ENABLING THE FULL CAPACITY OF MEERKAT FOR STUDYING X-RAY BINARIES

by

ANDREW KENNETH HUGHES

A thesis submitted in partial fulfilment of the requirements for the degree of

Doctor of Philosophy

Department of Physics
University of Alberta

Abstract

Accretion is ubiquitous in the Universe, observed in various systems spanning many orders of magnitude in their length and mass scales. The X-ray-bright inward-moving accretion flow is often accompanied by highly-relativistic radio-bright outflows that drive material away from the accretor, facilitating interactions with distant environments and the subsequent deposition of the outflowing energy and matter. This process, known as feedback, is critical in regulating the local evolution in these interaction regions. Thus, a comprehensive understanding of accretion flows and accretion-powered jets is crucial for building a coherent picture of our Universe. Despite the ubiquity of accretion and jets, we are still uncertain of the energetics of jets and the mechanisms that form and launch jets, warranting further high-quality monitoring. In this work, I present recent multi-wavelength monitoring campaigns of three outbursting X-ray binaries — accreting stellar-mass black holes and neutron stars. X-ray binaries function as naturally occurring time domain laboratories with which we can study accretion-jet coupling. These observations, taken as a part of the ThunderKAT and X-KAT collaborations, utilize quasi-simultaneous radio (taken with the radio interferometer MeerKAT) and X-ray observations (from various telescopes) to probe the

time-evolution of outflows and inflows. In addition to the targeted observations, I discuss source-agnostic tools I developed to maximize the astrometric and polarimetric information we can derive from MeerKAT observations. My first target, the neutron star X-ray binary SAX J1810.8−2609, underwent a bright outburst in 2021. ThunderKAT monitoring revealed that SAX J1810.8−2609 is a member of the sub-population of X-ray binaries that fail to undergo the standard outburst evolution through X-ray accretion states. Moreover, I detected persistent radio emission in the absence of X-ray emission, which I found was likely due to contamination from an unassociated foreground/background source. This thesis component revealed the need for high-precision astrometry, as the variable source position was the most substantial evidence for the unrelated background source, preventing over-interpretation of the ThunderKAT data. Thus, I developed (and later upgraded) a routine to quantify the astrometric precision in MeerKAT observations (ASTKAT); I found that the current *ad hoc* approach to astrometry significantly overestimates errors. The techniques in ASTKAT are not MeerKAT specific and can be applied to other radio interferometers. High-precision astrometry is critical for the accurate modelling of jet propagation and interactions as this type of kinematic modelling that provides the best estimate for jet energetics and, thus, jet-based feedback as a whole. My second target, 1A 1744−361, another neutron star X-ray binary, showed evidence for ballistic outflows during the canonical state transition. While commonly observed in black hole X-ray binaries, my work shows that 1A 1744-361 provides the best evidence that these ejections also occur in standard neutron star X-ray binary systems, implying that both neutron stars and black holes have properties or processes conducive to jet ejections. While this thought was widely adopted, until now it has lacked empirical evidence. My third and final target, Swift J1727.8−1613, is a black hole X-ray binary that, in late 2023, underwent one of the brightest outbursts in the past decade. At this point in my research, I improved the ThunderKAT / X-KAT calibration and analysis pipeline (POLKAT) to handle polarisation observations. Swift J1727.8−1613 demonstrated exotic signatures in its photometric, polarimetric, and spatial properties, tied to repeated ejections of ballistic outflows. The utility of polarisation became immediately apparent as I detected, for the first time, evidence for a transient increase in the Faraday rotation measure during a ballistic ejection event, favouring an electron-proton (rather than an electron-positron) internal composition in X-ray binary ejecta. Since protons are much more massive than electrons, electron-proton jets have significantly more kinetic energy for feedback (at a fixed bulk velocity). Moreover, POLKAT can be readily adapted to any synthesis imaging polarization observations that aim to measure polarization properties using short tracks on interferometric arrays with linear polarization feeds. In addition to the MeerKAT array, the Advanced Atacama Large Millimeter Array, Australian Square Kilometre

Array Pathfinder, and the Square Kilometer Array Mid are all examples of relatively new or soon-to-be-built interferometric arrays with linearly polarized feeds.

Preface

This thesis is an original work by Andrew K. Hughes. The MeerKAT observations were taken as a part of The Hunt for Dynamic and Explosive Radio transients with meerKAT (ThunderKAT; MeerKAT observations taken 2018–2023) collaboration and its followup X-KAT collaboration (MeerKAT observations taken 2023–). Both projects are led by Rob Fender, Patrick Woudt, and James Miller-Jones. The observing plans of all Thunder-KAT and X-KAT observations of X-ray binaries were created weekly by either Joe Bright, Lilia Tremou, Francesco Carotenuto, Payaswini Saikia, or myself. For each plan, an on-site team led by Sarah Buchner then scheduled the observations with MeerKAT. ThunderKAT and X-KAT have a sister program, *Swift*KAT, that schedules quasi-simultaneous X-ray observations using the X-Ray Telescope aboard The Neil Gehrels Swift Observatory. The *Swift*KAT observations are planned by Sara Motta and Gregory Sivakoff. Individual early-career scientists held responsibility for leading outbursts from different sources. I led observations of **SAX J1810−2609**, GRS 1747−312, **1A 1744−361**, MAXI J1807+132, and **Swift J1727−1613** although I only discuss the three most scientifically revealing outbursts in this thesis (the 'bolded' sources). The remaining sources will be published as sub-components of multi-wavelength analyses or as a part of a future ThunderKAT summary paper.

Chapter 3 was published in the Monthly Notices of the Royal Astronomical Society as "Andrew K. Hughes, et al., SAX J1810.8-2609: an outbursting neutron star X-ray binary with persistent spatially coincident radio emission, Monthly Notices of the Royal Astronomical Society, Volume 527, Issue 3, January 2024, Pages 9359-9377". For this thesis, the introduction component that discusses the properties of X-ray binaries was removed, expanded, and replaced with Chapter 1. The Appendices of the original paper were renamed and appear in Section 3.6.

Chapter 4 was part of a multi-collaboration effort between ThunderKAT and a team based out of the Massachusetts Institute of Technology (MIT) that uses The Neutron Star Interior Composition ExploreR (NICER) to study X-ray binaries. The NICER component, which focused on the X-ray timing properties, was led by (now-completed) MIT doctoral student Mason Ng; I led the ThunderKAT component, focusing on the signatures of ballistic outflows. Our results were published in The Astrophysical Journal as "Mason Ng, Andrew K. Hughes, et al., X-Ray and Radio Monitoring of the Neutron Star Low-mass X-Ray Binary 1A 1744-361: Quasiperiodic Oscillations, Transient Ejections, and a Disk Atmosphere, The Astrophysical Journal, Volume 966, Issue 2, May 2024, Page 232". Like Chapter 3, I replaced the broader introduction to X-ray binaries with text in Chapter 1. To accurately reflect my contributions to the joint paper as published, Chapter 4 only includes the research components I led. While I reference some critical results from the X-ray timing analysis, they have been largely omitted from this thesis. I include a more detailed spectral analysis of a single X-ray *Swift*KAT observation, but note that it does not change any of the interpretations from the published version of this work.

Chapter 5 is neither published nor (currently) submitted for publication, as monitoring of the target, Swift J1727.8−1613, is ongoing. I include observations up to 2024 July 1 and discuss the best current interpretations of the source properties. Moreover, I detail the new calibration and imaging pipeline designed to handle the full polarisation observations of X-KAT, i.e., POLKAT. Although this pipeline was made for X-KAT, and thus X-ray binary monitoring, it can handle any full polarisation observations with

MeerKAT.

In Chapter 6, I present an improvement on the astrometric routine introduced initially in Appendix A of "SAX J1810.8-2609: an outbursting neutron star X-ray binary with persistent spatially coincident radio emission" (Section 3.6.1 in this thesis) and a more comprehensive description of astrometric errors with MeerKAT (and other interferometers). This routine, ASKAT, is both source- and interferometer-agnostic and thus can be used for any time-domain data taken with a wide-field interferometer.

"I'm about to go TURBO 455!"

—Schteve-an Fall-man

Acknowledgements

First and foremost, I would like to thank my supervisor, Gregory Sivakoff. When I first began grad school, I had a lot of interest in astronomy but no experience. Greg, immediately recognising my masochism, directed me towards polarimetry with radio interferometers. As a result, I have developed a highly employable skill set, and fortunately, I now enjoy what I am good at. Next, I thank the supervisory committee and external examiners for their constructive feedback. I would also like to thank all members of the U of A astronomy group; our regular meetings and discussions allowed me to gain insight into the broader scope of astronomical research. I would also like to acknowledge every member of every collaboration I participate in. Unfortunately, if I began listing names, this thesis would double in length, so, as with papers, they will have to get relegated to an 'et al.' I have to give a shout-out to the first friends I made in Alberta: Steven Fahlman (aka Schteve), Coleman Dean (aka ColeMAESTROeX), and Mario Ivanov (aka... Mario — but pronounced funny). Our discussions about whether you eat or drink soup were *definitely not* wastes of time.

Among the non-astronomers, I would like to first thank my lovely past girlfriend Kayla Carson — she is only my *past* girlfriend because she is my *current* fiance and *future* wife. Next, I would like to thank my mother, Nicki Hughes; father, Ken Hughes; step-mother, Nicole Brebner; step-father, Mark Samis; and *maybe* my little brother, James Hughes. Lastly, I thank my long-time friends Tyler Rogerson and Piraaveenan Nirmalen; I hope we stay in touch as I am not returning to Toronto. In the case of a failed defence, replace all the instances of 'thank' with 'blame' in this second paragraph.

Contents

List of Tables

List of Figures

Acronyms

AGN Active Galactic Nuclei

ATCA The Australia Telescope Compact Array

BAT Burst Alert Telescope

BH Black Hole

BHXB Black Hole X-ray Binary

EVPA Electric Vector Position Angle

FBO Flaring Branch Oscillation

FDF Faraday Dispersion Function

FWHM Full-Width at Half-Maximum

GRB Gamma-ray Burst

HBO Horizontal Branch Oscillation

HID Hardness Intensity Diagram

HMXB High Mass X-ray Binary

HR Hardness Ratio

IMXB Intermediate Mass X-ray Binary

LBA The Long Baseline Array

LMXB Low Mass X-ray Binary

MAXI/GSC The Monitor of All-sky X-ray Image Gas Slit Camera

MSP Millisecond Pulsar

NBO Normal Branch Oscillation

NICER The Neutron Star Interior Composition Explorer Mission

NS Neutron Star

NSXB Neutron Star X-ray Binary

QPO Quasi-Periodic Oscillation

RM Rotation Measure

RMTF Rotation Measure Transfer Function

S/N Signal-to-Noise Ratio

SKA Square Kilometer Array

SMBH Supermassive Black Hole

tMSP Transitional Millisecond Pulsar

VLA Very Large Array

VLASS Very Large Array Sky Survey

VLBA Very Long Baseline Array

WD White Dwarf

XRB X-ray Binary

XRT X-ray Telescope

Chapter 1

X-ray Binaries: Laboratories for Accretion and Jet Physicss

Accretion, the gravitational-driven accumulation of material, is ubiquitous in the Universe. The 'accretor' is usually small — relative to the size of the accretion flow — and gravitationally bound. Some examples include: young stars; white dwarfs (remnants of *dead* low-mass stars); neutron stars (NSs; remnants of *dead* high-mass stars); and black holes (BHs, both stellar-mass BHs and the super-massive BHs found in the centre of galaxies (Livio, 2002)). During accretion, the inflow of matter powers simultaneous bipolar outflows in the form of winds and astrophysical jets, the latter being the primary topic of this thesis. These jets are highly collimated and, when powered by compact objects (a shorthand for BHs and NSs), highly relativistic—such jets also dwarf in size both the accretion flow and the accretor. As the outflowing material in jets is known to interact with the surrounding environment, depositing matter and significant amounts of energy, it regulates the local evolution through a process known as (jet-based) feedback (e.g., Hardcastle & Croston, 2020). Therefore, understanding accretion and accretion-driven outflows is critical for understanding the Universe.

The archetypal (continuously) accreting systems with compact object accretors are X-ray binaries (XRBs) — binary systems composed of stellar-mass BH or NSs accreting material from a non-compact companion star (henceforth the *donor* star) — and active galactic nuclei (AGN) — accreting supermassive BHs. Additionally, short-lived cataclysmic events, namely gamma-ray bursts (GRBs), also produce accretion discs and drive relativistic jets, both short (NS-mergers) and long (exploding, rapidly-rotating massive stars) GRBs (Ruiz et al., 2021). Earth-based laboratories cannot reproduce the extreme conditions associated with these systems, so we use XRBs, AGN, and GRBs as natural laboratories to study accretion-jet coupling. Due to their size, AGN jets arguably have the largest impact on the Universe. However, a by-product of their large sizes are longer evolutionary timescales; AGN have *complete* evolutionary timescales (i.e., the timescales for a jet to turn "on", evolve and turn "off") of $> 10^3$ years (e.g., Czerny et al., 2009), fundamentally limiting the time-domain processes that can be studied with individual AGN. GRBs are stellar-sized objects and, thus, are rapidly evolving. However, they occupy the other extreme, as the timescale for active accretion is minutes (Woosley & Bloom, 2006; Gottlieb et al., 2023), and the systems destroy themselves afterwards, preventing repeat events. XRBs occupy a 'sweet spot' and are often considered a canonical time-domain laboratory for accretion physics. These stellar systems undergo (repeated) outbursts that last weeks to years (Tetarenko et al., 2016a; Corral-Santana et al., 2016). Moreover, XRBs are unique as they contain both black holes and neutron stars, allowing us to explore the difference between the two accretors, answering a greater range of questions regarding accretion and jet launching (e.g., what effect does an event horizon have when compared to a solid surface)?

It is widely theorized that accreting sources that drive relativistic jets have similar geometries and, to some degree, size-scale invariance. Therefore, inferences derived from one of a class of accreting systems sources should provide insight into *all* systems. Despite the ubiquity of accretion-jet coupling, many crucial questions remain unanswered; for example, what is the total energy content and composition of astrophysical jets, and, as a result, what is the total energy available for jet-based feedback? What are the jet formation, launching, and disruption mechanisms, and how do they differ between compact jets and jet ejecta? And how are these properties connected to the accretion flow? Progress towards answering these questions requires high-quality multi-frequency observations capable of simultaneously measuring the evolution of the intensity, polarisation, spatial, and spectral properties for both the accretion flow and jets. This thesis presents multi-wavelength observations for three outbursting X-ray binaries taken as a part of The Hunt for Dynamic and Explosive Radio transients with MeerKAT (ThunderKAT; Fender et al., 2016, MeerKAT observations taken 2018–2023) collaboration and its follow-up X-KAT collaboration (MeerKAT observations taken 2023–). The remainder of this Chapter introduces X-ray binaries.

1.1 Classifying X-ray Binaries

Like many astronomical systems, XRBs are classified into distinct subpopulations. The primary classifier is the compact object type: black hole (BHXBs) and neutron star (NSXBs) X-ray binaries. The next classification uses the mass of the donor star: low- (LMXB; $M_d \lesssim 1.5 M_\odot$; $M_d \lesssim M_c$), intermediate- (IMXB; $1.5 M_\odot \lesssim M_d \lesssim 8 M_\odot$), and high-mass (HMXB; $M_d \gtrsim 8 M_\odot$; $M_d > M_c$) systems[1]. These "mass classes" typically have distinct accretion mechanisms, where LMXBs (and likely, IMXBs) are fed via Roche-lobe overflow. In contrast, HMXBs have massive donors (O, A, or B spectral class) and tend to transfer mass through strong stellar winds[2]. While the XRB subpopulations

[1]Here M_d, M_c, and $M_\odot$ correspond to the masses of the donor, compact object, and our Sun, respectively

[2]There is another type of NS HMXB, the Be/X-ray binary, where the binary is on a highly elliptical orbit, and the Be-type donor is rapidly rotating resulting in a decretion

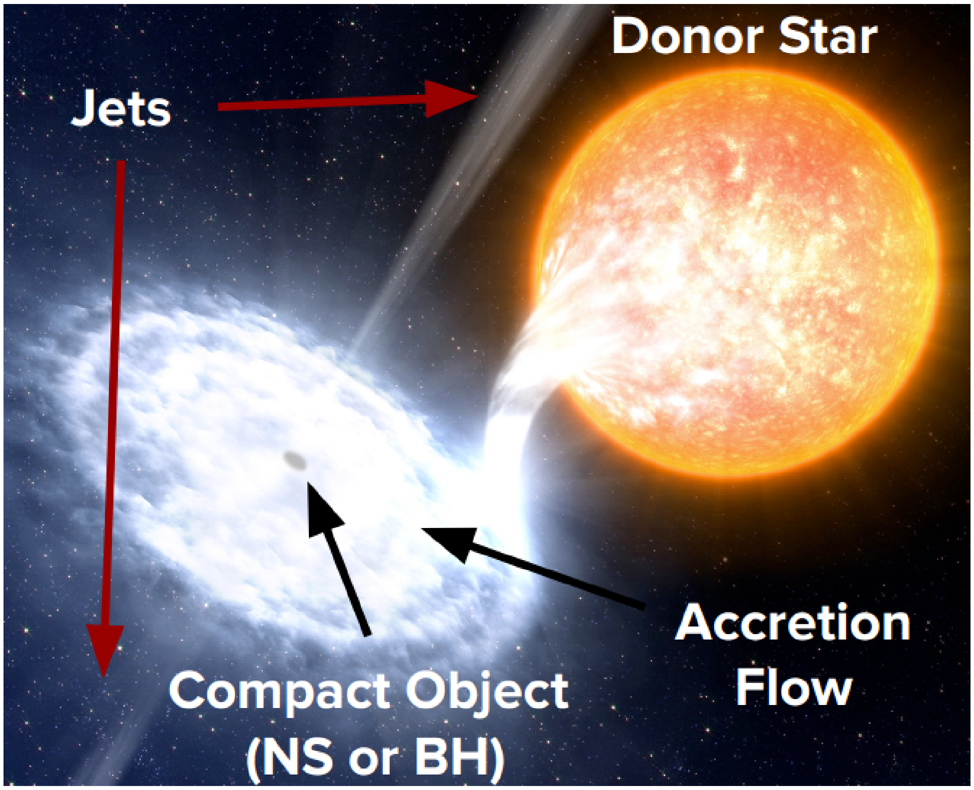

Figure 1.1: An artist's rendition of the LMXB geometry. The strong gravitational pull of the compact object *steals* material from a non-degenerate companion/donor star. Due to the conservation of angular momentum, the inflowing material forms an accretion disc, powering simultaneous jetted outflows perpendicular to the plane of the disc. Permission to include the figure in this thesis was given by the original creator, ESO/L. Calçada.

exhibit a broad range of phenomenology, each is fascinating in its own right. Since focusing on the total population of XRBs is too wide of a scope for a single thesis, this work, and the remainder of this chapter, concentrates on LMXBs (both BHs and NSs, see Fig. 1.1 for an artist's impression of the geometry).

1.2 Accretion Flows and X-ray Outbursts

Most LMXBs are transient sources, spending the majority of their lifetimes in quiescence, accreting small amounts of matter at low X-ray luminosities ($L_X \lesssim 10^{-5} L_{\rm Edd}$) before sporadically entering into bright outbursts that

disc. Mass transfer occurs when the neutron star passes through the decretion disc.

reach an appreciable fraction of the Eddington Luminosity, $L_{\rm Edd}$[3]. During the week-to-year-long outbursts (McClintock & Remillard, 2006; Tetarenko et al., 2016a), LMXBs occupy several "accretion states" categorized by the properties of the accretion flow (best observed at X-ray frequencies, e.g., Belloni et al., 1999; Plant et al., 2014; Ingram et al., 2023) and, to a lesser extent, the astrophysical jet(s) (best observed at **radio**-to-infrared frequencies, e.g., Corbel & Fender, 2002; Russell et al., 2013; Tetarenko et al., 2017). The accretion state definitions were originally motivated by the evolving X-ray properties of the canonical BHXB Cygnus X-1[4] (e.g., Zhang et al., 1997; Gies et al., 2003). As a result, it is useful to introduce accretion state nomenclature in the context of outbursting BHXBs. Section 1.4 will discuss outbursting NSXBs.

The accretion states are empirically defined based on the observed X-ray spectral and timing properties (see, e.g., Homan & Belloni, 2005; van der Klis, 2006; McClintock & Remillard, 2006; Belloni, 2010; Belloni & Motta, 2016, for comprehensive reviews of accretion states). For spectral properties, state classification is based on the model describing the X-ray spectrum, e.g., thermal (disc-)blackbody or non-thermal power-laws. Another (and more straightforward) metric is the X-ray hardness ratio (HR), which quantifies the ratio between the amount of high-energy (i.e., hard) and low-energy (i.e., soft) X-ray photons. The timing properties were added later to accretion state characterization. These properties are measured from structures in the X-ray power spectral density[5] (PSD). The fractional root-mean-square (rms[6]) variability quantifies the PSD continuum, measuring the amount

[3] $L_{\rm Edd} \sim 1.3 \times 10^{38} (M/M_\odot)\,{\rm erg\,s^{-1}}$ is the luminosity where radiation pressure becomes stronger than gravity in the limit of a steady, spherical accretion of hydrogen; i.e., the luminosity that would disrupt the accretion flow.

[4] Ironically, Cygnus X$-$1 is an HMXB.

[5] The power spectral density is a Fourier analysis that quantifies the power as a function of frequency for time-domain data; here, frequency refers to the frequency of the Fourier components, not spectral frequency.

[6] I employ a lower case acronym to avoid potential confusion with the abbreviation used throughout this thesis for rotation measures — RMs.

of stochastic variability over the expected (Poisson) noise levels (see, e.g., Belloni et al., 2002, and references therein, for a quantitative description of X-ray timing analyses with XRBs). Furthermore, the different accretion states exhibit periodic timing phenomena called quasi-periodic oscillations (QPOs). QPOs, which appear as Lorentzian peaks in the PSD, are classified as Type-A, -B, and -C based on their amplitudes, widths, and central frequencies (see, e.g., Lewin et al., 1988; Ingram & Motta, 2019, for a comprehensive review of QPO phenomena). I will not go into detail about the physical origins of each QPO type (as these origins are a matter of active debate), but I will note that these are X-ray timing features and, thus, originate from the accretion flow. The observed QPO type depends strongly on the accretion state; the appearance (or disappearance) of QPOs provides some of the most robust evidence for accretion state transitions.

Black hole X-ray binary outbursts are often described as an evolution in the X-ray hardness-intensity diagram (HID). Observations of multiple XRBs revealed a common outburst track in the HID (the q-shaped track of Fig. 1.2), which was then adopted as the *standard* outburst. These standard outbursts begin with a rapid increase in the X-ray luminosity, as the system transitions from quiescence into the *hard* state[7]. In the hard state, the X-ray spectrum is dominated by high-energy, non-thermal photons comptonized by an optically-thin, geometrically thick coronal flow. The X-ray spectrum is traditionally modelled as a power law distribution with a photon index of $\Gamma \sim 1.7$ (where the X-ray flux at a given frequency ν is given by $F_{X,\nu} \propto \nu^{-(\Gamma-1)}$). The X-ray timing properties show a large fractional rms variability $\gtrsim 20\%$ (in the 2-20 keV energy band; van der Klis, 2006) and Type-C QPOs.

As the X-ray luminosity increases, and thus the accretion rate, the system transitions towards the *soft* state after briefly occupying the *hard interme-*

[7]Depending on the literature, the hard state can also be referenced as the 'hard/low' or 'comptonized' state

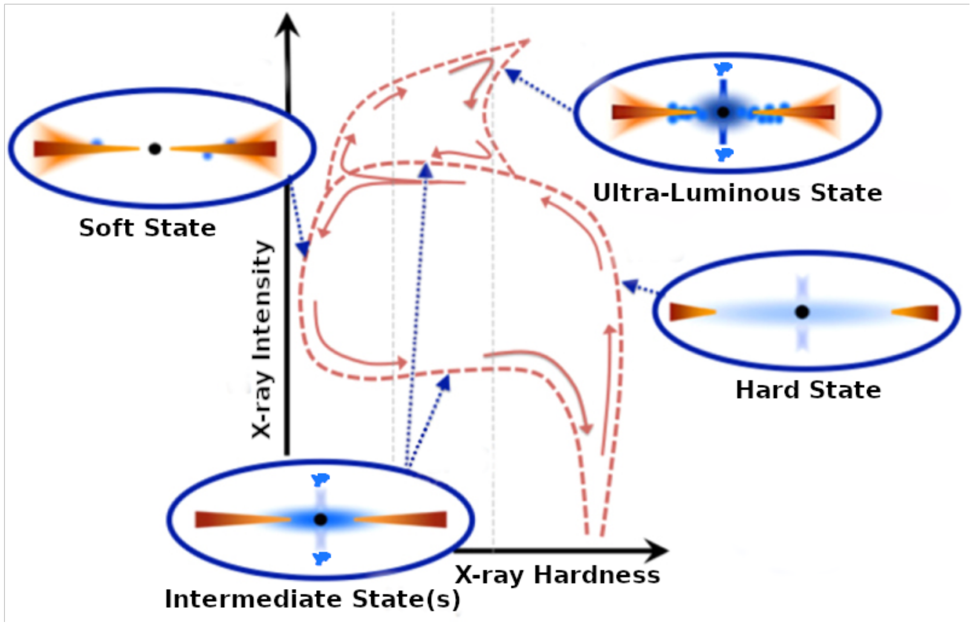

Figure 1.2: Schematic representation of a standard HID track for a BHXB outburst. The inset figures show the geometry of the accretion flow and jets, and the colours of the accretion flow highlight the characteristic X-ray energies (i.e., red and blue correspond to softer and harder X-rays, respectively). The red arrows represent the evolution of forward time starting from the bottom right corner of the diagram (i.e., quiescence). While the standard outburst typically includes just the lower q-shaped evolution, this diagram also shows some outbursts that reach the ultra-luminous state before the decaying intensity portion of the soft state. This figure is a modified version of Figure 1 from Tetarenko et al. (2016a), which is licensed under CC BY 4.0.

diate followed by the *soft intermediate* states (the transitional luminosity is typically ~ 0.01–$0.05 L_{\mathrm{Edd}}$; e.g., Maccarone, 2003). In the intermediate states, the X-ray emission becomes progressively softer, becoming dominated by thermal photons, and the rms variability decreases to $< 3\%$. In the soft intermediate state, a Type-B QPO replaces the Type-C QPO; the existence of a Type-B QPO is the definition of the soft intermediate state (Ingram & Motta, 2019). Following the intermediate states, the system enters the *soft* state, where the X-ray emission is dominated by an optically thick, geometrically thin accretion disc, typically modelled as a multi-coloured disc black body with a characteristic temperature of $k_B T \sim 1\,\mathrm{keV}$ at the inner disc radius (McClintock & Remillard, 2006). The Type-B QPO disappears, and another Type-C QPO may emerge. BHXBs typically remain in the soft state for weeks to months before following a reverse, lower-luminosity track through the intermediate states back to the hard state and, finally, returning to quiescence.

While the abovementioned behaviour is considered the standard outburst, many systems deviate from this picture. Some BHXBs enter an *ultra-luminous* state[8] before or shortly after the soft state. In the ultra-luminous state, the X-ray luminosity approaches, or in rare cases, exceeds the Eddington Luminosity, and the system shows short time-scale hardness fluctuations, analogous to multiple erratic state transitions (e.g., Uttley & Klein-Wolt, 2015; Kajava et al., 2020). The X-ray timing properties have moderate fractional rms variability ($\sim 10\%$) and may exhibit any (or all) of the QPO types. Other less-luminous sources will exhibit transient softenings/hardenings of their X-ray spectra, commonly interpreted as short excursions from the soft state back into the intermediate state (e.g., Carotenuto et al., 2021). Lastly, (Tetarenko et al., 2016a) showed that nearly $\sim 40\%$ of BHXBs never transition to soft state, undergoing hard state only or *failed* outbursts. The frequency of failed outbursts is the most robust evidence against the concept of a 'standard' outburst.

1.3 Relativistic Jets

Throughout an outburst, the properties of the relativistic jet evolve in sequence with the accretion states (see, e.g., Fender et al., 2004, 2009; Fender, 2010, for detailed reviews). In the hard state, the jet is continuously fueled by the coronal accretion flow and is typically described with the highly collimated, conical outflow model of Blandford & Königl (1979). This 'compact jet' is *usually* optically thick with a partially self-absorbed synchrotron spectrum and an inverted (or flat) spectral index ($\alpha \gtrsim 0$; for a

[8] also known as the steep-power law, very high, or ultra-soft state

flux density of $F_{R,\nu} \propto \nu^{\alpha}$) up to a break frequency. At higher frequencies, the spectrum of the compact jet becomes optically thin ($\alpha \sim -0.7$, with a break frequency typically in the sub-mm or near-infrared regimes; Migliari et al., 2010; Russell et al., 2013). According to Blandford & Königl (1979), the partial self-absorption results from a superposition of multiple populations of synchrotron-emitting particles, with properties that vary as a function of distance from the BH (i.e., lower frequency emission originates further from the BH).

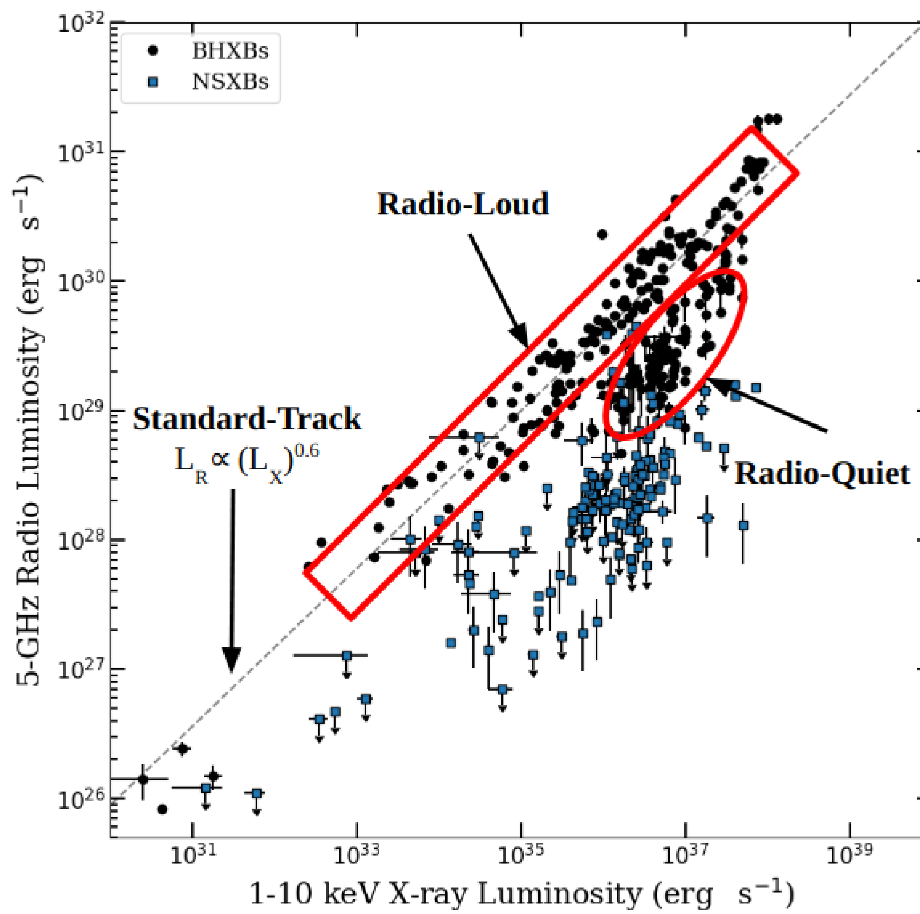

Figure 1.3: L_R–L_X observations for the BHXB (black circles) and NSXB (blue square) populations. The red region identifies the radio-loud and radio-quiet tracks for BHXB, and the dashed line shows the standard track with $L_R \propto L_X^{0.6}$. Data taken from Bahramian & Rushton (2022) and van den Eijnden et al. (2021).

The radio luminosity (L_R) in the hard state shows a positive correlation with the X-ray luminosity (henceforth, the L_R–L_X relation, $L_R \propto L_X^{\beta}$; Gallo et al., 2003; Corbel et al., 2013; Gallo et al., 2018). After including a mass-scaling term, the L_R–L_X relation was extended to include both AGN and BHXBs (Merloni et al., 2003), providing empirical evidence for scale-invariance in accreting systems. However, this 'standard-track' for BHXBs, $L_X \propto L_R^{0.6}$, was predominantly derived from observations of a single well-monitored source (GX 339-4 Hannikainen et al., 1998). While some BHXBs followed the 'standard' relationship (e.g., V404 Cyg and GRO J0422+32, see, Gallo et al., 2003, and reference therein), others exhibit significantly lower radio luminosities at a fixed L_X (e.g., H1743−32 Coriat et al., 2011; Williams et al., 2020). The existence of multiple populations led to the introduction of the concept of 'radio-loud' (i.e., standard-track) and 'radio-quiet' sources (see Fig. 1.3[9]. The physical mechanism responsible for the differences, and whether it is due to variations in the jet or accretion flow (see, Coriat et al., 2011, and reference therein), has yet to be understood — necessitating additional monitoring of both radio-quiet and radio-loud sources. Moreover, there may be no dichotomy; instead, BHXBs may form a continuum from radio-loud to radio-quiet (Gallo et al., 2014, 2018).

In the soft state, the compact jet is quenched (i.e., strongly suppressed), decreasing in luminosity by ≥ 3 orders of magnitude, to the point of non-detection (e.g., Coriat et al., 2011; Russell et al., 2020). During the transition through the intermediate or ultra-luminous states, the source may exhibit one or more bright flaring events, commonly associated with the launching of bipolar jet ejecta, decoupling the jet from the accretion flow (e.g., Han & Hjellming, 1992; Hjellming & Rupen, 1995; Fender et al., 1999a; Rushton

[9]The term 'radio-quiet' does not comment on the physical origin of the different tracks, just that at a fixed X-ray luminosity some sources have lower radio luminosities. The reason for the difference may, as some authors have proposed, be due to enhanced X-ray emission from the accretion flow rather than suppressed radio emission from the jet. Regardless, we adopt the 'radio-quiet'/'radio-loud' dichotomy as it is a standard naming convention.

et al., 2017; Tetarenko et al., 2017; Miller-Jones et al., 2019; Russell et al., 2019b; Bright et al., 2020; Carotenuto et al., 2021). As the ejecta propagate at mild-to-extremely relativistic velocities, they undergo expansion, resulting in temporally evolving spectra (see, e.g., the van der Laan — vdL — model; van der Laan, 1966; Hjellming & Johnston, 1988). In contrast to compact jets, the spectra of jet ejecta are consistent with a single population of synchrotron-emitting electrons that become optically thin to progressively lower frequencies during propagation. Jet ejecta can remain at detectable flux densities anywhere from hours to years post-launch (e.g., Miller-Jones et al., 2019; Bahramian et al., 2023). The longest-lasting ejecta are thought to remain bright due to continual particle acceleration driven by interactions (i.e., collisions) with the ambient interstellar medium (ISM; Espinasse et al., 2020; Bright et al., 2020). In extreme cases, after fading below the detection threshold, some jet ejecta can remain undetected for months or years before re-brightening during these jet-ISM interactions. Tracking the kinematics of long-lasting ejecta, particularly when deceleration is observed, transforms the system into a calorimeter and provides the best estimates of the energetics of an ejecta (e.g., Carotenuto et al., 2021, 2024).

Historically, LMXB jets have been studied using their photometric, spectral, timing, and, when available, spatial properties, with a sparse sample including polarimetry despite synchrotron-emission being linearly polarised. For a uniform magnetic field, the percent linear polarisation (hereafter linear polarisation fraction) is $m_l = (3p+3)/(3p+7) \times 100\,\% \approx 70\,\%$ and $m_l = 3/(6p+13) \times 100\,\% \approx 10\,\%$, for optically thin and thin synchrotron emission, respectively (assuming a typical value for the electron energy distribution index, $p = 2.2$; Ginzburg & Syrovatskii, 1969; Longair, 2011). In practice, linear polarisation detections of BHXB jets are typically measured at the $\sim$0.1–10% level (e.g., Gallo et al., 2004; Rushton et al., 2010; Brocksopp et al., 2007, 2013; Curran et al., 2015; Hughes et al., 2023), with only a few systems reaching fractions comparable to the theoretical maximum (e.g., the $\sim 50\%$ polarised jets of Swift J1745-26; Curran et al., 2014).

The electric vector position angle (EVPA), or "linear polarisation angle", is expected to be either perpendicular to or parallel to the sky-projected orientation of the magnetic field for the optically thin and thick cases, respectively. In the simplest interpretations, the dominant magnetic field direction will be established by shock compression, or velocity shearing, aligning the magnetic field perpendicularly, or parallel, to the jet direction of propagation (i.e., the jet position angle — PA; Laing, 1980). However, complex magnetic field geometries or overlapping, unresolved components with (partially) orthogonal magnetic fields can significantly reduce the linear polarisation fraction and rotate the EVPA. Therefore, polarimetry is the most direct probe of underlying magnetic field structure, where the angle measures orientation and the polarisation fraction measures coherence. Given that evolving magnetic fields are thought to be present during jet launching, collimation, and collisions (Blandford & Znajek, 1977; Blandford & Payne, 1982; Meier et al., 2001; Brocksopp et al., 2007, 2013; Contopoulos et al., 2012), the inclusion of polarimetry is crucial.

1.4 NS LMXBs Outbursts

NSXBs share a similar geometry with their black hole counterparts while having the added complexity of intrinsic magnetic fields and solid surfaces originating from the neutron star. Historically, strong magnetic fields were thought to prohibit the formation of jets; however, recent radio observations have unambiguously revealed the accretion-powered jets in strongly magnetic NSXBs (e.g., van den Eijnden et al., 2018; van den Eijnden et al., 2021). Regardless, the evolution of weakly magnetic NSXBs ($< 10^{10}$ G) is known to be more analogous to BHXBs, making them a powerful *in situ* control population to determine what role BH-specific phenomena (e.g., the event horizon) play in accretion and accretion-related phenomena. Strongly magnetic NSXBs will not be discussed further.

Weakly magnetic NSXBs are typically classified as one of two subpopulations: *atoll*- and Z-sources (named for the shape of the tracks in the X-ray colour-colour diagrams, see van der Klis, 2006, for a review). Of the subpopulations, atoll sources tend to be fainter ($L_{X,\max} \lesssim 0.5\,L_{\rm Edd}$) and transient. Atolls follow the 'standard' outburst progression, exhibiting similar hard/soft states as BHXBs (see, Migliari & Fender, 2006; Muñoz-Darias et al., 2014, for a review). In NSXB nomenclature, the soft state is referred to as *banana* state, whereas the hard state is the *island* state. However, the neutron star surface is another source of X-ray photons, and, as a result, each accretion state has an additional thermal X-ray component (often modelled as a black body Lin et al., 2007). Similar to the hard-only outbursts observed in BHXBs, some atoll sources do not show state transitions, existing only in the hard state (Tarana et al., 2006; Gladstone et al., 2007; van Straaten et al., 2005), although the fraction of hard-only NSXB outbursts is yet to be explored. Additionally, and unlike BHXBs, some atolls are persistently detected in a soft accretion state (i.e., "bright atolls" such as GX 9+9: Hasinger & van der Klis 1989; Kong et al. 2006; Fridriksson 2011).

The other subpopulation of low-magnetic field NSXBs — Z-sources — are bright ($\gtrsim 0.5 L_{\rm Edd}$), persistent, and highly variable, cycling through their accretion states in $\lesssim 1$ dy. The Z sources' states — the *normal*, *horizontal*, and *flaring* branches — are ubiquitously softer than the atoll island state, suggesting that Z sources do not have a hard state equivalent (Muno et al., 2002). A small but growing fraction of NSXBs appear to transition from atoll-to-Z source behaviour at high X-ray luminosities and, thus, accretion rates. Indeed, for some sources, the transition has been observed directly (e.g., Oosterbroek et al., 1995; Shirey et al., 1998; Homan et al., 2007; Lin

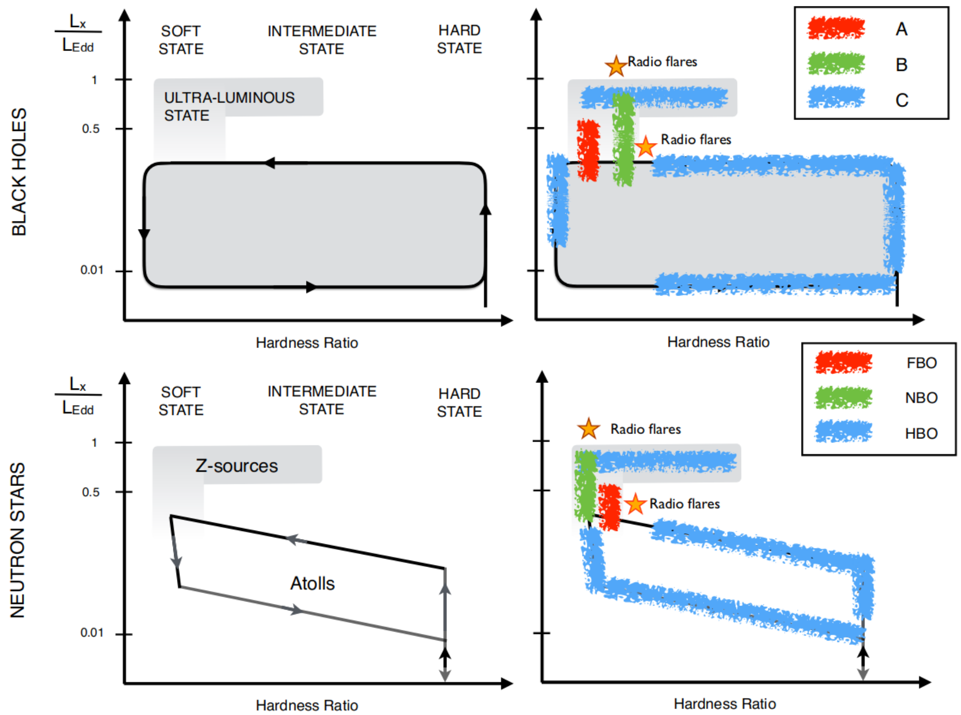

Figure 1.4: Schematic representation of the BHXB (top) and NSXB (bottom) HID track. The right panels highlight the regions of the HID and, therefore, the accretion state where each type of QPO is observed. Note the similarities between the compact object types. Reprinted from Ingram & Motta (2019), with permission from Elsevier.

et al., 2009; Altamirano et al., 2010; Chakraborty et al., 2011). These transitional sources strongly suggest that the initial description of the atoll- and Z-sources as distinct subpopulations is incorrect. Instead, like BHXBs, a single population exists, and the atoll/Z classification corresponds to different accretion states ("bright atolls" are an intermediary), with a transitional luminosity of $\sim 0.5\,L_{\rm Edd}$. In this framework, Z-source behaviour is equivalent to the ultra-luminous state observed in BHXBs, and Z-sources are neutron star analogues to semi-persistent, flaring BHXBs like GRS 1915+105 (Migliari & Fender, 2006). The analogy between the black hole and neutron star subpopulations is consistent with QPO behaviour. NSXBs exhibit three main types of QPOs: horizontal branch oscillations (HBOs), normal branch oscillations (NBOs), and flaring branch oscillations (FBOs), that follow a similar state-dependency as Type -C, -B, and -A QPOs, respectively (see Fig. 1.4; Ingram & Motta, 2019).

While the re-interpretation of the NSXB subpopulations into states is more consistent with the BHXB picture, several significant differences still exist:

1. NSXBs radio luminosities are fainter by a factor of ~ 20 at comparable X-ray luminosities. This difference in luminosity cannot be attributed to a difference in compact object mass (Gallo et al., 2018), suggesting some process intrinsically weakens the jet (or strengthens the accretion flow). Moreover, the NSXB population does not show a similarly strong L_R–L_X correlation, spanning a much broader range of L_X for a fixed L_R (see Fig. 1.3).

2. Compact jet radio emission has been seen to persist into the soft state (e.g., Migliari et al., 2004; Gusinskaia et al., 2017; van den Eijnden et al., 2021), suggesting a weaker quenching or an alternative jet mechanism exclusive to neutron stars.

3. Jet ejections during state transitions have only been observed in Z-sources (e.g., Fomalont et al., 2001; Spencer et al., 2013), and thus, only in an ultra-luminous state, and never during the canonical hard-to-soft state transition. While some atolls have exhibited radio flares (e.g., Rhodes et al., 2022), these have not occurred alongside state transitions, nor have jet ejecta been resolved. (At least until the work presented here.)

Understanding the origins of the differences above may provide critical insight into jet physics. For instance, some processes could be inhibited (or amplified) by the solid surface and/or magnetic fields of the accreting neutron star (as opposed to the horizons and ergosphere of a black hole). Historically, studies of NSXBs have suffered from their weaker radio emission and more rapid evolution. Indeed, most NS L_R–L_X monitoring has only led to a single datum point or few (e.g., $\lesssim 3$) data points per source. In contrast, individual BHXBs have been monitored for many tens of observations, and, as a result, we can distinguish between 'radio-loud' and 'radio-quiet' sources. The original motivation behind this thesis was to add high-quality NSXB monitoring, significantly improving the sample of neutron stars in the L_R–L_X plane, investigating whether NSXBs show similarly variable tracks, as seen with BHXBs.

As a result, I monitored two outbursting NSXBs, SAX J1810.8−2609 and 1A 1744−361, discussed in Chapter 3 and Chapter 4, respectively. These monitoring campaigns successfully identified fascinating outburst evolution including, but not limited to, a hard-state-only NSXB with (the possibility) for persistent radio emission and evidence for jet ejecta from an atoll NSXB; the observed signatures of jet ejecta are the first of their kind, and strongly support the wide held belief that a neutron star accretor can launch jet ejecta despite the differences between the compact object types (i.e., the event horizons and ergospheres of black holes compared to solid surfaces and internal magnetic fields of neutron stars). However, several collaborators and I quickly realized that the ThunderKAT sensitivity and cadence were insufficient to improve the NS L_R–L_X. Recently, the X-KAT collaboration that succeeded ThunderKAT, which this thesis includes data from, modified its observation strategy in response to the issues faced during the ThunderKAT NS program.

Moreover, X-KAT began focusing on full polarisation observations, and, as a result, developing polarimetric calibration, imaging, and analysis routines was identified as the highest priority. Given the expertise developed with radio polarimetry during my M.Sc. project, I switched focus to building a semi-automated polarisation analysis pipeline for X-KAT. With our new polarisation measurement capabilities, our X-KAT observations are sensitive to the magnetic field evolution of the jets, which would otherwise be 'invisible' to total flux density observations alone. In addition, given that X-KAT primarily used the MeerKAT L-band receivers (with an observing frequency of $\sim$ 1.28 GHz that is lower than often used in XRB radio observations), we can get high-precision measurements of key polarimetric quantities that require a larger range in the square of the wavelength. This allows us to probe how much the intervening plasma rotates the polarization angle as it travels to the observer, and we can correct this to derive the intrinsic polarization angle. In Chapter 5, I describe the polarimetric pipeline, the critical polarimetric properties, and the polarimetric evolution of the BHXB Swift J1727.8−163 (including the total intensity evolution). While this source exhibited a wide range of exotic jet phenomena, the new polarimetric capabilities revealed a novel time-domain evolution that suggests the jet ejecta were comprised of an electron-proton plasma (rather than an electron-positron pair-plasma). Understanding the composition of jets is critical to understanding the total kinetic energy content available for feedback interactions (as protons are $\sim 1000\times$ more massive than positrons); X-KAT can now look for similar signatures (or their absence) in future outbursts, investigating whether the majority of jet ejecta have similar compositions.

The astrometric evolution of SAX J1810.8−2609 and Swift J1727.8−163, alongside broader discussion amongst the collaboration, motivated my investigation into the astrometric precision of our observations. I present the results of that investigation in Chapter 6 and show that the current *ad hoc* assumptions overestimate the astrometric errors significantly for MeerKAT snapshot observations, with the potential that many radio observations across facilities overestimate their astrometric errors. Furthermore, in both Chapter 5 and Chapter 6, I discuss how accurate astrometry is critical for the kinematic modelling of jet ejecta as they propagate through and interact with the ISM — kinematic modelling seems to be the best estimator of the total energy content in the jets (if their compositions are known).

Finally, in Chapter 7, I combine the results from the outbursts of SAX J1810.8−2609, 1A 1744−361, and Swift J1727.8−163, as well as the newly developed polarization and astrometric techniques I pioneered. I conclude with some perspectives on futurefhard studies.

Chapter 2

Radio Interferometry and Imaging

The primary instrument for this thesis is the MeerKAT Radio Telescope (Jonas & MeerKAT Team, 2016, hereafter MeerKAT). MeerKAT is a radio interferometer that consists of a 64-element array of 12.5 m dishes used to synthesize a single 8 km aperture. Throughout this work, I frequently reference standard radio interferometric *jargon* when describing the calibration and imaging necessary to extract science-ready products. These processes share few similarities with standard semiconductor-based (e.g., CCDs and CMOS detectors) observatories and single-dish radio telescopes. Thus, the remainder of this chapter will introduce radio interferometry and its terminology.

2.1 Fundamentals of Radio Interferometry

2.1.1 The 2-Element Interferometer

In contrast to 'classical' (e.g., optical) astronomy, the energies of radio photons are far below the photon ionization thresholds of semiconductors.

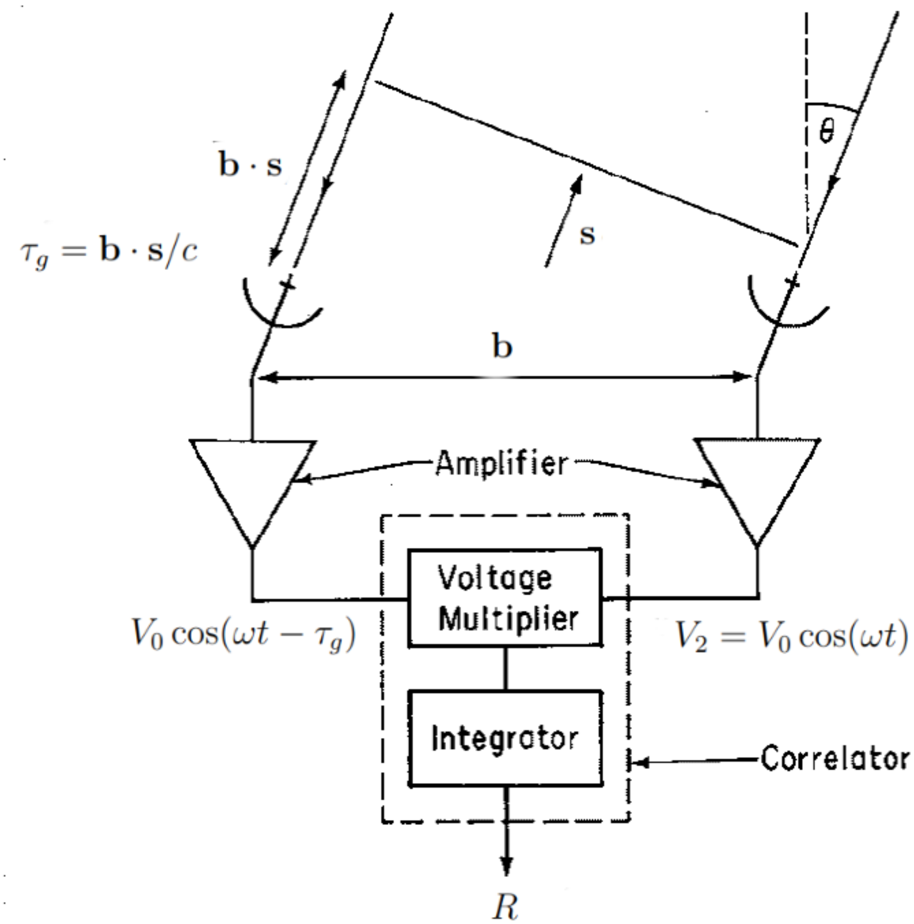

Figure 2.1: Simple schematic representation of the geometry of a 2-element interferometer (Adapted from; Taylor et al., 1999).

As a result, individual radio telescopes (hereafter referred to as antennas) measure source properties utilizing the wave-like nature of light. Consider the 2-element array shown in Fig. 2.1, and for simplicity, assume that the incident radiation is a monochromatic plane wave. The electric field of the incident radiation will induce a time-varying voltage in each antenna such that

$$V_1 = V_0\cos(\omega t - \tau_g), \tag{2.1}$$

$$V_2 = V_0\cos(\omega t), \tag{2.2}$$

where the angular frequency, $\omega = 2\pi\nu$, is dependant on the frequency (ν) of the incident radiation, and $\tau_g = \mathbf{b} \cdot \mathbf{s}/c$ is the geometric delay caused by the difference in path length between the source and the two antennas.

For single-dish telescopes, the radio intensity (within the field of view) can be measured from the amplitude of the incident voltage. However, interferometers reconstruct the radio sky using the coherence (or 'similarity') of the signals received by each antenna. The incident signals are amplified, and the coherence is measured in a correlator, calculating a time-averaged product of the two voltages. The correlator output or *response* (R) is

$$R = \langle V_1 V_2 \rangle_t = V_0^2 \langle \cos(\omega t - \tau_g) \cos(\omega t) \rangle_t \tag{2.3}$$

and

$$R = V_0^2 \langle \cos^2(\omega t) \cos(\omega \tau_g) + \cos(\omega t) \sin(\omega t) \sin(\omega \tau_g) \rangle_t, \tag{2.4}$$

where $\langle \cdot \rangle_t$ is the time-average operator. For simplicity, we assume a common amplitude of V_0, although this need not be the case. The correlator averages the input over a time interval, T, such that $\tau_g \gg T \gg \nu^{-1}$. As a result, $\cos^2(\omega t) \to 1/2$ and $\sin(\omega t)\cos(\omega t) \to 0$, and R becomes a function of τ_g:

$$R = V_0^2 \cos(\omega \tau_g). \tag{2.5}$$

The total interferometric response for a collection of plane waves originating from different sky locations can be written as a function of the sky intensity distribution, $I(\mathbf{s})$. As the signals from each direction are spatially and temporally incoherent, the only non-zero coherence comes from signals along the same $\mathbf{s}$-direction, and thus, the response is an integral over the celestial sphere:

$$R \propto \Delta\nu \int_{\Omega} \mathcal{A}(\mathbf{s}) I(\mathbf{s}) \cos(\omega\tau_g) d\Omega, \tag{2.6}$$

where $\Delta\nu$ ($\ll \nu$, for the mono-chromatic assumption) is the observing bandwidth, and $\mathcal{A}(\mathbf{s})$ is the antenna response or normalized *primary beam* response. While the integral is over the entire celestial sphere, $\mathcal{A}(\mathbf{s}) \to 0$ outside of a small angular region. The centre of this region, $\mathbf{s_0}$, known as the *phase centre*, is the pointing direction of the individual antennas; $\mathcal{A}(\mathbf{s}) \equiv 1$ for $\mathbf{s} = \mathbf{s_0}$.

Equation (2.6) was left as a proportionality, as it is only sensitive to one component of the sky intensity distribution. Any real function can be separated into the sum of an odd and even function, i.e., $I(\mathbf{s}) = I_O(\mathbf{s}) + I_E(\mathbf{s})$. The above derivation implicitly assumed a *cos*-correlator, which measures the even component of the sky brightness distribution. The odd component can be measured using a *sin*-correlator, a *cos*-correlator with a $\frac{\pi}{2}$-phase offset introduced into one of the signal paths:

$$R_{\cos} = V_0^2 \cos(\omega\tau_g), \tag{2.7}$$

$$R_{\sin} = V_0^2 \sin(\omega\tau_g). \tag{2.8}$$

Modern interferometers use *complex*-correlators, simultaneously measuring the even and odd components:

$$\bar{R} = R_{\cos} - iR_{\sin} \tag{2.9}$$

$$\bar{R} = \Delta\nu \int_{\Omega} \mathcal{A}(\mathbf{s}) I(\mathbf{s}) e^{-i\omega\tau_g} d\Omega. \tag{2.10}$$

From here, let $\mathbf{s} = \mathbf{s}_0 + \sigma$, so that the position of any point on the celestial sphere is written as an offset with respect to the direction to the phase centre (see Fig. 2.2). Substituting in $\omega\tau_g = 2\pi\frac{\mathbf{b}\cdot\mathbf{s}}{\lambda}$, Equation (2.10) adopts

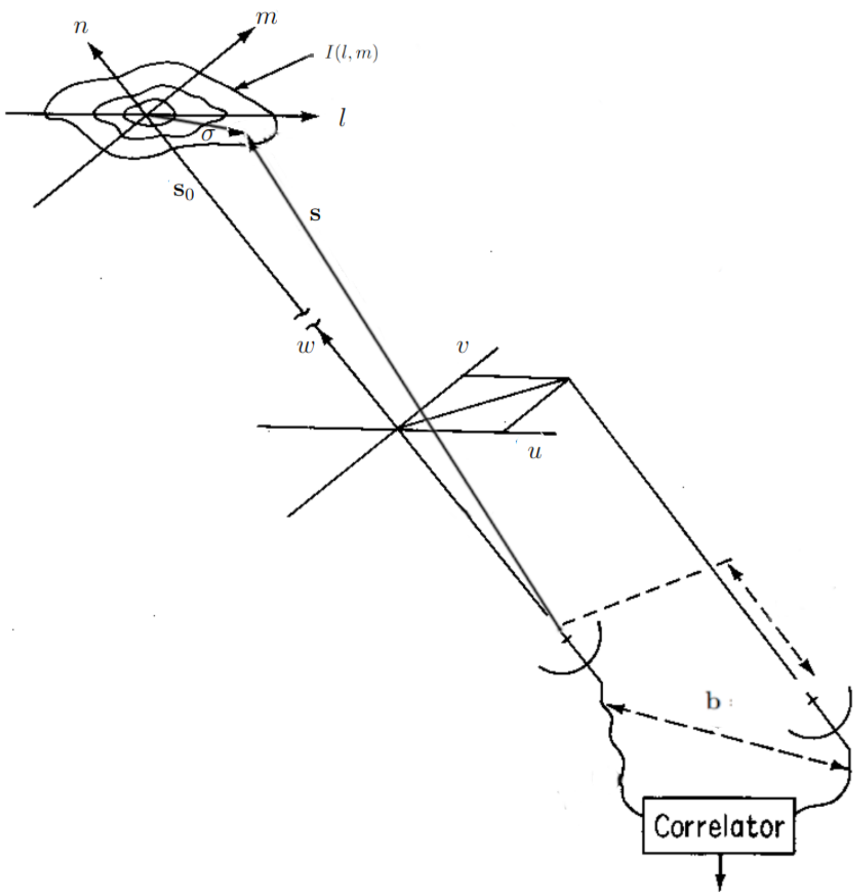

Figure 2.2: Representation of standard interferometric coordinate system. The intensity distribution is described in the (l, m)-plane, whereas the baseline vector is described in the (u, v, w)-plane (Adapted from; Taylor et al., 1999).

its commonly reported form:

$$\bar{R} = \Delta\nu e^{-2\pi i \frac{\mathbf{b}\cdot\mathbf{s}_0}{\lambda}} \int_\Omega \mathcal{A}(\sigma) I(\sigma) e^{-2\pi i \frac{\mathbf{b}\cdot\sigma}{\lambda}} d\Omega. \tag{2.11}$$

The integrand corresponds to the **key** interferometric quantity, the complex *visibility* $\mathcal{V}$:

$$\mathcal{V} \equiv \int_\Omega \mathcal{A}(\sigma) I(\sigma) e^{-2\pi i \frac{\mathbf{b}\cdot\sigma}{\lambda}} d\Omega, \tag{2.12}$$

which is related to the response of a *complex*-correlator,

$$\bar{R} = \Delta\nu |\mathcal{V}| e^{i\left(\phi_{\mathcal{V}} - 2\pi \frac{\mathbf{b}\cdot\mathbf{s}_0}{\lambda}\right)}, \tag{2.13}$$

modified by the characteristics of the interferometer and observing conditions. The interferometer measures the response interference pattern, which is used to solve for the visibility amplitude ($|\mathcal{V}|$) and phase ($\phi_{\mathcal{V}}$) during calibration (see Section 2.2).

As an instrument, a 2-element interferometer behaves like a spatial filter sensitive to intensity variations on particular angular scales. The characteristic size of the angular scale is determined by the observing wavelength and the projected length of the baseline vector on the plane of the sky ($\mathbf{b}_{\mathrm{proj}}$),

$$\theta_s \sim \frac{\lambda}{|\mathbf{b}_{\mathrm{proj}}|}. \tag{2.14}$$

Here, θ_s represents the smallest angular scale over which a 2-element interferometer can measure intensity variations. Recognizing that $\mathbf{b}$ can be made arbitrarily large by moving the antenna further apart, interferometers can achieve sub-arcsecond angular resolutions without requiring kilometre-scale apertures. Modern arrays combine many tens of antennas corresponding to hundreds to thousands of baselines, and thus thousands of 2-element interferometers, each sensitive to a different spatial scale.

2.1.2 N-Element Interferometer

For the N-element interferometer, we define the standard coordinate system used to describe each baseline ($\mathbf{b}$) and sky position vector ($\mathbf{s}_0\,, \mathbf{s}\,, \sigma$):

$$\mathbf{b} = \begin{pmatrix} u \\ v \\ w \end{pmatrix} \quad \mathbf{s} = \begin{pmatrix} l \\ m \\ n \end{pmatrix} \quad \mathbf{s}_0 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}. \tag{2.15}$$

The baseline coordinates (u, v, w) are expressed in units of observing wavelength; the sky coordinates $(l,\ m)$ are the direction cosines along the u, v-axis; both n and w point towards the phase centre (see Fig. 2.2). By definition, $n = \sqrt{1 - l^2 - m^2}$. Given that $\sigma = \mathbf{s} - \mathbf{s_0}$, equation (2.12) can be rewritten as

$$\mathcal{V}(u, v, w) = \int_{\Omega} \mathcal{A}(l, m) I(l, m) e^{-2\pi i \left[ul + vm + w(\sqrt{1-l^2-m^2}-1)\right]} d\Omega, \tag{2.16}$$

where a small interval of the celestial sphere is:

$$d\Omega = \frac{dldm}{n} = \frac{dldm}{\sqrt{1 - l^2 - m^2}}. \tag{2.17}$$

Thus,

$$\mathcal{V}(u, v, w) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathcal{A}(l, m) I(l, m) e^{-2\pi i \left[ul + vm + w(\sqrt{1-l^2-m^2}-1)\right]} \frac{dldm}{\sqrt{1 - l^2 - m^2}}. \tag{2.18}$$

Applying the small-field approximation (l, $m \ll 1$) $\sqrt{1-l^2-m^2} \sim 1$, Equation (2.18) becomes part of a two-dimensional Fourier transform pair:

$$\mathcal{V}(u,v) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \mathcal{A}(l,m)I(l,m)e^{-2\pi i(ul+vm)}dldm, \text{ and} \tag{2.19}$$

$$I(l,m) = \frac{1}{\mathcal{A}(l,m)}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \mathcal{V}(u,v)e^{2\pi i(ul+vm)}dudv. \tag{2.20}$$

Therefore, an 'image' with an interferometer (hereafter *synthesis imaging*) is the Fourier transformation of the measured visibilities. Note, the quantity, *intensity*,

$$I \approx \frac{\mathrm{d}P}{\mathrm{d}\nu\mathrm{d}\Omega\mathrm{d}A}, \tag{2.21}$$

is the power per unit frequency, per unit area, per unit solid angle, radiating from a sky region, typically expressed in units of *Jansky per Beam* ($\mathrm{Jy\,beam^{-1}}$). The unit 'beam' quantifies the two-dimensional response of the instrument (i.e., the angular resolution discussed in Section 2.3). A related quantity, the *Flux Density*, $F = \int I\mathrm{d}\Omega$, removes the solid angle dependence by integration over a (subregion) of the celestial sphere. Only flux densities can be measured for unresolved sources.

2.1.3 Interferometry and the Vector Properties of Light

The preceding sections have implicitly assumed scalar electric fields; in reality, electromagnetic radiation is a vector phenomenon, and any arbitrarily polarised wave can be described as a superposition of two orthogonally polarised components. As a result, the antennas in most interferometers are equipped with dual, orthogonally polarised receivers. The properties of polarised radiation can be characterized with four intensities called the *Stokes* parameters (Stokes I, Q, U, and V), which are functions of the amplitudes and phase offsets between the electric field vectors measured in the orthogonal receivers. For a qualitative description, Stokes I is equivalent to the total intensity, Q and U are related to the linear polarisation properties, and V is the circularly polarised intensity. The latter three Stokes parameters are used to calculate commonly reported polarisation metrics:

$$P = \sqrt{Q^2 + U^2}, \tag{2.22}$$

$$m_l = P/I, \tag{2.23}$$

$$m_c = V/I, \text{ and} \tag{2.24}$$

$$\text{EVPA} = \frac{1}{2}\arctan(U/Q), \tag{2.25}$$

where P is the linearly polarised intensity, m_l is the linear polarisation fraction, m_c is the circular polarisation fraction, and EVPA is the electric vector position angle (otherwise known as the linear polarisation angle).

Dual-feed interferometers have four unique visibilities, as each receiver has a complex conjugate, and the Stokes parameters are linear combinations of these four visibilities. For circularly polarised receivers,

$$\begin{pmatrix} I & Q \\ U & V \end{pmatrix} = \begin{pmatrix} \frac{1}{2}(\mathcal{V}_{RR^*} + \mathcal{V}_{LL^*}) & \frac{1}{2}(\mathcal{V}_{RL^*} + \mathcal{V}_{LR^*}) \\ \frac{1}{2i}(\mathcal{V}_{RL^*} - \mathcal{V}_{LR^*}) & \frac{1}{2}(\mathcal{V}_{RR^*} - \mathcal{V}_{LL^*}) \end{pmatrix}, \tag{2.26}$$

and for linearly polarised receivers,

$$\begin{pmatrix} I & Q \\ U & V \end{pmatrix} = \begin{pmatrix} \frac{1}{2}(\mathcal{V}_{XX^*} + \mathcal{V}_{YY^*}) & \frac{1}{2}(\mathcal{V}_{XX^*} - \mathcal{V}_{YY^*}) \\ \frac{1}{2}(\mathcal{V}_{XY^*} + \mathcal{V}_{YX^*}) & \frac{1}{2i}(\mathcal{V}_{XY^*} - \mathcal{V}_{YX^*}) \end{pmatrix}, \tag{2.27}$$

where the subscripts denote the polarisation of the feeds (e.g., R is right-hand circular polarised), and the asterisks denote a complex conjugate. MeerKAT is a linear-feed interferometer (i.e., it uses linearly polarised receivers).

2.2 Calibration and Flagging

For any observation, the *observed* visibilities ($\mathcal{V}^{\mathbf{obs}}$) are modified from the *true* visibilities ($\mathcal{V}^{\mathbf{true}}$). The extent of the modification is determined by the properties of the interferometric array (e.g., artificial spectral shapes introduced by the electronics) and the local observing conditions (e.g., atmospheric phase variations or ionospheric rotation of the EVPA). Synthesis imaging of the observed visibilities will deviate from the true sky intensity distribution, 'smearing' the signal across the field of view (top panel, Fig. 2.3). Recovering the true visibilities is the purpose of *calibration*, although strong external radio signals (radio frequency interference; RFI) will render some fraction of visibilities irrecoverable. The first step in any calibration procedure is to excise (i.e., *flag*) the RFI-affected visibilities; subsequent iterations of flagging can be necessary if partially calibrated data uncovers additional RFI.

Assuming the irrecoverable data have been removed, the primary calibration method, first-generation (1GC) or *reference* calibration, utilizes bright ($\sim$ Jy) calibrator sources with known visibility models ($\mathcal{V}^{\mathrm{mod}} \sim \mathcal{V}^{\mathrm{true}}$) to correct the visibilites of the target source. Comparing the observed visibilities of the calibrators to their respective models, we solve for multiplicative *gain terms* (C_{ij}) for an arbitrary baseline containing antenna i and j:

$$\mathcal{V}_{\mathrm{cal}}^{\mathrm{obs}} = C_{ij} \mathcal{V}_{\mathrm{cal}}^{\mathrm{mod}}. \tag{2.28}$$

recovering the true visibilities:

$$\mathcal{V}_{\mathrm{src}}^{\mathrm{true}} = C_{ij}^{-1} \mathcal{V}_{\mathrm{src}}^{\mathrm{obs}}. \tag{2.29}$$

Most calibration effects occur before signal correlation, allowing for the decomposition of the baseline terms into a product of antenna-based terms,

$C_{ij} = C_i C_j^*$, significantly reducing the dimensionality of calibration. The generalized gain terms (C_i and C_j) include multiple separable effects. The solution order is critical as each subsequent gain term is derived from the partially-calibrated visibilities. The gains terms can be broken into parallel- and cross-hand effects; $\mathcal{V}_{XX^*}$ and $\mathcal{V}_{XY^*}$ are examples of parallel-hand and cross-hand visibilities, respectively. If polarisation is not a scientific interest, only parallel-hand calibration is required. While there is no universal calibration routine, below is a standard order for solving the gain terms:

- **Delay** (K-term): A linear-in-frequency phase term (i.e., 'delay') between the parallel-hand feeds caused by the different signal propagation times. Delay is typically stable over many hours. Like all other phase terms, the solutions are measured with respect to a (user-specified) reference antenna. The calibrator can be any bright source.

- **Bandpass** (B-term): Phase and amplitude corrections, which remove higher-order (i.e., non-linear) frequency effects. For most interferometers, the shape of the bandpass is typically stable over many hours. The calibrator must have a well-known spectrum and spatial intensity distribution.

- **Complex Gain** (G-term): Phase and (relative) amplitude corrections, which remove higher-order (i.e., non-linear) temporal drifts. Complex gain is solved using any bright, non-variable calibrator source (often point-like); complex gain solutions do not require a spectral model. Temporal variability affects phase much more rapidly than amplitude and can require second-scale solution intervals.

- **Flux** (F-term): Absolute amplitude correction. Suppose a G-term calibrator does not have a spectral model. In that case, most software assumes an unpolarised ~ 1 Jy point-source, and thus, the amplitude G-terms require a time-independent absolute correction. The amplitude solutions are scaled to their absolute value using a calibrator with a known spectrum (almost always the B-term calibrator).

These four gain terms are sufficient for parallel-hand calibration. Due to the primary-beam sensitivity drop-off, calibrators are rarely visible when pointing at the target field. Therefore, interferometric observations will consist of multiple pointings or *scans*, which separately point at the calibrator and target fields. Typical total-intensity-only observations will consist of two calibrators, a bandpass (i.e., *primary*) and complex gain (i.e., *secondary*) calibrator. While a primary alone can derive all calibration terms, these sources are few and far between ($\sim$10 total over the celestial sphere). The G-terms applied to the target must be derived from a nearby line of sight, as atmospheric effects are anisotropic and time-variable. As a result, when choosing a secondary, proximity to the source outweighs *a priori* spectral information. A typical observation would adopt a cadence of primary $\rightarrow$ secondary $\rightarrow$ target $\rightarrow$ secondary($\rightarrow$ target $\rightarrow$ secondary...); the target–secondary cycling is dependant on the gain decorrelation timescales. The gain terms are (linearly) interpolated onto the target field using the solutions derived from the calibrators.

For polarisation observations, cross-hand calibration follows the parallel-hand solutions:

- **Leakage** (D-term): Slight non-orthogonalities between the feeds in each antenna result in the artificial conversion of unpolarised-to-polarised intensity (i.e., $I \leftrightarrow Q$, U, or V). Leakage corrections are typically stable over many hours. The leakage is corrected using an unpolarised calibrator or a calibrator observed over a large range of parallactic angles (See P-term). The observations presented in this thesis are short-duration (i.e., $\lesssim 1\,$hour) and, thus, do not have the parallactic angle coverage for the latter correction.

- **Cross-hand Delay** (KCROSS): Cross-hand analogue of the K-term. It measures a linear-in-frequency phase slope between the cross-hand correlations and should be baseline-independent after applying the K-term solution. Cross-hand delays are typically stable over many hours and are corrected using a strongly polarised calibrator.

- **Cross-hand Phase** (X-term): Higher-order (in frequency) phase-offset between the cross-hand receivers. After applying the preceding calibration terms, the X-term is antenna-independent. X-term corrections are typically stable over many hours and are corrected using a strongly polarised calibrator with well-modelled circular polarisation properties (often $V = 0$).

- **Parallactic Angle** (P-term): Most antennas are AltAz-mounted. As a result, the feed orientations do not rotate with the celestial sphere, and thus, the measured polarisation angle rotates with the parallactic angle. This correction does not require a calibrator source, only *a priori* knowledge of antenna locations, pointing direction, and observing time(s).

From here, the visibilities are ready for polarisation imaging. However, interpolating gain terms from the calibrator to target visibilities will leave residual phase and amplitude errors if, for example, the phase exhibits non-linear variability on timescales shorter than the target-secondary cycling. In the third panel of Fig. 2.3, the bright central point source has clear striations due to these residual calibration errors. The limitations of reference calibration gave rise to next-generation calibration techniques that use the target to calibrate itself, i.e., 'self-calibration'. While paradoxical at face value, self-calibration is crucial for high dynamic range imaging in nearly all cases. The standard self-calibration approach adopts the images produced from reference calibrated visibilities as models of the target, improving the gain terms by solving on finer frequency binning or, more commonly, by solving the gain terms on time intervals shorter than the scan-cycle time. The bottom panel of Fig. 2.3 shows how 'second-generation self-calibration' (2GC) improves image fidelity.

Both reference calibration and 2GC self-calibration assume direction-independence. Given that $\mathcal{V}$ is a superposition of interfering signals from many sources, direction independence suggests that the per-antennas gain terms are sufficient for any sky location. More precisely, we assume that each antenna and each baseline view the same sky intensity distribution. In reality, effects such as antennas pointing errors, complex (asymmetric) primary beam shapes, and variable path lengths through the ionosphere violate the *same-sky* assumption. These direction-dependent effects are most prominent for off-axis sources (i.e., sources significantly offset from the phase centre). The worsening amplitude and phase errors manifest as nonphysical structures and artificially induced polarisation. Treating direction-dependent effects (DDEs) has led to third-generation calibration (3GC). As these effects are multiplicative in the image plane, they are convolutions in visibility-space with time-dependent, antenna-dependant convolution functions, severely complicating the calibration procedure (see, e.g., Smirnov, 2011a; Smirnov & Tasse, 2015; de Villiers, 2023, for a discussion on DDEs and 3GC). Fortunately, the work in this thesis focuses on near-field sources well-calibrated with reference and 2GC self-calibration. Aside from correcting non-zero w-terms in Equation (2.18), we exclude 3GC from this work.

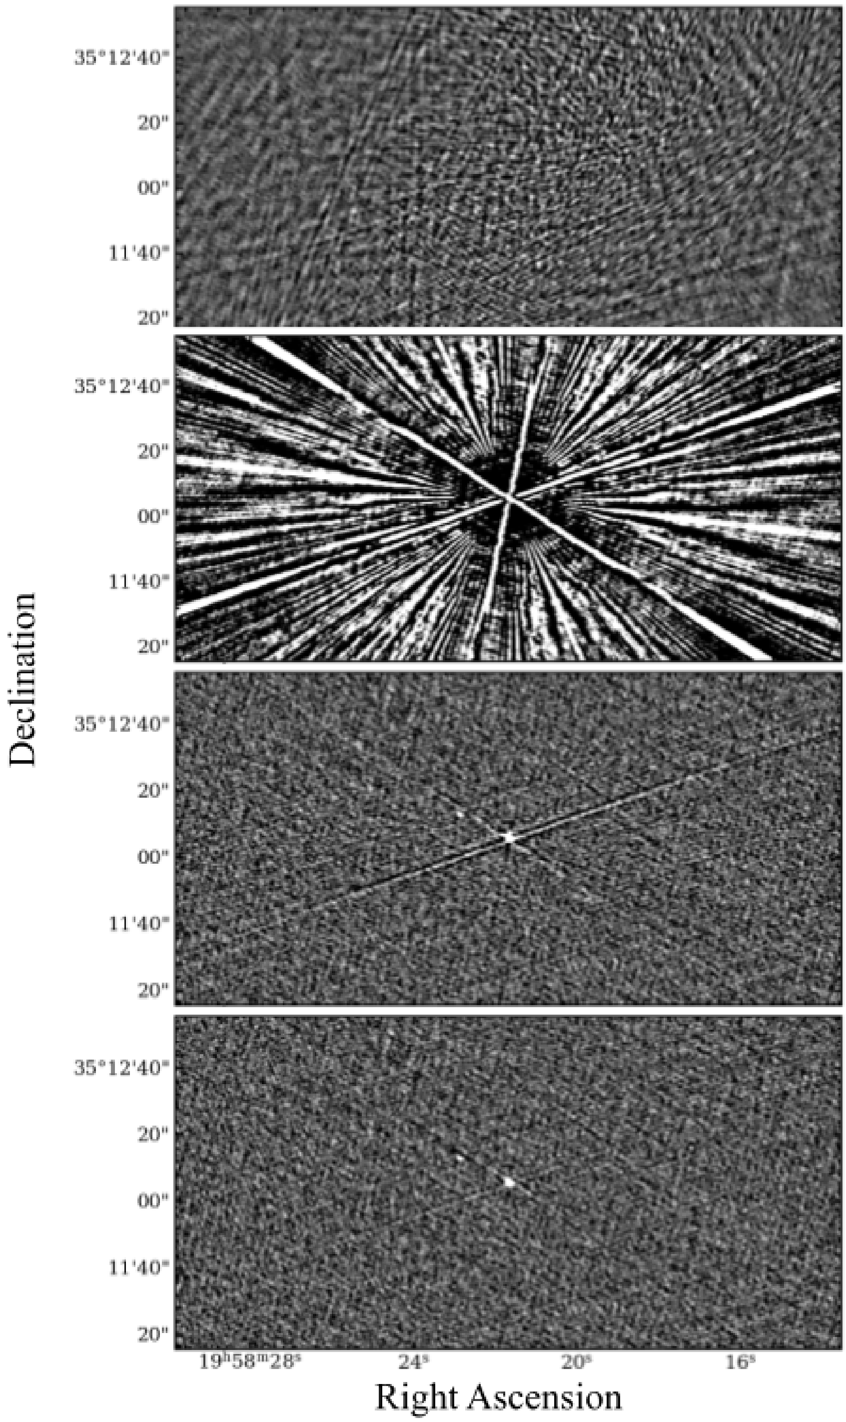

Figure 2.3: Sample images of a target field dominated by a central point source: (top panel) uncalibrated data; (second panel) reference calibrated data, dirty image; (third panel) reference calibrated data, restored image; (fourth panel) self-calibrated data, restored image. The maximum/minimum scale of the top panel is 10% of the remaining three for visualization purposes. Note that the uncalibrated data has no central source, as intensity is smeared across the entire field of view. These data were taken with the Very Large Array (VLA).

2.3 Synthesis Imaging

2.3.1 Fundamental Angular Scales

Ignoring finite sampling effects, the idealized case of continuous uv-coverage has baselines up to some maximum baseline length $\sqrt{u^2+v^2} \leq d_{\mathrm{max}}$. The angular resolution is dependent on the maximum baseline length and the observing frequencies:

$$\Theta_{\mathrm{PSF}} \sim \frac{\lambda}{d_{\mathrm{max}}}, \tag{2.30}$$

adopting the same form as the Rayleigh criterion. The angular scale, Θ_{PSF}, is the Full-Width at Half-Maximum (FWHM) of the point-spread function (PSF), which is the instrument response to a point source. Mathematically, the idealized PSF adopts the form of a two-dimensional *sinc*-function.

Interferometers also have a largest resolvable angular scale (LAS), dependent on the smallest baseline in the array ($b_{\mathrm{min}} \leq \sqrt{u^2+v^2}$),

$$\Theta_{\mathrm{LAS}} \sim \frac{\lambda}{d_{\mathrm{min}}}. \tag{2.31}$$

The diameter of the dish for a single antenna, D_{app}, is the physical limit of b_{min} and also describes the field of view (FOV):

$$\Theta_{\mathrm{FOV}} \sim \frac{\lambda}{D_{\mathrm{app}}}, \tag{2.32}$$

where Θ_{FOV} is the FWHM of the primary beam response. Many arrays have $\Theta_{\mathrm{LAS}} < \Theta_{\mathrm{FOV}}$, becoming insensitive to the signals from sources with angular scales $> \Theta_{\mathrm{LAS}}$, even if they are contained within the field of view. Sources with angular sizes that are $\gg \Theta_{\mathrm{LAS}}$, are completely undetectable or 'resolved out'.

2.3.2 Finite Sampling in the uv-plane

Continuous uv-plane coverage is *effectively* impossible over the distances used in radio astronomy; instead, interferometers sample the uv-plane with a finite number of baselines. The sample visibilities, $\mathcal{V}^{\text{samp}}$, are a product of the *true* visibilities, $\mathcal{V}^{\text{true}}$, and the uv-sampling function, $\mathcal{S}$:

$$\mathcal{V}^{\text{samp}}(u,v) = \mathcal{S}(u,v) \cdot \mathcal{V}^{\text{true}}(u,v), \tag{2.33}$$

where $\mathcal{S}(u,v)$ is a sum of δ-functions for M samples:

$$\mathcal{S}(u,v) = \sum_{n=1}^{M} \delta(u - u_n, v - v_n). \tag{2.34}$$

Given that the visibilities and the sky intensity distribution are related through the Fourier transform (represented with $\mathcal{F}$), a product in visibility space becomes a convolution in imaging space:

$$\mathcal{F}\left[\mathcal{V}^{\text{samp}}(u,v)\right] = \mathcal{F}\left[\mathcal{S}(u,v)\right] * \left[\mathcal{V}^{\text{true}}(u,v)\right] \text{ or} \tag{2.35}$$

$$I^{\text{dirty}}(l,m) = B^{\text{dirty}}(l,m) * I^{\text{true}}(l,m). \tag{2.36}$$

The observed sky brightness distribution or *dirty* image, $I^{\text{dirty}}(l,m)$, is the convolution of the *true* sky brightness distribution, $I^{\text{true}}(l,m)$, and the *dirty beam*, $B^{\text{dirty}}(l,m)$. The dirty beam is the colloquial term used to describe the PSF, including finite sampling effects. The gaps in the sampling cause the PSF to deviate from a two-dimensional *sinc*-function, causing bright non-physical structures, or 'artefacts', in dirty images (second panel, Fig. 2.3).

The most straightforward approach for artefact mitigation is an improvement in uv-plane coverage; $\mathcal{S}(u,v) \to 1$, $B^{\text{dirty}}(l,m) \to \delta(l,m)$ and, as a result, $I^{\text{dirty}}(l,m) \equiv I^{\text{true}}(l,m)$. Construction of more antennas is often expensive and unfeasible, so uv-coverage is increased by combining observation at different times ('Earth rotation synthesis') and frequencies ('Multi-frequency synthesis'). The applicability of the former is the result of $\mathbf{b}$ being a vector quantity. As the Earth rotates $|\mathbf{b}| = \sqrt{v^2 + u^2}$ will remain (approximately) constant, but the relative contributions from u and v will vary, sweeping out tracks in the uv-plane. One caveat: combining visibilities observed at different times implicitly assumes a static sky. Earth rotation synthesis can worsen artefacts if sources have variable timescales shorter than the length of the observation.

Additionally, as $\mathbf{b}$ is measured in units of λ, uv-plane sampling can be improved by combining visibilities from different wavelengths. Like temporal variability, the visibilities can show frequency structure(s) worsening artefacts. However, unlike temporal changes, spectra are almost guaranteed to be smoothly varying functions, well described by an n^{th}-order polynomial is linear or logarithmic space. Nearly all modern imaging software includes multi-frequency synthesis capabilities, utilizing the smoothness of spectral variations to remove frequency artefacts.

For most observations, improvements in uv-plane coverage are necessary yet insufficient for artefact removal; the work-horse of high-fidelity synthesis imaging comes from the deconvolution algorithms.

2.3.3 Imaging and Deconvolution

The majority of deconvolution software implements either a CLEAN (Högbom, 1974) or Maximum Entropy Method (MEM; Gull & Daniell, 1978; Frieden, 1972) algorithm. This work uses the WSCLEAN imager (Offringa et al., 2014), which adopts the former algorithm. As a result, the remainder of this section will introduce the CLEAN deconvolution algorithm, and MEM will not be discussed further.

Modern CLEAN-based imagers adopt the approach introduced by Schwab

(1984), which separates visibility and image plane manipulations into *major* and *minor* cycles, respectively. Deconvolution begins with interpolating the unevenly sampled visibilities onto a uniformly spaced *uv*-grid. A uniform grid enables Fast Fourier Transforms (FFTs), significantly decreasing the computational requirements. After gridding, the visibilities are Fourier transformed, producing the dirty image and initiating minor cycles that,

Step 1. Identify the intensity and position of the brightness pixel in the image.

Step 2. Subtract, from the position of the peak, a δ-function, convolved with the dirty beam and multiplied by some gain damping factor (typically ~ 0.1). This source-subtracted image is called the *residual* image.

Step 3. Record the position and magnitude of the subtracted component in a model image.

Step 4. Go to Step 1 unless the peak pixel reaches the minor cycle-stopping threshold.

The stopping threshold occurs when the strength of the peak component during a minor cycle is $\sim 20\%$ of its starting value, initiating a major cycle. Major cycles consist of:

Step 1. De-grid the observed and model visibilities to their original, unevenly sampled *uv*-values.

Step 2. Subtract the model visibilities from the observed visibilities (performing this operation in visibility space prevents gridding errors and aliasing).

Step 3. Re-grid the subtracted visibilities and initiate a new round of minor cycles unless the major-cycle stopping threshold is reached.

Reaching the major-cycle stopping threshold marks the end of the CLEAN algorithm. It occurs after a user-specified number of minor-cycle iterations or when the peak intensity drops below a user-specified value. The latter stopping condition is more widely used, and the stopping intensity is typically chosen to be (1–3)$\times$ the expected image plane rms noise. Ideally, the residual image will be artefact-free, the remaining signal being thermal noise or confusion of faint, unresolved sources. The model image is then convolved with an idealized PSF (also known as the *synthesized* beam) and added to the residual images, creating the final or *restored* image. The idealized PSF is a two-dimensional Gaussian fit to the main lobe of the dirty beam (with an FWHM similar to Equation 2.30).

The CLEAN algorithm is procedural. As a result, the restored image is highly dependent on user-defined imaging parameters (e.g., stopping conditions, pixel size, visibility weights). Some of these parameters follow observation-independent recommendations. For example, no physical pixels exist, so the image dimensions are free parameters. However, the FWHM of the synthesized beam should contain $\gtrsim 3$ pixels (to avoid significant sub-pixel intensity variations), and the restored image should span an angular area $\gtrsim 2 \times \theta_{\rm FOV}$ (to avoid aliasing bright objects near the edge of the primary beam). The most critical observation-dependent parameters are the choice of visibility gridding and weighting schemes.

The gridding scheme controls the re-mapping of the visibility samples onto an evenly-spaced uv-grid. For targets consistent with the small-field approximation, and thus Equation (2.19), visibilities can be trivially mapped onto a single grid under the assumption of a co-planar array ($w \sim 0$). However, the design of modern interferometers prioritizes wide field sensitivity, enabling single-pointing observations to image increasingly larger regions of the radio sky at the cost of the co-planar assumption. Imaging algorithms can no longer ignore the 'w-term' in Equation (2.18), complicating the transformation between the visibility and the imaging planes. An accurate treatment of the w-terms is one of the most crucial aspects of high-fidelity imaging with wide-field interferometers. There are several methods (of varying accuracy and computation efficiency) to correct for w-terms, such as faceting (Perley, 1999), w-projection (Cornwell et al., 2008), w-stacking (Offringa et al., 2014), and, the scheme used for this work, w-gridding (Ye et al., 2022; Arras et al., 2021). While a complete description of the w-gridding algorithm is far beyond the scope of this work, in (relatively) simplistic terms, using an optimal convolution function, gridding can be extended along the w-direction, mapping each visibility onto a narrow range of w-values based on the exact (or un-gridded) w-value.

The second key observation-dependent consideration is the choice of visibility weighting function(s). The sampling function shown in Equation (2.34) can be generalized into a *weighted* sampling function:

$$W(u,v) = \sum_{n=1}^{M} D_n T_n \delta(u - u_n, v - v_n), \tag{2.37}$$

comprising a density weighting, D_n, and tapering, T_n, function. For D_n, the simplest cases are *natural* (i.e., un-weighted),

$$D_n = 1 \tag{2.38}$$

and *uniform* weightings,

$$D_n = \frac{1}{N_k(n)}, \tag{2.39}$$

where $N_k(n)$ is the number of points in the uv-plane within an arbitrary region of size k, centred on visibility n. Most uv-sampling comes from shorter baselines, so uniform weighting down-weights the visibilities from large-scale structures. Moreover, the down-weighting reduces the significance of the dirty beam side lobes and optimizes for smaller PSFs at the cost of an increase in the rms noise (compared to natural weights). The most commonly used weighting type is the hybrid 'Briggs' weightings (Briggs, 1995a). Briggs weightings minimize the sum of the power in the dirty beam sidelobes and the thermal noise. Which of the two components is more important is controlled by the *robustness* parameter Briggs (1995a). The robustness is a real number $\in(-2, 2)$ that behaves like a scale between the extremes; values of -2 and 2 approximate uniform and natural weightings, respectively. More natural weightings are meant to optimize signal-to-noise (by minimizing rms noise) but are susceptible to residual artefacts from incomplete deconvolution; they also lead to larger PSFs. Bright extended sources are more prone to incomplete deconvolution, as CLEAN models the sky intensity distribution as a collection of point sources. In extreme cases, uniform weightings that down-weight side-lobe power can improve signal-to-noise by mitigating artefacts. An example of the effects of different weightings is shown in Fig. 2.4 and Fig. 2.5.

Taper functions are another control of the interferometer's sensitivity to specific angular scales. Inner and outer tapers down-weight short- and long-baselines, making the interferometer less sensitive to large- and small-scale structures, respectively; this research does not utilize outer tapers. With MeerKAT, low-frequency ($\lesssim 1\,\mathrm{GHz}$) observations are prone to solar contamination, even for observations with solar separations of $> 10°$. While there are routines to subtract out Solar effects in visibility space (i.e., *peeling*; Samboco et al., 2024) without the need for down-weighting short baselines. In this thesis, I apply a 'quick-and-dirty' correction. For the few observations with Solar interference, I found that an inner Tukey taper mitigates solar interference and improves the rms noise; see Fig. 2.6.

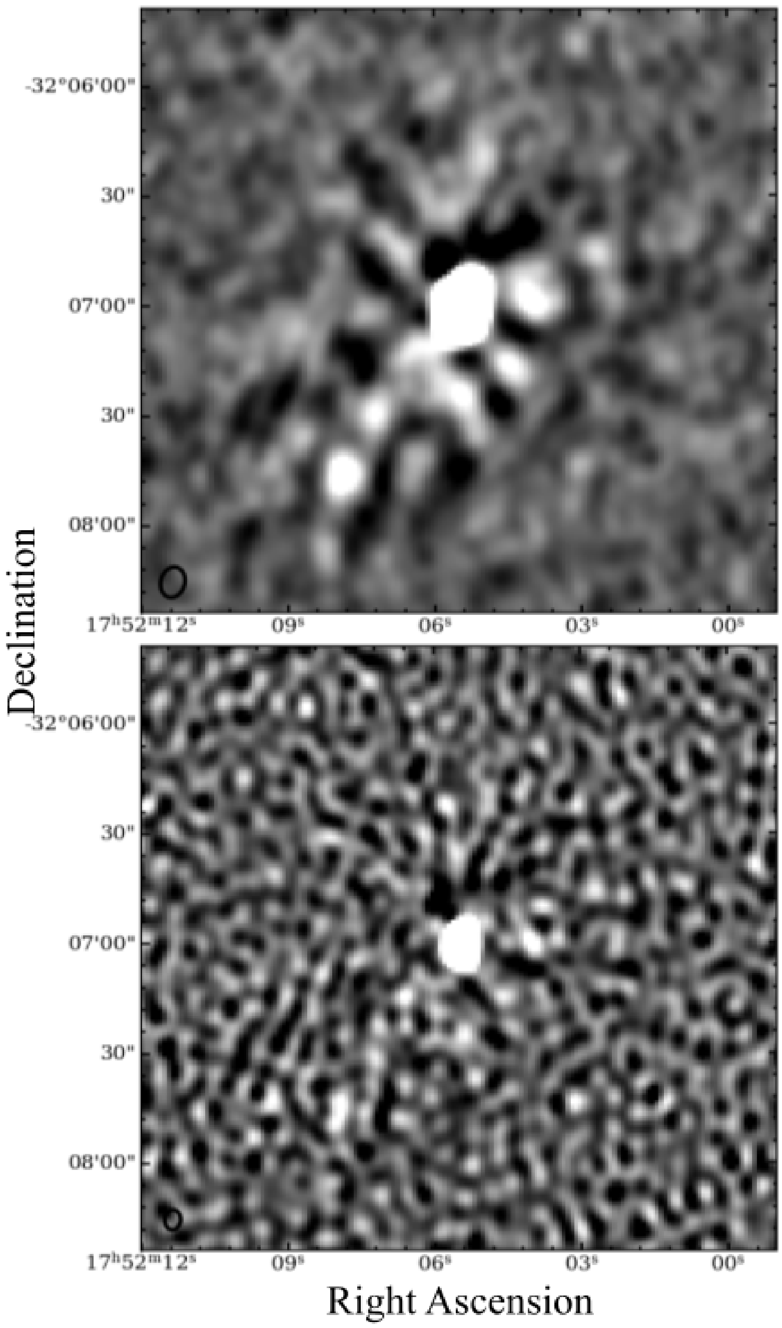

Figure 2.4: Example MeerKAT images of a point source observed with Briggs robustness of 0 (top) and uniform (bottom) weighting. For MeerKAT, the high density of short spacing baselines results in a PSF main lobe that is significantly non-Gaussian when $R > 0$, violating a fundamental assumption for CLEAN-based imaging. As a result, we exclude natural weighting from our example. Note the non-physical structure around the $R = 0$ image from residual artefacts. A uniform weighting results in (i) a smaller points-source response and, thus, better angular resolution and (ii) less pronounced imaging artefacts.

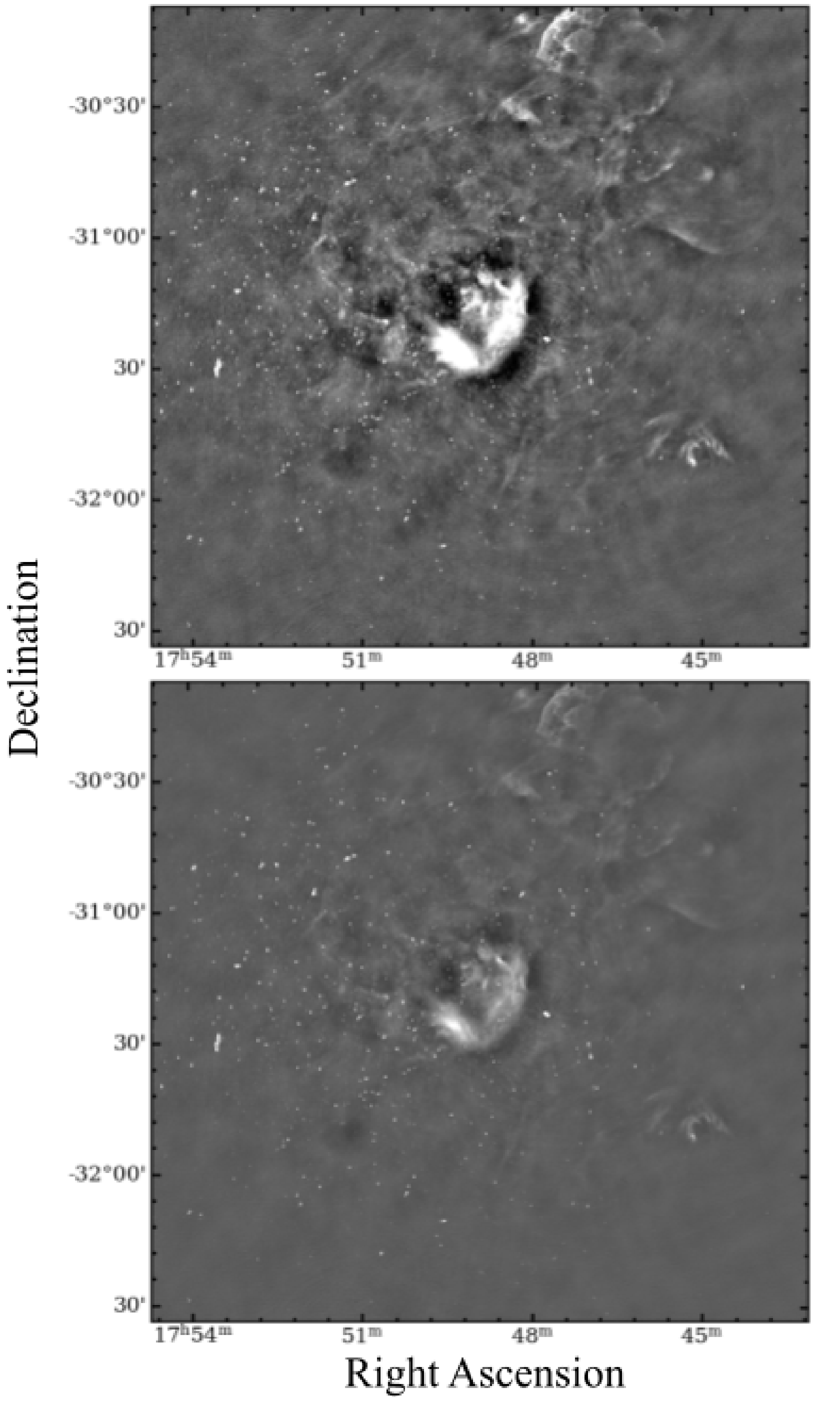

Figure 2.5: Example MeerKAT wide-field images with Briggs robustness of 0 (top) and uniform (bottom) weighting. Incomplete deconvolution of the bright supernova remnant near the centre of the field results in pronounced regions of negative intensity (the dark region around the remnant) and large 'ripples' throughout the image. The uniformly weighted image is less sensitive to extended emission and suppressed side-lobe power, resulting in a noticeable decrease in non-physical and physical extended structures.

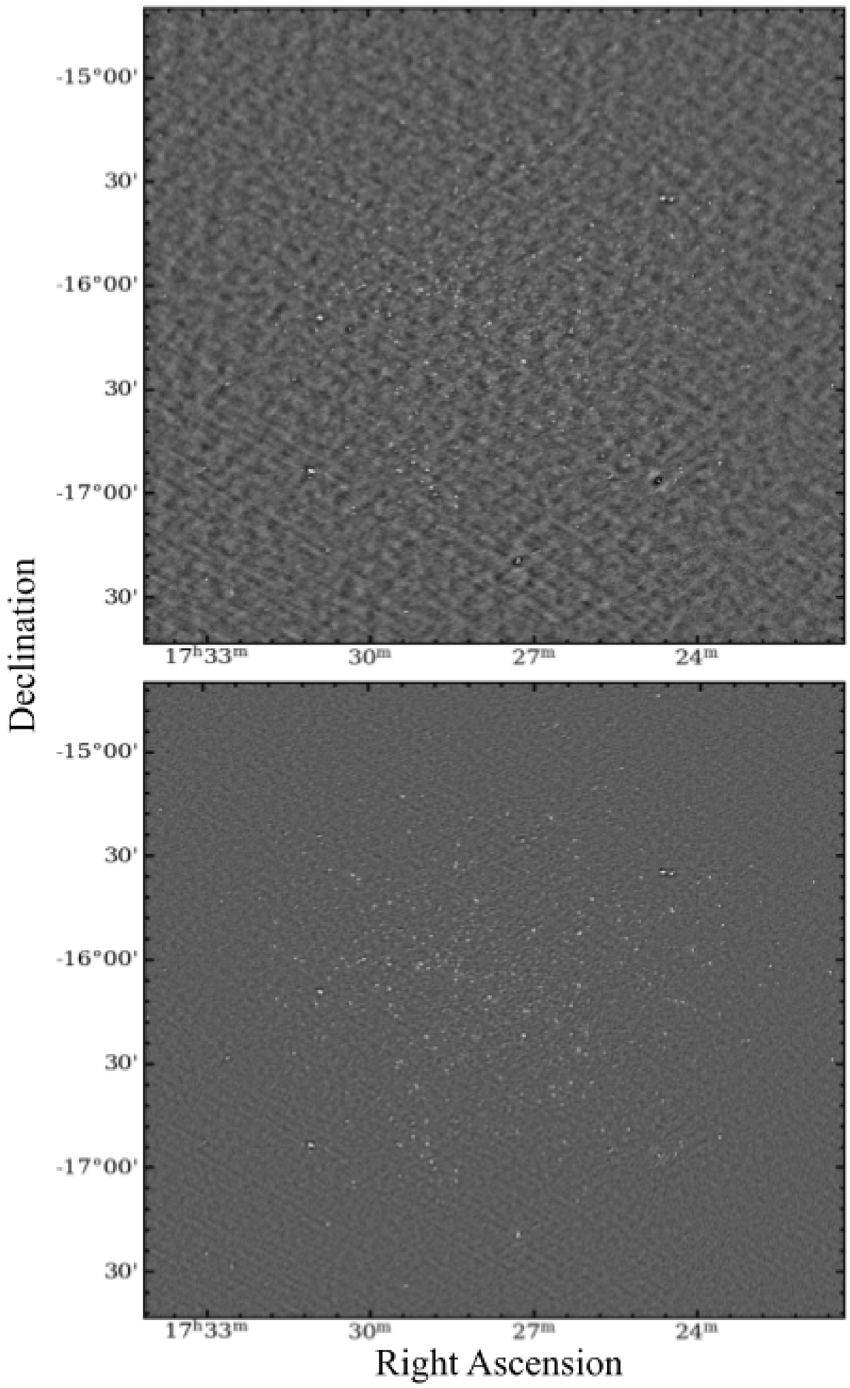

Figure 2.6: Example MeerKAT wide-field images with Briggs (robustness 0) weighting. The large-scale fluctuations (top) are caused by solar interference. Implementing a Tukey taper (bottom) effectively removes the interference effects.

Chapter 3

SAX J1810.8-2609: An Outbursting Neutron Star X-ray Binary with Persistent Spatially Coincident Radio Emission

The following chapter presents our 2021–2023 monitoring of the NSXB SAX J1810.8−2609. The published version incorrectly listed the date of the Type I X-ray burst that collaborators and I observed as 2022 August 7 (and not 2021 August 7); that mistake is corrected here. Small typographical changes that originally evaded the copy-edit process at the journal. I also made changes to maintain consistency of the variations related to the word 'ejecta' across the thesis. I also corrected a y-axis label for one of the subplots of a figure. Additionally, I include an additional line in the caption of all figures and tables following the copyright permission requirements of the Monthly Notices of the Royal Astronomical Society. The only other modifications from the published version, (Hughes et al., 2024), are the renaming of Appendixes A, B, and C as Section 3.6.1, 3.6.2, and 3.6.3. In Chapter 6, I present an updated version of the astrometry routine from Section 3.6.1.

3.1 SAX J1810.8−2609

SAX J1810.8−2609 (henceforth SAX J1810) is a NSXB that was initially discovered in 1998 by the wide-field X-ray cameras aboard the *BeppoSAX* satellite (Ubertini et al., 1998). Since its discovery, there have been four subsequent (detected) outbursts that occurred in 2007 (Degenaar et al., 2007), 2012 (Degenaar & Wijnands, 2013), 2018 (Negoro et al., 2018), and 2021 (Iwakiri et al., 2021). A Type I X-ray burst (i.e., the runaway thermonuclear detonation of a hot-dense surface layer of accreted matter, see Galloway & Keek, 2021, for a review) revealed the presence of a solid surface, identifying the accreting object as a neutron star (Natalucci et al., 2000). Furthermore, X-ray modelling of the burst showed a clear signature of photospheric radius expansion (PRE), where the burst luminosity exceeds the local Eddington limit, causing a radial expansion of the neutron star photosphere. The PRE X-ray burst was used to estimate the source distance of 4.9 ± 0.3 kpc (see, Kuulkers et al., 2003, for a review of PRE bursts as standard candles). However, we note that the quoted distance error is purely statistical, as it does not take into consideration any systematic effects, such as the potential for the neutron star to deviate from the assumed mass of $1.4M_{\odot}$ or the potential for accreting elements besides hydrogen. Therefore, the error on the distance is likely an underestimation. An analysis of multiple Type I X-ray bursts detected during the 2007 outburst showed timing signals consistent with a neutron star spin frequency of 531.8 Hz (Bilous et al., 2018). These 'millisecond burst oscillations' are thought to be caused by anisotropic X-ray emission (i.e., 'hot spots'; Watts, 2012) and allow for the determination of the neutron star spin frequency without the need for consistent pulsations.

The source has not been classified as an atoll or Z source; instead, it has adopted the broader label of neutron star 'soft X-ray transient', which encompasses both sub-classes. However, given its moderate peak X-ray luminosity ($L_X \leq 4 \times 10^{36}\,\mathrm{erg\,s^{-1}}$) and transient behaviour, it is likely to be an atoll source. The majority of Z sources are persistent and bright, with maximum X-ray luminosities reaching appreciable fractions of the Eddington limit (e.g., $L_X \sim 2 \times 10^{38}\,\mathrm{erg\,s^{-1}}$).

On 2021 May 13 (MJD 59347), the gas slit camera (GSC) aboard The Monitor of All-sky X-ray Image (i.e, MAXI; Matsuoka et al., 2009) satellite detected the X-ray brightening of SAX J1810 as it entered its fifth recorded outburst (Iwakiri et al., 2021). Following the X-ray detection, radio observations with the MeerKAT radio telescope on 2021 May 21 (MJD 59356) revealed a spatially coincident radio source, constituting the first radio detection of this source (Motta et al., 2021). Here we present our multi-instrument radio/X-ray monitoring campaign of SAX J1810. Our monitoring includes the 2021 outburst and 2023 follow-up that revealed the existence of a spatially coincident, persistent steep spectrum radio source. The remainder of this paper is structured as follows: in Section 3.2, we introduce our observation and analysis procedure, while in Sections 3.3 and 3.4, we present and discuss our results. Finally, we summarize our findings in Section 3.5.

3.2 Observations and Data Analysis

3.2.1 MeerKAT

Weekly Monitoring

We observed SAX J1810 with MeerKAT (a radio interferometer; Camilo, 2018) as a part of the large survey project ThunderKAT (Fender et al., 2016). We began a weekly monitoring campaign on 2021 May 22 (MJD 59356), nine days after the outburst's initial detection, and continued until 2021 October 23 (MJD 59508) for a total of 21 observations. Each observation consisted of a single scan of 15 minutes on-source flanked by two 2-minute scans of a nearby gain calibrator (J1830-3602). Each epoch also included a 5-minute scan of PKS B1934-638 (J1939-6342) for flux and bandpass calibration. In addition to the weekly monitoring, we observed two deep (1-hour) epochs on 2023 May 22 (MJD 60086) and 2023 August 16 (MJD 60172) when the source was in (X-ray) quiescence. The deep epochs followed the same observing strategy, except the source monitoring was broken into two 30-minute scans. All MeerKAT observations used the L-band receiver, with a central frequency of 1.3 GHz, and a total (un-flagged) bandwidth of 856MHz split evenly into 32768 frequency channels. To decrease the size of each data set, we averaged together every 32 channels (resulting in 1024 total channels) before data reduction and imaging. This averaging will not affect our final results as we are focused on radio continuum emission (as opposed to spectral lines).

We performed flagging, calibration, and imaging using a modified version of the semi-automated routine OXKAT[1] (Heywood, 2020), which breaks the process into three steps. Here we will briefly outline the workflow and direct readers to Heywood et al. (2022) for a more comprehensive description.

[1] Found at: `https://github.com/IanHeywood/oxkat`

The first step (1GC) uses CASA (v5.6; CASA Team et al., 2022) to remove data corrupted by radio frequency interference (RFI). After removing RFI, the data is corrected with standard calibration solutions (i.e., flux density, bandpass, and complex gain). The second step (FLAG) applies a second round of flagging using `tricolor` (Hugo et al., 2022) before creating a preliminary image of the source field using WSCLEAN (v2.9; Offringa et al., 2014). This preliminary image is then used to create an imaging mask. The final step (2GC) begins with a masked deconvolution before using the model image for direction-independent (DI) self-calibration with CubiCal (Kenyon et al., 2018). Following self-cal, the pipeline ends with a second round of masked deconvolution using the DI self-calibrated visibilities. We adopted the 2GC images as our final data products. We maximize our sensitivity by weighting each image with a Briggs' robustness of 0 (Briggs, 1995b)[2]. We note that OxKAT has the functionality to solve for direction-dependant (DD) self-calibration solutions if needed (i.e., the 3GC step). However, for SAX J1810, DI self-calibration was sufficient, and thus we omitted the 3GC step.

We measured the source properties in each epoch using the CASA task `imfit`, fitting an elliptical Gaussian component in a small sub-region around the source to measure the position and flux density. As the source was unresolved, we set the component shape to be the synthesized beam of each image. We quantified the (1σ) uncertainty on the flux measurement using the local root-mean-square (rms) noise. We extracted the rms from an annular region for each epoch using the CASA task `imstat`. Each annulus was centred on the position of the Gaussian component. We fixed the inner radius as the major axis of the synthesized beam and scaled the outer radius such that the annular area comprises the area of 100 synthesized beams. We quantified astrometric errors using the method detailed in Appendix 3.6.1.

[2]MeerKAT's synthesized beam becomes significantly non-Gaussian for robustness weightings >0, inhibiting accurate deconvolution and raising the image-plane rms noise.

3.2.2 Very Large Array

We were approved for a single director's discretionary time observation (Project Code: 23A–417) with the Very Large Array (VLA) as a follow-up of our initial 2023 MeerKAT observation. SAX J1810 was observed on 2023 July 17 (MJD 60142) in the 2–4 GHz (S-band) and 4–8 GHz bands (C-band). For S-band, the observations used the 8-bit sampler comprised of two base-bands, with eight spectral windows of sixty-four 2 MHz channels each, giving a total (unflagged) bandwidth of 2.048 GHz. The 3-bit sampler was used for C-band, which has four base-bands, and thus a 4.096 GHz bandwidth. In each band, we included a single 1-minute scan of the flux calibrator (3C286). For source monitoring the array cycled between SAX J1810, observed for ~ 8 minutes per cycle in S-band and ~ 5 minutes in C-band. Each source scan is flanked by ~ 1 minute observations of a nearby gain calibrator (J1820$-$2528). The total time on source was ~ 16 minutes in both bands. We performed flagging, calibration, and imaging using the most recent release of the CASA VLA pipeline (v6.4). We imaged the source using WSCLEAN but did not detect the source in either band. As a result, we extract the rms noise from each image to place (3σ) upper limits on the flux density. We used a circular extraction region (with an area equal to 100 synthesized beams) centred on the archival position of SAX J1810 to measure the rms. The radio flux densities from both MeerKAT and the VLA are presented in Table 3.4

3.2.3 *Swift*-XRT

Weekly Monitoring

We monitored SAX J1810 with the X-ray telescope (XRT; Burrows et al., 2005) aboard the Neil Gehrels *Swift* Observatory (Gehrels et al., 2004), capturing the quasi-simultaneous evolution of the X-ray flux (i.e., within

~ 3 days of a MeerKAT observation). During the outburst, we observed 21 epochs (target ID: 32459) between 2021 May 20 (MJD 59364) and 2021 November 6 (MJD 59524) at an approximately weekly cadence. To accompany our deep MeerKAT epochs, we were approved for two Target-of-Opportunity observations on 2023 May 25 (MJD 60089) and 2023 August 16 (MJD 60172). During the initial stages of the outburst, we monitored the source in Windowed Timing (WT) mode, where SAX J1810 exhibited a maximum count rate of $\sim 20\,\mathrm{count\,s^{-1}}$ during the first epoch. We transitioned to Photon Counting (PC) mode when the sources count rate decayed to $\lesssim 1$ count $\mathrm{s^{-1}}$ on 2021 October 9 (MJD 59496), although there was a single intermittent PC epoch on 2021 September 5 (MJD 59462).

We used the Python API version of the *Swift*-XRT pipeline, `swifttools` (Evans et al., 2007, 2009), to extract the source and background spectra for all epochs except 2021 August 7 (MJD 59433), where the source exhibited a Type I X-ray burst (see Section 3.2.3). We used the HEASOFT package (version 6.25) for our spectral analysis. For observations that had a sufficiently large number of counts (i.e., MJD 59364–59496), we used a modified `grppha` script to bin the spectra on 25-count intervals and performed spectral fitting using χ^2 statistics. Towards the end of our 2021 monitoring (i.e., the MJD 59504 and 59511), we used Cash statistics (i.e., `cstat`; Cash, 1979) with single-count binning intervals, due to the small number of counts collected in each observation. The final two epochs of the 2021 monitoring (MJD 59518 and 59524) and the late-time follow-up (MJD 60089 and 60172) were non-detections and thus were omitted from the spectral fitting routine.

Using XSPEC (Arnaud, 1996a), we performed our spectral fitting twice, once for the 0.5–10 keV energy range and again for 1–10 keV. As expected, changing the energy range had a negligible effect on the best-fit spectral parameters. We modelled the spectra using an absorbed power-law model with an added blackbody component; i.e., `tbabs` $\times$ (`pegpwrlw` + `bbody`), where `tbabs` models the interstellar absorption using an equivalent hydrogen column density (N_H) following the abundances from Wilms et al. (2000). The power-law accounts for the X-ray emission from the dominant component (i.e., the hard X-ray corona), and the blackbody accounts for any excess soft X-ray emission from a faint accretion disk, neutron star surface, or boundary layer. Initially, we fit each spectrum individually, allowing N_H to vary epoch by epoch. We then adopted the single epoch fitting as our starting parameters, linking the N_H values across all epochs and fitting the spectra simultaneously, resulting in a single time-independent value of N_H. When calculating the degrees of freedom, we treated the linked N_H as frozen (i.e., each spectrum has four free parameters). The epochs that utilized Cash statistics were omitted from the fitting procedure detailed above. Instead, we fit each of those spectra with a simple absorbed power-law model (i.e., `tbabs` $\times$ `pegpwrlw`), fixing N_H to our best-fit value of $3.88 \times 10^{21}\,\mathrm{cm}^{-2}$ and the power-law photon index (Γ) to the average value of 1.61 from the χ^2 fitting. As a result, the X-ray flux was the only free parameter in the Cash statistic modelling. The *Swift*-XRT monitoring and spectral parameters during the 2021 outburst are presented in Table 3.5. The quoted uncertainties on the X-ray parameters represent the standard 90% confidence intervals.

Type I X-ray Burst

On 2021 August 7 (MJD 59433), SAX J1810 underwent a Type I X-ray burst, and, as a result, we performed manual data reduction on the *Swift*-XRT (WT) observations. First, we ran the task `xrtpipeline` to produce cleaned event files and exposure maps. Second, using `barycorr`, we applied the barycentric timing correction. Lastly, we extracted source and background spectra using `xselect`. For the pre-burst times, we used a circular source extraction region with a radius of 30 pixels (1 pixel = 2.36 arcsec) and an annular background extraction region with an inner radius of 70 pixels and an outer radius of 130 pixels. The pre-burst spectrum was then processed using χ^2 statistics and the routine mentioned in 3.2.3.

During the burst, we broke the event file into multiple time bins to analyze the time evolution of the spectral parameters. Due to high count rates during the burst (i.e., maximum count rates $\gtrsim 400\,\mathrm{count\,s^{-1}}$), the observations are affected by systematic effects caused by photon pile-up. As a result, we used an annular source extraction region with an inner (exclusionary) radius that increases with an increasing count rate (ranging from 0 to 3 pixels). Following the *Swift*-XRT pipeline procedure (see, Evans et al., 2007, 2009), we choose inner radii that reduce the maximum count rate in a given time bin to $< 150\,\mathrm{count\,s^{-1}}$. The time ranges were chosen so each bin has $\gtrsim 300$ counts corresponding to 21 bins across the 1.5 minute burst. To model the burst parameters in XSPEC, we added a second blackbody component to the pre-burst spectrum, fixing the pre-burst parameters and thereby allowing only the second blackbody to vary. We used the `bbodyrad` model to fit for the normalization (which is related to the radius of the blackbody) and temperature before using the XSPEC convolution model `cflux` to calculate the flux.

For the timing analysis, we extracted two light curves. The first light curve was binned on 1 s intervals and was used to model the decay timescales of the burst. We extracted an initial light curve using the circular extraction region. For any time bins with a count rate $> 150\,\mathrm{count\,s^{-1}}$, we replaced their count rates with the count rate measured by the annular region with a 3-pixel exclusionary inner radius. We corrected for background and annular extraction region effects with `lcmath` and `xrtlccorr`, respectively. Following the prescription outlined in (Galloway et al., 2020) we fit an exponential decay function,

$$R(t) = Ae^{-\frac{t}{\tau}} + R_0, \tag{3.1}$$

where t is the time after the burst maximum, $R(t)$ is the count rate at a given t, R_0 is the constant background rate, τ is the e-folding decay time, and A is the peak count rate of the bursting component (excluding the contribution from a constant background). We fit for τ, R_0, and A with a Markov-Chain Monte Carlo (MCMC) routine using Python's EMCEE package (Goodman & Weare, 2010; Foreman-Mackey et al., 2013), assuming the sampled count rates were independently distributed normal random variables. The number of (sampling) walkers was fixed at five times the number of dimensions (i.e., 15). We chose three flat priors to ensure an unbiased analysis. To ensure convergence, we manually inspected the walkers over many autocorrelation times. Additionally, we analyzed the evolution of the autocorrelation time as a function of the number of MCMC steps following the routine outlined in the EMCEE documentation[3].

The second light curve was extracted using the circular extraction region and binned on 1.8 ms intervals (the minimum bin size possible for WT mode). We used the short timescale light curve to search for millisecond burst oscillations. Given the short timescale binning, no corrections were applied to the 1.8 ms light curves. Appendix 3.6.2 presents the X-ray burst properties.

3.2.4 The WATCHDOG Pipeline

We calculated the X-ray hardness ratio (HR) using a modified version of the pipeline developed for the Whole-sky Alberta Time-resolved Comprehensive black hole Database Of the Galaxy (WATCHDOG; see Tetarenko et al., 2016b, for a comprehensive description of the pipeline). The hardness ratio is the ratio between the number of counts in the hard and soft X-ray bands. We used the MAXI/GSC 4–10 keV band as the soft band and 15–50 keV

[3]The documentation can be found here: `https://emcee.readthedocs.io/en/stable/tutorials/autocorr/`

observations from the Burst Alert Telescope (BAT; Barthelmy et al., 2005) aboard *Swift* as the hard band. Both sets of observations are publicly available[4]. We modified the pipeline to average daily observations, ensuring the hard X-ray band had a $\geq 3\sigma$ detection. For data where the soft X-ray band detection significance was $< 3\sigma$, we replaced the measured count rate with $3\times$ the noise value to estimate a conservative 3σ lower limit. The source appears to have undergone a hard-only outburst, and, as a result, to get meaningful constraints, we needed to measure either a lower limit or detection on the hardness ratio. No further modifications were applied to the WATCHDOG pipeline.

WATCHDOG defined empirical HR limits that corresponded to the different X-ray states: (i) $C_{\rm hard} = 0.3204$; and (ii) $C_{\rm soft} = 0.2846$. A hardness ratio is considered consistent with the hard (soft) state if its lower (upper) error bars are above (below) the $C_{\rm hard}$ ($C_{\rm soft}$) limits. If neither criterion is met, the source is classified as being in an intermediate state. We note that the values of $C_{\rm hard}/C_{\rm soft}$ were calculated for BHXBs; in Section 3.4.1, we investigate whether it is valid to apply the same standard NSXBs.

3.3 Results

3.3.1 Radio Position

In Fig. 3.1, we show the offset in right ascension and declination between the MeerKAT position and the archival X-ray position of 18h10m44.47s $-26°09'01.2''$ from (Jonker et al., 2004). The average radio position is 18h10m44.34s $-26°09'02.1''$ ($\pm 0.1''$). The per-epoch declinations are consistent with the average radio position with a reduced $\chi^2 = 0.75$ (22 degrees of freedom), although the average radio position is offset by $\sim 1''$ from the

[4]MAXI/GSC: `http://maxi.riken.jp`
Swift-BAT: `https://swift.gsfc.nasa.gov/results/transients/`

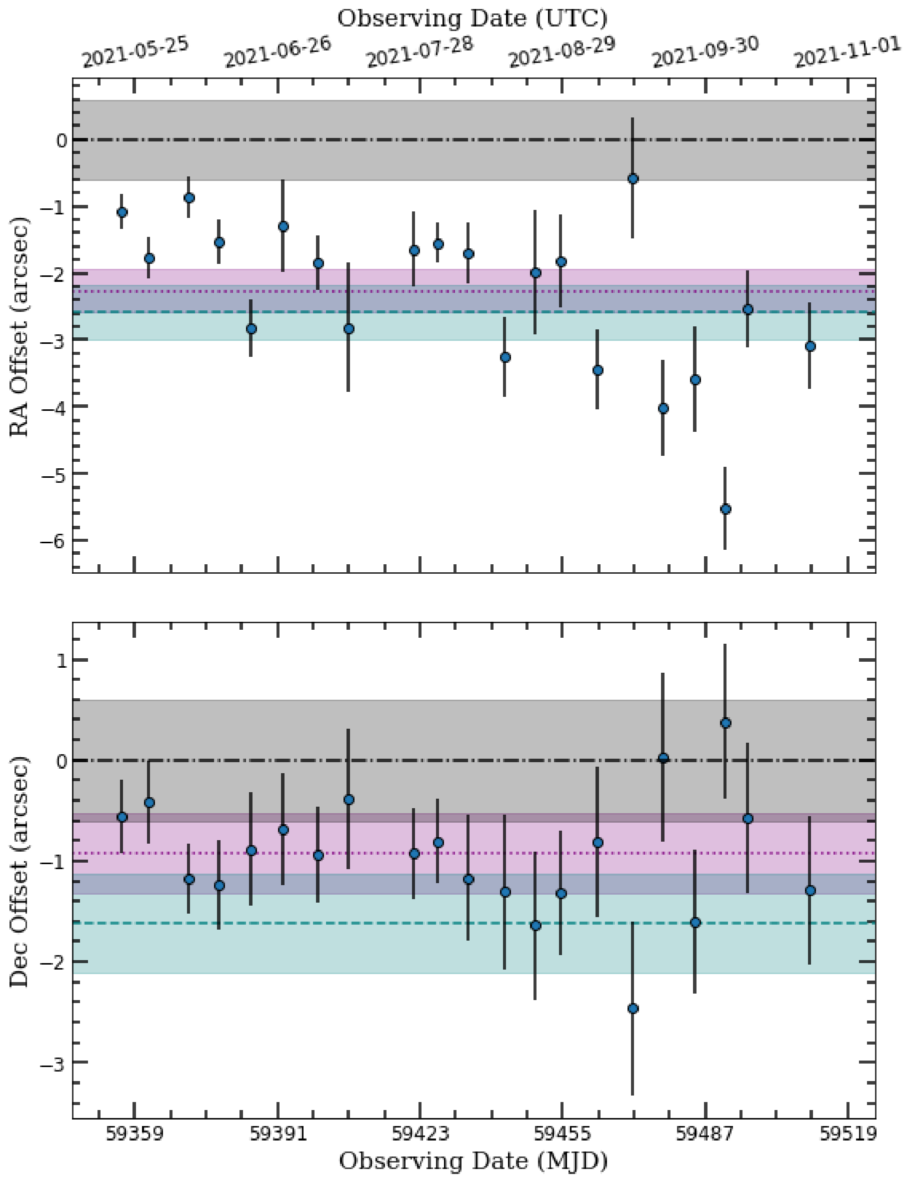

Figure 3.1: The right ascension (top panel) and declination (bottom panel) offsets for the best-fit SAX J1810 positions. The filled blue circles are the offsets of the source. The purple dotted line and cyan dashed line are the 2023 May 22 and 2023 August 13 offsets, respectively. The dashed-dotted black line is the archival X-ray position from Jonker et al. (2004), and the grey shaded area is the error on the archival position ($\pm 0.6''$). Note the clear offset and temporal variability in the right ascension of the source. Originally Figure 1 in (Hughes et al., 2024).

X-ray position. In contrast, the right ascensions show significantly larger offsets ranging from ~ 1–$5''$. Moreover, the measured right ascensions show temporal variability. Adopting the weighted mean offset in right ascension as a model and computing the reduced χ^2 results in a value of $\chi^2 = 4.4$ (22 degrees of freedom), suggesting that the variability is not the result of stochastic error fluctuations. We tested the right ascension offsets against a linearly increasing model (i.e., ballistic motion), which resulted in a negligible improvement in the reduced χ^2 (4.2; 21 degrees of freedom), and thus, we found no evidence of ballistic motion.

3.3.2 Outburst Light Curves

In Fig. 3.2 we show the MeerKAT (1.3 GHz; top panel), *Swift*-XRT (0.5-10 keV; second panel), MAXI/GSC (4-10 keV; third panel), and *Swift*-BAT (15-50 keV; bottom panel) outburst light curves. For our MeerKAT observations, 18 (out of 21) epochs were $\geq 5\sigma$ detections (blue circles). The remaining three epochs (blue diamonds) do not meet the typical reporting threshold of 5σ, with detection significance of ~ 4.3–4.9σ. Given the spatial coincidences between the low ($< 5\sigma$) and high-significance detections ($\geq 5\sigma$), it is likely that we are detecting a source in all of our MeerKAT observations. For the *Swift*-XRT light curve, we adopted the total fluxes from our spectral fits using the joint power-law and blackbody model components (filled black circles). The last two data points (open black circles) correspond to the epochs where the source was too faint for multi-component spectral modelling; instead, we fit the source with a single power-law component. The *Swift*-BAT and MAXI/GSC light curves display the data at a daily binning frequency.

The observed flux of SAX J1810 displays a common temporal evolution across all observing frequencies. At early times ($\sim$ MJD 59340–59370), all four instruments recorded the brightest signal of the outburst. Following the

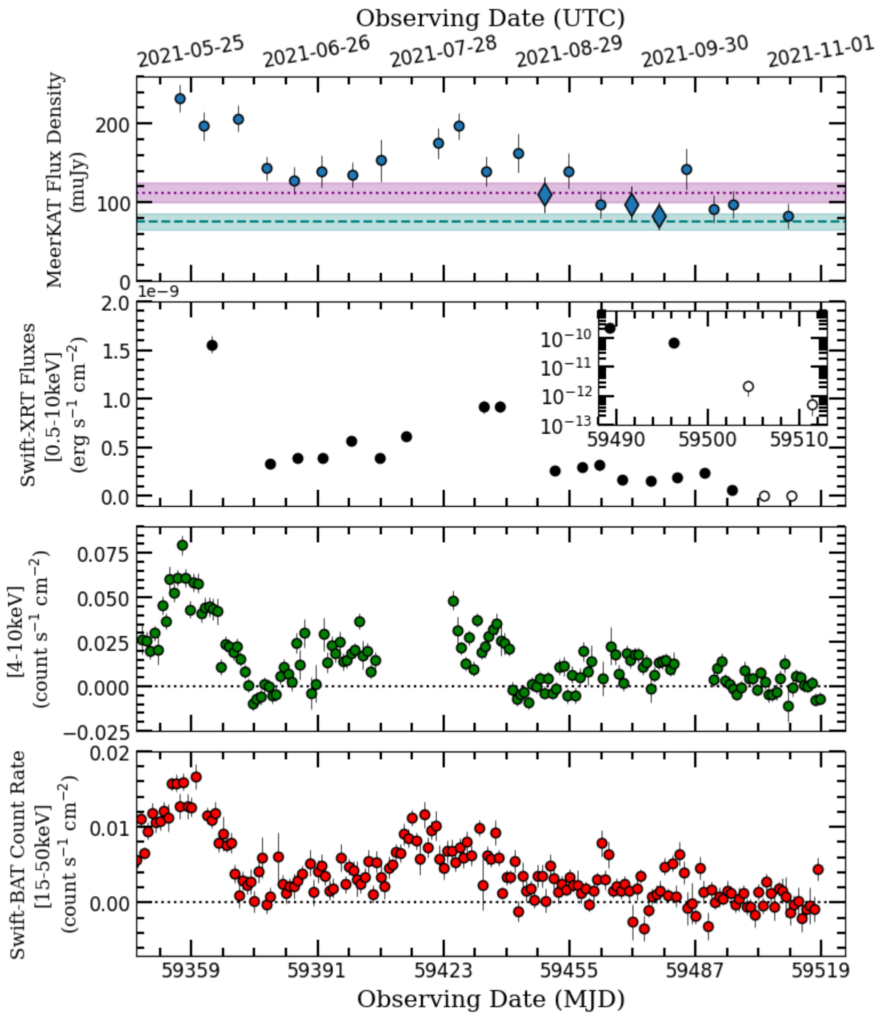

Figure 3.2: Multi-instrument light curves of the 2021 outburst of SAX J1810. The top panel is the MeerKAT 1.3 GHz radio light curves showing both $\geq 5\sigma$ (blue circles) and 4–5σ (blue diamonds) detections during the 2021 outburst. The horizontal lines show the 2023 May 22 (purple dotted) and 2023 August 13 (cyan dashed) flux densities. The second panel is the *Swift*-XRT (0.5–10.0 keV) light curves. The filled and open circles correspond to the epochs fit with χ^2 and Cash statistics, respectively. The bottom two panels show the MAXI/GSC (third panel) and *Swift*-BAT (bottom panel) daily-binned light curves. All four instruments show a common temporal evolution characteristic of the correlation between radio and X-ray emission in the hard accretion state. Originally Figure 2 in (Hughes et al., 2024).

maxima, the source flux began decreasing, showing a rebrightening between $\sim$ MJD 59410 and 59440, before the source flux continued to decrease, returning to X-ray quiescence and plateauing at $\sim 90\,\mu$Jy in the radio. We find no evidence for additional intra-observation variability beyond the Type I outburst discussed in this paper.

Although the radio and X-ray light curves share a similar evolution in time, the magnitude of the variability is significantly different. In radio, the source exhibits modest variability with a maximum ($\sim 230\,\mu$Jy) and minimum ($\sim 80\,\mu$Jy) flux density separated by a factor of only ~ 3. In contrast, when only considering the epochs with multi-component spectral modelling, the *Swift*-XRT fluxes show a factor of ~ 20 in variability, with a maximum and minimum flux of $\sim 1.6 \times 10^{-9}$ and $6.8 \times 10^{-11}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$, respectively. Including the final two *Swift*-XRT epochs during the source's return to quiescence, the minimum flux is $\sim 5 \times 10^{-13}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$, which corresponds to a factor of ~ 2000 decrease from the maximum. The plateauing radio emission at MJD 59463 (and beyond) is consistent with a spatially coincident, persistent radio source (see Section 3.4.2).

3.3.3 X-ray Spectra

The X-ray modelling parameters are shown in Fig. 3.3. The best fit equivalent hydrogen column density is $N_H = 3.9^{+0.1}_{-0.2} \times 10^{21}\,\mathrm{cm^{-2}}$. The Colden: Galactic Neutral Hydrogen Density Calculator[5] estimates a value of $N_H \sim (3.2–4.3) \times 10^{21}\,\mathrm{cm^{-2}}$ along the SAX J1810 line of sight (depending on the choice of neutral hydrogen data set — NRAO or Bell), making the measured N_H consistent with expectation.

To investigate the relative contributions of each model component, we calculated the *power-law flux fraction* (third panel, Fig. 3.3); i.e., $F_{X,\mathrm{PL}}/F_{X,\mathrm{tot}}$, where $F_{X,\mathrm{PL}}$ is the X-ray flux of the power-law component and $F_{X,\mathrm{tot}}$ is the

[5] The webtool can be found here: `https://cxc.harvard.edu/toolkit/colden.jsp`

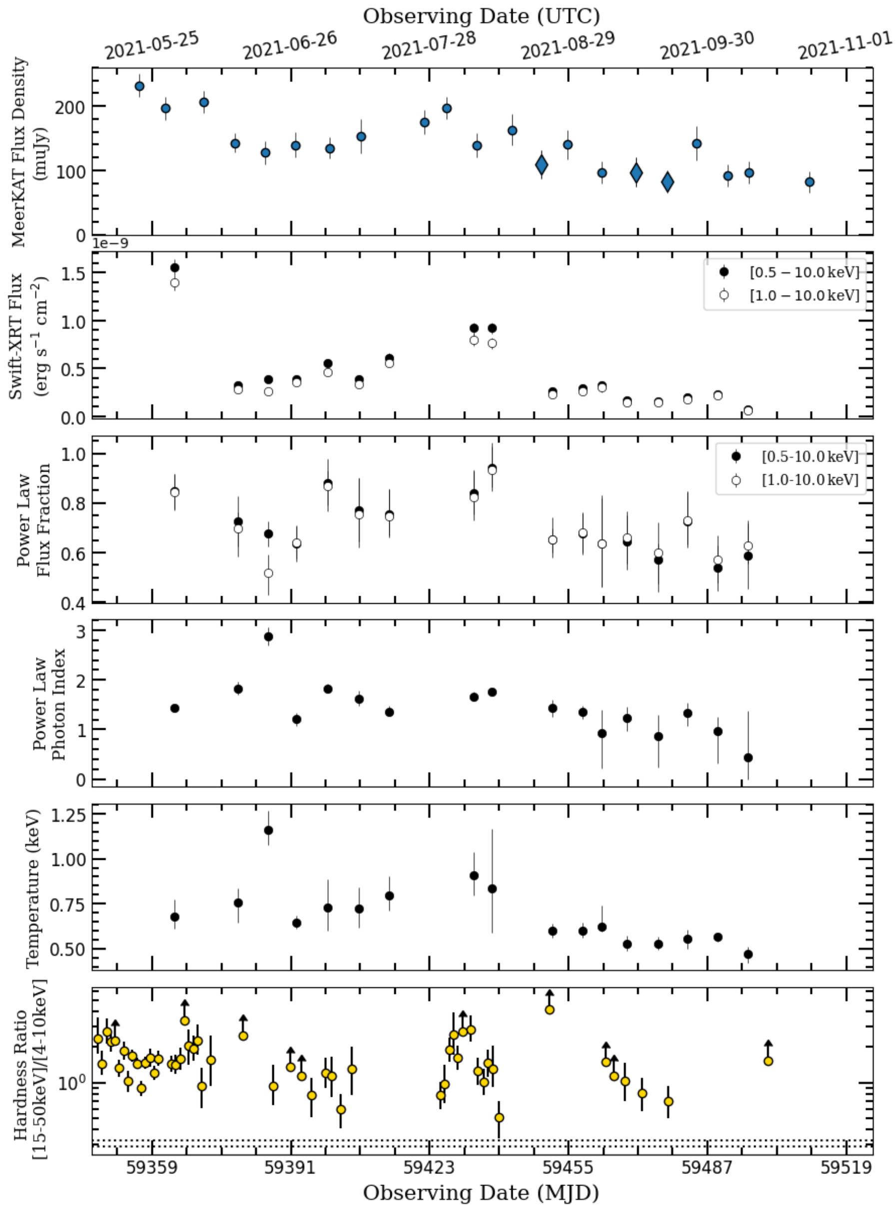

Figure 3.3: A summary of the spectral properties of SAX J1810. The top panel shows the MeerKAT radio flux density. The next two panels show the *Swift*-XRT X-ray flux (second panel) and the power-law flux fraction (third panel) in the 0.5–10.0 keV (filled circles) and 1.0–10.0 keV (open circles) energy bands. The fourth panel shows the power-law photon index, and the fifth shows the temperature of the black body component. The bottom panel shows the hardness ratio between the MAXI/GSC (4.0–10.0 keV) and *Swift*-BAT (15.0–50.0 keV) energy bands. The upper (C_{hard}) and lower (C_{soft}) dotted lines show the empirically defined state boundaries from WATCHDOG (Tetarenko et al., 2016b). These spectral properties are characteristic of the hard accretion state. Originally Figure 3 in (Hughes et al., 2024).

Table 3.1: Three component fit of the *Swift*-XRT observation on MJD 59385

model Component	Γ	kT (keV)	F_X ($10^{-11}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$)
`pegpwrlw`	$1.8^{+0.7}_{-1.5}$	—	18^{+7}_{-10}
`bbody`	—	$0.9^{+0.2}_{-0.1}$	9^{+2}_{-2}
`diskbb`	—	$0.22^{+0.03}_{-0.03}$	12^{+2}_{-3}

Notes. I fixed the best fit value for the absorption column density to $N_H = 3.88 \times 10^{21}\,\mathrm{cm^{-2}}$, and left all other parameters free. After including the `diskbb` component, the `pegpwrlw` becomes the subdominant component, and the fit becomes insensitive to both the flux and the photon index of the power-law component. SAX J1810 may have briefly entered a thermal X-ray-dominated accretion state before returning to the hard state. Originally Table 1 in (Hughes et al., 2024).

total X-ray flux of the model. In all epochs, the power-law component is dominant with a flux fraction ranging from ~ 0.53 to 0.94 with a (variance-weighted) average of 0.72 ± 0.02. The power-law photon index (Γ; fourth panel, Fig. 3.3) shows moderate variability with $0.4^{+0.93}_{-0.43} \leq \Gamma \leq 2.88^{+0.18}_{-0.08}$ and an average value of 1.61 ± 0.03. The average value is typical of comptonized hard state X-ray emission from (black hole) X-ray binaries (Remillard & McClintock, 2006). Moreover, if we exclude the anomalously steep photon index, the maximum photon index becomes $\Gamma = 1.83^{+0.10}_{-0.08}$. The blackbody temperature (kT; third panel, Fig. 3.3) varied between $0.5^{+0.18}_{-0.08} \leq kT \leq 1.2^{+0.18}_{-0.08}$ keV, with an average blackbody temperate of $kT = 0.60 \pm 0.01$ keV. Black body temperatures $\lesssim 1$ keV are consistent with past analyses of hard state neutron star X-ray binaries (e.g., Lin et al., 2007). The bottom panel of Fig. 3.3 displays the hardness ratio calculated from the daily *Swift*-BAT and MAXI/GSC light curves. We observe a moderate degree of variability in hardness ratio, with detections ranging from $\sim$0.5–2.8, and an average value of 1.19 ± 0.06. Including the lower limits increases the maximum hardness ratio to ~ 4.

The largest single epoch evolution occurs on MJD 59385, where the black body temperature reaches its maximum value of ~ 1.2 keV, alongside the

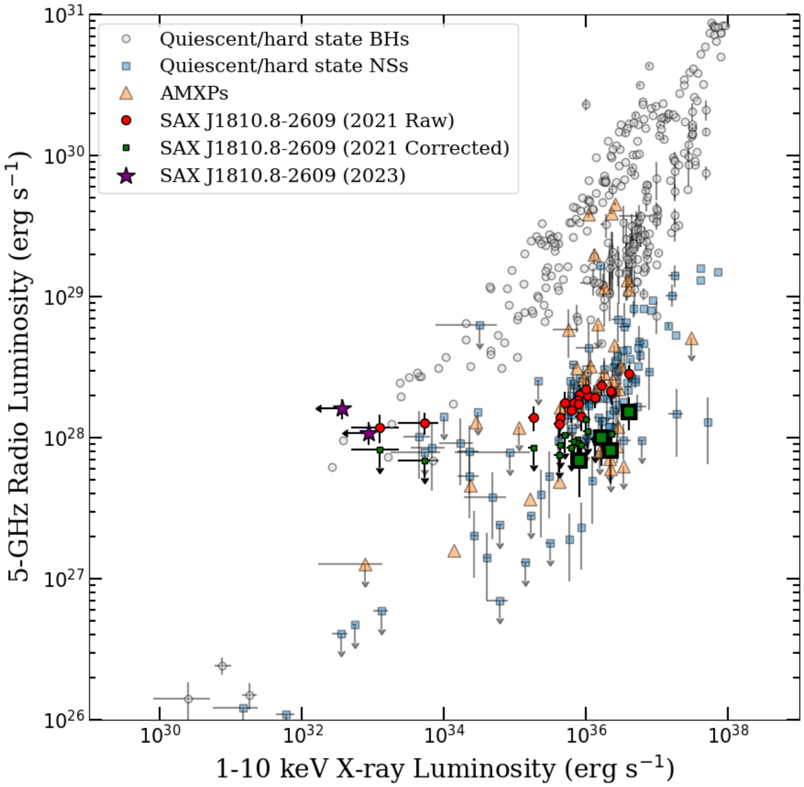

Figure 3.4: The radio (5 GHz) and hard state X-ray (1–10 keV) luminosity (L_R-L_X) relation. The archival values for black holes (grey circles), neutron stars (blue squares), and accreting millisecond X-ray pulsars (orange triangles) were taken from van den Eijnden et al. (2022), which is based on the Bahramian & Rushton (2022) catalog. The luminosities for SAX J1810 during the 2021 outburst are represented as red circles. The purple stars are the 2023 values if we assume that all radio emission originates from a hard state jet. The late-time plateau in the radio flux density strongly suggested the existence of a second radio source is uncorrelated with the X-ray emission. The green squares show the 2021 outburst values after subtracting 89μJy from each radio flux density (i.e., the average contribution from the persistent source). Even after subtracting off the persistent source, the two 3σ radio detections (large green squares) of SAX J1810 remain consistent with the general population of hard-state NSXB jets. Originally Figure 4 in (Hughes et al., 2024).

extreme softening of the power-law component ($\Gamma \sim 2.9$). During this epoch, the two-component fit had a reduced χ^2 value of ~ 1.17 (216.5/186). To investigate whether we were observing a transition to an intermediate or soft state, we added a multi-colour disk to the two-component model; i.e., `tbabs` $\times$ (`pegpwrlw` + `bbody` + `diskbb`). The inclusion of the third component moderately reduces the χ^2 to ~ 1.12 (206.1/184) and decreases both the power-law photon index and blackbody temperature to levels consistent with the other epochs (See Table 3.1 for the full model parameters). Moreover, the power-law component becomes sub-dominant, suggesting that the source may have briefly transitioned into an intermediate or soft state. The observations on MJD 59413 and 59462 show similarly large reduced χ^2 values of ~ 1.22 (237/194) and ~ 1.52 (50/33), respectively. As a result, we attempted to fit these spectra with the same three-component model. However, the fitting resulted in a negligible improvement of the χ^2 statistic. We note that, for the latter epochs, both have reduced χ^2 deviations that are consistent (at the $< 3\sigma$ level) with the expected value of 1. Therefore, the poor fits may result from statistical effects rather than a physical change in the X-ray spectrum.

3.3.4 Persistent Emission and the L_R–L_X relation

Our 2023 follow-up MeerKAT observations revealed a $112 \pm 12\,\mu$Jy radio (point) source on 2023 May 22 (MJD 60086) and another $75 \pm 11\,\mu$Jy radio source three months later on 2023 August 13 (MJD 60169). The best-fit positions of both 2023 detections are consistent with the 2021 outburst (see Fig. 3.1). Therefore, we confidently detect a persistent radio source spatially coincident with SAX J1810. We calculated an (intra-band) spectral index of the persistent source using the brighter of the two MeerKAT follow-up observations (MJD 60086). We broke our observations into four evenly spaced sub-bands, ensuring a $\geq 5\sigma$ detection in each sub-band. Applying a simple linear least squares fit, we measured a spectral index of $\alpha = -0.7 \pm 0.5$. In addition to the large statistical error, we note that intra-band spectral indices are known to bias towards flatness ($\alpha \sim 0$) at detection significances $\lesssim 35\sigma$ (Heywood et al., 2016). Given our source was only detected at $\sim 10\sigma$ and the relatively large error bar, we do not apply any strong physical inference based on this intra-band spectral index.

During the last seven epochs of 2021 monitoring (MJD 59463 to 59511) – after the radio flux density had plateaued – the average radio flux density is $93 \pm 7\mu$Jy. This value is consistent with our 2023 observations (at the $\sim 2\sigma$ level), suggesting the persistent emission is, at most, weakly variable with a $\sim 20\%$ excess variance. Combining the late-time 2021 and 2023 observations results in a (weighted) average flux density of $89 \pm 5\mu$Jy. The quasi-simultaneous *Swift*-XRT follow-up on MJD 60089 and 60172 did not detect any spatially coincident X-ray source in either epoch setting 3σ upper limits on the 1–10 keV X-ray flux of $< 1.3 \times 10^{-13}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$ and $< 3.0 \times 10^{-13}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$, respectively. Furthermore, our scheduled VLA follow-up at 3 GHz and 6 GHz, taken between our two MeerKAT observations on 2023 July 17 (MJD 60142), did not detect the source. The 3σ upper limits on the 3 GHz and 6 GHz were 30 μJy and 18 μJy, respectively. Adopting a 1.3 GHz flux density of 78 μJy (conservatively assuming a 3σ drop in flux caused by intrinsic variability), we use the 3 GHz non-detection to calculate a conservative upper limit of $\alpha < -1.1$.

Figure 3.4 presents the L_R–L_X relation. The plot includes archival hard state BHXBs (grey circles), hard state NSXBs (blue squares), and accreting millisecond X-ray pulsars (AMXPs; orange triangles). The archival sources were adapted from Fig. 4 of van den Eijnden et al. (2022), an updated version of the Bahramian & Rushton (2022) catalogue. As our *Swift*-XRT and MeerKAT observations were quasi-simultaneous, we applied a one-dimensional linear interpolation to map the radio observations onto the

X-ray times for our 2021 observations. We did not apply any interpolation for our 2023 follow-up observations. Instead, we grouped the MeerKAT observations with the nearest *Swift*-XRT follow-up. We present the L_R–L_X relation from the 2021 outburst as red circles. Fitting the 2021 results with a simple power-law results in a shallow exponent of $\beta = 0.09 \pm 0.03$ (for $L_X \propto L_R^{\beta}$). If we assume that the 2023 MeerKAT detections originate from a persistent hard state jet (purple stars on Fig. 3.4) and thus should follow the L_R–L_X relation, the measured power index becomes an upper limit (due to the X-ray non-detections) adopting a value of $\beta < 0.06$. Given that our results strongly suggest the existence of a persistent radio source that is unrelated to the hard state jet of SAX J1810, we present a secondary set of L_R–L_X data points (green squares) after subtracting off $93\,\mu$Jy from each of the radio flux densities from our 2021 outburst. Post-subtraction, there are only four epochs (MJD 59364, 59378, 59413, and 59437) that show a $> 3\sigma$ excess flux density when compared to the persistent level. For the rest of the epochs, we set the radio flux density to be $3\times$ the rms noise and displayed them as upper limits. The subtracted values are unconstraining but consistent with the broader population of NSXBs. The implications of SAX J1810 L_R–L_X evolution and the origin of the persistent radio source are discussed in Section 3.4.2.

3.4 Discussion

We monitored the NSXB SAX J1810 during its 2021 outburst. The X-ray and radio properties suggest that the source underwent a 'hard-only' outburst, never fully transitioning to a soft accretion state. Moreover, the late-time plateau of radio flux density in 2021, combined with our follow-up in 2023, suggests the existence of a persistent radio source. In the following subsections, we present the evidence of a 'hard-only' outburst and discuss the possible origins of the persistent radio emission.

3.4.1 Hard-Only Outburst

Our observations suggest that SAX J1810 exhibited a 'hard-only' outburst in 2021. We justify this claim with three points of evidence:

1. The hardness ratio between the *Swift*-BAT and MAXI/GSC observations is above the hard state limit throughout the monitoring. Although the limit was empirically defined using outbursting BHXBs, we expect that the persistent source of thermal X-ray photons (from the neutron star surface or boundary layer) would make all X-ray states softer, thereby decreasing the hard state limit for NSXBs. We investigate this proposition by analyzing the best-studied outbursting (atoll) NSXB, Aql X-1. In Fig. 3.5, we have plotted a sample light curve of Aql X-1 during its 2016 outburst. The source exhibits a rapid transition of its hardness ratio, with a large fraction of the outburst remaining at a steady value of ~ 0.05 well below the soft state limit derived for BHXBs. Díaz Trigo et al. (2018) performed an X-ray spectral analysis of four separate observations; the authors identified that the source was in the hard accretion state on 2016 Aug 3 (MJD 57603) and 2016 Sep 19 (MJD 57650) and in the soft accretion on 2016 Aug 5 (MJD 57605) and 2016 Aug 7 (MJD 57607). The hard and soft state epochs are shown with the dashed and dashed-dotted lines in Fig. 3.5. As expected, the soft and hard state epochs are temporally consistent with small and large hardness ratios. The final (Sep 19) hard state epoch shows a hardness ratio below the BHXB hard state limit, consistent with our prediction that the thermal photons from neutron stars will lower the hard state limits. We note that other outbursts of Aql X-1 (e.g., the 2009 outburst; Miller-Jones et al., 2010) show a similar 'softening' of the hard state limit. Therefore, we

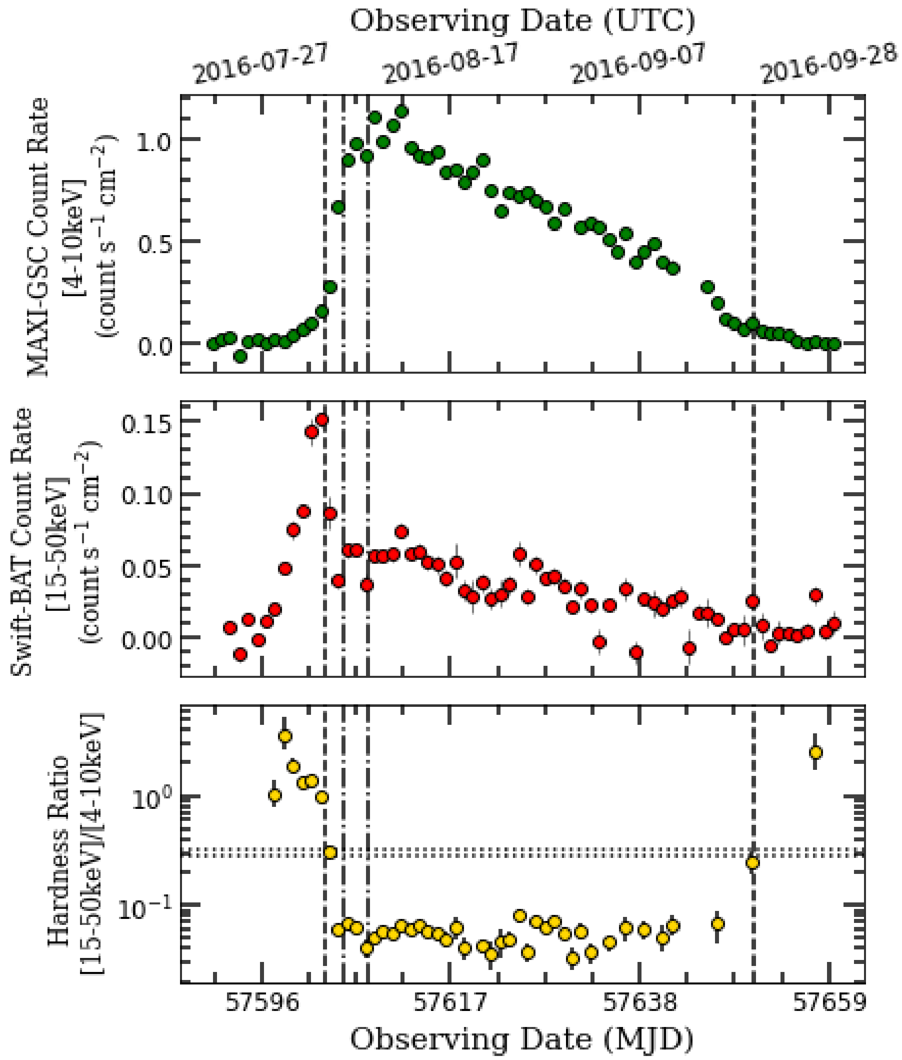

Figure 3.5: The X-ray evolution of the NSXB Aql X-1's 2016 outburst as seen by MAXI/GSC (top panel), *Swift-BAT* (middle panel), and the hardness ratio between the two instruments (third panel). The horizontal dotted lines adopt the same definition as BHXBs in Fig. 3.3. The vertical dashed lines and dashed-dotted lines show when the source was independently identified as in the soft and hard accretion states, respectively (Tasse et al., 2018). Both soft accretion states occur at an HR$\sim$0.05, well below the empirically defined transition values. This suggests one can use the BHXB transition hardness ratio to conservatively estimate if a NSXB undergoes a 'hard-only' outburst. Originally Figure 5 in (Hughes et al., 2024).

are confident that the *Swift*-BAT and MAXI/GSC hardness ratio for SAX J1810 is consistent with hard state emission throughout the 2021 outburst, and our adoption of the WATCHDOG limits is most likely appropriate (if not a conservative approximation).

2. Our *Swift*-XRT spectral modelling is consistent with hard state emission in nearly all epochs. The X-ray photon indices ($\Gamma_{\rm avg}\sim 1.6$) and low-energy black body temperatures ($kT_{\rm avg}\sim 0.6$) are typical of hard state X-ray emission from an NSXB (Lin et al., 2007). Moreover, the power-law component is the dominant flux component in all epochs (i.e., power-law flux fraction $\geq 50\%$). Although some epochs show approximately equal contributions between the blackbody and power-law components, the narrow ($0.5 - 10.0\,{\rm keV}$) energy range favours the black body component when calculating band limit flux, as the power-law component will dominate at higher energies ($\geq 10\,{\rm keV}$). The bolometric X-ray flux is more strongly dominated ($> 90\%$) by the power-law component than our observations would suggest, consistent with hard state emission. The anomalous epoch (MJD 59385; Table 3.1) that shows a clear softening of the X-ray spectrum suggests the source may have exhibited a brief deviation from a hard accretion state. Assuming a successful transition to the soft state, and given the cadence of our observations and the bracketed hard state epochs, the source would have gone through a full cycle (i.e., hard $\rightarrow$ soft $\rightarrow$ hard) in $\leq 14\,$dy before remaining in the hard state for the remaining $\sim 120\,$dy of outburst (atypical behaviour for an outbursting NSXB, see, Muñoz-Darias et al., 2014, for a review of outburst timescales). We find it more likely that the source briefly entered an intermediate state, failed to complete a transition to the soft state, and transitioned back to the hard state.

3. The evolution of our radio observations is consistent with the hard state. First, the radio and X-ray light curves show a correlated temporal evolution characteristic of hard state emission. Second, we do not detect any significant jet-quenching. Although radio emission from NSXBs has been observed in the soft state, when both hard and soft state (compact jet) radio emission has been detected, the jet emission is brighter in the hard state (at a fixed X-ray flux, e.g., Gusinskaia et al., 2017). Therefore, without a significant increase in the X-ray flux (which was never observed), we would expect a decrease in radio flux after transitioning to the soft state. We recognize that the spatially coincident, persistent radio source contaminates our ability to detect jet-quenching. However, the persistent source can not explain the joint radio–X-ray time evolution, as we would expect the radio flux to drop to the persistent level ($\sim 90\,\mu$Jy) without a similar decrease in X-ray flux. Whenever we observed an increasing X-ray flux, we observed a simultaneous increase in the radio flux density.

Comprehensive monitoring campaigns of future outbursts of SAX J1810 will be critical for confirming whether the source consistently exhibits 'hard-only' outbursts or shows a broader outburst phenomenology that sometimes results in successful transitions to the soft state (as observed in some BHXBs, e.g., H1743-322; Coriat et al., 2011; Williams et al., 2020).

3.4.2 The Origin of the Persistent Radio Emission

Our observations strongly support the existence of an unresolved, persistent, steep-spectrum radio source spatially coincident with the position of SAX J1810 ($\pm\,3''$). Considering the source exhibited a 'hard-only' outburst in 2021, we expect the radio emission to (partially) originate from a hard state jet (i.e., compact jet). The temporal coincidence between the flares at X-ray and radio frequencies is strong evidence for the existence of a steady jet. Moreover, the persistent source is weakly variable with an average flux density of $\sim 90\mu$Jy. Considering that we have multiple detections at $\gtrsim 200\mu$Jy, we have clearly detected radio emission from the compact jet.

However, a hard state jet associated with SAX J1810 cannot be the source of the persistent radio emission. Hard state jets are stationary and, therefore, would not exhibit the proper motion that we have observed (Fig. 3.1) Moreover, the locations of its luminosities on the L_R–L_X plane (red circles Fig. 3.4) are inconsistent with a hard state jet. At early times and high X-ray luminosities, the radio/X-ray luminosities are positively correlated, as expected from a compact, steady jet. Towards the end of the outburst (at $L_X \lesssim 5 \times 10^{35}\,\mathrm{ergs\,s^{-1}}$), there is a clear flattening of the correlation resulting in a $\beta < 0.06$ due to the radio luminosity remaining approximately constant while the X-ray luminosity decreased by over three orders of magnitude. The 2023 follow-up, in particular, would make SAX J1810 exceptionally radio-loud for a NSXB, consistent with the population of BHXBs. Recent analyses estimate a value of $\beta = 0.44^{+0.05}_{-0.04}$ for the total population NSXBs, with the atoll sub-population (which SAX J1810 is likely a member of) having $\beta = 0.71^{+0.11}_{-0.09}$ (Gallo et al., 2018). Both values of β reject our measurements at the $> 3\sigma$ level. Therefore, the observed radio emission likely originates from two components, with the most likely candidates of the persistent emission being either discrete jet ejecta or an unrelated, spatially coincident source.

We disfavour an origin due to jet ejecta. First, the average decay timescale for an ejectum to become undetected is $\ll 1$ year, and thus jet ejecta persisting for ~ 2 years and showing no significant decrease in the measured flux density is, in itself, unlikely. Long-lasting jet ejecta have been observed from BHXBs and are thought to be the result of jet-ISM interactions driving in situ particle acceleration and long-term synchrotron emission (e.g., Corbel et al., 2005; Bright et al., 2020; Carotenuto et al., 2022; Bahramian &

Rushton, 2022). However, such long-lasting ejecta have never been observed in NSXB (likely due to their weaker, lower-luminosity jets being unable to power such long-term emission), and when observed in BHXBs, the radio emission of long-lived ejecta is strongly variable. Second, our VLA follow-up observations suggest a 3σ upper limit on the radio spectral index of $\alpha < -1.1$, significantly steeper than expected from optically-thin synchrotron emission from jet ejecta ($\alpha \sim -0.7$). Lastly, our observations show no evidence of ballistic motion despite the source persisting for ~ 2 years, which would be the strongest evidence for a jet ejecta origin of the persistent emission. If the persistent emission originated from jet ejecta, we would have had to observe a long-lasting, non-variable, spectrally steep ejecta showing no motion on the sky. Therefore, we can rule out a jet ejecta origin with high confidence.

To estimate the probability of a spurious spatial coincidence with an unrelated source in the field we used the Python Blob Detector and Source Finder (PyBDSF; Mohan & Rafferty, 2015) to make a catalogue of all sources (in each image) with a flux density $> 74\,\mu$Jy (3σ lower than the average persistent radio flux density). We use the deep 2023 observations as their lower rms noise ($10\,\mu$Jy vs. $20\,\mu$Jy in 2021) makes PyBSDF less prone to mistaking spurious noise spikes as real sources. Due to flux variability, each image catalogue has a different number of sources. As a result, we conservatively use the 2023 May 22 image as it has more sources than the August observation and, therefore, a larger source density. We calculate the source density and then convert it to the expected number of sources within a $3''$ radius. The choice of $3''$ was motivated by the scatter of our best-fit positions. Using the expected number of sources, we then calculate the Poissonian probability of a chance coincidence of one or more unrelated background sources. The instrument's sensitivity decreases as a function of radial distance from the phase centre of the array, and thus, there is a progressively smaller number of sources catalogued at larger separations from

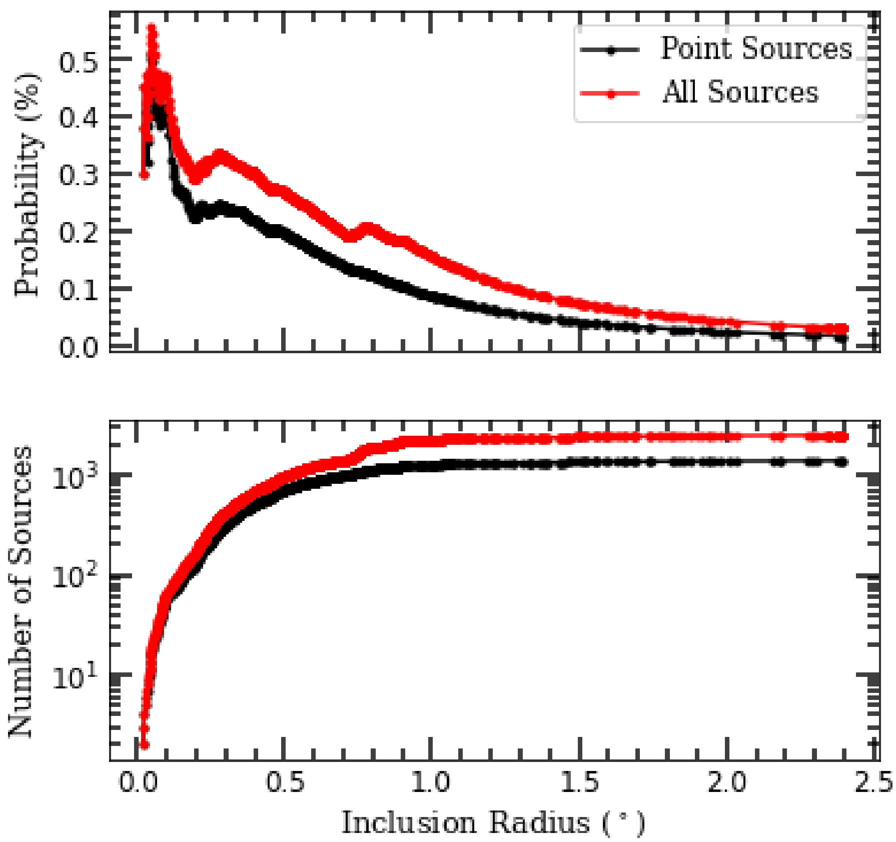

Figure 3.6: (top panel) The probability of a chance spatial coincidence between SAX J1810 and an unrelated background source as a function of the 'inclusion radius'. (bottom panel) The number of sources within the inclusion radius. We include data for the unresolved (black line) and unresolved + extended source populations (red line). We note that the sharp increase and peaks close to SAX J1810 correspond to a regime susceptible to low-count statistics. Regardless, we adopt the peak of the red curve as the most conservative estimate of the chance coincidence probability. Originally Figure 6 in (Hughes et al., 2024).

the phase centre (decreasing the source density). We applied a cut when calculating the probability to investigate this potential bias, only including sources within a certain distance from the phase centre in our calculations. In Fig. 3.6, we show the chance coincidence probability as a function of the aforementioned 'inclusion radius' for only unresolved sources (following the criteria from Appendix 3.6.1) and for both unresolved and extended sources (all sources). We adopt the peak value for all sources as our conservative estimate of the chance coincidence probability (i.e., $\sim 0.6\%$).

Radio-bright active galactic nuclei (AGN) are the dominant population of unresolved background sources. However, background AGN have an average spectral index of $\alpha \sim -0.7$. We use two recent surveys of background AGN spectral indices to estimate the probability of finding a steep spectrum AGN. Randall et al. (2012) calculated the spectral index of 166 AGN using 325, 610, and 1400 MHz flux densities. Only 43 sources had an $\alpha < -1.1$ corresponding to a probability of $\sim 26\%$. In a more recent, larger sample size survey, de Gasperin et al. (2018) measured the spectral indices of ~ 540000 radio sources (using 147 and 1400 MHz flux densities), with only a subset of ~ 32000 having an appropriately steep α. The corresponding probability is $\sim 6\%$. Adopting the older catalogue probability as a conservative estimate, we calculate the total probability of finding a spurious radio AGN with a sufficiently steep spectral index as $\sim 0.16\%$ (a $\sim 3.2\sigma$ event). Alternatively, the spectral index could suggest an origin from a class of sources known to have steep spectral indices. The most common steep spectrum source is pulsars, with average spectral indices of ~ -1.6 (Jankowski et al., 2018). We searched the Australian Telescope National Facility pulsar catalogue (Manchester et al., 2005) for any nearby known radio pulsars but found no pulsars within a radius of $0.6°$. Given that there are only 3000 known radio pulsars (corresponding to an expectation value of $\sim 2 \times 10^{-7}$ pulsars within a $3''$ radius), there is a chance coincidence probability of $\sim 0.002\%$.

When considering that pulsars tend to be distributed in the Galactic plane ($\sim$20% of the sky), and SAX J1810 is also in the Galactic plane, the chance coincidence probability would increase by a factor of $\sim$5 but is still less likely than the AGN scenario. We note that the persistent emission would correspond to a time-averaged flux of a pulsar; as a result, recent surveys that looked at this part of the sky would have detected a pulsed source (e.g., Keith et al., 2010). Moreover, MeerKAT's pulsar timing backend (i.e., MeerTRAP Sanidas et al., 2018) was operational during all of our observations but did not detect any pulsed emission from the source. Therefore, our estimated coincidence probability between SAX J1810 and an unknown pulsar is most likely an overestimate.

There is a small possibility that the persistent radio emission is local to SAX J1810. Transitional millisecond pulsars (tMSPs) — accreting neutron stars that transition between accretion-powered (i.e., NSXB-like) and radio pulsar behaviour — have shown anomalously bright radio emission while actively accreting. For instance, the tMSP, 3FGL J0427.9−6704, was measured at a point on the L_R–L_X relation that was also more consistent with the population of black hole X-ray binaries; however, its X-ray luminosities were a factor of $\gtrsim 3$ larger than our upper limits on MJD 60086 (e.g., Li et al., 2020). Other tMSPs (i.e., PSR J1023+0038) have even exhibited anti-correlations between radio and X-ray luminosities, which could allow for bright radio emission absent any X-ray detections (Bogdanov et al., 2018).

However, the properties of SAX J1810 are inconsistent with what is expected from tMSPs. Firstly, SAX J1810 does not show radio pulsations during X-ray quiescence (although eclipses or highly compact, elliptical binary orbits can prevent the detection of pulsations from tMSPs Lorimer & Kramer, 2004; Papitto et al., 2013). Second, at X-ray luminosities $\leq 10^{33}\,\mathrm{erg\,s^{-1}}$, tMSPs spectra are non-thermal ($\Gamma \leq 1.7$ Linares, 2014; Bogdanov et al., 2018; Li et al., 2020), whereas SAX J1810 is thermally dominated

($\Gamma \geq 3$ Jonker et al., 2004; Allen et al., 2018). Lastly, SAX J1810 does not exhibit any of the rapid X-ray variability that results from switching between different accretion modes (during outburst), showing, at most, modest variability (Allen et al., 2018). Although it cannot be conclusively ruled out, we find it unlikely that the persistent radio emission results from SAX J1810 being a tMSP.

Local emission, tMSP or otherwise, is difficult to reconcile with the variability in the position, as the source is spatially unresolved. Using the scatter in the measured position ($\sim 3''$) as a proxy for the expected separation of the two-source scenarios (i.e., the persistent emission is non-local), then observations by an instrument with sufficient angular resolution and sensitivity (e.g., the VLA in A-configuration or the Square Kilometer Array) during future outbursts when the compact jet is 'on' should be able to spatially resolve two distinct components. If only a single source is observed, and there continues to be temporally correlated evolution in the radio/X-ray light curves, this would strongly support the scenario where the persistent radio emission is local to SAX J1810.

3.5 Summary and Conclusions

We have presented our ~ 2 year joint radio and X-ray monitoring of the neutron star X-ray binary SAX J1810.8$-$2609. Our observations include dense (i.e., weekly cadence) observations during the source's 2021 outburst and a collection of late-time observations in 2023. The X-ray spectral properties suggested that the source remained in the hard state throughout the entire 2021 outburst. Moreover, the radio and X-ray luminosities show a temporally correlated evolution, characteristic of a hard state radio jet. We discovered a spatially coincident, persistent steep-spectrum radio source that shows no correlation with the simultaneous X-ray flux. Therefore, during the outburst, the radio emission originated from a superposition of two components: a variable hard state compact jet ($\lesssim 100\mu$Jy), and the unknown persistent source ($\sim 90\mu$Jy). The spectral index and evolution of the persistent source are inconsistent with jet ejecta. We conservatively estimated the probability of a chance coincidence with an unrelated spectrally steep background source, and although low ($\sim 0.16\%$), a background AGN seems to be the most plausible scenario.

SAX J1810.8−2609 is known to go into outburst every ~ 5 years, and future outbursts should focus on identifying the source of the persistent emission. Of the current generation of radio telescopes, the VLA (A-configuration) and the Very Long Baseline Array (VLBA) both have sufficient angular resolution and sensitivity to resolve two $\sim 100\,\mu$Jy sources (assuming a separation of $\sim 3''$). Moreover, next-generation radio interferometers, such as the Square Kilometer Array (SKA; of which MeerKAT is a pathfinder), would be able to reach the desired sensitivity with a fraction of the observing time (i.e., $\sim 10\,\mu$Jy rms for $\lesssim 3$ minutes on source; Braun et al., 2019). During the next outburst, if a second unrelated source is ruled out, follow-up observations should focus on understanding what physical mechanism is driving the persistent radio emission, whether the source is a tMSP or otherwise.

3.6 Appendices

3.6.1 Radio Astrometry

Our observations constitute the first radio detections of SAX J1810. Therefore, we designed a novel astrometric routine to test whether the radio emission is spatially coincident with the archival X-ray position of 18:10:44.47 −26:09:01.2 (with its 0.6 arcsec error; Jonker et al., 2004). We divided our astrometric analysis into two components; the first measures the random inter-epoch variability of each source position, quantifying the effects of noise fluctuations (*relative* astrometry), and the second measures the global offsets due to systematic effects in the instrumentation (*absolute* astrometry). The following section outlines our astrometry routine.

For unresolved sources (i.e., point sources) in synthesis radio images, the relative astrometric error is most often determined by the centroiding accuracy of the Gaussian fitting following deconvolution routines. As the shape of a point source adopts the shape of the synthesized beam in the absence of noise, the astrometric precision decreases with an increasing beam size. The error on the relative astrometry is often described as a function of two components: a signal-to-noise (SNR) dependency and a lower limit set by a systematic threshold. The most commonly assumed signal-to-noise scalings are, $1/\mathrm{SNR}$, or $1/(2\cdot\mathrm{SNR})$. The systematic threshold is assumed to be some fraction of the synthesized beam size. A common assumption is a lower limit of 10% of the synthesized beam size (e.g., for standard observing with the VLA[6]). We define a generalized (relative) astrometric error with the following functional form,

$$\sigma = \sqrt{(A\cdot\mathrm{SNR})^2 + B^2}, \tag{3.2}$$

where σ is the relative astrometric error expressed in units of synthesized-beam full widths at half-maxima (FWHM); and A and B are dimensionless variables that describe the SNR scaling and systematic threshold, respectively. Using PyBDSF, we generated a catalogue of (elliptical Gaussian) sources in each image; our parameters of interest were the right ascension (RA), declination (Dec), major axis FWHM of the source, minor axis FWHM of

[6]see the VLA astrometric performance summary: `https://science.nrao.edu/facilities/vla/docs/manuals/oss/performance/positional-accuracy`

the source, peak flux density (F_p), total island flux density[7] (F_i), and local rms. As SAX J1810 is isolated and unresolved, we trimmed the PyBDSF catalogue to include only similarly unresolved and isolated sources. We defined a source as unresolved if the source FWHMs deviated by $\leq 25\%$ from the synthesized beam shape. Similarly, a source is classified as isolated if the peak flux is within 25% of the island flux (e.g., $|F_p/F_i - 1| \leq 0.25$). Our routine calculates the average signal-to-noise of each source in the catalogue, and, therefore, we exclude bright transients and strongly variable sources, as their SNR ratio will vary drastically epoch-to-epoch. A source is classified as transient/variable and omitted from the sample if the source is missing from $> 25\%$ of the epochs or has a maximum and minimum flux density separated by a factor ≥ 2. Lastly, to mitigate biasing from off-axis calibration errors (e.g., from antenna pointing errors), we fit the sources that are within the inner $\sim 50\%$ of the primary beam FWHM (i.e., sources within 0.3° of the phase centre).

As the MeerKAT synthesized beam is an elliptical Gaussian, we solve for A and B independently along the RA and Dec directions. Below, we outline our fitting routine:

1. For each source, calculate an average SNR and an average position. Using the average position, calculate the RA/Dec offset from the average position for every source in each epoch.

2. Estimate the error in the astrometric precision of each source by bootstrapping the offsets, adopting the median value of the bootstrapped sample as an initial guess for σ and the ranges between the median and the 15^{th}/ 85^{th} percentiles as the 1σ $(-)/(+)$ uncertainties (Δ_σ).

[7] PyBDSF groups sources into islands, where an island is defined as a continuous region of pixels with a flux value above a user-defined threshold and at least one pixel has a flux larger than a higher (also user-defined) threshold. For large islands (i.e., extended emission), PyBDSF will fit multiple sources to a single island. For our fitting, we used 3σ and 4σ for our thresholding.

Table 3.2: Relative astrometry parameters in radio of the SAX J1810 field

Fit Type[a]	Dir.	A (%)	B (%)[b]	Pop.[c]	χ^2(dof)
Uncorrected	RA	$50.0^{+1.0}_{-1.0}$	$1.69^{+0.05}_{-0.05}$	1	159(120)
				2	1189/556
	DEC	$46.1^{+0.9}_{-0.9}$	$1.39^{+0.05}_{-0.05}$	1	164(120)
				2	1010/556
Corrected	RA	$47.3^{+0.7}_{-0.7}$	$0.39^{+0.03}_{-0.03}$	1	287(120)
				2	3084(556)
	DEC	$47.6^{+0.9}_{-0.9}$	$0.18^{+0.04}_{-0.04}$	1	159(120)
				2	2273(556)

Notes. The contrast between the χ^2 values of Population 1 and 2 highlights the effects of off-axis errors.
[a] This column indicates whether the fitting omitted (uncorrected) or used (corrected) the epoch-to-epoch astrometry correction.
[b] The fitting parameters A and B are expressed as a fraction of the synthesized beam FWHM for both the RA and Dec directions.
[c] This column indicates the population of sources used for the corresponding χ^2 calculations. Population 1 is the nearby ($< 0.3°$) isolated point sources used in the fittings. Population 2 includes all isolated point sources, regardless of distance from the phase centre. Originally Table A1 in (Hughes et al., 2024).

3. Using the σ estimates and the average SNR, solve for the scaling parameters A and B (i.e., the *uncorrected* fit). The fit implements an MCMC routine and follows the same approach detailed in 3.2.3.

4. Solve for the (inverse-variance weighted) average offset of all sources in each epoch (i.e., the *epoch-to-epoch correction*) weighting each offset using the uncorrected fit.

5. Correct the source offsets with the epoch-to-epoch correction and re-solve for A and B with the updated – corrected – offsets.

6. Repeat (ii)→(v) until the fitting converges on solutions for A and B. We defined a convergence parameter $C = (\sigma_i - \sigma_{i-1})/\Delta_\sigma$; i.e., the difference between the astrometric error of a source for the current (i) and previous ($i-1$) iterations in units of Δ_σ. The fit is said to have converged after three consecutive iterations with a mean value of

$C < 0.1$. The post-convergence fit is the *corrected* fit. Record the final epoch corrections.

The relative astrometric fitting is shown in Fig. 3.7 and the best-fit parameters are tabulated in Table 3.2. The uncorrected fits have reduced χ^2 values of ~ 1.3 (123 degrees of freedom) in both RA and Dec. Applying the epoch-to-epoch corrections (i.e., the corrected fit) shows a significant worsening of the fit quality with a reduced $\chi^2 > 2$, suggesting that a single per-epoch correction is not accurately capturing the time-dependent systematics in our observations, and a more complex epoch correction may be appropriate (e.g., one that accounts for distance and direction with respect to the phase center). We intend to expand upon this preliminary work to investigate whether the relative astrometric error is similar across a range of ThunderKAT fields.

The fits show that (for MeerKAT), the systematic threshold of the relative error is significantly lower than the commonly assumed limit of 10% the size of the synthesized beam. Moreover, the signal-to-noise dependency is similar to the commonly assumed $1/(2 \cdot \mathrm{SNR})$ scaling. Due to the residual issues in our modelling, for our SAX J1810 analysis, we conservatively rounded our uncorrected fit values, adopting $A = 0.5$ and $B = 0.02$ to quantify the relative astrometric errors.

To correct for absolute astrometry effects, we identified nine sources[8] within our field of view that are used as phase calibrators for very long baseline interferometry (i.e., with positions measured at < 10 milliarcsecond precision). Eight of the nine sources met our unresolved and isolated requirement, and we used this sub-sample for absolute astrometric corrections. After applying the epoch-to-epoch correction from the relative astrometric fitting, we measured the offsets of the eight calibrators with respect to their known positions. We then calculate each epoch's weighted mean (weighting

[8] `http://astrogeo.org/calib/search.html`

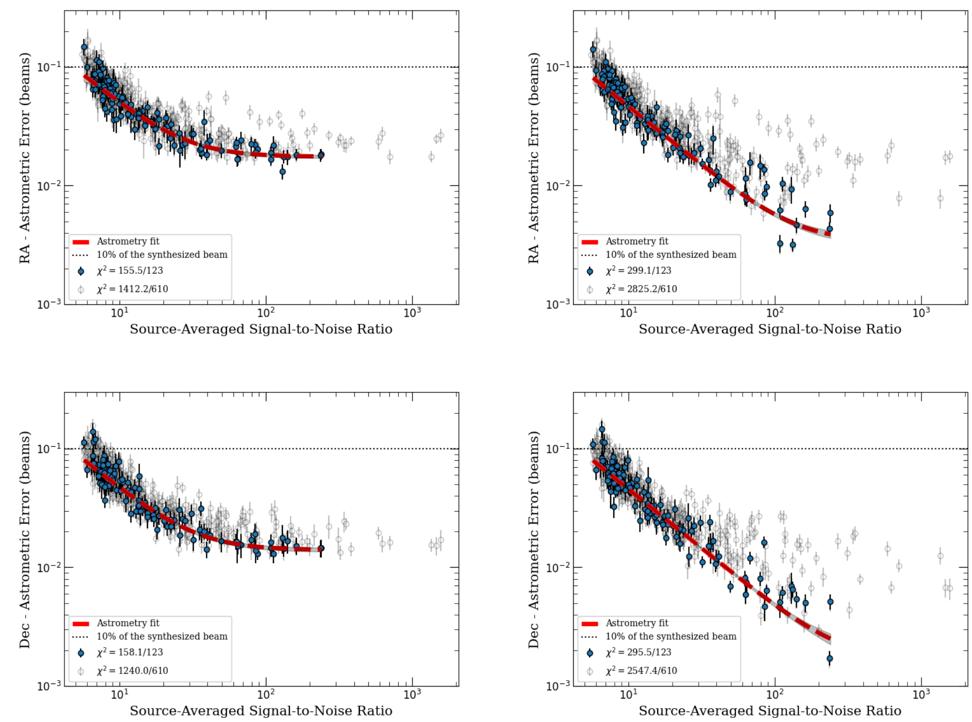

Figure 3.7: The relative astrometric fits for the population of isolated point-like sources: (top left) uncorrected Dec; (top right) corrected Dec; (bottom left) uncorrected RA; (bottom right) corrected RA. The sources used for fitting (i.e., 120 sources within 0.3° of the phase center) are given by the solid blue circles, and the total population of isolated point-like sources (612 sources) is shown as hollow black circles. The uncorrected fits are marginally acceptable for the fitted population (reduced $\chi^2 \sim 1.3$ for 123 degrees of freedom), although the corrected fits are poor (reduced $\chi^2 > 2.0$ for 610 degrees of freedom). Furthermore, the fits (uncorrected and corrected) are poor matches to the total population of isolated point-like sources, suggesting that off-axis effects (especially at high signal-to-noise ratios) are significant. Overall, the fits show that the systematic limit is well below 10% of the synthesized beam and that at SNR > 20, the global epoch effects (i.e., affecting every source in a given epoch) are the dominant astrometric error. Originally Figure A1 in (Hughes et al., 2024).

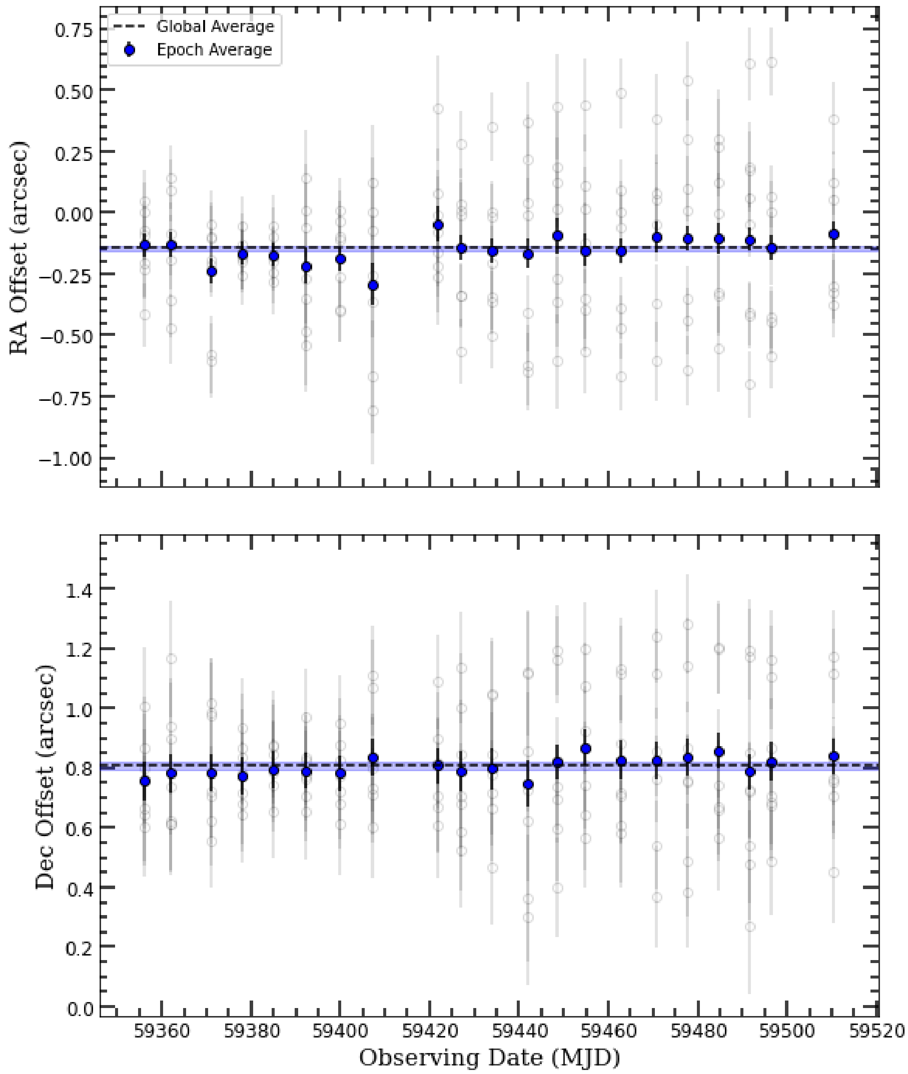

Figure 3.8: The absolute astrometric corrections from the very long baseline interferometry calibrators in the SAX J1810 field-of-view. The open black circles are the offsets of each calibrator in each epoch. The closed blue circles are the average offset in each epoch. The dashed black line is the time-independent average offset across all sources and all epochs. The blue region shows the 1σ errors on the average offset. The per-epoch RA (Dec) average offsets are consistent with the time-independent value at a reduced χ^2 of ~ 0.65 (~ 0.24) for 22 degrees of freedom. These low χ^2 values suggest we may overestimate the relative astrometric error. Originally Figure A2 in (Hughes et al., 2024).

Table 3.3: Timing parameters of the 2021 August 7 Type I X-ray burst. Originally Table B1 in (Hughes et al., 2024).

τ (s)	R_0 (counts s^{-1})	A (counts s^{-1})	χ^2(dof)
15.8 ± 0.2	27.0 ± 0.5	332 ± 3	86(89)

each source by their relative astrometric errors). Lastly, we calculated a single time-independent absolute astrometric correction (see Fig. 3.8). The epoch-to-epoch correction removed any (substantial) temporal variability, and, as a result, the per-epoch average offsets are consistent with a single (time-independent) RA/Dec offset.

The final astrometric error ($\sigma_{\rm tot}$) was calculated by adding (in quadrature) the relative astrometric precision (σ), the error on the epoch-correction ($\sigma_{\rm epoch}$), and the error on the absolute offset ($\sigma_{\rm abs}$),

$$\sigma_{\rm tot} = \sqrt{\sigma^2 + \sigma_{\rm epoch}^2 + \sigma_{\rm abs}^2}. \tag{3.3}$$

These are the errors shown in Fig. 3.1. We note that given the signal-to-noise ratio of our SAX J1810 detections (SNR $\lesssim 10$), the relative astrometry term, σ, dominates the quoted errors.

3.6.2 Type I X-Ray Burst

Figure 3.9 shows the parameters of the 2021 August 7 (MJD 59433) Type I X-ray burst. The top panel shows the 1-second binned light curves and the timing fits; the second panel shows the bolometric X-ray flux of the blackbody component; the third panel shows the temperature of the blackbody component; and the bottom panel shows the normalization of the `bbodyrad` component, defined as R^2/D^2, where R is the source radius in units of km and D is the distance to the source in units of 10 kpc.

The burst began its rise at 14:14:12 on 2021 August 7 (MJD 59433.59319), reaching a peak count rate of ~ 400 counts s^{-1} with a rapid 7 ± 1 s rise time

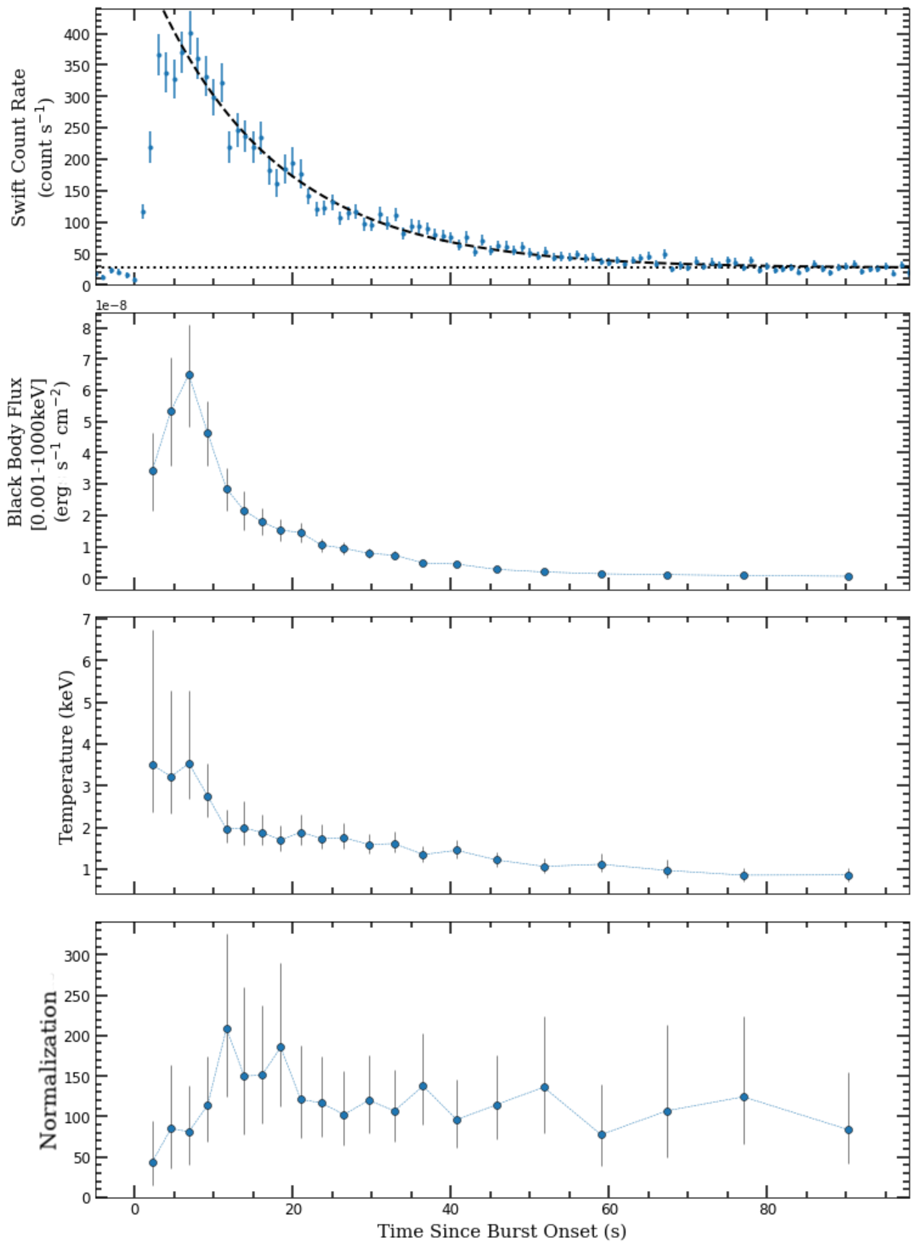

Figure 3.9: Spectral and timing fits from the Type I X-ray burst observed on 2021 August 7. The top panel shows the 1-second binned light curves. We overlayed the best-fit exponential decay (dashed line) and the constant (pre-burst) count rate(dotted line); the timing fit parameters are tabulated in Table 3.3. The bolometric X-ray flux (second panel), temperature (third panel), and `bbodyrad` normalization (bottom panel) of the black body component do not show conclusive evidence of PRE. Originally Figure B1 in (Hughes et al., 2024).

before decaying for the remainder of our observations. The timing fit converged on an e-folding decay time of $\tau = 15.8 \pm 0.2$ s (full fit parameters in Table 3.3). The burst parameters are consistent with those in the MINBAR burst catalogue Galloway et al. (2020) in both rise ($3.4^{+5.6}_{-2.4}$ s) and e-folding decay times (8^{+21}_{-4} s).

During its 2007 outburst, SAX J1810 exhibited 531.8 Hz oscillations in the light curves of a Type I X-ray burst, likely the result of the spin frequency of the neutron star (Bilous et al., 2018). Following the prescription outlined in Bilous et al. (2018) we searched for burst oscillations in our (1.8 ms resolution) light curves by calculating the power spectrum in sliding windows with widths of 0.5, 1, 2, and 4 s, where each subsequent window is offset by 0.5 s from the previous one. We found no evidence of burst oscillations. However, the temporal resolution of *Swift*-XRT WT mode (1.8 ms) makes our power spectra insensitive to frequencies above $\sim$280 Hz. Assuming the oscillations result from the spin period of the neutron star, we do not expect the oscillation frequency to evolve drastically between the 2007 and 2021 outbursts.

Furthermore, SAX J1810 is known to exhibit PRE (i.e., during the 1998 outburst a PRE signature provided the current distance constraint of 4.9 ± 0.3 kpc; Natalucci et al., 2000). Therefore, we performed time-resolved intra-epoch spectral modelling to search for evidence of PRE. We observe some radius (as seen from the `bbodyrad` normalization) and temperature evolution, although the large errors greatly reduce their significance. Assuming a distance of 4.9 kpc, the radius of the blackbody component ranges from $3.2^{+3.5}_{-2.6}$ to $6.7^{+5.0}_{-4.2}$ km (i.e., from $\sim$5-to-100% of the neutron star's surface assuming a 10 km stellar radius). However, the radius and temperature evolution does not occur alongside a period of (approximately) constant X-ray flux; thus, we do not detect PRE.

3.6.3 Data Tables

Table 3.4: Radio properties of SAX J1810 during and after its 2021 outburst. Originally Table C1 in (Hughes et al., 2024).

MJD	Date	Instrument	Central Frequency [GHz]	Flux Density [μJy]
59356	2021-05-22	MeerKAT	1.3	232 ± 18
59362	2021-05-27	MeerKAT	1.3	197 ± 18
59371	2021-06-05	MeerKAT	1.3	207 ± 17
59378	2021-06-12	MeerKAT	1.3	143 ± 15
59385	2021-06-19	MeerKAT	1.3	128 ± 18
59392	2021-06-27	MeerKAT	1.3	139 ± 20
59400	2021-07-04	MeerKAT	1.3	135 ± 16
59407	2021-07-12	MeerKAT	1.3	153 ± 26
59422	2021-07-26	MeerKAT	1.3	175 ± 20
59427	2021-07-31	MeerKAT	1.3	197 ± 17
59434	2021-08-07	MeerKAT	1.3	140 ± 19
59442	2021-08-15	MeerKAT	1.3	163 ± 24
59449	2021-08-22	MeerKAT	1.3	110 ± 22
59455	2021-08-28	MeerKAT	1.3	140 ± 23
59463	2021-09-05	MeerKAT	1.3	97 ± 17
59471	2021-09-13	MeerKAT	1.3	98 ± 22
59478	2021-09-20	MeerKAT	1.3	83 ± 17
59485	2021-09-27	MeerKAT	1.3	142 ± 26
59492	2021-10-04	MeerKAT	1.3	92 ± 17
59497	2021-10-09	MeerKAT	1.3	97 ± 17
59511	2021-10-23	MeerKAT	1.3	82 ± 16
60086	2023-05-22	MeerKAT	1.3	112 ± 12
60142	2023-07-17	VLA	3.0	<30
60142	2023-07-17	VLA	6.0	<18
60169	2023-08-13	MeerKAT	1.3	75 ± 11

Table 3.5: X-ray spectral properties from the *Swift*-XRT monitoring of SAX J1810. Originally Table C2 in (Hughes et al., 2024).

MJD	Date	E (keV)	Γ	$F_{\rm PL}$ [10^{-11} erg s^{-1} cm^{-2}]	kT (keV)	$F_{\rm BB}$ [10^{-11} erg s^{-1} cm^{-2}]	$F_{\rm tot}$ [10^{-11} erg s^{-1} cm^{-2}]	χ^2(dof)
59364	2021-05-30	0.5–10	$1.42^{+0.09}_{-0.08}$	$131.9^{+8.5}_{-8.2}$	$0.68^{+0.09}_{-0.07}$	$23.6^{+2.0}_{-2.7}$	$155.5^{+8.8}_{-8.6}$	257(278)
		1.0–10	$1.42^{+0.08}_{-0.04}$	$118.0^{+7.6}_{-7.5}$	$0.68^{+0.09}_{-0.07}$	$21.7^{+1.9}_{-2.5}$	$139.7^{+7.9}_{-7.9}$	257(278)
59378	2021-06-13	0.5–10	$1.82^{+0.16}_{-0.11}$	$24.0^{+2.7}_{-2.7}$	$0.76^{+0.08}_{-0.07}$	$9.0^{+0.8}_{-0.9}$	$33.0^{+2.8}_{-2.8}$	191(190)
		1.0–10	$1.82^{+0.13}_{-0.11}$	$19.5^{+2.5}_{-1.4}$	$0.76^{+0.08}_{-0.11}$	$8.5^{+0.8}_{-0.9}$	$28.0^{+2.7}_{-1.6}$	191(190)
59385	2021-06-20	0.5–10	$2.88^{+0.18}_{-0.08}$	$26.6^{+1.7}_{-1.5}$	$1.16^{+0.10}_{-0.09}$	$12.7^{+0.7}_{-0.8}$	$39.3^{+1.8}_{-1.7}$	217(186)
		1.0–10	$2.89^{+0.18}_{-0.18}$	$13.5^{+2.0}_{-1.7}$	$1.16^{+0.10}_{-0.08}$	$12.5^{+0.7}_{-0.4}$	$26.0^{+2.1}_{-1.8}$	216(186)
59392	2021-06-27	0.5–10	$1.20^{+0.08}_{-0.14}$	$24.9^{+2.4}_{-2.2}$	$0.64^{+0.04}_{-0.03}$	$14.1^{+0.6}_{-0.7}$	$39.0^{+2.4}_{-2.3}$	237(244)
		1.0–10	$1.22^{+0.12}_{-0.15}$	$23.0^{+2.0}_{-2.0}$	$0.65^{+0.03}_{-0.03}$	$12.9^{+0.5}_{-0.7}$	$35.9^{+2.1}_{-2.1}$	237(244)
59399	2021-07-04	0.5–10	$1.83^{+0.10}_{-0.08}$	$49.6^{+3.8}_{-1.8}$	$0.72^{+0.16}_{-0.13}$	$6.7^{+1.0}_{-1.3}$	$56.2^{+3.9}_{-2.2}$	238(231)
		1.0–10	$1.83^{+0.10}_{-0.04}$	$40.3^{+3.4}_{-3.8}$	$0.73^{+0.16}_{-0.13}$	$6.2^{+1.0}_{-1.3}$	$46.5^{+3.5}_{-4.0}$	238(231)
59406	2021-07-11	0.5–10	$1.61^{+0.16}_{-0.14}$	$30.0^{+3.8}_{-4.0}$	$0.72^{+0.12}_{-0.11}$	$8.9^{+1.0}_{-1.3}$	$39.0^{+4.0}_{-4.2}$	95(105)
		1.0–10	$1.62^{+0.15}_{-0.14}$	$25.8^{+3.6}_{-3.9}$	$0.72^{+0.12}_{-0.10}$	$8.4^{+1.0}_{-1.3}$	$34.1^{+3.7}_{-4.1}$	95(105)
59413	2021-07-18	0.5–10	$1.35^{+0.09}_{-0.11}$	$46.4^{+4.7}_{-4.8}$	$0.79^{+0.11}_{-0.08}$	$14.9^{+1.2}_{-1.4}$	$61.3^{+4.9}_{-5.0}$	237(194)
		1.0–10	$1.35^{+0.12}_{-0.06}$	$42.0^{+3.5}_{-4.8}$	$0.80^{+0.11}_{-0.09}$	$14.2^{+1.3}_{-1.4}$	$56.2^{+3.7}_{-5.0}$	237(194)
59433	2021-08-07	0.5–10	$1.65^{+0.10}_{-0.08}$	$77.5^{+5.9}_{-6.3}$	$0.91^{+0.13}_{-0.11}$	$14.6^{+1.9}_{-2.0}$	$92.1^{+6.2}_{-6.6}$	318(340)
		1.0–10	$1.66^{+0.09}_{-0.08}$	$65.8^{+5.7}_{-6.5}$	$0.91^{+0.13}_{-0.11}$	$14.1^{+1.8}_{-2.1}$	$80.0^{+6.0}_{-6.8}$	318(340)
59437	2021-08-11	0.5–10	$1.75^{+0.07}_{-0.06}$	$87.0^{+5.2}_{-5.4}$	$0.82^{+0.23}_{-0.24}$	$5.3^{+1.4}_{-2.4}$	$92.3^{+5.4}_{-5.9}$	291(324)
		1.0–10	$1.75^{+0.09}_{-0.03}$	$72.1^{+4.8}_{-5.9}$	$0.84^{+0.33}_{-0.25}$	$5.1^{+1.4}_{-2.4}$	$77.2^{+5.0}_{-6.3}$	291(324)
59451	2021-08-25	0.5–10	$1.43^{+0.16}_{-0.17}$	$16.8^{+1.1}_{-0.9}$	$0.60^{+0.05}_{-0.04}$	$9.0^{+0.5}_{-0.7}$	$25.8^{+1.2}_{-1.1}$	135(142)
		1.0–10	$1.44^{+0.15}_{-0.18}$	$15.0^{+1.4}_{-1.7}$	$0.60^{+0.04}_{-0.04}$	$8.0^{+0.5}_{-0.6}$	$23.1^{+1.4}_{-1.8}$	135(142)

Continued on next page

Table 3.5 continued

MJD	Date	E (keV)	Γ	$F_{\rm PL}$ [10^{-11} erg s^{-1} cm^{-2}]	kT (keV)	$F_{\rm BB}$ [10^{-11} erg s^{-1} cm^{-2}]	$F_{\rm tot}$ [10^{-11} erg s^{-1} cm^{-2}]	χ^2(dof)
59458	2021-09-01	0.5–10	$1.34^{+0.14}_{-0.15}$	$20.1^{+2.1}_{-1.9}$	$0.60^{+0.05}_{-0.04}$	$9.6^{+0.6}_{-0.6}$	$29.6^{+2.2}_{-2.0}$	169(181)
		1.0–10	$1.35^{+0.13}_{-0.14}$	$18.2^{+1.8}_{-1.7}$	$0.60^{+0.04}_{-0.04}$	$8.5^{+0.5}_{-0.6}$	$26.7^{+1.8}_{-1.8}$	169(181)
59462	2021-09-05	0.5–10	$0.95^{+0.42}_{-0.71}$	$20.4^{+4.6}_{-5.2}$	$0.62^{+0.11}_{-0.06}$	$11.7^{+1.7}_{-1.9}$	$32.1^{+4.9}_{-5.5}$	50(33)
		1.0–10	$0.92^{+0.47}_{-0.71}$	$19.1^{+4.3}_{-4.5}$	$0.62^{+0.12}_{-0.03}$	$10.9^{+1.4}_{-1.9}$	$30.0^{+4.6}_{-4.9}$	50(33)
59468	2021-09-11	0.5–10	$1.23^{+0.23}_{-0.29}$	$10.8^{+1.6}_{-1.5}$	$0.53^{+0.04}_{-0.02}$	$5.9^{+0.5}_{-0.3}$	$16.8^{+1.7}_{-1.5}$	103(94)
		1.0–10	$1.23^{+0.23}_{-0.26}$	$10.0^{+1.3}_{-1.3}$	$0.53^{+0.04}_{-0.04}$	$5.1^{+0.4}_{-0.5}$	$15.1^{+1.4}_{-1.4}$	103(94)
59475	2021-09-18	0.5–10	$0.85^{+0.44}_{-0.61}$	$9.0^{+1.8}_{-0.8}$	$0.53^{+0.04}_{-0.04}$	$6.7^{+0.6}_{-0.8}$	$15.7^{+1.9}_{-1.1}$	77(69)
		1.0–10	$0.87^{+0.43}_{-0.65}$	$8.6^{+1.5}_{-1.5}$	$0.53^{+0.04}_{-0.04}$	$5.8^{+0.6}_{-0.7}$	$14.4^{+1.6}_{-1.6}$	77(69)
59482	2021-09-25	0.5–10	$1.34^{+0.19}_{-0.28}$	$14.1^{+1.5}_{-1.8}$	$0.56^{+0.07}_{-0.06}$	$5.3^{+0.5}_{-0.7}$	$19.5^{+1.6}_{-2.0}$	59(87)
		1.0–10	$1.34^{+0.19}_{-0.27}$	$12.8^{+1.6}_{-1.6}$	$0.56^{+0.05}_{-0.06}$	$4.7^{+0.4}_{-0.6}$	$17.5^{+1.6}_{-1.7}$	59(87)
59489	2021-10-02	0.5–10	$0.84^{+0.39}_{-0.53}$	$12.7^{+1.9}_{-1.8}$	$0.56^{+0.02}_{-0.03}$	$10.9^{+1.1}_{-1.0}$	$23.5^{+2.2}_{-2.1}$	107(110)
		1.0–10	$0.96^{+0.29}_{-0.65}$	$12.3^{+1.7}_{-1.8}$	$0.57^{+0.02}_{-0.03}$	$9.2^{+0.9}_{-0.7}$	$21.5^{+1.9}_{-1.9}$	107(110)
59496	2021-10-09	0.5–10	$0.43^{+0.93}_{-0.43}$	$4.0^{+0.8}_{-0.8}$	$0.47^{+0.05}_{-0.05}$	$2.8^{+0.2}_{-0.7}$	$6.8^{+0.8}_{-1.0}$	29(35)
		1.0–10	$0.45^{+0.93}_{-0.45}$	$3.9^{+0.4}_{-0.4}$	$0.47^{+0.04}_{-0.05}$	$2.3^{+0.2}_{-0.6}$	$6.2^{+0.4}_{-0.7}$	29(35)
59504	2021-10-17	0.5–10	—	—	—	—	$0.21^{+0.12}_{-0.12}$	59(95)
		1.0–10	—	—	—	—	$0.18^{+0.11}_{-0.10}$	59(95)
59511	2021-10-24	0.5–10	—	—	—	—	$0.05^{+0.04}_{-0.03}$	7(7)
		1.0–10	—	—	—	—	$0.04^{+0.04}_{-0.03}$	7(7)

Note. The model parameters are: power-law photon index (Γ), power-law flux ($F_{\rm PL}$), black body temperature (kT), black body flux ($F_{\rm BB}$), and the total X-ray flux ($F_{\rm tot}$).

Chapter 4

X-ray and Radio Monitoring of the Neutron Star Low Mass X-ray Binary 1A 1744−361: Quasi Periodic Oscillations, Transient Ejections, and a Disc Atmosphere

The following chapter details our 2022 monitoring of the NSXB 1A 1744−361. This research was part of a more extensive collaboration led by myself and a former doctoral student, Mason Ng, from the Massachusetts Institute of Technology. To accurately represent my contributions to this research, I will only discuss the components focused on state transitions and the evolving jet. The full paper, Ng et al. (2024), contains a substantially more comprehensive summary of the source's X-ray properties.

4.1 1A 1744−361

1A 1744−361 (also known as XTE J1748-361; Remillard et al., 2003), is an outbursting X-ray binary discovered in 1976 by the Arial V satellite (Carpenter et al., 1977). The source has undergone several subsequent outbursts (e.g., Bahramian et al., 2013), and has exhibited Type I X-ray bursts (Bhattacharyya et al., 2006), resulting in its classification as an atoll NSXB. 1A 1744−361 is a known 'dipper' (Bhattacharyya et al., 2006), exhibiting periodic decreases in its X-ray luminosity due to X-ray absorption by the accretion flow as it passes through our line of sight (Church et al., 1998, 2005; Díaz Trigo et al., 2006; Bałucińska-Church et al., 2011). The properties of the X-ray 'dips' have provided estimates of the inclination angle ($i \sim 60-75°$; where $i = 90°$ is 'edge-on') and binary period of the system (97 ± 22 min; although a period of ~ 48 or 194 min has not been ruled out).

The source's most recent outburst (the subject of this chapter) began in May 2022, when MAXI/GSC detected a rise in X-ray luminosity. Following the detection, we began a monitoring campaign that tracked the source's X-ray and radio evolution during the three-month-long outburst. 1A 1744−361 exhibited a broad range of X-ray timing and spectral phenomenology extensively detailed in Ng et al. (2024). Most significantly, we observed rapid X-ray hardness variability and the presence of NBO-type QPOs when the source was near its peak X-ray luminosity. These properties identified 1A 1744−361 as the seventh NSXB to exhibit transitions from atoll- to Z-source behaviour, providing further evidence that atoll- and Z-sources are different accretion states (rather than distinct subpopulations). My co-author led the X-ray timing analysis, so its details will be excluded from the chapter. Instead, I will discuss the evolution of the jet during the observed state transitions, which provided the best evidence for jet ejections during the canonical hard-to-soft state transition in an NSXB.

The rest of the chapter will be outlined as follows: Section 4.2 outlines the observations and data analysis procedures; Section 4.3 presents the results; Section 4.4 discusses the outburst in the context of accretion-jet coupling; and 4.5 summarizes our findings and discusses future monitoring strategies.

4.2 Observations and Data Analysis

4.2.1 MeerKAT

We observed 1A 1744−361 with MeerKAT as a part of ThunderKAT (Fender et al., 2016). Our observing began with a single rapid response observation on 2022 May 31 (MJD 59730), ∼2 dy after the first MAXI/GSC X-ray detection (and ∼ 8 hours after the outburst's initial reporting; Kobayashi et al., 2022). Following this rapid response, we began a monitoring campaign on 2022 June 3 (MJD 59733), observing the source every ∼7 dy until 2022 August 27 (MJD 59818). Each observation consisted of a single scan with 15 minutes on-source flanked by two 2-minute scans of a nearby gain calibrator (J1830−3602). Each epoch also included a 5-minute scan of PKS B1934−638 (J1939−6342) for flux and bandpass calibration. Our observations used MeerKAT's L-band receiver, with a central frequency of 1.28 GHz and a total (un-flagged) bandwidth of 856 MHz. I followed the calibration and imaging scheme from Section 3.2.1.

4.2.2 NICER

The majority of the X-ray properties were inferred from observations taken with the Neutron star Interior Composition Explorer (NICER; Gendreau et al., 2016). While the NICER analysis is not an explicit component of this thesis, the measured X-ray properties are critical for understanding the evolution of 1A 1744−361. Here, I quote the original paper and direct the reader to (Ng et al., 2024) for more information on the methodology used to analyze the NICER data:

> NICER, an external payload on the International Space Station, consists of 52 operating co-aligned X-ray concentrator optics and silicon drift detectors in focal plane modules (FPMs). NICER has fast-timing capabilities in the 0.2–12.0 keV energy range, allowing for a GPS time-tagging accuracy of 100 ns.
>
> NICER observed 1A 1744−361 starting from 2022 June 3 through 2022 August 31, with ObsIDs starting with 5202 and 5406. We processed the NICER observations with HEASoft version 6.31.1 and the NICER Data Analysis Software (NICERDAS) version 10a (2022-12-16_V010a) with calibration version `xti20221001`. Our data processing criteria included the following: a source angular offset of `ANG_DIST` $< 0^{\circ}.015$; elevation angle from the Earth limb `ELV` $> 20^{\circ}$; NICER being outside the South Atlantic anomaly; bright Earth limb angle `BR_EARTH` $> 30^{\circ}$; undershoot rate (dark current; per FPM) range of `underonly_range` = 0–500; overshoot rate (charged particle saturation; per FPM) range of `overonly_range` = 0–1.5. We also applied `COR_SAX` (magnetic cut-off rigidity in GeV/c) filtering of `COR_SAX` > 1.5 to filter out background flares. This resulted in 53.5 ks of filtered exposure for scientific analysis out of 95.3 ks of unfiltered exposure

4.2.3 *Swift*-XRT

As a part of our ThunderKAT monitoring, for each radio epoch, we typically acquired an X-ray observation with *Swift*-XRT that was quasi-simultaneous (within $\sim$2 dy) of each of our weekly radio observations. The weekly cadence monitoring began on 2022 June 1 (MJD 59731) and continued until 2022 August 8 (MJD 59798), totaling ten individual epochs (target ID: 31222). We increased the cadence of our *Swift*-XRT monitoring (one observation every $\leq$3 dy) after the source's X-ray flux began decaying. We acquired an additional 12 epochs of this 'high-cadence' monitoring that began on 2022

August 10 (MJD 59801) and ended on 2022 August 24 (MJD 59815). There was no further *Swift*-XRT monitoring.

I created 0.5–10 keV light curves using SWIFTTOOLS (Evans et al., 2007, 2009), extracting a single data point for each observation. Most X-ray spectral properties were derived from our extensive and more sensitive NICER coverage. However, we extracted a spectrum (using SWIFTTOOLS) from the first *Swift*-XRT observation on 2022 June 1 (MJD 59731) as it was taken two days before the first NICER observation.

I used the HEASOFT package (version 6.25) for my spectral analysis, binning the spectrum on 25-count intervals (allowing for χ^2 statistics during spectral fitting) with a modified GRPPHA script[1]. Using XSPEC (Arnaud, 1996a), we fit two spectral models; (i) a hard state model consisting of a power-law component and a black body (`tbabs` $\times$ (`pegpwrlw` + `bbody`), in XSPEC parlance); (ii) a soft state model consisting of a black body and multi-colour disc black body (`tbabs` $\times$ (`diskbb` + `bbody`), in XSPEC parlance). I modelled absorption from the interstellar medium following the abundances from Wilms et al. (2000), the equivalent hydrogen column density (N_H), fixed at the value inferred from the NICER modelling $N_H \equiv 0.44$. These two models were chosen to remain consistent with the hard/soft state modelling with NICER and were motivated by the findings of Lin et al. (2007). In Ng et al. (2024), we fit a single power-law model to the *Swift*-XRT observations. Here, I provide a more detailed spectral analysis of the *Swift*-XRT, but note that it does not change the interpretations.

[1]In contrast to the GRPPHA function built-into HEASOFT, my custom script does not omit edge bins if they do not meet the 25-count interval. Instead, it include these counts in the first/last bin, i.e., the edge bins have ≥ 25 counts. This will have a minor effect for most cases but ensures I model the entire spectral range.

4.2.4 MAXI/GSC

I include daily MAXI/GSC X-ray light curves extracted using the MAXI on-demand product service[2]. Our light curves include the full, 2–20 keV, energy range, and the X-ray hardness ratio (HR)

$$\mathrm{HR} \equiv \frac{R_{6\text{–}20}}{R_{2\text{–}6\,keV}}. \tag{4.1}$$

calculated from the rates in the 2–6 ($R_{2\text{–}6\,keV}$) and 6–20 keV ($R_{6\text{–}20\,keV}$) sub-bands.

4.3 Results

4.3.1 X-ray Evolution

The X-ray count rate (which tracks the X-ray luminosity) as measured by *Swift*-XRT, NICER, and MAXI/GSC is shown in the top panel of Fig. 4.1. 1A 1744−361 exhibited a fast rise in X-ray luminosity, reaching its peak level ~ 3 dy after the outburst onset. The X-ray emission remained bright with a NICER count rate of 600–800 count s^{-1} for ~ 30 dy before decaying to a quiescence level over the next ~ 60 dy. X-ray modelling revealed the source was in the soft state at the start of our NICER monitoring, with an X-ray continuum well-fit by two thermal components: a multi-colour disc black body and a black body. 1A 1744-361 transitioned to the hard state on 2022 August 10 (MJD 59801; dashed line, Fig. 4.1), with an X-ray spectrum well-fit by a (broken) power-law and black body (Ng et al., 2024). Following the hard state transition, 1A 1744-361 returned to quiescence, showing no evidence of re-flaring.

As discussed in 4.2.3, I performed a supplementary spectral analysis

[2] https://maxi.riken.jp/mxondem/

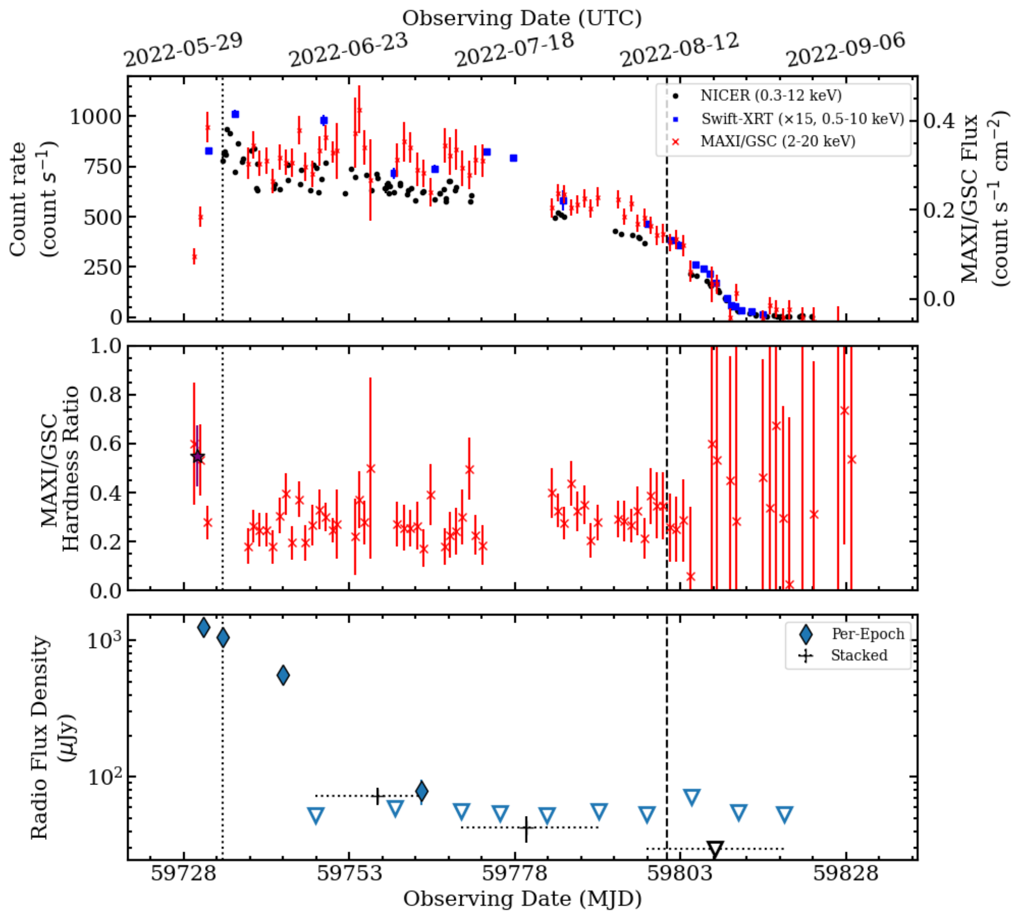

Figure 4.1: X-ray and radio light curves tracking the 2022 outburst of 1A 1744−361. (top panel) The X-ray count rate evolution of the rise and decay of the outburst from NICER (black circles; 0.3–12.0 keV), *Swift*-XRT (blue squares; 0.5–10 keV), and MAXI/GSC (red 'x' markers; 2–20 keV). The *Swift*-XRT count rates were re-scaled by a factor of 15 for plotting purposes. The vertical dotted lines correspond to the start of the NICER monitoring (which revealed the source to be in the soft state), and the vertical dashed line corresponds to the observed soft-to-hard state transition. (middle panel) MAXI/GSC hardness ratio. The red 'x' markers are the daily observations, and the purple star is the average hardness ratio value from the first two observations. (bottom panel) 1.3 GHz flux densities from MeerKAT radio evolution. Blue diamonds and triangles are the flux density per observation of the single epoch detections and 3σ upper limits, respectively. Black crosses denote the stacked observation detections, with horizontal error bars showing the corresponding stacked epochs. The last stacked data point (black triangle) is a non-detection with a 3σ upper limit of $\sim 30\,\mu$Jy. This figure is a modified version of Figure 1 from Ng et al. (2024), which is licensed under CC BY 4.0.

of the first *Swift*-XRT observation taken on 2022 June 1 (MJD 59801). The X-ray spectrum was well-fit by both the soft (`bbody+diskbb`) and hard state (`bbody+pegpwrlw`) models with formally acceptable reduced χ^2[3] values of ~ 1.009 and ~ 1.005 (592 degrees of freedom), respectively. However, the characteristic temperatures of the soft state components were $kT_{\rm BB} \sim 2.8^{+2.9}_{-0.7}$ keV (`bbody`) and $kT_{\rm disc} \sim 1.4^{+0.16}_{-0.12}$ keV (`diskbb`), substantially higher energy than the values measured during the first NICER observation ($kT_{\rm BB} \sim 1.5$ keV; $kT_{\rm disc} \sim 0.8$ keV). The excess high energy X-ray emission suggests that the source had not fully transitioned to the soft state during the first *Swift*-XRT observation.

4.3.2 Radio Evolution

I show the MeerKAT flux densities in the bottom panel of Figure 4.1. The measured flux densities in each of the 15-minute observations are shown in blue, where I plot both $>4\sigma$ detections (diamonds) and 3σ upper limits (triangles). During the first three observations, the source was radio-bright with a flux density $\gtrsim 500\,\mu$Jy. As a result, I was able to measure the (intra-band) spectral index in the bright epochs by breaking the bandwidth in half (i.e., 856–1284 and 1284–1712 MHz) and measuring the flux density in each sub-band. We measured radio spectral indices of $\alpha = -0.2 \pm 0.1$, -0.3 ± 0.1, and -0.6 ± 0.2 during the 2022 May 31 (MJD 59730), 2022 June 3 (MJD 59733), and 2022 June 12 (MJD 59742) observations, respectively. Following the bright detections, the source flux density rapidly decayed over ~ 7 dy, dropping below our single observation detection threshold of $\sim 80\,\mu$Jy for all but one of our remaining observations. I performed image stacking to increase the S/N, grouping the last eleven epochs into three stacked images of three, four, and four independent observations. The stacked flux densities

[3]Throughout this thesis I represent a reduced χ^2 statistic as $\chi^2_{\rm red}$, as opposed to the standard nomenclature of χ^2_ν. Given that ν is also used to represent the frequency of an electromagnetic wave, I replaced ν with 'red' to avoid confusion.

are shown in black, where I detected the source at $\gtrsim 4.5\sigma$ in the first two stacked images (black crosses). However, even with the improved S/N, I still did not detect the source in the final stacked image, thus representing it as a 3σ upper limit. I did not observe any correlated evolution (in time) when comparing the X-ray and radio light curves, as would be expected from hard-state jet emission.

To constrain the position of 1A 1744−361 on the L_R–L_X plane after it transitioned to the hard state, I calculated the radio luminosity at 5 GHz ($L_{R,5}$) through

$$L_{R,5} = 4\pi F_{R,\nu} \nu d^2 \approx 3.9 \times 10^{26}\,\mathrm{erg\,s^{-1}} \left(\frac{F_{R,\nu}}{1\,\mu\mathrm{Jy}}\right) \left(\frac{d}{8\,\mathrm{kpc}}\right)^2, \tag{4.2}$$

where $1\,\mathrm{Jy} = 10^{-23}\,\mathrm{erg\,s^{-1}\,cm^{-2}\,Hz^{-1}}$, $\nu = 5\,\mathrm{GHz}$, and adopting the standard assumption of a flat radio spectral index ($\alpha = 0$) in $F_{R,\nu} \propto \nu^{\alpha}$. The source's position on the L_R–L_X plane is shown in Fig. 4.2.

4.4 Discussion

The X-ray spectral and timing properties of 1A 1744−361 showed clear signatures of spectral state transitions. When we began our NICER monitoring, the X-ray spectra were best fit with thermal components (i.e., a blackbody and multi-colour disc blackbody), characteristic of a soft state (Lin et al., 2007); therefore, NICER did not capture the initial hard-to-soft transition. Looking at the MAXI/GSC evolution, the hardness ratio showed a drop from $\sim$0.6 on 2022 May 30 (the highest — and thus the hardest — hardness ratio of the entire outburst) to $\sim$0.2 on 2022 June 1 (a day after the initial radio detection), where it remained until MAXI/GSC could no longer constrain the hardness ratio. While the error bars are large, this behaviour is consistent with an evolution towards a soft(er) accretion state at early times. Additionally, the spectral properties measured by *Swift*-XRT (on the same day) showed an excess of high-energy X-ray emission compared to the NICER soft state modelling, suggesting the source had not yet transitioned to the soft state. Given the expected behaviour and the evolution of our early X-ray monitoring, we find it likely the source underwent a hard-to-soft state transition between 2022 May 29 (the onset of the outburst) and 2022 June 3 (the first NICER observation, shown as a dotted line on Fig. 4.1).

Initially, the bright radio emission was thought to be associated with a (hard state) compact jet (Hughes et al., 2022). However, the source's occupancy of the soft state, the absence of correlated radio and X-ray luminosities, and the temporal evolution of the radio lightcurves are more consistent with a jet ejection event that faded over time. The steepening of the radio spectral index ($\alpha = -0.2 \pm 0.1$ to -0.6 ± 0.2; optically thick-to-thin synchrotron emission) also suggests a transient jet ejection occurred.

An alternative explanation has been proposed for optically thin radio emission in NSXBs. Russell et al. (2021) observed optically thin emission in the 'high X-ray mode' (equivalent to the soft state) for the ultracompact NSXB 4U 1820$-$30. In this system, the radio emission was attributed to a heavily quenched compact jet with a break frequency (from synchrotron self-absorption) that had moved into the radio regime. Applying a similar explanation to our observations of 1A 1744$-$361 is difficult as the optically thin radio emission in 4U 1820$-$30 is persistent and modestly variable (a factor of ~ 2), whereas 1A 1744$-$361 showed a rapid decay corresponding to, at a minimum, a factor of ~ 100 decrease in the radio flux density. Overall, the decay properties and the fact that the radio flare occurred in the lead-up to a state transition favour the jet ejection scenario.

I looked for evidence of ballistic motion as ejecta have been spatially resolved in several BH (e.g., Hjellming & Rupen, 1995; Bright et al., 2020) and (Z-source) NSXBs (e.g., Fomalont et al., 2001; Miller-Jones et al., 2012;

Motta & Fender, 2019). While I saw no evidence of movement of the radio emission in 1A 1744−361, MeerKAT's beam size ($\sim 6''$) and the decay timescale (undetectable after ~ 40 dy) lead to undetectable proper motion for distances > 1 kpc. Therefore, the absence of proper motion does not provide significant evidence against the jet ejection scenario. I use the maximum flux density (~ 1.2 mJy) and observing frequency ($\sim$1.3 GHz) to derive a distance-scaling relation for the minimum energy ($E_{\rm min}$) of the jet ejecta (following the method described in Fender & Bright, 2019, which assumes incoherent synchrotron radiation and equipartition of energy in particles and magnetic fields in the adiabatically expanding jet ejecta):

$$E_{\rm min} \approx 2.2 \times 10^{37}\,{\rm erg} \left(\frac{d}{8\,{\rm kpc}}\right)^{40/17}, \tag{4.3}$$

where d is the distance to the source.

A minimum energy of $\sim 10^{37}$ erg is consistent with past observations of another flaring atoll-type NSXB (Swift J1858.6-0814, although a state transition was not observed in this source; Rhodes et al., 2022). Measuring an accurate distance (e.g., through the photospheric radial expansion of a Type-I X-ray burst; Kuulkers et al., 2003) will remove the distance ambiguities and confirm (or refute) the apparent consistency between the minimum energy calculations. Using BHXBs as a reference, ejection events in NSXBs have been long predicted to occur during hard-to-soft state transitions of atoll-type sources (Migliari & Fender, 2006; Muñoz-Darias et al., 2014). However, jet ejections have been exclusively observed in Z sources, analogous to ultra-luminous state flaring in BHXBs. Our observations provided the best evidence to date that jet ejections can occur in NSXBs during the standard hard-to-soft state transition.

In contrast to the initial hard-to-soft state transition, we directly observed the soft-to-hard transition as the source returns to quiescence. After MJD 59800 (dashed lines, Fig. 4.1), the X-ray spectra favoured power-law models,

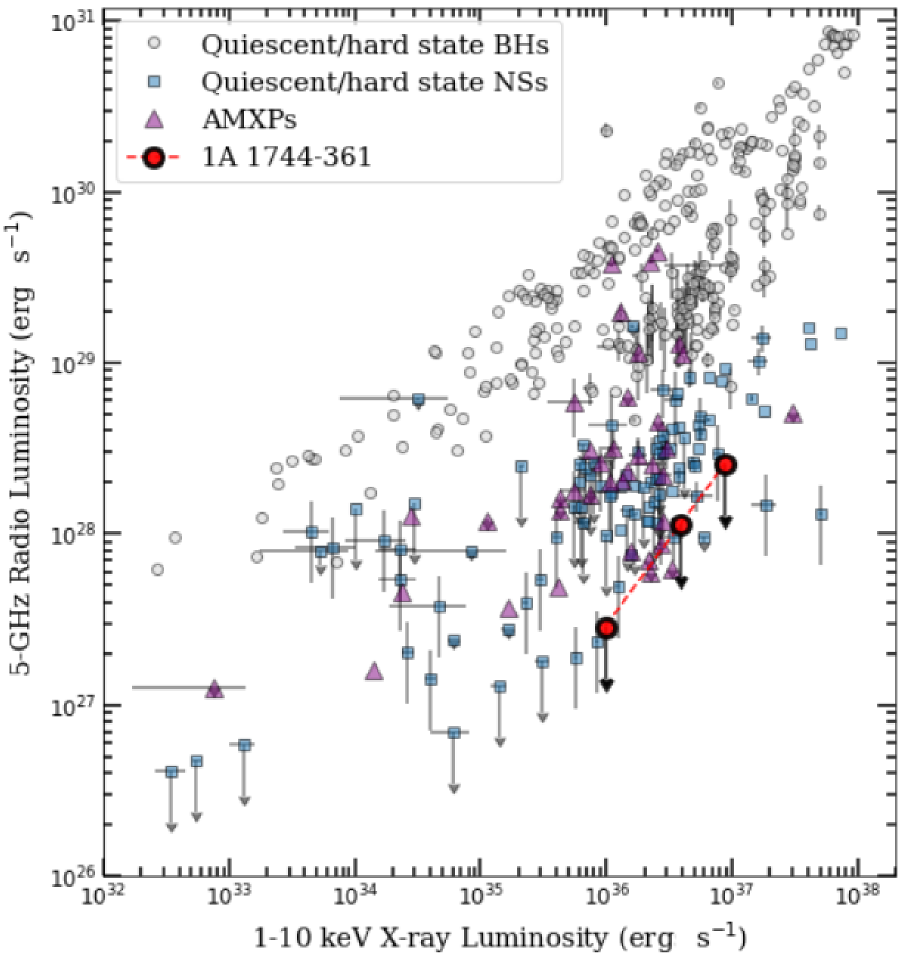

Figure 4.2: 1–10 keV X-ray and 5 GHz radio luminosity plane (i.e., L_R–L_X plane). The plot includes archival hard state BH LMXBs (gray circles), hard state NS LMXBs (blue squares), and accreting millisecond X-ray pulsars (AMXPs; purple triangles) as seen in Figure 4 of van den Eijnden et al. 2021, which contains a large number of sources from Bahramian & Rushton (2022). As 1A 1744−361 was never detected in the hard state, its position on the L_R–L_X plane is limited to upper limits (red circles). Furthermore, due to the unknown distance of 1A 1744−361, the plot includes three data points at distances of 4, 8, and 12 kpc. The upper limits reside within the lower end of radio luminosities for known hard-state NS LMXBs. Future, higher sensitivity observations (and an accurate distance measurement) are critical for determining the source's exact position on the L_R–L_X plane and whether it has an anomalously radio-quiet compact jet. This figure is a modified version of Figure 13 from Ng et al. (2024), which is licensed under CC BY 4.0.

characteristic of the hard state. This change in the best-fit X-ray model coincided with the evolution of the timing parameters, namely an increase in the average (continuum) fractional rms variability (from 3% to 10%; Ng et al., 2024), which is typical of sources that have undergone a soft-to-hard state transition (Hasinger & van der Klis, 1989; van der Klis, 2004). I did not detect any radio re-brightening after the source transitioned back to the hard state at the end of the outburst. Assuming 1A 1744−361 did re-form the compact jet, I used the final stacked image upper limit ($\sim 30\,\mu$Jy) to place constraints on the position of the source on the L_R–L_X diagram as this observation coincided with hard state X-ray emission. I present these results for three assumed distances (4, 8, and 12 kpc) in Figure 4.2. The luminosities at all three distances put 1A 1744−361 among the lower-end of radio luminosities for NSXBs at our measured X-ray luminosity. Higher sensitivity radio monitoring of future outbursts will allow us to directly detect (or better constrain) the radio flux density of the compact jet, thereby determining whether 1A 1744−361 is anomalously radio-quiet for an NSXB.

4.5 Summary and Conclusions

This chapter summarized recent observations of the NSXB 1A 1744−361 throughout its three-month-long 2022 outburst. Here, I focus on the properties of the accretion-jet coupling as the source evolved through different accretion states, omitting much of the X-ray spectral and timing analysis that revealed the source to be capable of transitioning between atoll- and Z-source behaviour (Ng et al., 2024). Looking at the totality of our observations, I develop a coherent picture of the outburst timeline:

1. On $\sim$MJD 59728, 1A 1744−361 entered into an outbursting state, its X-ray flux rapidly increased, and it likely transitioned from the hard to the soft spectral state within $\sim$5 dy of the onset of the outburst,

launching transient jet ejecta.

2. The source remained in the soft state from $\sim$MJD 59728 to 59800; it briefly exhibited properties consistent with a Z-type NSXB. The residual radio emission originated from the fading jet ejecta in the soft state.

3. On $\sim$MJD 59800, the source transitioned back to the island state. I detected no radio emission, and, therefore, if a compact jet formed, its radio emission was below my luminosity detection threshold of $L_{R,5} < 1.2 \times 10^{28}\,(d/8\,\mathrm{kpc})^2\,\mathrm{erg\,s^{-1}}$

4. For the remainder of our monitoring, the X-ray luminosity decreased, dipping below NICER's sensitivity limit on $\sim$MJD 59820, as the source returned to quiescence.

The joint radio and X-ray properties of 1A 1744−361 provide the best observational evidence of jet ejections occurring during the hard-to-soft state transition in NSXB, as I have captured, for the first time, radio flaring preceding a soft-state transition. Future monitoring of the source should prioritize more rapid follow-up to capture the rise of the radio flares and more dense coverage to compare the temporal evolution to the ejecta observed in Z-sources (e.g., Motta & Fender, 2019). Furthermore, very long baseline interferometry could measure any positional shifts due to propagation, thereby confirming the existence of jet ejecta. While such a comprehensive monitoring campaign may be challenging to schedule given the rapidity of the rise and decay of the radio flare, XRBs are known to brighten at optical frequencies before the onset of their X-ray bright outbursts. Appropriate coordination with optical monitoring programs (e.g., Russell et al., 2019a) would make rapid response follow-up possible (or even pre-emptive observations).

Chapter 5

The 2023/2024 Outburst of the BHXB Swift J1727.8−1613

The ThunderKAT program formally ended in 2023 and was succeeded by the MeerKAT open-time program X-KAT, with the latter focusing only on XRBs. Moving from ThunderKAT to X-KAT, our collaboration expanded our observing strategy to include polarimetric monitoring and updated our calibration pipeline. I took the lead in developing the calibration pipeline and imaging procedures for X-KAT polarimetric observing.

Shortly after the start of X-KAT, Swift J1727.8−1613, a (candidate) BHXB, underwent a bright outburst ($\sim 100\,\mathrm{mJy}$) that enabled high precision polarimetry (at the $\sim 0.1\%$ level). The following chapter presents the current state of X-KAT observations of Swift J1727.8−1613. Since monitoring is ongoing, I only include observations until 2024 July 1 (MJD 60492).

5.1 Swift J1727.8−1613

Swift J1727.8−1613 is a new X-ray transient, initially discovered by *Swift*-BAT on 2023 Aug 24 (MJD 60180). After discovery, the source's X-ray flux increased rapidly, reaching $\sim 7\,\mathrm{Crab}$ (at 2–20 keV) within days of the outburst onset. The quick X-ray rise resulted in the source's initial classification as a gamma-ray burst (GRB 230824A; Page et al., 2023). Follow-up X-ray observations refuted a GRB origin, reclassifying (and renaming) the source as the Galactic transient Swift J1727.8−1613 (hereafter Swift J1727, Kennea & Swift Team, 2023). The rapid rise in X-ray luminosity motivated an extensive multi-wavelength monitoring campaign, which identified Swift J1727 as a candidate BHXB (e.g., Bollemeijer et al., 2023a,b; Miller-Jones et al., 2023b,a; Castro-Tirado et al., 2023; Negoro et al., 2023; Nakajima et al., 2023).

Radio follow-up with the VLA localized the source with sub-arcsecond precision, 17h27m43.31($\pm$ 0.04s) $-16°12'19.23(\pm 0.02'')$ (Miller-Jones et al., 2023b). Multi-epoch optical spectroscopy by Mata Sánchez et al. (2024) estimated the properties of the companion star, favouring an early K-Type companion star with an ~ 8 hr orbital period. Furthermore, Mata Sánchez et al. (2024) used the orbital period and observed optical extinction to estimate a distance of 2.7 ± 0.3 kpc, although the authors point out that there are several unaccounted-for systematics that may affect this distance estimate. X-ray spectral modelling revealed relativistically broadened iron lines, suggestive of a high-spin ($a\sim 0.98$), medium-inclination system ($\sim$ 40–50° inclination angle; Draghis et al., 2023; Peng et al., 2024). Modelling of the soft-state X-ray polarisation found a similar inclination angle ($\sim$ 30–50°), but a lower spin ($a\sim 0.87$; Svoboda et al., 2024).

The initial rise in X-ray luminosity occurred as Swift J1727 entered a (bright) hard state, exhibiting a hard-dominated X-ray spectrum (e.g., Peng et al., 2024; Liu et al., 2024), Type-C QPOs (e.g., Bollemeijer et al., 2023a; Mereminskiy et al., 2024) and optically-thick radio emission consistent with a compact jet (Miller-Jones et al., 2023b). As Swift J1727 evolved (through the hard intermediate state), the Type-C QPO frequency monotonically increased (Bollemeijer et al., 2023a), before disappearing on 2024 October

5 (MJD 60222; Bollemeijer et al., 2023b), suggesting a transition to the intermediate of soft-state had occurred. Alongside the X-ray transition, Swift J1727 exhibited radio quenching, followed by bright radio flaring, consistent with the disruption of the compact jet and the launching of jet ejecta (Miller-Jones et al., 2023a). Subsequent X-ray observations, noting variability in the hardness ratio, were suggestive of multiple returns to the intermediate states (e.g., Yu, 2023). Swift J1727 took ~ 6 months to transition back to the hard state after its X-ray luminosity decreased by more than two orders of magnitude (on $\sim$ MJD 60385, 2024 March 15; Podgorny et al., 2024; Russell et al., 2024). As of 2024 July 1 (MJD 60481), the X-ray properties of Swift J1727 suggest it has returned to quiescence.

In this chapter, I report the MeerKAT view of the outburst, which includes extreme polarisation variability, anomalously optically thin and radio-quiet compact jets, and multiple spatially resolved jet ejecta. I connect the radio properties to the simultaneous evolution of the accretion flow by including multi-facility X-ray observations. The rest of the chapter is structured as follows: Section 5.2 describes my observation and analysis routine; Section 5.3 presents the results and a preliminary discussion of the observed behaviour; and Section 5.4 details X-KAT plans for the source, focusing on combining the totality of our (multi-facility) observations.

5.2 Observations and Analysis

5.2.1 MeerKAT

We observed Swift J1727 as a part of X-KAT (Proposal ID: SCI-20230907-RF-01). Our monitoring campaign began on 2023 August 27 (MJD 60183), three days after the initial detection, and consists of 52 epochs separated by an approximately weekly cadence. We used the L-band (~ 1.28 GHz;

856 MHz un-flagged bandwidth) receiver for 47 observations. Each L-band observation consisted of a single scan of either 15 or 30 minutes on-source, flanked by two 2-minute scans of a nearby gain calibrator (J1733−1304). We included two 5-minute scans of the J1939−6342 (PKS B1934−638) at the beginning and end of each observation for parallel-hand delay, bandpass, leakage, and flux scale calibration, as well as a single 10-minute scan of 3C286 (J1331+3030) for polarisation angle calibration. The observation on 2023 October 16 (MJD 60233) did not include a scan of 3C286; thus, we cannot measure a polarisation angle. In addition to our weekly L-band monitoring, we acquired five high-angular resolution epochs using MeerKAT's S-band receivers. One of the S-band observations used the S2 ($\sim$2.62 GHz; 875 MHz un-flagged bandwidth) sub-band, and the other four used the S4 ($\sim$3.06 GHz; 875 MHz un-flagged bandwidth) sub-band. Our S-band monitoring followed the L-band observing strategy.

The previous chapters used OXKAT, a pipeline designed to automate the processing of total intensity (i.e., Stokes I) observations (as described in, Heywood et al., 2022). While adopting the ThunderKAT monitoring strategy, X-KAT expanded its scope to include full polarisation observations, necessitating an updated calibration pipeline. Most archival polarisation calibration routines for linear feed interferometers require broad parallactic angle coverage, making them inapplicable for sub-hour observations. Fortunately, recent investigations (e.g., EVLA memo 209[1]; EVLA memo 219[2]; Taylor & Legodi, 2024; Hugo & Perley, 2024) have provided sufficient information to develop CASA-based calibration routines for short-track observations. Using these studies as a reference and through extensive testing, I led the development of an OXKAT-inspired pipeline capable of processing full polarisation (i.e., I, Q, U, and V) MeerKAT observations;

[1] https://library.nrao.edu/public/memos/evla/EVLAM_209.pdf
[2] https://library.nrao.edu/public/memos/evla/EVLAM_219.pdf

the POLKAT[3] pipeline. Our X-KAT targets are point sources positioned within 30 arcseconds of the pointing centre (i.e., $\mathcal{A} > 0.99$), making our observations unaffected by *off-axis* instrumental polarisation during imaging (see, e.g., de Villiers, 2023). Therefore, the limiting factor for the accuracy of the polarisation properties is the quality of the flagging and reference calibration.

While POLKAT was designed to analyse XRB data, it is source-agnostic and, thus, can be applied to any MeerKAT data and, in fact, snapshots observations with any linear-feed interferometer (e.g., the Atacama Large Millimeter Array; ALMA). The three major caveats are: (i) the primary/bandpass calibrator **must** be unpolarised; (ii) the analysis scripts (RMSYNTH as described below) are only valid if the flux density distribution can be well-described with an ensemble of point sources; (iii) MeerKAT is known to have significant off-axis instrumental polarisation, increasing rapidly with distance from the phase centre. Therefore, special care must be taken when interpreting the polarised signals in off-axis sources. In the text below, I describe how POLKAT applies to the Swift J1727 observations, calling out choices that users might need to make to analyse different sources.

Flagging and Reference Calibration (1GC)

By default, the MeerKAT visibility products "mislabel" the linear feeds compared to the standard convention (e.g., the 'X'-feed is labelled 'Y'; EVLA Memo 219). If uncorrected, the mislabelling reverses the signs of Stokes Q and V. Therefore, the first step for polarisation calibration (and thus POLKAT) is a label correction. Following the label correction, any frequency channels corrupted by persistent RFI (e.g., from satellites) must be removed manually. POLKAT removes stochastic RFI with the CASA auto-flaggers TFCROP and RFLAG; we use the default flag parameters for the

[3] `https://github.com/AKHughes1994/polkat`

auto-flaggers. OXKAT flags the persistent RFI from short baselines (<600 m) exclusively, as they were found to be most significantly affected. I found that while similarly selective flagging only targeting short baselines preserves some bandwidth, this flagging could induce artificial polarisation signals on the order of $\sim 0.5\%$. In POLKAT, we sacrifice the bandwidth, flagging the frequencies with known, persistent RFI from all baselines.

Reference calibration begins by initializing a (per-channel) model of the primary, J1939−6342. For L-band observations, POLKAT initialize full sky models[4] with CRYSTALBALL(Serra et al., 2022), including the visibilities of background sources in the calibrator field. Full-sky models are unavailable at S-band, but the higher frequency and smaller field of view attenuate the effects from background sources. For S-band, POLKAT uses the CASA task setjy to initialize a primary model using the Stevens-Reynolds 2016 standard. With the primary model, POLKAT solves for the parallel-hand delay and bandpass solutions with the CASA tasks GAINCAL (`gaintype='K'`) and BANDPASS (`gaintype='B'`), respectively. As J1939−6342 is unpolarised, POLKAT also solves for the on-axis instrumental polarisation (i.e., leakage term) with the CASA task POLCAL (`gaintype='Df'`). Currently, POLKAT will automatically identify and initialize models for the two recommended MeerKAT primaries (J1939−6342 and J0408−6545) and will accept manual models for any non-standard primaries (following the standard SETJY format).

After applying the primary solutions to the secondary (J1733−1304, in this case) and polarisation angle calibrator (3C286, in this case), POLKAT derives per-scan complex gain terms (using GAINCAL). The complex gain amplitudes are feed-averaged (`gaintype='T'`), whereas the phases are solved per-feed (`gaintype='G'`). This change was necessary as linear polarised intensities (specifically Stokes Q) are a function of the parallel-hand visibili-

[4]`https://github.com/ska-sa/katsdpcal/tree/master/katsdpcal/conf/sky_models`

ties for linear-feed interferometers. By default, GAINCAL (`gaintype='G'`) assumes an unpolarised point source calibrator, equalizing the 'X'- and 'Y'-feed amplitudes and removing any intrinsic polarisation. Most secondary calibrators are polarised at the $\sim 2\%$ level, and equalized gain terms (using `gaintype='G'`) will induce a similar level of artificial polarisation in the target. Feed-averaged solutions using `gaintype='T'` correct any average temporal drifts in amplitude without modifying the relative amplitudes in each feed, preserving the target's intrinsic polarisation properties. POLKAT corrects these relative amplitudes during bandpass calibration by using `gaintype='T'` with GAINCAL.

The complex gain amplitudes are scaled according to the values derived by the primary, using the CASA task FLUXSCALE. Finally, POLKAT derives the cross-hand delay (GAINCAL, `gaintype='KCROSS'`) and cross-hand phase (POCAL, `gaintype='Xf'`) terms using the polarisation angle calibrator. For each calibration solution, the values closest in time are interpolated onto the target (Swift J1727, in this case), and POLKAT applies the auto-flagging routines TFCROP, RFLAG and TRICOLOUR (Hugo et al., 2022) to flag the calibrated target visibilities. The target visibilities are then plotted using SHADEMS (Smirnov et al., 2022), and visually inspected by the user (e.g., me, in this case) to determine if additional flagging is required. Once all flagging is complete, POLKAT proceeds to imaging and self-calibration.

For linear-feed interferometers, a linearly polarised, circularly unpolarised calibrator can derive cross-hand phase/delay solutions without needing a polarisation model (EVLA Memo 219). In this case, the cross-hand phase is given by $\arctan(-V'/U')$, where U' and V' are the Stokes parameters measured after applying all other calibration terms (i.e., excluding cross-hand delays and phases). However, there exists a $\pm\pi$ degeneracy in the cross-hand phase due to $-V'/U' \equiv -(-V')/(-U')$. The degeneracy requires *a priori* knowledge of the calibrator to pick the correct phase. For

3C286, POLCAL resolves the degeneracy for a set of quasi-arbitrary polarisation models. Using SETJY, we initialize a point-source model with Stokes I, Q, U, and V intensities of 1.0, 0.0, 0.5, and 0.0 Jy, respectively (`fluxdensity=[1.0,0.0,0.5,0.0]`). We tested other initialization values and found that any model where $V = 0.0$ and $U > Q$ would derive accurate cross-hand phase solutions. This model-agnosticism is extremely useful because a calibrator with well-constrained polarisation properties (like 3C286) functions as a built-in 'check source' to determine calibration accuracy.

Self-Calibration and Imaging (2GC)

Imaging and self-calibration adopt a multi-step approach. We generated a preliminary, un-masked, Stokes I image using WSCLEAN (v3.4; Offringa et al., 2014). Using the preliminary image, POLKAT creates a deconvolution mask with BREIZORRO[5]. POLKAT then performs a masked deconvolution of all four Stokes parameters, using the resulting images to create a sky model of the target. Using the sky model, POLKAT performs phase-only (direction-independent) self-calibration of the target complex gains using GAINCAL with solution intervals of 32 seconds; the user can modify this interval. Following self-calibration, POLKAT performs a final round of masked deconvolution. I adopt the self-calibrated images as my final data products for Swift J1727.

The self-calibrated images consist of 16 channelised images (evenly spaced in frequency) and a single frequency-averaged "Multi-Frequency-Synthesis" (MFS) image. I made two sets of images for each epoch, one maximizing sensitivity and another minimizing the size of the restoring beam (leading to finer angular resolution). The sensitivity-maximizing images used a Briggs' weighting with a robustness parameter of 0. I minimize the restoring beam size with uniform weighting, using these images to measure the proper

[5]https://github.com/ratt-ru/breizorro/tree/main/breizorro

motions of jet ejecta.

In addition to the target imaging, I made channelised and MFS images of the calibrators used to derive leakage (J1939−6342) and cross-hand (3C286) terms. I used these images to quantify the systematic errors from our reference calibration (see below). I did not self-calibrate the calibrator images, as the self-calibrated properties would be unrepresentative of the gain errors applied to the source.

Polarisation Property Extraction (RMSYNTH)

I extracted the source properties in each observation and for each Stokes parameter using the CASA task IMFIT, modelling the flux density distributions with PSF-shaped (i.e., synthesized beam-shaped) elliptical Gaussians. Where there were resolved (or partially resolved) jet ejecta, I used a multi-Gaussian fit for observations. In all cases, I quantified the 1σ uncertainties on the flux densities using the local image-plane rms noise. Using the CASA task IMSTAT, I extracted the rms noise from an emission-free region centred near the source(s), with an area equal to ~ 100 PSFs. For observations where Q, U, or V had signal-to-noise values $< 5\sigma$, I fixed their positions to the Stokes I position, performing forced aperture photometry. I measure intra-band spectral indices using the channelised images. Additionally, I broke the observing bandwidth in half for each observation to measure two point-spectral indices; I used the two-point spectral indices for lower signal-to-noise observations (where the target flux density is too low for significant detections in one-sixteenth of the bandwidth). However, intra-band spectral indices are known to bias toward flatness (i.e., $\alpha \sim 0.0$) for signal-to-noise ratios ≤ 30. I quantified astrometric errors using the method detailed in Chapter 6.

I created a linear polarisation intensity image for each pair of Q/U images ($P = \sqrt{Q^2 + U^2}$) and applied the above routine to extract the linearly polarised flux density. As linearly polarised flux is positive definite, it follows a Rice distribution (Rice, 1945). The measured polarised flux (P) is positively biased from intrinsic polarised flux (P_0) according to

$$P_0^2 = P^2 - \sigma_{QU}^2 \tag{5.1}$$

and

$$\sigma_{QU}^2 = A_Q\sigma_Q^2 + (1 - A_Q)\sigma_U^2, \tag{5.2}$$

where σ_Q and σ_U are the rms noise in the Q and U images, respectively, and $A_Q \leq 1$ (Vaillancourt, 2006; Hales et al., 2012). Following Hales et al. (2012), I conservatively adopt $A_Q = 0.8$ ($A_Q = 0.2$) for epochs where $\sigma_Q \geq \sigma_U$ ($\sigma_Q < \sigma_U$), but note that the measured rms values are consistent with the ideal case (i.e., $\sigma_Q \sim \sigma_U$). I only accept high-significance polarised detections $P/\sigma_{QU} \geq 4$. When that condition is not met, I solve for 3σ (99.73%) upper-limits on P_0 following the prescription from Vaillancourt (2006). I calculate the linear and circular polarisation fractions using their standard definitions: $m_l \equiv P_0/I$ and $m_c \equiv V/I$, and I approximate their uncertainties through Gaussian error propagation.

Owing to their long wavelengths, radio waves are susceptible to a propagation-induced rotation of their EVPA as they travel through magneto-ionic media like the ISM. This effect, *Faraday Rotation*, is as a linear function in λ^2-space:

$$\mathrm{EVPA} = \mathrm{EVPA}_0 + \mathrm{RM} \cdot \lambda^2, \tag{5.3}$$

where EVPA_0 is the intrinsic polarisation angle, and the magnitude of the rotation is quantified by the rotation measure (RM). The RM is a path integral related to the particle number density (n) and magnetic field strength parallel to our line of sight ($B_{||}$):

$$\mathrm{RM} = \frac{q^3}{2\pi m^2 c^4} \int_{\mathrm{source}}^{\mathrm{observer}} n B_{||} \mathrm{d}l, \tag{5.4}$$

where q and m are the charges and the mass of the particles in the plasma. If we assume Faraday rotation originating from a population of electrons, we can alternately express the rotation measure as,

$$\mathrm{RM} = \left[812 \int_{\mathrm{source}}^{\mathrm{observer}} n_e B_{||} \mathrm{d}l \right] \ \mathrm{rad\,m^{-2}}, \tag{5.5}$$

where n_e, $B_{||}$, and dl are in units of cm^{-3}, μG, and kpc, respectively. As demonstrated later, Faraday Rotation is not a purely corrupting effect; identifying the local RM component provides invaluable insight into the local environment.

To extract the linear polarisation angle and rotation measures, I used the RM-TOOLS package (Purcell et al., 2020), developed by the Canadian Initiative for Radio Astronomy Data Analysis (CIRADA). RM-TOOLS is based on the concept of *Rotation Measure Synthesis*. Here, I briefly outline the basics of RM synthesis. RM synthesis is based on the fact that the complex (linearly) polarised flux density,

$$\widetilde{P}(\lambda^2) = Q(\lambda^2) + iU(\lambda^2) \tag{5.6}$$

is related to a quantity known as the 'Faraday dispersion function' (FDF), $\widetilde{F}(\phi)$, through the Fourier relation:

$$\widetilde{P}(\lambda^2) = \int_{-\infty}^{+\infty} \widetilde{F}(\phi_f) e^{2i\phi_f(\lambda^2 - \lambda_0^2)} d\phi_f \tag{5.7}$$

$$\widetilde{F}(\phi_f) = \int_{-\infty}^{+\infty} \widetilde{P}(\lambda^2) e^{-2i\phi_f(\lambda^2 - \lambda_0^2)} d\lambda^2, \tag{5.8}$$

where ϕ_f is known as the 'Faraday depth' and is equivalent to a rotation measure. The Faraday dispersion function can be thought of as the complex polarised flux density that corresponds to a specific Faraday depth (and thus rotation measure). Conversely, $\widetilde{P}(\lambda^2)$ is the total polarised flux density (at some λ^2) considering emission from all Faraday depths.

The basic workflow of RM synthesis begins with extracting complex polarised flux densities from multiple λ^2 channels (I used the sixteen channelised images outputted by wsclean). Given the measured $\widetilde{P}(\lambda^2)$ values, the next step is to solve for $\widetilde{F}(\phi_f)$ for a range of *trial* Faraday depths. Each trial ϕ_f de-rotates the complex polarisation flux density to a common λ_0^2, where λ_0^2 is often taken to be the $1/\sigma_{QU}^2$-weighted average of the λ^2 channels. Faraday depths corresponding to the intrinsic rotation measure(s) (i.e., ϕ_f = RM) will coherently sum the (de-rotated) polarised intensities from each λ^2 channel, manifesting as peaks in the Faraday dispersion function.

The value of the Faraday dispersion function at the intrinsic RM is related to the EVPA at λ_0^2 according to

$$\mathrm{EVPA}_{\lambda_0} = \frac{1}{2}\arctan\left(\frac{\mathrm{Im}[\widetilde{F}(\mathrm{RM})]}{\mathrm{Re}[\widetilde{F}(\mathrm{RM})]}\right). \tag{5.9}$$

Thus,

$$\mathrm{EVPA}_0 = \mathrm{EVPA}_{\lambda_0} - \mathrm{RM}\cdot\lambda_0^2. \tag{5.10}$$

Like synthesis imaging, the finite sampling in λ^2-space results in non-physical structures in the FDF; more precisely, the true FDF is convolved with the 'Rotation Measure Transfer Function' (RMTF), which is an analogue to the dirty beam. As a result, most RM synthesis packages will utilize CLEAN to remove these sampling effects. Furthermore, the smallest and largest spacing between the λ^2 channels determines the largest and smallest detectable rotation measures. The latter — similar to the synthesized beam size — determines the 'Faraday resolution' as signals in the Faraday dispersion function will appear Gaussian-like with an FWHM of $\sim \frac{2\sqrt{3}}{\Delta\lambda^2}$, where $\Delta\lambda^2$ is the maximum difference between λ^2-channels (to demonstrate this I have included a sample FDF and RMTF in Figure 5.1). It should be noted that the lower-frequency observations of X-KAT have a clear advantage over historical monitoring strategies that have focused on higher-frequency radio observations (i.e., $\gtrsim 6\,\mathrm{GHz}$). Lower-frequency observations will have broader λ^2-bandwidths at a fixed frequency bandwidth and, thus, more precisely measure the rotation measure and polarisation angles. The transient rotation-measure flare discussed later in the chapter would be undetectable at higher frequencies with current facilities.

As an illustrative example, consider a line of sight containing two equally polarised sources with RM values of 6 and $7\,\mathrm{rad\,m^{-2}}$: (i) If $\Delta\lambda^2$ corresponds to a FWHM of $0.1\,\mathrm{rad\,m^{-2}}$ the FDF will have two clear peaks at 6 and $7\,\mathrm{rad\,m^{-2}}$; (ii) If $\Delta\lambda^2$ corresponds to a FWHM of $100\,\mathrm{rad\,m^{-2}}$ the two sources will blend into a single component centred at $6.5\,\mathrm{rad\,m^{-2}}$. For a more complete description of Rotation Measure Synthesis, we direct readers to, e.g., Burn (1966), Brentjens & de Bruyn (2005), George et al. (2012),Macquart et al. (2012), and Hales et al. (2012).

Accurate measurement with rotation measure synthesis requires that $P/\sigma_{QU} \geq 7$ (Hales et al., 2012; Macquart et al., 2012). In the latter half of our observations (after 2024 January 16; MJD 60325), for observations with $4 \leq P/\sigma_{QU} < 7$, we calculate the observed EVPA using the Q and U flux densities in the middle of the observing band:

$$\mathrm{EVPA} = \frac{1}{2}\arctan\left(\frac{U}{Q}\right), \tag{5.11}$$

de-rotating the observed polarisation angle according to Equation (5.10), where λ_0 corresponds to the central frequency of the MFS images (e.g.,

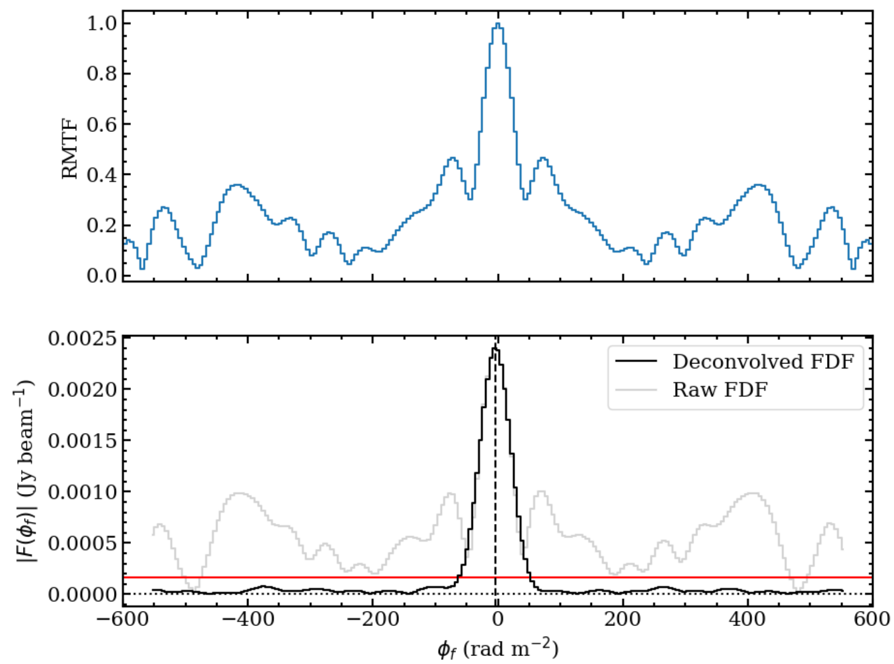

Figure 5.1: Sample Faraday Dispersion Function (FDF) and Rotation Measure Transfer Function from the observations of Swift J1727 taken on 2023 September 23 (MJD 60210). (top panel) Rotation Measure Transfer Function — the analogue to the dirty beam in synthesis imaging. (bottom panel) The deconvolved and 'raw' FDF (i.e., including the finite sampling effects) are shown in black and grey, respectively. The horizontal red line corresponds 5× the rms noise of the FDF, highlighting the significance of structures in the deconvolved FDF. The FDF has a single Gaussian-like feature, corresponding to the rotation measure of the source: $\mathrm{RM} = 1.4 \pm 0.4\,\mathrm{rad\,m^{-2}}$; here, the reported error is purely statistical (i.e., does not include the systematic contribution).

$\sim$ 1.28 GHz for L-band). During these epochs, I adopted the average rotation measure $\langle RM \rangle \sim 1.4\,\mathrm{rad\,m^{-2}}$ (assumed to be due to ISM interactions) from our observations with $P/\sigma_{QU} \geq 7$. I exclude earlier observations, as Swift J1727 showed signs of RM variability, and therefore, I cannot be confident that utilizing $\langle RM \rangle$ is appropriate. Lastly, each RM measurement has a partial contribution from the Ionosphere; I predict the Ionospheric contributions using ALBUS[6], subtracting the predicted Ionospheric RM from my measured value of RM.

I assessed the systematic errors in our data by extracting the polarisation properties from our calibrator fields. In Fig. 5.2, I show the absolute deviations from the expected values of EVPA_0 and RM for 3C286 (using the model from; Hugo & Perley, 2024); 3C286 is known to be circularly unpolarised so $m_c \equiv 0$. Furthermore, I calculate the total observed polarisation fraction, $m_t = \sqrt{Q^2 + U^2 + V^2}/I$, for our unpolarised leakage calibrator, J1939−6342, adopting $m_t = 0$ as our expected value — any residual polarization is (conservatively) attributed to systematic calibration errors. We adopt the rms of the deviations as systematic errors. For EVPA_0 and RM, the systematics are 1.1° and $0.9\,\mathrm{rad\,m^{-2}}$, respectively, from the 3C286 measurements. For the polarisation fraction, I adopt a systematic error of $\sqrt{(m_c^{\mathrm{3C286}})^2 + (m_t^{\mathrm{J1939-6342}})^2} \sim 0.07\%$. When reporting uncertainties in the target polarisation properties, I add these systematics to the measured errors in quadrature. Table 5.1 summarizes the calibration, imaging and analysis steps, comparing POLKAT to OXKAT. The MeerKAT-derived radio properties are presented in Tables 5.3–5.5.

5.2.2 *Swift*-XRT

X-KAT acquired quasi-simultaneous X-ray observations of Swift J1727 with the *Swift*-XRT as a part of the *Swift*KAT program (i.e., within $\sim$3 days of

[6] https://github.com/twillis449/ALBUS_ionosphere

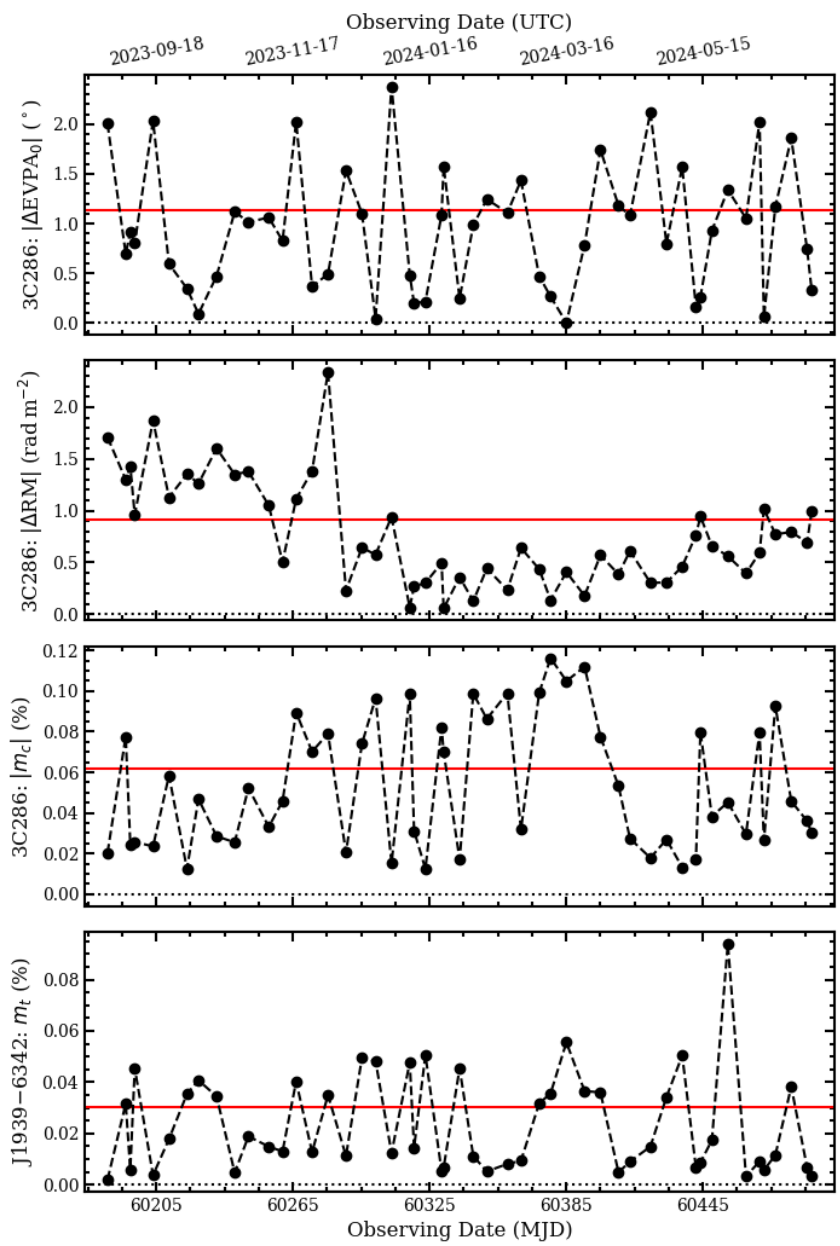

Figure 5.2: Absolute deviations in the polarisation properties when compared to the model values of calibrators: (top panel) the intrinsic linear polarisation angle of 3C286; (second panel) the rotation measure of 3C286; (third panel) the circular polarisation fraction of 3C286; (bottom panel) the total polarisation fraction of J1939−630. The solid (horizontal) red lines show the adopted systematic errors, which were set by the root-mean-square in each panel.

Table 5.1: Brief overview of how POLKAT implements polarisation calibration, building off of OXKAT Stokes I calibration, imaging, and analysis

Step	OXKAT	POLKAT
User Option Configuration	• Modify `config.py` (User)	• Modify extended `config.py` (User)
Reference Calibration and Flagging (1GC)	 • Deterministic flagging (FLAGDATA*; `mode='manual'`; $uv < 600\,\mathrm{m}$) • Stochastic flagging (FLAGDATA*; `mode='rflag'`, `mode='tfcrop'`) • Initialize primary model (SETJY*) • Parallel-hand delay (GAINCAL*; `gaintype='K'`) • Bandpass (BANDPASS*; `bandtype='B'`) • Complex gain: amplitude and phase (GAINCAL*; `gaintype='G'`) • Absolute flux scaling (FLUXSCALE*)	• Correct 'X'- and 'Y'-feed labelling (`correct_parang.py`[a]) • Deterministic flagging (FLAGDATA*; `mode='manual'`; all uv values) • Stochastic flagging (FLAGDATA*; `mode='rflag'`, `mode='tfcrop'`) • Initialize primary model (CRYSTALBALL for L-band and SETJY* for S-band) • Initialize (quasi-arbitrary) polarisation model (SETJY*) • Parallel-hand delay (GAINCAL*; `gaintype='K'`) • Bandpass (BANDPASS*; `bandtype='B'`) • Complex gain: amplitude (GAINCAL*; `gaintype='T'`) • Complex gain: phase (GAINCAL*; `gaintype='G'`) • Leakage (POLCAL*; `gaintype='Df'`) • Cross-hand delay (GAINCAL*; `gaintype='KCROSS'`) • Cross-hand phase (POLCAL*, `gaintype='Xf'`) • Absolute flux scaling (FLUXSCALE*) • Image calibrators (WSCLEAN; IQUV *optional*)

Continued on next page

Table 5.1 continued

Step	OXKAT	POLKAT
Self-Calibration and Imaging (2GC†)	• Target (Stochastic) flagging (TRICOLOUR) • Un-masked imaging (WSCLEAN; I) • Create deconvolution mask (`pyMakeMask.py`[b]) • Masked imaging (WSCLEAN; I) • Phase-only self-calibration (CubiCal) • Masked imaging of self-calibrated visibilities (WSCLEAN; I)	• Target (Stochastic) flagging (TRICOLOUR and FLAGDATA*) • Un-masked imaging (WSCLEAN; I) • Create deconvolution mask (BREIZORRO[c]) • Masked imaging (WSCLEAN; IQUV) • Phase-only self-calibration (GAINCAL*; `gaintype='G'`) • Masked imaging of self-calibrated visibilities (WSCLEAN; IQUV) • High resolution imaging of self-calibrated visibilities (WSCLEAN; I *optional*)
Polarisation Property Extraction (RM-SYNTH)		• Measure target polarisation properties from MFS and channelised images (`RMSYNTH_01_extract_fluxes.py`‡) • Measure calibrator polarisation properties from MFS and channelised images (`RMSYNTH_01B_systematics.py`‡ *optional*) • Run RM Synthesis on target and (*optional*) calibrators (RM-TOOLS) • Estimate ionospheric RM contribution (ALBUS[d])

Notes.
* Indicates a CASA task.
† In POLKAT the 2GC step has combined two of the OXKAT steps (FLAG and 2GC).
‡ Indicates a custom POLKAT analysis script.

References.

[a] Adapted from `https://github.com/bennahugo/LunaticPolarimetry`

CRYSTALBALL (Serra et al., 2022)

WSCLEAN (Offringa et al., 2014)

TRICOLOUR (Hugo et al., 2022)

[b]`https://github.com/IanHeywood/oxkat/blob/master/tools/pyMakeMask.py`

[c]`https://github.com/ratt-ru/breizorro`

RM-TOOLS (Purcell et al., 2020)

[d]`https://github.com/twillis449/ALBUS_ionosphere`

a MeerKAT epoch). We began our *Swift*-XRT monitoring on 2024 February 3 (MJD 60343) after the target was no longer in a Sun-constrained position. I note that while there are *Swift*-XRT observations between 27 August 2023 and 27 October 2023 (MJD 60183–60245) before the source became Sun-contained, these were taken during the peak of the outburst when the target's count rate was ($> 2000\,\mathrm{count\,s^{-1}}$), and thus, these observations are severely affected by photon pile-up. At early times, I characterize the X-ray properties of Swift J1727 using its daily MAXI/GSC light curves and *Swift*-BAT(Section 5.2.3), omitting the early *Swift*-XRT observations from our X-ray spectral analysis. During the outburst, we acquired 22 epochs of *Swift-XRT* observations (target ID: 89766/16584), ending on 2024 June 2 (MJD 60463). *Swift* experienced a severe malfunction on 2024 March 15 (MJD 60384), preventing observations until 2024 April 3 (MJD 60403); unfortunately, this malfunction occurred the day the source was reported to have transitioned back into the hard state.

Due to the high count rates, we observed the source in Windowed Timing (WT) at the start of our monitoring. We transitioned to Photon Counting (PC) mode on 2024 May 10 (MJD 60440) when the count rate dropped to $\lesssim 1\,\mathrm{count\ s^{-1}}$, as per standard procedures for *Swift*-XRT. I extracted the X-ray spectral files using SWIFTTOOLS; however, I performed spectral modelling manually with the HEASOFT package (version 6.25). For high-count observations ($\geq 300\,\mathrm{counts}$), I used a modified GRPPHA[7] script to bin the spectra on 25-count intervals, and used χ^2 as the fitting statistic. For low-count observations ($< 300\,\mathrm{counts}$), I used Cash statistics (i.e., CSTAT; Cash, 1979) and single-count binning intervals.

With the HEASOFT program XSPEC (Arnaud, 1996b), I fit each spectrum with two models, an absorbed power-law model (`tbabs` $\times$ `pegpwrlw`) and an absorbed disc blackbody (`tbabs` $\times$ `diskbb`). While these simple

[7] as described in Chapter 4

models are insensitive to fine spectral structure, they are sufficient for discrimination between hard and soft-state emissions. I modelled interstellar absorption with `tbabs` using an equivalent hydrogen column density (N_H) following the abundances from Wilms et al. (2000). Initially, I fit each spectrum individually, allowing N_H to vary epoch-to-epoch. I found no significant variation in N_H and, as a result, I then adopted the (variance-weighted) average value of $N_H = 1.99 \pm 0.03 \times 10^{21}\,\mathrm{cm}^{-2}$, fixed the hydrogen column density, and re-modelled each spectrum. This value of N_H was consistent with past X-ray modelling (e.g. Peng et al., 2024). It is also consistent with the expected column density through a line-of-sight towards Swift J1727 extending throughout the entire Galaxy ($1.94 \times 10^{21}\,\mathrm{cm}^{-2}$; Dickey & Lockman, 1990). Given the Galactic latitude of the source ($\sim 10.2°$) and assumed distance, most of the Galactic absorption should occur between Earth and the source; therefore, the consistency between the total Galactic absorption and the measured value is expected. The last four observations had ~ 5 target counts and were omitted from the abovementioned modelling. Instead, recognizing the target was already in a hard state, I measured the X-ray flux with the absorbed power-law model, fixing the photon index to its canonical value, $\Gamma \equiv 1.7$. The *Swift*-XRT monitoring and spectral parameters are presented in Table 5.6 and 5.7. Fixing the photon index decreases the error on the statistical error on the X-ray flux but results in an unmodelled systematic (i.e., if the true Γ deviates from 1.7). Regardless, fixing Γ is necessary to extract constraining flux from the low-count observations, and, given the prior observations, it is unlikely Γ deviates significantly from its canonical value. The quoted uncertainties on the X-ray parameters use the standard 90% confidence intervals.

5.2.3 MAXI/GSC and *Swift*-BAT

I utilize the publicly available[8] daily X-ray observations from *Swift*-BAT and MAXI/GSC to track the X-ray luminosity and calculate the X-ray HR. MAXI/GSC is sensitive to photon energies of 2–20 keV energy, and the standard data products include 2–4, 4–10, 10–20, and 2–20 keV light curves; *Swift*-BAT is sensitive to 15–50 keV. I apply adaptively binned averaging for both instruments and each energy range, grouping adjacent observations with $< 3\sigma$ significance; here, 3σ corresponds to a MAXI/GSC rate $3\times$ the quoted error. When averaging, I increase the bin size by a day until the averaged value is $\geq 3\sigma$ (adopting Gaussian statistics), or the subsequent daily observation meets one of the following conditions:

1. the detection is $\geq 3\sigma$;
2. the observing date is separated by > 1week; or
3. there are no more observations.

If any of these conditions are met, I calculate a 3σ upper limit on the adaptively sized bin using the error on the average from the current binning — these conditions avoid excessive averaging. To make use of the full MAXI/GSC energy range, I use the on-demand product service to extract the rates in the 2–6 ($R_{2\text{–}6\,\mathrm{keV}}$) and 6–20 keV ($R_{6\text{–}20\,\mathrm{keV}}$) sub-bands. This choice of sub-bands ensures a constant fractional bandwidth in the hard and soft X-ray bands. I then calculate the hardness ratio:

$$\mathrm{HR} \equiv \frac{R_{6\text{–}20\,\mathrm{keV}}}{R_{2\text{–}6\,\mathrm{keV}}}. \tag{5.12}$$

For the HR, I simultaneously enforced the 3σ detection threshold for both sub-bands.

[8] Standard data analysis steps, such as background subtraction, are performed by the MAXI/GSC and *Swift*-BAT teams. Aside from the averaging, I do not manipulate the data in any other way.

5.3 Results and Discussion

5.3.1 X-ray Evolution

I show the HID evolution of Swift J1727, as seen by MAXI/GSC, in Figure. 5.3. Comparing the evolution of Swift J1727 to the schematic shown in Chapter 1, it is evident that Swift J1727 followed the standard outburst track, as described in Chapter 1, with no evidence of an ultra-luminous state. The variability in the hardness ratio during the source's descent down the soft-state is consistent with the picture presented by Yu (2023), suggesting multiple excursions to (and from) intermediate states and, thus, multiple opportunities for jet ejections under the standard paradigm.

The rapid rise in X-ray luminosity reached its peak $\lesssim 1$ week after the onset of the outburst with a hard-dominated X-ray spectrum (HR ~ 0.45). Following the rise, the X-ray emission remained bright, softening as the source progressed through the intermediate states, reaching an HR ~ 0.1 when Bollemeijer et al. (2023b) reported the transition to the soft (or soft-intermediate) state had occurred (October 5, blue star, Fig. 5.3). Over the next ~ 6 months, the X-ray luminosity decreased by two orders of magnitude while maintaining a soft but variable (HR ~ 0.05–0.15) hardness ratio. Due to the significant drop in X-ray luminosity, the hardness ratio during the soft-to-hard transition and subsequent return to quiescence has substantial errors. Regardless, there is a clear hardening of the X-ray spectra coinciding with the reported date of the transition back to the hard state (2024 March 15, yellow star, Fig. 5.3; Podgorny et al., 2024).

The MAXI/GSC and *Swift*-BAT X-ray light curves (see Fig. 5.4) showed, in addition to the overarching decrease in X-ray luminosity, short time-scale variability. This variability is more significant at higher energies (i.e., MAXI/GSC 4–20 keV and *Swift*-BAT 15–50 keV). MAXI/GSC observed

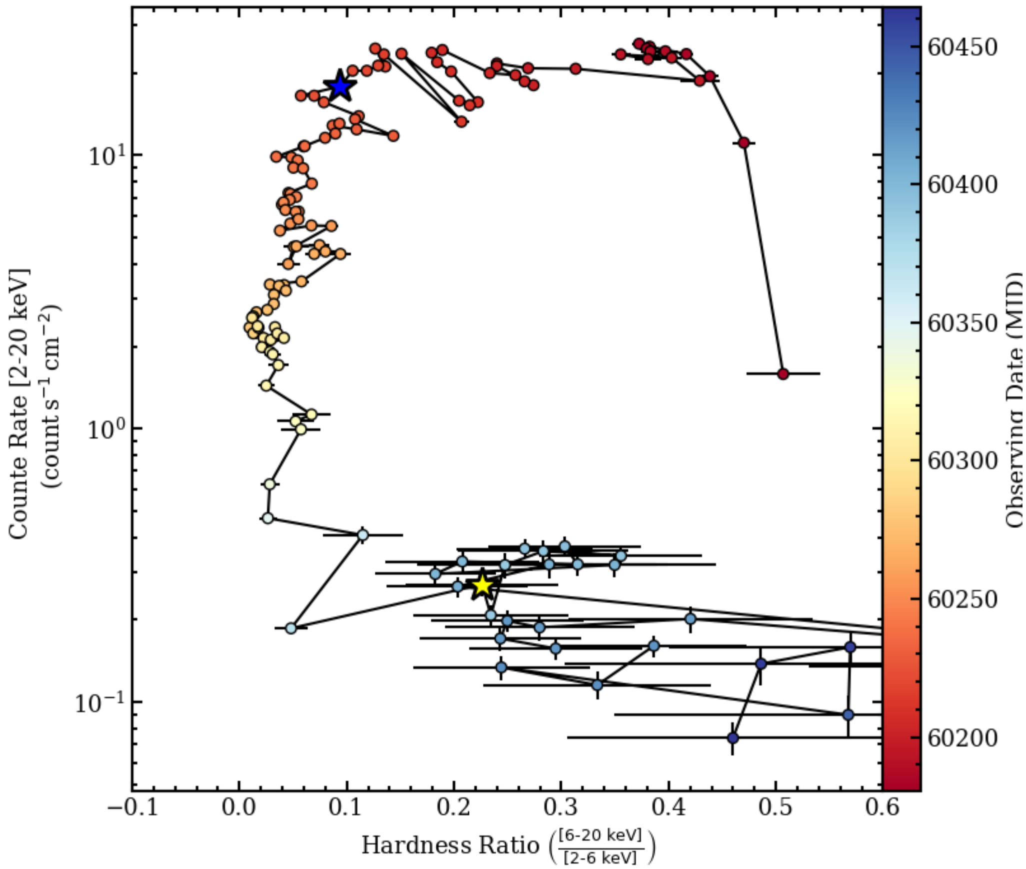

Figure 5.3: X-ray hardness intensity diagram of Swift J1727 when its X-ray luminosity was detectable by MAXI/GSC. Comparing the observed HID to the schematic in Figure 1.2, it is apparent that Swift J1727 closely follows the standard track through the X-ray hardness intensity diagram. I highlight the observations where the source was first identified to undergo a hard-to-soft (blue star) and a soft-to-hard (yellow star) state transition. Note that the hardness ratio shows clear variability during the decline in X-ray luminosity (i.e., from a rate of ~ 20 to $\sim 1\,\mathrm{count\,s^{-1}\,cm^{-2}}$) in the 'soft-state'.

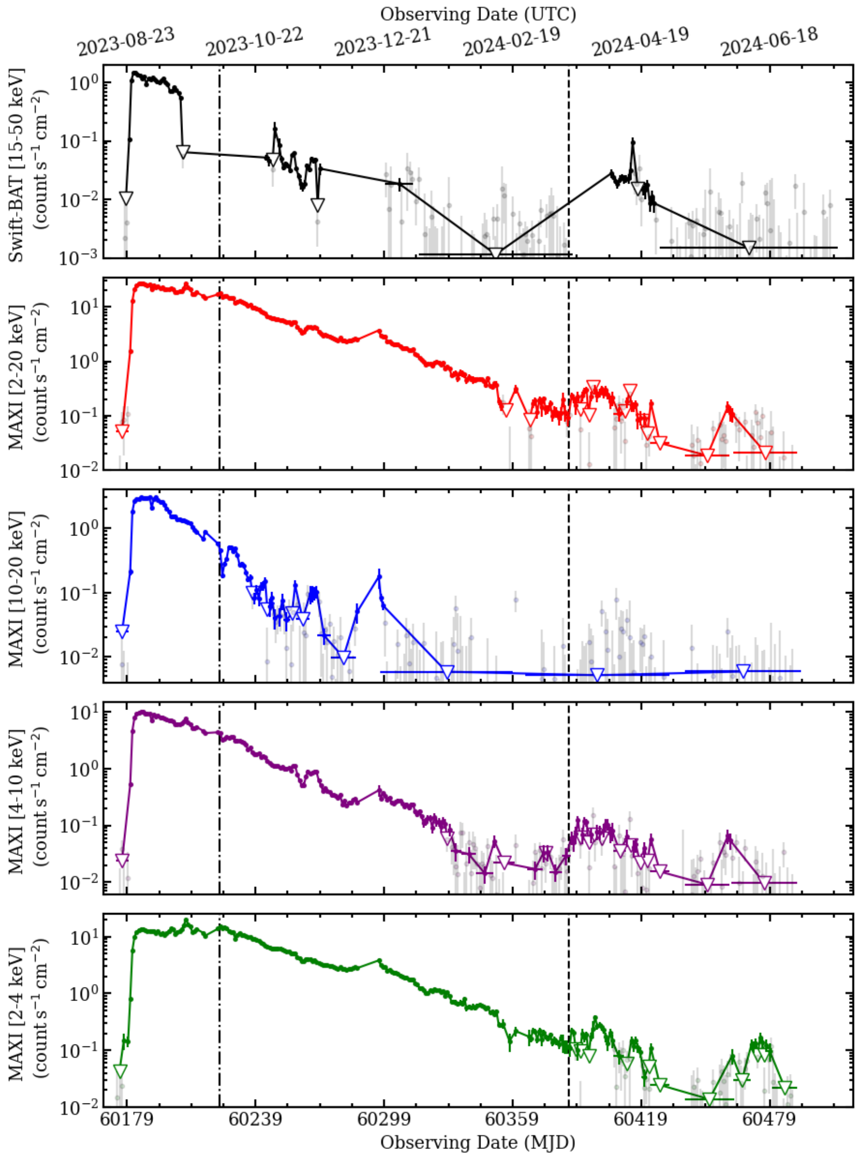

Figure 5.4: X-ray light curves from: (top panel) *Swift*-BAT, 15–50 keV; (second panel) MAXI/GSC, 2–20 keV; (third panel) MAXI/GSC, 10–20 keV; (fourth panel) MAXI/GSC 4–10 keV; (bottom panel) MAXI/GSC 2–4 keV. All light curves have been averaged, as discussed in Section 5.2.3. The dotted-dashed and dashed lines identify the reported dates of the hard-to-soft (Bollemeijer et al., 2023b) and soft-to-hard state transitions (Podgorny et al., 2024), respectively. The grey markers in each panel show the raw data before time averaging. Note the clear, short timescale variability that is more pronounced at higher energies.

a late-time uptick, which was more significant at 4–10 keV than 2–4 keV, after the source's return to the hard state (vertical dashed-line, Fig. 5.4). Furthermore, the source was again detected with *Swift*-BAT, consistent with a hardening of the X-ray spectrum. I analyzed the higher sensitivity *Swift*-XRT observations to verify that the X-ray spectral properties were consistent with a return to hard state emission.

Late-time X-ray Spectra

In Figure 5.5, we show the results of our *Swift*-XRT spectral modelling. At early times, the X-ray flux (top panel, Fig. 5.5) decayed slowly, decreasing in flux by a factor of ~ 15 over ~ 80 dy (MJD 60343 to 60423). Following the slow decline, the X-ray flux decreased by approximately three orders of magnitude in ~ 20 dy (MJD 60423 to 60441) before returning to a slower decline, dropping below the detection threshold on 2024 June 3 (MJD 60463).

While the *Swift*-XRT malfunction prevented real-time observations of the soft-to-hard state transition, the X-ray spectral evolution is consistent with expectation. The hardness ratio (second panel, Fig. 5.5) increased from $\lesssim 0.25$ to $\gtrsim 0.5$ as the spectral model favoured non-thermal power-laws over disc blackbodies (bottom panel, Fig. 5.5). The measured disc temperatures are consistent with the canonical value ($T_{\rm in} \sim 1$ keV; Remillard et al., 2003) and the observed decrease from ~ 0.6 to 0.5 keV is readily explained by the disc cooling as the source exhausts the accreted material, and thus, becomes less luminous. Similarly, in all but the first hard state observations, the power-law photon index is consistent with the canonical values of $\Gamma \sim 1.7$ (Remillard et al., 2003). The deviating observation with $\Gamma \sim 2.2$ likely had elevated soft X-ray emission from remnants of the thermal accretion disc as the source evolved through the intermediate states. Regardless, the evolution observed with *Swift*-XRT supports the state-transition interpretation drawn from the lower sensitivity MAXI/GSC data and claimed by (Podgorny et al.,

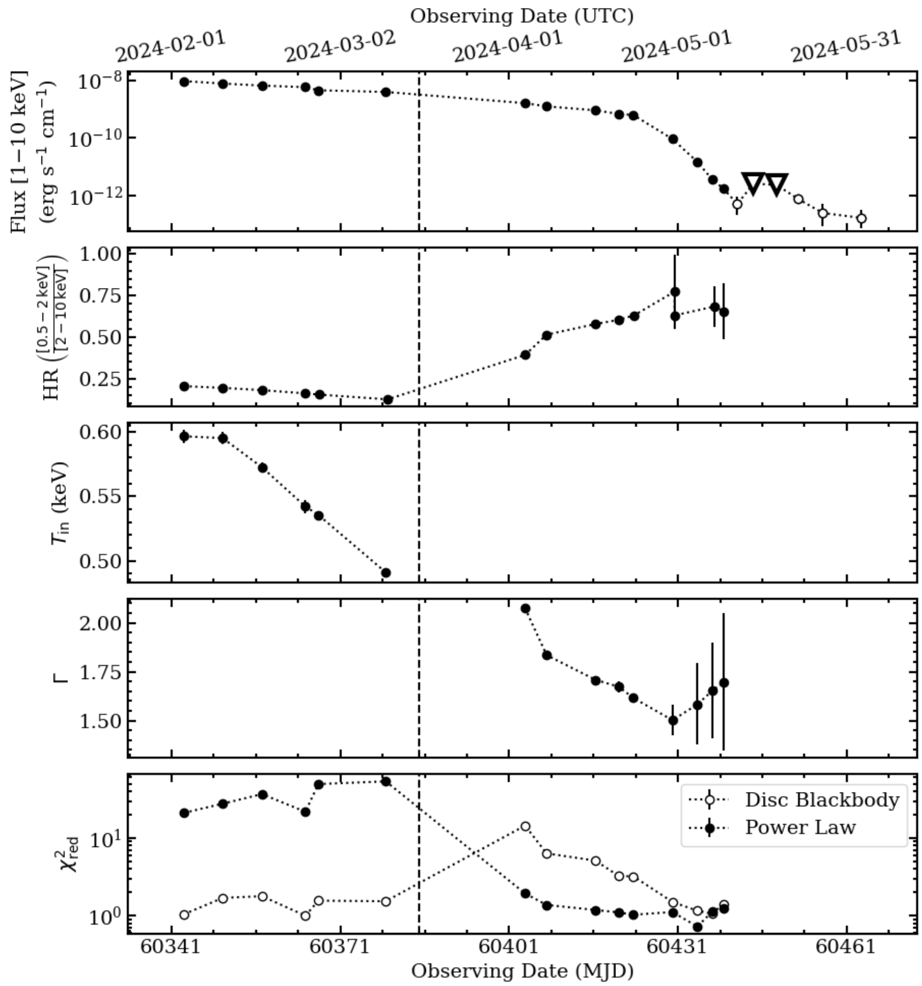

Figure 5.5: X-ray flux and spectral modelling with *Swift*-XRT; (top panel) X-ray flux in the 1–10 keV energy band; (second panel) X-ray hardness ratio using the 0.5–2 keV and 2–10 keV ranges for the soft and hard bands; (third panel) the inner disc temperature for the disc blackbody fits; (fourth panel) the power law photon index; (fifth panel) the reduced χ^2 statistics representing quantifying the goodness of fit. For fluxes above $2 \times 10^{-11}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$, there were between 41 and 552 degrees of freedom (see Table 5.6)Following the hard state transition (dashed vertical line), the χ^2_{red} values for the power law model decrease from $\gtrsim 100$ to $\lesssim 2$; a reverse evolution was observed for the disc blackbody. This is consistent with the claimed state transition leading to a change in which spectral model dominates the X-ray emission.

2024).

5.3.2 Stokes I Evolution

The radio properties of Swift J1727 are shown in Figure 5.6: (top panel) Stokes I flux density; (second panel) intra-band spectral index for $\geq 30\sigma$ Stokes I detected epochs; (third panel) linear polarisation fraction; (fourth panel) the intrinsic linear polarisation angle; (bottom panel) rotation measure. The vertical shaded regions highlight S-band observations — for clarity, I have also made these data points larger.

After the onset of the outburst, I detected a bright (~ 50 mJy) unresolved radio source on 2024 August 27 (MJD 60183), spatially coincident with the prior VLA detection (Miller-Jones et al., 2023b) and with radio spectral properties consistent with a hard state compact jet (i.e., spectrally flat). Henceforth, any unresolved radio emission at this position is called the 'core'. Over the next week, the core brightened to ~ 100 mJy, remaining at this flux density for ~ 1 month (black points, Fig. 5.6). On 2024 October 1 (MJD 60218), the flux density decreased to ~ 20 mJy before strongly flaring, reaching a maximum of ~ 1 Jy over the next two weeks. The radio flaring coincided with the reported times of the hard-to-soft-state transition (dotted-dashed vertical line; Bollemeijer et al., 2023a,b). This temporal coincidence suggests a jet ejection likely occurred; under the standard paradigm for radio quenching of the compact jet, the core emission that persisted throughout the soft-state was likely due to long-lasting jet ejecta close to the position of the core.

The radio flare quickly decayed to ~ 50 mJy, plateauing for ~ 3 weeks before declining to ~ 6 mJy by 2023 November 25 (MJD 60273). Swift J1727 then exhibited a re-brightening, with a peak value of ~ 25 mJy on 2024 January 7 (MJD 60316), subsequently continuing its decline in flux density. The core reached a low of ~ 0.5 mJy, before undergoing a second

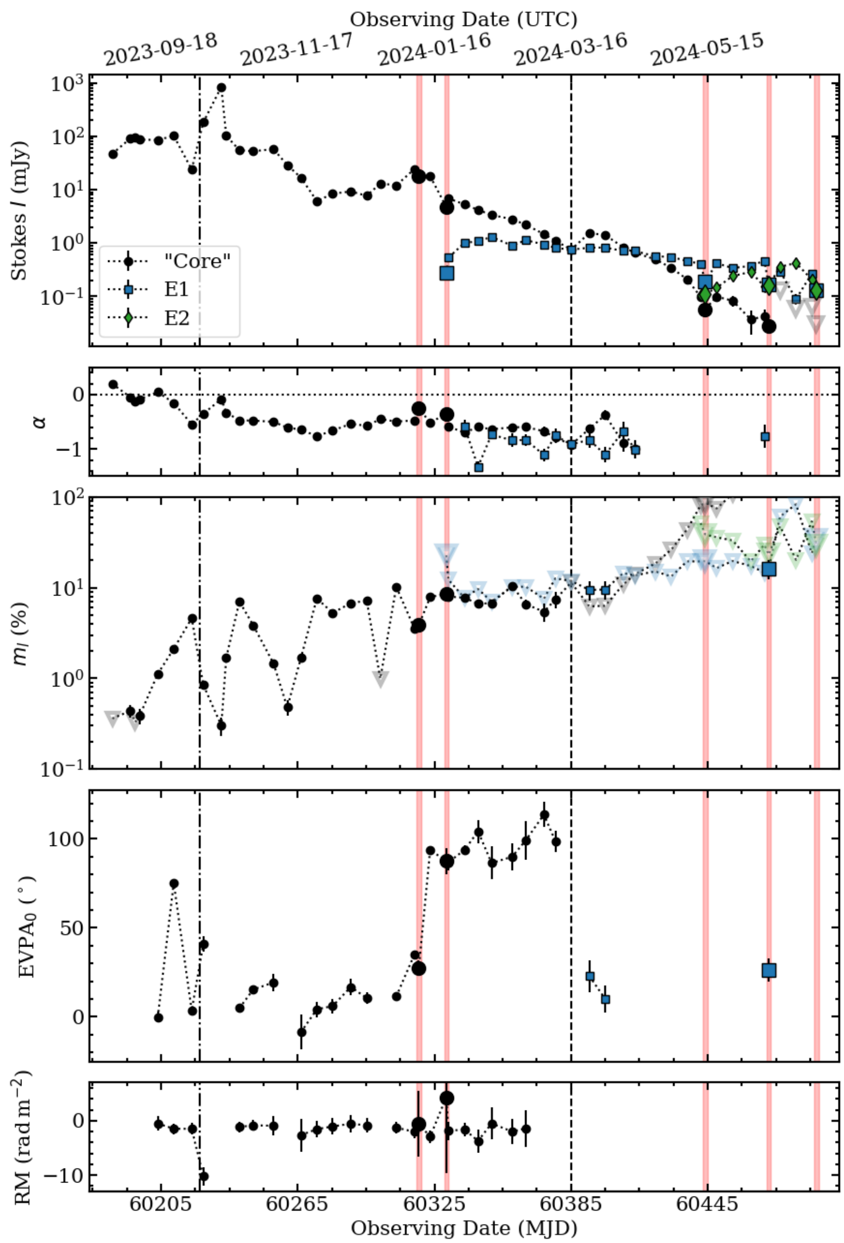

Figure 5.6: Radio evolution of Swift J1727 as seen by MeerKAT: (top panel) Stokes I flux density; (second panel) the intra-band spectral indices; (third panel) linear polarisation fraction; (fourth panel) intrinsic linear polarisation angle; (fifth panel) rotation measure. I mark the reported hard-to-soft and soft-to-hard state transition dates with a dotted-dashed and dashed vertical line, respectively. The shaded regions and larger markers denote the S-band observations. All other data were taken in L-band.

re-brightening at a time consistent with a return to the hard accretion state, and thus the re-formation of a compact jet (dashed line, Fig. 5.6; Podgorny et al., 2024; Russell et al., 2024). After that, the core declined alongside the X-ray luminosity, dropping below the 3σ-detection threshold of ~ 0.06 mJy on 2024 June 15 (MJD 60476).

On 2024 January 21 (MJD 60330), the radio core was marginally resolved at L-band due to the emergence of a single southward moving jet ejectum (E1; blue points, Fig. 5.6). I confirmed the presence of a resolved ejectum with an S-band observation (second panel, Fig. 5.7). The initial brightening of E1 (in L-band) is likely due to our multi-component Gaussian fitting misattributing some E1 flux density to the core. Once significantly resolved, E1 decayed slowly from ~ 1 mJy to ~ 0.2 mJy by 2024 July 1 (MJD 60492). On 2024 May 11 (MJD 60441), E1 appeared extended, motivating another S-band observation, revealing the 'single' ejectum consisted of (at least) two closely separated, southward-moving ejecta. The presence of two ejecta was confirmed by quasi-simultaneous VLA observations (T.D. Russell, private communication). The slower ejectum (E2) is shown with the green diamonds in Figure 5.6. Similar to the early rise of E1, the ejecta are marginally resolved at L-band, and thus, the apparent variability in the flux density is likely due to fitting errors. The higher resolution S-band observations components (see Fig. 5.7) that mitigate fitting errors revealed that E1 and E2 have similar flux densities.

5.3.3 Jet Ejections and Proper Motions

The proper motions of the core (black circles), E1 (blue square), and E2 (green diamonds) are shown in Figure 5.8. I measured the angular offsets with respect to the VLA position reported in Miller-Jones et al. (2023a). The middle panel presents the core declination offset. Given that the jet position angle is $\lesssim 10^{\circ}$, the declination offset shows the direction of the

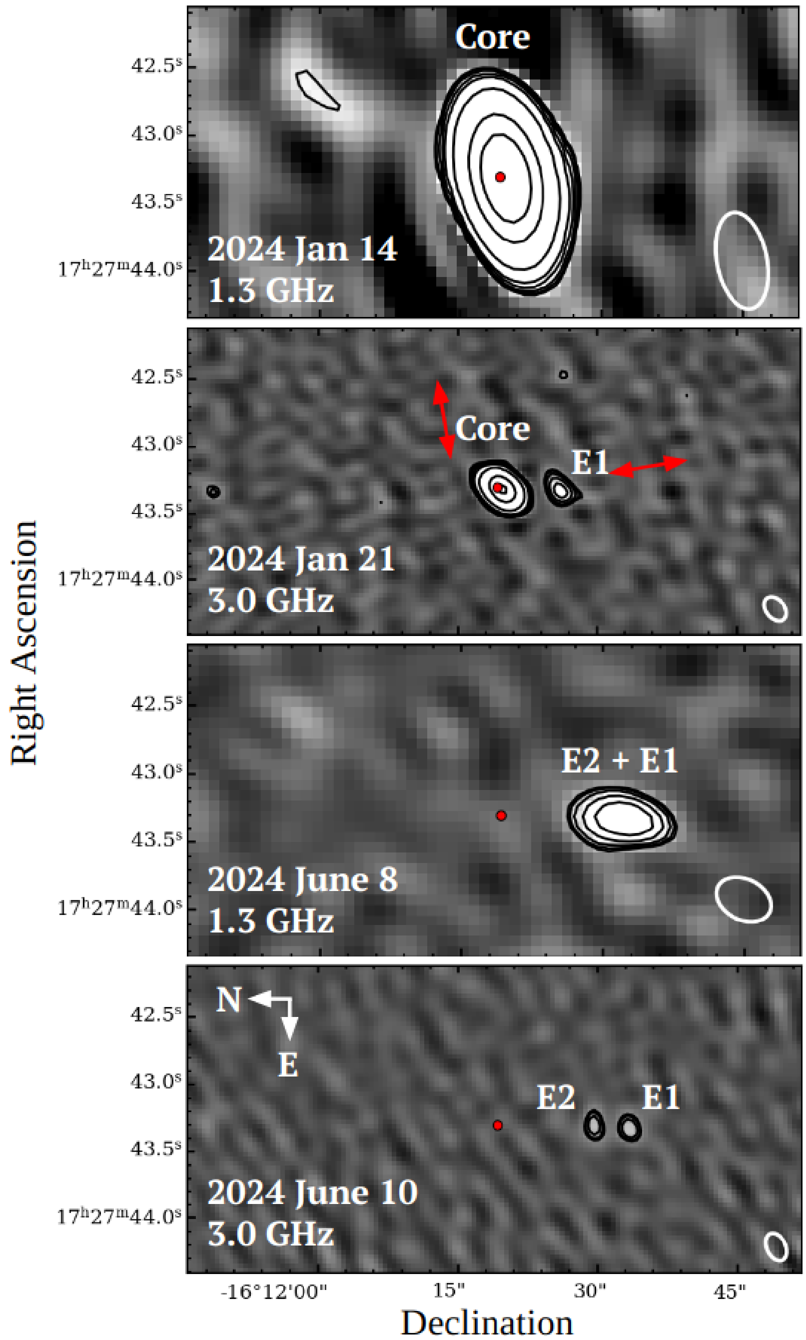

Figure 5.7: Multi-frequency radio imaging of Swift J1727 at four epochs. The images have been rotated by 90° for plotting purposes; i.e., southward, where jet ejecta are eventually observed, is to the right. The red circle marks the core position as reported by Miller-Jones et al. (2023b). The dates and frequencies of each image are shown in the bottom left-hand corner. The red arrows in the second panel show the intrinsic EVPA_0 of the core and E1. Looking at the bottom two panels, it is clear that L-band observations insufficiently resolve the ejecta, even with uniform weighting. To highlight the target, we include contours at the 3, 4, 5, 8, 10, 25, 100, and 200σ levels.

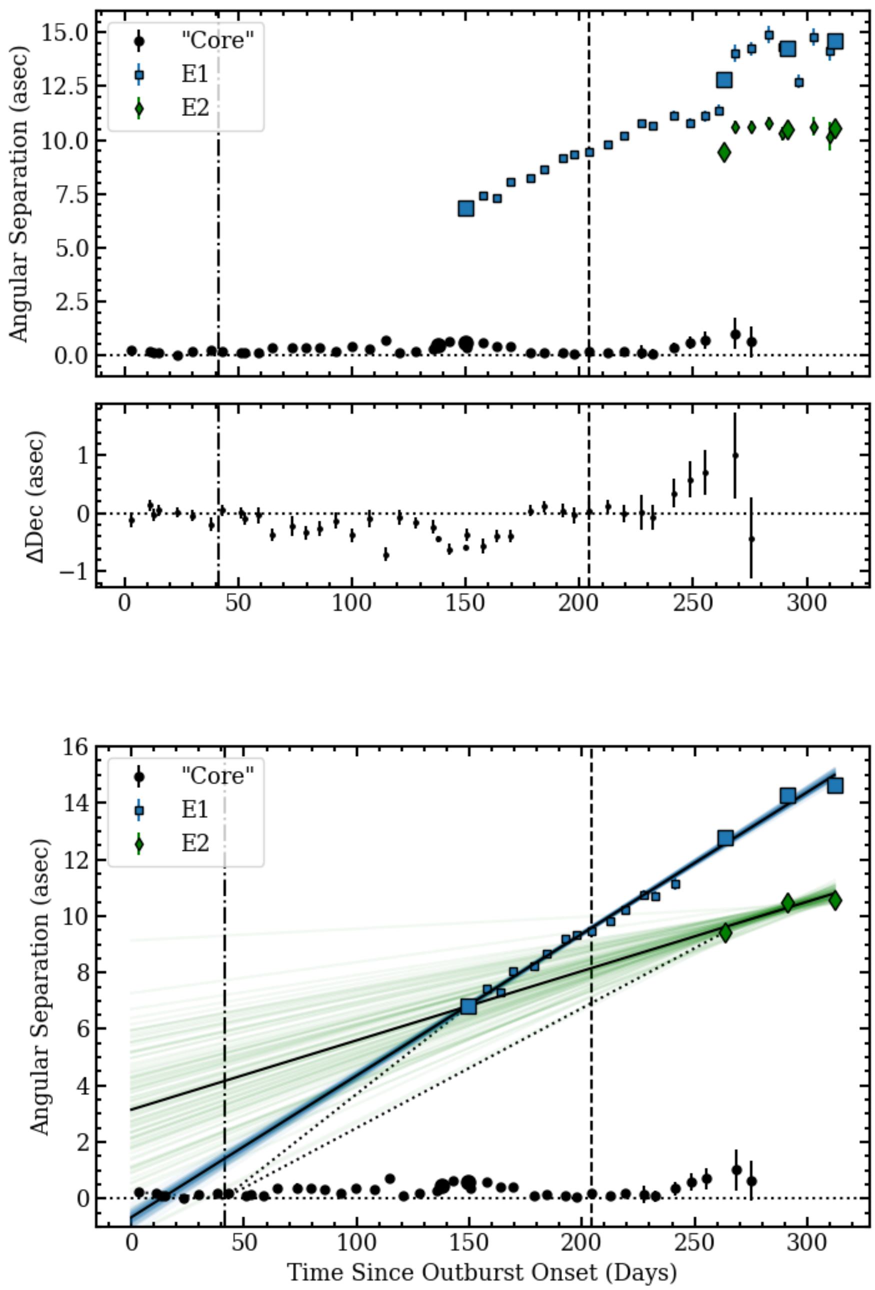

Figure 5.8: Angular offsets as a function of time. (top panel) all L+S-band observations. (middle panel) The declination offset of the core; given the largely North-South jet position angle, this panel shows the offset direction (negative is South and in the direction of E1/E2). (bottom panel) a sub-set of observations that only include the S-band after E1 became marginally resolved (i.e., after 2024 May 5, MJD 60435). For the bottom panel, we include two linear fits: (i) a 'physics-agnostic fit' assuming the ballistic motion of the two ejecta (solid diagonal lines) and (ii) a two-point fit (dotted diagonal lines), fixing the launch date as the hard-to-soft transition (vertical dashed-dotted line).

angular offset, where negative is South and in the direction of E1/E2.

Before the hard-to-soft-state transition (vertical dashed-dotted line; ~ 40 dy after outburst onset), the position of the radio core was consistent with its archival position. Following this state transition, the core appears more offset to the South. Resolving compact jet structure requires milliarcsecond angular resolution (e.g., Wood et al., 2024). In contrast, the relative motions and brightening of ejecta can produce offsets on the observed arcsecond scales. Evolving ejecta can produce 'jitter.' Therefore, the astrometric behaviour of the radio core is consistent with the interpretation that the soft-state 'core' is not a compact jet but unresolved (or marginally resolved) jet ejecta; this is also consistent with past BHXB outbursts, as compact jet emission has never been observed in the soft-state.

Well after the hard-to-soft-state transition, I resolve two southward-moving jet ejecta, E1 and E2. E1 first appears $\sim$150 dy after the onset of the outburst. At somewhat earlier times, it is sufficiently separated from the core that we could have detected E1 if it was relatively bright compared to the >5 mJy radio core. For E1, the angular separations evolved linearly at early times, deviating from the trend ~ 250 dy post-outburst. The deviation coincides with the separation of E2 from E1. Following this separation, the L-band positions of E1 and E2 show significant variability. As the components are marginally resolved in L-band (see third panel, Fig. 5.7), I am likely underestimating the astrometric errors of the multi-component fits; appropriate treatment of the errors would need to account for co-variances between the fit parameters of the overlapping Gaussians. Given these effects, I include a truncated sample (bottom panel, Fig. 5.7) that only consists of the S-band detections once E1 was resolved from E2. I use the truncated sample to analyse proper motion in this work. An investigation into the appropriate errors at L-band for the highly overlapping components is ongoing and left for future work.

On the bottom panel of Figure 5.7, I include a simple, relatively physics-agnostic, linear fit (i.e., ballistic motion) to the ejecta proper motion. Using SCIPY.`curve_fit` (a PYTHON package) I measure proper motions of $\mu_{\rm E1} = 50.1 \pm 0.9\,{\rm mas\,dy^{-1}}$ and $\mu_{\rm E2} = 24 \pm 7\,{\rm mas\,dy^{-1}}$, where 'mas' is the units of milliarcseconds. Following the prescription from (Mirabel & Rodríguez, 1994), the proper motions for an approaching (μ_a) and receding (μ_r) component are

$$\mu_{a,r} = \frac{\beta \sin i}{1 \mp \beta \sin i} \frac{c}{d} \longrightarrow \beta = \frac{\mu_{a,r} d}{c \sin i \pm \mu_{a,r} d \cos i}, \tag{5.13}$$

where d is the distance and β is the ejectum velocity as a fraction of the speed of light, assuming equal velocities from bipolar ejecta (solid diagonal lines, bottom panel, Fig. 5.8. Using the values for Swift J1727, $d \sim 2.7\,{\rm kpc}$ and $i \sim 40°$, I calculate moderately relativistic velocities for both ejecta: $\beta_{\rm E1} \sim 0.66$ and $\beta_{\rm E2} \sim 0.46$. However, this simple model predicts the ejection of E1 occurring before the hard-to-soft transition (something never observed in BHXBs) and, more prohibitively, an E2 ejection before the start of the outburst. Instead, I adopt a more physically motivated but equally simple picture, assuming the ejection of E1 and E2 occurred during the reported hard-to-soft transition ($\sim 41\,{\rm dy}$ post-onset) and adopting ballistic motion until the first detection. Here, the proper motions are $\mu_{\rm E1} = 63\,{\rm mas\,dy^{-1}}$ and $\mu_{\rm E2} = 43\,{\rm mas\,dy^{-1}}$ and the jet speeds are $\beta_{\rm E1} \sim 0.70$ and $\beta_{\rm E2} \sim 0.57$.

I do not spatially resolve a northward counter-jet, despite most BHXBs exhibiting bi-polar jet ejections (e.g., Fender et al., 1999a; Rushton et al., 2017; Miller-Jones et al., 2019; Carotenuto et al., 2021; Bahramian et al., 2023). However, I observe a consistently northward offset of the core towards the end of the observations ($\sim$240–270 dy post-onset). Comparing the four northward offset epochs to the non-offset case, I calculate a $\chi^2_{\rm red} = 3.53$ (3 degrees of freedom) corresponding to a null hypothesis probability of

$p \sim 1.4\%$. While this p-value cannot be taken as conclusive evidence, it is suggestive of a northern counter-jet. I note, however, that no trial penalties have been applied.

The X-ray modelling of Swift J1727 favours moderate disc-inclination angles of $i \sim 30$–$50°$ (Draghis et al., 2023; Peng et al., 2024; Svoboda et al., 2024), where i also corresponds to the jet viewing angle (i.e., $i = 0$ looks down the jet axis). At medium inclinations, like those seen here, light travel time will decrease the proper motion, and relativistic aberration will significantly affect the counter-jets, decreasing their flux densities. Here, I take a quantitative look at the expected counter-jet properties of E1 and E2.

Using the intrinsic speeds of E1 and E2 (inferred from the fixed-date fits), and, once again, assuming symmetric ejections, the proper motions of their counter-components are $\sim 19\,\mathrm{mas\,dy^{-1}}$ and $\sim 16\,\mathrm{mas\,dy^{-1}}$, respectively. These proper motions suggest that at the time when E1 was resolved (~ 150 dy post-onset; ~ 110 dy after the assumed launch date), its counter component would be at an angular offset of 2.2 arcseconds. While no clear ejecta are detected, at this time, I observe a southward core offset, inconsistent with an unresolved (northern) counter-jet ejecta. To determine whether relativistically-driven dimming accounts for non-detection of counter-jet ejecta, I follow the prescription from Miller-Jones et al. (2004), which expands on Mirabel & Rodríguez (1994) by considering light travel times. The flux ratio of the approaching and receding components (R_F) measured at the same observed time is given by

$$R_F = \frac{F_a}{F_r} = \left(\frac{1+\beta\cos i}{1-\beta\sin i}\right)^{k-p}, \tag{5.14}$$

where the $k - p \sim 1$ for jet ejectum (k=3) and a standard power-law distributed electron population with power-law index $p = 2$.[9] For E1, $R_F \sim 3.3$, and thus, for an E1 flux density of $\sim 1\,\mathrm{mJy}$ the counter-jet would be $\sim 0.3\,\mathrm{mJy}$. Given the brightness of the core ($\sim 5\,\mathrm{mJy}$), even with the improved angular resolution at of S-band (PSF FWHM of ~ 2.5 arcseconds), the convolution between the PSF and a ($\sim$ 5 mJy) point source would be $\sim 1.5\,\mathrm{mJy\,beam^{-1}}$ at an angular separation of ~ 2 arcseconds. Therefore, it is very plausible that, at early times, the bright core emission is 'washing out' any potential northern component from counter-jet ejecta. At later times, the explanation becomes unclear; ~ 270 dy post-onset (~ 230 dy post-launch), the combined flux density of E1 and E2 at L-band is $\sim 0.5\,\mathrm{mJy}$, while the core is $< 0.1\,\mathrm{mJy}$. Given their smaller proper motions, we would not expect the counter-jets of E1 and E2 to be separable from each other, and thus, we would measure integrated flux densities dimmed by $R_F \sim 3$ ($\sim 0.15\,\mathrm{mJy}$). The angular separation of the combined counter-ejecta would be ~ 4 arcseconds, and thus, even with the larger L-band PSF (~ 5 arcseconds), the counter-jet should be noticeably separated and easily detected at these times (considering our typical rms noise of $\sim 0.03\,\mathrm{mJy\,beam^{-1}}$).

The absence of counter-jet ejecta suggests that Swift J1727 occupies the extremes of the current proposed geometric properties. For example, if instead, we adopt the lower range on inclination, $i = 30°$, and the upper range on distance ($d = 3.6$ kpc, the 3σ distance upper limit), the counter-jet proper motion becomes $\sim 10\,\mathrm{mas\,dy^{-1}}$ with an $R_F \sim 6$ — leading to both smaller angular separations and weaker counter-jet flux densities, consistent with my non-detection of such ejecta. There have been alternative explanations for one-sided jets that invoke obscuration by strong disk winds (e.g., Fender et al., 1999b). However, these explanations were based on one-sidedness observed on milliarcsecond scales, thus close to the black hole. At arcsecond distances, it seems extremely unlikely that disk wind obscuration blocks

[9] This differs from the standard formula for ejecta at the same angular separation.

emission from the receding components. Another potential explanation could be that the localized ISM is asymmetric around the source and has a stronger interaction in the southern direction.

Setting aside issues related to counter-jet ejecta, even from this simplified picture, it is evident that E1 and E2 have undergone deceleration and, thus, energy exchange with the environment. Given the uncertain launch dates, the MeerKAT data alone are insufficient for accurately modelling the deceleration; Carotenuto et al. (2024) showed that prior information on jet ejection time is crucial for constraining parameter estimation during modelling. However, as discussed in Section 5.4, the X-KAT collaboration has access to multi-facility observations (including very long baseline interferometry) that should allow us to more accurately pin down the deceleration rates and ejection dates. After combining the totality of our observations, I intend to quantitatively model the kinematics of decelerating ejecta, extracting estimates of their energy and mass content (e.g., Carotenuto et al., 2022). Furthermore, a better understanding of the jet kinematics will allow for a deeper quantitative look at the issues regarding counter-jet ejecta.

5.3.4 Steep-Spectrum Compact Jet and The L_R–L_X Plane

The radio spectral properties of Swift J1727 exhibited both typical and atypical behaviours. At early times, the radio core behaved as expected; the intra-band spectral indices were flat or inverted ($\alpha \gtrsim 0$), characteristic of a hard state compact jet. The bright radio flare showed spectral variability, evolving from optically thin-to-thick-to-thin, matching expectations for 'optically thin flares' associated with ejection events (Fender & Bright, 2019). Following the flaring, the spectral index of the core was optically thin ($\alpha \lesssim -0.5$), consistent with the interpretation that the soft-state 'core' emission was

Table 5.2: Inter-band L–S-band spectral indices

Component[a]	L-Date	S-Date	α
Core	2024 January 7	2024 January 9	-0.426 ± 0.002
	2024 January 21	2024 January 22	-0.418 ± 0.005
	2024 May 11	2024 May 13	-0.6 ± 0.3
	2024 June 8	2024 June 10	-0.5 ± 0.5
E1	2024 January 21	2024 January 22	-0.76 ± 0.07
	2024 May 11	2024 May 13	-0.89 ± 0.08
	2024 June 8	2024 June 10	-1.16 ± 0.06
	2024 June 29	2024 July 1	-0.81 ± 0.13
E2	2024 May 11	2024 May 13	-0.1 ± 0.2
	2024 June 8	2024 June 10	0.03 ± 0.13
	2024 June 29	2024 July 1	-0.54 ± 0.15

Note. The flat spectral index observed in E2 is likely the result of L-band fitting errors, underestimating the flux in E2 (and instead attributing it to E1).

the result of spatially unresolved jet ejecta. However, after returning to the hard state, the radio core's intra-band spectral index remained surprisingly steep despite evidence suggesting a compact jet reformed. The spectrally steep core emission was consistent with our late-time inter-band spectral indices (see Table. 5.2, although the errors in late-time core spectral indices are substantial). Furthermore, the two-point spectral indices for the $< 30\sigma$ detections show a (variance-weighted) average of $\alpha = -0.8 \pm 0.2$. While these values will be biased, the bias tends towards flatness (Heywood et al., 2016). Lastly, higher frequency radio observations (at ~ 6 and 8 GHz) taken with the VLA (T.D. Russell, private communication) found a similarly steep spectral index. Therefore, it is likely that the core is, in fact, spectrally steep. While residual optically thin emission from jet ejecta could steepen the spectral index, Swift J1727 is already exceptionally radio-quiet, as seen from its position in the L_r–L_X plane (Figure 5.9).

I calculate the radio and X-ray luminosities before the initial hard-to-soft transition and after the source returns to the hard state. Before the source transitioned, when *Swift*-XRT was heavily affected by photon pile-up, I

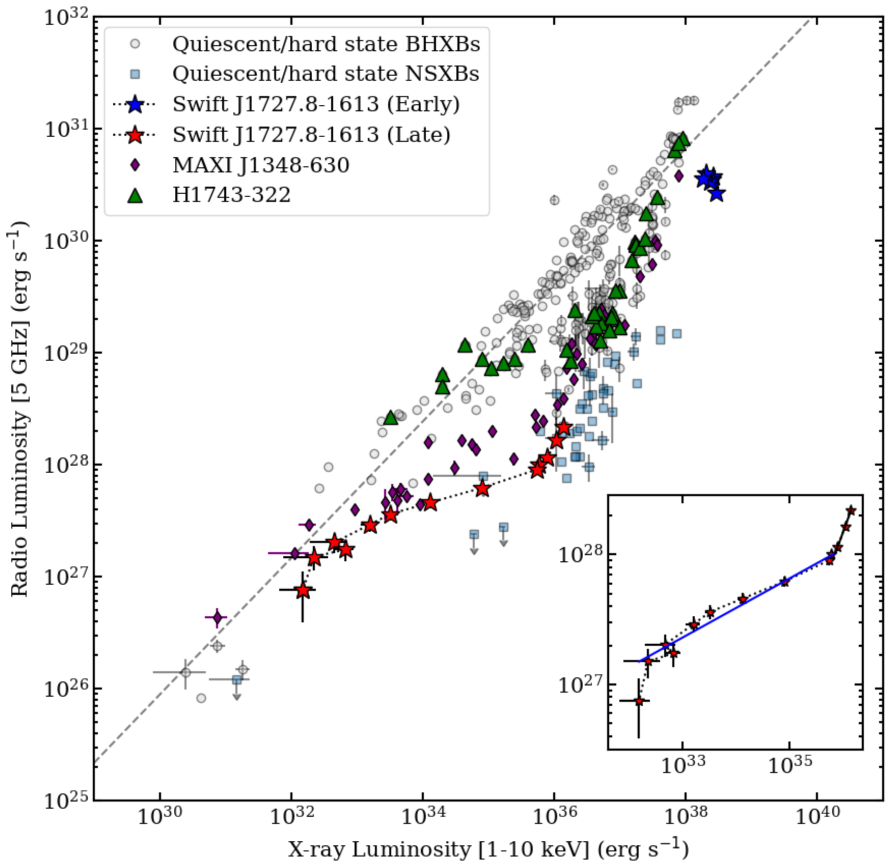

Figure 5.9: The L_R–L_X plane from Bahramian & Rushton (2022), including the data from Swift J1727 before the initial transition to the soft-state (blue stars) and after the return to the hard state (red stars). We highlight the data of MAXI J1348−630 and H1743−322, the two most comprehensively monitored radio-quiet BHXBs. The diagonal dashed line is the 'standard track' for radio-loud BHXBs ($L_R \propto L_X^{0.6}$). The inset panel shows the broken power law that fits the late-time data of Swift J1727. Swift J1727 appears to be the most radio-quiet BHXB to date.

converted the MAXI/GSC 2–20 keV count rate (on the same day as our radio observations) to a 1–10 keV X-ray flux using the NASA software WebPIMMS. I calculated the flux assuming a power-law X-ray spectrum ($\Gamma = 1.7$) and our *Swift*-XRT measured column density, $N_H = 0.1989\,\mathrm{cm}^{-2}$, converting to an X-ray luminosity at an assumed distance of 2.7 kpc:

$$L_{X,1\text{–}10\,\mathrm{keV}} = 4\pi d^2 F_{X,1\text{–}10\,\mathrm{keV}}. \tag{5.15}$$

After the return to the hard state, I calculate the X-ray luminosity using the fluxes derived from our *Swift*-XRT spectral modelling. Similarly, I calculate the radio luminosity from the MeerKAT Stokes I flux densities, scaling our fluxes to the 5 GHz standard according to

$$L_{R,5\,\mathrm{GHz}} = 4\pi F_{R,\nu}\nu_0 \left(\frac{\nu}{\nu_0}\right)^{\alpha} d^2, \tag{5.16}$$

where, $\nu_0 = 5\,\mathrm{GHz}$ and $\nu = 1.28\,\mathrm{GHz}$. Utilizing the more modest radio decay, I (linearly) interpolate the MeerKAT data onto the *Swift*-XRT times. I use the measured intra-band spectral indices to scale each observation before the soft-state transition, as each radio detection is $\gg 30\sigma$. At late times, I adopted the variance-weighted average spectral index ($\alpha_{\mathrm{avg}} = -0.65 \pm 0.05$) of the $> 30\sigma$ detections.

Figure 5.9 presents the L_R–L_X measurements of Swift J1727, which distinguishes itself as an extraordinarily radio-quiet BHXB[10], more akin to the population of NSXBs. Like other 'radio-quiet' BHXBs, I observed an evolution of radio–X-ray power-law index (i.e., β; $L_R \propto L_X^{\beta}$). Following the return to the hard state, Swift J1727 exhibited a steep β-value, transitioning to a shallow β as the source underwent the rapid three-order magnitude decrease in X-ray luminosity, re-joining the 'radio-loud' track. While this

[10]Here, radio-quiet is an empirical description and does not comment on the physical mechanism; i.e., whether the jet is suppressed or the accretion flow is enhanced ('X-ray bright').

multi-track evolution is considered the canonical behaviour of radio-quiet sources, most monitoring of radio-quiet BHXBs has only been able to catch one of the two tracks (e.g., Cao et al., 2014; Parikh et al., 2019; Xie et al., 2020). Indeed, only two BHXBs, MAXI J1348−630 (purple diamonds, Fig. 5.9; Carotenuto et al., 2021) and H1743−322 (green triangles, Fig. 5.9; Coriat et al., 2011; Williams et al., 2020), have been detected in both radio-quiet tracks. Therefore, Swift J1727 is only the third BHXB where we captured steep and flat radio-quiet tracks within a single outburst.

I quantified the L_R–L_X evolution with a broken power-law using the `curve_fit` function from the PYTHON package SCIPY,

$$F_R = \begin{cases} A(L_X/L_{X,\mathrm{tran}})^{\beta_{\mathrm{flat}}}, & \text{if } L_X < L_{X,\mathrm{tran}} \\ A(L_X/L_{X,\mathrm{tran}})^{\beta_{\mathrm{steep}}}, & \text{if } L_X \geq L_{X,\mathrm{tran}} \end{cases}, \tag{5.17}$$

measuring $\beta_{\mathrm{flat}} = 0.227 \pm 0.014$, $\beta_{\mathrm{steep}} = 1.12 \pm 0.19$, and $L_{X,\mathrm{trans}} = (7.1 \pm 0.8) \times 10^{35}\,\mathrm{erg\,s^{-1}}$. While necessary for fitting the amplitude ($A = (1.03 \pm 0.07) \times 10^{28}\,\mathrm{erg\,s^{-1}}$), it is not yet clear that it holds physical value. We also fit a simple power-law model to demonstrate the need for a break (i.e., $L_R = AL_X^{\beta}$; $\beta = 0.26 \pm 0.03$). Calculating the reduced χ^2 for both models, we get $\chi^2_{\mathrm{red}} = 4.43$ (11 degrees of freedom) and $\chi^2_{\mathrm{red}} = 1.04$ (9 degrees of freedom) for the simple power-law and the broken power-law fits, respectively. The broken-power law fits is clearly required. Due to the significant separation between the early (blue stars, Fig. 5.3.4) and late-time (red stars, Fig. 5.3.4) luminosities, we excluded the early-time data from the initial broken-power law fitting. As a check, we extrapolated our fit to the MAXI/GSC data, fixing β_{flat}, $L_{X,\mathrm{trans}}$, and A, and fitting β_{steep}. We measure $\beta_{\mathrm{steep}} = 1.0 \pm 0.1$, consistent with the 'late-time-only' value.

Comparing our broken power-law fit MAXI J1348−630 (another ThunderKAT / X-KAT target), we find remarkable agreement; $\beta_{\mathrm{flat}} = 0.24 \pm 0.05$,

$\beta_{\mathrm{steep}} = 0.95 \pm 0.04$, and $L_{X,\mathrm{trans}} = (6.3 \pm 1.5) \times 10^{35}\,\mathrm{erg\,s^{-1}}$ (Carotenuto et al., 2021). For H1743−322, we find similar agreement in the power-law indices ($\beta_{\mathrm{flat}} = 0.23 \pm 0.07$, $\beta_{\mathrm{steep}} = 0.95 \pm 0.04$). However, the transition luminosity is significantly higher with $L_{X,\mathrm{tran}} \sim 2 \times 10^{36}\,\mathrm{erg\,s^{-1}}$ at 3–9 keV (Coriat et al., 2011), corresponding to $L_{X,\mathrm{tran}} \sim 4 \times 10^{36}\,\mathrm{erg\,s^{-1}}$ at 1–10 keV (assuming a power-law X-ray spectrum with $\Gamma = 1.7$). MAXI J1348−630 and H1743−322 were observed to re-join the standard radio-loud track at low X-ray luminosities. Our last observation suggests a subsequent steepening of β; however, given that it is a single data point with large errors, we do not make any strong claim that Swift J1727 rejoined and evolved along the radio-loud track at low luminosities.

There is no consensus on the origin of the multiple tracks in the L_R–L_X plane. The possible explanations are broadly separated into jet-based and accretion-based origins, commonly referred to as the 'radio-quiet' and 'X-ray-loud' hypotheses, respectively. Motta et al. (2018) investigated whether the radio-loud/radio-quiet population could result from different viewing angles rather than intrinsic differences between individual BHXBs. The authors found that radio-loud BHXBs tended to have low(er) inclinations, and thus, the radio emission from the jet would be increased by relativistic beaming. Given that X-ray modelling of the accretion flow suggests Swift J1727 is a medium-inclined system, it is unlikely that an absence of beaming can explain its low radio luminosities. While misalignment between the jet and the accretion disc is possible (Maccarone, 2002), the absence of receding ejecta in our monitoring is consistent with significantly beamed emission from an approaching compact jet.

Espinasse & Fender (2018) found that the two populations showed a statistically significant difference in the radio spectral indices, with the radio-quiet sources having steeper spectral indices ($\alpha_{\mathrm{avg}} \sim -0.3$) than the radio-loud population ($\alpha_{\mathrm{avg}} \sim 0.2$). However, Swift J1727 was radio-quiet even when its spectral index was flat (i.e., at early times), suggesting its steep spectral index may be unrelated to its radio-quiet nature. Espinasse & Fender (2018) also discussed some potential origins of the radio-quiet/loud sub-populations, including differences in black hole spin, residual jet-ISM interactions, and compact jets with optically thin break frequency in the radio region; each was deemed an unlikely origin.

Pe'er & Casella (2009) showed that strong magnetic fields in compact jets would suppress their radio flux densities. We measured a linear polarisation fraction as high as $\sim 6\%$ when the radio emission was consistent with a compact jet, higher than the typical compact jet polarisations of $\lesssim 2\%$ for BHXBs (Han & Hjellming, 1992; Corbel et al., 2000; Russell et al., 2015) and AGN radio cores (Hodge et al., 2018). This suggests that the magnetic field in the compact jet of Swift J1727 may be more ordered (and perhaps stronger) than usual, which could have suppressed the radio flux. Like Swift J1727, the compact jets of BHXB XTE J1752$-$223 and Swift J1745$-$26 had polarisation fractions as high $\sim 8\%$, while also exhibiting the steep radio-quiet L_R–L_X track (Brocksopp et al., 2013; Curran et al., 2014)[11]. While only three examples are inconclusive, they indicate a possible connection between highly polarised jets and the radio-quiet subpopulation. Future polarisation monitoring of BHXB hard state jets will be well positioned to probe the proposed relation and to answer *why* some sources would have stronger magnetic fields.

In contrast, the X-ray-loud hypotheses suggest that the tracks originate from variations in the properties of the accretion flows. For example, some authors (e.g., Coriat et al., 2011) have suggested that radio-loud sources transition from radiatively inefficient to efficient accretion flows,

[11]For Swift J1745$-$26 the authors note that the variable hardness ratio makes it difficult to determine if the source was in the hard state, and thus, whether it was valid to measure its L_R–L_X relation. However, during these times, the radio spectral index was continuously flat, consistent with a compact jet, and thus a hard (or hard-intermediate) accretion state.

with the radiative efficiency being a function of the mass accretion rate, and thus, X-ray luminosity (see, e.g., Done et al., 2007, for a review of the different accretion flows). The transition luminosity is scaled by the source's Eddington accretion rate; a future determination of the black hole mass in Swift J1727 will allow us to, more quantitatively, investigate the consistency between our transition luminosity and an accretion efficiency transition.

Alternatively, Meyer-Hofmeister & Meyer (2014) have proposed the existence of a cool inner accretion disc for some hard state BHXBs, providing additional seed photons that are Compton up-scattered in the X-ray corona, thereby increasing the X-ray luminosity. Similar to MAXI J1348−630 (Carotenuto et al., 2021), Swift J1727 deviates from the radio-loud track at X-ray luminosities $\sim 10^{32}\,\mathrm{erg\,s^{-1}} \sim 10^{-6}(M_{\odot}/M_{\mathrm{BH}})L_{\mathrm{edd}}$. Even for the smallest stellar-mass black holes ($\sim 3M_{\odot}$), these luminosities are insufficient for inner disc formation, which requires $L_X \gtrsim 10^{-4}L_{\mathrm{edd}}$ (Meyer-Hofmeister & Meyer, 2014).

Regardless of the underlying physics driving the different tracks, it is clear that as we collect more data, it becomes clear the original notion of a 'standard' track was a misnomer. In 2020, MAXI J1348−630 became the most significant deviation from the radio-loud relation, and now, Swift J1727 has pushed further into the radio-quiet regime. This behaviour is similar to what was observed in NSXBs; thought initially to follow $L_R \propto L_X^{1.4}$ due to one well-monitored source, but later found to have more complex relations (Gallo et al., 2018; van den Eijnden et al., 2021). A larger sample of well-monitored sources (capturing both radio-quiet tracks) will be required to understand the physics behind the radio-loud/radio-quiet dichotomy or whether they are in fact a single population.

It should be noted that throughout this section, I have implicitly assumed that all of the core radio emission originates from a re-formed compact jet. If any (or all) of the core emission originates from residual jet ejecta, the conclusions remain unchanged, while the radio-quiet nature of the source becomes more severe.

5.3.5 The Polarisation Evolution

A Rotation Measure 'Flare'

Perhaps the most extreme evolution of Swift J1727 is seen in its radio polarisation properties. I did not detect any circular polarisation, and thus, I only discuss linear polarisation here. In the lead-up to the $\sim$Jy-level radio flare, the polarisation fraction increased steadily from $\lesssim 0.3\%$ to $\sim 6\%$, suggesting an increased ordering of the magnetic field lines in the compact jet. The polarisation fraction decreased to $\sim 0.3\%$ during the rise of the flare, returning to $\sim 2\%$ two days after the peak, and $\sim 7\%$ after a week.

When measurable, the rotation measures were consistent with a constant value of $\mathrm{RM_{avg}} = -1.4 \pm 0.3\,\mathrm{rad\,m^{-2}}$, for all observations except during the rise of the flare (2024 October 6; MJD 53284) where it decreased to $\mathrm{RM} = -10.2 \pm 1.6\,\mathrm{rad\,m^{-2}}$ (see Fig. 5.10). For rotation measures, a negative sign indicates that the magnetic field points away from the observer; from this point forward, we concentrate only on the magnitude of the change.

Galactic Faraday rotation from the ISM is not expected to be significantly variable on short timescales, and thus, the excess rotation measure $\Delta\mathrm{RM} \sim 8\,\mathrm{rad\,m^{-2}}$ strongly suggests the emergence of a local component. The rotation measure returned to the average value $\lesssim 3$ weeks after the increase. Unfortunately, we could not measure the rotation in the following two observations due to low (polarised) signal-to-noise and an omission of the polarisation angle calibrator. The increased magnitude in RM coincided with a MAXI/GSC count rate evolution: a substantial decrease at 10–20 keV, a moderate decrease at 4–10 keV, and no decrease at 2–4 keV. The transient reduction in hard X-ray luminosity —- a *softening* of the X-ray spectrum —

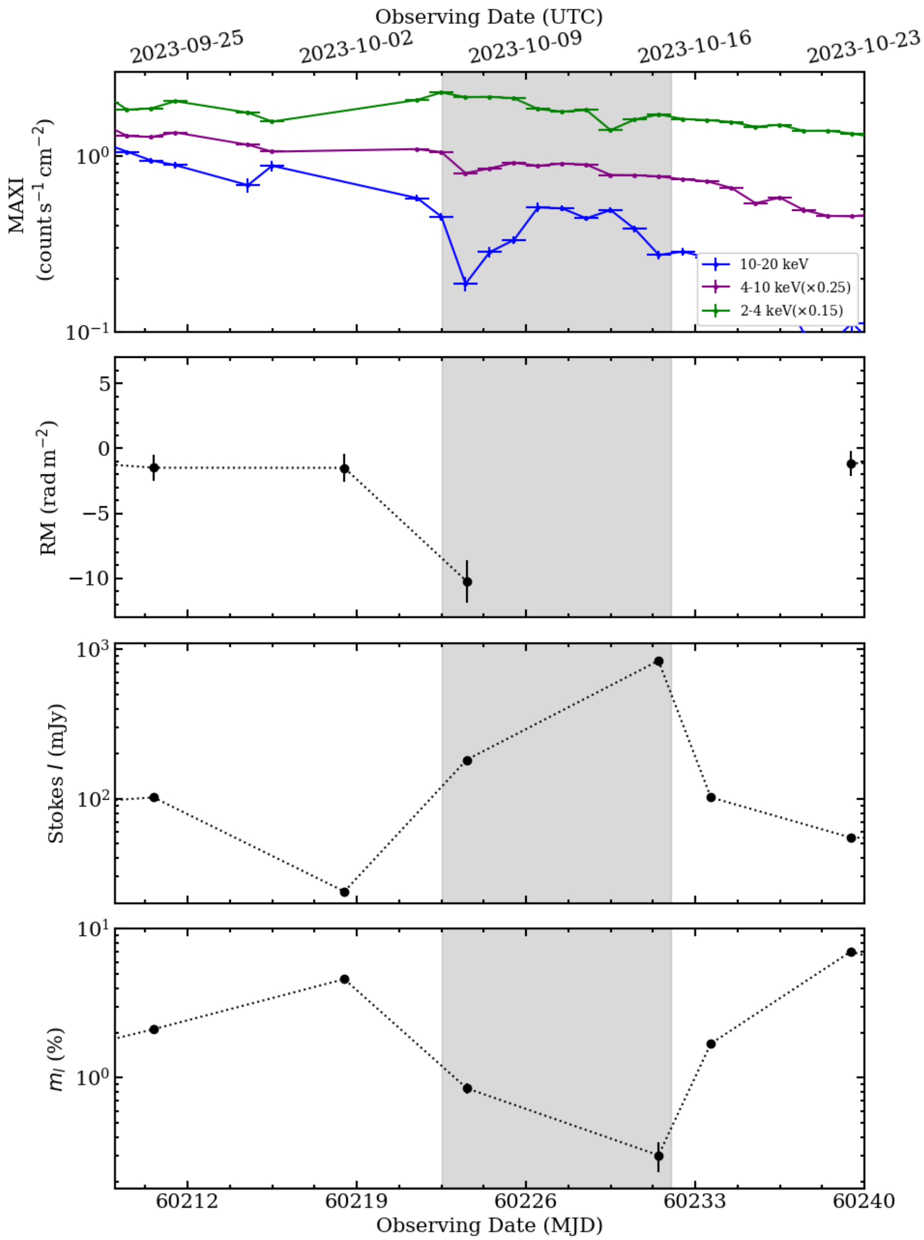

Figure 5.10: The X-ray and radio evolution of Swift J1727 near the first decrease in the polarisation fraction; i.e., a 'dip'. (top panel) the MAXI/GSC light curves; (second panel) the rotation measure; (third panel) the Stokes I flux density; (fourth panel) the linear polarisation fraction. The MAXI/GSC light curves in the 4–10 keV and 2–4 keV sub-bands have been rescaled for plotting purposes. The grey-shaded region highlights the observations that show the dip. We see a correlated evolution: a significant (but transient) decrease in the 10–20 keV X-ray count rate coinciding with the increase in the rotation measure, Stokes I flare, and decrease in polarisation fraction.

during the radio flaring is consistent with the launching of jet ejecta (e.g., one can think of this as akin to a 'mini' state transition). Internal Faraday rotation from internal ejecta plasma could produce the observed increase in the magnitude of the rotation measure and contribute to the observed de-polarisation (Burn, 1966; Sokoloff et al., 1998).

To investigate the feasibility of significant internal Faraday rotation, we adopt a simple model where the radio emission originates from a propagating (spherical) plasma 'blob' confined by the compact jet dimensions. We estimate the compact jet geometry and magnetic field strength using the model of MAXI J1820+70 from Zdziarski et al. (2022). The authors parameterize the magnetic field strength as $B = B_0 \xi^{-b}$, where $\xi = z_\nu / z_0 = (\nu/\nu_0)^{-q}$, z_ν is the down-stream distance along the jet axis at which the compact jet is optically-thin to frequency $\gtrsim \nu$, and $q = 0.88$. The radial width of the jet, $R = z_\nu \tan\theta$, at position z_ν, assuming a conical, Blandford & Königl (1979) geometry with opening angle θ. I adopt the following values from Zdziarski et al. (2022): $b = 1.1$; $q = 0.882$; $B_0 \sim 10^{10}\,\mu$G; $z_0 \sim 3 \times 10^{10}\,$cm; $\nu_0 \sim 2.3 \sim 10^4\,$GHz; $\theta \sim 1.5°$. Furthermore, I use the $\nu \sim 1.3\,$GHz when calculating B and R at z_ν. Given that the rotation measure returns to its average value in < 3 weeks, we calculate the radial width and magnetic field strength as the jet propagates (and expands), assuming ballistic motion with a bulk velocity equal to the speed of light (i.e., $z(t) = z_\nu + ct$).

The model predicts values of $B \sim 8 \times 10^5$ μG and $R \sim 4 \times 10^{12}$ at z_ν, with a radial expansion velocity of $\sim 0.03c$; the expansion velocity is consistent with prior observations of BHXB ejecta (Bright et al., 2020; Carotenuto et al., 2021). Modelling of radio flares associated with jet ejection events predicts similar strength internal magnetic fields ($\gtrsim 10^5\,\mu$G; Fender & Bright, 2019). To estimate the electron number density, we need an estimation of the ejecta mass, or, more specifically, the number of comprising electrons, $n_e = N_e/(4\pi R^3)$. Currently, the masses of ejecta are poorly constrained;

past analyses of BHXB ejection events span five orders of magnitude ($N_e = M_{\rm ej}/m_p$; $N_e \sim 10^{43}$–10^{48}, e.g., Fender et al., 1999a; Carotenuto et al., 2022, 2024; Espinasse et al., 2020), assuming an equal number of electrons and protons.

From Equation (5.5) and for the allowed values N_e, at z_ν the rotation measure is $\gtrsim 10^4\,\mathrm{rad\,m^{-2}}$ (and thus $\gg 8\ \mathrm{rad\,m^{-2}}$). Therefore, strong internal Faraday rotation originating from a jet ejection is plausible. However, the expansion of the blob will rapidly decrease the rotation measure. Maintaining an RM $\geq 8\ \mathrm{rad\,m^{-2}}$, even for a single week, requires $N_e > 10^{46}$. Future higher-cadence observations could better constrain the duration of the rotation measure elevation and, thus, the plausibility of internal Faraday rotation. To place these results in the broad context of jet physics, one of the major unknowns is composition, whether jets (or which jets) consist of electron-proton or electron-positron pairs, with the former having $1000\times$ the mass and, thus, significantly more kinetic energy at a fixed (bulk) velocity. Internal Faraday rotation from a pair-plasma requires an imbalance of electrons and positrons (favouring an electron-proton composition), as the direction of rotation is charge-dependent, and thus, equal numbers of electrons/positrons contribute equal rotations in opposite directions (Legg & Westfold, 1968; Wardle, 1977; Wardle et al., 1998). Therefore, detecting significant internal Faraday rotation may provide invaluable insight into the composition of jet ejecta and, thus, their total energy contents.

Long(er) Term Polarisation Fraction Evolution

Following the exotic polarisation behaviour during the bright Stokes I flare, the radio core continued 'dipping' in the polarisation fraction (see Fig. 5.11 and Fig. 5.12). Immediately following the flare, the core polarisation began decreasing again, reaching a minimum of $\sim 0.5\%$ on 2023 November 12 (MJD 60260; Fig. 5.11) before recovering to $\sim 7\%$ over the next two weeks.

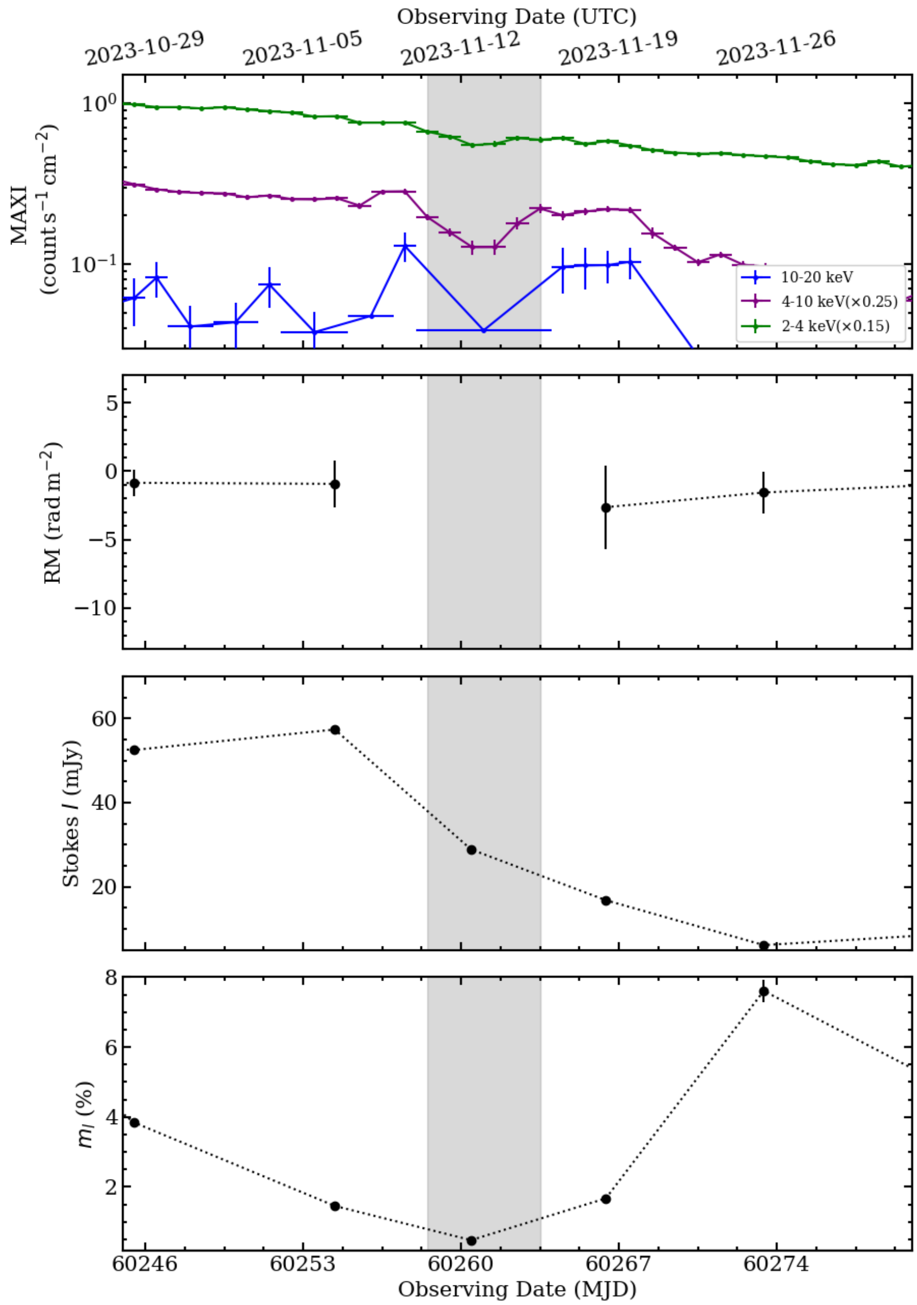

Figure 5.11: Same as Figure 5.10 except for the second dip. While there is no Stokes I flare, we still observe a softening of the X-ray spectrum during the decrease in polarisation fraction.

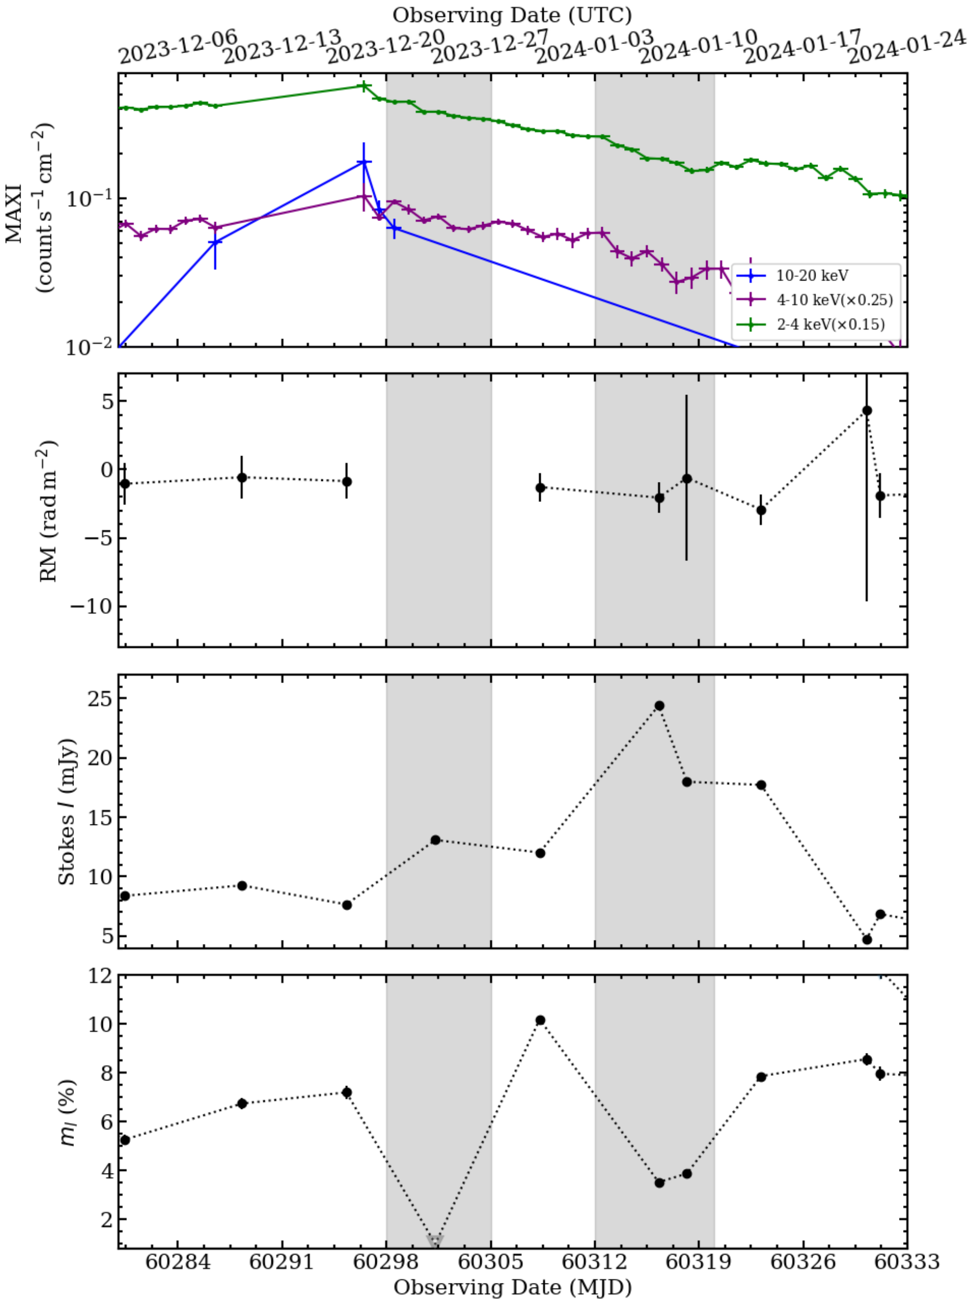

Figure 5.12: Same as Figure 5.10 except for the third and fourth dip.

While the decline in the source flux density (to $\sim 20\,\mathrm{mJy}$ for $m_l \sim 0.4\%$) prevented subsequent RM measurements (i.e., $P_0/\sigma_{QU} < 8$), we do, however, detect a drop in the X-ray luminosity that is more significant at higher energies (top panel, Fig. 5.11). While there is no clear flaring in Stokes I, the emission from bright ejecta launched during the $\sim$ Jy-flare could conceal any lower-level Stokes I evolution coincident with this second polarisation decrease.

The polarisation fraction dipped two more times (see Fig. 5.12), reaching minima of $\lesssim 1\%$ on 2023 December 23 (MJD 60301) and $\sim 4\%$ on 2024 January 9 (MJD 60318), both coinciding with increases in Stokes I, and the former followed the decay of an X-ray re-brightening at 2–20 keV. The final dip was more modest, decreasing to $\sim 4\%$ during the Stokes I increase. As a result, I could measure the rotation measure during this dip: RM$=$ $-2.1 \pm 1.0\,\mathrm{rad\,m^{-2}}$. While larger (in magnitude) than the average, the large error bars make the difference statistically insignificant. Given that the Stokes I flare is $\sim$40$\times$ dimmer than the large flare, I may be unable to detect a significant change in RM if, for example, the ejectum is less massive. Alternatively, the various polarisation dips could be caused by multiple mechanisms. Following the final dip, the core polarisation remained stable at $\sim 10\%$-level, until the soft-to-hard state transition, when the core no longer showed significantly polarised emission, with (unconstraining) upper limits of $\gtrsim 7\%$. The ejecta flux densities of $\lesssim 1\,\mathrm{mJy}$ did not meet our polarised detection criteria in most observations. The three observations with polarisation detections of E1 had polarisation fractions of $\sim 10\%$. The similarity between the core polarisation and the (few) ejecta detections is further evidence of the radio core being dominated by unresolved ejecta in the soft-state.

The linear polarisation angle also evolved significantly. In the lead-up to and during the bright flare, we saw multiple swings in the polarisation angle (i.e., magnetic field orientation) occurring on timescales $\lesssim 1$ week; $-10° \lesssim \mathrm{EVPA}_0 \lesssim 75°$. Similarly, the source showed an evolving polarisation angle during the second and fourth dips in the polarisation faction; for the former, the $\Delta\mathrm{EVPA}_0 \sim 50°$ between two consecutive observations. The dip in polarisation fraction that occurred on 2023 December 23 had insufficient polarised signal-to-noise to measure a polarisation angle. Alongside the short timescale variability, the polarisation angles are bi-modal, clustering around $\sim 15°$ values, and $\sim 100°$. Given the north-south position angle of the jet, these values are approximately parallel ($\sim 15°$) and perpendicular ($\sim 100°$) to the bulk velocity of the jet. The radio core showed a transition between the modes on (approximately) 2024 January 15 (MJD 60324), maintaining a consistent polarisation angle of $\sim 100°$ for the ~ 3 months before dropping below the detection threshold. Interestingly, the three detections of E1 were measured at $\mathrm{EVPA}_0 \sim 20°$, orthogonal to the core polarisation angle at late times but parallel to it early on. For emission originating from jet ejecta with similar absorption environments (i.e., both are optically-thin synchrotron emitters), the orthogonality in the polarisation angle would result from orthogonality in the magnetic field directions. Assuming a similar launching mechanism (perhaps naively), how two ejecta would acquire the orthogonal field directions remains unclear.

The repeated coincidence between the drops in polarisation and the X-ray *softening* (a known signature of jet ejections) strongly suggests these dips are associated with ejection events. As of yet, we have only discussed, in detail, a candidate explanation involving internal Faraday rotation when assuming large initial densities and strong magnetic fields. Past observations of anti-correlated m_l/Stokes I evolution have invoked alternative jet-interaction explanations. For example, Brocksopp et al. (2007) framed the behaviour in the context of the shock-in-jet framework introduced in Fender et al. (2004); collisions between transient jet ejecta with variables speeds temporarily disorder internal magnetic field lines (i.e., decrease polarisation fraction) but produce shock fronts that propagate through the ejecta, stimulating particle acceleration (i.e., increasing the Stokes I flux density), before re-establishing a dominant field direction. A similar explanation was invoked by Shahbaz et al. (2016) to explain the behaviour of the polarised optical evolution. Attempting to discern the origin of the polarisation behaviour is ongoing.

These results suggest that the observed polarisation behaviour may be utilized as signatures of jet ejections, thereby constraining ejection dates when contemporaneous X-ray observations are unavailable or insufficient. For Swift J1727, three (of the four) polarisation dips occurred during modest ($< 2\times$) or undetectable flux density evolution. Given that we have spatially resolved two ejecta and have strong evidence that the radio core is composed of additional ejecta, the singular Stokes I flare is unrepresentative of the number of ejection events. Accurate ejection dates are critical for modelling the ejecta energy, and thus mass, magnetic field strengths, among other fundamental jet properties (as shown in, e.g., Carotenuto et al., 2024). If confirmed, polarisation, as a tool for identifying ejection times, will prove invaluable for time-domain studies of astrophysical jets.

5.4 Swift J1727.8−1613: The Big Picture

The preceding chapter details X-KAT radio and X-ray monitoring of the BHXB Swift J1727.8−1613, describing the calibration, imaging, and analysis routines I developed and my measurements of systematic calibration accuracy and precision. I conclusively detect significant physical evolution(s). First, I observe Swift J1727 launching multiple jet ejecta that showed clear signs of deceleration; I defer higher detailed modelling of the proper motion of the ejecta to future work that will also integrate ATCA, e-Merlin, VLA, and VLBA data from projects led by collaborators. Second, I identify Swift

J1727 as a radio-quiet BHXB that shows a clear transition from a steep to flat relation in the L_R–L_X plane, it is only the BHXB third to have been identified with this behaviour. Third, I have identified, for the first time, extreme radio polarisation evolution coinciding with the softening of the X-ray spectrum, the latter being a commonly adopted signature of jet ejections. Moreover, I have detected, for the first time, a significant evolution of the rotation measure across a wide wavelength (squared) range. Moreover, this evolution coincides with an X-ray softening. If similar behaviour is established in other BHXBs, I will have established a new avenue to constrain jet composition.

While I have presented a substantial MeerKAT/*Swift*-XRT observing campaign, X-KAT and its affiliated collaborations have significantly more data on the source. In radio alone, these collaborations have multi-frequency observations with the VLA, ATCA, ATA, and AMI, with additional high angular resolution observations (i.e., at milliarcsecond precision) by very long baseline interferometers such as the VLBA, e-Merlin, and the LBA. As a result, we likely have $\geq 3\times$ the coverage in both time and frequency than what is presented in this thesis. Combining the totality of our radio observations, as well as the observations at other frequencies, we will be able to develop a coherent description of the evolution of Swift J1727.8−1613, improving constraints on ejection dates, our coverage in the L_R–L_X plane, and, with ATCA particularly, better understand the polarisation evolution.

5.5 Data Tables

Table 5.3: Properties of the 'radio core'

				Radio Core				
MJD	Date	ν (GHz)	Δt (min)	I (mJy)	α	m_l (%)	EVPA$_0$ (deg)	RM (rad m^{-2})
60183	2023 Sep 27	1.28	15	46.77 ± 0.03	$+0.199 \pm 0.003$	< 0.361	—	—
60191	2023 Sep 04	1.28	15	91.58 ± 0.042	-0.050 ± 0.002	0.44 ± 0.07	—	—
60193	2023 Sep 06	1.28	15	95.57 ± 0.05	-0.119 ± 0.003	< 0.33	—	—
60195	2023 Sep 08	1.28	15	86.64 ± 0.04	-0.085 ± 0.002	0.38 ± 0.07	—	—
60203	2023 Sep 16	1.28	15	85.15 ± 0.04	$+0.046 \pm 0.002$	1.12 ± 0.07	0 ± 3	-0.5 ± 1.3
60210	2023 Sep 23	1.28	15	101.87 ± 0.04	-0.152 ± 0.002	2.12 ± 0.07	75.2 ± 1.6	-1.4 ± 1.0
60218	2023 Oct 01	1.28	15	23.85 ± 0.02	-0.547 ± 0.006	4.60 ± 0.11	3 ± 2	-1.5 ± 1.1
60223	2023 Oct 06	1.28	15	182.03 ± 0.07	-0.352 ± 0.002	0.85 ± 0.07	41 ± 5	-10.2 ± 1.6
60231	2023 Oct 14	1.28	15	839.36 ± 0.3	-0.083 ± 0.002	0.30 ± 0.07	—	—
60233	2023 Oct 16	1.28	15	101.90 ± 0.05	-0.331 ± 0.003	1.69 ± 0.07	—	—
60239	2023 Oct 22	1.28	15	54.84 ± 0.03	-0.472 ± 0.003	7.03 ± 0.08	4.8 ± 1.4	-1.1 ± 1.0
60245	2023 Oct 28	1.28	15	52.48 ± 0.03	-0.478 ± 0.003	3.85 ± 0.08	15.3 ± 1.6	-0.8 ± 1.0
60254	2023 Nov 06	1.28	15	57.36 ± 0.03	-0.499 ± 0.003	1.458 ± 0.08	19 ± 5	-0.9 ± 1.7
60260	2023 Nov 12	1.28	15	28.85 ± 0.02	-0.607 ± 0.005	0.479 ± 0.09	—	—
60266	2023 Nov 18	1.28	15	16.78 ± 0.02	-0.643 ± 0.007	1.674 ± 0.12	-9 ± 10	-3 ± 3
60273	2023 Nov 25	1.28	15	6.12 ± 0.02	-0.77 ± 0.02	7.6 ± 0.3	4 ± 4	-1.6 ± 1.6
60280	2023 Dec 02	1.28	15	8.383 ± 0.018	-0.659 ± 0.013	5.3 ± 0.2	6 ± 4	-1.0 ± 1.5
60288	2023 Dec 10	1.28	15	9.26 ± 0.02	-0.537 ± 0.014	6.7 ± 0.2	16 ± 4	-0.6 ± 1.5

Continued on next page

Table 5.3 continued

Radio Core								
MJD	Date	ν (GHz)	Δt (min)	I (mJy)	α	m_l (%)	EVPA$_0$ (deg)	RM (rad m^{-2})
60295	2023 Dec 17	1.28	15	7.66 ± 0.02	-0.572 ± 0.017	7.2 ± 0.3	10 ± 3	-0.8 ± 1.3
60301	2023 Dec 23	1.28	15	13.08 ± 0.02	-0.45 ± 0.010	< 0.991	—	—
60308	2023 Dec 30	1.28	15	12.03 ± 0.02	-0.49 ± 0.010	10.19 ± 0.16	11 ± 2	-1.3 ± 1.0
60316	2024 Jan 07	1.28	15	24.39 ± 0.02	-0.476 ± 0.005	3.52 ± 0.10	35 ± 2	-2.1 ± 1.0
60318	2024 Jan 09	2.62	15	17.981 ± 0.014	-0.25 ± 0.010	3.88 ± 0.09	27 ± 4	-1 ± 6
60323	2024 Jan 14	1.28	15	17.72 ± 0.02	-0.520 ± 0.008	7.87 ± 0.14	94 ± 2	-2.9 ± 1.1
60330	2024 Jan 21	3.06	15	4.761 ± 0.012	-0.35 ± 0.04	8.6 ± 0.2	87 ± 7	4 ± 14
60331	2024 Jan 22	1.28	15	6.84 ± 0.02	-0.584 ± 0.019	8.0 ± 0.3	87 ± 5	-1.8 ± 1.6
60338	2024 Jan 29	1.28	15	5.303 ± 0.019	-0.70 ± 0.02	7.8 ± 0.4	94 ± 3	-1.6 ± 1.2
60344	2024 Feb 04	1.28	15	4.172 ± 0.019	-0.60 ± 0.03	6.6 ± 0.4	104 ± 7	-4 ± 2
60350	2024 Feb 10	1.28	15	3.35 ± 0.02	-0.64 ± 0.04	6.8 ± 0.6	86 ± 9	-1 ± 3
60359	2024 Feb 19	1.28	15	2.80 ± 0.02	-0.60 ± 0.04	10.5 ± 0.6	90 ± 7	-2 ± 2
60365	2024 Feb 25	1.28	15	2.20 ± 0.02	-0.60 ± 0.05	6.6 ± 0.8	99 ± 11	-1 ± 3
60373	2024 Mar 04	1.28	15	1.469 ± 0.019	-0.67 ± 0.08	5.4 ± 1.2	113 ± 7	—
60378	2024 Mar 09	1.28	15	1.106 ± 0.017	-0.79 ± 0.10	7.4 ± 1.5	99 ± 6	—
60385	2024 Mar 16	1.28	15	0.753 ± 0.019	-0.92 ± 0.15	< 11	—	—
60393	2024 Mar 24	1.28	15	1.54 ± 0.02	-0.63 ± 0.08	< 6	—	—
60399	2024 Mar 30	1.28	15	1.41 ± 0.02	-0.38 ± 0.09	< 6	—	—

Continued on next page

Table 5.3 continued

				Radio Core				
MJD	Date	ν (GHz)	Δt (min)	I (mJy)	α	m_l (%)	EVPA_0 (deg)	RM ($\mathrm{rad\,m^{-2}}$)
60407	2024 Apr 07	1.28	15	0.83 ± 0.02	-0.89 ± 0.15	< 11	—	—
60412	2024 Apr 12	1.28	15	0.66 ± 0.02	-1.01 ± 0.17	< 14	—	—
60421	2024 Apr 21	1.28	15	0.49 ± 0.02	—	< 18	—	—
60428	2024 Apr 28	1.28	15	0.343 ± 0.019	—	< 26	—	—
60435	2024 May 05	1.28	15	0.21 ± 0.02	—	< 44	—	—
60441	2024 May 11	1.28	30	0.098 ± 0.019	—	< 80	—	—
60443	2024 May 13	3.06	30	0.058 ± 0.010	—	< 86	—	—
60448	2024 May 18	1.28	15	0.099 ± 0.018	—	< 75	—	—
60455	2024 May 25	1.28	15	0.083 ± 0.017	—	< 100	—	—
60463	2024 Jun 02	1.28	15	0.051 ± 0.018	—	< 100	—	—
60469	2024 Jun 08	1.28	15	0.045 ± 0.015	—	< 100	—	—
60471	2024 Jun 10	3.06	15	0.028 ± 0.007	—	< 100	—	—
60476	2024 Jun 15	1.28	15	< 0.131	—	—	—	—
60483	2024 Jun 22	1.28	15	< 0.059	—	—	—	—
60490	2024 Jun 29	1.28	15	< 0.064	—	—	—	—
60492	2024 Jul 01	3.06	15	< 0.03	—	—	—	—

Notes . The column, Δt, corresponds to the on-source time for each observation; the remaining symbols retain their previous meanings from the main text.

Table 5.4: Radio properties of Ejectum 1 (E1)

				Ejectum 1 (E1)				
MJD	Date	ν	Δt	I	α	m_l	EVPA$_0$	RM
		(GHz)	(min)	(mJy)		(%)	(deg)	(rad m^{-2})
60330	2024 Jan 21	3.06	15	0.276 ± 0.012	—	< 23	—	—
60331	2024 Jan 22	1.28	15	0.532 ± 0.022	—	< 12	—	—
60338	2024 Jan 29	1.28	15	1.00 ± 0.019	-0.58 ± 0.12	< 8	—	—
60344	2024 Feb 04	1.28	15	1.081 ± 0.019	-1.33 ± 0.10	< 10	—	—
60350	2024 Feb 10	1.28	15	1.30 ± 0.02	-0.73 ± 0.09	< 7	—	—
60359	2024 Feb 19	1.28	15	0.89 ± 0.02	-0.84 ± 0.13	< 10	—	—
60365	2024 Feb 25	1.28	15	1.13 ± 0.02	-0.83 ± 0.10	< 10	—	—
60373	2024 Mar 04	1.28	15	0.941 ± 0.019	-1.11 ± 0.12	< 8	—	—
60378	2024 Mar 09	1.28	15	0.809 ± 0.017	-0.76 ± 0.14	< 13	—	—
60385	2024 Mar 16	1.28	15	0.751 ± 0.019	-0.92 ± 0.15	< 19	—	—
60393	2024 Mar 24	1.28	15	0.82 ± 0.02	-0.83 ± 0.15	9 ± 2	22 ± 9	—
60399	2024 Mar 30	1.28	15	0.82 ± 0.02	-1.11 ± 0.14	9 ± 2	10 ± 8	—
60407	2024 Apr 07	1.28	15	0.71 ± 0.02	-0.67 ± 0.17	< 14	—	—
60412	2024 Apr 12	1.28	15	0.71 ± 0.02	-1.01 ± 0.16	< 13	—	—

Continued on next page

Table 5.4 continued

Ejectum 1 (E1)								
MJD	Date	ν (GHz)	Δt (min)	I (mJy)	α	m_l (%)	EVPA_0 (deg)	RM ($\mathrm{rad\,m^{-2}}$)
60421	2024 Apr 21	1.28	15	0.57 ± 0.02	—	< 15	—	—
60428	2024 Apr 28	1.28	15	0.548 ± 0.019	—	< 13	—	—
60435	2024 May 05	1.28	15	0.46 ± 0.02	—	< 19	—	—
60441	2024 May 11	1.28	30	0.407 ± 0.019	—	< 19	—	—
60443	2024 May 13	3.06	30	0.188 ± 0.010	—	< 19	—	—
60448	2024 May 18	1.28	15	0.418 ± 0.018	—	< 16	—	—
60455	2024 May 25	1.28	15	0.339 ± 0.017	—	< 20	—	—
60463	2024 Jun 02	1.28	15	0.373 ± 0.018	—	< 17	—	—
60469	2024 Jun 08	1.28	15	0.465 ± 0.015	-0.8 ± 0.2	< 14	—	—
60471	2024 Jun 10	3.06	15	0.17 ± 0.007	—	16 ± 4	26 ± 6	—
60476	2024 Jun 15	1.28	15	0.29 ± 0.04	—	< 61	—	—
60483	2024 Jun 22	1.28	15	0.09 ± 0.02	—	< 83	—	—
60490	2024 Jun 29	1.28	15	0.27 ± 0.02	—	< 24	—	—
60492	2024 Jul 01	3.06	15	0.129 ± 0.010	—	< 34	—	—

Notes . The column, Δt, corresponds to the on-source time for each observation; the remaining symbols retain their previous meanings from the main text.

Table 5.5: Radio properties of Ejectum 2 (E2).

Ejectum 2 (E2)								
MJD	Date	ν (GHz)	Δt (min)	I (mJy)	α	m_l (%)	$EVPA_0$ (deg)	RM ($\mathrm{rad\,m^{-2}}$)
60441	2024 May 11	1.28	30	0.104 ± 0.019	—	< 51	—	—
60443	2024 May 13	3.06	30	0.111 ± 0.010	—	< 39	—	—
60448	2024 May 18	1.28	15	0.145 ± 0.018	—	< 36	—	—
60455	2024 May 25	1.28	15	0.241 ± 0.017	—	< 33	—	—
60463	2024 Jun 02	1.28	15	0.288 ± 0.018	—	< 20	—	—
60469	2024 Jun 08	1.28	15	0.154 ± 0.015	—	< 30	—	—
60471	2024 Jun 10	3.06	15	0.158 ± 0.007	—	< 23	—	—
60476	2024 Jun 15	1.28	15	0.36 ± 0.04	—	< 47	—	—
60483	2024 Jun 22	1.28	15	0.42 ± 0.02	—	< 20	—	—
60490	2024 Jun 29	1.28	15	0.21 ± 0.02	—	< 54	—	—
60492	2024 Jul 01	3.06	15	0.130 ± 0.010	—	< 29	—	—

Notes . The column, Δt, corresponds to the on-source time for each observation; the remaining symbols retain their previous meanings from the main text.

Table 5.6: X-ray spectral parameters from modelling the *Swift*-XRT observations.

Swift-XRT X-ray Properties							
MJD	Date	$F_{X,\mathrm{PL}}$ ($10^{-11}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$)	Γ	χ^2_{red}(dof)	$F_{X,\mathrm{BB}}$ ($10^{-11}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$)	T_{in} (keV)	χ^2_{red}(dof)
60343	2024 Feb 03	—	—	21.1(312)	923^{+9}_{-9}	$0.597^{+0.005}_{-0.005}$	1.05(312)
60350	2024 Feb 10	—	—	27.6(294)	760^{+8}_{-8}	$0.595^{+0.005}_{-0.005}$	1.70(294)
60357	2024 Feb 17	—	—	36.7(308)	644^{+5}_{-5}	$0.573^{+0.004}_{-0.004}$	1.79(308)
60364	2024 Feb 24	—	—	21.7(271)	569^{+6}_{-6}	$0.542^{+0.005}_{-0.005}$	1.01(271)
60367	2024 Feb 27	—	—	49.4(334)	445^{+3}_{-3}	$0.534^{+0.003}_{-0.003}$	1.57(334)
60379	2024 Mar 10	—	—	53.5(365)	389^{+2}_{-2}	$0.490^{+0.002}_{-0.002}$	1.54(365)
60403	2024 Apr 03	$159.5^{+1.7}_{-1.7}$	$2.075^{+0.013}_{-0.013}$	1.9(552)	—	—	14.5(552)
60407	2024 Apr 07	123^{+2}_{-2}	$1.83^{+0.02}_{-0.02}$	1.3(414)	—	—	6.355(414)
60416	2024 Apr 16	$90.3^{+1.3}_{-1.3}$	$1.708^{+0.017}_{-0.017}$	1.2(465)	—	—	5.13(465)
60420	2024 Apr 20	$67.5^{+1.7}_{-1.7}$	$1.767^{+0.03}_{-0.03}$	1.1(295)	—	—	3.25(295)
60423	2024 Apr 23	$62.2^{+1.1}_{-1.1}$	$1.62^{+0.02}_{-0.02}$	1.0(399)	—	—	3.22(399)
60430	2024 Apr 30	$9.2^{+0.7}_{-0.7}$	$1.50^{+0.07}_{-0.07}$	1.1(41)	—	—	1.5(41)
60434	2024 May 4	$1.5^{+0.2}_{-0.2}$	$1.6^{+0.2}_{-0.2}$	0.73(10)	—	—	1.2(10)
60437*	2024 May 7	$0.37^{+0.09}_{-0.08}$	$1.65^{+0.2}_{-0.2}$	1.1(103)	—	—	1.1(103)
60439*	2024 May 9	$0.18^{+0.07}_{-0.05}$	$1.69^{+0.4}_{-0.3}$	1.3(52)	—	—	1.4(52)

Notes. All errors represent the 90% confidence intervals. The MJD date with asterisks (*) used Cash statistics. Since Cash statistics use single-count binning and thus each bin has non-Gaussian errors, the reported χ^2_{red} statistics are not minimized during fitting. Regardless, I include the χ^2 statistic to qualitatively represent the goodness of fit even for the Cash statistic epochs.

Table 5.7: X-ray fluxes for the observations omitted from our spectral modelling

Swift-XRT Low Count Flux			
MJD	Date	F_X ($10^{-11}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$)	χ^2_{red}(dof)
60441*	2024 May 11	$0.05^{+0.04}_{-0.03}$	0.6(8)
60444*	2024 May 14	< 0.26	—
60450*	2024 May 18	< 0.21	—
60452*	2024 May 22	$0.08^{+0.03}_{-0.02}$	1.0(20)
60456*	2024 May 26	$0.025^{+0.030}_{-0.016}$	0.9(3)
60463*	2024 Jun 02	$0.017^{+0.014}_{-0.009}$	0.8(6)

Notes. These epochs were omitted from spectral fitting due to their low number of counts count rate. The reported fluxes assume a power-law model with $\Gamma \equiv 1.7$. All errors represent the 90% confidence intervals. The upper limits are 3σ significance. All fits use Cash statistics.

Chapter 6

Astrometric Precision with MeerKAT

This final Chapter is an expansion (and improvement) on the astrometric routine introduced in Section 3.6.1 (Originally presented in, Hughes et al., 2024). While the properties of SAX J1810 initiated this investigation, as discussed in the following sections, understanding the astrometric evolution of radio core emission Swift J1727 necessitated precise astrometry. Similar to POLKAT, despite the routine being developed to study XRBs, it applies to any point source observations.

6.1 Introduction

Section 2.3.1 introduced several fundamental angular scales, one of which was the angular resolution:

$$\Theta_{\mathrm{PSF}} \sim \frac{\lambda}{d_{\mathrm{max}}}.$$

Quantified by the FWHM (Θ_{PSF}) of the PSF[1], the angular resolution is related to the observing wavelength (λ) and the maximum baseline length ($d_{\max}$). While a reasonable first-order approximation, this equation assumes a circular Gaussian PSF and, thus, rotationally symmetric sampling of the *uv*-plane. Complex array configurations or non-Zenith observations can cause the projected baseline vectors to be shorter along one dimension, resulting in elliptical Gaussian PSFs. Appropriate *uv*-weightings can circularize the PSF, even for asymmetric array configurations, at the cost of angular resolution along the more precise direction. Therefore, utilising an interferometer's full astrometric capabilities necessitates considering PSF ellipticity.

By default, radio synthesis imaging software records the PSF parameters with the FWHM of the major axis (Θ_{maj}), the FWHM of the minor axis (Θ_{min}), and the beam position angle (θ_{bpa}) measured East of North. Given that the PSF is an elliptical Gaussian, the PSF axes can be alternatively expressed as standard deviations, σ_{maj} and σ_{min}, recalling that FWHM= $\sqrt{8\ln 2}\,\sigma$. The PSF is an instrument's point source response, so the ellipticity manifests as the shape of point sources (e.g., Figure 6.1). To maximize consistency with terminology used at other wavelengths, I adopted the more generic term of a point spread function (PSF) rather than the interferometry-specific terminology of a 'synthesized beam'. Depending on the resource, what I call a PSF may be referred to as the restoring beam, idealized beam, or synthesized beam; these all correspond to the same concept: the two-dimensional Gaussian fit to the main lobe of the Fourier transform of the *uv*-sampling function.

The position of a point source is given by the centre of its (PSF-shaped) intensity distribution. Adopting the location of the peak pixel can provide a crude position approximation. However, the peak pixel is affected by pixelization errors if the *actual* peak is situated at a pixel boundary. Pix-

[1]Here, I am referencing the idealized PSF, free from finite-sampling effects.

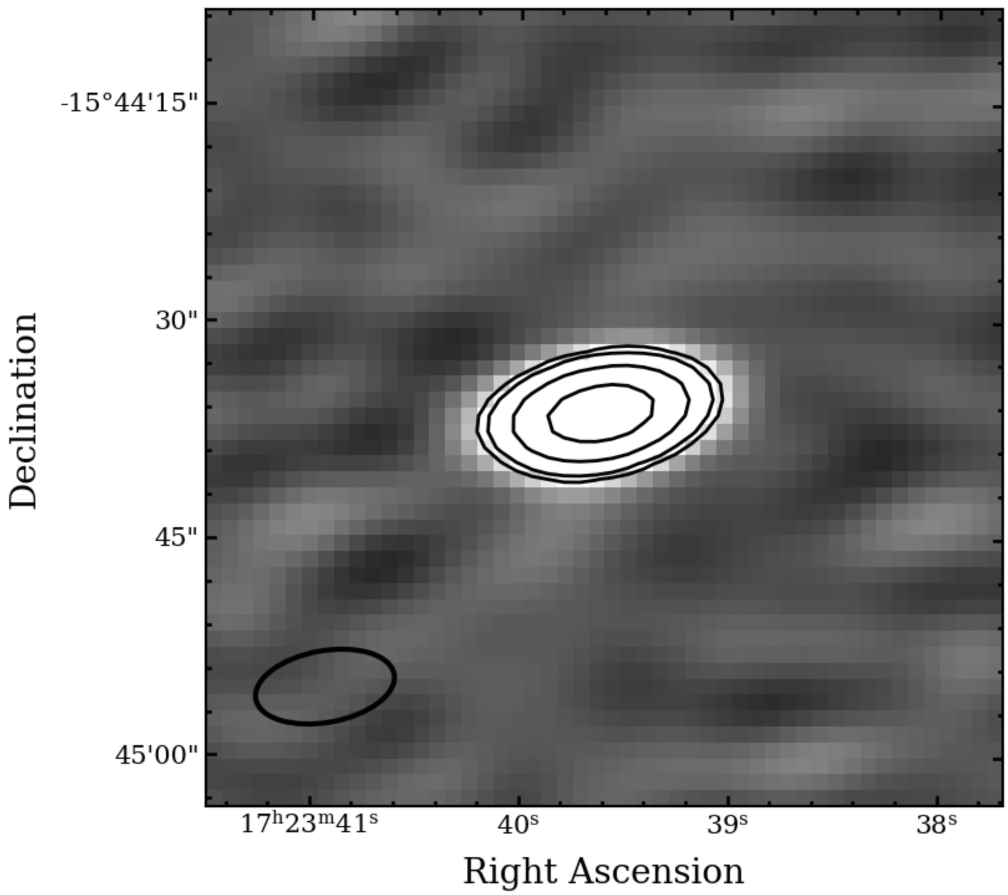

Figure 6.1: Example of a highly elliptical synthesized beam (hollow black ellipse) and its effect on the shape of a point source. In this example, the source was observed at a shallow elevation of $\sim 30°$ and an azimuth angle of $\sim 90°$. In this example, the relatively low elevation causes relatively high ellipticity, and the azimuth angle rotates the PSF to (roughly) align its major axis with the right ascension axis.

elization errors can be mitigated by decreasing the pixel size, although this approach is prohibitively computationally expensive for interferometers with large fields of view. As a result, the positions of point sources are typically measured by fitting Gaussian components to their intensity distributions, the approach adopted in the previous chapters. Assuming that the pixel size adequately samples the PSF (i.e., $\Theta_{\rm min} \sim 3$–5 pixels), Gaussian fitting also mitigates pixelization effects.

Kaper et al. (1966) derived analytic equations for the key parameters and their errors (peak amplitude, central position, and width) measured from one-dimensional Gaussian fitting, assuming uncorrelated Gaussian noise in each pixel. Later, Condon (1997, CR97) derived analytic solutions for two-dimensional data with uncorrelated noise and semi-analytic equations for spatially correlated noise (see also, Condon et al., 1998); the latter applies to synthesis imaging as the PSF acts as a smoothing kernel correlating adjacent pixels. CR97 showed that the variance in the position,

$$\sigma_k^2 = \left(\frac{\mathcal{D}_k}{S/N}\right)^2, \tag{6.1}$$

is a function of the signal-to-noise ratio (S/N) and the *extent* of the Gaussian ($\mathcal{D}_k$) along direction k (where k is RA or Dec). The extent is a function of the PSF shape, $\mathcal{D}_k = \mathcal{D}_k(\Theta_{\rm maj}, \Theta_{\rm min}, \theta_{\rm pa})$, where elliptical PSFs have $\mathcal{D}_{\rm RA} \neq \mathcal{D}_{\rm Dec}$ (see Equations 21 and 41 in CR97). For a circular PSF ($\Theta_{\rm maj} \equiv \Theta_{\rm min}$), Equation (6.1) simplifies to,

$$\sigma_k = \frac{\Theta_{\rm maj}}{\sqrt{8\ln 2}(S/N)} = \frac{\sigma_{\rm maj}}{S/N}. \tag{6.2}$$

Some source-finding software packages, like PYBDSF, adopt the CR97 equations to quantify astrometric errors. However, in CR97, the variance was assumed to originate from Gaussian noise, ignoring the effects of sparse uv-sampling or incomplete deconvolution of extended sources, where artefacts produce large-scale non-Gaussian “noise”. Moreover, Equations (6.1) and (6.2) neglect the potential for signal-to-noise independent systematics from, for example, calibration errors. As a result, many radio astronomers adopt a heuristic astrometric error of the form

$$\sigma_k = \mathcal{D}_k \sqrt{\left(\frac{A}{S/N}\right)^2 + B^2}. \tag{6.3}$$

Consistent with CR97, Equation (6.3) is proportional to the PSF size, scaling inversely with the fitted component's signal-to-noise ratio, while now including a systematic *floor* on the astrometric error. The next question is obvious: what are the correct choices for $\mathcal{D}_k$, A, and B?

Recommendations for A and $\mathcal{D}_k$ can vary depending on who is asked and often are provided without clear justification[2]. For $\mathcal{D}_k$, a common, conservative assumption is $\mathcal{D}_{\mathrm{RA}} \equiv \mathcal{D}_{\mathrm{Dec}} \equiv \Theta_{\mathrm{maj}}$. Alternatively, the ellipticity is typically accounted for by using the size of the rectangle bounding the PSF,

$$\mathcal{D}_{\mathrm{RA}} = \sqrt{(\Theta_{\mathrm{maj}} \sin\theta_{\mathrm{bpa}})^2 + (\Theta_{\mathrm{min}} \cos\theta_{\mathrm{bpa}})^2}, \tag{6.4}$$

$$\mathcal{D}_{\mathrm{Dec}} = \sqrt{(\Theta_{\mathrm{maj}} \cos\theta_{\mathrm{bpa}})^2 + (\Theta_{\mathrm{min}} \sin\theta_{\mathrm{bpa}})^2}, \tag{6.5}$$

or the lengths from the centre to the ‘edge’ of the PSF,

$$\mathcal{D}_{\mathrm{RA}} = \frac{\sigma_{\mathrm{maj}}\sigma_{\mathrm{min}}}{\sqrt{(\sigma_{\mathrm{maj}} \sin\theta_{\mathrm{bpa}})^2 + (\sigma_{\mathrm{min}} \cos\theta_{\mathrm{bpa}})^2}}, \tag{6.6}$$

$$\mathcal{D}_{\mathrm{Dec}} = \frac{\sigma_{\mathrm{maj}}\sigma_{\mathrm{min}}}{\sqrt{(\sigma_{\mathrm{maj}} \cos\theta_{\mathrm{bpa}})^2 + (\sigma_{\mathrm{min}} \sin\theta_{\mathrm{bpa}})^2}}. \tag{6.7}$$

For A, the most common values adopted in the literature are $A = 1$ and $A = 1/2$, where the latter likely originates from the fact that

[2]Moreover, many astronomers do not add the A and B components in quadrature, instead adopting the larger of the two.

$1/\sqrt{8\log 2} \approx 1/2.355 \sim 1/2$. To remain approximately consistent with CR97, $A = 1$ **should** be chosen if using *radial* $\mathcal{D}_k$ values (e.g., Equation 6.6 & 6.7) and $A = 1/2$ if using *diameter-like* $\mathcal{D}_k$ values (e.g., Equation 6.4 & 6.5). Incorrect pairings can overestimate errors, potentially leading to missed results, or underestimate errors, artificially increasing the significance of any astrometry-driven phenomenon.

The systematic limit, B, does not have the same theoretical motivation as A and thus can vary significantly from instrument to instrument. The VLA, possibly the best-known interferometer, recommends a systematic error of $\sim 10\%$ the PSF FWHM[3]. Newer facilities, like MeerKAT, do not have similar recommendations, leading a widespread *ad hoc* adoption of the $\sim 10\%$ VLA systematic. Investigating what systematic limit is appropriate for MeerKAT was the initial motivation for this research.

A last, often overlooked effect is the correlated RA and Dec errors that result from PSFs with principle axes misaligned to these primary directions (i.e., $\theta_{\rm pa} \neq 0$, $\pi/2$, π...). Consider an elliptical PSF with $\theta_{\rm pa} = \pi/4$; a same-sign deviation (e.g., towards North-East) in RA and Dec is more likely than an opposite-sign deviation (e.g., towards North-West) of equal magnitude. Framing this effect probabilistically, when sampling from a two-dimensional Gaussian, for a fixed radial separation ($r = \sqrt{x^2 + y^2}$), stochastic deviations along the beam position angle are the most likely to occur. While CR97 appropriately treated the co-variances, the heuristic forms of $\mathcal{D}_k$ given by Equations (6.4), (6.5), (6.6), & (6.7) will not enclose the same confidence region for all elliptical PSFs.

This chapter defines a simple yet powerful method to measure astrometric error. Wide-field interferometers like MeerKAT detect thousands of background or foreground point sources in a single observation. The variations in the positions observed during time-domain monitoring of a

[3]`https://science.nrao.edu/facilities/vla/docs/manuals/oss/performance/positional-accuracy`

single pointing (i.e., *field*) will sample from the underlying astrometric error distribution, allowing for its empirical determination. Ideally, with a large enough sample of targets and pointing, we should better understand the astrometric precision without needing time-domain monitoring of each source.

6.2 Methods

Section 3.6.1 presented an early version of this routine, originally designed to determine the astrometric variability of SAX J1810. Here, I assume the astrometric error adopts a similar functional form as the published version,

$$\hat{\sigma}_r = \sqrt{\left(\frac{A}{S/N}\right)^2 + B^2}. \tag{6.8}$$

The fundamental assumption (based on CR97) is that the astrometric error is proportional to some angular scale $\mathcal{D}_r(\Theta_{\mathrm{maj}}, \Theta_{\mathrm{min}}, \theta_{\mathrm{pa}})$. To account for variable PSFs, I solve for a dimensionless 'relative astrometric error' $\hat{\sigma}_r = \sigma_r/\mathcal{D}_r$; where $\hat{\sigma}_r$ is the astrometric error measured in units of 'PSFs', and can be converted to physical units given that $\mathcal{D}_r$ is always known. In contrast with the published version, I have refined the method to account for correlated errors by solving for total angular offsets rather than decomposing the offsets into RA and Dec components (hence the change from index $\mathcal{D}_k$ to $\mathcal{D}_r$).

6.2.1 Offset Geometry

A schematic of the offset geometry is shown in Fig. 6.2. For a single measurement of a point source, stochastic variability offsets the measured position (α, δ) from its true position (α_0, δ_0) by some angular distance r and an offset direction defined by the *offset position angle* (ϕ_{pa}) measured East

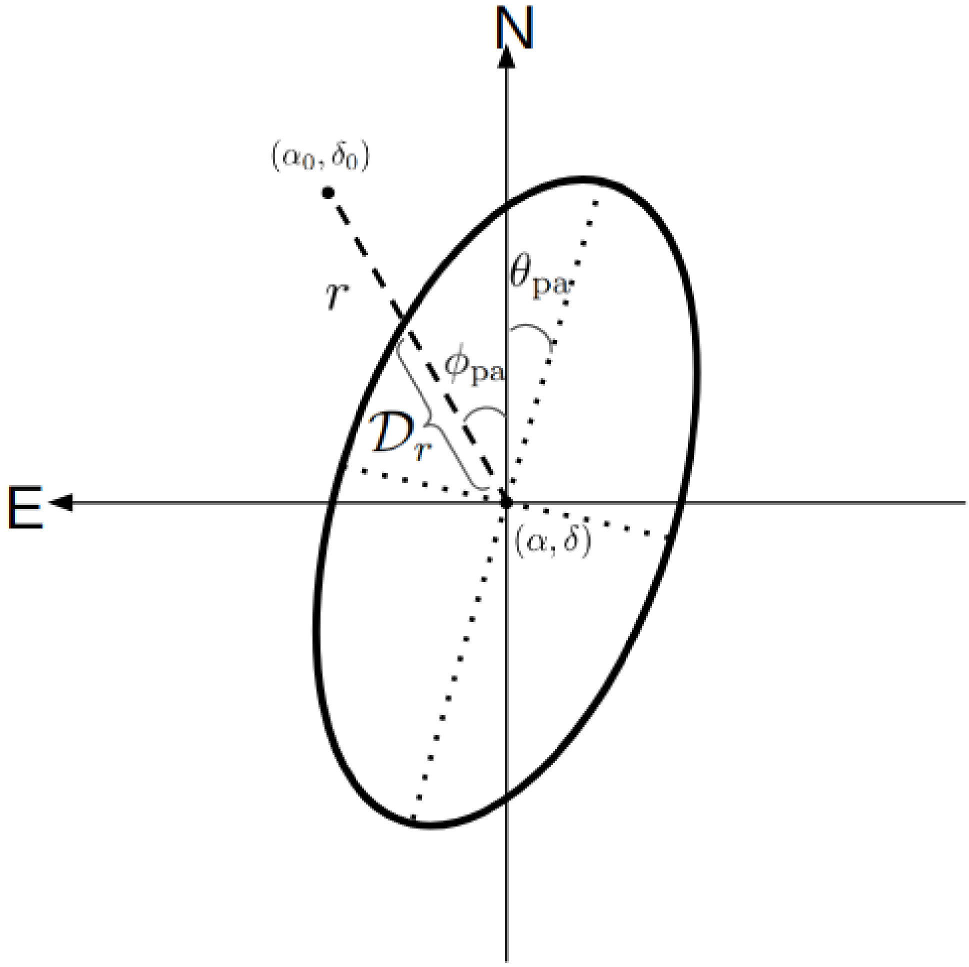

Figure 6.2: Schematic representation of the offset geometry. The measured position, $(\alpha,\ \delta)$, is offset from the true position, $(\alpha_0,\ \delta_0)$, by an angular separation $\hat{r}$. The r is rotated East of North by the offset position angle $\phi_{\rm pa}$. $\mathcal{D}_r$ is the distance from $(\alpha,\ \delta)$ to the elliptical confidence region defined by the PSF. The position angle of the elliptical confidence is given by $\theta_{\rm pa}$, but the major and minor axis (represented by the dotted lines) are dependent on the confidence interval of interest (e.g., 68%, 95%, 99.7%, ...)

of North. Similar to the dimensionless astrometric error, the dimensionless separation, $\hat{r} = r/\mathcal{D}_r$, is normalized by the distance from (α, δ) to the elliptical confidence region described by the PSF-shape (shown as a black ellipse in Fig. 6.2)

For an elliptical Gaussian with variances of σ^2_{maj} and σ^2_{min}, and no covariance, an arbitrary confidence region is described by,

$$s = \left(\frac{\Delta_{\text{maj}}}{\sigma_{\text{maj}}}\right)^2 + \left(\frac{\Delta_{\text{min}}}{\sigma_{\text{min}}}\right)^2. \tag{6.9}$$

here, Δ_{maj} and Δ_{min} are the zero-point offsets along the major and minor axes, and s depends on the confidence region of interest. Equation (6.9) is the sum of the squares of independent normally distributed variables, and thus s follows the χ^2_ν-distribution for two degrees of freedom ($\nu = 2$),

$$P(\chi^2_{\nu=2} \leq s) = P_0, \tag{6.10}$$

for a confidence region probability, P_0. We choose $P_0 = 0.68$ (1σ), and thus $s \approx 2.279$ given a χ^2_ν-distribution with two degrees of freedom. However, the choice is arbitrary, and different P_0 values will scale A and B appropriately, as long as P_0 remains consistent for observations with different PSF shapes. The confidence region is an ellipse with major and minor axes equal to $\sqrt{s}\,\sigma_{\text{maj}}$ and $\sqrt{s}\,\sigma_{\text{min}}$, respectively. Recognizing that the line connecting (α_0, δ_0) and (α, δ) passes through the ellipse at an angle of $\xi = \phi_{\text{pa}} - \theta_{\text{pa}}$,

$$\Delta_{\text{min}} = \Delta_{\text{maj}} \tan\xi, \tag{6.11}$$

where Δ_{maj} and Δ_{min} are the offset angles along the major and minor axis, respectively. Thus, Equation (6.9) becomes:

$$\Delta_{\text{maj}} = \frac{\sqrt{s}\,\sigma_{\text{maj}}\,\sigma_{\text{min}}}{\sqrt{\sigma^2_{\text{min}} + (\sigma_{\text{maj}}\tan\xi)^2}}, \tag{6.12}$$

and $\mathcal{D}_r$ is defined as,

$$\mathcal{D}_r = \sqrt{\Delta_{\mathrm{maj}}^2 + \Delta_{\mathrm{min}}^2}. \tag{6.13}$$

Knowledge of the measured position, true position, and PSF properties is sufficient to solve for r, $\mathcal{D}_r$, and thus, $\hat{r}$. As $\hat{r}$ is positive definite, repeated observations of an arbitrary source j will sample a Rayleigh distribution with a scale factor equivalent to the dimensionless astrometric error, $(\hat{\sigma}_r)_j$. For a source observed over N epochs, we can estimate its astrometric error from the sample distribution (Siddiqui, 1964),

$$(\hat{\sigma}_r)_j \approx \frac{\Gamma(N)\sqrt{N}}{\Gamma(N+\frac{1}{2})}\sqrt{\frac{1}{2N}\sum_{i=1}^{N}\hat{r}_{i,j}^2}. \tag{6.14}$$

Monitoring several sources at different signal-to-noise ratios is sufficient to derive A and B from Equation (6.8).

6.2.2 Workflow

I developed the following routine that uses the positional variability of field sources to determine an astrometric error in time-domain observations. Assuming a single image has been made for each observation, we use PYBDSF to make per-observation source catalogues. PYBDSF models the sky intensity distribution as a collection of Gaussian *components*. Components are then grouped into *sources* based on their relative separations and intensities[4]. The sources are then associated with *islands* — contiguous regions of pixels with values greater than some threshold value (`thresh_isl` in PYBSDF), and at least one pixel greater than the source detection threshold (`thresh_pix` in PYBSDF); both thresholds are expressed in units of rms devi-

[4]see `https://pybdsf.readthedocs.io/en/latest/algorithms.html` for more information

ations from the mean background intensity. I chose values of `thresh_isl=3` and `thresh_pix=4`.

The PYBDSF source catalogues include the peak intensity (F_p), the island flux density (F_i), and the source sizes quantified by major (Θ_M) and minor (Θ_m) axes. PYBDSF also classified each source as S- (single-component source, single-source island), C- (single-component source, multi-source island), or M-type (multi-component source).

I am interested in the astrometric error for the point source fitting. Thus, I remove resolved sources from the catalogues. To make it easier to identify point sources, the shapes of all Gaussian components are fixed to that of the PSF (`fix_to_beam=True` in PYBDSF), and components within an island are grouped into a single source (`group_by_isl=True` in PYBDSF). These flags make the classification binary, with S-types and (C or M)-types corresponding to point sources and resolved sources, respectively. By definition, all S-type sources will follow the strict definition of a point source; $F_i \sim F_p$[5], $\Theta_M \equiv \Theta_{\mathrm{maj}}$, and $\Theta_m \equiv \Theta_{\mathrm{min}}$. However, noise fluctuations or imaging artefacts can cause PYBDSF to fit additional sub-dominant components, causing some point sources to be misclassified as resolved. To avoid removing an excessive number of sources, I adopt an empirically motivated "point-like" definition,

$$\begin{aligned} \frac{F_i}{F_p} - 1 &< \delta_F, \\ \left| \frac{\Theta_M}{\Theta_{\mathrm{maj}}} - 1 \right| &< \delta_S, \text{ and} \\ \left| \frac{\Theta_m}{\Theta_{\mathrm{min}}} - 1 \right| &< \delta_S, \end{aligned}$$

allowing small tolerances, δ_F and δ_s. I chose to set $\delta_F = \delta_S = 0.1$ (i.e., PYBDSF measures the source to be within 10% of an ideal point source).

[5]The reason why the flux condition is left as an approximation is due to the definition of an Island. Pixel variations in each island can result in negligible (but finite) differences between the integrated island flux density and peak intensity

I apply an off-axis cutoff, excluding sources $> 0.3°$ from the phase centre, corresponding to a (normalized) primary beam power of $\mathcal{A} > 75\%$. This exclusion of these off-axis sources minimizes the errors from direction-dependent effects, which can affect the quality of both phase (i.e., position) and amplitude calibration (see, e.g., Smirnov, 2011a,b, and references therein). Lastly, I exclude highly variable and transient sources that would exhibit significant changes in their signal-to-noise ratios. Any sources with maximum and minimum peak intensities separated by a factor ≥ 2 or missing from $\geq 25\%$ of the epochs are removed from each catalogue. The astrometry routine proceeds as follows:

Step 1: With PYBDSF, extract a catalogue of each observation, removing sources that fail the point-like conditions. The catalogues are cross-matched to one another, adopting a single observation as the reference. Sources separated by less than $\Theta_{\mathrm{maj}}/3$ are taken as matches. The final catalog consists of n_{obs} observations of n_{src} sources, such that each source j has a position $(\alpha_{i,j},\, \delta_{i,j})$ and a signal-to-noise of $(S/N)_{i,j}$, in observation i.

Step 2: Calculate the average position of each source (average over index j), weighting each observation with the size of the PSF along the RA and Dec directions (to the 68% confidence region); i.e., inverse-variance weightings, $\mathcal{D}_{\mathrm{RA}}^{-2}$ and $\mathcal{D}_{\mathrm{Dec}}^{-2}$, respectively.

Step 3: Bootstrap the positions to estimate the uncertainty on $(\hat{\sigma}_r)_j$. For each source, create n_{boot} bootstrapped samples of length n_{obs}, resampling (with replacement) the positions from each observation. Assuming the average value from (2) represents the true position of each source, calculate the dimensionless astrometric errors in each bootstrap sample. I used $n_{\mathrm{boot}} = 1000$.

Step 4: For each source, adopt the median value from the bootstrapped samples as $(\hat{\sigma}_r)_j$, quantifying the 1σ confidence interval as half the range between the $16^{\rm th}$ and$84^{\rm th}$ percentile. The routine outputs the ratio of positive (median to $84^{\rm th}$ percentile) and negative ($16^{\rm th}$ to median) errors to investigate whether the errors are asymmetric; asymmetric errors (i.e., if the ratio is not close to unity) will invalidate the likelihood function as it assumes normally distributed variables.

Step 5: Calculate the median signal-to-noise ratio for each source, $(S/N)_j$.

Step 6: Solve for the astrometric error parameters A and B. We solve for A and B with a similar MCMC approach as applied in Section 3.2.3. Assume $(\hat{\sigma}_r)_j$ values are independent and normally distributed, adopting flat uninformative priors for A and B, with the only constraint being that $A \geq 0$ and $B \geq 0$ (we chose an arbitrary, large upper bound). These A and B values quantify the *uncorrected* astrometric error.

Step 7: Solve for the per-observation average astrometric correction (average across index i) in RA and Dec. The averages are inverse-variance weighted, using the parameters from Step 6 and Equation 6.8.

Step 8: Correct the positions with the per-observation corrections and repeat Steps 1–7 until the corrections converge (see Step 9). The uncertainty on the per-observation correction is included in each subsequent bootstrap subsample. This is done by adding a random offset drawn from a normal distribution with a standard deviation fixed to the uncertainty in the per-observation correction.

Step 9: I define the convergence parameter C between run n and $n+1$ as:

$$C = \frac{|(\hat{\sigma}_{r,j})_{n+1} - (\hat{\sigma}_{r,j})_n|}{\sqrt{(\Delta\hat{\sigma}_{r,j})^2_{n+1} + (\Delta\hat{\sigma}_{r,j})^2_n}}.$$

I consider convergence to have occurred after $C < 0.1$ for three consecutive iterations. The final values of A and B quantify the *corrected* astrometric error.

I made this routine (ASTKAT) publicly available[6]. It also allows the user to simulate images for code verification, as well as modify the parameters that control cataloguing, filtering, bootstrapping, and convergence. I solve for $\mathcal{D}_{RA}$ and $\mathcal{D}_{Dec}$ by marginalizing the PSF along the RA and Dec directions, respectively. I then use these values to estimate errors on individual RA and Dec measurements.

6.3 Results

The astrometric fit for the three XRB fields discussed in this thesis are presented in Figures 6.3– 6.5. Additionally, I tabulate the fit parameters in Table 6.1, including the reduced statistic χ^2_{red} and a 'fractional FWHM equivalent' ($B \times \sqrt{s/8\ln 2}$) to quantify the systematic astrometric limit as a fraction of the PSF FWHM.

[6] `https://github.com/AKHughes1994/AstKAT`

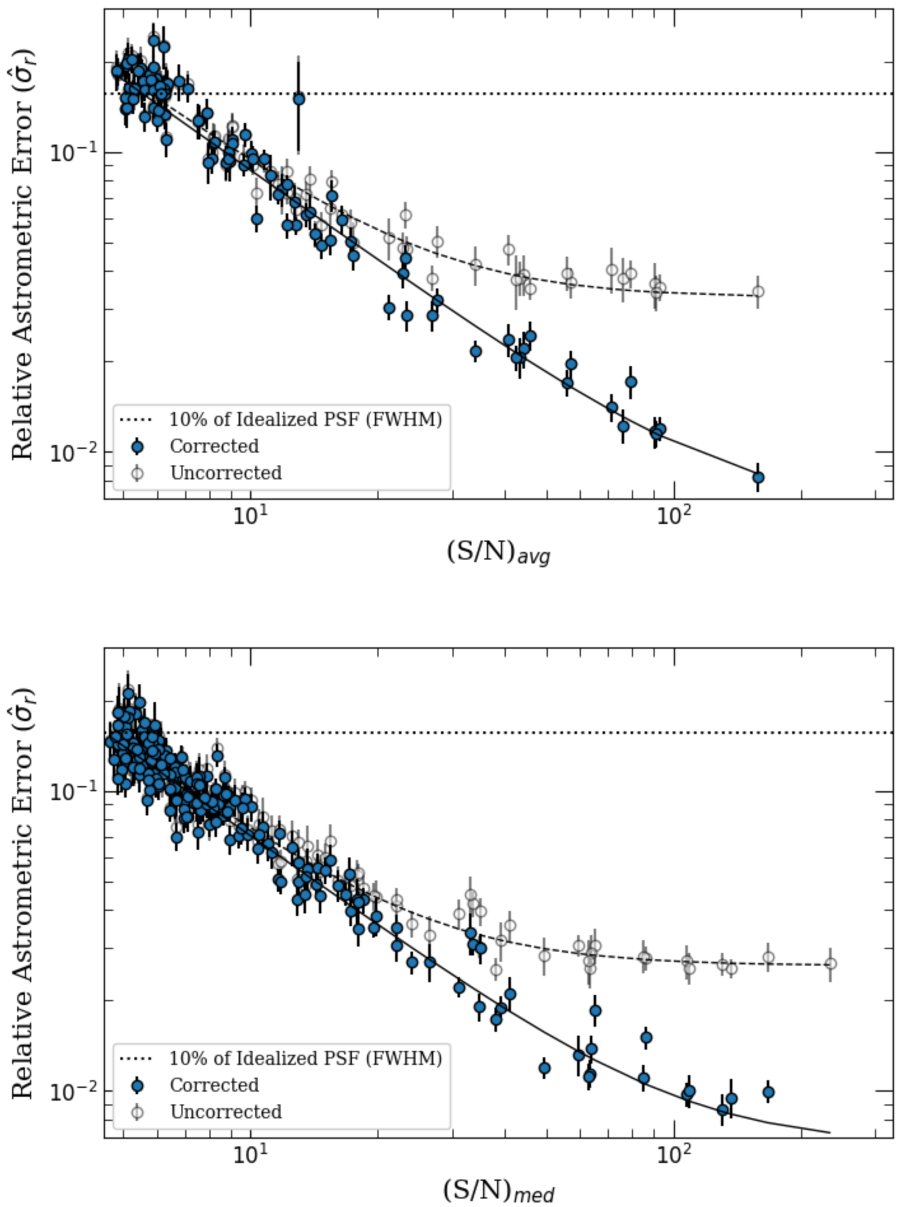

Figure 6.3: Relationship between the relative astrometric error ($\hat{\sigma}_r$) and the median signal-to-noise ratio $(\mathrm{S/N})_{med}$ for the field containing SAX J1810.8−2609. We include the plots for uniform (top) and Briggs' (bottom; robustness of 0) weightings. The hollow circles and filled blue circles correspond to astrometric errors before and after applying the per-observation correction. Similarly, the dashed and solid lines correspond to the best fit uncorrected and corrected astrometric error, respectively. The horizontal dotted line shows the $\sim 10\%$ FWHM; our results far exceed that systematic for both corrected and uncorrected fits.

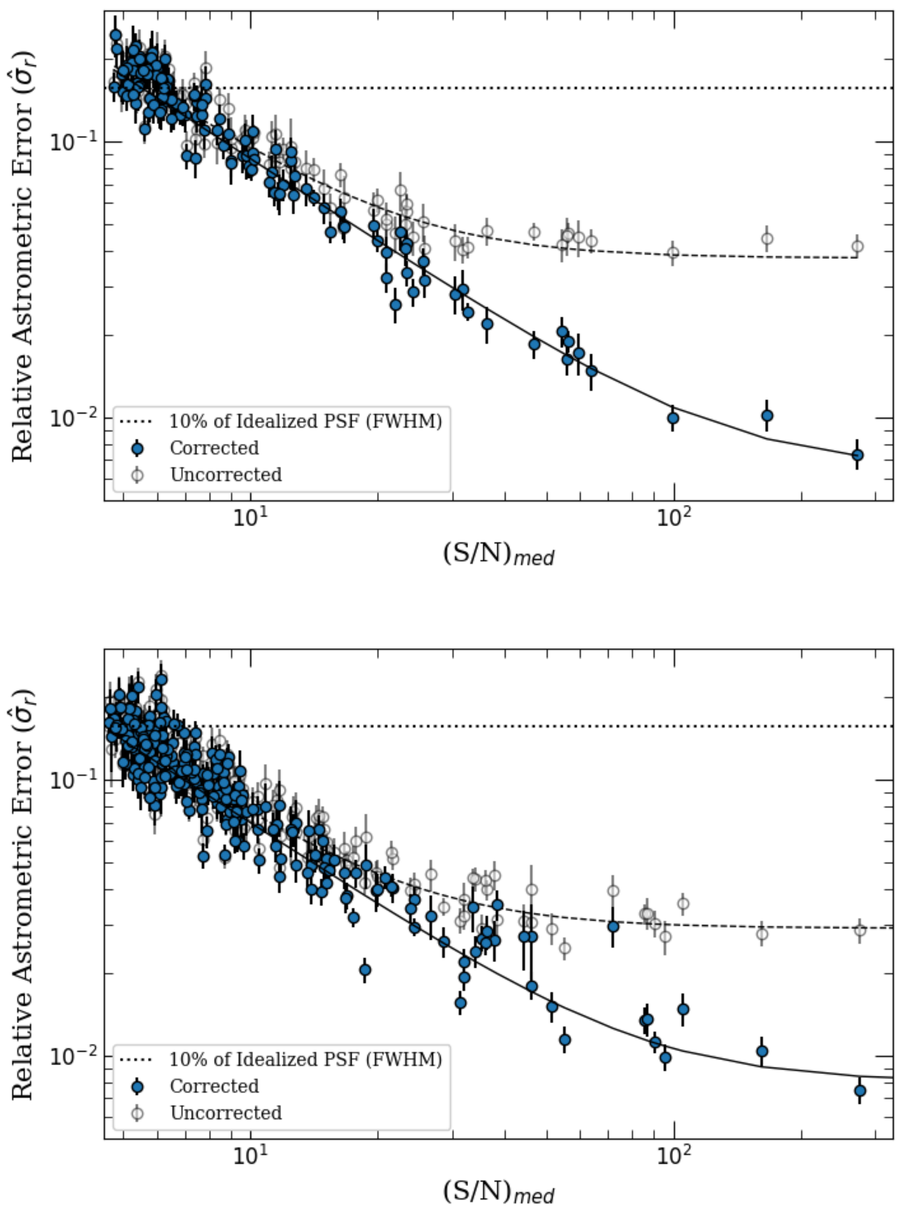

Figure 6.4: Same as Fig. 6.3 for the 1A 1744−361 field.

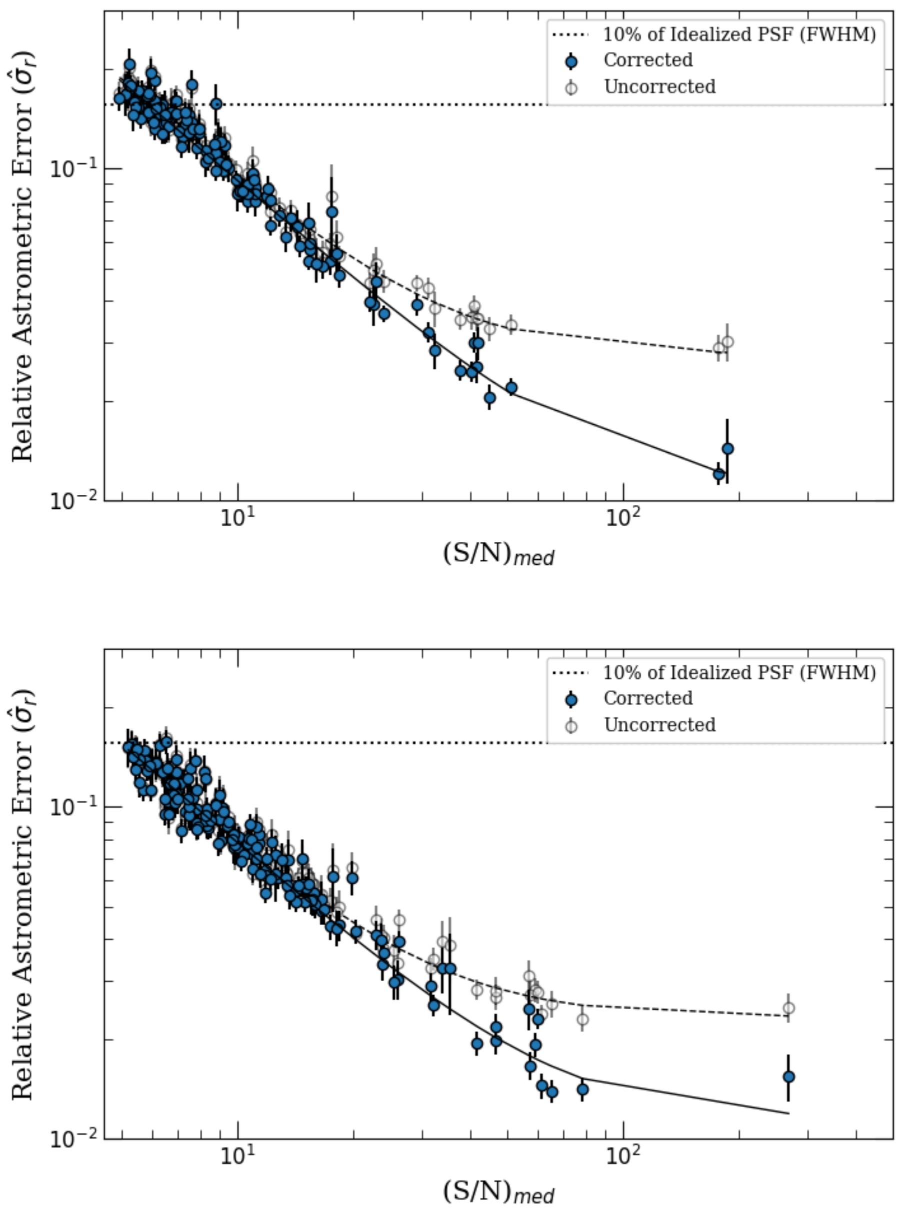

Figure 6.5: Same as Fig. 6.3 for the Swift J1727.8−163 field.

Table 6.1: Astrometric Error Parameters

Fit Type[a]	Target	Weighting	A	B	FWHM eq. (%)[b]	χ^2_{red}(dof)[c]
Uncorrected	SAX J1810.8-2609	Briggs 0	$0.707^{+0.006}_{-0.007}$	$0.0261^{+0.0007}_{-0.0006}$	$1.67^{+0.04}_{-0.04}$	1.54(206)
		Uniform	$0.884^{+0.012}_{-0.012}$	$0.0326^{+0.0012}_{-0.0013}$	$2.01^{+0.08}_{-0.08}$	1.48(104)
	1A 1744−361	Briggs 0	$0.710^{+0.007}_{-0.007}$	$0.0289^{+0.0007}_{-0.0007}$	$1.81^{+0.08}_{-0.08}$	2.13(242)
		Uniform	$0.879^{+0.012}_{-0.012}$	$0.0379^{+0.0013}_{-0.0013}$	$2.42^{+0.08}_{-0.08}$	1.52(127)
	Swift J1727.8−1613	Briggs 0	$0.766^{+0.007}_{-0.007}$	$0.0232^{+0.0007}_{-0.0007}$	$1.47^{+0.04}_{-0.04}$	1.38(158)
		Uniform[d]	$0.923^{+0.008}_{-0.009}$	$0.0272^{+0.0010}_{-0.0010}$	$1.73^{+0.06}_{-0.06}$	1.13(122)
Corrected	SAX J1810.8-2609	Briggs 0	$0.705^{+0.06}_{-0.006}$	$0.0066^{+0.0004}_{-0.0004}$	$0.42^{+0.03}_{-0.03}$	1.96(206)
		Uniform	$0.866^{+0.010}_{-0.010}$	$0.0064^{+0.0007}_{-0.0008}$	$0.41^{+0.05}_{-0.05}$	1.99(102)
	1A 1744−361	Briggs 0	$0.695^{+0.005}_{-0.005}$	$0.0080^{+0.0005}_{-0.0005}$	$0.50^{+0.05}_{-0.05}$	2.51(242)
		Uniform	$0.867^{+0.010}_{-0.010}$	$0.0065^{+0.0008}_{-0.0008}$	$0.42^{+0.05}_{-0.05}$	1.57(127)
	Swift J1727.8−1613	Briggs 0	$0.771^{+0.006}_{-0.006}$	$0.0115^{+0.0007}_{-0.0007}$	$0.64^{+0.04}_{-0.04}$	1.62(158)
		Uniform	$0.913^{+0.007}_{-0.007}$	$0.0114^{+0.0010}_{-0.0010}$	$0.64^{+0.06}_{-0.06}$	1.31(122)

Notes.
[a] This column indicates whether the fitting omitted (uncorrected) or used (corrected) the per-observation correction for RA and Dec;
[b] This column contains a fractional FWHM equivalent;
[c] This column contains the reduced χ^2_{red} and degrees of freedom (dof) for each fit.
[d] Marks the only formally acceptable fit.

6.4 Discussion and Future Improvements

The properties of the field sources are consistent with Equation (6.8); at low(er) signal-to-noise ratios, we observe a decrease in the astrometric error with an increasing signal-to-noise; at high(er) signal-to-noise ratios, we observe a flattened of the relation due to the systematic astrometric limit. Both the signal-to-noise dependent (A) and systematic (B) terms varied with the different XRB fields and the choice of the visibility weighting, adopting values of $A \in (0.70, 0.93)$ and $B \in (0.022, 0.039)$. Applying a per-epoch correction lowered the systematic limit to $B \in (0.0057, 0.012)$ but, as expected, had a negligible effect on A.

While the results demonstrate qualitative agreement, their $\chi^2_{\rm red}$ values are formally unacceptable given null-hypothesis probabilities of $p \lesssim 0.001$; the uniform-weighted Swift J1727.8−163 images are the exception as $\chi^2_{\rm red} \sim 1.13$ for 122 degrees of freedom, which corresponds to $p \sim 0.15$. The current iteration of the routine makes several simplifying assumptions that underestimate our measurement errors, thereby increasing the χ^2 (and thus $\chi^2_{\rm red}$) values. We neglect the error on the average position when calculating radial offsets, an approximation that only holds for many-epoch monitoring. For example, 1A 1744−361 was only monitored for 14 epochs; to first-order, the uncertainty on the average position is $\sim 30\%$ (i.e., $1/\sqrt{14}$) the uncertainty on position measured in a single epoch and should not be neglected in future versions of this analysis.

Furthermore, we neglect variance in the signal-to-noise of each source, instead adopting the median. Our variability exclusion criterion — sources need $< 200\%$ fractional variability — results in significant, non-zero variances in S/N. Stochastic RFI increases (or decreases) the required flagging, causing variations in rms noise of $\sim 30\%$, driving the signal-to-noise variance in otherwise persistent sources. Fitting with a sub-set of images with similar rms noise can minimize the signal-to-noise variance at the cost of (significantly) reducing the number of observations. In addition to RFI-driven noise changes, secular evolution(s) of the calibrator fields, which are assumed to be stable in time, is known to artificially induce $\sim 10\%$ intensity variability. I am in the process of improving this modelling to: (i) consider the two-dimensional uncertainties in both astrometric error and signal-to-noise ratio (using an approach similar to, e.g., Hogg et al., 2010); (ii) incorporate the error in the average position in our calculations (this can be as simple as increasing the astrometric uncertainty by a factor of $\sqrt{1 + n_{\rm obs}^{-1}}$); and (iii) solve for a per-observation intensity-correction to remove the calibrator induced variability.

I speculate that the worsening χ^2 values in the corrected fit result from residual direction-dependent effects. A single per-observation correction applied to all sources in the field may be insufficient. While I apply a off-axis source cutoff ($> 0.3°$, $\mathcal{A} < 75\%$) to minimize direction-dependant errors, DDEs can persist at separations corresponding to $\mathcal{A} \sim 90\%$ (Smirnov, 2011a; de Villiers, 2023). Applying a more restrictive off-axis condition (e.g., $\mathcal{A} \gtrsim 95\%$) would limit the catalogue to a prohibitively small number of sources. Given that MeerKAT has an elliptical primary beam response, I am investigating whether a more appropriate per-observation correction would factor in the source separation and orientation (e.g., North vs. East) with respect to the phase centre. I leave these improvements for future work.

Despite some deficiencies in the current workflow, my results robustly demonstrate that in these fields, the '10% of the FWHM" astrometric systematic is an over-correction. An extremely conservative assumption for these targets would be 3% systematic value. All comparable MeerKAT observations are likely capable of similarly precise astrometry. The ThunderKAT/X-KAT programs have multi-epoch monitoring of ~ 50 unique XRBs (and growing), which can be used to build an expansive sample for short-track observations, verifying the proposed astrometric errors. Furthermore, ThunderKAT/X-KAT has continuously monitored the BHXB GX 339−4, on a weekly-cadence for ~ 6.5 years (corresponding to ~ 300 epochs of monitoring). The large number of epochs makes this field the natural choice to select sub-sets of images at similar noise values and explore the relationship between the astrometric error and the choice of imaging/observation parameters, such as weightings, tapers, scan length, and pixel size. Finally, these results need to be tested with longer track observations. Such observations may have added systematic errors but also may have the potential to overcome these errors by breaking long tracks into 'multiple epochs' of short tracks.

I end this chapter with illustrative examples emphasizing the need for precise astrometry; in Figure 6.6, I show the core position of Swift J1727 applying the errors derived from this method (top panel) and with the standard $\sim 10\%$ systematic (bottom panel). Calculating a χ^2 assuming a model equal to the weighted average, I get values of ∼139 and ∼8 (with 42 degrees of freedom) from the two methods. Adopting the large systematic washes out any intrinsic variability in the core position, negating the most robust empirical evidence for the radio core being comprised of unresolved ejecta during the soft state. Moreover, the exceptionally low χ^2 of the 10% assumption strongly suggests that such errors are too large.

Furthermore, there has been a historical tendency to dismiss astrometric deviations from simple models, such as ballistic motion or constant deceleration of ejecta, by invoking arbitrary 'systematic effects'. Understanding these systematics has been a point of serious discussion (and, in some cases, disagreement) when interpreting our results in ThunderKAT and X-KAT. Recent relativistic hydrodynamic simulations have shown that the shocks produced in the ISM during ejecta propagation can produce quasi-stochastic *jitter* in the measured positions on the angular scales probable with MeerKAT (Savard et al., in prep). Therefore, dismissing the observed

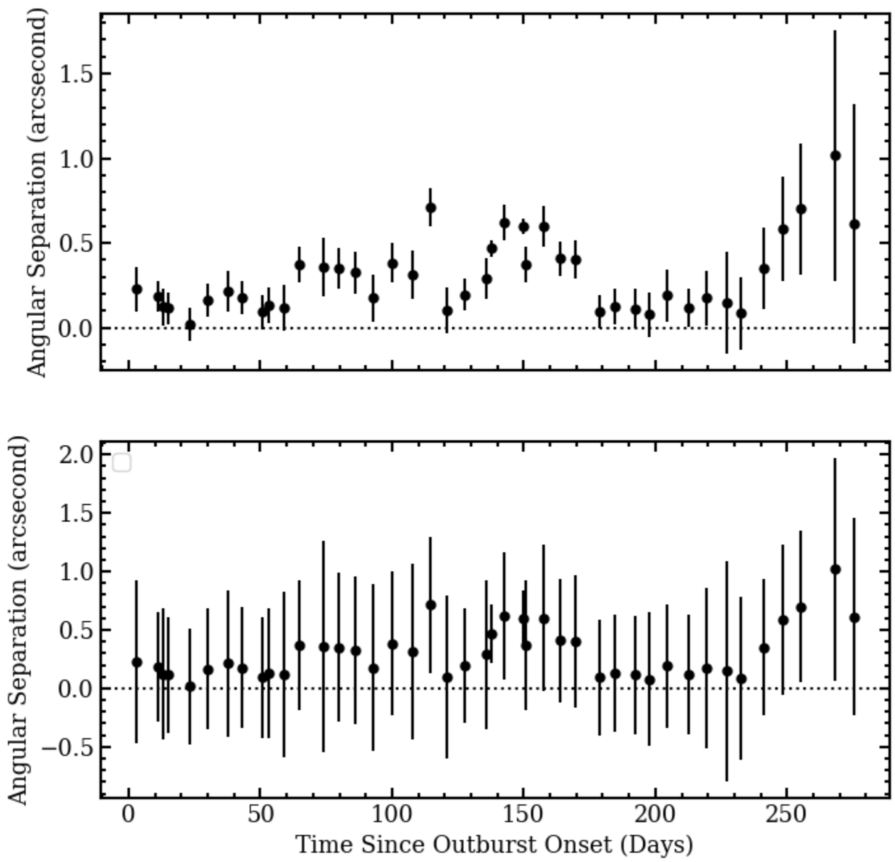

Figure 6.6: Angular offset of radio core of Swift J1727.8−1613: (top panel) The error derived using our empirical method; (bottom panel) astrometric errors with a $\sim$ 10% FWHM systematic.

motion as a systematic effect disregards potential physical phenomena that can occur during jet-ISM interactions. I argue that understanding these interactions is, instead, a primary purpose of studying astrophysical jets.

Chapter 7

Conclusion and Future Work

In this thesis, I presented the monitoring campaigns of three outbursting XRBs, two NSXB outbursts, and one (ongoing) BHXB outburst, monitored as part of the programmes ThunderKAT and, subsequently, X-KAT. Both programmes use the radio interferometer MeerKAT as their primary instrument. MeerKAT is a pathfinder for the upcoming Square Kilometer Array (SKA), and it will eventually be subsumed as the core of SKA-Mid. Once complete, the SKA will be one of the premiere radio observatories for the coming decades, combining milliarcsecond-scale angular resolutions with $\lesssim 5\,\mu\mathrm{Jy\,beam}^{-1}$ sensitivities from $\lesssim 10$ minute observations. While each component of my research had a scientific goal, an equally important motivation behind this research was developing new techniques for working with MeerKAT — and thus SKA — observations.

7.1 Conclusions

Scientifically, I focused on studying the time-domain evolution of astrophysical jets from XRBs, with the goal of deepening our understanding of the properties of jets and their relation to the simultaneous evolution of the coupled accretion (in)flow (Fender et al., 2009). Initially, the motivation for this work was to constrain better and, thus, understand the radio–X-ray luminosity correlation for NSXBs (Merloni et al., 2003; Gallo et al., 2003, 2014, 2018, i.e., the L_R–L_X plane) as the L_R–L_X correlations are the best-empirical evidence for the coupling between accretion flows and astrophysical jets; and comprehensively monitored NSXB are significantly underrepresented in our current sample (when compared to the rich monitoring of BHXBs). However, as is often the case, serendipitous results drove the research in slightly different directions.

My first NSXB and introduction to working with MeerKAT observations involved the analysis of the 2021 outburst of SAX J1810.8−2609. I found that the source exhibited a hard state only or 'failed' outburst. In hard-state-only outbursts, the system never transitions from the hard to the soft accretion state (e.g., Tarana et al., 2018; Alabarta et al., 2022). As a result, hard-state-only outbursts are ideal for constraining the L_R–L_X plane, as the compact jet is never disrupted during the hard-to-soft-state transition. During the early stages of the outburst, the source behaved as expected, showing positively correlated radio and X-ray emission. However, as the X-ray luminosity decayed by over four orders of magnitude, a spatially coincident, persistent radio source at a flux density of $\sim 100\,\mu$Jy, became apparent. While I observed some stochastic proper motion, it was inconsistent with ballistic motion, ruling out a long-lasting transient jet ejecta as the origin of the persistent radio emission. Such an interpretation is commonly invoked for decoupled X-ray and radio emission in XRBs. I investigated the possibility that SAX J1810.8−2609 is a transitional millisecond pulsar, as these sources are known to exhibit exotic radio–X-ray correlations (e.g., the anti-correlation of PSR J1023+0038; Bogdanov et al., 2018). However, SAX J1810.8−2609 has never exhibited radio pulsations nor rapid X-ray variability (i.e., accretion mode switching), and its quiescent X-ray spectrum is thermal. In contrast, tMSPs have non-thermal spectra at low X-ray luminosities. I calculated a chance coincidence probability, and while low ($\sim 0.16\%$), a background AGN seems the most likely explanation. In future outbursts, confirming the existence of a second radio component would be trivial with higher-resolution instruments like the VLA and SKA. Moreover, if the persistent radio emission is confirmed to be local, this discovery would require a novel emission mechanism (in the context of NSXBs).

The unusual proper motion observed from SAX J1810.8−2609 initiated an investigation into the astrometric precision of MeerKAT, as I needed to understand the likelihood that the observed variability was stochastic. During the investigation, I began the development of a simple yet novel empirical technique to determine the astrometric errors from time-domain observations with wide-field interferometers. My results robustly demonstrated that the astrometric variability of SAX J1810.8−2609 was significantly elevated when compared to the stochastic evolution of the background sources and, perhaps more importantly, provided evidence that the standard assumptions for astrometric limitations (i.e., errors greater than 10% the size of the PSF) is a severe overestimation — at least for MeerKAT.

Following the monitoring of SAX J1810.8−2609, I adopted a new role as one (of three, at the time) collaboration members who scheduled each observation. Shortly after becoming a scheduler, I was presented with the opportunity to lead the monitoring of the 2022 outburst of the NSXB 1A 1744−361. I created, submitted, and was approved for a rapid response observation of the source within days of the outburst onset (and hours of its initial reporting). Like SAX J1810.8−2609, the initial science-driven goal was to acquire comprehensive monitoring of its L_R–L_X evolution. The rapid response observation revealed the source to be at a radio flux density of $\sim 1\,\mathrm{mJy}$ (Hughes et al., 2022), with similarly bright X-ray emission. At these flux densities, 1A 1744−361 would be one of the (radio) brightest hard state NSXBs if hard state radio emission could be confirmed. Subsequent observations revealed a rapid decay of the radio emission and a transition to a steep, optically thin synchrotron spectrum, leading me to reinterpret the radio emission as originating from transient jet ejecta. After this realization, I collaborated with an external team that analysed the X-ray spectral and timing properties (using NICER), revealing that the X-ray properties supported the reinterpretation of the radio emission. Moreover, the X-ray timing properties identified 1A 1744−361 as the newest member of the class of NXSBs with X-ray properties that transition between Z- and atoll-subpopulations (Ng et al., 2024). Identifying these transitioning sources conclusively shows that atoll and Z-sources are not distinct populations. Instead, these 'subpopulations' are different accretion states at different mass accretion rates and, therefore, X-ray luminosities. This picture of evolution between accretion states makes NSXBs more consistent with the picture of BHXBs. While 1A 1744−361 never showed hard-state radio emission, thereby prohibiting our desired L_R–L_X monitoring, after parsing the literature, I realized that the behaviour of 1A 1744−361 was the best evidence for a transient jet ejection from an NSXB during the canonical hard-to-soft transition. This behaviour had been long-observed in BHXB, but until now, only predicted for NSXBs (Migliari & Fender, 2006; Muñoz-Darias et al., 2014).

The difficulties I faced during our NSXB monitoring as a part of ThunderKAT, particularly with 1A 1744−361, motivated the inclusion of a dedicated, higher cadence monitoring of NSXBs for the ThunderKAT successor, X-KAT. However, transitioning to X-KAT, the collaboration identified accurate polarimetry as its highest priority. As a result, I switched my focus and led the development of our polarisation calibration pipeline, POLKAT. Coinciding with the start of X-KAT, a new BHXB went into outburst in late 2023, Swift J1727.8−1613, quickly becoming one of the brightest outbursts in the last decade, and thus, becoming the ideal candidate for full polarisation monitoring. This source was sufficiently bright to measure $< 1\%$ polarisation fractions. In this thesis, I included our MeerKAT-centred results and discussed some of the exotic behaviours of this BHXB. Swift J1727.8$-$1613 exhibited clear signs of decelerated jet ejecta (i.e., due to jet-ISM interactions) and an absence of counter-jet ejecta (i.e., likely due to highly beamed emission), occupied an extraordinarily radio-quiet region of the L_R–L_X plane, and showed dramatic evolution of the polarization properties.

Regarding polarisation, which was my now primary interest, Swift J1727.8$-$1613 showed, for the first time, strong evidence of a 'rotation measure flare' from an XRB. Moreover, this flare coincided with a softening of the X-ray spectrum and a bright radio flare. After investigating, I found it plausible that the transient increase in rotation measure originated from the dense-magnetic plasma in a transient jet ejectum. This proposed origin carries significant insights into the composition of jet ejecta, a pressing question in jet physics, favouring electron-proton (over electron-positron) plasma. Relating this result to feedback, electron-proton jets have significantly more kinetic energy (and thus energy for feedback) than electron-positron plasmas (at the same bulk velocity) due to the larger masses of comprising protons; I intend to investigate whether the signatures of electron-proton plasma are ubiquitous for jet ejecta in X-ray binaries by looking for similar signatures in other X-KAT sources. In this work, I only presented the MeerKAT (radio) observations, omitting the multi-facility monitoring taken by other collaboration members. Indeed, the totality of our data will improve the breadth of our frequency coverage and the density of our temporal coverage by a factor $\geq 3\times$ when the full results are published. This will make Swift J1727.8$-$1613 and its 2023/2024 outburst one of the most densely monitored and thus best-understood XRBs.

The polarisation calibration pipeline POLKAT is now being applied to multiple X-KAT monitored XRBs as the collaboration transitions to extracting, at minimum, polarisation limits for all of its targets. Once completed, X-KAT will increase the number of XRBs with full polarisation monitoring by >100%, providing invaluable insights into the evolving magnetic fields and internal compositions of jets. More generally, POLKAT can be readily deployed for any (linear-feed) interferometric observation regardless of the science target. For instance, I am already affiliated with calibrations beyond X-KAT that will use the X-KAT pipeline when reducing their polarimetric observations.

Finally, motivated by proper motions observed from SAX J1810.8−2609 and Swift J1727.8−1613, as well as intra-collaboration discussions, I revisited the topic of astrometric precision, including a more comprehensive treatment of the statistics. My investigation led to the development of a source-agnostic astrometry software, ASTKAT, that empirically determines the astrometric precision using the background properties in the field for time-domain observations. While I have only applied the routine to my MeerKAT observations, the techniques can be readily applied to any field with sufficient sensitivity to far-field sources and, as a result, the ability to monitor tens to hundreds of sources in single pointings. The primary goal of this work was to quantify the systematic astrometric limits (of MeerKAT) that result from imperfect calibration and deconvolution. Moreover, I compared my results to the common *ad hoc* prescription of $\sim 10\%$ the PSF FWHM. As expected, applying the routine to all three fields supported what was found from SAX J1810.8−2609; MeerKAT achieves systematics of $\leq 2.5\%$ the PSF FWHM and can reach as low as $\sim 0.5\%$, if per-observation corrections are applied to the field sources. This floor to the astrometric precision is a factor of 4–20 improvement over the *ad hoc* systematic. Given that modern radio arrays can readily reach $10\,\mu\mathrm{Jy\,beam}^{-1}$ rms, point sources brighter than $\sim 500\,\mu$Jy can achieve this floor (corresponding to ~ 60 mas for MeerKAT). Kinematic modelling of ejecta motion is widely accepted to be the best estimator of the total energy content of jets. Therefore, higher precision astrometry either more precisely constrains parameters in current models or reveals additional physics (i.e., more complex shock-ISM interactions) that is being neglected in the current framework(s).

I end the chapter on astrometric accuracy by discussing some of the current shortcomings in our routine and plans to mitigate these issues. However, I emphasize that, despite the known flaws and the small sample, my results suggest the *ad hoc* systematic is a severe overcorrection; future monitoring should adopt these more precise errors to avoid missing physics-driven astrometric changes.

7.2 Future Outlooks: Utilizing the Full Capabilities of MeerKAT and SKA-Mid

As the community continues to push the bounds of astrophysical research, we will naturally require more powerful and, thus, more expensive facilities. Our ideas already require significant resources from multiple nations. Therefore, these newer facilities will become progressively more oversubscribed, resulting in fewer observations per researcher and, thus, the need to extract **all** of possible science from each observation. Throughout this thesis, I have developed and applied new techniques that maximize the astrometric and polarization information from MeerKAT observations. The former requires no changes to observing plans, while the latter only requires that observers include a polarization angle calibrator. Many observers already do so.

In the immediate future, my primary objective is to improve the systematics of POLKAT and ASTKAT, and thus, the science we can achieve from each observation as a part of X-KAT. To put this in context, for Swift J1727.8−1613, there are multiple observations where we are not signal-to-noise limited but, instead, limited by the systematics of our calibration. Most significantly, we lose rotation measure and polarization angle information immediately following the 'rotation-measure' flare as the measured polarization fraction is only $\sim 3.5\times$ the assumed systematic value (failing the $>7\sigma$ RM synthesis threshold). I am focused on improving our calibration routine (e.g., through the improvement of calibrator models with self-calibration) to reach the systematic polarization levels set by the intrinsic uncertainties of the calibrator source properties ($\sim 0.01\%$). Reaching these systematics would enable the full utilization of MeerKAT's polarimetric capabilities.

Bibliography

Alabarta, K., Altamirano, D., Mendez, M., et al. 2022, in 44th COSPAR Scientific Assembly. Held 16-24 July, Vol. 44, 1759

Allen, J. L., Homan, J., Chakrabarty, D., & Nowak, M. 2018, ApJ, 854, 58, doi: 10.3847/1538-4357/aaa6d3

Altamirano, D., Homan, J., Linares, M., et al. 2010, The Astronomer's Telegram, 2952, 1

Arnaud, K. A. 1996a, in Astronomical Society of the Pacific Conference Series, Vol. 101, Astronomical Data Analysis Software and Systems V, ed. G. H. Jacoby & J. Barnes, 17

Arnaud, K. A. 1996b, in Astronomical Society of the Pacific Conference Series, Vol. 101, Astronomical Data Analysis Software and Systems V, ed. G. H. Jacoby & J. Barnes, 17

Arras, P., Reinecke, M., Westermann, R., & Enßlin, T. A. 2021, A&A, 646, A58, doi: 10.1051/0004-6361/202039723

Bahramian, A., Heinke, C. O., & Sivakoff, G. R. 2013, The Astronomer's Telegram, 5301, 1

Bahramian, A., & Rushton, A. 2022, bersavosh/XRB-LrLx_pub: update 20220908, v220908, Zenodo, doi: 10.5281/zenodo.7059313

Bahramian, A., Tremou, E., Tetarenko, A. J., et al. 2023, ApJL, 948, L7, doi: 10.3847/2041-8213/accde1

Bałucińska-Church, M., Schulz, N. S., Wilms, J., et al. 2011, A&A, 530, A102, doi: 10.1051/0004-6361/201015931

Barthelmy, S. D., Barbier, L. M., Cummings, J. R., et al. 2005, SSRv, 120, 143, doi: 10.1007/s11214-005-5096-3

Belloni, T., Méndez, M., van der Klis, M., Lewin, W. H. G., & Dieters, S. 1999, ApJL, 519, L159, doi: 10.1086/312130

Belloni, T., Psaltis, D., & van der Klis, M. 2002, ApJ, 572, 392, doi: 10.1086/340290

Belloni, T. M. 2010, in Lecture Notes in Physics, ed. T. Belloni, Vol. 794 (Berlin: Springer Verlag), 53, doi: 10.1007/978-3-540-76937-83

Belloni, T. M., & Motta, S. E. 2016, in Astrophysics and Space Science Library, Vol. 440, Astrophysics of Black Holes: From Fundamental Aspects to Latest Developments, ed. C. Bambi, 61, doi: 10.1007/978-3-662-52859-4_2

Bhattacharyya, S., Strohmayer, T. E., Markwardt, C. B., & Swank, J. H. 2006, ApJL, 639, L31, doi: 10.1086/501438

Bilous, A. V., Watts, A. L., Galloway, D. K., & in 't Zand, J. J. M. 2018, ApJL, 862, L4, doi: 10.3847/2041-8213/aad09c

Blandford, R. D., & Königl, A. 1979, ApJ, 232, 34, doi: 10.1086/157262

Blandford, R. D., & Payne, D. G. 1982, MNRAS, 199, 883, doi: 10.1093/mnras/199.4.883

Blandford, R. D., & Znajek, R. L. 1977, MNRAS, 179, 433, doi: 10.1093/mnras/179.3.433

Bogdanov, S., Deller, A. T., Miller-Jones, J. C. A., et al. 2018, ApJ, 856, 54, doi: 10.3847/1538-4357/aaaeb9

Bollemeijer, N., Uttley, P., Buisson, D., et al. 2023a, The Astronomer's Telegram, 16247, 1

—. 2023b, The Astronomer's Telegram, 16273, 1

Braun, R., Bonaldi, A., Bourke, T., Keane, E., & Wagg, J. 2019, arXiv e-prints, arXiv:1912.12699, doi: 10.48550/arXiv.1912.12699

Brentjens, M. A., & de Bruyn, A. G. 2005, A&A, 441, 1217, doi: 10.1051/0004-6361:20052990

Briggs, D. 1995a, Ph. D. thesis, New Mexico Institute of Mining and Technology. Avaible via http://www.aoc.nrao.edu/dissertations/dbriggs/

Briggs, D. S. 1995b, in American Astronomical Society Meeting Abstracts, Vol. 187, American Astronomical Society Meeting Abstracts, 112.02

Bright, J. S., et al. 2020, Nature Astron., 4, 697, doi: 10.1038/s41550-020-1023-5

Brocksopp, C., Corbel, S., Tzioumis, A., et al. 2013, MNRAS, 432, 931, doi: 10.1093/mnras/stt493

Brocksopp, C., Miller-Jones, J. C. A., Fender, R. P., & Stappers, B. W. 2007, MNRAS, 378, 1111, doi: 10.1111/j.1365-2966.2007.11846.x

Burn, B. J. 1966, MNRAS, 133, 67, doi: 10.1093/mnras/133.1.67

Burrows, D. N., Hill, J. E., Nousek, J. A., et al. 2005, SSRv, 120, 165, doi: 10.1007/s11214-005-5097-2

Camilo, F. 2018, Nature Astronomy, 2, 594, doi: 10.1038/s41550-018-0516-y

Cao, X.-F., Wu, Q., & Dong, A.-J. 2014, ApJ, 788, 52, doi: 10.1088/0004-637X/788/1/52

Carotenuto, F., Fender, R., Tetarenko, A. J., et al. 2024, arXiv e-prints, arXiv:2405.16624, doi: 10.48550/arXiv.2405.16624

Carotenuto, F., Tetarenko, A. J., & Corbel, S. 2022, MNRAS, 511, 4826, doi: 10.1093/mnras/stac329

Carotenuto, F., Corbel, S., Tremou, E., et al. 2021, MNRAS, 504, 444, doi: 10.1093/mnras/stab864

Carotenuto, F., Corbel, S., Tremou, E., et al. 2021, Monthly Notices of the Royal Astronomical Society: Letters, 505, L58, doi: 10.1093/mnrasl/slab049

Carpenter, G. F., Eyles, C. J., Skinner, G. K., Wilson, A. M., & Willmore, A. P. 1977, MNRAS, 179, 27P, doi: 10.1093/mnras/179.1.27P

CASA Team, Bean, B., Bhatnagar, S., et al. 2022, PASP, 134, 114501, doi: 10.1088/1538-3873/ac9642

Cash, W. 1979, ApJ, 228, 939, doi: 10.1086/156922

Castro-Tirado, A. J., Sanchez-Ramirez, R., Caballero-Garcia, M. D., et al. 2023, The Astronomer's Telegram, 16208, 1

Chakraborty, M., Bhattacharyya, S., & Mukherjee, A. 2011, MNRAS, 418, 490, doi: 10.1111/j.1365-2966.2011.19499.x

Church, M. J., Parmar, A. N., Balucinska-Church, M., et al. 1998, A&A, 338, 556, doi: 10.48550/arXiv.astro-ph/9806223

Church, M. J., Reed, D., Dotani, T., Bałucińska-Church, M., & Smale, A. P. 2005, MNRAS, 359, 1336, doi: 10.1111/j.1365-2966.2005.08728.x

Condon, J. J. 1997, PASP, 109, 166, doi: 10.1086/133871

Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1998, AJ, 115, 1693, doi: 10.1086/300337

Contopoulos, I., Kylafis, N. D., Kazanas, D., & Christodoulou, D. M. 2012, in International Journal of Modern Physics Conference Series, Vol. 8, International Journal of Modern Physics Conference Series, 287–292, doi: 10.1142/S2010194512004722

Corbel, S., Coriat, M., Brocksopp, C., et al. 2013, MNRAS, 428, 2500, doi: 10.1093/mnras/sts215

Corbel, S., & Fender, R. P. 2002, ApJL, 573, L35, doi: 10.1086/341870

Corbel, S., Fender, R. P., Tzioumis, A. K., et al. 2000, A&A, 359, 251, doi: 10.48550/arXiv.astro-ph/0003460

Corbel, S., Kaaret, P., Fender, R. P., et al. 2005, ApJ, 632, 504, doi: 10.1086/432499

Coriat, M., Corbel, S., Prat, L., et al. 2011, MNRAS, 414, 677, doi: 10.1111/j.1365-2966.2011.18433.x

Cornwell, T. J., Golap, K., & Bhatnagar, S. 2008, IEEE Journal of Selected Topics in Signal Processing, 2, 647, doi: 10.1109/JSTSP.2008.2005290

Corral-Santana, J. M., Casares, J., Muñoz-Darias, T., et al. 2016, A&A, 587, A61, doi: 10.1051/0004-6361/201527130

Curran, P. A., Coriat, M., Miller-Jones, J. C. A., et al. 2014, MNRAS, 437, 3265, doi: 10.1093/mnras/stt2125

Curran, P. A., Miller-Jones, J. C. A., Rushton, A. P., et al. 2015, MNRAS, 451, 3975, doi: 10.1093/mnras/stv1252

Czerny, B., Siemiginowska, A., Janiuk, A., Nikiel-Wroczyński, B., & Stawarz, Ł. 2009, ApJ, 698, 840, doi: 10.1088/0004-637X/698/1/840

de Gasperin, F., Intema, H. T., & Frail, D. A. 2018, MNRAS, 474, 5008, doi: 10.1093/mnras/stx3125

de Villiers, M. S. 2023, AJ, 165, 78, doi: 10.3847/1538-3881/acabc3

Degenaar, N., Klein-Wolt, M., & Wijnands, R. 2007, The Astronomer's Telegram, 1175, 1

Degenaar, N., & Wijnands, R. 2013, in Neutron Stars and Pulsars: Challenges and Opportunities after 80 years, ed. J. van Leeuwen, Vol. 291, 141–144, doi: 10.1017/S174392131202340X

Díaz Trigo, M., Parmar, A. N., Boirin, L., Méndez, M., & Kaastra, J. S. 2006, A&A, 445, 179, doi: 10.1051/0004-6361:20053586

Díaz Trigo, M., Altamirano, D., Dinçer, T., et al. 2018, A&A, 616, A23, doi: 10.1051/0004-6361/201832693

Dickey, J. M., & Lockman, F. J. 1990, ARA&A, 28, 215, doi: 10.1146/annurev.aa.28.090190.001243

Done, C., Gierliński, M., & Kubota, A. 2007, A&A Rv, 15, 1, doi: 10.1007/s00159-007-0006-1

Draghis, P. A., Miller, J. M., Homan, J., et al. 2023, The Astronomer's Telegram, 16219, 1

Espinasse, M., & Fender, R. 2018, MNRAS, 473, 4122, doi: 10.1093/mnras/stx2467

Espinasse, M., Corbel, S., Kaaret, P., et al. 2020, ApJL, 895, L31, doi: 10.3847/2041-8213/ab88b6

Evans, P. A., Beardmore, A. P., Page, K. L., et al. 2007, A&A, 469, 379, doi: 10.1051/0004-6361:20077530

—. 2009, MNRAS, 397, 1177, doi: 10.1111/j.1365-2966.2009.14913.x

Fender, R. 2010, in Lecture Notes in Physics, ed. T. Belloni, Vol. 794 (Berlin: Springer Verlag), 115, doi: 10.1007/978-3-540-76937-85

Fender, R., & Bright, J. 2019, MNRAS, 489, 4836, doi: 10.1093/mnras/stz2000

Fender, R., Woudt, P. A., Corbel, S., et al. 2016, in MeerKAT Science: On the Pathway to the SKA, 13, doi: 10.22323/1.277.0013

Fender, R. P., Belloni, T. M., & Gallo, E. 2004, MNRAS, 355, 1105, doi: 10.1111/j.1365-2966.2004.08384.x

Fender, R. P., Garrington, S. T., McKay, D. J., et al. 1999a, MNRAS, 304, 865, doi: 10.1046/j.1365-8711.1999.02364.x

Fender, R. P., Hanson, M. M., & Pooley, G. G. 1999b, MNRAS, 308, 473, doi: 10.1046/j.1365-8711.1999.02726.x

Fender, R. P., Homan, J., & Belloni, T. M. 2009, MNRAS, 396, 1370, doi: 10.1111/j.1365-2966.2009.14841.x

Fomalont, E. B., Geldzahler, B. J., & Bradshaw, C. F. 2001, ApJ, 558, 283, doi: 10.1086/322479

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306, doi: 10.1086/670067

Fridriksson, J. K. 2011, PhD thesis, MIT, Center for Space Research/Kavli Institute

Frieden, B. R. 1972, Journal of the Optical Society of America (1917-1983), 62, 511

Gallo, E., Corbel, S., Fender, R. P., Maccarone, T. J., & Tzioumis, A. K. 2004, MNRAS, 347, L52, doi: 10.1111/j.1365-2966.2004.07435.x

Gallo, E., Degenaar, N., & van den Eijnden, J. 2018, MNRAS, 478, L132, doi: 10.1093/mnrasl/sly083

Gallo, E., Fender, R. P., & Pooley, G. G. 2003, MNRAS, 344, 60, doi: 10.1046/j.1365-8711.2003.06791.x

Gallo, E., Miller-Jones, J. C. A., Russell, D. M., et al. 2014, MNRAS, 445, 290, doi: 10.1093/mnras/stu1599

Galloway, D. K., & Keek, L. 2021, in Astrophysics and Space Science Library, Vol. 461, Timing Neutron Stars: Pulsations, Oscillations and Explosions, ed. T. M. Belloni, M. Méndez, & C. Zhang, 209–262, doi: 10.1007/978-3-662-62110-3_5

Galloway, D. K., in't Zand, J., Chenevez, J., et al. 2020, ApJS, 249, 32, doi: 10.3847/1538-4365/ab9f2e

Gehrels, N., Chincarini, G., Giommi, P., et al. 2004, ApJ, 611, 1005, doi: 10.1086/422091

Gendreau, K. C., Arzoumanian, Z., Adkins, P. W., et al. 2016, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9905, Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray, ed. J.-W. A. den Herder, T. Takahashi, & M. Bautz, 99051H, doi: 10.1117/12.2231304

George, S. J., Stil, J. M., & Keller, B. W. 2012, PASA, 29, 214, doi: 10.1071/AS11027

Gies, D. R., Bolton, C. T., Thomson, J. R., et al. 2003, ApJ, 583, 424, doi: 10.1086/345345

Ginzburg, V. L., & Syrovatskii, S. I. 1969, ARA&A, 7, 375, doi: 10.1146/annurev.aa.07.090169.002111

Gladstone, J., Done, C., & Gierliński, M. 2007, MNRAS, 378, 13, doi: 10.1111/j.1365-2966.2007.11675.x

Goodman, J., & Weare, J. 2010, Communications in Applied Mathematics and Computational Science, 5, 65, doi: 10.2140/camcos.2010.5.65

Gottlieb, O., Metzger, B. D., Quataert, E., et al. 2023, ApJL, 958, L33, doi: 10.3847/2041-8213/ad096e

Gull, S. F., & Daniell, G. J. 1978, Nature, 272, 686, doi: 10.1038/272686a0

Gusinskaia, N. V., Deller, A. T., Hessels, J. W. T., et al. 2017, MNRAS, 470, 1871, doi: 10.1093/mnras/stx1235

Hales, C. A., Gaensler, B. M., Norris, R. P., & Middelberg, E. 2012, MNRAS, 424, 2160, doi: 10.1111/j.1365-2966.2012.21372.x

Han, X., & Hjellming, R. M. 1992, ApJ, 400, 304, doi: 10.1086/171996

Hannikainen, D. C., Hunstead, R. W., Campbell-Wilson, D., & Sood, R. K. 1998, A&A, 337, 460, doi: 10.48550/arXiv.astro-ph/9805332

Hardcastle, M. J., & Croston, J. H. 2020, NewAR, 88, 101539, doi: 10.1016/j.newar.2020.101539

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

Heywood, I. 2020, oxkat: Semi-automated imaging of MeerKAT observations, Astrophysics Source Code Library, record ascl:2009.003. http://ascl.net/2009.003

Heywood, I., Jarvis, M. J., Baker, A. J., et al. 2016, MNRAS, 460, 4433, doi: 10.1093/mnras/stw1250

Heywood, I., Jarvis, M. J., Hale, C. L., et al. 2022, MNRAS, 509, 2150, doi: 10.1093/mnras/stab3021

Hjellming, R. M., & Johnston, K. J. 1988, ApJ, 328, 600, doi: 10.1086/166318

Hjellming, R. M., & Rupen, M. P. 1995, Nature, 375, 464, doi: 10.1038/375464a0

Hodge, M. A., Lister, M. L., Aller, M. F., et al. 2018, ApJ, 862, 151, doi: 10.3847/1538-4357/aacb2f

Högbom, J. A. 1974, A&AS, 15, 417

Hogg, D. W., Bovy, J., & Lang, D. 2010, arXiv e-prints, arXiv:1008.4686, doi: 10.48550/arXiv.1008.4686

Homan, J., & Belloni, T. 2005, Ap&SS, 300, 107, doi: 10.1007/s10509-005-1197-4

Homan, J., van der Klis, M., Wijnands, R., et al. 2007, ApJ, 656, 420, doi: 10.1086/510447

Hughes, A. K., Sivakoff, G. R., Fender, R., et al. 2022, The Astronomer's Telegram, 15410, 1

Hughes, A. K., Sivakoff, G. R., van den Eijnden, J., et al. 2024, MNRAS, 527, 9359, doi: 10.1093/mnras/stad3715

Hughes, A. K., Sivakoff, G. R., Macpherson, C. E., et al. 2023, MNRAS, 521, 185, doi: 10.1093/mnras/stad396

Hugo, B., & Perley, R. 2024, South African Radio Astronomy Observatory, doi: doi.org/10.48479/bqk7-aw53

Hugo, B. V., Perkins, S., Merry, B., Mauch, T., & Smirnov, O. M. 2022, in Astronomical Society of the Pacific Conference Series, Vol. 532, Astronomical Society of the Pacific Conference Series, ed. J. E. Ruiz, F. Pierfedereci, & P. Teuben, 541. https://arxiv.org/abs/2206.09179

Ingram, A., Bollemeijer, N., Veledina, A., et al. 2023, arXiv e-prints, arXiv:2311.05497, doi: 10.48550/arXiv.2311.05497

Ingram, A. R., & Motta, S. E. 2019, NewAR, 85, 101524, doi: 10.1016/j.newar.2020.101524

Iwakiri, W., Negoro, H., Nakajima, M., et al. 2021, The Astronomer's Telegram, 14649, 1

Jankowski, F., van Straten, W., Keane, E. F., et al. 2018, MNRAS, 473, 4436, doi: 10.1093/mnras/stx2476

Jonas, J., & MeerKAT Team. 2016, in MeerKAT Science: On the Pathway to the SKA, 1, doi: 10.22323/1.277.0001

Jonker, P. G., Wijnands, R., & van der Klis, M. 2004, MNRAS, 349, 94, doi: 10.1111/j.1365-2966.2004.07455.x

Kajava, J. J. E., Sánchez-Fernández, C., Alfonso-Garzón, J., Motta, S. E., & Veledina, A. 2020, A&A, 634, A94, doi: 10.1051/0004-6361/201937191

Kaper, H. G., Smits, D. W., Schwarz, U., Takakubo, K., & van Woerden, H. 1966, BAN, 18, 465

Keith, M. J., Jameson, A., van Straten, W., et al. 2010, MNRAS, 409, 619, doi: 10.1111/j.1365-2966.2010.17325.x

Kennea, J. A., & Swift Team. 2023, GRB Coordinates Network, 34540, 1

Kenyon, J. S., Smirnov, O. M., Grobler, T. L., & Perkins, S. J. 2018, MNRAS, 478, 2399, doi: 10.1093/mnras/sty1221

Kobayashi, K., Negoro, H., Nakajima, M., et al. 2022, The Astronomer's Telegram, 15407, 1

Kong, A. K. H., Charles, P. A., Homer, L., Kuulkers, E., & O'Donoghue, D. 2006, MNRAS, 368, 781, doi: 10.1111/j.1365-2966.2006.10157.x

Kuulkers, E., den Hartog, P. R., in't Zand, J. J. M., et al. 2003, A&A, 399, 663, doi: 10.1051/0004-6361:20021781

Laing, R. A. 1980, MNRAS, 193, 439, doi: 10.1093/mnras/193.3.439

Legg, M. P. C., & Westfold, K. C. 1968, ApJ, 154, 499, doi: 10.1086/149777

Lewin, W. H. G., van Paradijs, J., & van der Klis, M. 1988, SSRv, 46, 273, doi: 10.1007/BF00212242

Li, K.-L., Strader, J., Miller-Jones, J. C. A., Heinke, C. O., & Chomiuk, L. 2020, ApJ, 895, 89, doi: 10.3847/1538-4357/ab8f28

Lin, D., Remillard, R. A., & Homan, J. 2007, ApJ, 667, 1073, doi: 10.1086/521181

—. 2009, ApJ, 696, 1257, doi: 10.1088/0004-637X/696/2/1257

Linares, M. 2014, ApJ, 795, 72, doi: 10.1088/0004-637X/795/1/72

Liu, H.-X., Xu, Y.-J., Zhang, S.-N., et al. 2024, arXiv e-prints, arXiv:2406.03834, doi: 10.48550/arXiv.2406.03834

Livio, M. 2002, Nature, 417, 125, doi: 10.1038/417125a

Longair, M. S. 2011, High Energy Astrophysics (Cambridge University Press)

Lorimer, D. R., & Kramer, M. 2004, Handbook of Pulsar Astronomy, Vol. 4 (Cambridge University Press)

Maccarone, T. J. 2002, MNRAS, 336, 1371, doi: 10.1046/j.1365-8711.2002.05876.x

—. 2003, A&A, 409, 697, doi: 10.1051/0004-6361:20031146

Macquart, J. P., Ekers, R. D., Feain, I., & Johnston-Hollitt, M. 2012, ApJ, 750, 139, doi: 10.1088/0004-637X/750/2/139

Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, AJ, 129, 1993, doi: 10.1086/428488

Mata Sánchez, D., Muñoz-Darias, T., Armas Padilla, M., Casares, J., & Torres, M. A. P. 2024, A&A, 682, L1, doi: 10.1051/0004-6361/202348754

Matsuoka, M., Kawasaki, K., Ueno, S., et al. 2009, PASJ, 61, 999, doi: 10.1093/pasj/61.5.999

McClintock, J. E., & Remillard, R. A. 2006, Black hole binaries, Vol. 39 (Cambridge University Press), 157–213

Meier, D. L., Koide, S., & Uchida, Y. 2001, Science, 291, 84, doi: 10.1126/science.291.5501.84

Mereminskiy, I., Lutovinov, A., Molkov, S., et al. 2024, MNRAS, 531, 4893, doi: 10.1093/mnras/stae1393

Merloni, A., Heinz, S., & di Matteo, T. 2003, MNRAS, 345, 1057, doi: 10.1046/j.1365-2966.2003.07017.x

Meyer-Hofmeister, E., & Meyer, F. 2014, A&A, 562, A142, doi: 10.1051/0004-6361/201322423

Migliari, S., & Fender, R. P. 2006, MNRAS, 366, 79, doi: 10.1111/j.1365-2966.2005.09777.x

Migliari, S., Fender, R. P., Rupen, M., et al. 2004, Monthly Notices of the Royal Astronomical Society, 351, 186, doi: 10.1111/j.1365-2966.2004.07768.x

Migliari, S., Tomsick, J. A., Miller-Jones, J. C. A., et al. 2010, ApJ, 710, 117, doi: 10.1088/0004-637X/710/1/117

Miller-Jones, J. C. A., Bahramian, A., Altamirano, D., et al. 2023a, The Astronomer's Telegram, 16271, 1

Miller-Jones, J. C. A., Blundell, K. M., & Duffy, P. 2004, ApJL, 603, L21, doi: 10.1086/382954

Miller-Jones, J. C. A., Sivakoff, G. R., Bahramian, A., & Russell, T. D. 2023b, The Astronomer's Telegram, 16211, 1

Miller-Jones, J. C. A., Sivakoff, G. R., Altamirano, D., et al. 2010, ApJL, 716, L109, doi: 10.1088/2041-8205/716/2/L109

Miller-Jones, J. C. A., Moin, A., Tingay, S. J., et al. 2012, MNRAS, 419, L49, doi: 10.1111/j.1745-3933.2011.01176.x

Miller-Jones, J. C. A., Tetarenko, A. J., Sivakoff, G. R., et al. 2019, Nature, 569, 374, doi: 10.1038/s41586-019-1152-0

Mirabel, I. F., & Rodríguez, L. F. 1994, Nature, 371, 46, doi: 10.1038/371046a0

Mohan, N., & Rafferty, D. 2015, PyBDSF: Python Blob Detection and Source Finder, Astrophysics Source Code Library, record ascl:1502.007. http://ascl.net/1502.007

Motta, S. E., Casella, P., & Fender, R. P. 2018, MNRAS, 478, 5159, doi: 10.1093/mnras/sty1440

Motta, S. E., & Fender, R. P. 2019, MNRAS, 483, 3686, doi: 10.1093/mnras/sty3331

Motta, S. E., Tremou, E., R., F., et al. 2021, The Astronomer's Telegram, 14659, 1

Muñoz-Darias, T., Fender, R. P., Motta, S. E., & Belloni, T. M. 2014, MNRAS, 443, 3270, doi: 10.1093/mnras/stu1334

Muno, M. P., Remillard, R. A., & Chakrabarty, D. 2002, ApJL, 568, L35, doi: 10.1086/340269

Nakajima, M., Negoro, H., Serino, M., et al. 2023, The Astronomer's Telegram, 16206, 1

Natalucci, L., Bazzano, A., Cocchi, M., et al. 2000, ApJ, 536, 891, doi: 10.1086/308987

Negoro, H., Mihara, T., Nakahira, S., et al. 2018, The Astronomer's Telegram, 11593, 1

Negoro, H., Serino, M., Nakajima, M., et al. 2023, The Astronomer's Telegram, 16205, 1

Ng, M., Hughes, A. K., Homan, J., et al. 2024, ApJ, 966, 232, doi: 10.3847/1538-4357/ad35bd

Offringa, A. R., McKiconley, B., Hurley-Walker, N., et al. 2014, MNRAS, 444, 606, doi: 10.1093/mnras/stu1368

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

Page, K. L., Dichiara, S., Gropp, J. D., et al. 2023, GRB Coordinates Network, 34537, 1

Papitto, A., Ferrigno, C., Bozzo, E., et al. 2013, Nature, 501, 517, doi: 10.1038/nature12470

Parikh, A. S., Russell, T. D., Wijnands, R., et al. 2019, ApJL, 878, L28, doi: 10.3847/2041-8213/ab2636

Pe'er, A., & Casella, P. 2009, ApJ, 699, 1919, doi: 10.1088/0004-637X/699/2/1919

Peng, J.-Q., Zhang, S., Shui, Q.-C., et al. 2024, ApJL, 960, L17, doi: 10.3847/2041-8213/ad17ca

Perley, R. A. 1999, in Astronomical Society of the Pacific Conference Series, Vol. 180, Synthesis Imaging in Radio Astronomy II, ed. G. B. Taylor, C. L. Carilli, & R. A. Perley, 383

Plant, D. S., Fender, R. P., Ponti, G., Muñoz-Darias, T., & Coriat, M. 2014, MNRAS, 442, 1767, doi: 10.1093/mnras/stu867

Podgorny, J., Svoboda, J., & Dovciak, M. 2024, The Astronomer's Telegram, 16541, 1

Purcell, C. R., Van Eck, C. L., West, J., Sun, X. H., & Gaensler, B. M. 2020, RM-Tools: Rotation measure (RM) synthesis and Stokes QU-fitting, Astrophysics Source Code Library, record ascl:2005.003

Randall, K. E., Hopkins, A. M., Norris, R. P., et al. 2012, MNRAS, 421, 1644, doi: 10.1111/j.1365-2966.2012.20422.x

Remillard, R., Levine, A., Lin, D., & Smith, D. 2003, The Astronomer's Telegram, 204, 1

Remillard, R. A., & McClintock, J. E. 2006, ARA&A, 44, 49, doi: 10.1146/annurev.astro.44.051905.092532

Rhodes, L., Fender, R. P., Motta, S., et al. 2022, MNRAS, 513, 2708, doi: 10.1093/mnras/stac954

Rice, S. O. 1945, The Bell System Technical Journal, 24, 46, doi: 10.1002/j.1538-7305.1945.tb00453.x

Ruiz, M., Shapiro, S. L., & Tsokaros, A. 2021, Frontiers in Astronomy and Space Sciences, 8, 39, doi: 10.3389/fspas.2021.656907

Rushton, A., Spencer, R. E., Pooley, G., & Trushkin, S. 2010, MNRAS, 401, 2611, doi: 10.1111/j.1365-2966.2009.15838.x

Rushton, A. P., Miller-Jones, J. C. A., Curran, P. A., et al. 2017, MNRAS, 468, 2788, doi: 10.1093/mnras/stx526

Russell, D. M., Markoff, S., Casella, P., et al. 2013, MNRAS, 429, 815, doi: 10.1093/mnras/sts377

Russell, D. M., Bramich, D. M., Lewis, F., et al. 2019a, Astronomische Nachrichten, 340, 278, doi: 10.1002/asna.201913610

Russell, T. D., Carotenuto, F., Miller-Jones, J. C. A., et al. 2024, The Astronomer's Telegram, 16552, 1

Russell, T. D., Miller-Jones, J. C. A., Curran, P. A., et al. 2015, MNRAS, 450, 1745, doi: 10.1093/mnras/stv723

Russell, T. D., Tetarenko, A. J., Miller-Jones, J. C. A., et al. 2019b, ApJ, 883, 198, doi: 10.3847/1538-4357/ab3d36

Russell, T. D., Lucchini, M., Tetarenko, A. J., et al. 2020, MNRAS, 498, 5772, doi: 10.1093/mnras/staa2650

Russell, T. D., Degenaar, N., van den Eijnden, J., et al. 2021, MNRAS, 508, L6, doi: 10.1093/mnrasl/slab087

Samboco, V. G., Heywood, I., & Smirnov, O. 2024, SolarKAT: Solar Imaging Pipeline for MeerKAT, Astrophysics Source Code Library, record ascl:2401.013. http://ascl.net/2401.013

Sanidas, S., Caleb, M., Driessen, L., et al. 2018, in Pulsar Astrophysics the Next Fifty Years, ed. P. Weltevrede, B. B. P. Perera, L. L. Preston, & S. Sanidas, Vol. 337, 406–407, doi: 10.1017/S1743921317009310

Schwab, F. R. 1984, AJ, 89, 1076, doi: 10.1086/113605

Serra, P., Perkins, S., & Smirnov, O. 2022, in Astronomical Society of the Pacific Conference Series, Vol. 532, Astronomical Data Analysis Software and Systems XXX, ed. J. E. Ruiz, F. Pierfedereci, & P. Teuben, 409

Shahbaz, T., Russell, D. M., Covino, S., et al. 2016, MNRAS, 463, 1822, doi: 10.1093/mnras/stw2171

Shirey, R. E., Bradt, H. V., Levine, A. M., & Morgan, E. H. 1998, ApJ, 506, 374, doi: 10.1086/306247

Siddiqui, M. M. 1964, Journal of Research of the National Bureau of Standards, 68D, 1005

Smirnov, O. M. 2011a, A&A, 527, A107, doi: 10.1051/0004-6361/201116434

—. 2011b, A&A, 527, A108, doi: 10.1051/0004-6361/201116435

Smirnov, O. M., Heywood, I., Perkins, S. J., & van Rooyen, R. 2022, in Astronomical Society of the Pacific Conference Series, Vol. 532, Astronomical Data Analysis Software and Systems XXX, ed. J. E. Ruiz, F. Pierfedereci, & P. Teuben, 385

Smirnov, O. M., & Tasse, C. 2015, MNRAS, 449, 2668, doi: 10.1093/mnras/stv418

Sokoloff, D. D., Bykov, A. A., Shukurov, A., et al. 1998, MNRAS, 299, 189, doi: 10.1046/j.1365-8711.1998.01782.x

Spencer, R. E., Rushton, A. P., Bałucińska-Church, M., et al. 2013, MNRAS, 435, L48, doi: 10.1093/mnrasl/slt090

Svoboda, J., Dovčiak, M., Steiner, J. F., et al. 2024, arXiv e-prints, arXiv:2403.04689, doi: 10.48550/arXiv.2403.04689

Tarana, A., Bazzano, A., Ubertini, P., et al. 2006, A&A, 448, 335, doi: 10.1051/0004-6361:20053917

Tarana, A., Capitanio, F., & Cocchi, M. 2018, MNRAS, 477, 3353, doi: 10.1093/mnras/sty816

Tasse, C., Hugo, B., Mirmont, M., et al. 2018, A&A, 611, A87, doi: 10.1051/0004-6361/201731474

Taylor, A. R., & Legodi, L. S. 2024, AJ, 167, 273, doi: 10.3847/1538-3881/ad4150

Taylor, G., Carilli, C., Perley, R., & (U.S.), N. R. A. O. 1999, Synthesis Imaging in Radio Astronomy II: A Collection of Lectures from the Sixth NRAO/NMIMT Synthesis Imaging Summer School Held at Socorro, New Mexico, USA, 17-23 June, 1998, Astronomical Society of the Pacific conference series (Astronomical Society of the Pacific). https://books.google.ca/books?id=e23vAAAAMAAJ

Tetarenko, A. J., Sivakoff, G. R., Miller-Jones, J. C. A., et al. 2017, MNRAS, 469, 3141, doi: 10.1093/mnras/stx1048

Tetarenko, B. E., Sivakoff, G. R., Heinke, C. O., & Gladstone, J. C. 2016a, ApJS, 222, 15, doi: 10.3847/0067-0049/222/2/15

—. 2016b, ApJS, 222, 15, doi: 10.3847/0067-0049/222/2/15

Ubertini, P., in 't Zand, J., Tesseri, A., Ricci, D., & Piro, L. 1998, IAUC, 6838, 1

Uttley, P., & Klein-Wolt, M. 2015, MNRAS, 451, 475, doi: 10.1093/mnras/stv978

Vaillancourt, J. E. 2006, PASP, 118, 1340, doi: 10.1086/507472

van den Eijnden, J., Degenaar, N., Russell, T. D., et al. 2018, Nature, 562, 233, doi: 10.1038/s41586-018-0524-1

van den Eijnden, J., Fender, R., Miller-Jones, J. C. A., et al. 2022, MNRAS, 516, 2641, doi: 10.1093/mnras/stac2392

van der Klis, M. 2004, arXiv e-prints, astro, doi: 10.48550/arXiv.astro-ph/0410551

—. 2006, in Compact stellar X-ray sources, Vol. 39 (Cambridge University Press), 39–112

van der Laan, H. 1966, Nature, 211, 1131, doi: 10.1038/2111131a0

van Straaten, S., van der Klis, M., & Wijnands, R. 2005, ApJ, 619, 455, doi: 10.1086/426183

van den Eijnden, J., Degenaar, N., Russell, T. D., et al. 2021, Monthly Notices of the Royal Astronomical Society, 507, 3899, doi: 10.1093/mnras/stab1995

Wardle, J. F. C. 1977, Nature, 269, 563, doi: 10.1038/269563a0

Wardle, J. F. C., Homan, D. C., Ojha, R., & Roberts, D. H. 1998, Nature, 395, 457, doi: 10.1038/26675

Watts, A. L. 2012, ARA&A, 50, 609, doi: 10.1146/annurev-astro-040312-132617

Williams, D. R. A., Motta, S. E., Fender, R., et al. 2020, MNRAS, 491, L29, doi: 10.1093/mnrasl/slz152

Wilms, J., Allen, A., & McCray, R. 2000, ApJ, 542, 914, doi: 10.1086/317016

Wood, C. M., Miller-Jones, J. C. A., Bahramian, A., et al. 2024, arXiv e-prints, arXiv:2405.12370, doi: 10.48550/arXiv.2405.12370

Woosley, S. E., & Bloom, J. S. 2006, ARA&A, 44, 507, doi: 10.1146/annurev.astro.43.072103.150558

Xie, F.-G., Yan, Z., & Wu, Z. 2020, ApJ, 891, 31, doi: 10.3847/1538-4357/ab711f

Ye, H., Gull, S. F., Tan, S. M., & Nikolic, B. 2022, MNRAS, 510, 4110, doi: 10.1093/mnras/stab3548

Yu, W. 2023, The Astronomer's Telegram, 16276, 1

Zdziarski, A. A., Tetarenko, A. J., & Sikora, M. 2022, ApJ, 925, 189, doi: 10.3847/1538-4357/ac38a9

Zhang, S. N., Cui, W., Harmon, B. A., et al. 1997, ApJL, 477, L95, doi: 10.1086/310530