan, R., Callanan, P., Barret, D., & Murray, S.S., 2001, ApJ, 264, L49

Geldzahler, B.J., Fomalont, E.B., & Bradshaw, C.F., 1999, AAS, 194, 5007

Giaconni, R., Gursky, H., Paolini, F.R., & Rossi, B.B., 1962, Phys.Rev.Letters, 9, 493

Gierliński, M., & Done, C., 2002, MNRAS, 337, 1373

Glass, I.S., 1978, MNRAS, 183, 335

Goodall, P.T., Fathallah, A., & Blundell, K.M., 2011, MNRAS, 529

Hameury, J.-M., Barret, D., Lasota, J.-P., McClintock, J.E., Menou, K., Motch, C., Olive, J.-F., & Webb, N., 2003, A&A, 399, 631

Hannikainen, D.C., Hunstead, R.W., Campbell-Wilson, D., & Sood, R.K., 1998, A&A, 337, 460

Hannikainen, D., Campbell-Wilson, D., Hunstead, R., McIntyre, V., Lovell, J., Reynolds, J., Tzioumis, A.K., & Wu, K., 2001, Ap&SSS, 276, 45

Hargrave, P.J., & Ryle, M., 1974, MNRAS, 166, 305

Harmon, B.A., Wilson, C.A., Zhang, S.N., Paciesas, W.S., Fishman, G.J., Hjellming, R.M., Rupen, M.P., Scott, D.M., & 2 coauthors, 1995, Nature, 374, 703

Hasinger, G., 1988, Advances in Space Research, 8, 377

Hasinger, G., & van der Klis, M., 1989, A&A, 225, 79

Haynes, R.F., Jauncey, D.L., Murdin, P.G., Goss, W.M., Longmore, A.J., Simons, L.W.J., Milne, D.K., & Skellern, D.J., 1978, MNRAS, 185, 661

Heinz, S., Schulz, N.S., Brandt, W.N., & Galloway, D.K., 2007, ApJ, 663, L93

Hewish, A., & Okoye, S.E., 1965, Nature, 207, 59

Hewish, A., Bell, S.J., Pilkington, J.D.H., Scott, P.F., & Collins, R.A., 1968, Nature, 217, 709

Hjellming, R.M., & Wade, C.M., 1971, ApJ, 164, L1

Hjellming, R.M., & Rupen, M.P., 1995, Nature, 375, 464

Hjellming, R.M., Rupen, M.P., Mioduszewski, A.J., Smith, D.A., Harmon, B.A., Waltman, E.B., Ghigo, F.D., & Pooley, G.G., 1998, AAS, 193, 103.08

Högbom, J.A., 1974, A&AS, 15, 417

Homan, J., van der Klis, M., Wijnands, R., Belloni, T., Fender, R.P., Klein-Wolt, M, Casella, P., Mndez, M., & 3 coauthors, 2007, ApJ, 656, 420

Homan, J., van der Klis, M., Fridriksson, J.K., Remillard, R.A., Wijnands, R., Méndez, M., Lin, D., Altamirano, D., & 3 coauthors, 2010, ApJ, 719, 201

Iaria, R., Spanò, M., Di Salvo, T., Robba, N.R., Burderi, L., Fender, R., van der Klis, M., & Frontera, F., 2005, ApJ, 619, 503

Iaria, R., D'Aí, A., Lavagetto, G., Di Salvo, T., Robba, N.R., & Burderi, L., 2008, ApJ, 673, 1033

Ichimaru, S., 1977, ApJ, 214, 840

Jamil, O., Fender, R.P., & Kaiser, C.R.2010, MNRAS, 401, 394

Johnston, H.M., Fender, R., & Wu, K., 1999, MNRAS, 308, 415

Jonker, P.G., & Nelemans, G., 2004, MNRAS, 354, 355

Jonker, P.G., Nelemans, G., & Bassa, C.G., 2007, MNRAS, 374, 999

Kaiser, C.R., Sunyaev, R., & Spruit, H.C., 2000, A&A, 356, 975

Klein-Wolt, M., Fender, R.P., Pooley, G.G., Belloni, T., Migliari, S., Morgan, E.H., & van der Klis, M., 2002, MNRAS, 331, 745

Klein-Wolt, M., & van der Klis, M., 2008, ApJ, 675, 1407

Kong A.K.H., McClintock J.E., Garcia M.R., Murray S.S., Barret D., 2002, ApJ, 570, 277

Konopelko, A., Mastichiadis, A., Kirk, J., de Jager, O.C., & Stecker, F.W., 2003, ApJ, 597, 851

Körding, E.G., Jester, S., & Fender, R.P, 2006, MNRAS, 372, 1366

Körding, E.G., Fender, R.P., & Migliari, S., 2006, MNRAS, 369, 1451

Lai, D., 1999, ApJ, 524, 1030

Layton, J.T., Blondin, J.M., Owen, M.P., & Stevens, I.R., 1998, New Astronomy, 3, 111

Lewin, W.H.G., van Paradijs, J., & Taam, R.E., 1993, Space Science Reviews, 62, 223

Lewis, F., Russell, D.M., Fender, R.P., Roche, P., & Clark, J.S., 2008, arXiv, 0811, 2336

Linares, M., Watts, A., Altamirano, D., Soleri, P., Degenaar, N., Yang, Y., Wijnands, R., Casella, P., & 9 coauthors, 2010, ApJ, 719, L84

Livio, M., & Pringle, J.E., 1997, ApJ, 486, 835

Maccarone, T.J., 2002, MNRAS, 336, 1371

Margon, B., Lampton, M., Bowyer, S., & Cruddace, R., 1971, ApJ, 169, L23

Margon, B., 1984, ARA&A, 22, 507

Markoff, S., Falcke, H., & Fender R., 2001, A&A, 372, L25

Markoff, S., 2010, Lecture Notes in Physics, 794, 143

Matsuoka, M., Kawasaki, K., Ueno, S., Tomida, H., Kohama, M., Suzuki, M., Adachi, Y., Ishikawa, M., & 24 coauthors, 2009, PASJ, 61, 999

McClintock, J.E., Narayan R., & Rybicki, G.B., 2004, ApJ, 615, 402

McClintock, J.E., & Remillard, R.A., 2006, In: *Compact stellar X-ray sources*, Cambridge, UK: Cambridge University Press, 2006, Cambridge Astrophysics Series, No. 39, Edited by Walter Lewin & Michael van der Klis, ISBN 978-0-521-82659-4, arXiv:0303.456

McHardy, I., 2010, LNP, 794, 203

McNamara, B.R., & Nulsen, P.E.J., 2007, ARA&A, 45, 117

Meier, D.L., 2001, ApJ, 548, L9

Meier, D.L., Koide, S., & Uchida, Y., 2001, Science, 291, 84

Merloni, A., Heinz, S., & di Matteo, T., 2003, MNRAS, 345, 1057

Middelberg, E., Sault, R.J., & Kesteven, M.J., 2006, PASA, 23, 147

Migliari, S., Fender, R.P., Rupen, M., Jonker, P.G., Klein-Wolt, M., Hjellming, R.M., & van der Klis, M., 2003, MNRAS, 342, L67

Migliari, S., & Fender, R.P., 2006, MNRAS, 366, 79

Migliari, S., Tomsick, J.A., Miller-Jones, J.C.A., Heinz, S., Hynes, R.I., Fender, R.P., Gallo, E., Jonker, P.G., & Maccarone, T.J., 2010, ApJ, 710, 117

Mignani, R.P., De Luca, A., Caraveo, P.A., & Mirabel, I.F., 2002, A&A, 386, 487

Miller-Jones, J.C.A., Moin, A., Tingay, S.J., Reynolds, C., Phillips, C.J., Tzioumis, A.K., Fender, R.P., McCallum, J.N., & 2 coauthors, 2012, MNRAS, 419, L49

Mirabel, I.F., & Rodrguez, L.F., 1994, Nature, 371, 46

Mirabel, I.F., Bandyopadhyay, R., Charles, P.A., Shahbaz, T., & Rodriguez, L.F., 1997, ApJ, 477, L45

Mirabel, I.F., Dhawan, V., Chaty, S., Rodriguez, L.F., Marti, J., Robinson, C.R., Swank, J., & Geballe, T., 1998, A&A, 330, L9

Moin, A., Reynolds, C., Miller-Jones, J.C.A., Tingay, S.J., Phillips, C.J., Tzioumis, A.K., Nicolson, G.D., & Fender, R.P., 2011, MNRAS, 693

Murdin, P., Jauncey, D.L., Haynes, R.F., Lerche, I., Nicolson, G.D., Holt, S.S., & Kaluziensku, L.J., 1980, A&A, 87, 292

Nakajima, M., Matsuoka, M., Kawasaki, K., Ueno, S., Tomida, H., Kohama, M., Ishikawa, M., Mihara, T., & 24 coauthors, 2010, ATel, 2608

Nakamura, M., & Meier, D.L., 2004, ApJ, 617, 123

Narayan, R., & Yi, I., 1994, ApJ, 428, L13

Narayan, R., McClintock, J.E., & Yi, I., 1996, ApJ, 457, 821

Neilsen, J., Remillard, R.A., & Lee, J.C., 2011, ApJ, 737, 69

Nicolson, G.D., Glass, I.S., & Feast, M.W., 1980, MNRAS, 191, 293

Nicolson, G.D., 2007, ATel, 985, 1

Norman, M.L., Winkler, K.-H.A., Smarr, L., & Smith, M.D., 1982, A&A, 113, 285

Oke, J.B., 1963, Nature, 197, 1040

Oosterbroek, T., van der Klis, M., Kuulkers, E., van Paradijs, J., Lewin, W.H.G., 1995, A&A, 297, 141

Orosz, J.A., Groot, P.J., van der Klis, M., McClintock, J.E., Garcia, M.R., Zhao, P., Jain, R.K., Bailyn, C.D., & Remillard, R.A., 2002, ApJ, 568, 845

Papaloizou, J.C.B., & Pringle, J.E., 1985, MNRAS, 213, 799

Parkinson, P.M.S., Tournear, D.M., Bloom, E.D., Focke, W.B., Reilly, K.T., Wood, K.S., Ray, P.S., Wolff, M.T., Scargle, J.D., 2003, ApJ, 595, 333

Penzias, A.A., & Wilson, R.W., 1965, ApJ, 142, 419

Pszota, G., Zhang, H., Yuan, F., & Cui, W., 2008, MNRAS, 389, 423

Rau A., Greiner J., Filgas R., 2011, ATel, 3140, 1

Rees, M.J., Begelman, M.C., Blandford, R.D., & Phinney, E.S., 1982, Nature, 295, 17

Remillard, R., Morgan, E., McClintock, J., & Sobczak, G., 1998, IAUC, 7019, 1

Remillard, R.A., & McClintock, J.E., 2006, AAS, 209, 0705

Remillard, R.A., & McClintock, J.E., 2006, ARA&A, 44, 49

Russell, D.M., Fender, R.P., Hynes, R.I., Brocksopp, C., Homan, J., Jonker, P.G., & Buxton, M.M., 2006, MNRAS, 371, 1334

Ryle, M., & Hewish, A., 1960, MNRAS, 120, 220

Sánchez-Fernández, C., Castro-Tirado, A.J., Duerbeck, H.W., Mantegazza, L., Beckmann, V., Burwitz, V., Vanzi, L., Bianchini, A., & 6 coauthors, 1999, A & A, 348, L9

Sault, R.J., & Wieringa, M.H., 1994, A&AS, 108, 585

Sault, R.J., Teuben, P.J., & Wright, M.C.H., 1995, ASPC, 77, 433

Schandl, S., & Meyer, F., 1994, A&A, 289, 149

Schmidt, M., 1963, Nature, 197, 1040

Schwarz, U.J., 1978, A&A, 65, 345

Sell, P.H., Heinz, S., Calvelo, D.E., Tudose, V., Soleri, P., Fender, R.P., Jonker, P.G., Schulz, N.S., & 5 coauthors, 2010, ApJ, 719, L194

Shakura, N.I., & Sunyaev, R.A., 1973, A&A, 24, 337

Shirey, R.E., Bradt, H.V., & Levine, A.M., 1999, ApJ, 517, 472

Shirey, R.E., Levine, A.M., & Bradt, H.V., 1999, ApJ, 524, 1048

Smith, D.A., 1998, IAUC, 7008, 1

Sobczak, G.J., McClintock, J.E., Remillard, R.A., Cui, W., Levine, A.M., Morgan, E.H., Orosz, J.A., & Bailyn, C.D., 2000, ApJ, 544, 993

Soleri, P., Tudose, V., Fender, R.P., van der Klis, M., & Jonker, P.G., 2009a, MNRAS, 399, 453

Soleri, P., Heinz, S., Fender, R.P., Wijnands, R., Tudose, V., Altamirano, D., Jonker, P.G., van der Klis, M., & 3 coauthors, 2009b, MNRAS, 397, L1

Soleri, P., & Fender, R., 2011, MNRAS, 413, 2269

Stewart, R.T., Caswell, J.L., Haynes, R.F., & Nelson, G.J., 1993, MNRAS, 261, 593

Stirling, C.H., Esat, T.M., Lambeck, K., McCulloch, M.T., Blake, S.G., Lee, D.-C., & Halliday, A.N., 2001, Science, 291, 290

Swarup, G., Ananthakrishnan, S., Kapahi, V.K., Rao, A.,P., Subrahmanya, C.R., & Kulkarni, V.K., 1991, Current Science, 60, 95

Tarasov, A.E., Brocksopp, C., & Lyuty, V.M., 2003, A&A, 402, 237

Taylor, G.B., Carilli, C.L., & Perley, R.A., 1999, ASPC, 180

Taylor, A.R., 2008, Proceedings of the International Astronomical Union, 248, 164

Tennant, A.F., Fabian, A.C., & Shafer, R.A., 1986, MNRAS, 219, 871

Terzian, Y., & Parrish, A., 1970, ApL, 5, 261

Thomas, R.M., Duldig, M.L., Haynes, R.F., & Murdin, P., 1978, MNRAS, 185, P29

Thompson, A.R., Clark, B.G., Wade, C.M., & Napier, P.J., 1980, ApJS, 44, 151

Tigelaar, S.P., Fender, R.P., Tilanus, R.P.J., Gallo, E., & Pooley, G.G., 2004, MNRAS, 352, 1015

Tingay, S.J., Jauncey, D.L., Preston, R.A., Reynolds, J.E., Meier, D.L., Murphy, D.W., Tzioumis, A.K., McKay, D.J., & 12 coauthors, 1995, Nature, 374, 141

Tomsick, J.A., Corbel, S., Fender, R.P., Miller, J.M., Orosz, J.A., Rupen, M.P., Tzioumis, A.K., Wijnands, R., & Kaaret, P., 2003, ApJ, 597L, 133

Toor, A., 1977, ApJ, 215, L57

Tudose, V., Fender, R.P., Kaiser, C.R., Tzioumis, A.K., van der Klis, M., & Spencer, R.E., 2006, MNRAS, 372, 417

Tudose, V., Fender, R.P., Kaiser, C.R., Tzioumis, A.K., Spencer, R.E., & van der Klis, M., 2008, MNRAS, 390, 447

Tudose, V., Fender, R.P., Linares, M., Maitra, D., & van der Klis, M., 2009, MNRAS, 400, 2111

van der Laan, H., 1966, Nature, 211, 1131

van Paradijs, J., 1996, ApJ, 464, L139

de Vos, M., Gunst, A.W., & Nijboer, R., 2009, Proceedings of the IEEE, 97, 1431

Whelan, J.A.J., Mayo, S.K., Wickramasinghe, D.T., Murdin, P.G., Peterson, B.A., Hawarden, T.G., Longmore, A.J., Haynes, R.F., & 5 coauthors, 1977, MNRAS, 181, 259

Wilson, W.E., Ferris, R.H., Axtens, P., Brown, A., Davis, E., Hampson, G., Leach, M., Roberts, P., and 25 coauthors, 2011, MNRAS, 416, 832

Zhang, S.N., Wilson, C.A., Harmon, B.A., Fishman, G.J., Wilson, R.B., Paciesas, W.S., Scott, M., & Rubin, B.C., 1994, IAUC, 6046, 1