























Evidence for a dynamic corona in the short-term time lags of black hole X-ray binary MAXI J1820+070

Niek Bollemeijer ¹★, Phil Uttley ¹, Arkadip Basak ¹, Adam Ingram ^{2,3}, Jakob van den Eijnden ⁴, Kevin Alabarta ⁵, Diego Altamirano ⁶, Zaven Arzoumanian ⁷, Douglas J. K. Buisson ⁸, Andrew C. Fabian ⁹, Elizabeth Ferrara ^{7,10}, Keith Gendreau ^{11,12}, Jeroen Homan ¹³, Erin Kara ¹⁴, Craig Markwardt ⁷, Ronald A. Remillard ¹⁴, Andrea Sanna ¹⁵, James F. Steiner ¹⁶, Francesco Tombesi ^{10,12,17,18,19}, Jingyi Wang ¹⁴, Yanan Wang ^{6,20} and Abderahmen Zoghbi ^{10,21,11}

Affiliations are listed at the end of the paper

Accepted 2023 December 12. Received 2023 November 24; in original form 2023 September 15

ABSTRACT

In X-ray observations of hard state black hole X-ray binaries (BHXRBs), rapid variations in accretion disc and coronal power-law emission are correlated and show Fourier-frequency-dependent time lags. On short (~ 0.1 s) time-scales, these lags are thought to be due to reverberation and therefore may depend strongly on the geometry of the corona. Low-frequency quasi-periodic oscillations (QPOs) are variations in X-ray flux that have been suggested to arise because of geometric changes in the corona, possibly due to general relativistic Lense–Thirring precession. Therefore, one might expect the short-term time lags to vary on the QPO time-scale. We performed novel spectral-timing analyses on *Neutron Star Interior Composition ExploreR* observations of the BHXRB MAXI J1820+070 during the hard state of its outburst in 2018 to investigate how the short-term time lags between a disc-dominated and a coronal power-law-dominated energy band vary on different time-scales. Our method can distinguish between variability due to the QPO and broad-band noise, and we find a linear correlation between the power-law flux and lag amplitude that is strongest at the QPO frequency. We also introduce a new method to resolve the QPO signal and determine the QPO phase dependence of the flux and lag variations, finding that both are very similar. Our results are consistent with a geometric origin of QPOs, but also provide evidence for a dynamic corona with a geometry varying in a similar way over a broad range of time-scales, not just the QPO time-scale.

Key words: accretion, accretion discs – black hole physics – X-rays: binaries – X-rays: individual: MAXI J1820+070.

1 INTRODUCTION

Black hole X-ray binaries (BHXRBs) are stellar mass black holes that accrete matter from an orbiting star via an accretion disc. The resultant spectral shape of the disc emission in the X-rays is that of a multitemperature blackbody (Novikov & Thorne 1973; Shakura & Sunyaev 1973). Some fraction of the radiated power takes the form of X-ray power-law emission, thought to be produced by Compton upscattering of disc photons by an optically thin cloud of hot electrons close to the black hole, known as the corona (Sunyaev & Truemper 1979). The spectrum of this component is well described by a power law with a high-energy (~ 100 keV) cut-off (e.g. Motta, Belloni & Homan 2009), but the overall shape and geometry of the corona are as yet unknown. Different models have been proposed, with some assuming the corona to be a layer above the surface of the disc close to the black hole (Galeev, Rosner & Vaiana 1979; Svensson & Zdziarski 1994). The corona could also be a hot flow region between the (thin) disc and the black hole or the base of an astrophysical jet, or some combination of the above (e.g. Markoff, Nowak & Wilms 2005; Gilfanov 2010;

Marcel et al. 2018; Poutanen, Veledina & Zdziarski 2018; Zdziarski et al. 2021). It should be noted that recent X-ray polarization results favour a horizontal geometry in the plane of the accretion disc (Krawczynski et al. 2022). The third major component of the X-ray spectrum of BHXRBs originates from the reflection of high-energy photons on the disc. The reflection spectrum consists mainly of a continuum ‘Compton hump’ peaking around ~ 30 keV and emission lines, of which the Fe $K\alpha$ line is usually the strongest (George & Fabian 1991; Matt, Fabian & Reynolds 1997; Bambi et al. 2021).

Despite high data quality and advanced modelling, the spectral constraints on the properties of the corona are often degenerate. Adding time-domain information can break these degeneracies and provide better constraints on, for example, the geometry of the corona (Uttley et al. 2014). There are several timing properties that can be measured in light curves (see e.g. Uttley et al. 2014 for a detailed review of X-ray spectral-timing techniques). In this paper, we will focus on the time lag between variations in two X-ray energy bands (0.5–1 and 3–10 keV) for the BHXRB MAXI J1820+070 in its hard state. By calculating the phase of the mean Fourier cross-spectrum of two light curves in different energy bands, it is possible to measure the time lag between correlated variations in both light curves on different time-scales (Nowak et al. 1999). Lags between different

* E-mail: n.a.bollemeijer@uva.nl

energy bands are found in many XRBs. By probing the strong broadband noise variability observed in the hard state in a soft and a hard band, dominated respectively by emission from the disc and from the corona, two main types of lags are distinguished (Uttley et al. 2011; De Marco et al. 2015; Kara et al. 2019). At low Fourier frequencies, hard photon variability lags behind variations in the softer X-rays (typically by tenths of a second), so these are called hard lags. At higher Fourier frequencies, variations in the soft band are often found to arrive later than those in the hard band (typically by 1–10 ms), a signature known as soft lags.

Plots of lag versus energy demonstrate that most of the low-frequency hard lag can be associated with a delay between variations of the (lagging) coronal power-law and (leading) disc blackbody emission (Uttley et al. 2011). This delay can be naturally explained if the variability is produced by mass accretion fluctuations propagating inwards through the disc (Lyubarskii 1997) to a central coronal region.

When comparing two harder energy bands, a smaller energy-dependent hard lag is seen within the power-law emission, which has been variously attributed to Comptonisation light-travel delays in the corona (Kazanas, Hua & Titarchuk 1997; Bellavita et al. 2022) or jet (Reig et al. 2003; Kylafis et al. 2008; Wang et al. 2020a), accretion fluctuations propagating through a coronal hot flow (Kotov, Churazov & Gilfanov 2001; Arévalo & Uttley 2006; Veledina, Poutanen & Vurm 2013a) or coronal heating versus seed photon delays in response to accretion fluctuations propagating through the disc to the corona (Uttley & Malzac 2023). When measuring lags between 3–10 keV, which consists mainly of coronal power-law emission, and 0.5–1 keV, which is dominated by disc emission but also contains coronal emission, it is important to realize that the observed lags arise due to a combination of several mechanisms. The relative strength of different spectral components can influence the measured lag values.

To explain the soft lags between disc and coronal emission observed at higher Fourier frequencies, different mechanisms have been proposed. Mushtukov, Ingram & van der Klis (2018) argue that outward propagation of mass accretion fluctuations can cause soft lags and Kawamura, Done & Takahashi (2023) show that inwardly propagating fluctuations can lead to soft lags if the inner region has a higher low energy break in its spectrum than the outer region, which could happen in the hard-intermediate state (HIMS). Uttley et al. (2011), De Marco et al. (2015), and Alston et al. (2020) interpret soft lags as evidence of reverberation. In their view, photons are emitted by the corona in all directions, and some will be thermally reprocessed at lower energies by the disc. Due to, for example, the light travel time and the reprocessing time-scale between both components, variations in the soft band are observed later than those in the hard band, to produce soft lags. Reverberation lags should be highly dependent on coronal geometry, which influences what fraction of hard photons is reprocessed by the disc and which also determines the size of the lags and at what frequency the lags switch from hard to soft.¹ Kara et al. (2019) found that time lags on short time-scales (> 3 Hz) change during the 2018 outburst of BHXRB MAXI J1820+070, which is the source that is also being analysed in this paper. They interpreted these changes in lags as being related to changes in the vertical extent of the corona.

If high-frequency time lags depend on coronal geometry, they can potentially be used to study a different phenomenon observed




Figure 1. Examination of the mean power spectrum and time lags for all observations with a similar QPO frequency combined. We group ObsIDs into 10 data sets (A through J) with increasing QPO frequency (see Table 1 for more details.) Here, we demonstrate the lags and PSD for two data groups, C and G, where the mean QPO frequency is 0.18 and 0.34 Hz, respectively. The QPO is stronger in the hard band and a harmonic signal at twice the QPO fundamental frequency is clearly visible in both data sets. The lower panel shows the time lags for the same data sets. The short-term time lags around 4–20 Hz change during the outburst. The inset plot shows the soft lags on a linear scale.

in many BHXRBs, called quasi-periodic oscillations or QPOs (for recent reviews of QPOs, see Motta 2016 and Ingram & Motta 2019). Several subtypes of QPO are distinguished. During the hard state, low-frequency QPOs of type C, which have a centroid frequency ranging from a few mHz up to ~ 10 Hz, are typically present as a narrow peak in the power spectrum of the light curve. Wang et al. (2022) showed that the high-frequency lag behaviour is tightly linked to the type-C QPO frequency in MAXI J1820+070 and other BHXRBs.

The changes in both the power spectra and the short-term time lags of MAXI J1820+070 are visualized in Fig. 1, which shows power spectra and lag–frequency spectra for two different data groups. The upper panel of Fig. 1 shows power spectra for two subsets of data with different QPO-frequencies, with prominent peaks in the hard band power spectra demonstrating the presence of QPOs. In the power spectrum, QPO signals often consist of two clear peaks: One at the fundamental frequency and one at around twice that frequency, which is called the second harmonic. We will refer to the latter as the harmonic from now on. Variations at the fundamental and harmonic frequencies are linked in phase, and together they create a specific QPO waveform, which can be measured (see e.g. Ingram & van der Klis 2015; De Ruiter et al. 2019; Nathan et al. 2022).

A number of competing models have been proposed to explain low-frequency QPOs. Various models suppose that the QPO corresponds to intrinsic luminosity changes, linked to characteristic frequencies or instabilities in the accretion flow or jet, for example, corrugation modes (Tsang & Butsky 2013), the accretion–ejection instability (Varnière, Tagger & Rodriguez 2012), or an instability in the jet (Ferreira et al. 2022). However, a number of observational results support the idea that the QPO is primarily due to a geometric

¹The lags between two power-law-dominated energy bands remain positive or go to zero at high frequencies and do not show a turnover to soft lags.

Table 1. The combinations of ObsIDs used in different parts of this paper. Also, the total selected exposure time and the QPO frequency range for each data set of ObsIDs are shown. More information on the individual ObsIDs can be found in Table D1.

Data set	ObsIDs	Selected exposure time (s)	QPO frequency (Hz)	Flux range 3–10 keV (count s ⁻¹)
A	126, 127, 129, 130	6080	0.07–0.12	852–1932
B	131, 132, 133, 134	10 560	0.12–0.16	839–1861
C	135, 136, 137, 138, 139	14 080	0.16–0.20	856–1835
D	181, 182, 183, 184, 185	5248	0.20–0.24	243–524
E	140, 141, 142, 176, 177, 178, 180, 186	20 416	0.24–0.28	509–1047
F	143, 144, 169, 171, 173, 174, 175, 179, 187,	21 888	0.28–0.32	545–1108
G	145, 146, 158, 159, 160, 161 162, 167, 168, 170, 172, 188	32 384	0.32–0.36	620–1236
H	147, 165, 166	9152	0.36–0.40	675–1320
I	148, 151, 157, 163, 164, 189	26 560	0.40–0.44	560–1092
J	149, 150, 152, 153, 154, 155, 156	12 224	0.44–0.51	758–1452

Note. The rightmost column shows the range of hard fluxes (lowest and highest flux bin count rates) obtained by creating 5 flux bins of 0.25 s slices within each 64 s segment, as explained in Section 3.1.

variation in the appearance of the corona. Evidence that the QPO corresponds to a varying coronal geometry is provided by the binary system inclination dependence of QPO phase lags (Van den Eijnden et al. 2016b) and amplitudes (Heil, Uttley & Klein-Wolt 2015; Motta et al. 2015), as well as QPO phase-resolved spectroscopy (Ingram et al. 2016, 2017). The geometric interpretation is supported by the study of high-energy QPO lags (Ma et al. 2021) and the variability spectrum of the QPO (see e.g. Sobolewska & Życki 2006, Axelsson, Done & Hjalmarsdotter 2013 and Gao, Yan & Yu 2023 with *Insight-HXMT* data for MAXI J1820+070). The foremost explanation for these geometric variations is the relativistic precession model (Stella, Vietri & Morsink 1999; Schnittman, Homan & Miller 2006; Fragile et al. 2007; Ingram, Done & Fragile 2009). In the model, the hot inner flow and associated corona and/or jet precess due to the general relativistic Lense–Thirring effect (Lense & Thirring 1918). Recent general relativistic Magnetohydrodynamical (GRMHD) simulations also demonstrate the effect (see e.g. Liska et al. 2018, 2023; Musoke et al. 2022; Bollimpalli, Fragile & Kluźniak 2022).

If the coronal geometry changes during a QPO cycle, light travel times and reflection fractions will vary as well, and so we can expect that high-frequency time lags should also vary over the QPO cycle. Connecting QPO phase and short-term time lags is therefore an important way to test the geometric interpretation of QPOs. In this paper, we analyse the variability of high-frequency time lags in a hard state BHXRB on different time-scales, to determine whether there is evidence for variability in the lags that may be linked to the QPO or other types of variation.

To search for variations in the short-term time lags on the QPO time-scale, high-quality data are required in order to ensure that the signal at high Fourier frequencies is not overwhelmed by observational noise. Also, the studied source should show both broadband noise variability and QPOs. All these conditions are met by observations of BHXRB MAXI J1820+070 in its hard state with the *Neutron Star Interior Composition Explorer* (*NICER*) (Gendreau et al. 2016), obtained during the 2018 outburst. We describe these data and our basic methods in Section 2.

In Section 3, we present our analysis and results. We first look at how the lags change with flux, which is highly variable on a range of time-scales. We then use Fourier filtering methods to show how the correlation between lag and flux depends on the variability time-scale, to show that the strongest relation arises at the QPO time-scale. We also introduce a new method to identify the QPO phase at each point in time in a light curve, using a combination of Fourier filtering

and interpolation, to measure how the short-term time lags depend on QPO phase. In order to test our new method, we used a more established method, pioneered by Ingram & van der Klis (2015) and applied to many XRBs by De Ruiter et al. (2019), to see if it returned similar results in terms of QPO waveform reconstruction, which is presented in the Appendix, along with a description of a simulation approach to check the validity of our results. We end with a discussion of the implications of our results for different models for the QPO and the coronal geometry.

2 OBSERVATIONS AND METHODS

We made use of the extensive observations by *NICER* of the low-mass BHXRB MAXI J1820+070 (known in the optical as ASASSN-18ey, Torres et al. 2019). *NICER* has excellent timing capabilities and can observe bright sources with minimal instrumental deadtime and no spectral distortion such as due to pileup. It is located on the International Space Station (ISS). MAXI J1820+070 was one of the brightest X-ray sources ever observed during its outburst in 2018, which led to an extraordinarily rich data set (see e.g. Buisson et al. 2019; Kara et al. 2019; Fabian et al. 2020; Homan et al. 2020; Axelsson & Veledina 2021; De Marco et al. 2021; Wang et al. 2021; Zdziarski et al. 2021; You et al. 2021, 2023).

We analysed data collected between 2018 April 11 (ObsID 1200120126, MJD 58219) and 2018 June 28 (ObsID 1200120189, MJD 58297), while MAXI J1820+070 was in the hard state. During this part of the outburst, the QPO fundamental frequency varied between 0.077 and 0.51 Hz. For the remainder of this paper, data sets will be referred to by the last three digits of their ObsIDs. The full ObsID number is constructed by placing 1200120 before these three digits. One ObsID refers to observations that are obtained during one Earth day. The list of ObsIDs and their total observation time can be found in Table D1. Because data set D, corresponding to QPO frequencies between 0.2 and 0.24 Hz, only contains five relatively short observations with lower count rates and weak QPOs, results for this frequency range typically have large error bars and should be interpreted with caution.

All ObsIDs were reprocessed using the *nicerl2* pipeline from NASA’s HEASOFT (v6.28) (HEASARC 2014), version 6, released on 2020 July 22. Good Time Intervals (GTIs) were applied to all ObsIDs used. Any GTI shorter than 64 s was excluded. For GTIs longer than 64 s, only an integer number of 64 s segments was used and partially filled segments were discarded. In ObsIDs 130 and




Figure 2. To explain the three methods we use to analyse the behaviour of the short-term time lags, the three panels show how the 0.25 s slices are combined in different ways for each method with data from ObsID 145. The 3–10-keV light curve is shown in black and the vertical lines show the edges of each 0.25 s slice. The numbers and shades below the light curves each correspond to a single flux bin in the upper panels, while the lower panel also contains diagonal hatching to show the different phase bins. The cross-spectra of simultaneous 3–10 and 0.5–1-keV light-curve slices with the same shade and number are averaged to calculate the lags of each flux or phase bin. The edges of the flux bins are calculated for each individual 64 s segment and are visible as horizontal cyan lines in panel (1) for unfiltered flux binning and in panel (2) for QPO-filtered flux binning. For the green QPO-filtered curve, it is clear that both the phase and amplitude shift over time. Because the figure only shows 16 s of a 64 s segment, not all flux or phase bins are filled to the same extent, but there is an equal number of slices in each bin for the entire segment. In panel 2 and 3, the amplitude of the filtered light curves and the edges of the filtered flux bins were multiplied by a factor of 5 for clarity, so the actual flux difference between individual filtered flux bins is much smaller than shown in the figure.

131, from 2018 April 16 and 17, there were several GTIs in which an unusual flux decrease was observed. This was caused by the robot arm on ISS crossing in front of *NICER* during observations,² and could be identified through its characteristic shadow pattern crossing the 56 FPMs. Affected time intervals were omitted.

We created light curves with a time-resolution of 1/128 s, resulting in a Nyquist frequency of 64 Hz. Light curves were made for two energy bands. The soft band is defined to be 0.5–1 keV, where disc emission dominates for disc temperatures between ~ 0.2 and 0.3 keV, as observed in the hard state of MAXI J1820+070 (Dzielał, Marco & Zdziarski 2021; Wang et al. 2021). The hard band ranges from 3–10 keV, where the power-law component originating from the corona is the main contributor. To investigate how the short-term time lags vary on a time-scale of seconds, all 64 s segments are subdivided into 0.25 s slices. The choice for a certain slice length over which the lags are calculated is based on two properties. First, the available Fourier frequency is restricted to high frequencies for short slices. The lowest Fourier frequency measured for these short slices is close to the Poisson noise-dominated region, resulting in low signal-to-noise ratios (SNRs). However, later in our analysis, we want to be able to distinguish well between different QPO phases. Since the QPO frequency ranges from ~ 0.08 to 0.5 Hz for the data analysed here, corresponding to periods between 2 and 12 s, the slice length should be a fraction of these periods. Balancing between these two arguments, we use 0.25 s slices, so the lowest frequency measured is 4 Hz. At this frequency, Poisson noise does not dominate in the energy bands used. The frequency range over which the time lags are determined is 4–20 Hz. The measured lags in this frequency range arise due to a combination of hard lags at lower frequencies and soft lags at higher frequencies. Both the hard and the soft lags show strong long-term evolution (Kara et al. 2019; De Marco et al. 2021; Wang et al. 2022). Higher frequencies do not add much extra signal, because Poisson noise dominates the fast variability. To keep our results model independent, we do not investigate the hard and soft lags in the 4–20 Hz range separately. Instead, we show that the average 4–20-Hz time lags probe the evolution of the soft lags well (see e.g. Figs 3 and 4).

The Fourier cross-spectrum of the hard and the soft band is calculated for all simultaneous hard and soft band light-curve slices. In order to obtain good signal-to-noise, it is necessary to average the cross-spectra of many slices. The way we select which slices to average together is discussed in Section 3. To obtain the phase lag between both bands for those slices, the argument of the average of the cross-spectrum over the 4–20 Hz range for all included light-curve slices is then calculated. The phase lag can be converted to a time lag by dividing by $2\pi\bar{\nu}$, where $\bar{\nu}$ is the mean frequency of the range for which the lag was calculated. In our analysis, the mean frequency is $\bar{\nu} = 12$ Hz.

We also calculate the errors on the phase lags following Uttley et al. (2014), using

$$\Delta\phi(\nu_j) = \sqrt{\frac{1 - \gamma^2(\nu_j)}{2KM\gamma^2(\nu_j)}}, \quad (1)$$

where $\gamma^2(\nu_j)$ is the raw coherence between both light curves for frequency bin ν_j , K is the number of Fourier frequencies over which the lags are averaged (e.g. for 4–20 Hz and 0.25 s light-curve slices, $K = 5$) and M is the number of slices used to form the average cross-spectrum. The phase lag errors are converted to time lag errors by

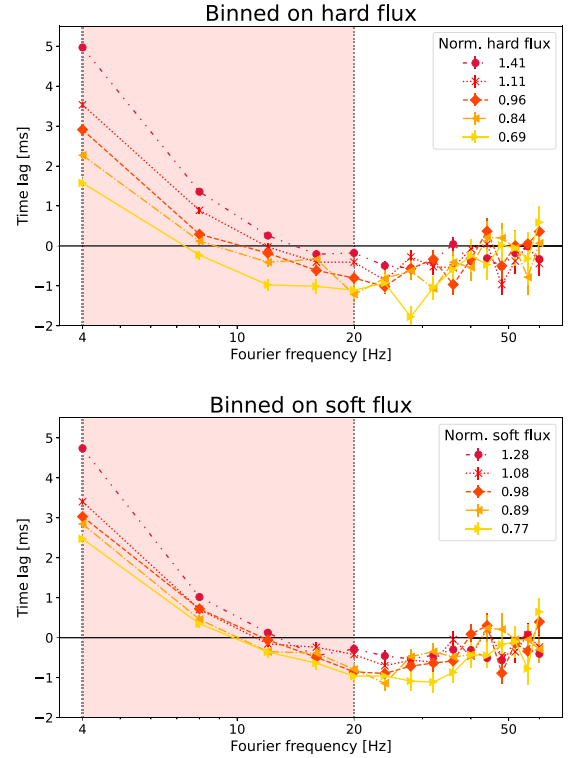


Figure 3. The upper panel shows the lag–frequency spectra of the five different hard flux bins, while the lower panel shows the same for soft flux bins. It is clear that the amplitude of the soft lag decreases and the frequency at which the lags switch from hard to soft increases with flux and that the effect of binning on hard flux is larger than for the soft flux. The 4–20 Hz range over which the lags are calculated for Fig. 4 is highlighted.

dividing by $2\pi\bar{\nu}$. For a higher SNR, synonymous with smaller error bars on the time lags, we combined multiple observations to obtain the cross-spectrum, and applied the described methods to all available data. The QPO frequency and with it the shape of the power spectrum evolves considerably during the outburst. The averaged cross-spectrum over a frequency range is weighted by the power in each frequency bin, so a changing power spectral shape can influence the lag measurement, even if the lags themselves do not change. We also perform the same analysis on 10 separate subsets of data, consisting of ObsIDs with similar QPO frequencies, which are listed in Table 1.

3 DATA ANALYSIS AND RESULTS

Our goal is to investigate how the short-term time lags change on a time-scale of seconds and how they are connected to QPOs, which might shed light on changing coronal-geometry on short time-scales. We can measure the short-term lags by averaging the cross-spectra obtained from many 0.25 s light-curve slices, but in order to determine their relation to specific variability time-scales or the QPO, we consider three different methods for ‘binning’ the slices together for cross-spectrum averaging purposes. The three methods are schematically explained in Fig. 2, which has one panel for each method, showing 16 s of hard band light curve and indicating how 0.25 s slices are combined. We briefly introduce each method below, before providing further details and results in separate sub-sections:

(i) First, we investigate whether the lags depend on the flux variability on a time-scale between 0.25 and 64 s, which are the lengths of a slice and of a segment. Both QPO and broad-band noise

²https://heasarc.gsfc.nasa.gov/docs/nicer/timelines/nicer_significant_events.html

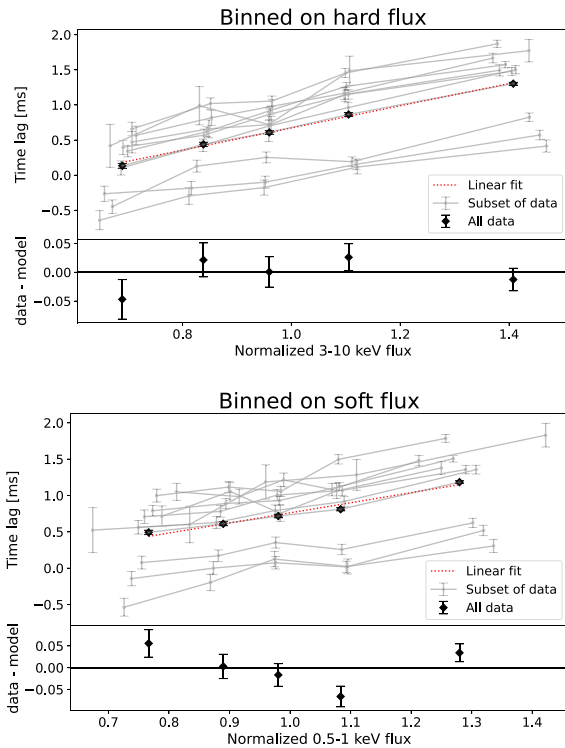


Figure 4. The upper panel shows the normalized hard flux versus the mean 4–20-Hz time lags between the hard and the soft band for all included ObsIDs as black diamonds. The lags and flux were calculated for 0.25 s slices. The lighter curves show the result of flux binning for the different subsets of data, grouped by their QPO frequency and shown in Table 1. All subsets show the same trend with hard flux, while the mean time lag differs between observations. The lower panel shows the same results for binning on the soft 0.5–1-keV flux.

variability play an important role on these time-scales. Short 0.25 s slices are grouped together based on their relative flux within 64 s segments, so any variability on longer time-scales (e.g. as the source evolves during the outburst) does not affect the flux-binning. The upper panel of Fig. 2 shows how the different slices are combined, which is discussed more extensively in Section 3.1.

(ii) In the framework of a geometric origin of QPOs, we expect the lag to vary more on the QPO time-scale than on other time-scales. We therefore devised a method using Fourier filtering to isolate the variability on a specific time-scale, using only the flux variations on that time-scale to bin slices according to their relative flux. We can then probe how variability on the QPO time-scale influences the short-term time lags and compare it to the effects of variability on other time-scales. The filtered flux binning is shown schematically in the middle panel of Fig. 2 and explained in more detail in Section 3.2.

(iii) Finally, we present a method to resolve the QPO signal and follow how the short-term time lags evolve with QPO phase. The main goal is to investigate whether changes in flux and lag happen simultaneously or that there is a delay between both, which could indicate different types of geometric change. We call it the filter-interpolate (FI) method, as shown in the lower panel of Fig. 2 and presented in Section 3.3.

In all of the methods described in this section, the lags are obtained from the average cross-spectrum of a large number of short slices. The short-term time lags evolve during the outburst and thus each group or bin will consist of slices with different intrinsic lags. The

absolute value of the average lags therefore does not have much physical meaning. However, we can still compare the lag values in different bins to measure a relation between lag and flux, as all the methods only take into account variability within a 64 s segment.

To test our methods and to be able to distinguish systematic effects from results that are intrinsic to the data, we simulated ‘null hypothesis’ light curves with constant lags based on Timmer & Koenig (1995), to which we added a method to simulate QPOs. The methods we used to simulate data are discussed in more detail in Appendix A.

3.1 Short-term lag response to flux variations on time-scales up to 64 s

We combine different 0.25 s slices according to their mean flux by calculating the mean flux of each slice and sorting the obtained values for each 64 s segment. Within each 64 s segment, the 51 slices with the lowest mean fluxes are assigned to the same flux bin, the next 51 slices to the next bin and so on, until we have five different flux bins. This is illustrated in the upper panel of Fig. 2. We decided to create five flux bins to have sufficient signal-to-noise in each bin, also for subsets of observations, while retaining a median flux bin. Because each segment consists of 256 slices of 0.25 s and we can only assign 255 of them to a flux bin, we exclude the last 0.25 s of each (unsorted) segment. The binning process is performed for all 64 s segments and the respective flux bins of each segment are combined to create five flux bins, each consisting of thousands of light-curve slices. It is important to note that the mean flux value that is used to combine different slices can be either the hard or the soft flux, i.e. 3–10 or 0.5–1 keV, as explained in the previous sections. We carried out the analysis and show the results for both cases, but Fig. 2 only shows an example for binning on 3–10-keV flux. Because dividing the 0.25 s slices over the five flux bins is only done according to the relative flux within each 64 s segment, the long-term evolution of the flux does not influence the binning.

We calculate the short-term time lags between the hard and soft energy band for each flux bin as described in Section 2. Once the mean short-term time lags for the five distinct flux bins are obtained, models can be fitted to the five data points to quantify the relation between the lags and the flux. We fit a constant and a linear model to determine whether there is a significant change of short-term time lags with flux.

In Fig. 3, the lag–frequency spectra for the five different hard and soft flux-binned 0.25 s light-curve slices are shown for all data. The lag–frequency dependence clearly evolves, as the cross-over frequency from hard to soft lags increases and the maximum soft lag decreases with flux. The net hard lag in the 4–20 Hz range therefore increases with flux, as is illustrated in Fig. 4, where we show the relation between the mean (normalized) hard and soft flux and the short-term (4–20 Hz) time lags. The normalization involves dividing the flux of each slice by the average flux of all slices in its segment. A linear fit and its corresponding residuals are also shown. It is clear that the linear fit describes the data reasonably well for the hard flux binned data, which is demonstrated by a goodness of fit (p -value) of 0.26 with Pearson’s χ^2 test. The slope of the linear fit for binning on hard flux is 1.58 ± 0.05 ms.

The lower panels of Figs 3 and 4 show the behaviour of the short-term time lags when binning on soft flux. Although the overall trend is the same, several discrepancies between the soft and hard flux-binned data can be discerned. The goodness of fit of the linear function to the soft flux binned data is 0.003, indicating a more complex relation. Also, the ranges of both normalized flux and time lags are smaller

when binning on soft flux. The slope of the linear fit is slightly shallower when binning on soft flux instead of hard flux with a value of 1.39 ± 0.14 ms. From the combination of these properties, we conclude that the short-term time lags are connected more closely to the hard band variations than to the soft band variations. Further results are only shown for binning on the hard band, even though similar effects can be distinguished when binned on soft flux. Since variability in both hard and soft light curves is correlated, it is not surprising that binning on either of them returns qualitatively similar effects.

The linear lag versus flux relation is not only seen when combining all observations, but also in the different subsets of data as presented in Table 1. The results for the subsets of data are shown as the grey lines in Fig. 4. The time lag range covered by binning on flux is very similar for all subsets of data. The subsets with higher QPO frequencies, corresponding to the lines with a more positive average lag in Fig. 4, display a reduced range in flux, which leads to slightly steeper slopes.

3.2 Time-scale dependence of short-term lag response to flux variations

In Section 3.1, we show that the short-term time lags vary linearly with flux on time-scales of 0.25–64 s, but it is unclear whether the relation is driven by the QPO or broad-band noise variability, or both. To distinguish between QPO and broad-band noise variability, we use a method that can successfully extract the QPO signal from a light curve. We follow, for example, Van den Eijnden, Ingram & Uttley (2016a) and Press et al. (1992) and apply an optimal filter to the Fourier transform (FT) of each light curve segment. To obtain the optimal filter, we first fit five Lorentzians to the mean power spectrum of all light curve segments in a single ObsID. Three broad Lorentzians account for the broad-band noise and two narrow Lorentzians are centred at the QPO fundamental and harmonic frequencies. Poisson noise is accounted for by adding a constant to the fit. The optimal filter consists of the narrow QPO fundamental Lorentzian function divided by the sum of all five Lorentzian functions, which leaves us with a filter that closely resembles the shape of the QPO peak in the power spectrum (see fig. 2 in Van den Eijnden, Ingram & Uttley 2016a for an example of an optimal filter). All frequencies outside twice the full width at half-maximum of the narrow fundamental Lorentzian are excluded. Because the QPO strength varies during the outburst and the fits return rather different QPO widths, we assume a Q-factor of 8 for all ObsIDs. $Q = \nu_{\text{QPO}}/\text{FWHM}$ and it is a measure of how many cycles, on average, the QPO signal stays in phase with itself. $Q \sim 8$ or higher for typical type-C QPOs observed at > 3 keV (Ingram & Motta 2019).

The inverse FT of the filtered FT returns a smooth and filtered light curve containing variations on the QPO time-scale. This is illustrated in the middle panel of Fig. 2, where the QPO filtered light curve and the mean value of that filtered light curve in each slice are shown in green and yellow, respectively. Slices of the original black light curve are combined according to the value of the filtered light curve, as illustrated by the shades and numbers below the x-axis. Slices with the same number belong to the same filtered flux bin. Only variability on the QPO time-scale is taken into account when creating the five different bins. The mean cross-spectrum between simultaneous slices in the (unfiltered) hard and soft band and the corresponding short-term time lags in the 4–20 Hz range are calculated for each of the five filtered flux bins.

To compare the results of filtering on QPO and non-QPO frequencies, we can move the centroid frequency of the Lorentzian fitting the QPO to other frequencies, while keeping the broad Lorentzians the same. We then only take into account variability from a chosen,




Figure 5. The result of measuring and fitting the time lags versus flux after binning 0.25 s light-curve slices on filtered flux for different Fourier frequency filters for all data combined. On the x-axis, the filtering frequency divided by the QPO frequency of each included ObsID is shown and the horizontal error bars indicate the width of the optimal filter that was used. The upper panel shows the time lag range, defined as the difference in time lags between the lowest and highest flux bin. The middle panel shows the standard deviation of the normalized filtered flux bin values, which indicates the variability on the filtered time-scale. The standard deviation $\sigma = 0.25$ for unfiltered flux binning. This is much larger than for filtered flux binning, as filtering removes a large part of the variability. In the lower panel, the slope of a linear fit to the lag–flux relation (see Fig. 4) is shown. The QPO clearly stands out in all three panels.

(non-)QPO time-scale to combine 0.25 s slices from the original (unfiltered) hard and soft light curves. The rest of the filtering process is identical to the QPO case. We always used an optimal filter with $Q = 8$ at both QPO and non-QPO frequencies. An alternative, applying a top-hat filter with the same width, returned similar results. It is important to note that even if we use an optimal filter at the QPO frequency, the broad-band noise will still contribute to the filtered light curve, as the optimal filter cannot distinguish between overlapping QPO and broad-band noise variability.

By shifting the centroid frequency of the optimal filter, choosing to include or exclude the QPO frequency, we can separate out the QPO time-scales from the broad-band noise continuum and map the dependence of the short-term time lags on flux variations at different time-scales. However, since the flux variability amplitude (fva) is time-scale dependent, the lags are likely to be influenced by the flux range covered by filtering on a given range of Fourier frequencies. To determine the strength of the connection between flux and lags on different time-scales, we fit a linear function to the short-term time lag as a function of hard flux, similar to what is visible in Fig. 4. We use the slope of the lag–flux relation to determine whether the QPO has a different relation to variations in short-term time lags than the broad-band noise. The result of the analysis for all data combined is shown in Fig. 5.

The upper panel shows the difference in the time lags between the lowest and highest (filtered) flux bin. There is a clear peak at the QPO frequency. In principle, the presence of the peak could be explained by the fact that the flux variability is large at the QPO frequency, as is visible in the middle panel. If there is a simple linear relation between flux and time lag, large flux variability leads to an increase in time lag range. However, in that case, the slope of the relation is

not expected to change. In the lower panel of Fig. 5, it is clear that the slope of the relation is steeper when using a light curve filtered on the QPO frequency than when using filters on non-QPO frequencies or the unfiltered flux as the binning property. This suggests that there is a profound connection between the QPO and the variations of the short-term time lags.

Fig. 5 shows 16 filter frequencies between 0.4 and 2.5 times the QPO frequency. Filter frequencies above this range lead to strong biases in the lag measurements, which are also observed in simulated light curves with a constant lag, while these biases are much weaker for lower filter frequencies. An explanation of the bias can be found in Appendix C.

3.3 QPO phase-resolved time lags: the filter-interpolate method

To further understand how short-term time lags and the QPO are related, we designed the filter-interpolate (FI) method for estimating the QPO phase at a given time. The first steps are the same as those used for measuring short-term time lag variations on different time-scales in Section 3.2. After creating the QPO filtered light curve by applying the same optimal filter to the hard band, we determine the time values of the extrema of the filtered curve. With these points in time, a triangular wave is created with value 1 at the maxima and value -1 at the minima. A linear interpolation connects these points, hence the name of the method. Interpolation is necessary because both the QPO frequency (or equivalently, phase) and amplitude vary from one cycle to the next (Van den Eijnden, Ingram & Uttley 2016a). The FI method is shown visually in the lower panel of Fig. 2. The original filtered line is shown as the solid green curve, while the result of the interpolation is visible as the dashed, red triangular wave in the same figure. The vertical lines indicate the start and stop times of each light-curve slice. Each slice is assigned a QPO phase based on the value of the triangular wave in the centre of the slice and the sign of its slope, which means we assume that the QPO frequency does not change during half a QPO cycle. Each QPO phase corresponds to a distinct shade and number below the x -axis. The inclusion of the sign of the slope of the interpolation is the main difference between the FI method and the method described in Section 3.2 and shown in panel 2 of Fig. 2. With the sign of the slope, we can distinguish between rising and falling parts of the QPO signal and resolve the QPO phase. Another difference between both methods is that the FI method only takes into account the phase of the filtered curve, while binning on filtered flux depends on the amplitude of the filtered curve in each slice.

Next, we can associate a QPO phase with the *unfiltered* hard and soft light-curve slices so that we can calculate a cross-spectrum for different QPO phase bins. Any light-curve slice shorter than $\sim 1/4$ of the QPO period can be assigned a QPO phase by the FI method. Longer slices will average out variations due to the QPO. The 0.25 s slices used in this work are not influenced by this effect, as even the highest QPO frequency in the data set (~ 0.5 Hz) has a period of about eight slices. Each hard and soft 0.25 s light-curve slice is assigned to 1 of 10 equal-width QPO phase bins and the mean cross-spectrum between the two energy bands for all slices in a phase bin is calculated and used to calculate the short-term (4–20 Hz) time lags. The 10 time lag values obtained constitute the time lag waveform, which informs us how the short-term time lags evolve with QPO phase. The flux and lag waveforms for the QPO for all data combined are visible in Fig. 6.

When we have the 10 QPO phase bins for both the hard flux and the short-term time lags, we can fit a function $f(\Theta)$ (where Θ is the phase in radians, defined at the central filter frequency) to the waveforms, which is the sum of two cosine waves corresponding to

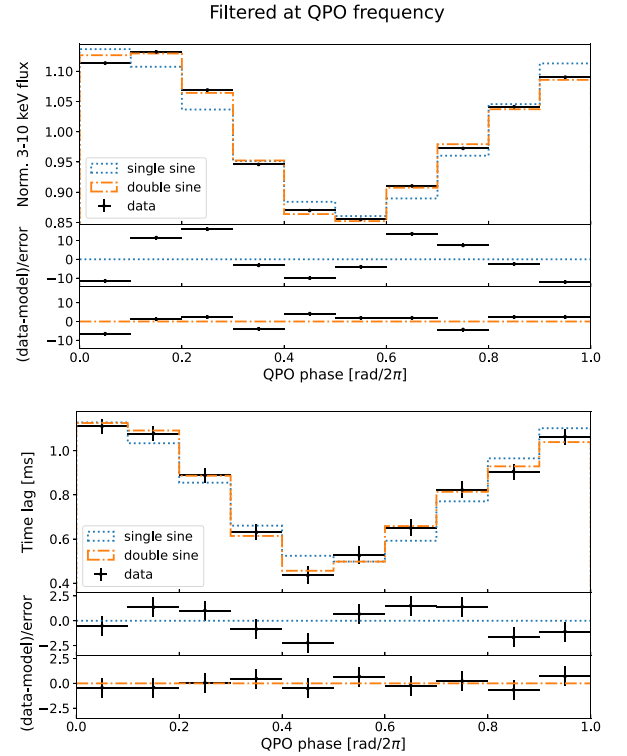


Figure 6. The hard flux and 4–20-Hz time lag (between the hard and the soft band) waveforms for all data, obtained by the FI method at the QPO frequency. Also, a single sine wave and a double sine wave fit are shown. It is clear that the time lag follows the shape of the flux. For both waveforms, the addition of a harmonic signal significantly improves the fit. The error bars on the flux are underestimated.

the fundamental and harmonic components:

$$f(\Theta) = 1 + A_f \cos(\Theta - \phi_f) + R_{\text{amp}} A_f \cos[2(\Theta - \phi_f - \Psi)]. \quad (2)$$

Here, A_f is the fundamental amplitude, R_{amp} is the ratio of the harmonic and the fundamental amplitudes, and Ψ is the phase difference that determines the waveform, defined by Ingram & van der Klis (2015) as

$$\Psi = (\phi_h/2 - \phi_f) \bmod \pi, \quad (3)$$

where ϕ_f and ϕ_h are the phase offsets of the fundamental and harmonic respectively. The division by 2 is due to the fact that the frequency of the harmonic is twice the frequency of the fundamental, so its phase changes twice as much for an equal shift in time. With the fit, we can quantify the waveforms, as is discussed in more detail in Appendices A and B.

In Fig. 6, the hard flux and time lag QPO waveforms are shown for all data combined. We fit both a single (same as equation 2 with $R_{\text{amp}} = 0$) and a double sine wave model to the waveforms, integrating the model over each flux bin. It is clear that the flux waveform requires a harmonic signal, at twice the filtering frequency, in order to be fit reasonably well. We used the presence of a harmonic signal in the QPO waveform to test our method, which is described in Appendix B. The parameters of the fits are shown in Table 2. The errors on the flux obtained from counting statistics (simply $1/\sqrt{N_{\text{counts}}}$, where N_{counts} is the number of counts in a flux bin) are too small, as they do not account for correlated errors between phase bins that are introduced by our FI method of phase reconstruction used to bin light-curve slices. Therefore we determined the errors on the

Table 2. The fitting parameters for a model consisting of a fundamental and harmonic sine wave to the flux and lag waveforms (equation 2) when filtering on the QPO frequency or 1.5 times the QPO frequency, which are shown in Figs 6 and 7.

Filter	QPO frequency		$1.5 \times$ QPO frequency	
	Normalized flux	Time lags	Normalized flux	Time lags
Offset	1 (fixed)	0.773 ± 0.013 ms	1 (fixed)	0.84 ± 0.012 ms
Fundamental amplitude (A_f)	0.1372 ± 0.0006	0.319 ± 0.016 ms	0.0977 ± 0.0005	0.121 ± 0.016 ms
Amplitude ratio (R_{amp})	0.186 ± 0.005	0.22 ± 0.06	0.074 ± 0.007	0.40 ± 0.17
ϕ_f [rad]	0.315 ± 0.004	0.12 ± 0.05	0.34 ± 0.005	0.16 ± 0.12
Ψ/π	0.320 ± 0.005	0.33 ± 0.05	0.206 ± 0.014	0.3 ± 0.3

Notes. The 1σ errors on the parameters were determined by bootstrapping the 64 s segments of all data 100 times and performing the analysis on the bootstrapped samples. Because the non-QPO lag waveform does not contain a significant harmonic, the amplitude ratio and phase difference Ψ are not well constrained and have large errors.

fitting parameters by bootstrapping 64 s segments of data 100 times and fitting double sine waves to the waveforms obtained from the bootstrapped samples. The standard deviations of the parameter distributions are reported in Table 2. The errors on the lags are more consistent with the scatter of the data points and we find the lag waveform contains a harmonic signal with a 3.5σ significance.³ The values for the harmonic-fundamental amplitude ratio and phase difference Ψ are consistent with those of the flux waveform, indicating that the short-term lags are tightly connected to the flux.

Analogous to what we show in Section 3.2, we can investigate both QPO and non-QPO variability and their relation to the short-term time lags by shifting the centroid of the optimal filter. We created optimally filtered hard band light curves for different Fourier frequency ranges and compare the results. The rest of the method is the same as described above, so we can compare the effect of including the QPO in the filtering frequency range. As an example, Fig. 7 shows the hard flux and time lag waveforms for broad-band noise at 1.5 times the QPO frequency of each ObsID. As can also be seen in Table 2, the broad-band noise flux waveform contains a harmonic component that is significantly weaker than in the QPO case and probably originates from the broad-band noise rms–flux relation, which causes Fourier phases to be correlated between frequencies (e.g. see Uttley, McHardy & Vaughan 2005). The lag waveform fit does not improve significantly (1.5σ)⁴ when including a harmonic signal, which could be due to the relatively large error bars on the lags (due to the lower rms in the flux).

From Table 2, it is clear that the values for the fundamental phase ϕ_f are different for the flux and lag waveforms. The non-zero phase difference⁵ indicates that the lags and the flux do not vary simultaneously, but that there is a delay between variations in the flux and in the short-term time lags. To investigate the delay further, we determined the phase difference between the lag and flux waveforms by fitting a single sine wave to both and comparing their phases, so we do not take into account any harmonic signal here. This simplified model is used to avoid degeneracies in the phase caused by the harmonic components, which can compensate for differences in the phase of the fundamental.

Fig. 8 shows the lag variability amplitude (lva), fva, their ratio and the phase difference between the lag and the flux waveforms

³The significance is calculated from the difference in χ^2 , which decreases from 17.5 for a single sine wave and 7 degrees of freedom (d.o.f.) to 2.2 for a double sine wave with 5 d.o.f.

⁴The χ^2 decreases from 10.5 for 7 d.o.f. to 6.3 for 5 d.o.f.

⁵To avoid confusion: Ψ is the phase difference between the fundamental and harmonic frequencies in a light curve (see Appendix A), while the phase difference referred to here is between the fundamental of the flux and lag waveform, which we will also call the delay between flux and lag.

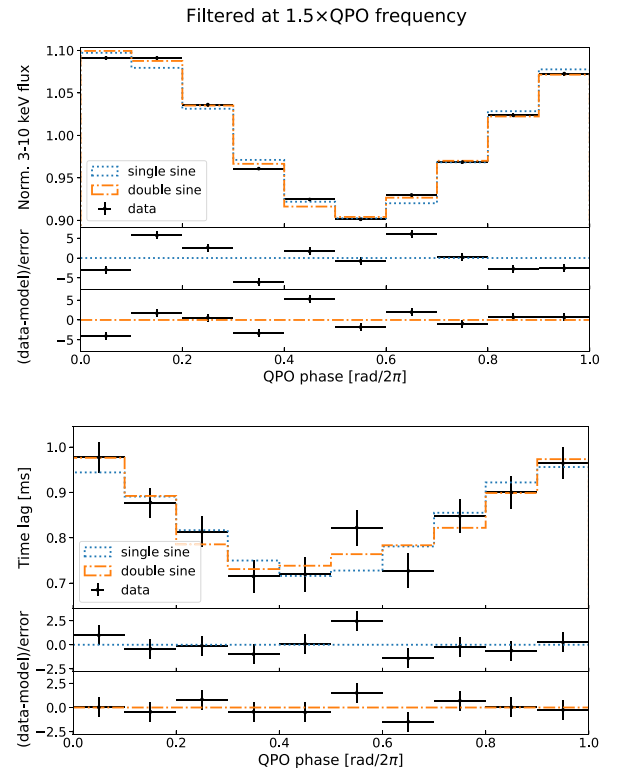


Figure 7. The hard flux and 4–20-Hz time lag (between the hard and the soft band) waveforms for all data, obtained by the FI method at 1.5 times the QPO frequency. The harmonic contribution to the flux waveform is much smaller than when filtering on the QPO frequency (Fig. 6) and the harmonic does not improve the fit to the lag waveform significantly.

for different filter frequencies. The lva and fva are defined as the amplitudes of the sine functions fitted to the lag and flux waveforms as presented in Figs 6 and 7. The errors were determined by bootstrapping the 64 s segments and applying the FI method to the bootstrapped samples. Bootstrapping the 64 s segments accounts for the fact that neighbouring filter frequency results are correlated, which is also visible in the turquoise bootstrapped curves. The lower right panel of Fig. 8 shows that the lag and flux waveforms are almost but not quite simultaneous and there is a delay (phase difference) between the lag and flux waveforms, especially for non-QPO frequencies. Negative delays correspond to the flux waveform following the lag waveform, i.e. the lags vary first. Performing the same analysis on subsets of data ordered by their mean QPO frequency results in significant scatter between different data sets, especially in the phase difference between lag and flux, although




Figure 8. In panel (a) the lva (4–20 Hz between 3–10 and 0.5–1 keV) is shown, in panel (b) the fva , the ratio of both is visible in panel (c) and panel (d) shows the phase difference between the short-term time lag and flux waveforms for all included data for different filter frequencies relative to the QPO frequency of each observation. The turquoise lines show the total of 50 bootstraps that were used to determine the errors, while the solid black lines represent the original data. Bootstrapping was used to account for the errors on different filter frequencies, which are correlated. There is a clear maximum for both variability amplitudes and their ratio when filtering at the QPO frequency. In the lower right panel, a negative phase difference indicates that the flux follows the short-term time lags, which is due to phase leakage also seen in null-hypothesis simulations.

almost all measured delays lie between 0 and -2 rad. Because the measured delays are difficult to understand, we investigated them using simulated light curves with constant lags, which is discussed in Section 3.4.

We confined the filter frequency in the FI method to between 0.6 and 2 times the QPO frequency, which is a slightly smaller range than for the filtered flux binning method introduced in Section 3.2. The FI method is slightly more sensitive to the Fourier leakage effects that show up in constant lag simulations, especially at higher frequencies, as is explained in Appendix C.

3.4 Comparison with constant-lag simulated light curves

Because the process of dividing the light curves into short slices and binning them according to flux or QPO phase can introduce different kinds of Fourier leakage and systematic biases (Uttley, McHardy & Papadakis 2002; Alston, Vaughan & Uttley 2013), we tested our methods on simulated light curves as introduced in Appendix A. These simulated light curves have a constant lag, so they test the null hypothesis that the observed changing lags are systematic effects introduced by the methods we used. There are several meaningful similarities and discrepancies between the observations and the simulated data, which we highlight here.

First, it is important to note that none of the constant lag simulations can reproduce the linear relation between short-term time lags and (hard) flux, as shown in Fig. 4, indicating that the observed relation cannot be due to the methods we used. We employed different methods of creating light curves consisting of a broad-band noise signal multiplied⁶ by a QPO signal and convolved with an impulse

response function (see Appendix A for more on the simulations) and included the ubiquitously observed rms–flux relation (Uttley, McHardy & Vaughan 2005; Heil, Vaughan & Uttley 2012) by exponentiation of the original light curve. The result of applying the FI method to simulated data is shown in Fig. 9, which shows for the simulated data the same measured parameters as Fig. 8, obtained from fitting single sine waves to the simulated lag and flux waveforms. Because the time lags are simulated in the Fourier domain, we expect them to stay constant in time. However, despite the lack of a direct relation between lag and flux, Fourier leakage effects can induce a waveform in the lags, especially at higher filter frequencies, albeit with a small amplitude (up to ~ 0.1 ms) when compared to the data (up to ~ 0.3 ms). The measured phase difference between the lag and flux waveform is not distributed uniformly, as would be the case if there were no Fourier leakage, but there is a clear preference for phase differences between -2 and -1 rad in the simulations, especially at higher filter frequencies. The negative phase differences signify that variations in flux follow variations in the lag, just like we observe in the data. By simulating light curves both with and without different rms–flux relations and for different lags, we found that the rms–flux relation is critical for understanding the effect.

Due to Fourier leakage, phase bins which correspond to a rising or falling part of the waveform (i.e. showing an overall trend), tend to have smaller (absolute) values for the measured lags, while peaks and troughs in the waveform do not suffer from leakage and have larger lags (see Appendix C). In simulated light curves without an rms–flux relation, the lag waveform due to leakage produces a harmonic signal with a low ($\lesssim 0.1$ ms) amplitude. Implementing an rms–flux relation by exponentiating the light-curve results in a lag waveform with a slightly higher amplitude, but more importantly, it introduces larger lags in the phase bins with a positive slope and lower lags in the phase bins with a negative slope. Due to this asymmetry, we measure an \sim quarter-cycle delay of the flux compared to the lags when fitting single sine waves to both waveforms. To confirm that the rms–flux

⁶For completeness: We also tested the methods with simulations in which the QPO was added to the broad-band noise or which did not contain a separate QPO signal. These simulations yielded similar results.




Figure 9. The 4–20-Hz I_{va} is shown in panel (a), panel (b) shows the f_{va} , (c) their ratio and panel (d) the phase difference between the short-term time lag and flux waveforms for 50 ‘null-hypothesis’ simulations consisting of 1000 segments of 64 s with a QPO frequency of 0.31 Hz. The black line shows the mean value of the parameters of the 50 simulations. Because there is no intrinsic lag–flux relation in the simulation, the I_{va} is much smaller than in the data. The lag waveform that arises due to Fourier leakage has a slightly larger amplitude for higher filter frequencies and tends to precede the flux waveform by $\sim \pi/2$ rad, as is visible in the lower right panel.

relation causes the delay, we also created light curves with an inverse rms–flux relation by subtracting the exponentiated light curve from a constant, increasing the amplitude of variability and finally setting all negative count rates equal to zero. The obtained light curves show a negative rms–flux relation and indeed we find that the lags follow the flux for these simulations, indicating that a larger rms increases the Fourier leakage and decreases the lags in subsequent phase bins.

The leakage effect described above could explain the lag versus flux phase-difference behaviour that we see in the data, where the phase difference becomes more negative at non-QPO filter frequencies. The value of the delay is much closer to zero in the data than in the simulations, which is expected if the general lag–flux relation is (close to) instantaneous but Fourier leakage effects push it to negative values. To better understand biases in the lag versus flux phase difference, simulations which can reproduce an instantaneous linear change in lag with flux are required. The Fourier-based simulation methods we employ here are efficient, but they are limited by the fact that it is difficult to vary the lags over short time-scales. In order to introduce time-dependent lags, the light curves should be simulated in the time-domain. This could be done with a more physics-based model, which we leave for future work.

4 DISCUSSION

We have shown that the short-term (4–20 Hz) 3–10 versus 0.5–1-keV time lags, observed in the broad-band noise of BHXR B MAXI J1820+070, are linearly correlated with flux over a broad range of time-scales, with the strongest correlation seen between the lags and the harder, power-law-dominated flux. Notably, the steepest lag versus flux relation occurs on the QPO time-scale. The flux and lag waveforms on the QPO time-scale are similar, with only a small (~ 0.2 rad or ~ 0.03 cycles, see Fig. 8) phase delay between them, which we attribute to bias caused by Fourier leakage effects (which probably also causes the delay between variations of lag and flux observed on other time-scales).

The linear lag–flux relation seems to be caused by the evolution of the lag versus frequency dependence, with the cross-over frequency from hard to soft lags increasing with flux and the maximum soft lag decreasing at the same time, so that the net hard lag in the 4–20 Hz range increases with flux. Similar changes in the high-frequency lags are seen on much longer time-scales, as the source evolves through the hard state (e.g. as seen in Fig. 1 and shown by Kara et al. 2019, Wang et al. 2021, and De Marco et al. 2021). Wang et al. (2022) analysed the high-frequency soft lags in many BHXR Bs that were observed by *NICER* and found that these lags evolve in a similar way in all systems during an outburst.

Wang et al. (2021) model the lags in MAXI J1820+070 using the RELTRANS reverberation model (Mastroserio, Ingram & van der Klis 2018; Ingram et al. 2019; Mastroserio et al. 2021) and find that the variations in soft lag at high frequencies can be explained by changes in coronal height. In their interpretation, larger heights produce greater light travel delays and correspondingly larger lags, with coronal height changing from ~ 30 to $> 300 R_g$ through the hard-to-soft state transition. For MAXI J1820+070 (Wang et al. 2021) and BHXR B GX 339-4 (Wang et al. 2020b), the required large coronal heights appear to be in tension with results from spectral modelling of relativistic reflection, which imply more compact coronae. Note that in a recent work, Lucchini et al. (2023) reconcile both reflection and timing results by including two lampposts to simulate a more vertically extended corona. In the context of these more vertically extended models, it is important to realize that the first polarization results from *IXPE* for Cyg X-1 in the hard state and for Swift J1727.8–1613 in the HIMS favour a more horizontally extended corona (Krawczynski et al. 2022; Ingram et al. 2023).

The observations we study here correspond to a relatively small range of lag variation during the bright hard state (compared to state changes), when the corona is likely more compact, at most a few tens of R_g in height (Wang et al. 2021). Coronal heights may be significantly smaller than derived from light travel times if disc mass-accretion propagation and seed photon effects are taken into

account (Uttley & Malzac 2023). In any case, coronal height changes may be required to explain the pattern of short-term lag variability seen in MAXI J1820+070, with the coronal height *decreasing* with increasing X-ray flux. The X-ray flux variations on non-QPO time-scales are thought to be linked to mass accretion fluctuations, which would further suggest that decreases in coronal height are linked to increases in mass accretion rate.

Following the idea that the observed changes in timing properties are due to variations in coronal geometry, we propose to distinguish two types of geometric change: observer dependent and intrinsic. If the geometric change is observer-dependent, this means that, for example, the angular size or beaming of the coronal emission towards the observer is different, but the actual shape and emission of the corona stays constant. In the case of a precessing corona, its orientation shifts over time, causing variations in the emission towards the observer due to a combination of solid-angle changes, relativistic beaming effects and angular-dependence of coronal emission (Veledina, Poutanen & Ingram 2013b). These are observer-dependent variations. Because we do not expect changes in, for example, orientation angles over a time-scale of weeks, this type of geometric variability is only feasible over short time-scales, such as due to a precessing hot flow or jet in the Lense–Thirring framework.

We define intrinsic coronal changes as variations in geometry that do influence the shape and/or emission of the corona, independent of the observer’s viewing angle. The shape or size of the corona actually change for this type of geometric change, instead of only the coronal orientation, as is the case in observer-dependent changes. To give an example, an increase in coronal height can take place over a broad range of time-scales, for example, as the accretion rate changes during an outburst, but probably also on a time-scale of seconds. The increased coronal height will affect the seed photon flux towards the corona, changing spectral-timing properties. Long-term lag differences, such as presented in Kara et al. (2019), De Marco et al. (2021), and Wang et al. (2022) will thus be due to intrinsic rather than observer-dependent changes in coronal geometry. Lense–Thirring precession also includes an intrinsic geometric change, in that it also causes the hot flow to *nutate*, i.e. there is a (quasi-periodic) change in the polar angle of the precessing hot flow. From the point of view of the disc, the height of one side of the precessing flow is varying at the QPO frequency, which affects the seed photon flux coming from the disc and associated spectral and timing properties.

As discussed in Section 1, there is substantial evidence that QPO timing properties depend on binary orbit inclination (Heil, Uttley & Klein-Wolt 2015; Motta et al. 2015; Van den Eijnden, Ingram & Uttley 2016a; De Ruiter et al. 2019), which, in turn, suggests that the QPO corresponds to some kind of changing inner-region geometry. It is possible that the QPO (in the flux) is itself produced by the same variations in coronal geometry (e.g. optical depth and corresponding height changes due to accretion rate fluctuations) that produce the lag variations and which are somehow enhanced at the QPO frequency. For example, if an increase in X-ray flux corresponds to a shrinking corona, then the QPO may correspond to a particular frequency at which the coronal height oscillates. This picture is inconsistent with the precessing hot flow model for the QPO, however, since in that case the quasi-periodic flux variation at a particular frequency is linked to the precession of a hot flow with otherwise constant shape.

In the precession model, lag changes might be produced due to, for example, light-travel delay variations as the hot flow preferentially illuminates the near and far side of the disc (as seen by the observer) at different precession phases (see e.g. Ingram & Done 2012; You, Bursa & Życki 2018; You et al. 2020, for examples, of varying coronal illumination of the disc, but note that none of these papers




Figure 10. The time lag versus frequency spectra for all data combined, using five different soft bands between 0.5 and 1 keV. The lag depends strongly on the choice of soft band energies and the difference is reminiscent of Fig. 4, where the soft band is 0.5–1 keV but the data are split into five different flux bins.

discuss the lags that are the focus of the current paper). Including nutation in the model could mimic coronal height changes at the QPO frequency. Coronal precession at the QPO frequency and coronal height changes linked to accretion rate fluctuations on a broad range of time-scales could then happen simultaneously and potentially explain our obtained results, as they would both correspond to changes in solid angle as seen from the disc and from an observer.

To understand the changing short-term lags at both QPO and broad-band-noise-dominated time-scales, we also looked the energy dependence of the time lags, which is demonstrated in Fig. 10. In the figure, we show the time lag versus frequency spectra for narrow soft bands between 0.5 and 1 keV, which can be compared to Fig. 4. The lag behaviour for slightly different soft bands is similar to what we see when binning on hard flux and the range of time lags covered is also comparable, so spectral changes might explain at least part of the lag versus flux relation. Reverberation lags are observed when comparing disc and power-law dominated bands, while the lags between different power-law energies do not become negative and are generally positive or consistent with zero lags at high Fourier frequencies. Spectroscopic fits show that the 0.5–1-keV energy band we use here consists of photons from the disc (between 50 and 70 per cent for all observations) and from the corona (50 to 30 per cent).⁷ A change in power-law normalization could impact the weighting of both components and their associated lags (i.e. disc versus power law and medium power law versus hard power law). For example, a larger normalization of the power law would lead to measuring harder lags as the power law versus power-law lags play a larger role. We observe a similar pattern in the data, which in this view does not require the intrinsic disc versus power-law lags to change as much. It remains to be explained why the relative strength of the spectral components in the 0.5–1-keV energy range changes, which could still be due to a change in geometry.

⁷These values are the result of fitting spectra in the 0.3–10-keV range with two models. The first spectral model, `tbabs*simpl*diskbb`, corresponds to the situation that the seed photons of the Comptonization component originate from the entire accretion disc, yielding a lower limit for the 0.5–1-keV disc contribution. The second model is `tbabs*(simpl*bb+diskbb)`, which corresponds to seed photons only arising from the inner parts of the disc and provides an upper limit for the soft band disc contribution. For all studied observations, the obtained values for the fraction of photons from the disc lie between 50 and 70 per cent.

From QPO-resolved spectroscopy, we know that the QPO is associated with a strong modulation of the power-law normalization (Ingram et al. 2016). A recent paper by Gao, Yan & Yu (2023) reached a similar conclusion for the hard state QPOs in MAXI J1820+070 with *Insight-HXMT* data. The large power-law flux modulation at the QPO frequency would explain why the largest change in short-term time lags happens at the QPO frequency. The normalization is expected to change on other time-scales as well, providing an explanation for the measured lag versus flux relation on broad-band noise frequencies. The idea that the main spectral property of the QPO is a varying power-law normalization, with a small change in photon index Γ , is consistent with a geometric origin (You et al. 2020). The coronal emission would not change (much) intrinsically, but a change in solid angle and relativistic boosting effects result in a modulation of the power-law flux. The same would go for, for example, coronal height changes linked to accretion rate fluctuations. Time- and energy-dependent simulations can shed light on the effects of changing the strength of different spectral components on the measured time lags.

The interpretation of the lag versus flux relation in this section depends strongly on the model used to explain the soft disc versus power-law lags routinely observed in BHXRBs, which we assume to be reverberation of power-law photons on the disc. Other models have been proposed (see e.g. Mushtukov, Ingram & van der Klis 2018; Kawamura, Done & Takahashi 2023) and those could lead to different conclusions. Because the measured short-term time lags arise due to several different mechanisms (e.g. reverberation and power-law pivoting), changes in the hard lags within the power-law component at those frequencies would cause similar observational results. However, these hard lags are probably also strongly affected by the geometry so they do not change our overall conclusion. Again, more detailed physical simulations could shed more light on the effects of a change in lags in different spectral components.

MAXI J1820+070 is a remarkably bright source, which spent a large amount of time in the hard state so that a detailed study of the kind presented here could be carried out using *NICER*'s exceptional data. Similar studies of other BHXRBs will be more challenging but could shed important light on whether the observed lag variability is common or specific to MAXI J1820+070 or sources like it. For example, the lag variability may also depend on system inclination, which might be expected given the enhanced lag variability at the frequency of the (inclination dependent) QPO signal.

In our analysis, we focused on changes in the short-term time lags on the QPO time-scale and broad-band-noise-dominated time-scales. If the fluctuating short-term lags can be attributed to geometric variations, our results sketch a picture of a dynamic inner X-ray emitting region. Assuming this, the strongest change in geometry is observer-dependent and takes place on the QPO time-scale, while slightly smaller intrinsic variations happen on both longer and shorter time-scales and are linked to accretion variability. We can connect these findings to recent GRMHD simulations of accreting black holes by, for example, Liska et al. (2018, 2022) and Musoke et al. (2022), which show that the component that could serve as the corona is far from being a static region. The simulations include a strong and dynamic magnetic field, which could be key to understanding coronal geometry variations on short time-scales.

5 CONCLUSIONS

Our main findings can be summarized as follows:

(i) The short-term (4–20 Hz) time lags between 0.5–1 and 3–10-keV light curves show a positive linear relation with the hard flux.

(ii) The relation between short-term time lags and hard flux exists for variability on a range of time-scales and is strongest at the QPO frequency.

(iii) In the framework of X-ray reverberation, we interpret the lag–flux coupling as being due to geometric changes on time-scales of seconds to tens of seconds, with the strongest variations taking place on the QPO time-scale.

We have introduced our new FI method, which is designed to perform QPO phase-resolved spectral-timing. We tested our method both by comparing its results with those from a more established method (Appendix B) and also by applying it to simulated ‘null hypothesis’ light curves with constant lags (Appendix A). Although these simulated light curves can replicate some important properties of both the QPO and the broad-band noise in real data, they cannot reproduce the short-term time lag–hard flux relation. In order to simulate the lag–flux relation, more physical time- and energy-dependent light-curve simulations are required, which can also be used to study the effects of changing the strength of different spectral components. Our results support a geometric origin of low-frequency QPOs, but also suggest that the corona is a dynamic structure over a broad range of time-scales, whose variable nature should be taken into account by any explanatory model.

Because the spectral-timing properties in BHXRBs vary on short time-scales, we should be somewhat cautious when interpreting time-averaged spectral-timing results. Large future X-ray observatories, such as *Athena* (Nandra et al. 2013), *eXTP* (Zhang et al. 2016), and *STROBE-X* (Ray et al. 2019), could address this issue by resolving spectral-timing variations down to short time-scales for a much larger sample of sources, providing important information on the dynamics of the corona.

ACKNOWLEDGEMENTS

The authors thank the anonymous referee for their helpful comments. JvdE acknowledges a Warwick Astrophysics prize post-doctoral fellowship made possible thanks to a generous philanthropic donation and was supported by a Lee Hysan Junior Research Fellowship awarded by St Hilda's College, Oxford during part of this work. AZ is supported by NASA under award number 80GSFC21M0002.

This research makes use of the SciServer science platform (www.sciserver.org). SciServer is a collaborative research environment for large-scale data-driven science. It is being developed at, and administered by, the Institute for Data Intensive Engineering and Science at Johns Hopkins University. SciServer is funded by the National Science Foundation through the Data Infrastructure Building Blocks (DIBBs) program and others, as well as by the Alfred P. Sloan Foundation and the Gordon and Betty Moore Foundation.

This work made use of *ASTROPY*:⁸ a community-developed core PYTHON package and an ecosystem of tools and resources for Astronomy (Astropy Collaboration 2013, 2018, 2022). Also, we made extensive use of *NUMPY* (Harris et al. 2020) and *SCIPY* (Virtanen et al. 2020).

DATA AVAILABILITY

This research has made use of data obtained through the High Energy Astrophysics Science Archive Research Center Online Service, provided by the NASA/Goddard Space Flight Center. The

⁸<http://www.astropy.org>

data underlying this paper are available in HEASARC, at <https://heasarc.gsfc.nasa.gov/docs/archive.html>. Upon publication, a basic reproduction package for the results and figures in this article will be made available on Zenodo.

REFERENCES

- Alston W. N., Vaughan S., Uttley P., 2013, *MNRAS*, 435, 1511
- Alston W. N. et al., 2020, *Nat. Astron.*, 4, 597
- Arévalo P., Uttley P., 2006, *MNRAS*, 367, 801
- Arur K., Maccarone T. J., 2019, *MNRAS*, 491, 313
- Astropy Collaboration, 2013, *A&A*, 558, A33
- Astropy Collaboration, 2018, *AJ*, 156, 123
- Astropy Collaboration, 2022, *ApJ*, 935, 167
- Axelsson M., Veledina A., 2021, *MNRAS*, 507, 2744
- Axelsson M., Done C., Hjalmarsdotter L., 2013, *MNRAS*, 438, 657
- Bambi C. et al., 2021, *Space Sci. Rev.*, 217
- Bellavita C., García F., Méndez M., Karpouzias K., 2022, *MNRAS*, 515, 2099
- Bollimpalli D. A., Fragile P. C., Kluźniak W., 2022, *MNRAS*, 520, L79
- Buisson D. J. K. et al., 2019, *MNRAS*, 490, 1350
- De Marco B., Ponti G., Muñoz-Darias T., Nandra K., 2015, *ApJ*, 814, 50
- De Marco B., Zdziarski A. A., Ponti G., Migliori G., Belloni T. M., Segovia Otero A., Dzielak M. A., Lai E. V., 2021, *A&A*, 654, A14
- De Ruiter L., van den Eijnden J., Ingram A., Uttley P., 2019, *MNRAS*, 485, 3834
- Dzielak M. A., Marco B. D., Zdziarski A. A., 2021, *MNRAS*, 506, 2020
- Fabian A. C. et al., 2020, *MNRAS*, 493, 5389
- Ferreira J. et al., 2022, *A&A*, 660, A66
- Fragile P. C., Blaes O. M., Anninos P., Salmonson J. D., 2007, *ApJ*, 668, 417
- Galeev A. A., Rosner R., Vaiana G. S., 1979, *ApJ*, 229, 318
- Gao C., Yan Z., Yu W., 2023, *MNRAS*
- Gendreau K. C. et al., 2016, in den Herder J.-W. A., Takahashi T., Bautz M., eds, Proc. SPIE Conf. Ser. Vol. 9905, Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray. SPIE, Bellingham. p. 99051H
- George I. M., Fabian A. C., 1991, *MNRAS*, 249, 352
- Gilfanov M., 2010, X-Ray Emission from Black-Hole Binaries. Springer, Berlin, p. 17
- HEASARC, 2014, Astrophysics Source Code Library, record ascl:1408.004
- Harris C. R. et al., 2020, *Nature*, 585, 357
- Heil L. M., Vaughan S., Uttley P., 2010, *MNRAS*, 411, L66
- Heil L. M., Vaughan S., Uttley P., 2012, *MNRAS*, 422, 2620
- Heil L. M., Uttley P., Klein-Wolt M., 2015, *MNRAS*, 448, 3348
- Homan J. et al., 2020, *ApJ*, 891, L29
- Ingram A., Done C., 2012, *MNRAS*, 427, 934
- Ingram A. R., Motta S. E., 2019, *New A Rev.*, 85, 101524
- 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*, 461, 1967
- Ingram A., van der Klis M., Middleton M., Altamirano D., Uttley P., 2017, *MNRAS*, 464, 2979
- Ingram A., Mastroserio G., Dauser T., Hovenkamp P., van der Klis M., García J. A., 2019, *MNRAS*, 488, 324
- Ingram A. et al., 2023, preprint (arXiv:2311.05497)
- Kara E. et al., 2019, *Nature*, 565, 198
- Kawamura T., Done C., Takahashi T., 2023, *MNRAS*, 525, 1280
- Kazanas D., Hua X.-M., Titarchuk L., 1997, *ApJ*, 480, 735
- Kotov O., Churazov E., Gilfanov M., 2001, *MNRAS*, 327, 799
- Krawczynski H. S. et al., 2016, *Astropart. Phys.*, 75, 8
- Krawczynski H. et al., 2022, *Science*, 378, 650
- Kylafis N. D., Papadakis I. E., Reig P., Giannios D., Pooley G. G., 2008, *A&A*, 489, 481
- Lense J., Thirring H., 1918, *Phys. Z.*, 19, 156
- Liska M., Hesp C., Tchekhovskoy A., Ingram A., van der Klis M., Markoff S., 2018, *MNRAS*, 474, L81
- Liska M. T. P., Musoke G., Tchekhovskoy A., Porth O., Beloborodov A. M., 2022, *ApJ*, 935, L1
- Liska M., Hesp C., Tchekhovskoy A., Ingram A., van der Klis M., Markoff S., 2023, *New Astron.*, 101, 102012
- Lucchini M. et al., 2023, *ApJ*, 951, 21
- Lyubarskii Y. E., 1997, *MNRAS*, 292, 679
- Ma X. et al., 2021, *Nat. Astron.*, 5, 94
- Maccarone T. J., Uttley P., van der Klis M., Wijnands R. A. D., Coppi P. S., 2011, *MNRAS*, 413, 1819
- Marcel G. et al., 2018, *A&A*, 617, A46
- Markoff S., Nowak M. A., Wilms J., 2005, *ApJ*, 635, 1203
- Mastroserio G., Ingram A., van der Klis M., 2018, *MNRAS*, 475, 4027
- Mastroserio G. et al., 2021, *MNRAS*, 507, 55
- Matt G., Fabian A. C., Reynolds C. S., 1997, *MNRAS*, 289, 175
- Motta S. E., 2016, *Astron. Nachr.*, 337, 398
- Motta S., Belloni T., Homan J., 2009, *MNRAS*, 400, 1603
- Motta S. E., Casella P., Henze M., Muñoz-Darias T., Sanna A., Fender R., Belloni T., 2015, *MNRAS*, 447, 2059
- Mushtukov A. A., Ingram A., van der Klis M., 2018, *MNRAS*, 474, 2259
- Musoke G., Liska M., Porth O., van der Klis M., Ingram A., 2022, *MNRAS*, 518, 1656
- Nandra K. et al., 2013, preprint (arXiv:1306.2307)
- Nathan E. et al., 2022, *MNRAS*, 511, 255
- Novikov I. D., Thorne K. S., 1973, Black Holes (Les Astres Occlus). Gordon and Breach, New York, p. 343
- Nowak M. A., Vaughan B. A., Wilms J., Dove J. B., Begelman M. C., 1999, *ApJ*, 510, 874
- Poutanen J., Veledina A., Zdziarski A. A., 2018, *A&A*, 614, A79
- Press W. H., Flannery B. P., Teukolsky S. A., Vetterling W. T., 1992, Numerical Recipes in FORTRAN 77: The Art of Scientific Computing, 2 edn. Cambridge Univ. Press, Cambridge
- Ray P. et al., 2019, *Bull. Am. Astron. Soc.*, 51, 231
- Reig P., Kylafis N. D., Giannios D., 2003, *A&A*, 403, L15
- Schnittman J. D., Homan J., Miller J. M., 2006, *ApJ*, 642, 420
- Shakura N. I., Sunyaev R. A., 1973, *A&A*, 500, 33
- Sobolewska M. A., Życki P. T., 2006, *MNRAS*, 370, 405
- Stella L., Vietri M., Morsink S. M., 1999, *ApJ*, 524, L63
- Sunyaev R. A., Truemper J., 1979, *Nature*, 279, 506
- Svensson R., Zdziarski A. A., 1994, *ApJ*, 436, 599
- Timmer J., Koenig M., 1995, *A&A*, 300, 707
- Van den Eijnden J., Ingram A., Uttley P., 2016a, *MNRAS*, 458, 3655
- Van den Eijnden J., Ingram A., Uttley P., Motta S. E., Belloni T. M., Gardenier D. W., 2016b, *MNRAS*, 464, 2643
- Torres M. A. P., Casares J., Jiménez-Ibarra F., Muñoz-Darias T., Padilla M. A., Jonker P. G., Heida M., 2019, *ApJ*, 882, L21
- Tsang D., Butts I., 2013, *MNRAS*, 435, 749
- Uttley P., Malzac J., 2023, preprint (arXiv:2312.08302)
- Uttley P., McHardy I. M., Papadakis I. E., 2002, *MNRAS*, 332, 231
- Uttley P., McHardy I. M., Vaughan S., 2005, *MNRAS*, 359, 345
- Uttley P., Wilkinson T., Cassatella P., Wilms J., Pottschmidt K., Hanke M., Böck M., 2011, *MNRAS*, 414, L60
- Uttley P., Cackett E. M., Fabian A. C., Kara E., Wilkins D. R., 2014, *A&AR*, 22, 72
- Varnière P., Tagger M., Rodriguez J., 2012, *A&A*, 545, A40
- Veledina A., Poutanen J., Vurm I., 2013a, *MNRAS*, 430, 3196
- Veledina A., Poutanen J., Ingram A., 2013b, *ApJ*, 778, 165
- Virtanen P. et al., 2020, *Nat. Methods*, 17, 261
- Wang Y. et al., 2020a, *ApJ*, 896, 33
- Wang J. et al., 2020b, *ApJ*, 899, 44
- Wang J. et al., 2021, *ApJ*, 910, L3
- Wang J. et al., 2022, *ApJ*, 930, 18
- You B., Bursa M., Życki P. T., 2018, *ApJ*, 858, 82
- You B., Życki P. T., Ingram A., Bursa M., Wang W., 2020, *ApJ*, 897, 27
- You B. et al., 2021, *Nature Commun.*, 12, 1025
- You B. et al., 2023, *Science*, 381, 961
- Zdziarski A. A., Dzielak M. A., Marco B. D., Szanecki M., Niedźwiecki A., 2021, *ApJ*, 909, L9
- Zhang S. N. et al., 2016, in den Herder J.-W. A., Takahashi T., Bautz M., eds, Proc. SPIE Conf. Ser. Vol. 9905, Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray. SPIE, Bellingham, p. 99051Q

APPENDIX A: SIMULATIONS

We tested our methods using ‘null hypothesis’ simulated light curves with a constant lag. The broad-band noise is simulated following the method of Timmer & Koenig (1995). We used the multi-Lorentzian fit of the hard band power spectrum of a representative observation, ObsID 144, taken on 2018 May 4, as input for the simulation power spectrum. We removed the narrow Lorentzian components which fitted the QPO fundamental and harmonic. The observed shape of the broad-band noise power spectrum in the data is approximated well in this way.

The original simulated light curve has mean zero and we obtain the final light curve by exponentiating it, $x_e(t) = e^{x_n(t)}$, where $x_e(t)$ is the exponentiated light curve and $x_n(t)$ the original (normally distributed) light curve. Exponentiation prevents light curve values from dropping below zero (which is unphysical) and accounts for the rms–flux relation, which was observed to be ubiquitously present by Heil, Vaughan & Uttley (2012) and which is naturally recreated by exponentiation (Uttley, McHardy & Vaughan 2005). The exponentiation does have a small effect on the width of the power spectrum, which is just visible in Fig. A1. The broadening effect can be significant when the model PSD contains sharp features, but for the broad power distributions we used, the effect is negligible.

The QPOs are simulated using a different approach. The most important aspect is that the QPO signal is interpreted as the sum of two phase-linked sine waves with a phase offset that follows a




Figure A1. The upper panel shows two power spectra. The upper blue spectrum is created by taking the average over 400 simulated segments of 64 s. It is very similar to the input power spectral shape, which is based on the three Lorentzians fitting the broad-band noise of ObsID 144, shown as the black dashed line, except for the peaks that resemble QPOs, which are due to multiplication with a simulated QPO signal. At high frequencies, there is a small positive offset from the input power spectrum, which is the result of exponentiating the light curve to include the rms–flux relation. The contribution of Poisson noise to the power spectrum has been subtracted. The lower red power spectrum is obtained from convolving the ‘hard’ light curve with the impulse response shown in Fig. A3. The lower panel shows the resultant lag–frequency spectrum. The simulations in this figure can be compared to the data in Fig. 1, which look very similar. The light curve corresponding to the ‘hard’ power spectrum is visible in the lower panel of Fig. A2.




Figure A2. Three simulated light curves with a QPO signal. The upper panel shows the QPO component without amplitude modulation, the middle panel contains a simulation with amplitude modulation. The random walk step size used here is 0.04, corresponding to $Q \approx 6$ for $\nu_{\text{QPO}} = 0.3$ Hz and 1/128 s time bins. $\Psi = 0.3$ rad/ π in both cases, which is similar to the value found in the data, and the amplitude of the fundamental is 1.5 times the amplitude of the harmonic. The product of the constant amplitude QPO component and a broad-band noise component is shown in the lower panel, where the light curve has been scaled to realistic values and also includes Poisson noise.




Figure A3. The impulse response is used for simulations of broad-band noise. If this impulse response is applied to the simulated hard band, the hard band will lag behind the soft band at low frequencies (due to the broad structure at negative times), while the sign of the lag is reversed for high frequencies, due to the narrow function at positive times. The dotted red line indicates zero delay. The impulse response has already been binned to the light-curve time-resolution.

random walk, which was also done by Krawczynski et al. (2016) and Ingram & Motta (2019) and is described by

$$x_{\text{QPO}}(t) = 1 + A \cos(2\pi\nu_c - \phi_{\text{rw}}(t)) + B \cos(2[2\pi\nu_c - \phi_{\text{rw}}(t) - \Psi]). \quad (\text{A1})$$

In this equation, $x_{\text{QPO}}(t)$ is the QPO signal, A and B are the amplitudes of the fundamental and the harmonic, respectively, ν_c is the fundamental centroid frequency, Ψ is the constant phase difference between the fundamental and harmonic phase which defines the waveform shape (see equation 3), and ϕ_{rw} is the random walk phase offset of the fundamental. The random walk is

created by $\phi(t_i) = \phi(t_{i-1}) + \epsilon_i$, where $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$. \mathcal{N} is a normal distribution with mean 0 and standard deviation σ equal to the step size of the random walk, which determines how fast the QPO signal goes out of phase with a purely periodic signal. In other words, it determines the width of the QPO peak in the power spectrum, the Q-value. In the figure, the step size is 0.04, while the time-resolution is 1/128 s and the QPO frequency 0.3 Hz, corresponding to $Q \sim 6$. The random walk starts between 0 and 2π .

The result of equation A1 is a double sine waveform with a fluctuating phase (or equivalently, frequency). The phase-linked harmonic signal at twice the centroid frequency enables us to simulate QPO waveforms that are more complicated than a single sine wave, which have been observed in many BHXRBs (De Ruiter et al. 2019). An example of such a waveform is shown in the upper panel of Fig. A2, with phase difference $\Psi = 0.3\pi$ rad.

Both amplitudes A and B can be made dependent on time, $A(t)$ and $B(t)$, to account for the fact that they might vary in observed QPOs as well. Van den Eijnden, Ingram & Uttley (2016a) found that in GRS 1915+105, coherent intervals exist for the QPO signal that seem to be connected to the Q-value. In our simulation, $A(t)$ and $B(t)$ can be described by another quasi-periodic sine wave with a mean period equal to the length of coherent intervals. An example of what this looks like can be seen in the middle panel of Fig. A2. In the rest of this paper, we only show results with a constant QPO amplitude (before being modulated by the broad-band noise, see below), mainly because amplitude modulation is an extra complication that does not change our conclusions for the types of tests we conducted.

We obtain the final combined light curve by multiplying (as opposed to adding) the generated broad-band noise by the QPO signal, which is more consistent with the QPO being due to geometric change and with the observed coupling of QPO and broad-band noise signals (e.g. Heil, Vaughan & Uttley 2010; Maccarone et al. 2011; Arur & Maccarone 2019). The mean of the QPO signal is, in this case, normalized to 1, with the sum of the fundamental and harmonic amplitudes < 1 , in order to prevent negative values. The FT of two multiplied signals is the convolution of their respective FTs. The final product of simulating light curves by multiplying the QPO and broad-band noise components, scaling to the required mean count rate and including Poisson noise is shown in the lower panel of Fig. A2.

The methods for QPO phase-resolved spectral-timing require two correlated light curves from different energy bands with frequency-dependent time lags. We use impulse response functions (henceforth ‘impulse response’) to introduce these lags between two simulated light curves (e.g. Uttley et al. 2014). In our simulations, we first simulate the broad-band noise and QPO signal of the ‘hard band’, because the convolution with the impulse response will lead to a reduction of power, especially in higher frequencies. Fig. 1 clearly shows that the soft band fractional rms power is lower than the hard power spectrum at almost all frequencies in the hard state of MAXI J1820+070. This is also visible in the upper panel of Fig. A1, where we show the input power spectrum and the simulated ‘hard’ and ‘soft’ bands, which can be compared to the upper panel of Fig. 1.

To simulate the ‘soft band’, the impulse response in Fig. A3 is convolved with the simulated ‘hard’ light curve, which consists of both the broad-band noise and the QPO. We use an empirical form of the impulse response given by

$$f(t) = \begin{cases} c_0|t|^{-\alpha} & \text{for } -t_s < t < 0 \\ c_1 & \text{for } t = 0 \\ c_2 & \text{for } t = \delta t \\ 0 & \text{for } t < -t_s \vee t > \delta t, \end{cases} \quad (\text{A2})$$

where c_0 is a normalizing constant, power-law index $\alpha = 0.9$, $t_s = 10$ s is the longest time-scale for which we define lags, c_1 and c_2 are constants determining the amplitude of all lags and soft lags, respectively, and δt is the time-resolution of the light curve. The power-law results in hard lags at low frequencies and the narrow positive component causes the soft lags at high frequencies. The resultant lag–frequency spectrum is very similar to those obtained from the data, as is visible in the lower panel of Fig. A1, which can be compared to the lower panel of Fig. 1.

APPENDIX B: TESTING THE FILTER-INTERPOLATE METHOD WITH QPO WAVEFORM RECONSTRUCTION

To test the FI method presented in Section 3.3, we used it to reconstruct the QPO waveform (modelled using equation 2) and compare the obtained phase difference between fundamental and harmonic Ψ to the value obtained by another method, designed by Ingram & van der Klis (2015) and applied to many systems by De Ruiter et al. (2019). We will refer to the latter approach as the Fourier phase offset (FPO) method from now on. It works by splitting the light curve into many small segments and then estimating the waveform phase difference Ψ from the statistical distribution of observed Fourier phase differences between the fundamental and harmonic frequencies. For a more extensive explanation of the FPO method and its use, we refer to Ingram & van der Klis (2015), Ingram et al. (2017), and De Ruiter et al. (2019). MAXI J1820+070 has not been studied using the FPO method before. For both methods, errors were determined by bootstrapping. More recently, Nathan et al. (2022) introduced a more sophisticated method to calculate the phase difference between the fundamental and the harmonic using the bispectrum. We also applied their method and obtained very similar results.

We applied the FI and FPO methods to the 10 QPO frequency-selected data sets listed in Table 1, to obtain the phase difference for the fundamental versus harmonic Ψ . The results are plotted in Fig. B1. There is a significant (3.7σ for the FI method and $> 5\sigma$ for the FPO method) rise in the phase difference Ψ with QPO frequency. The measured value of $\Psi \sim 0.3$ rad/ π is reasonably consistent with the results from De Ruiter et al. (2019) for high-inclination sources like MAXI J1820+070, even though the general trend for these




Figure B1. The phase difference between the fundamental and the harmonic for all data sets, as defined in Table 1. The x-axis value is the mean QPO frequency of each data set. The large error bars with a QPO frequency between 0.2 and 0.24 Hz arise due to the small amount of lower quality data in this range. It is clear both methods return comparable values for the phase difference between the fundamental and the harmonic. The mean phase difference Ψ when combining all data as measured with the interpolation method is shown as the dashed line.

sources is that the phase difference Ψ decreases with QPO frequency. In their results, however, the scatter is large and the range of QPO frequencies studied in our work, 0.07 to 0.51 Hz, is modest compared to the total range covered by De Ruiter et al. (2019). The phase difference values measured with both methods are consistent given the statistical errors,⁹ which gives us confidence that the FI method is correctly resolving the QPO signal, with no bias in the resulting waveform shape.

APPENDIX C: RED NOISE LEAKAGE

The frequency range for which filtered light curves were created in Section 3.2 is limited by red noise leakage, which creates a bias (Alston, Vaughan & Uttley 2013). In order to mitigate the effect of the bias, we only made filtered light curves between 0.4 and 2.5 times the QPO frequency of each observation, corresponding to a maximum central filter frequency of ~ 1.25 Hz. For the interpolation method, the

⁹Using Pearson's χ^2 -test to compare the results from both methods for all data sets, we obtain a probability that the differences are due to statistical fluctuations of 0.007 (2.7σ). The low probability arises due to the relatively small error bars obtained by bootstrapping for some of the data sets.

filter range was even smaller, ranging from 0.6 to $2\nu_{\text{QPO}}$. In Section 3.4, the consequences of the Fourier leakage effects described below for the lag waveform are reported in more detail.

The bias can be understood by thinking of the average shape of the light-curve slices in each bin. When binning short light-curve slices by their filtered flux, such as in Section 3.2, the slices with the lowest and highest filtered flux will, on average, contain a peak or trough in the flux and will not contain any strong overall flux-trend. The light-curve slices in other bins will, on average, contain a slope that is due to variability at frequencies below the sampled Fourier frequency range, which induces leakage of these lower frequencies and decreases the measured (absolute) time lags, since the trend will appear the same in both bands.

The rms–flux relation affects the leakage in such a way that the highest filtered flux bin has larger absolute lags than the lowest flux bin. The end result of the different light curve shapes in each flux bin is that there seems to be a linear relation between the short-term time lags and the hard flux, at least when the period of the filter is not at least a few times larger than the slice length (0.25 s). At the same time, light curves filtered at high frequencies are not well correlated with the total flux, leading to only a small range in hard flux. The combination of a modest range in lags and a small hard flux range introduces steep slopes. Simulated light curves show very similar results when using filtered light curves at high Fourier frequencies (>2 Hz for 0.25 s slices), while this is not the case for lower filter frequencies, indicating that the slopes arise due to systematics. We therefore do not show results for these higher frequency filters in this paper.

APPENDIX D: OBSERVATION LIST

Table D1. The ObsIDs used in this paper and the observation time obtained by excluding GTIs of less than 64 s. Also, the average QPO fundamental frequency in each observation and the mean count rates for the soft 0.5–1 and hard 3–10-keV energy bands are shown. The total observation time is 158 592 s.

ObsID	ν_{QPO} (Hz)	Selected exposure time (s)	Flux 0.5–1 keV (count s ⁻¹)	Flux 3–10 keV (count s ⁻¹)
126	0.077	1280	7143	1350
127	0.082	640	7011	1306
129	0.094	448	7212	1338
130	0.12	3712	6863	1303
131	0.12	1920	6141	1159
132	0.14	1920	7029	1310
133	0.14	1792	6972	1292
134	0.14	4928	7027	1307
135	0.17	3200	7118	1311
136	0.17	1344	7169	1304
137	0.17	5760	6963	1264
138	0.18	3264	6962	1260
139	0.17	512	6896	1244
140	0.25	2176	7187	1250
141	0.27	832	7132	1211
142	0.28	4800	7029	1184
143	0.29	3648	7030	1167
144	0.31	4608	7083	1170
145	0.34	5888	7149	1159
146	0.35	4928	7064	1142
147	0.39	4544	7272	1146
148	0.41	3776	7374	1142
149	0.45	2048	7517	1149
150	0.46	1600	7581	1130
151	0.42	704	7250	1119
152	0.46	1984	7356	1101
153	0.46	768	7465	1104
154	0.50	448	7520	1084
155	0.51	576	6899	1008
156	0.49	4800	6758	980
157	0.41	4608	6158	931
158	0.35	2176	5701	891
159	0.34	1664	5477	859
160	0.34	2048	5224	824
161	0.34	3264	5260	827
162	0.34	2240	5147	807
163	0.42	2816	5466	821
164	0.41	1856	5316	803
165	0.38	2624	5034	774
166	0.37	1984	4907	756
167	0.35	1856	4752	732
168	0.35	960	4699	729
169	0.31	1664	4242	669
170	0.34	3712	4302	668
171	0.32	2176	4011	630
172	0.35	2112	3996	617
173	0.28	3008	3642	576
174	0.28	2560	3427	546
175	0.31	2112	3260	516
176	0.27	5696	3097	490
177	0.24	704	2984	469
178	0.27	2432	2732	436
179	0.31	832	2662	420
180	0.25	2112	2536	403
181	0.24	384	2251	373
182	0.22	1856	2252	365
183	0.24	640	2262	367
184	0.24	640	2246	363
185	0.22	1728	2224	364
186	0.27	1664	2669	419
187	0.29	1280	2891	450
188	0.36	1536	3754	567
189	0.42	12 800	4226	615

¹*Anton Pannekoek Institute for Astronomy, Amsterdam, Science Park 904, NL-1098 NH, the Netherlands*

²*Department of Physics, Astrophysics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK*

³*School of Mathematics, Statistics and Physics, Newcastle University, Herschel Building, Newcastle upon Tyne NE1 7RU, UK*

⁴*Department of Physics, University of Warwick, Coventry CV4 7AL, UK*

⁵*Center for Astro, Particle and Planetary Physics, New York University Abu Dhabi, Abu Dhabi 129188, United Arab Emirates*

⁶*School of Physics and Astronomy, University of Southampton, Southampton SO17 1BJ, UK*

⁷*X-Ray Astrophysics Laboratory, NASA Goddard Space Flight Center, Code 662, Greenbelt, MD 20771, USA*

⁸*Independent*

⁹*Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK*

¹⁰*Department of Astronomy, University of Maryland, College Park, MD 20742, USA*

¹¹*CRESST II, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA*

¹²*NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA*

¹³*Eureka Scientific, Inc., 2452 Delmer Street, Oakland, CA 94602, USA*

¹⁴*MIT Kavli Institute for Astrophysics and Space Research, MIT, 77 Massachusetts Avenue, Cambridge, MA 02139, USA*

¹⁵*Dipartimento di Fisica, Universita' degli Studi di Cagliari, SP Monserrato-Sestu km 0.7, I-09042 Monserrato, Italy*

¹⁶*Center for Astrophysics, Harvard & Smithsonian, 60 Garden St, Cambridge, MA 02138, USA*

¹⁷*Physics Department, Tor Vergata University of Rome, Via della Ricerca Scientifica 1, I-00133 Rome, Italy*

¹⁸*INAF – Astronomical Observatory of Rome, Via Frascati 33, I-00040 Monte Porzio Catone, Italy*

¹⁹*INFN – Rome Tor Vergata, Via della Ricerca Scientifica 1, I-00133 Rome, Italy*

²⁰*National Astronomical Observatories, Chinese Academy of Sciences, 20A Datun Road, Beijing 100101, China*

²¹*HEASARC, Code 6601, NASA/GSFC, Greenbelt, MD 20771, USA*

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.