UNIVERSITY OF AMSTERDAM

MSc Astronomy & Astrophysics
GRAPPA

Master Thesis

Probing the short-timescale evolution of quasi-periodic
oscillations

by

Jakob van den Eijnden
10725423

July 2016

54 ECTS

June 2015 – July 2016

Supervisor:

Dr. P. Uttley
Dr. A. Ingram

Examiner:

Prof. dr. M. van der Klis

Anton Pannekoek Institute for Astronomy

*And suddenly as I gazed upon the night
Well I notice the stars
They began to shake and dance and burst
And fall into the darkness*

Jason Mraz – 'Galaxy'

Abstract

The X-ray emission of black hole X-ray binaries (BHXRBS) often shows quasi-periodic oscillations (QPOs), that originate from the close proximity of the black hole. As such, QPOs trace the behaviour of the accreted material in the most extreme gravitational environments. However, the physical origin of the QPO is still the subject of debate: it could either be geometric, arising from systematic changes in the accretion geometry, or intrinsic, resulting from modulations in intrinsic properties of the accretion flow. In this thesis, I present the results of two novel approaches to understanding the QPO origin, both investigating the observed lag between the QPO signal in different energy bands. After introducing the timing properties of black holes, I first discuss the short-timescale evolution of the QPO lag in the BHXRBS GRS 1915+105. Unexpectedly, I find that this lag increases systematically over 5–10 QPO cycles. Simultaneously, the QPO amplitude rises and falls in an enveloping fashion, dubbed coherent intervals. To interpret this result, I formulate a geometric toy model for the QPO: differential vertical precession of the inner accretion flow. In addition, I present a systematic analysis of the QPO lag in a sample of fifteen BHXRBS. I find that the lag associated with the Type-C QPO significantly depends on the source inclination, strongly suggesting a geometric QPO origin. The found inclination dependence is consistent with Lense-Thirring precession as the QPO origin, taking into account dominant relativistic effects. Finally, I also show preliminary results of a coherent-interval resolved analysis of the QPO in GRS 1915+105, which provides a promising new handle on the origin of QPOs.

CONTENTS

1	INTRODUCTION	1
1.1	black hole X-ray binaries in outburst	3
1.1.1	Geometrically thin, optically thick accretion disks . . .	3
1.1.2	BHXRBB states	4
1.2	X-ray variability and quasi-periodic oscillations	7
1.2.1	Visualizing variability	7
1.2.2	Types of quasi-periodic oscillations	9
1.2.3	Energy dependence of QPO properties	11
1.2.4	QPO models and Lense-Thirring precession	13
1.2.5	Recent observational evidence of Lense-Thirring precession	16
1.3	Thesis outline	17
2	THE SHORT-TIMESCALE EVOLUTION OF QPO PHASE LAGS IN GRS 1915+105	19
2.1	Observations and timing analysis	20
2.1.1	Data reduction and optimal filtering	23
2.1.2	Phase lags	25
2.2	Results	27
2.2.1	Test 1: frequency-amplitude correlations	27
2.2.2	Test 2: phase lag evolution	29
2.3	Discussion	34
2.3.1	Robustness of the method	34
2.3.2	Phase lags	34
2.3.3	Decoherence of the QPO	36
2.3.4	A unifying model: differential precession and spectral evolution	38
2.4	Conclusion	42
3	INCLINATION DEPENDENCE OF PHASE LAGS IN BLACK HOLE X-RAY BINARIES	45
3.1	Sample and Data Analysis	46
3.2	Results	48
3.2.1	Type-C QPO fundamental lags	51
3.2.2	Significance testing	54
3.2.3	Type-B QPO fundamental lags	56
3.2.4	Harmonics and broad-band noise	56
3.2.5	Energy-dependent frequency differences of the Type-C Fundamental	57
3.3	Discussion	59
3.3.1	Type-C QPO lags	59
3.3.2	Type-B QPO lags	62
3.3.3	Broad-band noise	62
3.3.4	Individual sources	63
3.3.5	QPO frequency differences and differential precession	65
3.4	Conclusions	67
4	COHERENT-INTERVAL RESOLVED QPO ANALYSIS	69
4.1	Method	70

4.1.1	Full energy band analysis	71
4.1.2	Broad energy band analysis	71
4.1.3	Narrow energy band analysis	72
4.2	Results	72
4.2.1	Basic coherent interval properties	72
4.2.2	Hardness modulations	75
4.2.3	Coherent-interval resolved lag-energy spectra	78
4.3	Discussion	79
4.3.1	Coherent intervals: connecting QPO and BBN?	79
4.3.2	The evolution of lag-energy spectra	82
4.4	Conclusion and future outlook	83
	SUMMARY	85
	POPULAR SCIENTIFIC SUMMARY	87
A	SPECTRAL-TIMING METHODS	95
A.1	A single time series: power-spectral density	95
A.2	Two time series: cross spectra and CCFs	96
A.3	Even more time series: rms spectra and lag energy spectra	97
A.4	Simulating noise processes	98
B	BINARY-ORBIT INCLINATION ESTIMATES	103
C	ADDITIONAL PLOTS FOR INDIVIDUAL SOURCES	105

INTRODUCTION

Ever since their discovery, black holes have been among the most captivating objects in the Universe. Black holes were first predicted from Einstein's theory of general relativity by Karl Schwarzschild in 1916 ([Schwarzschild, 1916](#)). However, it took over half a century for the first observational evidence of black holes to be found and accepted. The first X-ray observations of accreting compact objects (i.e. black holes and neutron stars) were made in the nineteen-sixties ([Giacconi et al., 1962](#)), although it took several more years before Cygnus X-1 was widely accepted as the first confirmed black hole ([Shipman, 1975](#)). Supermassive black holes, residing in the centers of galaxies, were first found as bright radio sources. The cosmological distances derived from redshift estimates by [Schmidt \(1963\)](#) implied intrinsic luminosities that could only be explained by accretion processes onto a super-massive black hole.

Since these first observations, stellar-mass black holes have been extensively studied in black-hole binaries (BHXBs): binary systems consisting of both a black hole and a companion star. When matter from this companion star is accreted onto the black hole, it forms an accretion disk and heats up, emitting light over a wide range of frequencies. This accretion flow, famously modeled by [Shakura & Sunyaev \(1973\)](#), emits brightly in X-rays from its innermost regions. However, the closer to the black hole, the less we understand the structure and properties of the accreting material. Fully grasping the physics of this innermost region of the accretion flow, or of accretion and outflows in general, is an important challenge in modern astronomy. It allows us to investigate the behaviour of matter in extreme gravitational fields and thus test predictions of general relativity. On larger scales, accretion and the feedback from it play a vital role in the formation of structures in the universe and in the evolution of galaxies hosting a super-massive black hole ([Done et al., 2007](#)).

One approach to understanding the accretion flow is through the study of variability. The flow's X-ray emission is highly variable on a wide range of timescales, from outbursts lasting months or years, to variations on subsecond timescales. While most of the variability is aperiodic, i.e. without characteristic repeating patterns, BHXBs often show *quasi-periodic oscillations* (QPOs): oscillations of varying frequency and amplitude, detected as finite-width peaks in X-ray power spectra on top of broad band noise (BBN) (see e.g. [Van der Klis, 2006](#)). An illustrative power spectrum of a BHXB, containing such a QPO, together with the corresponding lightcurve is shown in Figure 1. So-called low-frequency QPOs, with periods from 0.05 to 10 seconds, are believed to originate from the inner region of the accretion flow that are associated with the hardest X-ray emission ([Axelsson et al., 2013](#); [Sobolewska & Zychi, 2006](#)). Hence, QPOs are a powerful tool to study the direct surroundings of the accreting black hole. However, currently, the physical origin of the QPOs is still subject to debate.

Figure 1.: Illustrative example of a (white-noise-subtracted) power spectrum (top) and corresponding lightcurve segment (bottom) of a BHXR. The X-ray flux is evidently highly variable. A rough periodicity can be seen in the lightcurve, corresponding to the large, extended peak in the power spectrum: the QPO. The dotted lines show Lorentzian fits to the QPO fundamental and harmonic feature. Source: GRS 1915+105. RXTE ObsID: 10408-01-22-02.

Over the past years, novel spectral-timing methods have greatly increased our understanding of accretion flows onto black holes (see e.g. Uttley et al., 2014, for an overview of such methods). These methods combine variability and spectral information, instead of considering both separately. An example of such a technique is the investigation of reverberation lags: time lags between photons of different energies due to differences in pathlength between direct and reflected emission, which can be applied to map the geometry of the accretion flow (see e.g. Fabian et al., 2009). Another example is the detection of variations in the X-ray spectrum as a function of position in the quasi-periodic oscillation cycle, which can be linked directly to changes in the geometry of the accretion flow (Ingram & Van der Klis, 2015; Stevens & Uttley, 2016).

In this thesis, I present novel results from both newly developed and existing spectral-timing methods, applied to accreting black hole X-ray binaries. I have developed a new method to track the properties of the QPO on short timescales. This method allows me to investigate the evolution and stability of the underlying physical process causing the QPO. Our results suggest that the QPO originates from vertical *differential precession* of the innermost accretion flow, setting up and decaying within seconds. I also study the binary system inclination dependence of phase lags associated with the QPO. I find that these lags are strongly dependent on our line of sight, providing further evidence that the QPO has a geometric origin.

This thesis is structured as follows: in the remainder of this first Chapter, I will introduce the physics of accretion, and the phenomenological properties of and models for X-ray variability and QPOs. In Chapter two, I present the new method to track QPO properties on short timescales, and its results on phase lags associated with the QPO. In the third Chapter, I present my analysis of the inclination dependence of these phase lags by comparing fifteen BHXRBS. Lastly, in Chapter four, I discuss the short-timescale evolution of spectral(-timing) BHXRBS properties, using the new method. Chapter two has been adapted from Van den Eijnden et al. (2016); Chapter three has been submitted to MNRAS for publication.

1.1 BLACK HOLE X-RAY BINARIES IN OUTBURST

1.1.1 Geometrically thin, optically thick accretion disks

X-ray binaries (XRBS) consist of a compact object, either a black hole or a neutron star, accreting material from a companion star. Based on the mass of this companion star, these systems can be classified as either low-mass or high-mass XRBS. While in the latter stellar winds could contribute to the accretion of matter onto the compact object, for low-mass XRBS accretion occurs through Roche-lobe overflow. In this process, named after the Roche-potential of two orbiting masses, material in the outer layers of the companion star crosses the first Lagrangian point and starts to orbit the black hole instead. This process arises when the companion star fills its Roche lobe, either due to expansion after leaving the main sequence or due to a decrease in orbital separation, resulting for instance from the emission of gravitational radiation (Frank et al., 2002; Shakura & Sunyaev, 1973).

Due to the conservation of angular momentum, the accreted gas cannot directly fall onto the compact object. Instead, it will assume a circular orbit, as for a given angular momentum, a circular orbit is the lowest energy solution. Viscous dissipation, caused most probably by magnetic shear (the

Magnetorotational Instability, Balbus & Hawley 1991, 1992; Hawley & Balbus 1991, 1992) is responsible for the outward transportation of angular momentum, allowing the material to accrete onto the compact object. This gradual outward angular momentum transportation leads to the creation of a geometrically thin, optically thick disk with a temperature increasing towards smaller radii. Each radius emits a blackbody spectrum at the local temperature (Frank et al., 2002; Shakura & Sunyaev, 1973). The main difference between the two possible compact objects is the surface present in a neutron star. However, as we will see, this thin disk description does not always adequately describe the regions closest to the compact object.

X-ray binaries show large outbursts, which can be attributed to the so-called hydrogen ionisation instability. In quiescence, the temperature is low and the disk consists mostly of neutral hydrogen. However, if the temperature rises enough so that hydrogen starts to ionize, the opacity increases drastically. This increase traps photons in the disk, increasing the temperature and thus the ionization even more. This runaway process continues until most of the hydrogen is ionized, and the disk temperature has risen significantly. As the mass accretion rate increases with temperature, more material is accreted and thus more energy is liberated. The result is a quick and enormous increase in luminosity: the start of an outburst (Done et al., 2007; Frank et al., 2002).

As the mass accretion rate quickly increases, the disk becomes locally depleted of material. As a result, the pressure and consequently temperature decrease, and the same thermal runaway happens in reverse. Finally, what remains is again a low temperature, neutral disk, with a low accretion rate and thus low luminosity. Although this scenario is a local effect, the changes in mass accretion link different regions in the accretion flow together. Hence, the entire system can undergo this instability and show an outburst (Done et al., 2007; Frank et al., 2002).

1.1.1.2 BHXRb states

During an outburst, BHXRb move through different states, in which the geometry of the accretion flow changes drastically. This evolution is observed through spectral changes, as different regions in the accretion flow contribute differently to the observed X-ray spectrum. Generally speaking, the BHXRb X-ray spectrum consists of three components, depicted schematically in Figure 2: at low X-ray energies, the standard geometrically thin, optically thick disk emits a multicolour blackbody spectrum - the sum of the individual blackbody spectra emitted at radii of different temperature. The innermost region of the accretion flow consists of a hot corona, which is not fully understood. Energetic electrons in this corona Compton upscatter photons from the disk into a power-law spectral component at higher energies. The third component is the reflected spectrum, consisting of upscattered photons reflecting off the disk. This component carries fluorescence lines from the disk material, most famously the iron $K\alpha$ line around 6.4 keV. This geometry is depicted schematically in Figure 3, with colours corresponding to those in Figure 2 (Gilfanov, 2010).

A powerful interpretation of BHXRb behaviour during outbursts is the *truncated disk model* (see e.g. Done et al., 2007; Esin et al., 1997; Frank et al., 2002). The BHXRb starts out in the so-called hard state, where the spectrum is dominated by the high-energy power-law emission. Physically, this state corresponds to a large truncation radius, where the geometrically thin

Figure 2.: Cartoon depiction of the three main components of a BHXRB X-ray spectrum: the soft multi-colour blackbody disk spectrum in blue, the hard comptonized power law in red, and the reflected spectrum, with $K\alpha$ line, in green. Illustration by [Gilfanov \(2010\)](#)

Figure 3.: Cartoon depiction of the accretion geometry around a black hole. The colours of the emission components matches the colours in Figure 3. Illustration by [Gilfanov \(2010\)](#)

Figure 4.: Example of a hard and soft state spectrum in BHXRb Cygnus X-1.
Figure from [Gilfanov \(2010\)](#)

disk transitions into the corona. The prominence of this corona results in the dominant power-law in the spectrum. During the outburst, the mass accretion rate increases, causing the thin disk to move further inward. As a result, the BHXRb transitions into the soft state, where the spectrum is dominated by the soft multi-colour black body emission from the now prominent disk. Eventually, the high mass accretion rate depletes the disk of material, causing the disk to return to the initial, large truncation radius - the source returns to the hard state and quiescence.

Examples of a hard and soft state spectrum in Cygnus X-1 are shown in Figure 4. Inbetween the hard and soft state, several intermediate states are crossed, an overview of which can be found in [Belloni \(2010\)](#). It is also important to reiterate that the structure of the corona is still ill-understood: instead of inbetween the black hole and the disk, it might also be located above the disk ([Galeev et al., 1979](#); [Gilfanov, 2010](#)) or in a lamppost geometry above the black hole ([Fabian et al., 2009](#); [Markoff et al., 2005](#)).

Chapters 2 and 4 of this Thesis contain a study of the source GRS 1915+105. As this source shows peculiar spectral behaviour and has been in outburst since its discovery, I briefly introduce it here. GRS 1915+105 is a galactic low-mass BHXRb located at a distance of $8.6^{+2.0}_{-1.6}$ kpc ([Reid et al., 2014](#)), that was discovered in 1992 by [Castro-Tirado et al. \(1992\)](#). It shows a wide variety of variability properties ([Belloni et al. 1997a,b](#), also see Section 1.2), which are

difficult to interpret within the standard picture of BHXRB states (Belloni 2010, see Van Oers et al. 2010 for a spectral comparison). Belloni et al. (2000) report the presence of 12 accretion classes based on its observed properties, which can be interpreted as transitions between three main states: a hard state (C) where the disk is truncated, and two soft states (A & B, with a low and high flux respectively) where the inner disk extends further inwards. However, all three states show similarities to the canonical intermediate state of BHXRBs (Reig et al., 2003). The peculiarity of GRS 1915+105 is important when generalizing the results from Chapters 2 and 4.

1.2 X-RAY VARIABILITY AND QUASI-PERIODIC OSCILLATIONS

The X-ray emission of BHXRBs is variable on an enormous range of timescales. In the previous section, I introduced the variability on the longest timescales: the transient nature of BHXRBs with outbursts lasting up to months or even years. This timescale corresponds to the viscous timescale in the accretion flow - the timescale associated with for instance the propagation of changes in mass accretion rate (Done et al., 2007). Hence, as I discussed, this variability is associated with large scale changes in the geometry of the accretion flow. However, depending on the spectral state of the system, the X-ray emission of BHXRBs can also be variable on timescales down to milliseconds. Contrary to the slower outbursts, this fast variability does not correspond to global changes in the system. The fastest variability timescales (frequencies in the hundreds of Hz) are comparable to the dynamical timescale at the innermost edge of the accretion disk (Frank et al., 2002; Stella & Vietri, 1998). Hence, studying this fast X-ray variability allows us to probe the highly dynamical nature and properties of the (inner) accretion flow, within one of the aforementioned global BHXRB states.

In order to study these shortest timescales, X-ray observations with at least millisecond timing resolution are required. The field of X-ray timing was revolutionized by the launch of the *Rossi X-ray Timing Explorer* (RXTE) on the 31st of December 1995, a space observatory carrying the Proportional Counter Array (PCA): an X-ray detector, sensitive from 2 up to 75 keV (originally), with a timing resolution down to 1 μ s. The energy resolution of RXTE, although worse than more recent X-ray observatories such as XMM-Newton or NuSTAR, is more than sufficient to perform standard spectral analysis of BHXRBs¹. Despite ending observations in 2012, the RXTE data archive is still analysed today. This entire thesis is based on the analysis of archival RXTE observations.

1.2.1 Visualizing variability

The characteristics of the short-timescale variability of BHXRBs are best visualized in a power spectrum²: the power of variability at a given Fourier frequency (e.g. 1/timescale), plotted as a function of that frequency. In Figure 5, I show a set of example power spectra from several BHXRBs ranging from the hard to soft spectral state, adapted from Heil et al. (2015a). In all power spectra, a broad noise component, sometimes including a break, is visible: the so-called broad band noise (BBN). The origin of this BBN is most probably linked to mass accretion rate fluctuations propagating through the

¹ See <http://heasarc.gsfc.nasa.gov/docs/xte/xteof.html>

² In this thesis, we will interchangeably use both power spectrum and power-spectral density/PSD.

Figure 5.: Examples of power spectra for several different BHXRBS, showing the power of variability on different timescales. The numbering is based on power-colour properties as defined by [Heil et al. \(2015a\)](#). The BHXRBS binary states in the panels move from the hard state (start of outburst) to the soft state (peak of outburst). Figure from [Heil et al. \(2015a\)](#).

disk ([Arévalo & Uttley, 2006](#); [Ingram & Van der Klis, 2013](#); [Kotov et al., 2001](#); [Lyubarskii, 1997](#)). As a higher mass accretion rate results in a higher flux, this propagating fluctuations lead to aperiodic variability.

A more nuanced timing property of BHXRBS, which will be the most important observable in this thesis, is the phase or time lag. While a power spectrum is calculated from a single lightcurve, phase lags contain information on two lightcurves, usually at different energies. The phase or time lag is the lag between the arrival of a correlated signal in two different time series; the simplest example is the $\pi/2$ phase lag between a sine and a cosine of the same frequency. Time lags are expressed in seconds, while phase lags are expressed in radians. While a power spectrum merely contains information on the frequency-dependent amplitude of variability, lags can more directly provide information on the process causing the variability. For instance, lags can be introduced by differences in lightcrossing times, or by delays due to the reprocessing of a signal.

Phase and time lags between two time series can be calculated for a range of timescales simultaneously. This is a powerful tool: variability on different timescales will likely originate from different processes and thus lags at these timescales do not necessarily have the same values. I use this property of the lags extensively, and will thus often refer to the phase or time lag at/around a certain Fourier frequency. For a more detailed introduction to lags, including their numerical calculation, I refer the reader to Appendix A. This Appendix also contains background information on the calculation and properties of the power spectrum.

In several of the power spectra in Figure 5, narrow peaks can be identified on top of the BBN. Although not purely periodic, these peaks represent an

Figure 6.: Power spectral fit of the LFQPO in two observations of BHXRBS GRS 1915+105. The BBN is modeled by a broken power law, while the QPO fundamental and harmonic are fitted with Lorentzian peaks. Adapted from [Pahari et al. \(2013\)](#).

almost periodic phenomenon in the accretion flow. Given their extended shape, they are referred to as *quasi-periodic oscillations* (QPOs). As I will discuss further on, their origin is currently a subject of debate, and this thesis will focus on contributing to the search for the QPO mechanism. I will introduce this search in the next section, by introducing both the observed properties of QPOs and their possible models in detail.

1.2.2 Types of quasi-periodic oscillations

A large variety of QPOs exists in XRBs, both in BHXRBS and accreting neutron star systems. This variety is not surprising: a narrow, but extended power spectral peak simply corresponds to any oscillation with either a drifting frequency or variable amplitude (or both). Thus, not all QPO types necessarily represent the same process. In BHXRBS, QPOs are divided into two categories: Low-Frequency QPOs (LFQPOs), with frequencies ranging from ~ 0.1 to ~ 30 Hz, and High-Frequency QPOs (HFQPOs), with frequencies from ~ 67 to a few hundreds of Hz ([Belloni & Altamirano, 2013](#)). While the former are ubiquitous in BHXRBS, the latter are quite rare; HFQPOs have only been observed in a handful of BHXRBS ([Belloni et al., 2012](#)). In accreting neutron stars, LFQPOs are also observed. Moreover, neutron stars show kHz QPOs, with, as the name suggests, frequencies in the kHz range. Contrary to HFQPOs in BHXRBS, kHz QPOs are a common feature in accreting neutron stars (see e.g. [Van der Klis, 2006](#)). In this thesis, I will focus only on the LFQPOs in BHXRBS. Hence, I will refer to them simply as QPOs. Examples of such LFQPOs in BHXRBS are present in Figure 1 and several panels in Figure 5.

QPOs in BHXRBS are usually modelled in the power spectrum as a Lorentzian peak on top of the BBN. This description allows for the determination of several interesting characteristics: the centroid frequency of the QPO, its Q-factor, defined as the centroid frequency divided by the FWHM, and its variance: the integrated area below the fitted Lorentzian peak. In Figure 6, I show two more examples of power spectra containing a QPO. This Figure is adapted from [Pahari et al. \(2013\)](#) and shows the fitted components to the QPO and BBN for two observations of BHXRBS GRS 1915+105.

Clearly visible in Figure 6 is also a second, smaller power spectral peak. Similar features accompany many of the LFQPOs present in Figure 5, in

Table 1.: Overview of the basic properties of Type-A, B and C LFQPOs in BHXRBs. Reproduced from Table 1 in [Casella et al. \(2005\)](#). Hard lags imply that the QPO fundamental/(sub)harmonic signal in the hard energy band lags behind the signal in the soft band.

Property	Type-C	Type-B	Type-A
Frequency [Hz]	$\sim 0.1 - 15$	$\sim 5-6$	~ 8
Q-factor	$\sim 7 - 12$	≥ 6	≤ 3
Amplitude [% rms]	3–16	$\sim 2-4$	≤ 3
Noise	Strong flat top	Weak red	Weak red
Phase lag at fundamental	Soft & hard	Hard	Soft
Phase lag at harmonic	Hard	Soft	–
Phase lag at subharmonic	Soft	Soft	–

all cases at a frequency approximately double the centroid frequency of the primary peak. A third peak is sometimes observed as well, at half the frequency of the largest peak. These two or three features are all harmonics of the same physical phenomenon. However, for clarity, I will refer to them as the subharmonic, fundamental and harmonic, in order of increasing frequency.

QPOs can be further divided into three types, introduced by [Wijnands et al. \(1999\)](#) and investigated extensively by [Casella et al. \(2005\)](#): Type-A, B and C. Below, I list their most important properties:

- C: Type-C QPOs are the most common: narrow, strong features on top of a flat-top BBN, showing large drifts in frequency between observations. Type-C QPOs are often accompanied by harmonic features, and are located on top of so-called band-limited noise which can be modelled by one or multiple broad Lorentzians. In Figure 5, Type-C QPOs are visible in for example panels 8 to 12.
- B: Type-B QPOs are weaker and show less drift in frequency. The BBN associated with Type-B QPOs is weak and at most consists of weak flicker noise³ at low frequencies. Examples of Type-B QPOs are present in panels 13 and 14 in Figure 5.
- A: Finally, Type-A QPOs are the broadest and weakest category, and are rarely observed. In this thesis, I do not consider any Type-A QPO observations.

Table 1, reproduced from [Casella et al. \(2005\)](#), contains a comparison of basic QPO properties between the three types. The phase lags, introduced conceptually in the previous Section, are especially important, as these are the focus of Chapter 3. As I will discuss extensively in that Chapter as well, the different QPO types are thought to originate from different physical mechanisms ([Motta et al., 2012](#); [Stevens & Uttley, 2016](#)). In this thesis, the focus will primarily be on the Type-C QPO, as this QPO type occurs most often and has been most extensively studied. Only in Chapter 3, I also briefly consider Type-B QPOs.

The Type-C QPO shows a large drift in frequency between observations. This behaviour can be confirmed visually in Figure 5: the QPO frequency shifts by more than a factor of 10 between panels eight and sixteen. Although these plotted power spectra correspond to different BHXRBs, the

³ Power spectral noise corresponding to a random walk in the time domain; $P(\nu) \sim \nu^{-1}$.

same behaviour is present in individual sources. During this frequency drift, the BHXRB transitions further into the soft state. In the truncated disk model, this change corresponds to a decreasing truncation radius and a smaller corona. This suggests that the Type C QPO originates from this corona, as its characteristic frequencies will increase as its size decreases. This interpretation is complemented by the analyses of [Sobolewska & Zycki \(2006\)](#) and [Axelsson et al. \(2013\)](#), which show that the Comptonized spectral component shows much larger variability at the QPO frequency than the disk emission.

1.2.3 Energy dependence of QPO properties

An effective method to gain understanding of the physical origin of the QPO and the associated locations in the accretion flow, is to study the energy dependence of QPO properties. As I discussed in section 1.1, different energies are dominated by different spectral components. Hence, changes in QPO properties between lightcurves of different energies are a useful tool in probing the origin of the QPO. Here, I present a short overview of the literature on the energy dependence of the Type-C QPO. The interpretation of these relations within QPO models will be a recurring theme in the discussions in this thesis.

Easiest to interpret in terms of spectral components is the energy dependence of the QPO amplitude, i.e. the square root of the integral under the QPO peak for lightcurves of different energies. This dependence has been studied extensively in a range of BHXRBs. The QPO amplitude is consistently found to increase as a function of energy ([Lehr et al., 2000](#); [Morgan et al., 1997](#); [T. Belloni et al., 1997](#); [Yan et al., 2013](#)). In other words, the QPO is stronger in harder energy bands. This is consistent with the observation by [Sobolewska & Zycki \(2006\)](#) and [Axelsson et al. \(2013\)](#) that the Comptonized spectral component is more variable than the disk. Furthermore, it is consistent with the evolution of QPO frequency with spectral state in the truncated disk model.

Recently, the value of the QPO frequency has been found to change between energy bands in several BHXRBs. First, this energy dependence of the QPO frequency in the BHXRB GRS 1915+105 was studied by [Qu et al. \(2010\)](#) and [Yan et al. \(2012\)](#). [Yan et al. \(2012\)](#) analysed all RXTE observations of GRS 1915+105 up to 2010, and found a smooth evolution of the dependence of QPO frequency on photon energy. For observations with a low QPO frequency ($\sim 0.4 - 2.0$ Hz) in the full energy band, the QPO frequency decreases with energy, while for observations with a high QPO frequency ($\sim 2.0 - 8.0$ Hz) in the full energy band, the QPO frequency increases with energy. Similarly, [Li et al. \(2013a\)](#) found an increase in QPO frequency with energy in XTE J1550-564 for frequencies above ~ 3.3 Hz. However, below ~ 3.3 Hz no variations with energy were observed. [Li et al. \(2013b\)](#) reported comparable behaviour in the QPO in H1743-322, also exclusively showing frequency increases with photon energy. GRS 1915+105 is thus the only source to systematically show decreases of QPO frequency with energy. In Chapter 2, I will investigate the origin of the energy-dependent QPO frequency in GRS 1915+105.

Similarly, the phase lag at the QPO frequency⁴ is known to be energy dependent. This can be investigated in a so-called *lag-energy spectrum*: the lags between a fixed reference band lightcurve and lightcurves in different

⁴ Also referred to as the QPO (phase) lag.

Figure 7.: The Type-C QPO lag in GRS 1915+105 plotted as a function of the QPO frequency. The lag depends log-linearly on QPO frequency, switching sign at approximately 2 Hz. Adapted from [Reig et al. \(2000\)](#).

energy bands, as a function of the energy of the latter band. [Reig et al. \(2000\)](#), [Qu et al. \(2010\)](#) and [Pahari et al. \(2013\)](#) all report a smooth relation between this QPO phase lag and energy in GRS 1915+105. The same behaviour is present in XTE J1550-564 ([Wijnands et al., 1999](#)). However, the slope of this energy-dependence can vary wildly between observations. Furthermore, the QPO phase lags often show a power law break, which might be attributable to differences between the disk and the corona. Interestingly, the lag at the harmonic feature often shows a different dependence on energy than that associated with the QPO fundamental, suggesting differences in origin between the QPO fundamental and the harmonic.

Futhermore, by comparing multiple observations of GRS 1915+105, [Reig et al. \(2000\)](#), [Qu et al. \(2010\)](#) and [Pahari et al. \(2013\)](#) also show that the phase lag between hard and soft photons decreases approximately log-linearly as a function of QPO frequency in the full band, switching from a hard to soft lag around ~ 2 Hz. In Figure 7, I plot the Figure from [Reig et al. \(2000\)](#) showing this result. A similar sign change is also present in the QPO lags in several other BHXRBs (see Chapter 3). This sign change is particularly interesting, as it puts a clear constraint on possible models for the QPO phase lag. However, no systematic study of the relation between QPO frequency and QPO phase lag across a set of sources has been performed prior to this thesis. In Chapter 3, I present the results of such a study.

1.2.4 QPO models and Lense-Thirring precession

In the previous section, I have presented a short overview of observational properties of mostly Type-C QPOs. In this section, I will discuss theoretical models that have been developed to explain these observed Type-C QPO properties. At the end of the section, I will briefly touch on the origin of Type-B QPOs.

Currently, there is no complete consensus on the origin of Type-C QPOs. We can however divide the existing models into two broad categories: *geometric* and *intrinsic* models. In geometric models, the BHXRB's X-ray emission is intrinsically constant (barring the aperiodic BBN). However, an oscillating geometry causes quasi-periodicity in the observed flux. In intrinsic models, it is the intrinsic X-ray emission that is variable, resulting from non-geometric changes in the accretion flow.

Intrinsic variations in X-ray brightness could arise from the setup of magneto-acoustic waves in the corona (Cabanac et al., 2010) or magnetically-driven density waves in the disk (Tagger & Pellat, 1999). Alternatively, these could arise from the formation and presence of a standing shock in the accretion flow (Chakrabarti & Molteni, 1993). However, although the origin of Type-C QPOs is still debated, most recent results point towards geometric models. In addition, the result presented in this thesis are most consistent with a geometric QPO origin. Hence, in the remainder of this section, I will focus on the most prominent geometric model: *Lense-Thirring precession*.

In general relativity, space-time around a rotating black hole is described by the Kerr-metric. A test particle orbiting a black hole in this metric has multiple characteristic frequencies in addition to the Keplerian frequency ν_ϕ . One of these frequencies is the frequency of Lense-Thirring precession: the vertical precession of the test particle orbit due to a misalignment between the orbit's angular momentum and the black hole spin. In terms of the Keplerian frequency and the dimensionless spin parameter⁵ of the black hole a , the Lense-Thirring precession frequency is given by (Merloni et al., 1999, using natural units)

$$\nu_{\text{LT}} = \nu_\phi \left(1 - \sqrt{1 - \frac{4a}{r^{3/2}} + \frac{3a^2}{r^2}} \right) \quad (1)$$

If $a \ll 1$ or $r \gg 1$, the dependence of the Lense-Thirring precession frequency on radius reduces to $\nu_{\text{LT}} \propto r^{-3}$ (Van der Klis, 2006). This frequency thus falls off quickly as a function of radius, which is in agreement with the behaviour of the Type-C QPO frequency in the truncated disk model: a softer spectrum, corresponding to a smaller truncation radius, shows higher frequencies.

The connection between Lense-Thirring precession and Type-C QPOs was first made in the Relativistic Precession Model (RPM) by Stella & Vietri (1998) and Stella et al. (1999). However, in this model, the QPO is associated with precession at the inner edge of the disk extending to the innermost stable circular orbit. As I discussed, the QPO is known to be stronger in the harder X-ray emission from the corona, instead of the disk. Secondly, the Lense-Thirring precession frequency of a test particle greatly depends on the spin of the black hole. If the truncation radius of the disk is set by the properties of the accretion flow (for instance mass accretion rate),

⁵ Defined as the black hole angular momentum over the black hole mass, i.e. J/M . The maximum value of this parameter is 1, for a maximally spinning black hole and natural units.

Figure 8.: Both the Lense-Thirring precession frequency ν_{LT} and the Keplerian frequency ν_{ϕ} as a function of radius in gravitational radii. The black, red, green, blue and magenta lines correspond to different black hole spins: $a = 0.3, 0.5, 0.7, 0.9$ and 0.998 , respectively. The frequencies are plotted up until the innermost stable circular orbit for the given black hole spin. The dashed lines indicate the observed range of Type-C QPO frequencies. Figure from [Ingram et al. \(2009\)](#)

black holes with different spin should show wildly different QPO frequencies. However, QPO frequencies show remarkably little variations between BHXRBS. Moreover, the predicted Lense-Thirring precession frequencies easily surpass the highest observed QPO frequencies, as is shown in Figure 8. Hence, the Lense-Thirring precession of a test particle cannot be the full story.

[Fragile et al. \(2007\)](#) simulate a thick accretion disk in a system where the disk spin and black hole spin are misaligned. They do not find the creation of a global warp, where the inner disk spin aligns with the black hole spin, as predicted by [Bardeen & Petterson \(1975\)](#). Instead, the disk shows a global precession. This simulation suggests that Lense-Thirring precession could be invoked as an explanation of the Type-C QPO, by considering precession of the inner flow instead of the inner radius of the thin disk.

This notion of a precessing inner flow was developed in detail by [Ingram et al. \(2009\)](#), [Ingram & Done \(2011\)](#) and [Ingram & Done \(2012a\)](#). Figure 9, adapted from [Ingram et al. \(2009\)](#), depicts the considered geometry: a hot, geometrically thick inner flow, with its spin axis misaligned with respect to the black hole spin, is located within the truncation radius of the disk. As

Figure 9.: Schematic depiction of the geometry considered in the solid-body precession model by [Ingram et al. \(2009\)](#). A hot, geometrically thick inner flow, with its spin axis misaligned with respect to the black hole spin, is located within the truncation radius of the disk. Figure from [Ingram et al. \(2009\)](#)

the flow is geometrically thick, warps are communicated through pressure waves. The relevant communication timescale is thus the sound crossing timescale, which is shorter than the precession timescale. As a result, the entire inner flow precesses in a solid-body like fashion. The frequency of this precession is set by the surface density-averaged Lense-Thirring precession frequencies. As [Ingram et al. \(2009\)](#) show, this predicted frequency does not exceed the measured Type-C QPO frequency and does not depend on black hole spin as significantly as the test-particle orbit. As this thesis focusses on observational analysis of the QPO, I refer the reader to [Ingram et al. \(2009\)](#), [Ingram & Done \(2011\)](#) and [Ingram & Done \(2012a\)](#) for a thorough description of the theoretical model. In the following paragraphs, I will instead compare this model to the basic phenomenological properties of the Type-C QPO.

The Lense-Thirring precession model is able to reproduce most basic phenomenological properties of the Type-C QPO. As stated, the frequency range matches the observed frequencies for a precessing hot, geometrically thick flow (this fundamental requirement is actually met by *all* QPO models). In addition, [Motta et al. \(2014\)](#) show that the correlation between the frequencies of the Type-C QPO and the two simultaneously observed HFQPOs in GRO J1655-40 is as expected for a precession inner flow, where the HFQPOs correspond to the periastron precession and orbital frequencies. The evolution of QPO frequency during outbursts and dominance of hard photons in the QPO also naturally follows in this model, as the QPO originates from the hot inner flow. The same holds for the increasing amplitude of the QPO as a function of energy. Furthermore, [Ingram & Done \(2011\)](#) connect the QPO origin to the BBN through conservation of mass in the accretion flow, thus accounting for the remaining power spectral components.

An explanation for the energy dependence of the Type-C QPO frequency (an increasing trend in GRS 1915+105, XTE J1550-564 and H1743-322), is less evident. If the inner flow precesses as a solid body, only a single frequency is expected in all energy bands. This effect could be accounted for through selection effects in the surface density weighing of the precession frequency: when the surface density is relatively high at large radius, the emission is expected to be softer while the precession frequency is low. Alternatively, when this surface density enhancement has travelled closer to the black hole, the emission turns harder and the frequency increases. Thus, even

though the frequency is the same at all energies at any given time, averaging together large time intervals will create an apparent energy dependence in the QPO frequency. However, this cannot account for the decrease of QPO frequency with energy, observed in GRS 1915+105 as well. Chapter 2 of this thesis focuses completely on this problem and investigates whether the required hardness-QPO frequency correlations are present.

Secondly, the behaviour of QPO phase lags, as a function of energy and QPO frequency, is difficult to directly explain in the Lense-Thirring precession model. The model does not actually incorrectly predict the lag behaviour - the QPO lags are simply not fully understood yet. Recent efforts using phase-resolved spectroscopy of the QPO (Ingram & Van der Klis, 2015; Ingram et al., 2016; Stevens & Uttley, 2016) have provided new suggestions for the origins of QPO phase lags, related to changes in the spectral power law component during QPO cycles. In Chapter 3, I investigate the inclination dependence of these QPO phase lags as a handle on the underlying QPO model in general and on the origin of QPO lags specifically.

The presence of (sub)harmonic features in the QPO power spectra does not appear to follow naturally from the Lense-Thirring precession model. However, recent results obtained through phase-resolving spectral QPO properties (Ingram & Van der Klis, 2015; Ingram et al., 2016) might provide further insight into the harmonic. Phase-resolving the QPO reveals that, in addition to the flux, several spectral features, including the reflective component, show strong harmonic content. The presence of a reflective harmonic suggests that the harmonic might result from the simultaneous illumination of disk by both the top and under side of the precessing inner flow. However, this explanation still leaves the subharmonic unexplained. While I will discuss the nature of the Type-C harmonic in Section 3.3, the origin of the subharmonic is beyond the scope of this thesis.

As can be seen in Table 1, the basic properties of the Type-B QPO differ significantly from the Type-C QPO. Hence, the Lense-Thirring precession can most likely not be applied directly to the Type-B QPO. A simultaneous observation of a Type-B and C QPO in BHXR B GRO J1655-40 (Motta et al., 2012) also suggests a different origin of the Type-B QPO. Several authors have argued that the Type-B QPO is linked to the presence of a relativistic jet (Fender et al., 2009; Stevens & Uttley, 2016). Regardless of the precise underlying mechanism, the Type-B QPO origin appears to be geometric as well (Motta et al., 2015), as we'll briefly discuss in the next section. I will also discuss this in Chapter 3.

1.2.5 Recent observational evidence of Lense-Thirring precession

In the previous section, I have stated that geometric QPO models in general, and Lense-Thirring precession specifically, are currently the most promising models for Type-C QPOs. I have merely discussed these models, without referring to direct observational evidence obtained *after* formulating them - observational results that the model was not designed to match, as they were not obtained yet. In this section, I will present an overview of such recent observational results in the understanding of Type-C QPOs. I will first discuss results related to the question whether the QPO is a geometric or intrinsic effect. Afterwards, I will discuss results connected to the Lense-Thirring precession model.

In order to distinguish a geometric from an intrinsic QPO model, the inclination dependence of QPO properties is a powerful tool. Schnittman

et al. (2006) predict that a QPO resulting from a vertically precessing ring should show a clear inclination dependence in amplitude: edge-on sources are expected to show a systematically stronger QPO than face-on sources. However, the search for inclination dependence in any timing property of BHXRBS is difficult due to intrinsic differences between the systems - for instance the misalignment between the black hole spin and binary orbit, an essential ingredient in the Lense-Thirring precession model, can differ wildly between BHXRBS. Hence, a large source sample with adequate, independent inclination estimates is required.

Motta et al. (2015) search for the inclination dependence of QPO and BBN amplitudes in a sample of six edge-on, six face-on and two undetermined-inclination BHXRBS. They find that the amplitude of the Type-C QPO is systematically larger in edge-on sources. For Type-B QPOs, the result is opposite: face-on sources show larger QPO amplitudes. In both cases, this dependence on inclination implies a geometric underlying effect. However, the opposite behaviour of Type-B and Type-C sources implies different origins: while the inclination dependence of the Type-C QPOs matches the expectation for a vertically precessing flow formulated by Schnittman et al. (2006), the Type-B QPOs are more consistent with a jet-based origin. These results are thus consistent with the simultaneous observation of a Type-B and a Type-C QPO in GRO J1655-40 by Motta et al. (2012).

Using a completely different approach, Heil et al. (2015b) also find an inclination dependence in the Type-C QPO. Using a power-colour analysis (see Heil et al., 2015a, for an introduction to power-colours), they find differences in timing properties between edge-on and face-on sources as these transition from the hard to soft state. They show that these differences can mainly be attributed to the type-C QPO, providing further evidence that this QPO arises from a geometric effect.

The smoking gun for the Lense-Thirring precession model is the modulation of the iron line throughout each QPO cycle (Ingram & Done, 2012b): during each QPO cycle, a precessing inner flow illuminates consecutively the approaching and receding side of the disk, resulting in alternating blue- and redshifting of the reprocessed and reflected emission. Searching for such shifts is easiest in the most prominent reflection feature: the reflected iron $K\alpha$ line. Coupled with relativistic effects, the red- and blueshifts cause a characteristic modulation of the iron line profile. Hints of this effect were first observed in GRS 1915+105 by Ingram & Van der Klis (2015). Using phase-resolved spectroscopy, the equivalent width of the iron line was shown to change with QPO phase. Recently, Ingram et al. (2016) reported a modulation of the iron line centroid energy in H1743-322 as a function of QPO phase. These results constitute a strong indication that the Type-C QPO indeed originates from Lense-Thirring precession.

In this thesis, I will discuss new results on both the QPO inclination dependence and the Lense-Thirring precession model. In Chapter 2, I present the first observational evidence of *differential* Lense-Thirring precession in the accretion flow, where different radii precess at different rates. In Chapter 3, I report the discovery of an inclination dependence in Type-C QPO phase lags.

1.3 THESIS OUTLINE

In this thesis, I present the results of two distinct, but complementary, observational analyses aiming to better constrain the physical origin of QPOs

in BHXRBs. I aim to contribute observationally to the answer to the most fundamental question regarding QPOs, introduced in detail in this Chapter:

What is the origin of low-frequency QPOs in black hole X-ray binaries?

In this thesis, I will tackle this incredibly broad question by investigating two more-constrained, but ill-understood, properties of the QPO.

In Chapter 2, I perform an in-depth investigation of the energy dependence of the Type-C QPO frequency. As outlined in this introduction, this energy dependence is expected to be the result of the surface-density weighting of the solid-body precession frequency of the inner flow. Unexpectedly, I find that the Type-C QPO frequency is intrinsically different in different energy bands. I interpret this result as being due to *differential precession*, where different radii in the inner flow precess at different rates. In this second Chapter, I present a toy-model description of this interpretation and discuss the effect of the spectral shape on the observed energy dependence. This Chapter is adapted from [Van den Eijnden et al. \(2016\)](#). In Chapter 3, I turn to a sample of BHXRBs. In fifteen sources, I measure the dependence of QPO phase lag on QPO frequency, both for Type-B and Type-C QPOs. I find that the Type-C QPO phase lag significantly depends on the inclination of the binary orbit, both in relation to the QPO frequency and, most importantly, sign. This result constitutes arguably the most evident inclination dependence of a QPO property detected so far. The inclination dependence adds to the growing collection of observational evidence that the QPO has a geometric origin. I interpret the inclination dependence in the light of different dominant relativistic effects on the observational appearance of a QPO caused by Lense-Thirring precession. This Chapter has been submitted for publication in MNRAS.

Finally, in Chapter 4, I present preliminary results of a continued analysis of differential precession in GRS 1915+105. I investigate the evolution of spectral and lag properties of the BHXRB during coherent intervals, wherein the precession appears to set up and decay. After the completion of this thesis, I plan on extending this analysis and publishing these results as well. As each Chapter is either published or aimed for publication, all are written in plural form, except for the adapted introduction to each Chapter, which is written in the first person.

THE SHORT-TIMESCALE EVOLUTION OF QPO PHASE LAGS IN GRS 1915+105

Abstract

We present a model-independent analysis of the short-timescale energy dependence of low QPOs in the X-ray flux of GRS 1915+105. The QPO frequency in this source has previously been observed to depend on photon energy, with the frequency increasing with energy for observations with a high (> 2 Hz) QPO frequency, and decreasing with energy for observations with a low (< 2 Hz) QPO frequency. As this observed energy dependence is currently unexplained, we investigate if it is intrinsic to the QPO mechanism by tracking phase lags on (sub)second timescales. We find that the phase lag between two broad energy bands systematically increases for 5 - 10 QPO cycles, after which the QPO becomes decoherent, the phase lag resets and the pattern repeats. This shows that the band with the higher QPO frequency is running away from the other band on short timescales, providing strong evidence that the energy dependence of the QPO frequency is intrinsic. We also find that the faster the QPO decoheres, the faster the phase lag increases, suggesting that the intrinsic frequency difference contributes to the decoherence of the QPO. We interpret our results within a simple geometric QPO model, where different radii in the inner accretion flow experience Lense-Thirring precession at different frequencies, causing the decoherence of the oscillation. By varying the spectral shape of the inner accretion flow as a function of radius, we are able to qualitatively explain the energy-dependent behaviour of both QPO frequency and phase lag.

In the first Chapter, I introduced the energy dependence of the Type-C QPO frequency in GRS 1915+105 as reported by [Qu et al. \(2010\)](#) and [Pahari et al. \(2013\)](#): when the QPO frequency in the full, 2–13 keV energy band is higher than 2 Hz (*‘high frequency’*), the QPO frequency increases with energy. However, when the QPO frequency in this full energy band is lower than 2 Hz (*‘low frequency’*), the QPO frequency decreases with energy. At the dividing QPO frequency of approximately 2 Hz, no energy dependence is observed in the QPO frequency. Interestingly, the QPO phase lags appear to show the opposite effect: whenever the QPO frequency increases with energy, the QPO lags decrease with energy, and vice versa. Again, no energy dependence is observed at 2 Hz. Lastly, the QPO broad-band lag is found to depend apparently log-linearly on the QPO frequency. This relation crosses zero at a full-band QPO frequency of, again, 2 Hz. So what is causing these energy dependences? And what is the role of the QPO frequency of 2 Hz?

These recent results on the energy dependence of both the QPO frequency and phase lags pose several challenges for current QPO models: not all models predict an energy-dependent QPO frequency, and none can account for the decrease of frequency with energy observed in GRS 1915+105 ([Qu et al., 2010](#)). As stated in the introductory Chapter, the energy dependence of the

QPO frequency could be explained through a correlation between hardness and QPO frequency. In that interpretation, both the hard and soft band always show the same frequency jitter, but the changing hardness weights the QPO frequency differently in different energy bands. This would lead to observed differences in QPO frequency between different energy bands, even though there is only a single underlying QPO frequency at any time.

In this Chapter, adapted from [Van den Eijnden et al. \(2016\)](#), I test this hypothesis that the observed energy dependence of the QPO frequency arises due to hardness-frequency correlations in the QPO lightcurve. I have developed a novel, model-independent approach to investigate properties of the QPO, such as hardness, frequency and phase lag, on the timescale of single QPO cycles, by removing non-QPO variability from the observed lightcurves. This allows us to test for biases causing the observed energy dependence, as explained above, by tracking QPO frequencies and hardness on short timescales.

Unexpectedly, I find strong evidence that the observed frequency differences are a genuine property of the underlying QPO mechanism. I also find that the phase lag at the QPO frequency increases systemically on the timescale of 5 – 10 QPO cycles as a result of this frequency difference, before resetting once the QPO has become decoherent. I interpret our results in a geometric toy model where the innermost accretion flow is subject to differential precession. By varying the shape of the emitted X-ray spectrum as a function of radius, I am able to qualitatively explain the observed energy dependencies of the QPO properties.

2.1 OBSERVATIONS AND TIMING ANALYSIS

In this Chapter, we consider two plausible origins for the energy dependence of the QPO frequency, which are depicted schematically in Figure 10. In the left scenario, the QPO lightcurves in the two energy bands always have the same frequency. However, this frequency changes as a function of time. Whenever the frequency is high, the hard band lightcurve has a large amplitude compared to the soft band lightcurve (where amplitude refers to the maximum deviation from the mean and not the rms amplitude). When the frequency is low, the amplitudes are reversed, i.e. the amplitude is higher in the soft band. In this scenario, the power spectra in the two energy bands would show a QPO frequency weighted towards the high amplitude segments of the lightcurve. Thus the power spectra would show a different QPO frequency, even though the frequencies are always the same. In the alternative scenario, on the right, the QPO lightcurves in the two energy bands simply possess a different frequency. While this might seem to be a simpler explanation of the observed energy-dependence of the QPO frequency, the former scenario is more consistent with current models as there is only a single QPO frequency. Furthermore, in the latter scenario, the different QPO frequencies would cause a runaway between different energy bands over long timescales, which is contradicted by the coherent nature of the QPO.

There are two tests to distinguish between these two possible explanations: first, in the left scenario in Figure 1, the amplitude of the QPO lightcurve should be either correlated with the frequency in the hard band and anticorrelated with frequency in the soft band, or vice versa. In the right scenario, such (anti)correlations are not necessarily expected. Secondly, if the QPO frequency in both energy bands is always the same, the phase lag

Figure 10.: Cartoon depiction of the two considered scenarios for the observed energy dependence of the QPO frequency. In the scenario on the left, the frequencies in both energy bands are always the same, while in the scenario on the right, the frequencies are different. For clarity, the differences in frequency are exaggerated compared to actual observed frequency differences (listed in Table 2).

is expected to stay constant. However, if both energy bands possess a different QPO frequency, this phase lag would systematically change over time. As we know that the QPO is coherent on long timescales, these changes in phase lag would occur only on very short timescales.

We have developed a model-independent method to search both for correlations between QPO frequency and amplitude, and for short-timescale variations in phase lag. The method consists of broadly four steps: we (1) extract light curves in two broad energy bands and calculate their power spectra, (2) filter these light curves in order to conserve only the QPOs, (3) determine the frequency and amplitude of each QPO cycle, and (4) track the phase lag between the energy bands on the timescale of individual QPO cycles. Steps (1), (2) and (3) are described in section 2.1.1, while step (4) is described in section 2.1.2. Our method applies multiple standard spectral-timing techniques, which are explained in more detail in Appendix A.

Table 2.: Overview of analysed *RXTE* observations. Listed are the ObsID, observation date, fundamental QPO frequency in the reference band ν_0^{ref} , the difference in QPO frequency (hard - soft) $\Delta\nu_0$, the Q-factor in the full energy band ($\nu_{\text{QPO}}/\text{FWHM}$), the used energy bands (I: 1.94 – 6.89 and 6.89 – 12.99 keV; II: 1.94 – 6.54 and 6.54 – 12.99 keV; III: 2.13 – 6.72 and 6.72 – 12.63 keV), the reduced χ^2 of the power spectral fit $\chi^2/\text{d.o.f.}$, and the slope and offset of the phase lag evolution (see Section 2.2.2).

ObsID	Date	ν_0^{ref} (Hz)	$\Delta\nu$ (Hz)	Q-factor	Energy Bands	$\chi^2/\text{d.o.f.}$	Slope	Offset
10258-01-06-00a	29-08-1996	5.141 ± 0.024	0.078 ± 0.036	5.9	I	1.02	−0.20	−0.06
10408-01-21-02	07-07-1996	8.107 ± 0.043	0.313 ± 0.071	6.2	I	1.11	−0.40	0.01
10408-01-22-00	11-07-1996	3.480 ± 0.004	0.008 ± 0.008	8.5	I	1.22	−0.08	−0.06
10408-01-22-01	11-07-1996	2.777 ± 0.005	0.003 ± 0.007	6.8	I	1.24	−0.03	−0.03
10408-01-22-02	11-07-1996	2.560 ± 0.004	0.000 ± 0.007	6.3	I	1.36	−0.04	−0.02
10408-01-27-00	26-07-1996	0.632 ± 0.002	$−0.003 \pm 0.002$	5.3	I	1.08	0.08	0.13
10408-01-28-00	03-08-1996	0.966 ± 0.002	$−0.005 \pm 0.003$	5.3	I	1.34	0.04	0.10
10408-01-29-00a	10-08-1996	1.658 ± 0.004	$−0.004 \pm 0.006$	10.8	I	1.03	−0.01	0.06
10408-01-29-00b	10-08-1996	1.856 ± 0.004	0.006 ± 0.005	9.9	I	1.18	−0.07	0.07
10408-01-29-00c	10-08-1996	1.963 ± 0.004	$−0.010 \pm 0.006$	6.8	I	1.40	−0.03	0.02
10408-01-30-00	18-08-1996	4.944 ± 0.010	0.041 ± 0.017	3.0	I	1.07	−0.05	−0.11
10408-01-31-00a	25-08-1996	4.084 ± 0.008	0.013 ± 0.01	9.5	I	0.96	−0.05	−0.09
10408-01-31-00b	25-08-1996	4.439 ± 0.009	0.04 ± 0.014	5.0	I	1.17	−0.11	−0.07
10408-01-31-00c	25-08-1996	3.514 ± 0.006	0.008 ± 0.009	6.7	I	1.25	−0.0	−0.08
10408-01-32-00	31-08-1996	6.121 ± 0.018	0.296 ± 0.029	3.9	I	1.38	−0.39	−0.02
20402-01-48-00	29-09-1997	7.639 ± 0.034	0.242 ± 0.054	5.9	II	1.35	−0.43	−0.02
20402-01-50-01	16-10-1997	1.042 ± 0.003	$−0.005 \pm 0.004$	6.1	II	1.11	0.03	0.10
30182-01-01-00	08-07-1998	1.870 ± 0.009	0.007 ± 0.013	8.1	II	1.08	−0.08	0.04
30402-01-11-00a	20-04-1998	5.245 ± 0.032	0.127 ± 0.033	7.3	II	1.09	−0.33	−0.01
30402-01-11-00b	20-04-1998	5.857 ± 0.017	0.129 ± 0.039	3.6	II	1.23	−0.24	−0.06
30703-01-20-00	24-05-1998	0.696 ± 0.002	$−0.004 \pm 0.003$	5.3	II	1.01	0.03	0.14
30703-01-35-00	25-09-1998	2.464 ± 0.006	0.008 ± 0.009	5.9	II	1.14	−0.03	−0.05
40703-01-38-01	15-11-1999	7.114 ± 0.030	0.373 ± 0.044	4.1	III	1.09	−0.57	0.09
40703-01-38-02	15-11-1999	7.943 ± 0.032	0.282 ± 0.048	7.2	III	1.27	−0.89	0.22

2.1.1 Data reduction and optimal filtering

For our analysis, we select 24 RXTE PCA observations of GRS 1915+105, based on the observations discussed in [Qu et al. \(2010\)](#), [Pahari et al. \(2013\)](#) and [Yan et al. \(2013\)](#). The observations are selected to evenly span a range in QPO frequency from ~ 0.5 to 8 Hz. Table 2 summarizes the main properties of these observations. Using the standard FTOOLS¹ package, we extract binned mode data to produce light curves in three energy bands: a soft band from ~ 2 to ~ 6.7 keV, a hard band from ~ 6.7 to ~ 13 keV, and a reference band covering both energy ranges. Due to changes in the PCA gain, the exact energy bands differ slightly between observations. The exact energy bands are indicated in Table 2 for all observations. We extract all observations using a 1/128 s time resolution, which yields a Nyquist frequency of 64 Hz for the subsequent analysis.

We divide all lightcurves into 8 second segments and for each one calculate the power spectral density (PSD) with a 1/8 Hz resolution. After applying the rms-squared normalisation ([Belloni & Hasinger, 1990](#)) we average the separate power spectra into one PSD per lightcurve to reduce the standard errors. Using XSPEC v12², we fit the average power spectra with a model consisting of a constant white noise, two broad band noise (BBN) Lorentzians with a fixed centroid frequency of 0 Hz, and a Lorentzian for the QPO fundamental and each (sub)harmonic. We fit all energy bands within the same observation separately with the same model, as linking parameters between energy bands generally results in worse fits. Details of the resulting fits, including reduced χ^2 values and QPO frequencies, are listed in Table 2. The errors shown are the one sigma confidence intervals.

In order to study the behaviour of only the QPO, and remove broad band and Poisson noise contributions to the variability, we apply an optimal filtering technique based on the method in [Press et al. \(1997\)](#). We assume that the observed count rate $c(t)$ consists of the true QPO signal $q(t)$ and an added noise component $n(t)$:

$$c(t) = q(t) + n(t) \quad (2)$$

Our aim is to remove $n(t)$ in order to estimate $q(t)$ as accurately as possible. The optimal filter provides such an estimate of the true QPO signal, $\tilde{q}(t)$, by minimizing the squared difference between $q(t)$ and $\tilde{q}(t)$. In practice, the filter $F(\nu)$ is applied by multiplication with the Fourier transform of the count rate (capitalized variables indicate the Fourier transform):

$$\tilde{Q}(\nu) = F(\nu) \cdot C(\nu) \quad (3)$$

The filter in Fourier space is given by

$$F(\nu) = \frac{|Q(\nu)|^2}{|C(\nu)|^2} \quad (4)$$

and thus requires an estimate of the actual QPO power spectrum $|Q(\nu)|^2$, for which we apply the fitted QPO Lorentzian. The remaining time series $\tilde{q}(t)$ estimates the true QPO lightcurve, without other variability contributions.

The optimal filter assumes that the QPO signal $q(t)$ and the noise contribution $n(t)$ are uncorrelated. In our method, the noise consists of the white noise, the BBN and any (sub)harmonic QPO peaks. Since the (sub)harmonics

¹ <https://heasarc.gsfc.nasa.gov/ftools/ftools.menu.html>

² <https://heasarc.gsfc.nasa.gov/xanadu/xspec/>

Figure 11.: Explanatory example of the optimal filter in the frequency domain. The black points correspond to the observed power spectrum, the red and blue stars to the filtered power spectrum. We only use the inner (red) part of the filtered power spectrum, within the range shown by the dotted lines ($\nu_{\text{QPO}} \pm \text{FWHM}$), to produce QPO light curves.

are clearly related to the fundamental QPO and correlations between the BBN and the QPO are known to exist (Heil et al., 2011), the assumption of uncorrelated noise does not fully hold. This will especially spoil the filter at low frequencies, where the BBN is dominant, and at the (sub)harmonic frequencies. To cancel these effects, we apply an extra cut that removes all high and low frequencies outside the range $\nu_{\text{QPO}} \pm \text{FWHM}$, where ν_{QPO} is the fitted QPO frequency in the considered energy band. As this does not remove the correlations at the QPO frequency, the filter remains slightly less than optimal. Alternative filters, that do not make assumptions about noise correlations, exist: for example, the tophat filter simply removes all power outside a certain frequency range. These filters are less accurate than the optimal filter and do not use any known properties of the QPO peak. For this reason, we apply the optimal filter for the subsequent analysis. However, our main results, presented in the next section, do not differ significantly when using the tophat filter.

Both the optimal and alternative filters only affect the amplitude of the power spectrum, while leaving the phases unaltered. This implies that we can use the filtered lightcurves to measure phase lags in the subsequent analysis. However, this also means that while the BBN amplitude is removed, its phase lags are still present in the filtered lightcurve. This requires us to make the assumption that, at ν_{QPO} , the phase lags are dominated by the QPO and the contribution of the BBN is negligible. We will discuss the effects of our choice of filter and the validity of this assumption in section 2.3.1.

Figure 12.: Example of the optimal filter in the time domain. The black and red curves are the light curves corresponding to respectively the black and red power spectra in Figure 11. The filtered light curve (red) clearly picks out the QPO, while removing other variability present in the observation.

Figure 11 shows an example of an unfiltered and filtered power spectrum. The peak of the QPO is clearly sampled by the filtered power spectrum, while the power becomes zero outside the allowed frequency range. We use 64 second segments of the light curves to produce the power spectra that are filtered, causing the difference in frequency resolution in Figure 11. It is possible to select longer segments since the power spectra are already fitted, so there is no need to average many power spectra to reduce standard errors. Figure 12 shows the light curves corresponding to the power spectra in Figure 11 in the same colours. As intended, the filtered light curve tracks the large overall oscillations, but does not sample the added noise contributions.

In order to track each QPO cycle individually, we use simple linear interpolation to estimate the mean-crossings and extrema of the filtered reference band light curves. Defining a QPO cycle as a light curve segment including either three consecutive mean-crossings or two consecutive maxima, we can determine both the maximum amplitude and frequency of each individual cycle. This allows us to perform the first test of the energy-dependence of the QPO frequency: the aforementioned presence of (anti)correlations between amplitude and frequency.

2.1.2 Phase lags

The second method to test the energy-dependence of the QPO frequency is to track the phase lag over time: if the different energy bands harbor a different QPO frequency, this phase lag should change systematically. But as was already stated, these changes should occur on short timescales only,

Figure 13.: Categorizing method for the position of segments in the filtered lightcurve. The dotted lines indicate the edges of each segment. Segments containing the start of a coherent interval, equivalent to a minimum QPO amplitude, are indicated by a grey band. Segments are grouped together based on either absolute label (upper) or ratio of label to full coherent interval length (lower).

since the QPO is coherent on longer timescales. Hence, we first need to identify the timescale on which to search for variations in phase lag.

As is visible in Figure 12, the filtered lightcurves show that the QPO amplitude repeatedly rises and falls in an enveloping modulation. While the individual envelopes differ in length, the mean length of the envelopes in an observed lightcurve is tightly correlated with the Q-factor ($\nu_{\text{QPO}}/\text{FWHM}$) in that lightcurve. In other words, the more coherent the QPO, the longer the envelopes in the filtered QPO lightcurve. Hence, we name these envelopes *coherent intervals* and select these as the timescales within which we search for the changes in phase lag that might signal a genuine energy dependence. We designed our subsequent analysis such that it allows us to track the phase lag throughout these coherent intervals. As such, it consists of three independent steps: (1) splitting each observation in segments, (2) identifying the positions of these segments in coherent intervals, and (3) calculating the phase lag by combining lightcurve segments located at similar positions in their respective coherent interval.

First, we divide the time-axis of each observation into short segments of length $1/\nu_0^{\text{ref}}$, where ν_0^{ref} is the QPO frequency in the full 2 – 13 keV band (hereafter referred to as the *full-band QPO frequency*). Using the optimally-filtered reference band lightcurve, we search for minima in QPO amplitude and identify these as the start and end points of coherent intervals. For the second step, we label each segment with a number based on its position inside its respective coherent interval, as illustrated in Figure 13: starting with zero for the segment containing the start of the coherent interval, we increase the label of each successive segment by one, up until the end of the

coherent interval. We define the *absolute position* as the label of a segment, and the *fractional position* as the ratio of the label to the total amount of segments in the coherent interval. These positions are shown in Figure 13 as the upper and lower numbers, respectively. As we only used the filtered QPO amplitudes in this process of defining segment labels, this first step is completely independent of the phase or frequency properties of the QPO.

As the final step, we filter the hard and the soft band using their respective best fitting parameters. We normalise both lightcurves in each segment by subtracting their mean and dividing by their standard deviation, and cross correlate each set of simultaneous segments in the two energy bands. The position of the peak in the resulting cross correlation function (CCF) indicates the time lag between the two energy bands in that specific segment. In order to increase the signal-to-noise ratio in the CCFs, we average CCFs with either the same absolute or similar relative positions. In order to account for the uncertainty in our simple determination of the starts of the coherent intervals, we do not consider the zero-labeled segments in the subsequent analysis.

The time lag $\tau(\nu_0^{\text{ref}})$ between the hard and soft band can then be measured as the location of the peak of the averaged CCF. Due to the extra cuts in the optimal filter and the frequency difference between the two energy bands, no analytical description of this CCF exists. Therefore, we fit a skewed Gaussian model to the central peak of the averaged CCF. As the maximum of the skewed Gaussian model is not represented by a single parameter, we then apply bootstrapping to the averaging of the CCFs to determine the one-sigma uncertainty on the time lag (see e.g. [Press et al., 1997](#), section 15.6). Finally, we convert the time lags into phase lags $\phi(\nu_0^{\text{ref}})$ using $\phi(\nu_0^{\text{ref}}) = \tau(\nu_0^{\text{ref}}) \cdot 2\pi\nu_0^{\text{ref}}$. Thus, we obtain an estimate of the average phase lag between the hard and soft band at different absolute or fractional positions in the coherent intervals.

2.2 RESULTS

2.2.1 Test 1: frequency-amplitude correlations

In Figure 14, we show a representative example of the relation between the amplitude and frequency of single QPO cycles. The panels show a two-dimensional histogram of the amplitude and frequency of individual QPO cycles, rescaled such that the highest counts are equal to one. The left and right panel show the relation for the soft and hard band, respectively. Two immediate conclusions can be drawn from the Figure: first, both the soft and hard band show no apparent (anti)correlation between frequency and amplitude. Secondly, the panels of the energy bands appear strikingly similar, even though the center of the distribution lies at a different QPO frequency. This lack of correlation and similarity between the energy bands is present in all observations analysed. Thus, this first tests unexpectedly points towards different QPO frequencies in the different energy bands, i.e. a genuine energy frequency dependence of the QPO frequency.

Figure 14.: Representative example of the relation between amplitude and frequency of individual QPO cycles. The panels show a two-dimensional histogram of these amplitudes and frequencies, rescaled such that the highest counts are equal to one. The left and right panel show the relation for the soft and hard band, respectively.

As a test of our method, we converted the observed power spectra into maximally stochastic light curves using the method of [Timmer & Koenig \(1995\)](#), i.e. the best fitting power spectra were used to generate simulated lightcurves with the properties of Gaussian noise. After applying the optimal filtering and the tracking of QPO properties, these simulated light curves show a completely similar relation between amplitude and frequency as the observed ones shown in Figure 14. This similarity indicates that the triangular shapes in the amplitude-frequency plane are simply an inherent property of a noise process with the observed power spectra. The triangular shape itself is likely due to fact that at lower amplitudes, the determination of individual QPO cycles becomes less precise due to noise contributions.

2.2.2 Test 2: phase lag evolution

As a second test, we search for changes in the phase lag on short timescales. In panel 1 of Figure 15, we show a representative example of the evolution of the phase lag ϕ as a function of fractional position within a coherent interval, with $\phi > 0$ corresponding to hard lags. The phase lag clearly increases from no lag at the start of the coherent intervals (fractional position of 0.0) to a significant soft lag at the end (fractional position of 1.0). This shows that the hard band oscillation is systematically faster than, and hence ‘runs away from’, the soft band within a coherent interval, and thus that the two lightcurve possess a different QPO frequency. Furthermore, the phase lag resets at the start of each coherent interval, causing the QPO to stay coherent on longer timescales. This reset is already visible as the turnover around a cycle fraction of 0.9, which is present in almost all observations. We will discuss the presence and interpretation of this turnover in more detail in the next section.

In order to test our method for systematic effects and biases, we repeat the analysis, but randomize the measured start times of the coherent intervals. In other words, the assigned label of each segment is randomized, and does not represent the actual segment position. We show the result of this check in panel 3 of Figure 15. All phase lags are consistent with the grey band, which represents the average phase lag, independently calculated by averaging all CCFs without regard for their positions. This consistency indicates that the increase in panel 1 is inherently linked to the existence of and position within the coherent intervals.

Panels 2 and 4 show the phase lag behaviour as a function of absolute position. In order to compare phase lag behaviour between coherent intervals of different lengths, we apply an extra selection before averaging the CCFs: segments should not only have the same absolute position, but also be located in coherent intervals of the same total length. Interestingly, as can be seen in panel 2, shorter coherent intervals systematically show a steeper phase lag increase, but typically reach the same maximum lag before resetting. This suggests an underlying relation between the origin of the coherent intervals and the observed phase lag behaviour, that we will discuss in section 2.3. Similar to panel 3, panel 4 shows the results of randomizing the assigned labels. Again, the observed effect disappears and only the average phase lag remains.

Figure 15.: Evolution of the phase lag between the hard and soft QPO light curve throughout a coherent interval for a ~ 6.1 Hz QPO ($\Delta\nu_0 \sim 0.3$ Hz). Panels 1 and 2 show the phase lags versus fractional and absolute position in a coherent interval, respectively. The lag evidently increases and resets at the end. Panels 3 and 4 show the corresponding randomized checks of our method, and are expected to hover around the average phase lag, shown by the grey band. Positive values of ϕ indicate hard lags. See full text for details.

Figure 16.: Evolution of the phase lag between the hard and soft QPO light curve throughout a coherent interval for a ~ 0.6 Hz QPO ($\Delta\nu_0 \sim -0.003$ Hz). Same as Figure 15

In figure 16, we show the same relations for a lower frequency QPO, where the frequency difference is opposite (i.e. the QPO frequency decreases with energy). The general trends are similar to, but, as expected, in the opposite direction as those in Figure 15. The main difference is the presence of a non-zero phase lag at start of the coherent interval. The similarity of the phase lag behaviour at both high and low frequencies suggests that the runaway is a global characteristic of the observed QPO.

As a quantification of the increase in phase lag, we fit the phase lag as a function of fractional position with a straight line with non-zero offset:

$$\phi(x) = \alpha x + \beta \quad (5)$$

where x is the fractional position. We do not propose this as the correct empirical representation of the lag behaviour, but it allows us to characterize and compare all observations: the slope α quantifies the total change in phase lag in a full coherent interval and the offset β estimates the starting phase lag. In other words, the slope characterizes the runaway between the two energy bands. Figure 17 summarizes all observations by plotting the slope and offset of this straight line model against both full-band QPO frequency ν_0^{ref} and QPO frequency difference between the hard and soft band, $\Delta\nu_0$. The slope and offset corresponding to the first panel in Figures 15 and 16 are shown as the red triangle and blue square, respectively. Due to the relatively large uncertainties at the fractional position of 0.95, we find no significant difference between fitting all phase lags or disregarding the turnover.

Figure 17 shows that the runaway effect within coherent intervals is present in all observations with a non-zero frequency difference, and is small but significant for small values of $\Delta\nu_0$. The sign of the slope, indicating which energy band lags the other, is consistent within one sigma with the sign of $\Delta\nu_0$ in all observations. Panel 2 clearly shows that the slope increases as the frequency difference grows, as expected if the runaway is caused by a frequency difference. Panel 3 shows that the offset, i.e. the phase lag at the start of the coherent interval, follows a log-linear dependence on ν_0^{ref} , similar to the behaviour of the average phase lag in [Qu et al. \(2010\)](#), [Pahari et al. \(2013\)](#) and [Reig et al. \(2000\)](#), up to $\nu_0^{\text{ref}} \sim 3.5$ Hz. At higher frequencies, the offset diverges from the log-linear relation back to positive values. The observations with these diverging offsets correspond to the observations with large frequency differences in panel 4. Using a completely different approach, [Qu et al. \(2010\)](#) concluded that the frequency difference contributes only a small fraction to the average phase lag at low frequency but becomes dominant above $\nu_0^{\text{ref}} \sim 3.5$ Hz. In our analysis, this is visible in this divergence of β from a log-linear dependency on frequency above approximately the same value of 3.5 Hz.

Figure 17.: Fitted slope α and offset β of the phase lag evolution plotted against both QPO frequency ν_0^{ref} and frequency difference $\Delta\nu_0$. The vertical dotted line in panels 1 and 3 indicates $\nu_0^{\text{ref}} = 1.9$ Hz, where the average phase lag at the QPO frequency changes sign. The red triangle and blue square correspond to the slope and offset of the phase lag evolution in panel 1 of Figures 15 and 16, respectively. The inset in panel 1 shows a zoom in of the region below 5 Hz, showing a significant change from positive to negative slopes around 1.9 Hz.

2.3 DISCUSSION

2.3.1 Robustness of the method

We have tested our method for systematic biases at several important steps and have not found any significant effects. The choice of energy band to determine the start times of the coherent intervals has no effect on our results: our conclusions hold when using either the full, hard or soft band in this process. The same holds for the definition of a QPO cycle used while tracking the filtered lightcurves: we find no difference between defining a QPO cycle as two consecutive maxima or crossings of the mean photon rate. We find slight changes in our results if we replace the optimal filter by a tophat filter, which simply removes all variability outside a frequency range of $\nu_{\text{QPO}} \pm \text{FWHM}$. The phase lags do still increase in the coherent interval, but with a shallower slope. However, these changes can be expected since the tophat filter is less efficient than the optimal filter. This causes a decrease in the accuracy of our filtered QPO amplitude estimation and thus the determination of the coherent intervals. Finally, we find no significant changes in our results if we fit the CCFs with other functions, such as a damped sine wave, instead of a skewed Gaussian model. However, we apply the skewed Gaussian model to more accurately account for the asymmetry in the CCF when $\Delta\nu_0$ is significant.

The robustness of our method also follows from the effect of randomizing the coherent interval start times (panels 3 and 4 in Figures 15 and 16). In all observations, the phase lag evolution disappears and only the average phase lag remains. This shows that our method is capable of correctly estimating known phase lags. It also indicates that the observed phase lag evolution is inherently connected to the coherent intervals and cannot be attributed to biases in our method. Alternatively, if we distribute the start times evenly over the full lightcurve, the phase lag evolution disappears as well. This reinforces the notion that the evolution does not simply occur in any short time scale segment, but only in the actual coherent cycles.

In order to interpret our results as a property of the QPO, we make the assumption that the phase lag at ν_{QPO} is dominated by the QPO. For observations with a full-band QPO frequency $> 5\text{Hz}$, the lag spectrum shows a clear, distinct feature at the QPO frequency, while at other frequencies it appears featureless and hovers around zero. This implies that the phase lag is dominated by the QPO. At lower full-band QPO frequencies, these features in the lag spectrum become more diluted. Hence, we apply another test to our method: instead of calculating the CCF from optimally filtered lightcurves, we calculate it from lightcurves produced using a narrow tophat filter. If this narrow tophat filter is not centered on the QPO frequency, this method simply searches for phase lag evolution associated with the BBN. We perform this test several times per observation, sliding the tophat filter from low to high frequency. We find that the observed phase lag evolution is only present when applying the filter at the QPO frequency, and disappears otherwise. This indicates that the phase lag evolution is fundamentally linked to the QPO mechanism, and not to the BBN.

2.3.2 Phase lags

As previously shown by [Qu et al. \(2010\)](#), [Pahari et al. \(2013\)](#) and [Reig et al. \(2000\)](#), the average phase lag measured for an entire observation evolves

smoothly with the QPO frequency of that observation. We demonstrate this relation in Figure 18, where we plot average phase lag determined from the CCF versus QPO frequency for each analysed observation. Our results from the CCF confirm the aforementioned smooth evolution, which was previously only measured using the cross spectrum. To quantify this evolution, we fit the average phase lags with a simple log-linear model

$$\langle \phi \rangle = k \log(\nu_0^{\text{ref}}) + \phi_0 \quad (6)$$

which yields $\phi_0 = 0.096 \pm 0.001$ rad and $k = -0.348 \pm 0.003$ rad for ν_0^{ref} in Hz. These values correspond to a switch from hard to soft lags at a QPO frequency of approximately 1.9 Hz. Interestingly, this is consistent with the frequency where $\Delta\nu_0$ changes sign.

It seems remarkable that the phase lags show such complexity within each coherent interval, yet the average phase lag falls on such a neat relation. An apparent discrepancy is clearly visible in Figure 17. As long as the slope α is small ($\nu_0^{\text{ref}} < 3.5$ Hz), the offset β follows a log-linear relation with frequency. This is expected, since in the absence of any significant phase lag evolution, the offset will simply track the average phase difference. However, for higher QPO frequencies, where $\Delta\nu_0$ is large and the hard band quickly runs away from the soft band in each coherent interval ($\nu_0^{\text{ref}} > 3.5$ Hz, α becomes significantly less than 0), the offset β diverges from the log-linear relation. This behaviour of the offset allows the average phase lags to remain on a log-linear relation with QPO frequency, by compensating for the short timescale phase lag evolution.

To demonstrate this compensating behaviour, we plot the average phase lags measured from the CCF against those measured from the phase lag evolution (Equation 5) in Figure 19. The latter are simply given by $\langle \phi \rangle = \langle \alpha x + \beta \rangle = \alpha/2 + \beta$. All observations clearly scatter around the dotted line, which represents the diagonal. Even at the most negative phase lags, corresponding to the highest QPO frequencies and largest values of the slope α , the relation holds. The relatively large errorbars in the y-direction are caused by correlations between α and β in fitting the phase lag evolution, which are not large enough to explain the relation over the entire range of observations. Thus, this tight agreement over the entire range indicates that the slope α and offset β of each observation indeed compensate such that the average phase lag depends log-linearly on QPO frequency. This shows that the phase lag evolution simply pivots around the centre of each coherent interval. Of course, this does not explain why on long timescales, fundamentally, the average phase lag should follow a log-linear dependence on QPO frequency.

A possible way to explain the smooth log-linear relation between the frequency and phase lag on long timescales, could be by considering the dependence of phase lag on photon energy. [Pahari et al. \(2013\)](#) studied this energy dependence and found that the phase lag at the QPO frequency follows a log-linear dependence on photon energy. A similar relation is found by [Qu et al. \(2010\)](#), and hints of it are present in [Reig et al. \(2000\)](#). The exact slope of this relation varies between different observations, and several observations also show hints of a break around $\sim 4 - 6$ keV. Next to GRS 1915+105, [Wijnands et al. \(1999\)](#) showed log-linear relations between phase lag and energy in XTE J1550-564. The smooth relation between phase lag and QPO frequency might be equivalent to a simple rotation of the dependence of phase lag on energy, where the rotational angle depends on the frequency of the QPO. This scenario is encouraged by the presence of the

Figure 18.: The average phase lag $\langle \phi \rangle$ as a function of fundamental QPO frequency ν_0^{ref} . Positive values of $\langle \phi \rangle$ correspond to hard lags. The straight line shows a log-linear model fit of $\langle \phi \rangle = k \log \nu_0^{\text{ref}} + \phi_0$.

same phase lag-frequency relation in [Pahari et al. \(2013\)](#), [Qu et al. \(2010\)](#) and our results, despite differences between the compared energy bands, both in width of the bands and centroid energies. In section 2.3.4, we propose a possible mechanism that can explain these two observed dependencies of the phase lag and their possible equivalence based on a simple toy model.

2.3.3 Decoherence of the QPO

As part of our analysis, we have identified the envelopes in the QPO amplitude as so-called coherent intervals in all observations. These successive coherent intervals can be interpreted as subsequent, independent excitations of the QPO mechanism with a lifetime of approximately a coherence time. In this scenario, each single QPO excitation starts with the same initial phase lag between the hard and soft band. Subsequently, the frequency difference leads to the observed phase lag evolution as the excitation decays through the coherent interval. The observed turnover in the phase lag evolution could arise from the emergence of the next QPO excitation, which takes over from the previous one. In this interpretation, the frequency difference would play a vital role in the decay of the coherence through each interval. Here, we will discuss earlier evidence for separate QPO excitations and the relation between QPO frequency difference and (de)coherence. In section 2.3.4, we will discuss our results in the light of a simple physical toy model.

Using adaptive decomposition in a single ~ 4 Hz QPO observation of XTE J1550-564, [Su et al. \(2015\)](#) identify *intermittent oscillations* of ~ 3 seconds, where the QPO is coherent. Similar to our coherent intervals, these intermittent oscillations consist of high amplitude segments of the QPO lightcurve and are separated by short, lower amplitude segments. [Su et al. \(2015\)](#) interpret this behaviour within the Lense-Thirring precession model as the

Figure 19.: The measured average phase lag, $\langle \phi \rangle$, versus the predicted average phase lag based on the fitted phase lag evolution, $\alpha/2 + \beta$. The dotted line shows the diagonal. The tight correlation indicates that the latter is an adequate predictor of the average phase lag at all QPO frequencies.

alternation between steady precession (high amplitude) and unstable precession (low amplitude) of the inner hot flow. The ~ 3 second timescale is interpreted as the viscous timescale at the outer radius of the inner hot flow. BHXR B XTE J1550-564 shows energy-dependent QPO frequencies similar to GRS 1915+105 (Li et al., 2013a), suggesting that these intermittent oscillations are in fact the same as coherent intervals.

Lachowicz & Done (2010) expanded the standard Fourier methods using wavelet analysis and the Matching Pursuit algorithm to track the QPO properties over time. They are able to reconstruct the observed properties of the QPO using multiple short timescale QPO signals with random amplitudes. Given these existing results in XTE J1550-564, the similarity between intermittent oscillations and coherent intervals, and the simple explanation of the short timescale phase lag evolution, we adopt the interpretation of coherent intervals as a manifestation of distinct, independent excitations of the QPO mechanism.

Identifying coherent intervals as separate QPO excitations does not provide an explanation for the decoherence of the QPO through the interval. However, the systematic phase lag evolution might be accountable for that. Interestingly, we find that the speed of the phase lag evolution is linked to the length of the coherent interval: as is visible in panel 2 in Figures 15 and 16, shorter coherent intervals have a significantly quicker phase lag evolution. Moreover, for all coherence lengths, the phase lag reaches approximately the same maximum value for the same QPO frequency before turning over. Only when we account for the length of the coherent interval, are all phase lag evolutions mapped onto one relation (panel 1 in Figures 15 and 16). In other words: the length of the coherent interval appears to be

Figure 20.: The quality factor Q of the QPO in the hard band of each observation as a function of the relative QPO frequency difference $\Delta\nu_0/\nu_0^{\text{ref}}$.

set at the start of each interval and the phase lag evolves accordingly. This suggests a causal relation between the length of a coherent interval and the frequency difference in that particular coherent interval: a larger frequency difference leads more quickly to a large phase lag, which in turn reduces the coherence of the QPO and thus the length of the coherent interval. This is consistent with the interpretation of coherent intervals as independent QPO excitations, since in this case each individual excitation can have a distinct frequency difference. Even though the hard and soft band are not completely out of phase at the end of each coherent interval, the phase lag could still contribute to the decoherence.

A final property of the coherent intervals is not directly accounted for by the observed phase lag evolution: the QPO shows coherent intervals in all observations, even when no phase lag evolution is observed. In other words, the oscillations are always *quasi*-periodic, regardless of the presence of a frequency difference between energy bands. This can be seen clearly in Figure 20, where we plot the Q -value of the hard band as a function of relative frequency difference $\Delta\nu_0/\nu_0^{\text{ref}}$. In the Figure, no clear relation between $\Delta\nu_0$ and Q exists, contrary to what is expected if the frequency difference and resulting phase lag evolution causes the decoherence. In the next section, we offer a toy model based on our results that accounts for this last difficulty in explaining the decoherence through the phase lag evolution.

2.3.4 A unifying model: differential precession and spectral evolution

Any model to explain the decoherence of QPO excitations should account for both the sign of the frequency difference at low QPO frequencies and the link between phase lag evolution, QPO frequency difference and length

of the coherence interval. The former is especially challenging: phenomenologically, an increase of frequency with energy can be understood by considering different radii in the accretion flow. If the QPO frequency can be associated with characteristic timescales in the accretion flow, a higher frequency would correspond to a smaller radius. Simultaneously, from these smaller radii, a harder spectrum is expected to be emitted. In this simple picture, it is possible to account for a positive correlation between frequency and energy. However, this simple scenario is unable to explain the anti-correlation between frequency and energy that is observed at low QPO frequencies. This does not even touch upon the fact that the slope of the relation between energy and QPO frequency changes very systematically with ν_0^{ref} , which adds another layer of difficulty to this problem. Since we conclude that the frequency difference is intrinsic, any viable interpretation should relax either the assumption that a smaller radius corresponds to shorter timescales, or to a harder emission. Moreover, any model should also account for the observed properties of the phase lags on long timescales, both as a function of frequency and energy.

One possible interpretation of our results and the observed sign of $\Delta\nu_0$ incorporates spectral changes of the inner accretion flow as the full-band QPO frequency decreases. In this approach, we stick to the anticorrelation between oscillation frequency and radius and instead loosen the view that hardness tracks proximity to the black hole. As the QPO amplitude is known to depend on inclination, favoring a geometric QPO mechanism (Heil et al., 2015b; Motta et al., 2015), we assume a toy model where the QPO is caused by vertical precession of the inner flow (as in e.g. Ingram et al., 2009). In this toy model, we consider this inner flow to consist of two separate halves, which differentially precess at slightly different frequencies. In agreement with the expected dependence of frequency on radius, we assume that the inner half *always* causes a higher QPO frequency than the outer half: $\Delta\nu_0(\text{Inner} - \text{Outer}) > 0$. Dividing the inner flow into two halves is of course an extremely simplified picture. More realistically, the inner flow might show a more continuous evolution in precession frequency over radius. However, as our analysis only compares two broad energy bands, the simplified approach is sufficient for a quantitative comparison with our results.

To qualitatively explain our results in the light of this simplified model, we introduce an evolution of the spectral properties of the inner and outer half of this differentially precessing flow. At high full-band QPO frequencies (> 2 Hz), the inner half has a harder spectrum than the outer half, as expected. As the full-band frequency decreases, the spectral shapes change, such that at ~ 2 Hz, both the inner and outer spectrum have the same hardness. Below ~ 2 Hz, the spectral hardness of the inner and outer parts have switched around so that the outer half is harder than the inner half. In this scenario, the hard and soft band are not always accurate tracers of the inner and outer parts of the inner flow. This causes the observed $\Delta\nu_0(\text{hard} - \text{soft})$ to change from positive to negative as frequency evolves from high to low, while the physical frequency difference $\Delta\nu_0(\text{Inner} - \text{Outer})$ stays positive, as expected. In Figure 21, we show a simplified sketch of the model geometry and associated spectral evolution.

This differential precession model is able to explain several of the observed features of the QPO frequency and phase lag. A first advantage is the fact that the difference in QPO frequency is always intrinsically present, even when it is not observed. Thus, the phase lag between the inner and

Figure 21.: Cartoon depiction of the proposed geometry and spectral evolution in the differential precession model. Combined, the inner and outer part constitute the inner accretion flow, while the disk is not shown. ν_0^{ref} corresponds to the QPO frequency in the full energy band, while ν_{qpo} refers to the QPO frequency observed in a narrow energy band.

outer half can always evolve and decohere the precession. As each QPO excitation $\Delta\nu_0$ can change stochastically, coherent intervals will vary in length. Thus, this mechanism can explain both why the QPO is always quasi-periodic, even when no frequency difference is observed, and why the phase lags evolve quicker as the QPO decoheres faster. While intrinsically, the inner and outer half might get completely out of phase as the QPO decoheres ($\phi_{\text{intrinsic}} = \pm\pi$), we observe the spectrally weighted energy bands. Thus, the maximum observed phase lags are actually smaller than $\phi = \pm\pi$, as in our results.

Secondly, this model can explain the observed properties of the average phase lags as a function of both energy and frequency. On average, the outer half will lag behind the inner half with a lag of $\sim \pi/2$, assuming that the two halves get completely out of phase during a coherent interval. When the inner half has a harder spectrum (high full-band QPO frequencies), this causes an observed soft lag. However, when the spectrum is reversed (low full-band QPO frequencies), the observed phase lag becomes hard. Thus, this model can explain the change from hard to soft lags as the QPO frequency increases. This intrinsically constant average phase lag will also display an energy dependence similar to the QPO frequency, as is observed. An important argument for our interpretation is the peculiarity that the energy dependence of both the QPO frequency and phase lag flip over at ~ 2 Hz, where the lags change from hard to soft. In this model, this frequency of ~ 2 Hz is simply the frequency where the spectra of the inner and outer halves have the same shape.

This interpretation with a differentially precessing inner flow introduces a fairly complicated extra geometry, with different intrinsic precession frequencies. [Nixon et al. \(2012\)](#) have shown that in cases with a sufficient misalignment between the outer disk and the black holes spin axis, Lense-Thirring torques can break the inner flow into separately rotating rings. However, to explain the coherent intervals as subsequent QPO excitations, these rings should set up and decay on very short timescales. Thus, it could be envisioned more as the introduction of a gradual warp in the inner accretion flow. However, due to our limited energy resolution of two broad energy bands, it is not possible to distinguish between a gradual warp or two independently precessing parts.

Several possible explanations for the required spectral evolution could exist. Here, we discuss two options, although others might be possible as well. First, the reflected spectrum could contribute: in the differential precession model, this reflected spectrum is expected to be dominated by photons emitted by the outer precessing half, as this blocks photons from the inner half. Thus, the reflected spectrum will show a relatively low precession frequency. However, the hardness of the reflected spectrum can change: based on observed relations between spectral and timing properties of accreting black holes, the disk ionisation is expected to increase with full-band QPO frequency (see e.g. [Gilfanov 2010](#) for an overview or [Shaposhnikov & Titarchuk 2009](#) and [Garcia et al. 2015](#) for specific examples). This would imply a relatively soft reflected spectrum at high full-band QPO frequencies, and a relatively hard reflected spectrum at low full-band QPO frequencies. As the reflected spectrum shows a slow precession, this effect could possibly cause the QPO frequency to change from increasing to decreasing as a function of energy.

A second explanation could be the presence of an extra cooling process very close to the black hole. At high full-band QPO frequency, the differ-

entially precessing inner flow is expected to be small. In this scenario, the cooling process would be present in the entire flow, leading to a steepening of the spectrum everywhere. At lower full-band QPO frequencies, the inner flow extends further out - if the extra cooling term is only present in the innermost part, this might cause a steepening of the spectrum only for the regions precessing most quickly. Thus, at these low frequencies, this additional cooling term could cause the outer part to be harder than the inner part. Such an extra cooling process could for instance be pair-production and annihilation, which, if present, would be expected close to the black hole due to the high densities in the flow (Fabian et al., 2015).

The origin of the spectral evolution as a function of full-band QPO frequency might also account for the differences in energy dependence of the QPO frequency between GRS 1915+105, XTE J1550-564 and H1743-322, and the lack of observed energy dependence in other BHXRBS. For example, if the reflected spectrum plays a significant role, its inclination-dependence could contribute to differences in behaviour between these sources. The role of reflection could be tested by considering the energy-dependence of the QPO frequency below ~ 2 keV: if this role is indeed significant, reprocessing of photons by the disk is also expected at lower energies. As the reprocessed photons also originate from the slower oscillating outer half, this would cause a low observed QPO frequency below ~ 2 keV, similar to the QPO frequency in our hard band. This energy range could in the future be probed by for instance the *Neutron star Interior Composition ExploreR* (NICER; Gendreau et al., 2012) to search for this effect.

2.4 CONCLUSION

We have developed and implemented a new model-independent method to compare QPO properties between different energy bands on short timescales. This has allowed us to test whether the observed energy dependence of the QPO frequency corresponds to a genuine difference in this frequency between energy bands. We find that no clear (anti)correlations exist between QPO frequency and amplitude in filtered QPO lightcurves, and that the phase lag between two energy bands increases throughout each coherent interval ($\sim 5 - 10$ QPO cycles), before resetting at the end. We also find that the speed of this phase lag evolution is larger when the coherent interval is shorter. These results lead us to conclude that the QPO possesses a genuinely different frequency in different energy bands. We interpret this in the context of a geometric toy model, where the QPO is caused by differential precession of the innermost accretion flow. This toy model allows us to qualitatively explain (1) the energy-dependence of the QPO frequency and phase lag, (2) the smooth relation between full-band QPO frequency and phase lag and (3) the decoherence of the QPO on the timescale of $\sim 5 - 10$ QPO cycles.

Our newly developed method to track the evolution of QPO properties on short timescales shows promise for a more general *coherent-interval resolved* analysis of the QPO. For instance, next to the phase lag, it is possible to track the QPO frequency and frequency difference and compare these within and between coherent intervals. This would yield a much more detailed look at the relations between the length of coherent intervals, the QPO frequency difference and the speed of the phase lag evolution. Such an analysis could shed more light on the intricacies of the underlying QPO mechanism. Applying this method using narrower energy bands could also provide a useful

handle on the physical location where the phase lag evolution and the frequency difference originate. In Chapter 4, we present the results of such extension to our method in order to investigate differential precession more extensively.

INCLINATION DEPENDENCE OF PHASE LAGS IN BLACK HOLE X-RAY BINARIES

Abstract

Quasi-periodic oscillations (QPOs) with frequencies from ~ 0.05 -30 Hz are a common feature in the X-ray emission of accreting black hole binaries. As the QPOs originate from the innermost accretion flow, they provide the opportunity to probe the behaviour of matter in extreme gravity. In this Chapter, we present a systematic analysis of the inclination dependence of phase lags associated with both Type-B and Type-C QPOs in a sample of 15 Galactic black hole binaries. We find that the phase lag at the Type-C QPO frequency strongly depends on inclination, both in evolution with QPO frequency and sign. Although we find that the Type-B QPO soft lags are associated with high inclination sources, the source sample is too small to confirm this as a significant inclination dependence. These results are consistent with a geometrical origin of Type-C QPOs and a different origin for Type-B and Type-C QPOs. We discuss the possibility that the phase lags originate from a pivoting spectral power law during each QPO cycle, while the inclination dependence arises from differences in dominant relativistic effects. We also search for energy dependences in the Type-C QPO frequency. We confirm this effect in the three known sources (GRS 1915+105, H1743-322 and XTE J1550-564) and newly detect it in XTE J1859+226. Lastly, our results indicate that the unknown inclination sources XTE J1859+226 and MAXI J1543-564 are most consistent with a high inclination.

In-depth studies of the QPO in individual sources, such as presented in the previous Chapter, can be complemented by the simultaneous investigation of multiple BHXRBS. Sample studies of the QPO across many BHXRBS have recently played an important role in distinguishing between geometric and intrinsic QPO models: both [Motta et al. \(2015\)](#) and [Heil et al. \(2015b\)](#) present an inclination dependence in the timing properties of BHXRBS, as expected in geometric QPO models ([Schnittman et al., 2006](#)). As other recent evidence also points to a geometric QPO origin (see Sections 1.2.5 and 2.3.4), more properties of the QPO can be expected to depend on binary orbit inclination.

The currently known inclination dependences (i.e. [Heil et al. 2015b](#); [Motta et al. 2015](#)) merely provide evidence for geometric QPO models, but not on the detailed properties of this geometry. In other words, they do not present direct evidence for, for instance, Lense-Thirring precession. Since different regions in the accretion flow dominate different parts of the observed BHXRBS X-ray spectrum ([Done et al., 2007](#)), the inclination dependence of more energy-resolved QPO properties could provide new insights into the origin of both type-B and type-C QPOs. One such property is the phase lag at different Fourier frequencies, especially the QPO fundamental and (sub)harmonics. The different types of QPOs are known to show different

phase lag behaviour at the QPO fundamental and (sub)harmonic (Casella et al., 2005; Pahari et al., 2013). Furthermore, the relation between QPO frequency and QPO phase lag differs between sources (see e.g. Figure 4 in Remillard et al. 2002 and Figure 3 in Reig et al. 2000) - most noteworthy is the apparently log-linear dependence of QPO phase lags on QPO frequency in GRS 1915+105 (Pahari et al., 2013; Qu et al., 2010; Reig et al., 2000). These known results make phase lags an interesting QPO property to compare between sources of different inclination.

In this Chapter, submitted in slightly altered form to MNRAS, I present a systematic, model-independent analysis of the inclination dependence of phase lags using archival RXTE observations of fifteen BHXRBS. I determine these phase lags at various Fourier frequencies: the QPO fundamental, (sub)harmonics, and a range dominated by the broad band noise (BBN). I also measure the difference in Type-C QPO frequency between energy bands for all sources and investigate the inclination dependence of this frequency difference. This analysis allows us to compare QPO properties between both sources of different inclinations and different QPO types.

I find that the phase lags show a clear dependence on inclination for the type-C QPO fundamental and the Type-C BBN. Furthermore, I find significant QPO frequency differences in four sources, none of which are consistent with a low inclination. Lastly, our results indicate that XTE J1859+226 and MAXI J1543-564, both sources of undetermined inclination, are consistent with a high inclination. This Chapter is structured as follows: in the next Section, I present the source sample, the set of analysed RXTE observations, and our analysis method. In Section 3.2, I present our results on both phase lags and frequency differences, which are subsequently discussed in Section 3.3. Finally, I present our conclusions and future outlook in Section 3.4.

3.1 SAMPLE AND DATA ANALYSIS

To compare QPO and BBN phase lag properties between black hole binaries of different inclinations, we analyse archival RXTE observations of 15 Galactic BHXRBS: the fourteen sources analysed in Motta et al. (2015) and GRS 1915+105. For all sources except GRS 1915+105, we analyse a selection of observations based on the samples in Motta et al. (2015), Casella et al. (2004) and Shaposhnikov & Titarchuk (2009). We apply the same classification of sources into high (*more edge-on binary orbit*), low (*more face-on binary orbit*) and undetermined inclination, and of QPO observations into type-B and type-C, as Motta et al. (2015). The sample includes all sources analysed by Heil et al. (2015b) except for XTE J1118+480, GS 1354-64 and XTE J1720-318, since as these sources were not analysed by Motta et al. (2015), no extensive overview of QPO observations is available. The classification of inclinations as low or high is consistent between Motta et al. (2015) and Heil et al. (2015b), except for XTE J1817-330: this source is treated as low inclination in the former, but as undetermined in the latter.

As stated, we adopt the inclination classifications as applied in Motta et al. (2015). For ten sources, a direct inclination estimate or an upper/lower limit on inclination is available. References for these estimates are listed in Table 3. Besides direct estimates, the inclination classification is mostly based on the presence of lightcurve dips, observed in all sources identified as high inclination except GRS 1915+105. Additional constraints are provided by the source's tracks in their colour-luminosity diagram, as investigated by

Table 3.: Summary of the source sample, based on [Motta et al. \(2015\)](#). Type-B and Type-C indicate the amount of analysed observations containing the respective QPO type. In addition to direct inclination estimates, the inclination classification is based primarily on lightcurve dips, as well as the presence of equatorial winds and the source’s colour-luminosity track ([Muñoz-Darias et al., 2013](#)). References for the direct inclination measure (4th column): [1] [Neustroev et al. \(2014\)](#), [2] [Orosz et al. \(2002\)](#), [3] [Orosz et al. \(2004\)](#), [4] [Munoz-Darias et al. \(2008\)](#); [Zdziarski et al. \(1998\)](#), [5] [Miller-Jones et al. \(2011\)](#), [6] [Corral-Santana et al. \(2011\)](#), [7] [Orosz et al. \(2011\)](#), [8] [Greene et al. \(2001\)](#), [9] [Steiner et al. \(2012\)](#), [10] [Mirabel & Rodriguez \(1994\)](#)

Source	C	B	i	Sample	Ref.
Swift J1753.5-01	31	–	$\sim 40 - 55^\circ$	Low	[1]
4U 1543-47	5	3	$20.7 \pm 1.5^\circ$	Low	[2]
XTE J1650-500	21	1	$> 47^\circ$	Low	[3]
GX 339-4	46	17	$40 \leq i \leq 60^\circ$	Low	[4]
XTE J1752-223	3	1	$\leq 49^\circ$	Low	[5]
XTE J1817-330	–	9	–	Low	–
XTE J1859+226	26	14	$\geq 60^\circ$	–	[6]
MAXI J1543-564	4	–	–	–	–
XTE J1550-564	63	15	$74.7 \pm 3.8^\circ$	High	[7]
4U 1630-47	5	–	–	High	–
GRO J1655-40	27	1	$70.2 \pm 1^\circ$	High	[8]
H1743-322	104	36	$75 \pm 3^\circ$	High	[9]
GRS 1915+105	27	–	$70 \pm 2^\circ$	High	[10]
MAXI J1659-152	37	6	–	High	–
XTE J1748-288	6	–	–	High	–
<i>Total</i>	405	103			

[Muñoz-Darias et al. \(2013\)](#), and the observation of winds in high inclination sources. For more background on this inclination classification, we refer the reader to the extensive discussion in the main text and second appendix in [Motta et al. \(2015\)](#).

For GRS 1915+105, we consider the same observations as [Van den Eijnden et al. \(2016\)](#), which are specifically selected on type-C QPOs spanning a range in frequency from ~ 0.5 to ~ 8.0 Hz. No accurate estimates of the binary orbit inclination exist for this system. However, its relativistic jet has an inclination of $70 \pm 2^\circ$ ([Mirabel & Rodriguez, 1994](#)). Assuming that any misalignment between the jet and the binary orbit axis is small, we thus classify GRS 1915+105 as a high inclination system. In Table 3, we list our complete sample, including inclination estimates and classifications, and number of type-B and type-C QPO observations. In Appendix B, we list all analysed observations with corresponding QPO frequency and measured phase lag.

A small fraction of the observations in [Motta et al. \(2015\)](#) are not analysed in this work, for several reasons. In some cases, the QPO possesses a negligible rms amplitude at either the soft or hard energies. Alternatively, for some observations, the short exposure results in an inaccurate estimate of phase lags. In both cases, this leads to barely constrained phase lag values around the QPO fundamental and especially (sub)harmonic frequencies. In a minority of cases, we were unable to identify the QPO in the considered observation due to it possessing a low amplitude in all energy bands. In all three cases, we disregard these observations from our sample.

Depending on the RXTE observation mode, we extract binned, event or GoodXenon lightcurves with a 1/128 s resolution in three energy bands: a full band from 2 to 13 keV, and a soft and hard band from 2 to 7 and 7 to 13 keV respectively. To account for gain changes in the *Proportional Counter Array* (PCA) detector, we select the absolute PCA channels most closely matching the aforementioned energy ranges for each observation individually. We calculate the power density spectrum (PDS) in the full band with a frequency resolution varying from 1/8 Hz to 1/64 Hz, in order to adequately sample the QPO peak while keeping the power uncertainties small. We apply an rms-squared normalisation (Belloni & Hasinger, 1990) and fit the resulting PDS using XSPEC v12¹ with a model consisting of a constant white noise and a combination of Lorentzians to account for the BBN, the QPO fundamental and possible QPO (sub)harmonic.

To calculate phase lags, we determine the complex cross spectrum between the soft and hard band lightcurve and average the real and imaginary part over the considered frequency ranges. We then calculate the phase lag from this averaged cross spectrum, while we determine the errors through the raw coherence (see e.g. Uttley et al., 2014, section 2). The cross spectrum at the QPO fundamental and (sub)harmonic are averaged over a symmetric range around the fitted Lorentzian centroid frequency with a width of the fitted FWHM. The BBN cross spectrum is averaged over the frequency range between 0.5 and 1.5 Hz, provided that both the QPO and subharmonic frequencies are larger than 2 Hz to ensure that the lags are not directly contaminated by the QPO. In this paper, following common convention, we define hard lags (i.e. hard photons lagging soft photons) as positive.

In addition to phase lags between a broad soft and hard band, we also calculate lag-energy spectra of the Type-C QPO: the Type-C QPO lag between narrow bands and a soft reference band, as function of energy. For this purpose, we extract lightcurves at the highest possible energy resolution given the observation's data mode. We calculate the cross spectrum between each narrow energy band lightcurve and the softest energy band lightcurve (generally between ~ 2 and $\sim 3 - 4$ keV, depending on the RXTE data mode), using the same method as detailed above (see Uttley et al., 2014, for a more extensive description of this procedure). Subsequently, we average the real and imaginary part of the cross spectra at the QPO frequency in each cross spectrum, calculate the QPO lag from this averaged cross spectrum, and plot the lag as a function of energy.

Finally, for each Type-C QPO observation, we also calculate the difference between the QPO centroid frequency in the broad hard and soft band. We calculate the normalised PDS of the two broad energy band lightcurves, and fit the QPO frequency for each lightcurve using the multi-Lorentzian fit detailed above. We define the frequency difference $\Delta\nu_0$ as the QPO frequency in the hard band minus the QPO frequency in the soft band. When plotting this frequency difference, we rescale it by the QPO frequency in the full energy band ν_0 .

3.2 RESULTS

In Figures 22 and 23, we show several examples of logarithmically rebinned lag-frequency spectra of observations showing Type-B and Type-C QPOs, respectively. The QPO fundamental phase lags are averaged over the grey frequency range. The lag at the QPO fundamental is clearly well constrained in

¹ <https://heasarc.gsfc.nasa.gov/xanadu/xspec/>

Figure 22.: Examples of geometrically binned Type-B QPO lag-frequency spectra. The grey band indicates the range over which the QPO lag is averaged ($\nu_0 \pm \text{FWHM}/2$). The vertical dotted lines indicate (sub)harmonic frequencies.

Figure 23.: Examples of geometrically binned Type-C QPO lag-frequency spectra. The grey band indicates the range over which the QPO lag is averaged ($\nu_0 \pm \text{FWHM}/2$). The vertical dotted lines indicate (sub)harmonic frequencies.

this range, regardless of the amount of noise at surrounding frequencies. If present, (sub)harmonic frequencies are indicated with a vertical dashed line. These (sub)harmonic frequencies generally show larger uncertainties in the lags, due to the smaller amplitude of these features. In none of the shown Type-C observations is a (sub)harmonic present in the power spectrum. Generally, the Type-B lag-frequency spectra show larger lag uncertainties than Type-C lag-frequency spectra.

3.2.1 Type-C QPO fundamental lags

Figure 24 summarizes the main result of this paper: we plot the Type-C QPO fundamental phase lag as a function of QPO frequency for all high (grey points) and low (black points) inclination BHXRBs. For plotting clarity, the high and low inclination lags are binned separately in logarithmic QPO frequency bins and plotted as the blue squares and red circles, respectively. The different markers in the unbinned datapoints represent different sources. All subsequent analysis has been performed on the unbinned data.

Clear differences are present between the lag behaviour of low and high inclination sources: the dependency of the lags on QPO frequency appears mirrored around a marginally hard average lag: both samples possess a slightly hard lag at low QPO frequencies, at high frequencies high inclination sources turn to soft lags while lags in low inclination sources become harder. In Section 4, we will discuss the peculiar behaviour of GRO J1655-40, which shows no significant changes in lag as a function of QPO frequency, and the lag behaviour of the two sources of unclear inclination (XTE J1752+226 and MAXI J1543-564). In addition, we will also discuss GRS 1915+105, which is the only source showing increasingly hard lags towards low (< 2 Hz) frequency. Plots of the Type-C QPO lags as a function of QPO frequency for each individual source are available in Appendix A.

To better understand the differences in phase lags between low and high inclination sources, we compare the calculated higher resolution lag-energy spectra for observations with a low ($\sim 0.5 - 1.5$ Hz) and high (~ 6 Hz) Type-C QPO frequency. In Figure 25, we show representative lag-energy spectra for three low and three high QPO frequency observations. The plotted sources include low, high and undetermined inclination sources. For some lag-energy spectra, the uncertainties are relatively large as the QPO is not fully spectrally coherent between all energy bands (Vaughan & Nowak, 1997). The reference energy band for each observation (generally between ~ 2 and $\sim 3 - 4$ keV, depending on the RXTE data mode) is indicated by the softest energy range with zero phase lag.

Figure 24.: Type-C QPO phase lags as a function of QPO frequency. Black and grey points indicate low and high inclination sources, respectively. Different sources within the low or high inclination sample are shown with different markers. The red and blue points show the average phase lag and QPO frequency in logarithmic frequency bins. The plots for individual sources, including the two sources of unknown inclination, can be found in Appendix A.

Figure 25.: Representative examples of high and low inclination lag-energy spectra for observations with a low ($\sim 0.5 - 1.5$ Hz, left) and high (~ 6 Hz, right) Type-C QPO frequency. The sources, RXTE ObsId and inclination are listed in both panels. The reference energy bands are indicated with zero phase lag.

At low QPO frequency (left panel), the lag-energy spectra of high and low inclination sources are shaped similarly, as expected based on the similar broad-band lags visible in Figure 24. GRS 1915+105 however shows distinct behaviour at these low frequencies, as its broad-band QPO lag does not tend to zero but becomes increasingly harder. At high QPO frequency (right panel), the high and low inclination lag-energy spectra are opposite in sign but similar in shape, following a log-linear dependence on energy. Hints of a break, well-known in GRS 1915+105 (Pahari et al., 2013), are visible as well in several observations around an energy of 6–7 keV. The similarity in shape between low and high inclination sources at the same QPO frequency (barring GRS 1915+105 below 2 Hz) suggests that the underlying mechanism could be the same for hard and soft lags; in that scenario, the inclination is merely responsible for the sign and not the shape of the observed lag and lag-energy spectrum.

3.2.2 Significance testing

The Type-C QPO phase lags show a clear difference between low and high inclination sources. However, the number of sources in our sample is quite small. Thus, we want to know how likely it is that the observed differences between low and high inclination could have arisen by chance. In other words, we want to test the null hypothesis that the Type-C QPO phase lags do not depend on the inclination of the source.

As our test statistic for the difference between phase lags in the two samples, we adopt the absolute difference between the error-weighted average phase lags of the samples:

$$D = \left| \sum_{\text{high}} \frac{\phi}{\sigma_{\phi}^2} \bigg/ \sum_{\text{high}} \frac{1}{\sigma_{\phi}^2} - \sum_{\text{low}} \frac{\phi}{\sigma_{\phi}^2} \bigg/ \sum_{\text{low}} \frac{1}{\sigma_{\phi}^2} \right| \quad (7)$$

where *high* and *low* indicate a sum over the high and low inclination observations, respectively. This test statistic offers a simple parameterization of any phase lag differences, without making any assumptions about the dependence on QPO frequency. Hence, it will also be applicable to the lags calculated in other frequency ranges.

We test our null hypothesis, that the Type-C QPO phase lags do not depend on source inclination, by redistributing the source inclinations and consequently recalculating D 10^4 times. We redistribute the source inclinations by appointing to each source a new inclination classification, drawn from the set of original classifications ($5 \times \text{low}$, $7 \times \text{high}$) without replacement. We then estimate the p-value as the fraction of simulated D -values equal to or larger than the D -value of the observed data.

The resulting p-values are listed in Table 4. As the Type-C QPO lags appear to show the same behaviour in all sources at low QPO frequency, we calculate two p-values: one for the full set of observations, and one with a minimum frequency cutoff of 1 Hz. At a confidence level of 99%, we reject our null hypothesis for the Type-C QPO lags with frequency cutoff. In addition, the Table includes p-values for various other frequency ranges – these will be discussed in the next sections.

Figure 26.: Phase lags averaged over different Fourier frequency ranges for observations containing either Type-B or Type-C QPOs. The left column contains phase lags at the Type-B fundamental, harmonic and BBN. The right column contains phase lags at the Type-C subharmonic, harmonic and BBN. The frequency on the horizontal axes refers to the averaged frequency range, except for the BBN: there it refers to the QPO fundamental frequency. The binning, colors and markers of the datapoints are the same as in Figure 24. The significance of any differences between the distribution of high and low inclination lags is listed in Table 4.

Table 4.: Summary of the significance tests of the inclination dependence of QPO lags. BBN refers to the frequency range 0.5–1.5 Hz, when no QPO/subharmonic is present below 2 Hz. S and NS refer to the significance/non-significance of the result at a confidence level of 1%.

	D_{data}	P-value	Significance
Type-C QPO (no cut)	0.072	0.0408	NS
Type-C QPO (> 1 Hz)	0.236	0.0056	S
Type-B QPO	0.321	0.0600	NS
Type-C harmonic	0.180	0.0490	NS
Type-B harmonic	0.012	0.6728	NS
Type-C subharmonic	0.155	0.1090	NS
Type-C BBN	0.143	0.0016	S
Type-B BBN	0.169	0.0467	S

3.2.3 *Type-B QPO fundamental lags*

As the Type-C QPO phase lags show a clear dependence on inclination, it is interesting to consider other frequency ranges and types of QPOs. Plots of these additional phase lag results are shown in Figure 26. All plots are structured and binned in the same manner as Figure 3, including the same markers for the individual sources. Note that XTE J1817-330, which is not included in Figure 24 as there are no Type-C QPO observations of this source in our sample (see Table 3), is indicated in Figure 26 by the black squares. The lags in the upper four panels are plotted as a function of their associated frequency (i.e. QPO fundamental, harmonic or subharmonic frequency). The lower two panels (the BBN lags) are plotted against the QPO fundamental frequency, although all BBN lags are calculated in the same frequency range. For consistency, we have binned the lags in all panels in a similar fashion as for Figure 24. However, in several panels (especially the top left and bottom right), the scatter in lags renders these bin averages physically uninteresting. All statistical analysis and interpretations in this paper are based on the unbinned measurements.

The Type-B fundamental lags are shown in the upper left panel. The lags do not reveal a dependence on inclination as clearly as the Type-C fundamental lags: as is shown in Table 4, we cannot conclude that the Type-B fundamental lags depend on inclination at a 99% confidence level. However, this might be partially attributable to the smaller number of sources showing Type-B QPOs, and the larger uncertainties on the lags. The observed soft lags are dominated by high inclination sources, namely XTE J1550-564, H1743-322, MAXI J1659-152 and 4U 1630-47. These soft lags are unexpected, as the Type-B QPO classification as defined in [Casella et al. \(2005\)](#) states that the Type-B fundamental lags are hard. We will discuss this discrepancy further in section 3.3.2.

3.2.4 *Harmonics and broad-band noise*

The remaining panels of Figure 26 show the phase lags at the Type-C harmonic, subharmonic and BBN, and at the Type-B harmonic and BBN. As stated earlier, the BBN lags are calculated over a frequency range from 0.5–1.5 Hz. We do not consider the Type-B subharmonic, as our sample did not contain significant detections of such a feature. The phase lags at the harmonic are predominantly soft and hard for the Type-B and Type-C QPO respectively, as expected based on the QPO Type-Characteristics listed in [Casella et al. \(2005\)](#). Interestingly, the Type-C subharmonic phase lags show both hard and soft lags, with all soft lags observed in high inclination sources. Lastly, the BBN phase lags cover a wide range in values for both QPO types, without an apparent dependence on the QPO frequency.

Interestingly, as Table 4 shows, the BBN lags of the Type-C QPO depends significantly on inclination at a 99% confidence level. Although differences between the high and low inclination sample appear to be present in some other cases, especially at the Type-C subharmonic, the number of low inclination sources is too low to find a significant inclination dependence. It is however interesting to note that for the Type-C subharmonic, only high inclination sources show soft lags. Additionally, it is notable that the Type-C harmonic lags are remarkably constant for high inclination sources, while the lags in low inclination sources do show variations between sources.

Figure 27.: The relative frequency difference between the Type-C fundamental centroid in the hard and soft band $\Delta\nu_0/\nu_0$, as a function of QPO frequency in the full band ν_0 . The binning, markers and colors of the datapoints are the same as in Figure 24.

3.2.5 Energy-dependent frequency differences of the Type-C Fundamental

In addition to the inclination dependence of the average phase lag, we also investigate the short-timescale evolution of the phase lag by measuring differences in Type-C fundamental frequency between the soft and hard band. We plot these frequency differences, divided by the QPO frequency itself, as a function of Type-C QPO frequency in Figure 27, in a similar fashion to Figures 24 and 26. Again, we bin low and high inclination sources separately for clarity, while plotting the unbinned data in black and grey, respectively. Low inclination source XTE J1748-288 is not included, as the error on the fitted QPO centroid frequency is too large to accurately determine frequency differences.

There is no immediately clear difference between high and low inclination sources, contrary to the average Type-C QPO phase lags. At low QPO frequency, few significantly non-zero frequency differences are present, as was known from earlier investigations (Li et al., 2013a,b; Qu et al., 2010). At higher QPO frequencies, non-zero frequency differences become more prevalent, but due to large uncertainties, these are rarely significant. Hence, no clear difference can be found when grouping the sources as low and high inclination.

In order to investigate the presence of QPO frequency differences in more detail, we consider them for all individual sources, independent of inclination. For each source, we use a χ^2 -test to test whether the distribution of frequency differences is consistent with zero within the 99% confidence limit. This returns four sources with a significant relative frequency difference, which are shown in Figure 28. The three high inclination sources (GRS 1915+105, XTE J1550-564 and H1743-322) were already known to show frequency differences (see respectively Li et al., 2013a,b; Qu et al., 2010). The

Figure 28.: Relative frequency difference between the Type-C centroid in the hard and soft band $\Delta\nu_0/\nu_0$, as a function of QPO frequency in the full band ν_0 , for four sources: GRS 1915+105, XTE J1550-564, H1743-322 and XTE J1859+226. Only these four sources show significant non-zero frequency differences. Similar plots for all sources individually can be found in Appendix A.

fourth source, with unknown inclination (XTE J1859+226), also shows clear non-zero frequency differences at high QPO frequencies. In addition, XTE J1752-223 (low inclination) shows marginally significant frequency differences ($p = 0.02$). However, we analysed only three Type-C observations for this source, so the significance of this result should be regarded with caution. No other low inclination sources show hints of significant frequency differences. In appendix A, we plot the frequency difference as a function of QPO frequency for each individual source.

3.3 DISCUSSION

In this paper, we present the results of a systematic, model-independent analysis of the inclination dependence of QPO phase lags and frequency differences. We report the following two main results:

- (1) the phase lag at the Type-C fundamental QPO frequency depends significantly on source inclination, both in sign and relation to QPO frequency, and
- (2) four sources (three high inclination and one undetermined inclination) show significant frequency differences in the Type-C QPO.

In this section, we will discuss and interpret all Type-C and Type-B results, the absence and presence of frequency differences, and the behaviour of several individual sources.

3.3.1 Type-C QPO lags

For the Type-C QPOs, we conclude that phase lags at the fundamental and the BBN depend significantly on inclination. This finding is consistent with the result in [Motta et al. \(2015\)](#) that the amplitude of both the QPO fundamental and the BBN depend on inclination. Furthermore, we find that the phase lags at the harmonic are, with a handful of exceptions, hard - high inclination sources even show a remarkably constant lag as a function of QPO frequency, as was known for GRS 1915+105 ([Pahari et al., 2013](#)). Lastly, it is important to note that for the QPO fundamental, inclination does not suppress or magnify the value of the observed phase lags. Instead, the signs are completely flipped, while the amplitudes remain comparable. Any physical interpretation of these results should be able to explain this behaviour.

Calculating the cross spectrum between two lightcurves does not allow for the direct distinction between QPO and BBN lags. In the example Type-C QPO lag-frequency spectra (Figure 23), lag features are visible at other frequencies than only the QPO frequency; for instance, in the top-right lag-frequency spectrum (H1743-322), the QPO lag is located on the slope of a larger lag feature at lower frequency. In the remainder of this discussion, we will assume that the lag at the QPO frequency is dominated by the QPO itself. However, we note that the BBN might also influence the measured lags. Hence, we merely discuss the general differences between the high and low inclination sources, such as Type-C QPO lag sign, and do not discuss more detailed lag properties.

We consider a physical explanation in the context of the Lense-Thirring precession model, wherein the inclination dependence follows from differences in dominant general-relativistic effects on the observed flux along different lines of sight. In this interpretation, the lags themselves (independent

Figure 29.: Sketch of Figure 3 from [Veledina et al. \(2013\)](#): the observed flux of the simulated precessing ring as a function of the cosine of the viewing angle θ . Large θ , shown on the left, corresponds to a high inclination (i.e. edge on) of the ring. The flux, fully modeled by [Veledina et al. \(2013\)](#), is sketched for a small and large ring.

of inclination) originate from spectral pivoting during each precession cycle. [Ingram & Van der Klis \(2015\)](#), [Ingram et al. \(2016\)](#) and Stevens & Uttley (in prep.) use phase resolved spectroscopy of GRS 1915+105, H1743-322 and GX 339-4 respectively to show that the power law spectral index Γ oscillates as a function of Type-C QPO phase. Although QPO phases cannot be mapped directly onto precession phases, this implies that the hardness changes during a precession cycle. If the Γ variation peaks after the QPO flux peaks, this causes a soft lag at the QPO frequency, and vice versa. Making the assumption that the spectral pivoting as a function of precession phase is independent of source inclination, both the sign and size of the observed lag will simply depend on the QPO flux as a function of precession phase.

[Veledina et al. \(2013\)](#) model a precessing inner flow as a precessing Comptonizing ring, taking relativistic effects into account. In Figure 29, we show a sketch of their Figure 3: the dependence of the observed flux on inclination for emission from different radii. The angle θ is the instantaneous inclination angle of the flow as seen by an observer – small θ corresponds to a face on precessing flow. The angle θ changes systematically during a

precession cycle. Figure 29 shows that, for small radii, this has the opposite effect on the observed flux for low and high inclination sources: high inclination sources become brighter as θ decreases ($\cos \theta$ increases), while for low inclination sources this effect is opposite. This difference arises due to differences in dominant relativistic processes: Doppler effects at low inclination versus solid angle effects at high inclination².

As mentioned earlier, the dependence of flux on precession phase could set both the sign and size of the observed QPO lag. Hence, this difference in dominant relativistic effects between inclinations could result in the observed inclination dependence of the Type-C QPO phase lag. As is visible in Figure 29, this effect might become less important at larger radii, as the curve's maximum shifts to high values of $\cos \theta$ and differences in relativistic effects become smaller (i.e. solid angle effects dominate for nearly all values of $\cos \theta$). This might explain the similarity in the Type-C QPO lags at low QPO frequencies, which correspond to larger radii.

The origin of the QPO phase lags and their inclination dependence can be probed more accurately using detailed lag-energy spectra. At low QPO frequency, high and low inclination sources display similar lag-energy spectra, just as they show similar average lags. At high QPO frequency, low and high inclination sources still show similarly shaped lag-energy spectra, but of opposite sign. This similarity in shape is consistent with our proposed interpretation of the inclination dependence of the broad-band lags: the lags themselves always originate from the same mechanism (the pivoting power law), while the sign is set by only the inclination through relativistic processes. In this scenario, it is expected that the sign of the lag-energy spectra would change, but not the shape. At low QPO frequencies, GRS 1915+105 is an exception: its lag-energy spectrum is not flat but similar to low-inclination high-QPO-frequency lag-energy spectra, suggesting an additional source of lags. We will discuss this behaviour in Section 3.3.4.3.

The phase lags at the Type-C harmonic show a completely different behaviour than those at the fundamental. The difference between high and low inclination sources is not significant at a 99% confidence level, as the sample contains a relatively small number of low inclination observations. More importantly, the lags are clearly hard, with little dependence on QPO frequency. Thus, contrary to the fundamental, in our sample the inclination does not have an influence on the sign of the phase lags at the harmonic. This points to a difference in the exact physical mechanism creating the harmonic signal.

Using frequency-resolved spectroscopy, [Axelsson et al. \(2014\)](#) show that the spectrum of the harmonic is softer than both the spectrum of the fundamental and the time-averaged spectrum. This indicates that the origin of the harmonic might be related to the thermal reprocessing and reflection of hard photons by the disk. Furthermore, [Ingram et al. \(2016\)](#) show that in H1743-322, several spectral parameters vary during each QPO cycle. However, these parameters show strong harmonic features (i.e. a second minimum and maximum) within each QPO cycle. This effect is interpreted as resulting from a precessing inner flow, which is inclined with respect to the disk and thus illuminates the disk on two sides: both the front and the back. In such an interpretation, this behaviour could also give rise to a harmonic at twice the QPO frequency. These results on the possible role of reflection in the harmonic tie in to our lag results, since the inclination

² Note that the limb darkening law also plays a role – however, as it does not depend on radius, we do not consider it in this interpretation.

of the disk with respect to the precessing inner flow is independent of the position of the observer. Hence, the position of the observer will not have a direct influence on the observed sign of the phase lag, as is the case with the fundamental. Instead, the lags might be set by the time difference between illuminating the receding and approaching side of the disk. In that scenario, red and blueshifting of the reflected harmonic could cause a lag that is independent of the observer's position.

3.3.2 *Type-B QPO lags*

Despite apparent differences in Figure 26, the Type-B fundamental phase lags do not show a dependence on source inclination at a 99% confidence level. Furthermore, these lags show a different dependence on QPO frequency than the Type-C lags. This adds to the evidence that the Type-B QPO arises from a different mechanism, (Fender et al., 2009; Motta et al., 2012), such as the precessing base of a jet (Motta et al., 2015; Stevens & Uttley, 2016). However, it is important to note that the data quality for the Type-B QPOs is less than for the Type-C QPOs: the uncertainties are larger and the number of observations and sources is smaller. If the Type-B QPO indeed originates from a jet-based mechanism, an inclination dependence in the phase lags might be expected. It is possible that the data quality in our analysis is simply too low to significantly detect such a dependence.

The soft lags we observe in several high inclination sources are not expected based on the Type-B QPO classification as defined by Wijnands et al. (1999) and Casella et al. (2005), which describes only hard lags in Type-B QPOs. We observe these unexpected lags in XTE J1550-564, H1743-322, 4U 1630-47 and MAXI J1659-152. From these sources, only XTE J1550-564 was included in the analysis of the original definition of QPO types. However, our sample of Type-B observations in XTE J1550-564 is more extensive than the original sample. Furthermore, Wijnands et al. (1999) define the fundamental QPO as the ~ 3 Hz peak, while we define this as the subharmonic. Thus, it appears to be too simple to state that all Type-B QPO fundamentals have hard lags. Based on our results, these lags could be both hard and soft.

3.3.3 *Broad-band noise*

The phase lags of the BBN associated with the Type-C QPO (in the 0.5–1.5 Hz range) show a significant inclination dependence at the 99% confidence level, while those associated with the Type-B QPO do not. In both cases, the soft lags appear to be dominated by the high inclination sources. For Type-B QPO observations, marginally hard lags are seen in low inclination sources. For Type-C QPO observations however, hard lags are observed in both high and low inclination sources. For both the Type-B and Type-C observations, the BBN lag does not show any clear dependence on the QPO frequency (as is clearly present in for instance the Type-C QPO lag).

Existing results on the inclination dependence of the BBN are not in full agreement: Heil et al. (2015b) find that the BBN, unlike the QPO, does not show a dependence on inclination. Motta et al. (2015) however show that the amplitude of the BBN associated with the Type-C QPO does depend on inclination. In comparing these results with ours it is important to note that Heil et al. (2015b) analyse observations regardless of the presence of the QPO. Our results on the other hand are derived from observations selected on QPO type. The bias introduced by this selection might partially explain

the inclination dependence observed for the Type-C BBN: if the physical origin of the QPO and BBN are related, the presence of a QPO might result in inclination dependent properties in the BBN. Alternatively, the BBN lags might simply be contaminated by QPO lags, despite our cut on QPO and subharmonic frequency.

Furthermore, the 0.5–1.5 Hz range is quite arbitrary, chosen merely for its relatively large BBN amplitude. Thus, we have repeated the analysis on the BBN lags averaged over the range 0.1–0.5 Hz in the same observations (i.e. those with no QPO or subharmonic below 2 Hz). As this range covers lower frequencies, possible contamination by the QPO lags is less likely. In this new range, no significant inclination dependence is present in both QPO types (Type-C: $p = 0.0729$, Type-B: $p = 0.73$), despite similar data quality. Thus, as the inclination dependence of Type-C BBN lags between 0.5–1.5 Hz appears to be closely related to the arbitrary frequency range, we do not draw general conclusions about the nature of the BBN associated with Type-C QPOs.

3.3.4 Individual sources

In our analysis, we have grouped sources into high, low or undetermined inclination and assumed that any significant differences in lag behaviour originate from differences in inclination. In this section, we will discuss two sets of individual sources: sources of undetermined inclination (XTE J1859+226 and MAXI J1543-564) and sources that do not follow the standard trend in Type-C fundamental lags as function QPO frequency (GRO J1655-40 and GRS 1915+105).

3.3.4.1 XTE J1859+226 and MAXI J1543-564

For both XTE J1859+226 and MAXI J1543-564, no definite classification into low or high inclination could be determined. Thus, both sources were excluded from the comparison between the high and low inclination samples. In Figure 30, we plot the Type-C QPO phase lag as a function of QPO frequency for these two sources. XTE J1859+226 clearly follows the same trend as the high inclination sources. Due to the small number of observations and larger uncertainties, the MAXI J1543-564 lags are less evident. However, the lags are clearly soft, consistent with the high inclination sample. Based on the QPO and noise amplitudes, [Motta et al. \(2015\)](#) also found that XTE J1859+226 and MAXI J1543-564 are most consistent with a high inclination. The clear presence of QPO frequency differences in XTE J1859+226 (Figure 28), previously only found in three high inclination sources, reinforces the interpretation that XTE J1859+226 is indeed a high inclination source.

3.3.4.2 GRO J1655-40

The high inclination source GRO J1655-40 does not seem to show any significantly non-zero phase lags at the Type-C QPO frequency. To demonstrate this, we plot these phase lags versus QPO frequency, together with all other high inclination sources in grey, in Figure 31; GRO J1655-40's behaviour is clearly distinct from the other sources above approximately 1 Hz. However, GRO J1655-40 is known to show remarkable timing properties, and is thought to accrete close to or above the Eddington rate ([Neilsen et al., 2016](#); [Uttley & Klein-Wolt, 2015](#)). [Uttley & Klein-Wolt \(2015\)](#) discuss the possibility that GRO J1655-40 has a slim disk, that could partially obscure the

inner flow. In such a scenario, the QPO could still arise from a precessing inner flow, although the changes in its solid angle would be dominated by the shielding by the disk. This could lead to the observed lags, which are consistent with zero, as the slim disk would block photons of all energies in the same manner.

3.3.4.3 GRS 1915+105

GRS 1915+105 was not considered by [Motta et al. \(2015\)](#) and [Heil et al. \(2015b\)](#), as it does not follow a standard q-shaped hardness-intensity diagram and possesses unique timing properties ([Belloni, 2010](#)). We have included this source in our sample, both because of the peculiar log-linear dependence of QPO phase lag on QPO frequency, and the clear presence of QPO frequency differences. Removing GRS 1915+105 decreases the significance of the inclination dependence of the Type-C QPO phase lag: the p-value with the 1 Hz cut increases to $p = 0.0138$. This is both the result of the small uncertainties on the QPO phase lags in GRS 1915+105, and most importantly the small number of sources in our sample. Hence, removing a source with relatively many, well-constrained lag measurements has a clear impact on the significance of the inclination dependence.

We find that, while the frequency differences are present in a selection of other sources, the log-linear phase lag behaviour is not seen in any other source. This is especially interesting towards the lowest QPO frequencies, where all other sources tend to zero phase lags, while the hard phase lags in GRS 1915+105 continue increasing. In our interpretation of the inclination dependence, this corresponds to a process where the flux increases as the precessing flow increases its inclination. For low inclination sources, at high QPO frequency, we suggested that this process could be caused by Doppler effects. However, this is less expected at the lower QPO frequencies where the QPO lags in GRS 1915+105 are hard, as this corresponds to a larger truncation radius. This is also shown by the model of [Veledina et al. \(2013\)](#): our sketch of their results (see Figure 29) shows that the influence of these Doppler effects becomes less prominent for precessing rings of larger radii.

An alternative interpretation is proposed in [Van den Eijnden et al. \(2016\)](#) through differential precession. GRS 1915+105 is the only known source showing negative frequency differences (for QPO frequencies below 2 Hz) ([Qu et al., 2010](#); [Yan et al., 2012](#)). In the differential precession model, this is interpreted as a QPO frequency that increases towards the inner radius, spectrally weighted by emission that becomes softer towards smaller radii. In this model, hard lags could correspond to the average lag of the slower precessing, harder regions with respect to the faster precessing, softer regions. This interpretation could explain why only GRS 1915+105 shows hard lags at low QPO frequencies, as it is the only source showing negative frequency differences. However, why the hardness would decrease towards smaller radii for low QPO frequencies only in GRS 1915+105, remains unanswered.

The log-linear dependence of QPO phase lag on frequency, including a sign change, was first reported by [Reig et al. \(2000\)](#). There, the change in lag sign was interpreted as the result of a change in optical depth of the Comptonizing region: a low optical depth, assumed to be present at lower count rates and thus low QPO frequency, results in a hard lag. Higher optical depths, present at high count rates i.e. high QPO frequency, lead to a soft lag. This scenario does not assume a specific model for the QPO, and hence would be consistent with geometric QPO models. However, the

Figure 30.: Type-C fundamental phase lags as a function of QPO frequency for XTE J1859+226 and MAXI J1543-564, both sources of undetermined inclination. Both sources show significant hard lags, consistent with the high inclination sample.

difference with other high inclination sources, not showing hard lags, is not directly accounted for.

3.3.5 QPO frequency differences and differential precession

We observe a significant non-zero frequency difference $\Delta\nu_0$ in four out of fifteen sources in our sample. In all cases, the trend as a function of QPO frequency is comparable, with a sudden increase in $\Delta\nu_0/\nu_0$ at the highest QPO frequencies ($> 6-7$ Hz). As already stated, three of these sources (GRS 1915+105, XTE J1550-564 and H1743-322) are high inclination. As discussed in Section 3.3.4.1, the fourth source, XTE J1859+226, shows very similar Type-C QPO lag behaviour to the high inclination sample. No low inclination source shows a non-zero $\Delta\nu_0$; XTE J1650-500 for instance covers the same range in QPO frequencies and phase lags (with opposite sign) as XTE J1859+226, while its $\Delta\nu_0$ does not show non-zero values³. This is not due to the size of the uncertainties in XTE J1650-500 – these are much smaller than many of the observed non-zero values of $\Delta\nu_0$ in other sources.

In the high inclination source sample, three sources do not show significant frequency differences: GRO J1655-40, 4U 1630-47 and MAXI J1659-152. As discussed in section 3.3.4.2, the timing properties of GRO J1655-40 might be significantly influenced by the presence of a slim disk. In the latter two sources, the lack of frequency differences could simply result from the data quality: for 4U 1630-47, only few QPO observations, with large phase lag uncertainties, are analysed. In MAXI J1659-152, the QPO frequency barely passes 7 Hz, after which most non-zero frequency differences in other sources are observed. Hence, despite the lack of a clear inclination

³ See Appendix A for plots of $\Delta\nu_0$ versus ν_0 for each source.

Figure 31.: Type-C fundamental phase lags as a function of QPO frequency for GRO J1655-40 (black) and the remaining high inclination sources (grey). GRO J1655-40 shows constant phase lags, consistent with zero, contrary to the hard lags in other high inclination sources. This difference could be linked to the presence of a slim disc partially obscuring the precessing region (see text for details).

dependence, the frequency difference might only be observable in high inclination sources. The analysis of new transient BHXRBS or outbursts is thus vital to investigate this further.

Analysing $\Delta\nu_0$ sheds light on the short-time-scale evolution of phase lags in Type-C QPOs. Such lag evolution was reported by [Van den Eijnden et al. \(2016\)](#) in GRS 1915+105 and interpreted as a sign of differential precession of the inner accretion flow. In this interpretation, different radii in the inner flow precess at different rates, resulting in the gradual build-up of a warp over the course of 5 – 10 precession cycles. This warp in turn destabilises the precession, leading to the decoherence of the QPO. This process results in the presence of an enveloping structure (*coherent intervals*) in the QPO amplitude, in which the phase lag systematically increases, and might be related to the Q-factor and *quasi*-periodic nature of QPOs. Although not all sources show a significant $\Delta\nu_0$, the coherent intervals appear to be present in all sources, as expected in this interpretation: even if no $\Delta\nu_0$ is observed, the precession is still destabilised by its differential nature.

We measure $\Delta\nu_0$ as a probe of this effect, as the actual measurement of evolving phase lag requires too high a signal-to-noise to be applicable to all sources in our sample. This approach thus assumes that the link between differential precession and phase lag evolution is generally applicable to BHXRBS. A more problematic downside of this method is that the observed $\Delta\nu_0$ is not a direct proxy of the intrinsic difference in precession rate between different radii in the inner accretion flow: we can only observe the spectrally weighted $\Delta\nu_0$. If each radius emits with the same spectral shape, all energy bands will consist of the same weighted combination of precession frequencies. Thus, no $\Delta\nu_0$ will be observed.

As any observed $\Delta\nu_0$ is always spectrally weighted, investigating its inclination dependence could provide further information on the emission processes of the differentially precessing flow. If frequency differences are indeed only visible in edge-on sources, this would imply that the inner flow's spectral appearance in edge-on sources generally differs more at different radii than in face-on sources. This could for instance result from differences in the amount of scatterings between different inclinations: in face-on sources, photons received from all radii have scattered approximately the same amount of times. In edge-on sources, the innermost regions might be more obscured. As a result, photons from the smallest radii have scattered more, resulting in a harder observed spectrum from these smaller radii. To test such a scenario, it is important to test the hypothesis that Type-C QPO frequency differences are only visible at high inclination, in a new sample of sources.

3.4 CONCLUSIONS

In this paper, we present the results of a systematic investigation of the inclination dependence of both phase lags and QPO frequency differences in BHXRBS. We conclude that

1. the Type-C QPO lag depends clearly on source inclination, both in sign and behaviour as a function of QPO frequency,
2. The Type-B QPO lag does not depend significantly on source inclination, but shows distinct behaviour as a function of QPO frequency compared to the Type-C QPO,

3. the Type-C BBN lags show a dependence on inclination, but no clear dependence on QPO frequency,
4. four BHXRBS show a difference in Type-C QPO frequency between the hard and soft band, none of which are of low inclination,
5. XTE J1859+226, and to a lesser extent MAXI J1543-564, are consistent with being high inclination sources.

These results provide a clear indication that both the Type-B and Type-C QPOs originate from two distinct geometric processes, consistent with recent results. It is important that the inclination-dependence hypothesis, both for the lags and QPO frequency differences, will be tested in the future in new sources and/or outbursts. The Type-C QPO lags show a very clear and simple difference between high and low inclination: the sign of the phase lag. Hence, this will be easy to test with observations of new black hole transients which show Type-C QPOs above a few Hz and for which the inclination can be determined. Furthermore, systematically investigating energy-dependent frequency differences in the Type-C harmonic or Type-B fundamental could provide new insights in the differences between QPO types and harmonics.

COHERENT-INTERVAL RESOLVED QPO ANALYSIS

Abstract

We present the first coherent-interval resolved analysis of Type-C QPOs, in GRS 1915+105. Applying the method detailed in Chapter 2, we determine the coherent intervals in 27 observations of GRS 1915+105. Subsequently, we measure the lengths and average amplitude profile of the coherent intervals in each observation. This reveals that the lengths of successive coherent intervals are independent, suggesting that they originate from independent excitations of the same underlying mechanism. We also track the unfiltered hardness and lag-energy spectra throughout the coherent intervals. The former shows that the hardness changes significantly within the interval, suggesting a new connection between QPO and BBN. Finally, the coherent-interval resolved lag-energy spectra show that primarily the lags above $\sim 5\text{--}6$ keV evolve during the coherent interval. This result is consistent with the differential precession model presented in Chapter 2, and suggests that the QPO lag-energy spectrum might consist of multiple, physically distinct features.

In Chapter 2 of this thesis, I report the identification of a new phenomenon connected to the Type-C QPO: *coherent intervals*. Over such an interval's associated *coherence timescale*, on average comparable to the reciprocal of the QPO FWHM, the QPO cycle amplitude increases and decreases in a fashion enveloping the QPO cycles. Similar amplitude modulations were also observed in XTE J1550-564 by [Lachowicz & Done \(2010\)](#) and [Su et al. \(2015\)](#), although their properties and physical origin were not discussed in more detail. I have found that the phase lag associated with the QPO in GRS 1915+105 systematically increases during these coherent intervals, in a manner consistent with the presence/absence in the observation of an energy-dependent QPO frequency. As discussed in detail in Chapter 2, I interpret this result as the first observational sign of differential Lense-Thirring precession as the Type-C QPO origin.

In the differential precession interpretation, coherent intervals correspond to the setup and decay of the precession: the precession starts out coherent, but its differential nature destabilises the precession in a mere five to ten cycles. However, I have not yet investigated the basic properties of the coherent interval, besides the amplitude modulation and phase lag evolution. Investigating other properties could shed light on the viability of the interpretation offered above. For instance, the presence of the coherent intervals appears to be ubiquitous across both Type-C QPO frequency and sources (see Chapter 3). However, there might be subtle differences between the QPO amplitude profiles within the intervals for those different observations. Alternatively, the connection between successive coherent intervals is inter-

esting to consider: are these intervals related to each other, or does each constitute an independent excitation of the same process?

Differential precession also suggests that QPO observables (in addition to the amplitude) might vary on the coherence timescale. I have observed a first example of such evolution, in the phase lag, but other properties could evolve as well. Such evolution could also test the viability of differential precession: in Chapter 3, I have extended the already compelling evidence that the QPO originates from a geometric origin. However, differential precession is not the only geometric model. In order to compare it to other models, especially solid-body like precession, investigating the behaviour of other QPO properties throughout coherent intervals is vital. This distinction between geometric models is also important since an energy-dependent Type-C QPO frequency appears to occur more generally in BHXRBs and might possibly be inclination dependent. In other words, the observed hints of differential precession are not merely another peculiarity of GRS 1915+105.

In order to answer these questions, I present a coherent-interval resolved study of the Type-C QPO in GRS 1915+105. Analysing a slightly extended version of the set of observations in Chapter 2, I focus on the two issues detailed in the previous two paragraphs: the nature of the coherent intervals themselves, and the behaviour of spectral-timing QPO properties throughout individual coherent intervals. I study the former through the average profile of Type-C QPO amplitudes within a coherent interval and through the power spectral properties of the length of successive coherent intervals. In order to study the Type-C QPO itself, I present coherent-interval resolved measurements of the unfiltered lightcurve's hardness and of lag-energy spectra.

This Chapter is structured as follows: in the next Section, I very briefly reiterate the method detailed in Chapter 2 and introduce our extended approach to resolving the coherent intervals. In Sections 4.2 and 4.3, I present and discuss the results of this analysis, respectively. Finally, I present our concluding remarks in Section 4.4. The contents in the Chapter are a work in progress – in the final Section, I will also note the planned future steps with this coherent-resolved analysis. As I do not consider the Type-B QPO in this Chapter, I will refer to the Type-C QPO simply as the QPO.

4.1 METHOD

For our coherent-interval resolved analysis, we select the same RXTE observations of GRS 1915+105 as in Chapter 2, supplemented with three extra observations with a high energy resolution data mode¹. We extract binned mode lightcurves with a time resolution of 1/128 s, across multiple energy bands²: a full energy band from $\sim 2 - 13$ keV, a broad soft and hard band from $\sim 2 - 7$ and $\sim 7 - 13$ keV respectively, and narrow energy bands up to ~ 13 keV with the maximum energy resolution given the observational data mode. Subsequently, we calculate and fit the PSDs in the same manner as in Chapter 2.

After extracting the lightcurves and fitting their PSDs, our newly developed method, introduced in Chapter 2, consists of broadly three steps: first, we apply the optimal filter in order to obtain estimated QPO lightcurves.

¹ These observations are: 40403-01-11-00 ($\nu_0 = 4.3$ Hz), 40703-01-02-00 ($\nu_0 = 3.6$ Hz), and 40703-01-05-00 ($\nu_0 = 4.2$ Hz).

² Due to the slowly changing instrumental response RXTE PCA instrument, the energy ranges are given with approximate borders.

Subsequently, we determine the frequency and amplitude of each QPO cycle and identify coherent intervals by searching for local minima in QPO amplitude. Finally, we track the phase lag between the hard and soft lightcurve, as a function of both relative and absolute position in the coherent interval. For our coherent interval resolved analysis, we apply the first two steps to the full energy band in order to determine the coherent intervals. The subsequent analysis depends on the energy resolution of the considered lightcurves, and will be detailed in the following three subsections.

4.1.1 Full energy band analysis

We analyse the full energy-band lightcurve ($\sim 2 - 13$ keV) to probe the most basic properties of the coherent intervals themselves: their size and shape. First, we measure the length of each coherent interval simply as the time difference between two successive local minima in QPO amplitude, provided that this lightcurve segment does not contain gaps. Subsequently, we compute the PSD of the resulting sequence of coherent interval sizes, by averaging together the individual PSDs for segments of 64 coherent intervals (without applying the rms-normalisation). Technically, this sequence of lengths is not a time series with a fixed time resolution. Hence, the associated ‘frequencies’ do not correspond to timescales as they do for a regular lightcurve. Instead, the unit of these frequencies is *per coherent interval*. Thus, the shape of the PSD contains information on the possible connections between successive coherent intervals.

Furthermore, we also determine the average amplitude profile of the coherent intervals in each observation. In other words, we determine how the amplitude of QPO cycles changes as a function of the fractional position³ in the coherent interval. First, for each QPO cycle, we calculate the fractional position Ω of the cycle’s maximum in the coherent interval as

$$\Omega_{\max} = \frac{t_{\max} - T_{\text{start}}}{T_{\text{end}} - T_{\text{start}}} \quad (8)$$

where T_{start} and T_{end} are the start and end time of the coherent interval, respectively, and $t_{\max}$ corresponds to the time of the flux maximum of the considered QPO cycle. Subsequently, we divide the Ω -axis into fifteen bins between 0 and 1 and average the QPO amplitudes with a $\Omega_{\max}$ within the same Ω -bin. This procedure provides an estimate of the average coherent interval profile for each observation, and allows for the comparison of this profile between observations.

4.1.2 Broad energy band analysis

Secondly, we turn to the broad soft and hard energy bands ($\sim 2 - 7$ and $\sim 7 - 13$ keV) to investigate the evolution of the hardness of unfiltered lightcurves throughout coherent intervals. After finding the start and end time of each coherent interval in the filtered QPO lightcurve (T_{start} and T_{end}), we assign a fractional position to each individual time bin in similar fashion to Equation 8 (replacing $t_{\max}$ by the time of the considered timebin). For each timebin, we also calculate the hardness as the fraction of hard over soft counts in the *unfiltered* lightcurve. Finally, we again bin the fractional positions into ten bins, and average the hardness measurements in each bin. This approach

³ See Section 2.1 for the introduction and definition of the fractional position.

provides us with a profile of the unfiltered hardness within the coherent intervals. Although we plan on implementing full coherent-interval resolved spectroscopy, this hardness measurement provides the first crude estimate of the presence of spectral evolution on the coherence timescale.

4.1.3 *Narrow energy band analysis*

Finally, we also calculate coherent-interval resolved lag-energy spectra for the observations with binned data mode containing sixteen energy bands between RXTE absolute channels 0 to 35 (approximately $\sim 2 - 13$ keV). In other words, at different positions throughout the coherent interval, we calculate the phase lag between narrow energy bands and plot it as a function of energy. This analysis is nearly the same as present in Chapter 2, with only two important differences:

- (i) Instead of calculating lags between a broad soft and hard band, we now calculate lags between a fixed narrow, soft reference band and narrow bands of increasing energy. Thus, we obtain not one but multiple relations per observation showing the evolution of phase lags throughout a coherent interval. Since coherent-interval resolving reduces the number of counts and thus data quality in the individual lag-energy spectra, we combine the two softest narrow energy bands into one soft reference band.
- (ii) Instead of dividing the coherent interval into ten bins in fractional position, we select only four bins. This binning allows for an easier comparison between different fractional positions, and results in higher data quality in the four resulting lag-energy spectra, as more lag estimates are averaged per bin.

All other steps in the analysis are the same: the optimal filtering of the QPO, the determination of the coherent interval start and end times using the full energy band, the assignment of segment positions within the coherent interval, and the calculation of time and phase lags using the CCF. As we have seen that the lag evolution depends on the fractional position in the coherent interval, we do not compare the lag-energy spectra between absolute positions.

All in all, this coherent-interval resolved analysis returns three or four data products per observation: the PSD of coherent-interval lengths, the QPO amplitude profile within the coherent interval, and the evolution of hardness as a function of fractional position. For observations with 16 energy bands between the RXTE absolute channels 0 and 35 we calculate the coherent-interval resolved lag-energy spectra as the fourth data product. This criterion of 16 energy bands is met by 10 observations.

4.2 RESULTS

4.2.1 *Basic coherent interval properties*

In Figure 32, we show three representative examples of the PSD of coherent interval lengths. The three PSDs shown correspond to three different observations, spanning the full range of QPO frequencies in our sample. Interestingly, the shape of the PSDs is consistent with a constant power and is strikingly similar in all cases. In other words, at least on the scale of 64

Figure 32.: Representative power-spectral density of coherent interval lengths for three observations. The QPO frequency ν_0 of the observations is listed in the upper left. The frequency on the horizontal axis does not have a unit, as the sequence of coherent interval lengths is not a time series.

coherent intervals, their lengths appear to be uncorrelated. If the coherent intervals correspond to the setup and decay of differential precession, the simplest interpretation of this result is that the individual intervals occur independently of one another⁴.

The increase in power with decreasing QPO frequency is the result of not explicitly normalising the PSDs: as we have seen in Chapter 2, the coherent intervals length in terms of QPO cycles correlates with the observation's Q value. Hence, coherent intervals at different QPO frequencies span roughly the same range in amount of QPO cycles (assuming similar Q factors). However, as the QPO cycle size equals approximately $1/\nu_0$, coherent intervals in low QPO frequency observations span a much greater range in absolute length. Normalising using the rms^2 -normalisation (Belloni & Hasinger, 1990, see also Appendix A) would correct for this effect. However, as we do not normalise, this difference is clearly visible in the larger power in lower QPO frequency observations.

The average QPO amplitude profile as a function of fractional position Ω in the coherent interval is shown in Figure 33 for six representative observations. For each observation, we plot all amplitudes rescaled by the maximum amplitude A_{max} , in order to also show the relative increase in amplitude. The overall shape of the amplitude profiles appears to be consistent between observations, without notable features. Furthermore, the unscaled amplitude profiles, not shown here, show a decrease in QPO amplitude with increasing QPO frequency. This outcome is unsurprising as the QPO rms-amplitude is also known to be lower for higher QPO frequencies.

⁴ Other effects resulting in a similar power spectral shape, such as deadtime effects, could play a role as well; however, investigating that possibility is beyond the scope of this thesis.

Figure 33.: The average QPO amplitude profile within a coherent interval for a range of QPO frequencies. Ω represents the fractional position in the coherent interval. The amplitudes have been rescaled by the maximum amplitude $A_{\max}$ for each observation separately. Although the featureless shape appears to be the same between different observations, the relative size of amplitude variations increases with QPO frequency.

Interestingly, the rescaled amplitude profiles reveal that high QPO frequency observations show relatively larger changes in amplitude throughout the coherent interval than low QPO frequency observations. The QPOs in the plotted observations have comparable Q-values, so this difference does not originate from differences in the average amount of QPO cycles per coherent interval. These differences in relative amplitude variations could also be due to the different coherence timescales *in seconds* between observations. However, this seems unlikely, as the coherence timescale in seconds changes more than a factor of ten between the lowest and highest QPO frequencies, while the differences in relative amplitude variations are much smaller.

4.2.2 Hardness modulations

In Figure 34, we plot the relative hardness evolution within the coherent interval for the same three observations as plotted in Figure 32. The hardness H , calculated as the ratio between the broad hard ($\sim 7\text{--}13$ keV) and soft ($\sim 2\text{--}7$ keV) band flux, has been rescaled by its mean value $\langle H \rangle$ in order to compare the relative changes between observations. These relative changes amount to maximally $\sim 2\%$, but are significant due to the much smaller uncertainties on the measured hardnesses. For comparison, we also show the hardness of a representative filtered QPO lightcurve. This reveals no modulation throughout the coherent interval. The mean hardnesses are listed in the bottom left corner; as expected in the truncated disk model, the mean hardness decreases with increasing QPO frequency, as this corresponds to a smaller Comptonizing region.

Interestingly, the hardness varies systematically throughout the coherent interval with similar relative amplitudes in all observations (including those not shown). The location of the maximum hardness in the coherent interval appears to shift cyclically as a function of QPO frequency⁵. Such ubiquitous hardness evolution does not automatically appear in the differential precession model: if the hardness variations are caused by different precessing regions dominating the flux at different fractional positions, the relative size of variations would not be the same in all observations. Instead, in that scenario, no hardness evolution would be expected around 2 Hz: at this frequency, the spectrum proposed by the toy-model is the same at all radii, as no frequency difference is observed (see Section 2.3.4). Hence, the observed hardness would not change due to changes in the precessing region dominating the flux.

As discussed above, the hardness variations in all observations are of similar size, despite wildly different QPO properties; at low QPO frequency, the lightcurve's variability is dominated by the QPO, while towards higher frequencies, the QPO's relative contribution decreases. This indicates that the hardness variations are not directly caused by the QPO, but that instead X-ray emission not attributable to the QPO is aware of the presence of the coherent intervals as well. We discuss this surprising connection between QPO and non-QPO properties in more detail in Section 4.3.

Figure 34.: The unfiltered broad-band hardness H as a function of fractional position Ω , rescaled by the average hardness $\langle H \rangle$, for three representative observations. The hardness clearly evolves throughout the coherent interval, with the phase of the evolution within the coherent interval changing between observations. For comparison, we also show the hardness of the filtered QPO lightcurve, which shows no modulation.

Figure 35.: Coherent-interval resolved lag-energy spectrum of a 7.11 Hz QPO. Different colours/markers indicate different fractional positions within the coherent interval.

Figure 36.: Same as Figure 35 for a 2.46 Hz QPO.

Figure 37.: Same as Figure 35 for a 0.69 Hz QPO.

Figure 38.: Same as Figure 35 for a 7.94 Hz QPO.

4.2.3 Coherent-interval resolved lag-energy spectra

In Figures 35 to 37, we show representative coherent-interval resolved lag-energy spectra for the three observations shown in previous figures. As the phase lag evolution is most significant at the highest QPO frequency differences, and thus highest QPO frequencies, we show an extra example at a QPO frequency of 7.94 Hz in Figure 38. Each colour/marker represents one of the four bins in fractional position, shown in the legend by each bin's centre position. Each vertical slice of four lag measurements basically corresponds to a plot of phase lag evolution as a function of fractional position (i.e. to panel 1 in Figure 15 or 16 with four instead of ten datapoints) at the corresponding energy.

The overall, unresolved shape of all lag-energy spectra is as expected: at high QPO frequencies (> 2 Hz), the QPO lags become softer as a function of *energy*, while at low QPO frequencies (< 2 Hz) the QPO lags become increasingly harder as a function of energy. In Figure 36), where the QPO frequency is just above 2 Hz, the QPO lag barely changes as a function of energy. As we calculate the lags between narrow energy bands instead of a broad soft and hard band, the signal-to-noise in each lag measurement decreases significantly. This results in the relatively large uncertainties, especially seen in the high QPO frequency observations at low energies.

More interesting are of course changes in shape as a function of fractional position. In Chapter 2, we find that the lag between a broad hard and soft band changes systematically within the coherent interval, consistent with the observed frequency differences; below 2 Hz, the lag becomes increasingly hard, while above 2 Hz, the lag grows softer during the coherent interval. Such lag evolution should also be visible in the lag-energy spectra.

At low and middle QPO frequency (Figures 37 and 36), such changes in the shape of the lag-energy spectrum are marginal and not systematic. This

⁵ This holds in the full sample of observations, not only the three shown here.

lack of clear changes is expected as the broad-band lag evolution in both cases is small as well. Thus, coherent-interval resolving the data reduces the data quality too much to detect any systematic changes in the lag-energy spectra. In Figure 37, the lags appears to be softer at the end of the coherent interval ($\Omega = 7/8$), opposite to the expected lag evolution. This might be related to the turnover in broad-band lags observed at the end of the coherent interval in many observations.

Evolution in the shape of the lag-energy spectra is more apparent at high QPO frequencies (especially in Figure 38), despite the large uncertainties at low energies. At these low energies, the values of the QPO lag are consistent between the different fractional positions and show little variation as a function of energy. At higher energies however (above ~ 5 -6 keV), the shape changes clearly between different fractional positions: the uncertainties decrease significantly and a break in the lag-energy starts to appear towards higher fractional positions. Since these shape changes occur only in the hard energy bands, dominated by emission from the inner flow, this implies that the break could be directly related to differential precession.

4.3 DISCUSSION

In the previous Section, we presented four main results on the basic properties of coherent intervals and spectral variations within these intervals:

- (i) the number of QPO cycles in each coherent interval is independent of the preceding intervals,
- (ii) while the QPO amplitude profile of coherent intervals appears similar in all observations, the relative change in QPO amplitude increases as a function of QPO frequency,
- (iii) the (unfiltered) hardness varies throughout the coherent interval in all observations, and
- (iv) at high QPO frequency, the shape of the lag-energy spectrum changes significantly throughout the coherent interval, possibly connecting the lag-energy spectral break to differential precession.

In this Section, we will discuss these results in the framework of two general questions: what is the nature of the coherent intervals, especially in relation to the QPO origin? And secondly, what information do the coherent-interval resolved lag-energy spectra provide on the origin of the phase lags, and of the QPO itself?

4.3.1 Coherent intervals: connecting QPO and BBN?

Our interpretation of the coherent intervals in the differential precession model assumes that these intervals are only connected to the QPO: the amplitude rises and falls because the differential precession sets up and quickly decoheres. As expected, these amplitude variations only become visible after removing the BBN by optimal filtering. After all, the BBN is thought not to be caused directly by this process (see Chapter 1 or e.g. [Arévalo & Uttley 2006](#)). Furthermore, the similarity of the amplitude profiles in different observations is consistent with the differential precession interpretation: the setup and decay of precession occurs independently of the presence or absence of observed QPO frequency differences. Finally, in this simple

interpretation, the successive coherent intervals could be independent excitations of the same process. This notion of independent coherent intervals is observed in the power spectrum of coherent interval lengths, which is consistent with a white noise time series. So how do our other results tie into this? Do these results support the notion that the coherent intervals are purely related to the QPO, or is the picture more complicated?

Interestingly, we find that the hardness of the *unfiltered* lightcurve varies systematically as a function of fractional position in the coherent interval. The relative hardness variations are comparable in all observations ($\sim 2\%$), although the location of the peak hardness shifts as a function of QPO frequency. This result is most remarkable because the hardness is not inherently connected to the QPO: calculating the ratio of the unfiltered soft and hard band flux largely removes such correlated variability (although not always fully, as we will discuss below). Moreover, the similarity of the relative variations in hardness between observations, despite significant changes in QPO properties (rms amplitude, frequency), also indicates that the hardness evolution is not connected to the QPO. Finally, the hardness is not found to change on the timescale of a single QPO cycle, but multiple successive QPO cycles. These findings raise the question how such a general property of the observed lightcurve can vary on a timescale that only shows up after removing all non-QPO variability. Can we incorporate this result in the picture of the setup and decay of differential precession?

As stated in Section 4.2.2, the observed hardness variations are not likely to originate from the differentially precessing region itself. Such hardness variations could be explained if they were observed in the filtered (QPO) flux: the differential nature of the precession frequency in the inner flow could result in a variable rate of decay of the precession. In other words, if the gradient in the precession rate as a function of radius is not constant, the speed of decoherence of the QPO due to this gradient might depend on radius as well. Coupled with a varying emitted spectrum as a function of radius, this scenario could lead to hardness fluctuations in the QPO: if the innermost radii decohere quickest and emit the hardest spectrum, the hardness would decrease towards the end of the coherent interval. However, we do not observe hardness variations in the filtered QPO lightcurve. Furthermore, this interpretation does not hold for the unfiltered hardness for two reasons: first, the observed relative variations in hardness are always comparable in size, independent of the QPO rms amplitude, i.e. the contribution of the QPO to the lightcurve. Secondly, the differential precession model requires a constant emitted spectrum as a function of radius around 2 Hz. This would lead to a non-variable hardness at 2 Hz as well, contrary to our observed results.

The presence of coherent interval properties in the BBN variability is, in itself, not surprising. [Van der Klis \(1989\)](#) already discusses how the extended (Lorentzian) power-spectral shape of the QPO can correspond to amplitude variations on the coherence timescale. [Van der Klis \(1989\)](#) then also states that these amplitude modulations can contribute to a noise component in the PSD as well; for instance, an exponentially decaying amplitude results in additional red noise in the PSD. Similarly, it turns out that the enveloping structure of coherent intervals can result in a flat-noise component, as we show below.

This statement follows from estimating the QPO amplitude as a function of time at a 1/128 s time resolution, by calculating a cubic spline connecting the measured maxima in the filtered QPO flux. The resulting QPO

Figure 39.: *Upper*: power-spectral density of the QPO amplitude time series $A(t)$ and the multiplication of this time series by a cosine function, $A(t) \cos(2\pi\nu_0 t)$. *Lower*: segment of the time series $A(t)$ and $A(t) \cos(2\pi\nu_0 t)$ in black and red, respectively. The amplitudes and QPO frequency are calculated from the full band lightcurve of observation 10408-01-22-00 ($\nu_0 = 3.48$ Hz).

amplitude time series (which does not contain the QPO cycles themselves), is shown in black the lower panel of Figure 39. The corresponding power spectrum is shown in the Figure's upper panel – a clear flat-top broad-band noise profile is visible in the black circles. Making the assumption that the QPO frequency does not drift significantly during an observation, we can subsequently model the QPO lightcurve $q(t)$ as

$$q(t) = A(t) \cos(2\pi\nu_0 t) \quad (9)$$

where $A(t)$ is the estimated amplitude time series and ν_0 is the QPO centroid frequency. The time series and power spectrum corresponding to this simple description of the QPO lightcurve is shown in Figure 39 as well, as the red line and triangles. A clear, extended QPO peak, on top of a broad noise component, appears in the power spectrum.

This simple demonstration shows that the coherent intervals need not only be present in the QPO variability, but also in what we classify as the BBN; this might (partially) account for the connection between QPO and non-QPO properties. The BBN is already known to be correlated to the

QPO as well (Heil et al., 2011), as can for instance be seen in the bispectrum of GRS 1915+105 (Maccarone et al., 2011). However, it is also important to note that obviously, the BBN can not fully be recreated by these amplitude modulations alone – the BBN is present independently of the QPO, sources not showing QPOs (Cyg X-1, AGN) still show BBN, and the BBN can be adequately modeled by the propagation of mass accretion rate fluctuations (see Chapter 1 and e.g. Lyubarskii 1997). Furthermore, this possible presence of QPO amplitude variation in the BBN does not physically explain why the hardness varies on the coherence timescale. Moreover, the issue detailed above, that the maximum size of relative hardness variations is independent of the properties of the QPO, is not solved.

The conclusion of these reflections on the fluctuation of both QPO and non-QPO properties on the coherence timescale is simple: currently, we do not fully understand what is causing this connection between non-QPO properties and the coherence timescale. As this investigation is a work in progress, we aim to further investigate the origin of this unexpected connection after the completion of this thesis. Full coherent-interval resolved spectroscopy will be a vital step in this process: a full spectral analysis as a function of fractional position can shed light on the origin of the observed hardness variations.

4.3.2 *The evolution of lag-energy spectra*

As stated in this Chapter’s introduction, our coherent-interval resolved analysis does not only provide insight into the nature of coherent intervals themselves, but also provides an in-depth look at the phase lag evolution reported in Chapter 2. The coherent-interval resolved lag-energy spectra show no clear evolution in shape at low and intermediate QPO frequencies; in those observations, the frequency differences are small and thus our method is not accurate enough to distinguish these possible changes. However, at high QPO frequency, the lag-energy spectral shape is actually observed to change, despite the large lag uncertainties at low energies. Most notably, the break appears only towards high fractional positions. This evolution of shape has several implications for our interpretation of the QPO phase lags offered in the previous two Chapters.

First of all, the prominent evolution of phase lags at the highest energies, suggests that the observed lag evolution (Chapter 2) is indeed coupled to the inner flow. As optimal filtering a power spectrum only separates the QPO power from the BBN power, but not the QPO lags from the BBN lags, the lag-evolution’s origin can not completely be inferred from the evolution of the broad-band lag alone. Similarly, the lack of clear shape changes at low energy could indicate that these lags are associated with a different mechanism, for instance related to the disk instead of the inner flow.

Secondly, this result emphasizes the possibility of a prominent role of differential precession in the origin of the observed QPO lags. The coherent-interval resolved lag-energy spectra reveal that the break, that is known from the time-averaged lag-energy spectrum, is not always present. In fact, it only appears towards the middle and end of the coherent interval. At that point, the various precessing radii in the differential precession model have already build up a significant phase lag. This is also consistent with the zero QPO phase lag around 2 Hz: in Figures 35 and 38, we can see that the phase lags at low Ω hover around zero phase lag. In the absence of further

observed phase lag evolution, due to a constant spectral shape as a function of radius, this lag will simply remain around zero.

Inferring differential precession as the origin of the QPO phase lag conflicts with our interpretation of the Type-C QPO phase lag inclination dependence in Chapter 3. In that interpretation, we propose a pivoting power law during QPO cycles as a possible origin for the QPO lag, instead of the average lag between regions precessing at different rates. However, the behaviour of average QPO lags as a function of QPO frequency is known to be peculiar in GRS 1915+105: only this source shows a significant hard lag below a QPO frequency of 2 Hz and a log-linear relation between QPO lag and frequency (Pahari et al., 2013; Qu et al., 2010; Reig et al., 2000). As noted in Section 3.3.4.3, this peculiar behaviour hints towards a more complicated origin of QPO lags in GRS 1915+105. It would be interesting to test this idea by calculating coherent-interval resolved lag-energy spectra for another source showing QPO frequency differences, such as XTE J1550-564. However, this calculation requires a very high data quality, which is often not available without the large count rates of GRS 1915+105.

4.4 CONCLUSION AND FUTURE OUTLOOK

In this Chapter, we have presented the results of a first coherent-interval resolved analysis of the QPO in GRS 1915+105. We find that the length of coherent intervals is independent of previous intervals, and that the amplitude profile is similar in all observations. The relative increase in QPO amplitude however, is larger at higher QPO frequencies. Unexpectedly, we find that the hardness of the full band lightcurve, which is not directly determined by the QPO, varies on the coherence timescale as well. Although difficult to interpret, this result presents an interesting link between QPO and non-QPO properties. Finally, we find that for high QPO frequency observations, the well-known break in the lag-energy spectrum only appears towards high fractional positions. We interpret this result as additional evidence for differential precession as a possible origin of the Type-C QPO.

The detected evolution of hardness and lag-energy spectral shape show the promise of detailed coherent-interval resolved spectroscopy. This approach could for instance reveal the origin of the hardness evolution, and show whether this origin is the same for all observations, despite differences in the fractional positions of maximum hardness. We have started the process of implementing this spectroscopy and aim to finish it soon after the completion of the thesis.

SUMMARY

Black hole X-ray binaries (BHXRBS), which consist of a stellar-mass black hole accreting material from an orbiting star, emit bright and highly variable X-ray emission. Most notably, the X-ray flux shows *quasi-periodic oscillations* (QPOs), known to originate from the innermost accretion flow. Hence, QPOs are a powerful tracer of the behaviour of matter in the strongest gravitational regimes. However, their physical origin is ill-understood. A broad range of models describing the origin of the QPO exists, which can be divided into two broad categories: *geometric models*, attributing the X-ray oscillations to systematic changes in the accretion geometry, or *intrinsic models*, invoking variations in intrinsic properties of the accretion flow. While currently a geometric model, the Lense-Thirring precession model by [Ingram et al. \(2009\)](#), appears most feasible, this debate is not settled yet. In this thesis, I present the result of two distinct, but complementary approaches to shedding light on this problem: an in-depth investigation of the QPO in a single BHXRBS, and a sample study comparing QPOs in BHXRBS of different inclination.

In the former approach, which appeared as [Van den Eijnden et al. \(2016\)](#), I investigate the peculiar energy-dependent QPO frequency in BHXRBS GRS 1915+105 using archival RXTE observations. The QPO frequency in this source differs between lightcurves of different X-ray energies, while the Lense-Thirring precession model predicts only a single frequency. In this thesis, I present a newly developed method to remove aperiodic variability noise and track phase lags between a soft and hard QPO lightcurve. This allows me to test whether the different QPO frequencies are an intrinsic property of the QPO mechanism, unaccounted for in current models. I find that the QPO amplitude repeatedly rises and falls over an interval of 5–10 QPO cycles. Within these amplitude envelopes, dubbed *coherent intervals*, the lag between the hard and soft band increases. This lag evolution implies that the different observed frequencies are indeed a fundamental property of the QPO.

Based on this first result, I formulate a geometrical toy model for the QPO origin, based on the Lense-Thirring precession model: differential precession of the inner accretion flow. In this interpretation, the precession rate of the flow decreases as a function of radius. As a result, a lag builds up between different radii in the inner flow, destabilising the precession. With the precession effectively decohering itself, I interpret the observed coherent intervals as the setup and decay of this differential precession. However, we do not observe different radii in the accretion flow, as BHXRBS are not spatially resolvable. Instead, I analyse various X-ray energy bands. Hence, the actual observed QPO lags and frequencies as a function of energy, will depend on the emission processes in the inner accretion flow. Thus, it is possible that a single source (GRS 1915+105) can show different energy dependencies of the QPO frequency in different observations.

The second approach, the results of which have also been submitted to MNRAS, consists of a detailed comparison between QPO phase lags in high and low inclination sources. I calculate the phase lag between a broad soft and hard energy band at the QPO frequency in 504 archival RXTE observations of 15 BHXRBS. I find that the sign of the lag associated with the Type-C QPO depends clearly on the source inclination: edge-on sources show soft

lags, while face-on sources show hard lags. I do not find a significant inclination dependence in the Type-B QPO lags, adding to the evidence that both QPO types originate from different mechanisms. In the same set of observations, I search for energy dependences of the Type-C QPO frequency similar to that in GRS 1915+105. This search strongly suggests that such differences in QPO frequency are only observed in edge-on sources. Finally, I classify XTE J1859+226 and MAXI J1543-564, both sources of undetermined inclination, as edge-on sources.

The inclination dependence of the Type-C QPO lag is strong evidence that the QPO arises from a geometric origin. The observed dependence is consistent with Lense-Thirring precession (either solid-body like or differential): differences in dominant relativistic effects between an edge-on or face-on view of the BHXRBB, lead to a different phase of maximum X-ray flux in the precession cycle. This translates into a different sign of the lag: hard photons lagging soft photons in face-on sources, and vice-versa in edge-on sources. At low QPO frequency, corresponding to a large inner flow, these relativistic effects play a smaller role – differences between edge and face-on sources should thus decrease. Interestingly, I observe exactly that.

The final chapter of this thesis contains a preliminary look at the variation of several spectral-timing properties of GRS 1915+105 during coherent intervals. Most notably, both the hardness of the full lightcurve and the shape of the QPO lag-energy spectrum systematically vary systematically as a function of position within the coherent interval. The former implies an interesting connection between the QPO and non-QPO variability properties of the BHXRBB, which I can currently not fully explain. The evolution of the lag-energy spectrum indicates that its well-known break around ~ 5 – 6 keV could result from differential precession. This analysis constitutes a stepping stone to full coherent-interval resolved spectroscopy, which I aim to develop in the near future.

The results presented in this thesis all support a geometric origin of the Type-C QPO. Multiple other recent analyses, based on either phase-resolved spectroscopy of the QPO (Ingram & Van der Klis, 2015; Ingram et al., 2016; Stevens & Uttley, 2016) or a comparison of both QPO and spectral properties between BHXRBBs of different inclinations (Heil et al., 2015b; Motta et al., 2015), point towards the same conclusion. In other words, the debate about the fundamental nature of QPOs as either a geometric or intrinsic accretion phenomenon, appears to have been settled in favour of geometric models through this surge of recent observational evidence.

POPULAR SCIENTIFIC SUMMARY

Note: this popular summary will appear in a slightly altered version in Amsterdam Science, a magazine about research performed at the University of Amsterdam, aimed at scientists of all disciplines (see www.amsterdamscience.org).

How big is a black hole in the sky? Is it the size of the Moon, a star, a grain of sand? The answer is troublesome for astronomers: its about the size of a virus, in New York, as seen from Amsterdam. Troublesome, because the extreme properties of black holes make their surroundings an astrophysical playground. So how to study a region appearing so small? Surprisingly, music guides the way.

This thesis describes a study of the way matter behaves when it is subjected to the boldest predictions of general relativity: Einstein's ground breaking description of gravity. In everyday situations, gravity's force is relatively unremarkable not being enough to make a small fridge magnet fall to the ground. However, when this force becomes extremely strong, matter is expected to behave differently than on Earth. As such strong gravitational conditions are impossible for humans to recreate, we search the universe for locations where they occur naturally. As it turns out, the neighbourhood of a black hole is the perfect place to look.

In essence, a black hole is nothing but an object with enough gravitational pull to prevent even light from escaping. As a result, it is impossible to visually see a black hole itself. However, when a normal star orbits around a black hole its outer layers are stripped off and fall towards the black hole. Just as a spinning top wants to keep spinning, this gas keeps rotating and forms a disk a giant whirlpool with the black hole as its drain. However, although these whirlpools themselves are well understood, it is still debated how the disk material behaves as it disappears down the drain where gravity is at its strongest. Zooming in on this process is, due to the size issue alluded to above, an insurmountable task. Therefore to understand what happens right before the gas falls into the black hole, we need another approach.

Historically, star-swallowing black holes were studied using two distinct approaches: spectral and timing analysis. In the former, astronomers investigate the observed spectrum: the amount of light received of each wavelength, or colour if you will. These different colours originate from different physical processes, just as the light from a laser and a campfire also look very different. Hence, the observed colour combination is a handle on the relevant processes occurring as matter approaches the black hole: for instance, closer to the black hole the material is hotter and emits more light at the blue end of spectrum. Conceptually, these spectra are actually quite similar to musical chords, just made up of colours instead of tones. Alternatively, timing analysis studies variations in the brightness of the gas travelling towards the black hole. Hence, it tracks how the system changes on slow, month-to-month to swift, sub-second timescales. These brightness fluctuations can be considered represent the rhythm of the whirlpool. And while chords and rhythm might be interesting on their own, it is their com-

bination that elevates music to a higher level. Similarly, combining these spectral and timing techniques can elevate our knowledge of the behaviour of matter just before it is swallowed by a black hole.

Often, the brightness of the gas disk contains fluctuations that are close to, but not perfectly periodic – like a drummer continually just missing the beat. By measuring which colours fluctuate most extremely, it becomes apparent that this modulation originates from the matter closest to the black hole. Hence, these fluctuations provide insight into the matter in extreme gravitational conditions whose behaviour we want to understand. The variations are thought to be created by precession of this very inner part of the whirlpool: as it rotates, it wobbles, causing the brightness observed here on earth to fluctuate.

In my master's project, I have compared the speed of these precessional wobbles, analysing different colours for one specific black hole. This reveals that the material does not precess in unison. Instead, gas close to the black hole precesses faster than gas further out, and as a result the innermost regions get spun up. But just as continuing to wind up a spring, this increasingly spun-up state is unstable and eventually stops the precession. Afterwards, this entire process repeats. So, instead of peacefully wobbling continuously, as was thought previously, the rotating gas appears to experience precession in violent spasms just before it falls down the black hole.

In addition, I have also compared another property of the almost-periodic brightness fluctuations, between many different black holes. The fluctuations do not always arrive at the same time in all colours – sometimes, there can be a delay: an echo in the light. I have compared these echos between black holes where we see the whirlpool from the side, and black holes where we see it from the top. It turns out that the order in which bluer and redder fluctuations arrive at our telescope, flips between these viewing points. This reinforces the idea that the fluctuations are indeed caused by the wobbling of the whirlpool. It is through combining the rhythm and melody of the emission from the whirlpool of gas that we are able to unravel how matter behaves under such extreme gravitational conditions.

ACKNOWLEDGEMENTS

On the advice of Adam, many parts of this thesis are written in the first person. However, each of those I's actually refers to a large group of people without whom I would not have finished this work. Of course, first of all I want to thank my supervisors Phil and Adam. Your enthusiasm about basically anything we discussed has been incredibly contagious. Thank you, for all the meetings where I may not have understood every detail at once, but have many times felt like new science and insight were being created. Thank you, for your trust and guidance when writing the papers and preparing presentations. And, of course, for all advice and conversations during my PhD-hunt. I look forward to continuing collaborating with you in the coming years. And Phil, when our second paper is accepted, I'll bring authentic Limburgse riestevloai.

I'd also like to thank Michiel van der Klis, for our conversations about unexpected results and your tips about interpreting them, and for reading my thesis. The same goes for the rest of the X-ray Timing group - thank you, especially for all the feedback on my PhD application presentation. Furthermore, I want to thank Patrick Decowski, for all your advice on my PhD applications, and Volkert van der Willigen, for granting me the opportunity to travel to the United States for my first conference.

The past thirteen months, I have probably spend more waking hours in the API MSc room than in my student room. But due to my fellow students, this is definitely not as bad as it sounds. Thanks to all of you for the great atmosphere, the silly drawings on the white board, the countless pizza or kapsalon-dinners, the help by listening to me complaining about unwilling code, your tolerance towards my awfully ugly coffee mug, and the way-to-many late evenings working and writing. Special thanks to David, for your help with spectral-timing struggles, your aesthetic tips on my presentations, your forgiveness towards my messy desk, and for letting me use your code so that I did not have to reduce all data by hand. Also special thanks to Ziggy and Coen, together with David my partners in crime during the seemingly endless PhD-application season.

Thank you, Nina, for always making me smile, lifting me out of astronomy's bubble, and always distracting me with your latest sog-material. Thank you, Ruben, for all the hamburgers and frisbees to distract me from thesis stress. Thanks to all of you, Heftige Gasten, for the unrivaled friendship after leaving Maastricht. Thank you, Max, for five years of friendship, for all our bets on grades, and for indirectly convincing me that astronomy is objectively better than theoretical physics.

Als laatste, pap, mam, Simone, Saskia en Jesper: over de afgelopen twee jaar is teveel te zeggen om hier op te schrijven. Dat zeg ik dan ook liever dan dat ik het opschrijf. Maar weet dat zonder jullie steun, kracht en liefde, me dit nooit gelukt was. Jullie zijn de reuzen op wiens schouders ik sta.

July 2016, Amsterdam

BIBLIOGRAPHY

- Arévalo P., Uttley P., 2006, *MNRAS*, 367, 801
- Axelsson M., Hjalmarsdotter L., Done C., 2013, *MNRAS*, 431, 1987
- Axelsson M., Done C., Hjalmarsdotter L., 2014, *MNRAS*, 438, 657
- Balbus S. A., Hawley J. F., 1991, *ApJ*, 376, 214
- Balbus S. A., Hawley J. F., 1992, *ApJ*, 400, 610
- Bardeen J. M., Petterson J. A., 1975, *ApJ*, 195, L65
- Belloni T., 2010, in Belloni T., ed., *Lecture Notes in Physics*, Berlin Springer Verlag Vol. 794, The Jet Paradigm. pp 53–84
- Belloni T. M., Altamirano D., 2013, *MNRAS*, 432, 10
- Belloni T., Hasinger G., 1990, *A&A*, 230, 103
- Belloni T., Mendez M., King A. R., Van der Klis M., Van Paradijs J., 1997a, *ApJ*, 479, L145
- Belloni T., Mendez M., King A. R., Van der Klis M., Van Paradijs J., 1997b, *ApJ*, 488, L109
- Belloni T., Klein-Wolt M., Mendez M., Van der Klis M., Van Paradijs J., 2000, *A&A*, 355, 271
- Belloni T., Sanna A., Mendez M., 2012, *MNRAS*, 426, 1701
- Cabanac C., Henri G., Petrucci P.-O., Malzac J., Ferreira J., Belloni T. M., 2010, *MNRAS*, 404, 738
- Casella P., Belloni T., Homan J., Stella L., 2004, *A&A*, 426, 587
- Casella P., Belloni T., Stella L., 2005, *ApJ*, 629, 403
- Castro-Tirado A. J., Brandt S., Lund N., Lapshov I., Sunyaev R. A., Shlyapnikov A. A., Guziy S., Pavlenko E. P., 1992, *ApJS*, 92, 469
- Chakrabarti S. K., Molteni D., 1993, *ApJ*, 417, 671
- Corral-Santana J. M., Casares J., Shahbaz T., Zurita C., Martinez-Pais I. G., Rodríguez-Gil P., 2011, *MNRAS*, 413, L15
- Done C., Madejski G. M., Mushotzky R. F., Turner T. J., Koyama K., Kunieda H., 1992, *ApJ*, 400, 138
- Done C., Gierlinski M., Kubota A., 2007, *A&ARv*, 15, 1
- Esin A. A., McClintock J. E., Narayan R., 1997, *ApJ*, 489, 865
- Fabian A. C., et al., 2009, *Nature*, 459, 540
- Fabian A. C., Lohfink A., Kara E., Parker M. L., Vasudevan R., Reynolds C. S., 2015, *MNRAS*, 451, 4375

- Fender R. P., Homan J., Belloni T. M., 2009, *MNRAS*, 396, 1370
- Fragile C., Blaes O. M., Anninos P., Salmonson J. D., 2007, *ApJ*, 668, 417
- Frank J., King A., Raine D., 2002, *Accretion power in astrophysics*, 3 edn. Cambridge University Press
- Galeev A. A., Rosner R., Vaiana G. S., 1979, *ApJ*, 229, 318
- Garcia J. A., Steiner J. F., McClintock J. E., Remillard R. A., Grinberg V., Dauser T., 2015, *ApJ*, 813, 84
- Gendreau K. C., Arzoumanian Z., Okajima T., 2012, in *SPIE Astronomical Telescopes + Instrumentation*.
- Giacconi R., Gursky H., Paolini F. R., Rossi B. B., 1962, *Phys. Rev. Lett.*, 9, 439
- Gilfanov M., 2010, in Belloni T., ed., *Lecture Notes in Physics*, Berlin Springer Verlag Vol. 794, The Jet Paradigm. p. 17
- Greene J., Bailyn C. D., Orosz J. A., 2001, *ApJ*, 554, 1290
- Hawley J. F., Balbus S. A., 1991, *ApJ*, 376, 223
- Hawley J. F., Balbus S. A., 1992, *ApJ*, 400, 595
- Heil L. M., Vaughan S., Uttley P., 2011, *MNRAS*, 411, L66
- Heil L. M., Uttley P., Klein-Wolt M., 2015a, *MNRAS*, 448, 3339
- Heil L. M., Uttley P., Klein-Wolt M., 2015b, *MNRAS*, 448, 3348
- Ingram A., Done C., 2011, *MNRAS*, 415, 2323
- Ingram A., Done C., 2012a, *MNRAS*, 419, 2369
- Ingram A., Done C., 2012b, *MNRAS*, 427, 934
- Ingram A., Van der Klis M., 2013, *MNRAS*, 434, 1476
- Ingram A., Van der Klis M., 2015, *MNRAS*, 446, 3516
- Ingram A., Done C., Fragile P. C., 2009, *MNRAS*, 397, L101
- Ingram A., Van der Klis M., Middleton M., Done C., Altamirano D., Heil L., Uttley P., Axelsson M., 2016, *MNRAS*
- Kara E., Fabian A. C., Cackett E. M., Steiner J. F., Uttley P., Wilkins D. R., Zoghbi A., 2013, *MNRAS*, 428, 2795
- Kotov O., Churazov E., Gilfanov M., 2001, *MNRAS*, 327, 799
- Lachowicz P., Done C., 2010, *A&A*, 515, A65
- Lehr D. E., Wagoner R. V., Wilms J., 2000
- Li Z. B., Qu J. L., Song L. M., Ding G. Q., Zhang C. M., 2013a, *MNRAS*, 428, 1704
- Li Z. B., Zhang S., Qu J. L., Gao H. Q., Zhao H. H., Huang C. P., M. S. L., 2013b, *MNRAS*, 433, 412
- Lyubarskii Y. E., 1997, *MNRAS*, 292, 679

- Maccarone T., Uttley P., Van der Klis M., Wijnands R. A. D., Coppi P. S., 2011, *MNRAS*, 413, 1819
- Markoff S., Nowak M. A., Wilms J., 2005, *ApJ*, 635, 1203
- Merloni A., Vietri M., Stella L., Bini D., 1999, *MNRAS*, 304, 155
- Miller-Jones J. C. A., Jonker P. G., Ratti E. M., Torres M. A. P., Brocksopp C., Yang J., Morrell N. I., 2011, *MNRAS*, 415, 306
- Mirabel I. F., Rodriguez L. F., 1994, *Nature*, 371, 46
- Morgan E. H., Remillard R. A., Greiner J., 1997, *ApJ*, 482, 993
- Motta S. E., Homan J., Munoz Darias T., Casella P., Belloni T. M., Hiemstra B., Mendez M., 2012, *MNRAS*, 427, 595
- Motta S. E., Belloni T. M., Stella L., Muñoz Darias T., Fender R., 2014, *MNRAS*, 437, 2554
- Motta S. E., Casella P., Henze M., Munoz-Darias T., Sanna A., Fender R., Belloni T., 2015, *MNRAS*, 447, 2059
- Munoz-Darias T., Casares J., Martinez-Pais I. G., 2008, *MNRAS*, 385, 2205
- Muñoz-Darias T., Coriat M., Plant D. S., Ponti G., Fender R. P., Dunn R. J. H., 2013, *MNRAS*, 432, 1330
- Neilsen J., Rahoui F., Homan J., Buxton M., 2016, *ApJ*, 822, 20
- Neustroev V. V., Veledina A., J. P., V. Z. S., Tsygankov S. S., G. S., E. K. J. J., 2014, *MNRAS*, 445, 2424
- Nixon C., King A., Price D., Frank J., 2012, *ApJ*, 757, L24
- Orosz J. A., et al., 2002, *ApJ*, 568, 845
- Orosz J. A., McClintock J. E., Remillard R. A., Corbel S., 2004, *ApJ*, 616, 376
- Orosz J. A., Steiner J. F., McClintock J. E., Torres M. A. P., Remillard R. A., Bailyn C. D., Miller J. M., 2011, *ApJ*, 730, 75
- Pahari M., Neilsen J., Yadav J., Misra R., Uttley P., 2013, *ApJ*, 778, 136
- Ponti G., Fender R. P., Begelman M. C., Dunn R. J. H., Neilsen J., Coriat M., 2012, *MNRAS*, 422, L11
- Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1997, *Numerical Recipes in Fortran 77 : the art of scientific computing*, 2nd edn. Press Syndicate of the University of Cambridge
- Qu J. L., Lu F. J., Lu Y., Song L. M., Zhang S., Ding G. Q., Wang J. M., 2010, *ApJ*, 710, 836
- Reid M. J., McClintock J. E., Steiner J. F., Steeghs D., Remillard R. A., Dhawan V., Narayan R., 2014, *ApJ*, 796, 2
- Reig P., Belloni T., Van der Klis M., Mendez M., Kylafis N. D., Ford E. C., 2000, *ApJ*, 541, 883
- Reig P., Belloni T., Van der Klis M., 2003, *A&A*, 412, 229

- Remillard R. A., Sobczak G. J., Muno M. P., McClintock J. E., 2002, *ApJ*, 564, 962
- Schmidt M., 1963, *Nature*, 197, 1040
- Schnittman J. D., Homan J., Miller J. M., 2006, *ApJ*, 642, 420
- Schwarzschild K., 1916, *Sitzungsberichte der Kniglich-Preussischen Akademie der Wissenschaften*, pp 189–196
- Shakura N. I., Sunyaev R. A., 1973, *A&A*, 24, 337
- Shaposhnikov N., Titarchuk L., 2009, *ApJ*, 699, 453
- Shipman H. L., 1975, *ApL*, 16, 9
- Sobolewska M. A., Zycki P. T., 2006, *MNRAS*, 370, 405
- Steiner J. F., McClintock J. E., Reid M. J., 2012, *ApJ*, 745, L7
- Stella L., Vietri M., 1998, *ApJ*, 492, L59
- Stella L., Vietri M., Morsink S. M., 1999, *ApJL*, 524, L63
- Stevens A. L., Uttley P., 2016, *MNRAS*, 460, 2796
- Su Y. H., Chou Y., Hu C. P., Yang T. C., 2015, *ApJ*, 815, 74
- T. Belloni T., Van der Klis M., Lewin W. H. G., Van Paradijn J., Dotani T., Mitsuda K., Miyamoto S., 1997, *A&A*, 322, 857
- Tagger M., Pellat R., 1999, *A&A*, 349, 1003
- Timmer J., Koenig M., 1995, *A&A*, 300, 707
- Uttley P., Klein-Wolt M., 2015, *MNRAS*, 451, 475
- Uttley P., Cackett E., Fabian A., Kara E., Wilkins D., 2014, *A&ARv*, 22, 1
- Van Oers P., et al., 2010, *MNRAS*, 409, 763
- Van den Eijnden J., Ingram A., Uttley P., 2016, *MNRAS*, 458, 3655
- Van der Klis M., 1989, *ARA&A*, 27, 517
- Van der Klis M., 1997, in Babu G. J., Feigelson E. D., eds, *Statistical Challenges in Modern Astronomy II*. p. 321 (arXiv:astro-ph/9704273)
- Van der Klis M., 2006, in Lewin W., Van der Klis M., eds, *Compact stellar X-ray sources*. Cambridge University Press, pp 39–112
- Vaughan B. A., Nowak M. A., 1997, *ApJL*, 474, L43
- Veledina A., Poutanen J., Ingram A., 2013, *ApJ*, 778, 165
- White N. E., Swank J. H., 1982, *ApJL*, 253, L61
- Wijnands R., Homan J., van der Klis M., 1999, *ApJL*, 526, L33
- Yan S. P., et al., 2012, *Ap&SS*, 337, 137
- Yan S. P., Ding G. Q., Wang N., Qu J. L., Song L. M., 2013, *MNRAS*, 434, 59
- Zdziarski A. A., Poutanen J., Mikolajewska J., Gierlinski M., Ebisawa K., Johnson W. N., 1998, *MNRAS*, 301, 435

SPECTRAL-TIMING METHODS

In this appendix, we introduce the mathematics and background of several basic (spectral-)timing methods used in the study of X-ray variability. The basic methods discussed in this section form the building blocks of our method (Chapter 2) and novel analyses by others (see e.g. [Fabian et al. 2009](#), [Ingram & Van der Klis 2015](#) or [Stevens & Uttley 2016](#) for other recent examples of applications). In this appendix, we first discuss the power spectral density, the cross spectrum and the cross-correlation function (CCF). Subsequently, we discuss more energy-resolved techniques, namely the lag energy spectrum and the rms spectrum. Finally, we introduce the [Timmer & Koenig \(1995\)](#) algorithm to simulate time series. The overview in this appendix is primarily based on the second section in [Uttley et al. \(2014\)](#), unless indicated otherwise.

A.1 A SINGLE TIME SERIES: POWER-SPECTRAL DENSITY

The most fundamental description of variability in a time series is given by the power-spectral density (PSD): the modulus squared of the Fourier transform of the time series (For example plots of the PSD: see the Figures in section 1.2). For an observed lightcurve f_j with index j referring to the corresponding discrete timebins t_j , the discrete Fourier transform F_n at frequency ν_n is numerically calculated as

$$F_n = \sum_{j=0}^{N-1} f_j e^{2\pi i n j / N} \quad (10)$$

The frequencies are given by $\nu_n = n / Ndt$, where dt is the time resolution of the transformed lightcurve and n runs from 0 to $N/2$ in integer steps (assuming N is even) with N the number of timebins. This implies that the frequency resolution $\Delta\nu$ is equal to the inverse of the length of the transformed lightcurve, and that the maximum frequency $\nu_{\max}$ (the so-called Nyquist frequency) is equal to $1/2dt$. The discrete Fourier transform also returns values at negative frequencies (up to $\nu_{\min} = -\nu_{\max} + \Delta\nu$ if N is even). These frequencies are not necessarily interesting from a physical point of view, although they do play an important role in simulating lightcurves with a given PSD ([Timmer & Koenig, 1995](#), see Section A.4 in this appendix).

As stated above, the PSD is estimated with the modulus squared of the Fourier transform F_n :

$$P_n = |F_n|^2 = F_n^* F_n \quad (11)$$

In this thesis (except for Section 4.2.1), the PSD estimate is normalised according to the so-called rms²-normalisation (Belloni & Hasinger, 1990):

$$P_n^{\text{normalised}} = \frac{2dt}{\langle f \rangle^2 N} P_n \quad (12)$$

As the normalisation factor is constant, this normalisation simply rescales the units of the estimated PSD to rms²Hz⁻¹ without adjusting the shape. Thus, integrating this normalised PSD estimate returns the square of the fractional root-mean squared variability of the variations in the integrated frequency range. This fractional rms corresponds to the fraction of the total flux that is modulated by variability in the integrated frequency range.

This method of estimating the PSD returns one specific random manifestation of the underlying PSD, which follows a $\chi^2_{\nu=2}$ distribution with the underlying PSD as its mean value at each frequency. As the standard deviation of this distribution equals the mean, a single PSD estimate is highly uncertain. Simply analysing a longer lightcurve will not improve these uncertainties, as this merely increases the frequency resolution. Instead, one can split a lightcurve into smaller segments, estimate the PSD for each, and average those together. In this case, the uncertainties will decrease by a factor of the square root of the number of segments. However, this does introduce a trade-off between the resolution in frequency and the uncertainty in estimated power.

A.2 TWO TIME SERIES: CROSS SPECTRA AND CCFs

The PSD describes the variability properties of a single time series. In order to compare two time series, in X-ray timing most probably two lightcurves in different energy bands, the cross spectrum and the CCF are powerful tools. Both methods allow one to measure the lag between the two lightcurves: the cross spectrum provides the phase lag at individual Fourier frequencies, while the CCF measures the time lag without selecting any timescales. These phase and time lags, measured using either of the two methods, are the most important observable analysed in this thesis.

The numerical cross spectrum between two time series f_j and g_j is defined as

$$C_n = F_n^* G_n \quad (13)$$

where F_n and G_n correspond to the Fourier transform of f_j and g_j calculated with Equation 10. The cross spectrum contains phase information since C_n , F_n and G_n are complex quantities. Writing the individual Fourier transforms as $F_n = R_{F,n} \exp(i\alpha)$ and $G_n = R_{G,n} \exp(i(\alpha + \phi_n))$, it is evident that the cross spectrum contains only the difference in complex phase

$$C_n = R_{F,n} R_{G,n} e^{i(-\alpha + \alpha + \phi_n)} = R_{F,n} R_{G,n} e^{i\phi_n} \quad (14)$$

Hence, the phase lag between two lightcurves at each Fourier frequency is simply calculated as

$$\phi_n = \arctan \left[\frac{\Im(C_n)}{\Re(C_n)} \right] \quad (15)$$

These phase lags can easily be converted into timelags τ_n using the relation $\phi_n = 2\pi\nu_n \tau_n$. Examples of lag spectra, i.e. the phase lag as function of Fourier frequency, are shown in Figures 22 and 23 in Chapter 3.

Figure 40.: Illustrative example of the cross correlation function of a sine and cosine with a 1 second period. The decay towards large τ is caused by the limited length of the cross-correlated time series. The time-lag of a quarter period is indicated by the blue dashed line.

The CCF is the time-domain equivalent, i.e. the inverse Fourier transform, of the cross spectrum. Just as the time-domain equivalent of the PSD, the actual lightcurve, contains all frequencies simultaneously, the CCF does not allow one to find the lag at individual frequencies¹. As the cross spectrum is defined as a complex multiplication in the Fourier domain, the CCF is given by a discrete convolution of f_j^* and g_j in the time-domain:

$$(f \star g)_k = \sum_{j=-\infty}^{+\infty} f_j^* g_{j+k} \quad (16)$$

In this equation, both index j and k run over the timebins. The time lag between the two lightcurves is defined as the maximum of the cross correlation. Visually, this corresponds to the timeshift where both lightcurves overlap the most, resulting in the maximum summed value of their product at each timebin. A simple example of the CCF of a sine and cosine, with a trivial time lag of a quarter period, is shown in Figure 40.

A.3 EVEN MORE TIME SERIES: RMS SPECTRA AND LAG ENERGY SPECTRA

Although the methods detailed above can provide a basic phenomenological understanding of the variability properties of BHXRBs, they rarely enable a direct interpretation in terms of physical processes; for instance, the cross spectrum and CCF allow the determination of the phase or time lag, but

¹ Unless a Fourier-domain filter is applied to the lightcurves, as we do in our newly developed method – see Chapter 2.

do not provide any more direct information about the underlying process. A higher energy resolution, resulting in multiple lightcurves in different energy bands, provides more information on the actual physical origin of the variability. A straightforward example would be simply comparing the PSD in many energy bands, linking the observed variability to the dominant spectral component in the considered energy band. Two more sophisticated energy-resolved methods, discussed or referred to in this thesis, are the rms spectrum and the lag energy spectrum. Both will be briefly introduced in this section.

The rms spectrum is a spectrum of the variability in a certain Fourier frequency range. For instance, an rms spectrum of the QPO is a spectrum containing only the flux that is attributable to the QPO variability². As a spectrum directly carries information on the original emission processes, the rms spectrum can be used to connect specific emission processes to individual timing features of the BHXRB.

The rms spectrum is measured by first calculating a PSD for the light curve in each energy band, normalised with the rms-squared normalisation according to Equation 12. Subsequently, these PSDs are all integrated over the required frequency range, e.g. the QPO, or whatever other interesting range in frequency. This returns the rms variability in the specified frequency range, in each energy band. Due to the choice of normalisation, plotting this as rms variability as function of energy returns the spectrum of only the variability in the required frequency range.

An illustrative example of an rms-spectrum, adapted from [Kara et al. \(2013\)](#), is shown in Figure 41. In this Figure, the rms-spectrum of active galactic nucleus 1H0707-495 is plotted for three frequency ranges³. Comparing the spectral properties between these frequency ranges allows for the investigation of the origin of the variability.

Secondly, also the lag energy spectrum can provide more information on the physical origin of observed time or phase lags. A lag energy spectrum is simply the lag between lightcurves in individual energy bands and a reference band lightcurve. When calculated using the cross correlation function, this lag is the average lag between the lightcurves on all timescales. Determining the lags using the cross spectrum instead, makes it possible to investigate the energy dependence of lags associated to any specific timing feature.

The higher energy resolution allows for a possible association of the lags with physical components in the BHXRB; for instance, the lag energy spectrum of the QPO in GRS 1915+105 is known to show a different slope at soft and harder X-ray energies, possibly linked to the an origin in the disk and the corona, respectively ([Pahari et al., 2013](#)). [Pahari et al. \(2013\)](#) also compare the lag energy spectrum of the QPO fundamental and harmonic, and find significant differences. This way, it is possible to investigate similarities and differences between, in this case, the different QPO harmonics.

A.4 SIMULATING NOISE PROCESSES

Self-evidently, the broad-band noise in the PSD of accreting black holes is a noise process. However, [Van der Klis \(1997\)](#) shows that the QPO itself

² To be precise: the flux that is attributable to the variability at the QPO frequency. This could also contain contributions by the BBN, unless a fit of the QPO feature is integrated.

³ As characteristic timescales of accreting black holes increase with black hole mass, the frequency ranges are significantly lower than those for LMXRBs.

Figure 41.: Illustrative example of an application of the rms spectrum in the AGN 1Ho707495. The spectra are calculated through the integrated power spectra in three different frequency ranges. Adapted from [Kara et al. \(2013\)](#)

is a noise process as well (i.e. that the QPO possesses the same statistical properties as the aperiodic noise). Hence, when developing new techniques to investigate the QPO, it is important to verify whether any observed behaviour is not simply attributable to this noise process nature. In this section, we briefly review how to test for this possibility using the [Timmer & Koenig](#) algorithm. [Timmer & Koenig \(1995\)](#) developed an algorithm to simulate a time series of a noise process with a given underlying power-law PSD. Where older algorithms (for instance in [Done et al., 1992](#)) only randomize the phases in the simulated power spectrum, the [Timmer & Koenig \(1995\)](#) method also randomizes the amplitude. As is discussed in section A.1, this is essential due to the statistical properties of the PSD: the amplitude follows a $\chi^2_{\nu=2}$ -distribution, which has a standard deviation equal to its mean. Hence, a wide range of simulated power spectra are consistent with the same underlying PSD. As we observe only a single manifestation of the underlying PSD, it is important to also take this amplitude-randomization into account.

The algorithm consists of the following five steps:

- (i) Choose the required underlying PSD $S(\omega)$. While the original algorithm only discusses power law functions, other noise processes can be modeled as well. Hence, we can choose a QPO-containing PSD as the underlying process.
- (ii) For each nonnegative frequency, draw two random numbers from a Gaussian distribution. To obtain the real and imaginary part of the simulated power spectrum at that frequency, multiply both by $\sqrt{S(\omega)/2}$.
- (iii) Neglect the imaginary part of the power at the Nyquist frequency if the number of simulated datapoints is even.
- (iv) Calculate the power at negative frequencies through $f(-\omega) = f^*(\omega)$ to ensure that the resulting time series is real valued. Here, f denotes the random manifestation of the underlying PSD $S(\omega)$, and the star indicates complex conjugation.
- (v) Apply the inverse Fourier transform to obtain the simulated time series.

This algorithm provides a powerful approach to test whether detected properties of the QPO arise from its noise nature or, alternatively, the detailed properties of the method. However, it does not reproduce every fundamental variability property of accreting black holes. For instance, it assumes that the simulated time series will have a zero-valued bispectrum, i.e. that there are no correlations between different Fourier frequencies. This assumption is invalid for QPOs, as shown by for instance [Maccarone et al. \(2011\)](#). [Maccarone et al.](#) calculate the bispectrum of several GRS 1915+105 observations containing a Type-C QPO, and find various structures indicating correlations between the QPO, harmonic and BBN. These will not be reproduced by the algorithm by [Timmer & Koenig \(1995\)](#).

Despite such downsides of the algorithm, we apply it in several cases to test the robustness of our conclusions. In Chapter 2, we discuss how we compare the shape of the QPO amplitude-frequency relation between observed and simulated lightcurves. This reveals that the shape of this relation arises due to the accuracy of our method, and does not contain information on the underlying physical QPO origin. Additionally, we briefly discuss the

presence of phase lag evolution in simulated lightcurves in Chapter 4. The small, but essential differences between this evolution in actual observations and simulated data show that we are probing the actual QPO mechanism. This way, the [Timmer & Koenig](#) algorithm has proved to be vital in testing the reliability of our method and the robustness of our interpretations.

BINARY-ORBIT INCLINATION ESTIMATES

In this Appendix, I provide a brief summary of the binary-orbit inclination classification of the BHXR sample in [Motta et al. \(2015\)](#). I will not discuss the direct inclination measurements (although references are listed), but instead focus on the indirect observables of the inclination. As the sample in [Motta et al. \(2015\)](#) does not contain GRS 1915+105, I will not discuss it in this Appendix – as explained in Chapter 3, its inclination is estimated from the observed inclination of its relativistic jet ($\sim 70^\circ$, [Mirabel & Rodriguez 1994](#)).

Table 5 is an extended version of Table 3 in Chapter 3: in addition to the inclination classification and possible direct inclination measurement with reference, the Table lists for each source whether it shows absorption dips, X-ray flux spikes, or equatorial winds, and what type of disk temperature evolution it displays. Below, I expand on these observables and their connection to the binary-orbit evolution.

- (i) *Absorption dips* in the X-ray flux are the result of the presence of a bulge at the outer edge of the accretion disk, where it interacts with the companion star ([White & Swank, 1982](#)). Hence, these dips are distinct from regular eclipses. However, they do still track the orbital timescale of the system. Although the exact size of this bulge and the disk influence the range of inclinations where dips are observed, the presence of dips is generally interpreted as a sign of an edge-on (i.e. *high inclination*) binary orbit. All sources showing dips in our sample are identified as high inclination sources.
- (ii) *Flux spikes* refers to the observation by [Muñoz-Darias et al. \(2013\)](#) of spikes in both the X-ray lightcurve and hardness-intensity diagram of high inclination sources. Although no definite explanation of this feature is provided, it appears to be at least phenomenologically related to a high inclination.
- (iii) *Equatorial winds* refers to the presence of accretion disk winds. [Ponti et al. \(2012\)](#) show that these winds have a geometry with only a small opening angle from the disk plane. Hence, the observation of such disk winds implies that the BHXBs are seen edge-on.
- (iv) *Disk temperatures* are shown by [Muñoz-Darias et al. \(2013\)](#) to be systematically higher in edge-on sources. This effect can be fully explained through inclination-dependent relativistic effects, and thus the evolution of the disk temperature can be applied as a probe of binary orbit inclination.

Table 5.: Extended version of Table 3, showing the presence/absence of observables related to a high (edge-on) inclination. See text in Appendix B and Motta et al. (2015) for details. References for the direct inclination measure (4th column): [1] Neustroev et al. (2014), [2] Orosz et al. (2002), [3] Orosz et al. (2004), [4] Munoz-Darias et al. (2008); Zdziarski et al. (1998), [5] Miller-Jones et al. (2011), [6] Corral-Santana et al. (2011), [7] Orosz et al. (2011), [8] Greene et al. (2001), [9] Steiner et al. (2012).

Source	<i>i</i> classification	Direct <i>i</i> measurement	Ref.	Dipping	Spikes	Equatorial winds	Disk temperature
Swift J1753.5-01	Low	$\sim 40 - 55^\circ$	[1]	✗	✗	✗	–
4U 1543-47	Low	$20.7 \pm 1.5^\circ$	[2]	✗	✗	✗	Low- <i>i</i> evolution
XTE J1650-500	Low	$> 47^\circ$	[3]	✗	✗	✗	Low- <i>i</i> evolution
GX 339-4	Low	$40 \leq i \leq 60^\circ$	[4]	✗	✗	✗	Low- <i>i</i> evolution
XTE J1752-223	Low	$\leq 49^\circ$	[5]	✗	✗	✗	Low- <i>i</i> evolution
XTE J1817-330	Low	–	–	✗	✗	✗	Low- <i>i</i> evolution
XTE J1859+226	Undetermined	$\geq 60^\circ$	[6]	✗	✓	✗	Undetermined- <i>i</i> evolution
MAXI J1543-564	Undetermined	–	–	✗	✓	✗	–
XTE J1550-564	High	$74.7 \pm 3.8^\circ$	[7]	✓	✓	✗	High- <i>i</i> evolution
4U 1630-47	High	–	–	✓	✓	✓	High- <i>i</i> evolution
GRO J1655-40	High	$70.2 \pm 1^\circ$	[8]	✓	✓	✓	High- <i>i</i> evolution
H1743-322	High	$75 \pm 3^\circ$	[9]	✓	✓	✓	High- <i>i</i> evolution
MAXI J1659-152	High	–	–	✓	✓	✗	–
XTE J1748-288	High	–	–	✓	✓	✗	–

ADDITIONAL PLOTS FOR INDIVIDUAL SOURCES

In Figures 42 – 55, we plot the Type-C fundamental phase lag and relative QPO frequency difference as function of QPO frequency for all individual sources showing Type-C QPOs. This excludes the source XTE J1817-330, for which only Type-B QPOs were analysed.

Figure 42.: Type-C QPO phase lag and $\Delta\nu_0/\nu_0$ as function of QPO frequency for SWIFT J1753.5-0127

Figure 43.: Type-C QPO phase lag and $\Delta\nu_0/\nu_0$ as function of QPO frequency for 4U 1543-47

Figure 44.: Type-C QPO phase lag and $\Delta\nu_0/\nu_0$ as function of QPO frequency for XTE J1650-500

Figure 45.: Type-C QPO phase lag and $\Delta\nu_0/\nu_0$ as function of QPO frequency for GX 339-4

Figure 46.: Type-C QPO phase lag and $\Delta\nu_0/\nu_0$ as function of QPO frequency for XTE J1752–223

Figure 47.: Type-C QPO phase lag and $\Delta\nu_0/\nu_0$ as function of QPO frequency for XTE J1859+226.

Figure 48.: Type-C QPO phase lag and $\Delta\nu_0/\nu_0$ as function of QPO frequency for MAXI J1543–564

Figure 49.: Type-C QPO phase lag and $\Delta\nu_0/\nu_0$ as function of QPO frequency for XTE J1550–564.

Figure 50.: Type-C QPO phase lag and $\Delta\nu_0/\nu_0$ as function of QPO frequency for 4U 1630-47

Figure 51.: Type-C QPO phase lag and $\Delta\nu_0/\nu_0$ as function of QPO frequency for GRO J1655-40

Figure 52.: Type-C QPO phase lag and $\Delta\nu_0/\nu_0$ as function of QPO frequency for H1743–322.

Figure 53.: Type-C QPO phase lag and $\Delta\nu_0/\nu_0$ as function of QPO frequency for GRS 1915+105.

Figure 54.: Type-C QPO phase lag and $\Delta\nu_0/\nu_0$ as function of QPO frequency for MAXI J1659-152

Figure 55.: Type-C QPO phase lag and $\Delta\nu_0/\nu_0$ as function of QPO frequency for XTE J1748-288