



| | |
|----------------------------------|---|
| Publication Year | 2025-08-18 |
| Acceptance in OA | 2025-09-05T13:58:42Z |
| Title | A NuSTAR Study of Quasiperiodic Oscillations from the Ultraluminous X-Ray Sources in M82 |
| Authors | EL BYAD, Hamza, BACHETTI, Matteo, Silvia Columbu, Giuseppe Rodriguez, PILIA, Maura, Matthew J. Middleton, Dominic J. Walton, Murray Brightman, Hannah Earnshaw, Karl Forster, Brian Grefenstette, Felix Fürst, Marianne Heida, IMBROGNO, Matteo, LAI, Eleonora Veronica, Thomas Maccarone |
| Publisher's version (DOI) | 10.3847/1538-4357/aded8a |
| Handle | http://hdl.handle.net/20.500.12386/43578 |
| Journal | THE ASTROPHYSICAL JOURNAL |
| Volume | 989 |



A NuSTAR Study of Quasiperiodic Oscillations from the Ultraluminous X-Ray Sources in M82

Hamza El Byad^{1,2}, Matteo Bachetti¹, Silvia Columbu², Giuseppe Rodriguez², Maura Pilia¹, Matthew J. Middleton³, Dominic J. Walton⁴, Murray Brightman⁵, Hannah Earnshaw⁵, Karl Forster⁵, Brian Grefenstette⁵, Felix Fürst⁶, Marianne Heida⁹, Matteo Imbrogno⁷, Eleonora Veronica Lai¹, and Thomas Maccarone⁸

¹INAF-Osservatorio Astronomico di Cagliari, Italy

²Department of Mathematics and Computer Science, University of Cagliari, Italy

³University of Southampton, UK

⁴Centre for Astrophysics Research, University of Hertfordshire, UK

⁵California Institute of Technology, USA

⁶ESA/ESAC, Spain

⁷INAF-Osservatorio Astronomico di Roma, Italy

⁸Department of Physics & Astronomy, Texas Tech University, Box 41051, Lubbock, TX 79409-1051, USA

Received 2025 May 19; revised 2025 July 4; accepted 2025 July 7; published 2025 August 18

Abstract

The study of quasiperiodic oscillations (QPOs) in X-ray binaries provides valuable insights into the physics of accretion around compact objects. The M82 galaxy hosts multiple ultraluminous X-ray sources, including two prominent ones—X-1 and X-2—where X-1 is suspected to harbor an intermediate-mass black hole. In this work, we analyze data from 39 NuSTAR observations acquired between 2014 and 2024 to investigate the aperiodic X-ray variability in M82. In particular, we study in detail the evolution of the QPOs from M82 X-1 in the range 20–300 mHz. We do not find additional timing features in the data, besides a frequently present broad noise component at lower frequencies. The QPO behaves similarly to other classes of low-frequency oscillations in accreting compact objects, both black holes and neutron stars.

Unified Astronomy Thesaurus concepts: [Ultraluminous x-ray sources \(2164\)](#); [X-ray binary stars \(1811\)](#); [X-ray astronomy \(1810\)](#); [Compact objects \(288\)](#); [Accretion \(14\)](#); [Time domain astronomy \(2109\)](#)

1. Introduction

One of the most intriguing phenomena in the accretion processes of X-ray binaries is the presence of quasiperiodic oscillations (QPOs; see M. van der Klis et al. 1989; A. R. Ingram & S. E. Motta 2019 for reviews). They are nearly periodic variations manifesting in the form of relatively coherent oscillations in the X-ray emissions from accretion disks surrounding neutron stars (NSs) and black holes (BH). These oscillations have been observed across a wide range of frequencies in X-ray binary systems. They are generally categorized into two types: low-frequency QPOs (LFQPOs) and high-frequency QPOs (HFQPOs). LFQPOs have been detected in the majority of BH binaries (BHBs) observed with the Rossi X-ray Timing Explorer (RXTE; H. V. Bradt et al. 1993), with frequencies reaching up to ~ 30 Hz (T. Belloni et al. 2005), and up to ~ 60 Hz in NS binaries (NSBs; M. van der Klis 2006). HFQPOs have centroid frequencies spanning approximately 60–500 Hz in stellar-mass BHBs, ranging from the ~ 60 to 70 Hz features detected in GRS 1915+105 (T. M. Belloni & D. Altamirano 2013) and IGR J17091–3624 (D. Altamirano & T. Belloni 2012) to the $\gtrsim 400$ Hz signals observed in sources such as GRO J1655–40 and XTE J1550–564, and from several hundred hertz to over 1 kHz in NSBs with weak magnetic fields

(M. van der Klis 2006). These variations highlight some of the differences in the variability of BHBs and NSBs, reflected in the distinct variability patterns of these systems. By studying these oscillations, researchers can gain insight into the fundamental physics of compact objects, including their masses, radii, and spins (W. Kluźniak 2006), and the differentiation between different types of compact objects, such as BHs and NSs. For example, certain frequency patterns may indicate specific properties unique to NSs or quark stars (N. Shaposhnikov 2012). The characteristics of QPOs can also offer critical information about the dynamics of matter accretion, the structure of the inner accretion disk, and the strong gravitational environments near compact objects (R. A. Remillard & J. E. McClintock 2006; C. M. Zhang et al. 2006; T. M. Belloni & L. Stella 2012). HFQPOs, in particular, have been suggested to originate from the relativistic precession model (L. Stella & M. Vietri 1998, 1999) or from nonlinear resonances between orbital epicyclic modes—specifically the radial and vertical oscillation frequencies—within the inner accretion disk (M. A. Abramowicz & W. Kluźniak 2001; W. Kluźniak & M. A. Abramowicz 2001; S. E. Motta 2016). Thus, these oscillations serve as a valuable tool for probing the physics of accretion and the fundamental properties of compact objects.

The M82 galaxy harbors a number of ultraluminous X-ray sources (ULXs), off-nuclear X-ray sources where the accretion process implied by their luminosity exceeds the theoretical Eddington limit (P. Kaaret et al. 2017; A. King et al. 2023). M82 X-2 was revealed as the first ULX pulsar when NuSTAR detected coherent 1.37 s pulsations at $L_X \simeq 10^{40}$ erg s⁻¹, and a 2.5 day orbit with a semimajor axis of 22 lt-s (M. Bachetti et al. 2014, hereafter B14). A long-term tracking of

⁹ Current affiliation: Rijksinstituut voor Volksgezondheid en Milieu, Anthonie van Leeuwenhoeklaan 9, 3721 MA Bilthoven, The Netherlands.



the pulsar revealed that the orbit is decaying at a rate of $-5.69(24) \times 10^{-8}$ days days $^{-1}$ (M. Bachetti et al. 2022). Only $\sim 5''$ away lies M82 X-1, long considered an intermediate-mass BH (IMBH) candidate: its persistent $L_X \gtrsim 10^{40}$ erg s $^{-1}$ and 50–300 mHz QPOs (T. E. Strohmayer & R. F. Mushotzky 2003; G. C. Dewangan et al. 2006; P. Mucciarelli et al. 2006), plus a tentative 3.3–5.1 Hz 3:2 pair (D. R. Pasham et al. 2014), imply a mass of $\sim 10^2$ – $10^3 M_\odot$ (see also H. Feng & P. Kaaret 2007; M. Brightman et al. 2020). Both ULXs lie in a crowded field that includes X-3 and X-4, making their subarcsecond Chandra imaging essential for disentangling their spectra and timing.

M82 is a nearly edge-on, “cigar-shaped” starburst with a star formation rate of $\sim 10 M_\odot \text{ yr}^{-1}$ (K. Iwasawa 2021) at a distance of 3.6 Mpc (W. L. Freedman et al. 1994). Although it hosts a low-luminosity active galactic nucleus (H. Matsumoto & T. G. Tsuru 2001), its hard-X-ray output is dominated by the ULXs. X-1 can reach $L_X \sim 10^{41}$ erg s $^{-1}$, yet the proximity of X-2 complicates broadband studies; only Chandra ($\sim 0.5''$) angular resolution cleanly separates them, but its 0.3–10 keV bandpass limits hard-X-ray spectroscopy of these sources.

Timing analysis of M82’s X-ray emission has provided strong evidence that M82 X-1 is a promising IMBH candidate. Initial evidence emerged from the detection of a QPO at 54 mHz associated with M82 X-1 using XMM-Newton data from T. E. Strohmayer & R. F. Mushotzky (2003). Building on this research, P. Mucciarelli et al. (2006) and G. C. Dewangan et al. (2006) reported QPOs between 50 and 166 mHz using combined XMM-Newton and RXTE archival observations. By examining the relationship between this QPO frequency and the photon index of the energy spectrum, they estimated the BH mass of M82 X-1 to be between 25 and 520 M_\odot . The most compelling evidence for M82 X-1 harboring an IMBH came from D. R. Pasham et al. (2014), who reported stable twin-peak QPOs with a 3:2 frequency ratio at 3.3 and 5.1 Hz using RXTE data. This frequency pattern, which is characteristic of HFQPOs in stellar-mass BHs but scaled down to lower frequencies, is potentially a signature of a BH of approximately 400 M_\odot . This would place M82 X-1 firmly in the intermediate-mass range—a critical “missing link” between stellar-mass and supermassive BHs. K. Atapin et al. (2019) analyzed the X-ray power density spectra (PDSs) of several ULXs, including M82 X-1, and found that the QPO frequencies are anticorrelated with the level of flat-topped noise, suggesting that mass accretion rate variations influence the observed variability (see also M. J. Middleton et al. 2015a).

Adding to the complexity in the high-energy regime in M82, Chandra observations revealed LFQPOs (3–4 mHz) from M82 X-2 (H. Feng et al. 2010). These findings underscore the diversity of ULXs within M82, each with different timing properties, thereby offering additional avenues to study varied accretion mechanisms.

In this work, we analyze observational data acquired with NASA’s NuSTAR (F. A. Harrison et al. 2013) satellite spanning from 2014 to 2024, to investigate the variability of the X-ray flux in the M82 galaxy.

2. Data Reduction

2.1. NuSTAR

We used data from all available NuSTAR observations of the M82 galaxy between 2014 and 2024, reduced with the

standard pipeline (`nupipeline` from NuSTARDAS)¹⁰ from the High Energy Astrophysics Science Archive Research Center.¹¹ Events from both focal-plane modules (FPMA and FPMB) were retained to maximize the signal-to-noise ratio, and photons were extracted from a circular region of 70'' radius centered on the source position. The full list of observations is provided in Table 1. Our data consist of individual photon events, each with a detection time (time stamp) and associated properties such as photon energy. Each time stamp was corrected from local time to the barycenter of the solar system using the `barycorr` FTOOL, using the International Celestial Reference System coordinates of M82 X-2, 09:55:51.040, +69:40:45.49¹² (P. Kaaret et al. 2006), and using the DE430 JPL ephemeris.¹³

We first computed the raw power-spectral density for each observation to check for telemetry issues and to define an adequate frequency grid for subsequent fitting. For every ObsID, we combined the barycentre-corrected FPMA + FPMB events into a 3–79 keV light curve binned at 0.1 s, sliced it into contiguous 512 s segments, and Leahy-normalized the periodogram of each segment; the average of all segments yields a minimum Fourier frequency of $\simeq 1.95$ mHz and a Nyquist frequency of 5 Hz. The primary analysis used only events recorded inside good-time intervals (GTIs) with full star-tracker visibility (mode-01 data). For a limited part of the study, we made use of photons from intervals with limited star-tracker coverage (the SCIENCE_SC mode, or mode-06, data). We used the tool `nusplitsc` to split the mode-06 data into intervals with single star-tracker combinations; we moved the extraction region for events in each subinterval to adapt to the position of the point-spread function (PSF) centroid; we extracted the events from the source region; and finally, we merged the events together with the mode-01 data. Given the circumpolar position of the source and NuSTAR’s near-equatorial orbit, this procedure nearly doubled the number of photons available for sensitive QPO searches, at the cost of increasing the red-noise level. Consequently, we used the merged data set only to increase the sensitivity of our search for new features at high frequencies where the red noise was negligible.

We ran the analysis using `Stingray` (D. Huppenkothen et al. 2019), a Python library built to perform time series analysis, providing implementations of the most advanced spectral timing techniques available in the literature.

Raw data do not provide the photon energy directly. However, in NuSTAR the energy channel number (PI) and the central energy E_p of the channel are related by the simple formula $E_p(\text{keV}) = 1.62 + 0.04 \text{ PI}$. This is accounted for automatically when loading data in `Stingray`.

We plotted light curves of the source region of all observations, and background light curves containing all photons further than 100'' from the source. This criterion is different from the one usually employed for spectral analysis, where a background region is chosen in a large circular region of the field of view (FOV) devoid of sources. In our case, these background light curves were meant to catch flares in the

¹⁰ https://heasarc.gsfc.nasa.gov/docs/nustar/analysis/nustar_swguide.pdf

¹¹ <https://heasarc.gsfc.nasa.gov>

¹² Data were barycentered in order to allow the study of aperiodic variability from M82 X-1 and pulsations from M82 X-2. A 5'' mismatch is irrelevant for the study of slow variability, but it might be detectable in precise pulsar timing.

¹³ https://naif.jpl.nasa.gov/pub/naif/generic_kernels/spk/planets/aareadme_de430-de431.txt

Table 1Summary Table with All the NuSTAR Data Sets in This Work, Containing Best-fit Values and 1σ Uncertainties for the Red-noise and Quasiperiodic Oscillation Components Obtained in Section 3.3 and Information on the Observing Conditions

| ObsID | Date (MJD) | Exposure (ks) | rms_{rn} (counts s^{-1}) | $\nu_{\text{max,m}}$ (mHz) | rms_{qpo} (counts s^{-1}) | $\nu_{0,\text{qpo}}$ (mHz) | w_{qpo} (mHz) | Q_{qpo} | Notes |
|-------------|------------|---------------|--|----------------------------|---|----------------------------|------------------------|------------------|--------------|
| 30101045002 | 57493.3 | 195.5 | 0.122(4) | 3.48(29) | 0.161(6) | 114.3(14) | 42(4) | 5.4(6) | Q, P |
| 30202022002 | 57542.9 | 40.3 | ... | ... | ... | ... | ... | ... | BF |
| 30202022004 | 57570.7 | 48.2 | 0.069(30) | 2.1(12) | ... | ... | ... | ... | ... |
| 30202022008 | 57599.0 | 43.7 | 0.055(24) | 0.9(7) | ... | ... | ... | ... | BF, SF |
| 30202022010 | 57619.4 | 44.1 | 0.055(20) | 1.3(6) | ... | ... | ... | ... | BF |
| 30502020002 | 58691.9 | 93.0 | ... | ... | -0.256(11) | 275(16) | 400(50) | 1.39(27) | Q |
| 30502020004 | 58701.9 | 91.5 | 0.135(9) | 0.98(20) | ... | ... | ... | ... | SF |
| 30502021002 | 58918.1 | 86.3 | ... | ... | -0.194(8) | 116(5) | 116(15) | 2.01(34) | Q, P |
| 30502021004 | 58929.2 | 79.8 | 0.162(21) | 128(35) | 0.114(19) | 185.7(25) | 36(10) | 10.1(29) | BF, Q |
| 30502022002 | 59000.8 | 90.9 | 0.120(8) | 6.8(11) | 0.117(8) | 46.6(6) | 10.9(19) | 8.6(16) | SF, Q |
| 30502022004 | 59012.6 | 99.2 | 0.096(7) | 2.8(6) | 0.184(9) | 64(4) | 69(11) | 1.8(4) | BF, Q |
| 30602027002 | 59311.9 | 73.6 | 0.074(11) | 5.1(23) | 0.196(12) | 165(8) | 140(27) | 2.4(6) | Q, P |
| 30602027004 | 59325.3 | 71.6 | 0.159(19) | 67(17) | 0.115(12) | 91.3(8) | 15.2(33) | 11.0(27) | BF, Q, P |
| 30602028002 | 59215.2 | 69.0 | ... | ... | -0.192(8) | 159(6) | 136(18) | 2.3(4) | BF, Q |
| 30602028004 | 59226.9 | 70.8 | 0.038(14) | 0.52(25) | ... | ... | ... | ... | BF |
| 30702012002 | 59504.0 | 128.3 | 0.106(6) | 4.1(6) | 0.163(9) | 135(4) | 75(11) | 3.6(6) | Q, P |
| 30702012004 | 59674.3 | 124.4 | 0.13(5) | 130(50) | 0.15(4) | 143(9) | 110(40) | 2.7(12) | Q |
| 30901038002 | 60110.6 | 128.1 | 0.108(7) | 0.92(18) | ... | ... | ... | ... | BF, SF |
| 31001019002 | 60657.4 | 133.9 | 0.089(5) | 3.0(4) | 0.215(11) | 171(10) | 240(40) | 1.46(31) | BF, Q, P |
| 50002019002 | 57037.9 | 32.8 | ... | ... | -0.130(8) | 31.9(16) | 22(4) | 2.8(7) | BF |
| 50002019004 | 57041.8 | 168.2 | 0.101(11) | 12.3(28) | 0.114(9) | 36(14) | 14.1(20) | 5.2(27) | BF, SF, Q |
| 80002092002 | 56680.5 | 67.4 | 0.130(18) | 12(4) | 0.095(19) | 43(16) | 6.7(15) | 12(8) | Q, P |
| 80002092004 | 56682.8 | 92.3 | 0.122(11) | 11.7(28) | 0.103(11) | 50.4(6) | 9.8(17) | 10.3(20) | BF, Q, P |
| 80002092006 | 56685.5 | 321.1 | 0.103(4) | 6.7(7) | 0.127(4) | 59.7(7) | 26.3(22) | 4.5(4) | BF, Q, P |
| 80002092007 | 56692.2 | 319.3 | 0.086(6) | 8.4(14) | 0.152(5) | 71.9(16) | 56(5) | 2.56(29) | BF, Q, P |
| 80002092008 | 56698.8 | 35.2 | 0.172(19) | 68(19) | ... | ... | ... | ... | BF, Q, P |
| 80002092009 | 56699.5 | 119.6 | 0.114(12) | 30(8) | 0.124(8) | 98.3(12) | 23(4) | 8.4(15) | BF, Q, P |
| 80002092011 | 56719.7 | 114.5 | 0.118(15) | 49(17) | 0.099(9) | 65.6(5) | 11.8(17) | 11.1(17) | BF, SF, Q, P |
| 80202020002 | 57413.8 | 37.9 | 0.079(18) | 1.6(6) | ... | ... | ... | ... | SF |
| 80202020004 | 57441.7 | 32.4 | 0.269(35) | 230(60) | ... | ... | ... | ... | BF |
| 80202020006 | 57483.4 | 31.8 | ... | ... | -0.184(9) | 47.6(30) | 56(9) | 1.7(4) | BF, SF |
| 80202020008 | 57502.8 | 41.7 | ... | ... | -0.153(12) | 154.1(22) | 28(6) | 10.0(26) | BF |
| 80202020008 | 57502.8 | 41.7 | ... | ... | -0.153(12) | 154.1(22) | 28(6) | 10.0(26) | BF, SF |
| 90101005002 | 57193.6 | 38.8 | 0.063(32) | 0.8(5) | ... | ... | ... | ... | BF, SF |
| 90201037002 | 57641.5 | 82.5 | 0.265(35) | 440(130) | ... | ... | ... | ... | P |
| 90202038002 | 57668.8 | 46.0 | ... | ... | -0.23(4) | 30(100) | 960(320) | 0.07(23) | BF, Q |
| 90202038004 | 57722.6 | 45.0 | ... | ... | -0.213(13) | 222(10) | 166(31) | 2.7(6) | BF, Q |
| 90901332002 | 60263.9 | 71.9 | 0.117(10) | 4.9(10) | ... | ... | ... | ... | ... |
| 90901333002 | 60275.3 | 53.7 | 0.139(8) | 6.2(9) | 0.157(10) | 131.0(19) | 33(5) | 7.7(13) | BF, Q |

Note. The Notes column indicates when observations are affected by background flaring (BF), when the flux is highly variable during the observation (HV), and when QPOs (Q) and/or pulsations (P) are detected by M. Bachetti et al. (2022) and M. Bachetti et al. (2025, in preparation).

background, corresponding to anything from particles hitting the telescope or other increased environmental background (e.g., approaching the South Atlantic Anomaly (SAA)), and we needed to gather all the photons we could at a reasonable distance from the source region. We rescaled the background light curve by multiplying by the ratio of pixels inside the source and background regions. Whenever the rescaled background light curve had flares reaching more than $\sim 10\%$ of the source mean flux, we rejected that time interval and excluded it from the GTIs. See Figure 1 for an example.

2.2. Chandra

We used a script based on `astroquery` (A. Ginsburg et al. 2019) to select all Chandra observations of M82 executed




Figure 1. Example of the GTI cleaning process. The light curves shown (background and source) are constructed from the coadded FPMA + FPMB data, binned at $\Delta t = 0.1$ s. Intervals where the background exceeds 10% of the mean source level are excluded (black bands). Standard Earth-occultation or poor star-tracker intervals are shown with blue vertical bands, and data taken during limited star-tracker coverage are plotted in red (these red data are not used in PDS modeling).

Table 2
Quasi-simultaneous Chandra and NuSTAR Observations Used for the Quasiperiodic Oscillation Identification

| ObsID | Instrument | Date (MJD) | Exposure (ks) | Simultaneous NuSTAR ID |
|-------|------------|------------|---------------|------------------------|
| 16580 | ACIS-S | 56691.8 | 47.5 | 80002092007 |
| 17578 | ACIS-S | 57038.6 | 10.1 | 50002019002 |
| 16023 | ACIS-S | 57042.0 | 10.1 | 50002019004 |
| 18064 | ACIS-I | 57483.7 | 25.1 | 80202020006 |
| 18068 | ACIS-I | 57502.8 | 25.1 | 80202020008 |
| 18070 | ACIS-I | 57669.0 | 25.1 | 90202038002 |
| 18072 | ACIS-I | 57723.4 | 25.6 | 90202038004 |
| 26664 | ACIS-S | 60274.5 | 40.1 | 90901333002 |

within 1 day of a NuSTAR ObsID in the same field. The search returned the coincidences shown in Table 2. These observations are part of the Chandra Data Collection at DOI: [10.25574/cdc.427](https://doi.org/10.25574/cdc.427). We downloaded the Chandra data of these observations and ran the `chandra_repro` script distributed with the CIAO software (A. Fruscione et al. 2006) to produce level-2 cleaned event files. We barycentered the event files similarly to what we did for NuSTAR data. We selected photons from circular or oval regions approximating the shape of M82 X-1 and M82 X-2 in the images, typically 1'' to 3'' wide. The change of shape was particularly accentuated in off-axis observations, that in turn were less affected by pileup. M82 X-1 and M82 X-2 never overlapped, allowing easy separation of the photons from the two sources. M82 X-2 had a slight overlap with the very nearby M82 X-3 and X-4, which are however usually fainter than M82 X-2 if the latter is not in a low state (M. Brightman et al. 2019).

3. Timing Analysis

3.1. Statistical Properties of the Periodogram

Periodogram-based methods are commonly used as a nonparametric approximation of the PDS. Given the discrete Fourier transform components a_i of a light curve x_n ($n = 0, \dots, N$), where

$$a_i = \frac{1}{N} \sum_{n=0}^{N-1} x_n e^{-j2\pi ni/N}, \quad (i = -N/2, \dots, N/2 - 1), \quad (1)$$

the periodogram is defined as $P_i = |a_i|^2$.

In X-ray astronomy, it is common to use the normalization from D. Leahy et al. (1983), where the periodogram defined above is multiplied by a factor $2/N_{\text{ph}}$, where N_{ph} is the number of photons in the light curve. With this normalization, the powers of a periodogram of pure white noise follow a χ^2_2 distribution, allowing for easy identification of outliers. A common procedure to limit the noise of the periodogram is based on power averaging, either of W nearby bins from the same periodogram, or M periodograms from different segments of the data (the so-called Bartlett periodogram; M. S. Bartlett 1950), or using both methods. It is easy to demonstrate that the effect of averaging MW noise powers leads to a normalized χ^2_{2MW}/MW distribution, which resembles increasingly more of a Gaussian distribution with width $\sigma = 2/\sqrt{MW}$ (M. van der Klis et al. 1989) as the number of averaged powers increases.

The Bartlett periodogram has the additional major advantage of being applicable to observations containing missing

data, for example due to Earth occultation, SAA passes, high background, and so on. In X-ray observations, these ‘‘bad intervals’’ are usually eliminated from the observations during the data reduction procedure, and good observing conditions are encoded in a GTI list in the same FITS files of the data. The Bartlett periodogram can then be chosen so that one or more intervals of duration t_{seg} fit inside the typical length of GTIs.

The powers I_j of a periodogram containing signal are distributed following a χ^2_{2MW} around the real spectrum S_j (D. Barret & S. Vaughan 2012):

$$I(f_j) = \frac{S(f_j)}{2MW} X, \quad (2)$$

where X follows a χ^2_{2MW} distribution.

From the properties of the χ^2_n distribution, the log-likelihood formula in this general case can be derived as follows:

$$\log \mathcal{L} = \frac{-\nu}{2} \sum_{j=1}^{N-1} \left\{ \frac{I_j}{S_j} + \ln S_j + \left(\frac{2}{\nu} - 1 \right) \ln I_j + c(\nu) \right\}, \quad (3)$$

where $\nu = 2MW$ is the number of degrees of freedom and $c(\nu)$ is a constant for fixed ν . Since minimization algorithms are more common than maximization ones, maximum-likelihood fitting procedures usually consist of minimizing the quantity $-2 \log \mathcal{L}$.

The Bartlett periodogram calculated from a typical fast Fourier transform (FFT; J. W. Cooley & J. W. Tukey 1965) has a major limitation in terms of frequency resolution, limited to $\Delta\nu = 1/t_{\text{seg}}$. This also means that we cannot investigate frequencies lower than $\Delta\nu$. Some techniques to improve the frequency resolution, such as interbinning or Fourier interpolation (S. M. Ransom et al. 2002), can be used, but at the cost of altering the statistical properties of the data and making the fitting and interpretation of the results less robust.

On the other hand, a single periodogram of the whole time series has a much better frequency resolution and sensitivity to low frequencies (as now the resolution is $\Delta\nu = 1/t_{\text{obs}}$), but it is affected by missing data and contains a large number of (typically) low-frequency peaks that correspond to the variability introduced by the visibility windows. If the bad time intervals (BTIs) are very small (e.g., less than 1% of the data), it is customary to add some white noise to fill up the intervals and have a final periodogram with all the desired statistical properties, at the expense of some minor loss of sensitivity. However, in our observations, the bad intervals are comparable in length with the good ones, which would imply simulating about half the data, which is unacceptable for our purposes.

For this work, we devised a treatment for the periodogram that limits the effects of windowing while maintaining most of the statistical properties of the periodogram (see Appendix A for details). An example of one of our cleaned periodograms is depicted in Figure 2.

3.2. Model Construction

The periodograms of X-ray binaries show a variety of variable phenomena, and can be conveniently modeled through a composition of Lorentzian components (T. Belloni et al. 2002).




Figure 2. Example analysis using ObsID 80002092006. Data are cleaned as described in Appendix A. The light curve (FPMA + FPMB, binned at 0.1 s) was split into 512 s segments to compute Leahy-normalized periodograms. The top panel shows the periodogram before filling gaps, and the bottom panel shows the periodogram after filling gaps; in both panels, the orange bands mark frequencies excluded by notch filtering (orbital harmonics). The blue curves are the best-fit Lorentzian models (broad noise component + QPO). The frequency axis runs from approximately 1.95 mHz ($1/512$ s) up to the Nyquist frequency of 5 Hz ($\frac{1}{2 \times 0.1}$ s).

For these functions, we use the definition

$$P(\nu) = \frac{A_0 \left(\frac{w}{2}\right)^2}{(\nu - \nu_0)^2 + \left(\frac{w}{2}\right)^2}, \quad (4)$$

where ν_0 is the centroid frequency or the frequency at the peak of the signal, w represents the FWHM, and A_0 is the amplitude of the signal.

The use of Lorentzian components is mostly phenomenological given the symmetry properties of these functions, even though it is originally rooted in the fact that Lorentzians are the Fourier transforms of exponentially decaying oscillations, a common phenomenon in nature.

These Lorentzian components can be characterized through three main quantities. First is the characteristic frequency (ν_{\max}):

$$\nu_{\max} = \sqrt{\nu_0^2 + \left(\frac{w}{2}\right)^2},$$

which represents the peak of the Lorentzian component in a νP versus ν plot and the frequency at which the Lorentzian contains the most power per logarithmic frequency interval (T. Belloni et al. 2002). This quantity is very close to ν_0 for coherent QPOs, while it departs considerably from it for broadband red-noise components. It is particularly useful when considering the evolution of timing features, as it is common to observe a broadband noise component evolve into a QPO. The characteristic frequency, in this case, makes a smooth transition that would not be as clear when using the central frequency (e.g., S. E. Motta 2016). The second important quantity is the quality factor (Q): defined as $Q = \frac{\nu_0}{\text{FWHM}} = \frac{2\nu_0}{w}$, it measures the signal's coherence.¹⁴ Commonly, Lorentzian components with $Q > Q_{\text{lim}}$

¹⁴ Many works in the literature divide by the FWHM, so their values of Q would be 50% of the ones calculated here.

are classified as QPOs, while those with $Q < Q_{\text{lim}}$ are considered broadband peaked noise, with different choices of Q_{lim} in different papers. In this work we will use $Q_{\text{lim}} = 2$ as the boundary between a QPO and a broadband component. However, it will become clear that the feature we identify as a QPO can sometimes have low coherence. Finally, we can define the (fractional or absolute) rms amplitude: a measure of the signal's strength, which depends on the source flux. It is proportional to the square root of the integrated power contributed by the QPO to the periodogram. Using Equation (4), the rms amplitude can be calculated as the square root of the integral of $P(\nu)$ normalized in the desired rms units (fractional; e.g., T. Belloni & G. Hasinger 1990; or absolute, in counts per second). Since power is only calculated at positive frequencies, assuming $P(\nu)$ was fit in Leahy normalization, the rms amplitude can be calculated as

$$\text{rms} = \sqrt{F \int_0^{\infty} P(\nu) d\nu} = \sqrt{A_0 F \left(\pi/2 - \tan^{-1} \frac{-\nu_0}{w/2} \right)}, \quad (5)$$

where F is a conversion factor between the Leahy normalization and the desired rms units.

3.3. Inference

Our periodogram modeling consists of two main steps: a maximum-likelihood estimation (MLE) of the best-fitting model, following D. Barret & S. Vaughan (2012), and a parametric bootstrap technique to evaluate uncertainties (B. Efron 1979); see also B. Efron & R. J. Tibshirani (1994). The MLE is done with the `Stingray.modeling` package (D. Huppenkothen et al. 2019). The procedure consists of minimizing the negative log-likelihood in Equation (3), using a model power spectrum. We employ an optimization algorithm that supports bounds, like the limited-memory Broyden–Fletcher–Goldfarb–Shanno scheme (LR. H. Byrd et al. 1995) to maximize the likelihood function, thereby obtaining parameter estimates. We fit the periodograms with one or two Lorentzians. When in doubt about the most appropriate number of components, we use the Akaike information criterion (AIC; H. Akaike 1974) to determine whether a model is best described by one or two Lorentzians (with threshold $\Delta\text{AIC} = 2$). For the Lorentzian components, we use an `Astropy` model (`Lorentz1D`) with the same parameters as Equation (4), and we fit an additional constant with a starting value of two, the expected white noise level in the Leahy normalization. For the initial values of the Lorentzians, we use an interactive interface to build a reasonable starting model, but then we do not set boundaries to the parameters other than being positive definite, and the FWHM being more than 0.001 Hz.

Once we obtain a best-fit model, we use parametric bootstrapping to evaluate the uncertainties, which involves the following steps.

1. *Generate random powers.* Randomly simulate powers from the best-fit model, ensuring that they follow a distribution of $\chi^2_{2M}/2M$, where M is the number of averaged powers in each bin of the original periodogram. These simulated powers should scatter around the best-fit model.

2. *Model fitting.* Fit a model to the random powers, starting from random parameters distributed within 10% of the real parameters.
3. *Bootstrap procedure.* Repeat steps 1–2 1000 times, recording the fit parameters in every iteration (a bootstrap distribution is obtained for each parameter).
4. *Parameter estimates.* Use the appropriate percentiles from the bootstrap results of the parameters to compute relevant statistics such as means, two-sided standard errors, and confidence intervals around the estimates. These provide information on the uncertainty associated with the parameter estimates based on the simulated powers. To take into account cases where the fit swaps the two Lorentzian components, we always order them by central frequency ν_0 and, if they both go to zero, by characteristic frequency ν_{\max} . We consider an outlier any point at more than five median absolute deviations from the median.

This simulation-based strategy is statistically more reliable than a profile-likelihood approach, which presumes a locally quadratic and unimodal likelihood surface. Those assumptions often break down for PDS fits, where the underlying noise is χ^2 distributed and multiple local maxima are common (D. Barret & S. Vaughan 2012). By resampling directly from the model, our bootstrapping captures any bias or skewness and yields realistic—often asymmetric—confidence intervals, even for complex or multimodal surfaces. For well-behaved cases, we find that the 16th–84th percentile bootstrap intervals are numerically close to those obtained from the classic profile-likelihood rule $\Delta \ln L = -0.5$ (approximately 1σ ; see Figure 6 of D. Huppenkothen et al. 2019), but the bootstrapping remains valid when that rule does not. Because the *Stingray* library (D. Huppenkothen et al. 2019) provides native routines for generating and refitting synthetic periodograms, the method is straightforward to apply. We therefore adopt the bootstrap ensemble as the sole source of both point estimates and uncertainties reported throughout this paper.

Table 1 presents the complete set of best-fit parameters derived from our modeling of the PDS described above. Figure 3 shows the variation of the total absolute rms of the features with frequency. The points with upper limits refer to models where the AIC suggested an advantage in adding a model component, but the bootstrap procedure returned a 3σ confidence interval including zero. We use the total rms and not the more customary fractional rms because the total X-ray flux of M82 is the combination of many X-ray sources, and in particular of both M82 X-1 and M82 X-2. Therefore, the fractional rms of the features will change randomly based on which source is more luminous during each observation, while the absolute rms retains a physical meaning in terms of total variable luminosity, in units of counts s^{-1} .

From this visualization, two distinct features are clearly identified. One is at frequencies below ~ 0.02 Hz whose rms does not depend on frequency, and one above, with a clear correlation between rms and frequency, which can confidently be identified with the QPO from M82 X-1 studied by T. E. Strohmayer & R. F. Mushotzky (2003), G. C. Dewangan et al. (2006), and P. Mucciarelli et al. (2006; see below for an additional test of the association). The increase of absolute rms with frequency follows an approximate $\nu^{1/2}$ law, which might imply that the frequency is increasing with increasing flux of M82 X-1 (i.e., implying an approximately constant fractional

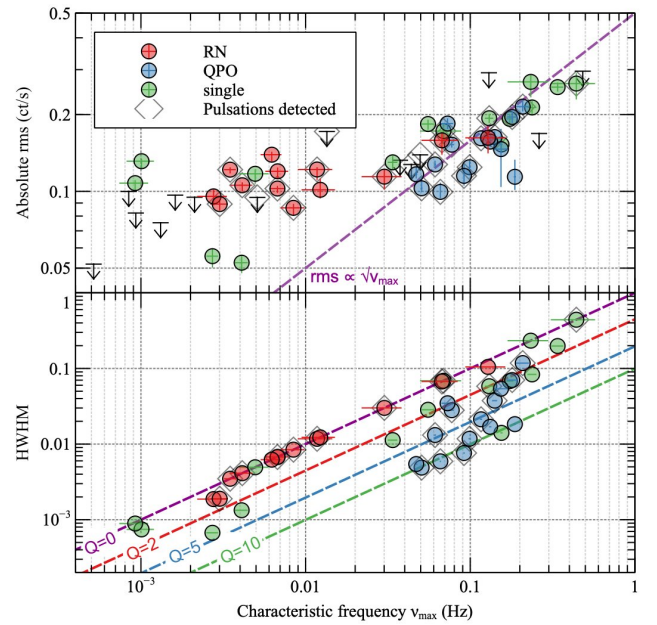


Figure 3. Top: rms vs. frequency for the power-spectral features fit in Section 1. Bottom: half-width at half-maximum vs. frequency for the same features. There does not seem to be a systematic pattern in the appearance of these timing features and pulsations. We calculated 1σ uncertainties and 3σ upper limits through the bootstrap procedure in Section 3.3. The identification of the two features is often difficult if only one of them is present in the data, but a simple criterion to distinguish them seems to emerge from this visualization, with the QPO, regardless of its Q factor in a given observation, following an $\text{rms} \propto \nu^{1/2}$ law (dashed line) and generally having a frequency above 0.02 Hz.

amplitude). There seems to be no evident correlation of any of the features with the appearance of pulsations from M82 X-2,¹⁵ which is probably due to the fact that the red noise is influenced by both sources.

3.4. Energy Dependence

It is interesting to investigate how the features evolve with energy. We divide the 3–80 keV energy band into eight intervals whose width follows approximately a geometric sequence but with larger intervals at higher energies to account for the very small number of high-energy counts. We refit the model in each energy band; however, in most observations, the fit is not robust enough to leave all parameters free. Since there is no evidence of changes in the shape of the Lorentzian components with energy, only their normalization, and the Poisson noise level does not depart significantly from two, we fix all parameters but the amplitudes of the one/two Lorentzians to their best-fit values from the total-flux analysis. Through repeating the bootstrap procedure, we get sensible values for the amplitudes and their uncertainties, and we can calculate the rms at different energies. Since we are proceeding on an observation-by-observation basis, we calculate the fractional rms, as we are only interested in the relative variation of rms with energy. The results are shown in Figure 4, and clearly show that the rms increases with energy for all components.

¹⁵ The pulsation information is mostly extracted from M. Bachetti et al. (2022), plus an additional detection that will be published by M. Bachetti et al. (2025, in preparation).




Figure 4. Fractional rms of the red noise and the QPO in different ObsIDs. Horizontal bars with caps represent exact energy ranges, and vertical bars are 1σ uncertainties. The rms increases steadily toward higher energies, a trend also seen for the ~ 0.5 mHz QPO in M51 ULX-7 (M. Imbrogno et al. 2024; Q.-C. Shui et al. 2025).

3.5. Search for Other Quasiperiodic Oscillations and Harmonics

D. R. Pasham et al. (2014) reported the detection of 3 and 5 Hz QPOs from M82, most probably from M82 X-1, using RXTE data. We now know that there are ~ 15 ULXs between the galaxies M82, M81, and the satellite Holmberg IX (D. J. Walton et al. 2022), all within the 1° FOV of the RXTE/Proportional Counter Array (PCA) instrument used to detect the QPOs. Despite M82 X-1 being the brightest at its maximum, all bright ULXs in the field, including M82 X-1, M82 X-2, Holmberg IX X-1, are known to be transient. This raises the possibility that M82 X-1 is not the source of the twin QPOs, and more generally, implies that the 3%–5% fractional-rms upper limit derived by D. R. Pasham et al. (2014) from the blended light curve almost certainly underestimates the intrinsic rms amplitude that M82 X-1 itself would need to produce the twin QPOs.

Despite the very long total exposure by NuSTAR, reproducing the result with NuSTAR data might not be straightforward, for various reasons. Despite the focusing capabilities that reduce considerably the contamination compared to RXTE, the NuSTAR PSF still includes M82 X-1, M82 X-2, and at least one other ULX, M82 X-3. Also, NuSTAR has a much lower effective area than RXTE. With these caveats in mind, we set out to look for the twin QPOs by averaging all M82 observations analyzed above. This time, we were less worried about low-frequency leakage given the relatively clean part of the periodogram above 1 Hz, so we used all data, including those from intervals with poor star-tracker coverage (see Section 2). Given the circumpolar position of M82, this allowed us to almost double the exposure, improving the detection sensitivity by a factor $\sim \sqrt{2}$. As can be seen in Figure 5, there is no significant feature at frequencies above the pulsation frequency of M82 X-2 with an rms reaching as low as 2%.

Additionally, we looked for harmonics of the 50 mHz QPO, in order to compare it to other known classes of QPOs such as Type C from BH LMXBs. We used the shift-and-add technique, originally introduced by M. Méndez et al. (1998) and later applied in similar contexts by D. Barret et al. (2005),

on the longest observing span available, the series of ObsIDs 80002092002–80002092011 from 2014. We split the observation in 128 s intervals, calculated the periodogram in each, then averaged 25 such intervals to gain in signal-to-noise ratio, creating a series of periodograms, each from 3.2 ks of noncontiguous data. These were sufficient to track the evolution of the QPO frequency ν_{QPO} during the observations. We calculated the average QPO frequency $\bar{\nu}_{\text{QPO}}$ and then shifted each periodogram by an amount corresponding to $\nu_{\text{QPO}} - \bar{\nu}_{\text{QPO}}$, obtaining the average shape of the QPO. This technique was used successfully to characterize kilohertz QPOs in NS LMXBs, and in some cases even to discover the upper kilohertz QPOs (D. Barret et al. 2005). We also used a variation of this technique; since we were not looking for a frequency following a parallel track, but for a harmonic, we shifted the periodogram around $2\bar{\nu}_{\text{QPO}}$ by a factor of 2. We did not find evidence for a harmonic of the QPO.

4. Discussion

Our analysis is based on a multi-Lorentzian fit of the periodogram, as is often done in similar studies of accreting sources (T. Belloni & G. Hasinger 1990). One difference is that we cannot rely on fractional variability when we study the evolution of the QPO, as we cannot get a clean view of M82 X-1 using NuSTAR due to the presence of M82 X-2, and the source flux is unknown. Hereafter, we will use the fractional rms only to compare the strength of the QPO at different energy bands in a given observation, while we will use the absolute rms (in counts per second) when discussing the evolution over time.

This source confusion, on a related note, also hinders us from doing detailed modeling of the spectral break and its relation with the QPO frequency (à la K. Atapin et al. 2019), because the red-noise component is produced by both sources.

The identification of the QPO with M82 X-1 is tricky by itself. Most detections come from missions that do not resolve the two ULXs (T. E. Strohmayer & R. F. Mushotzky 2003), such as XMM-Newton and RXTE, although with a good degree of confidence given the higher flux of M82 X-1 and the change of the QPO strength when carefully selecting data closer to M82 X-1 in XMM-Newton (H. Feng & P. Kaaret 2007). We looked for observations having quasi-simultaneous Chandra observations showing a low state from M82 X-2 and the QPO in the data, and searched the Chandra data themselves for QPO detections. Most on-axis Chandra observations of M82 are often plagued by pileup and the sensitivity to any variability was low. We found that the QPO was present in NuSTAR ObsID 90202038004, with Chandra ObsID 18072 showing a low state of M82 X-2, and, for the first time, a tentative detection in off-axis Chandra ObsIDs 17578 and 18064, this time using events firmly associated with M82 X-1, at a frequency compatible with the detection from the simultaneous NuSTAR ObsIDs (Figures 6 and 7).

The evolution of the 20–300 mHz QPO from M82 X-1 observed in our 10 yr NuSTAR campaign shows behavior compatible with what was previously observed for this source (P. Mucciarelli et al. 2006; H. Feng & P. Kaaret 2007; K. Atapin et al. 2019), extending the range of observed frequencies. The low background and hard response of NuSTAR, together with our power spectrum cleaning procedure, allows for a better modeling of the red-noise

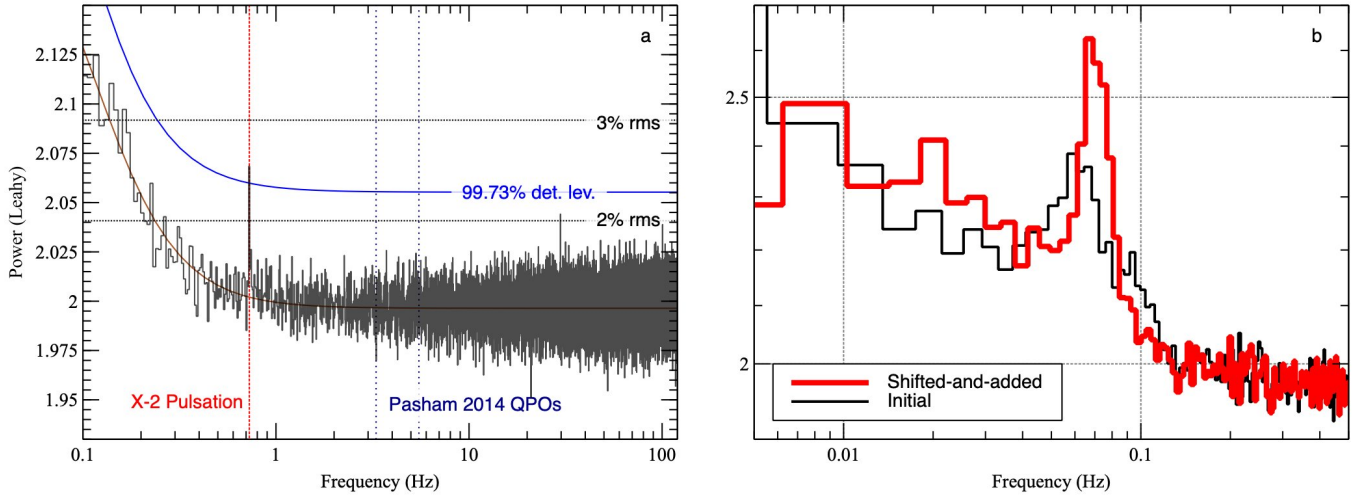


Figure 5. (a) Periodogram of the full ~ 5 Ms of NuSTAR observations of M82, including data intervals with poor star-tracker coverage. The peak at ~ 0.7 Hz is the pulsation of M82 X-2, while the peak at ~ 30 Hz is not significant and sets the upper limit to any quasi-coherent variability to $\sim 2\%$. (b) Result of the shift-and-add technique applied to the 50 mHz QPO. The resulting Q factor is ~ 8 , and there is no significant harmonic at twice the frequency.

component of the power spectrum and a systematic analysis of the evolution of the QPO using almost 3 Ms of exposure.

LFQPOs are observed in all classes of accreting sources (R. Wijnands & M. van der Klis 1999). In particular, NSBs and BHs have a number of low-frequency features spanning the frequency range ~ 0.001 –50 Hz. Some evolution of the QPO frequency over time is a hallmark of almost all classes of QPOs (M. van der Klis 2006; S. E. Motta 2016), and in particular, all classes of LFQPOs. Typically, the absolute rms is linearly correlated with flux, with the fractional rms being stable or slightly decreasing. Also the QPO frequency is often seen correlating with the flux on short timescales, while the correlation is broken on long timescales (the so-called “parallel tracks”). At higher energies, the characteristic frequency does not change significantly, while their fractional rms generally increases (M. van der Klis 2006).

One class of QPOs that can naturally be compared with ours is Type C QPOs from BH LMXBs, the most common oscillatory pattern in this class of sources. They start to appear in the low-hard state and their frequency generally increases as the luminosity (and probably the mass accretion rate) increases, going through the intermediate states that lead to the high-soft state and sometimes the so-called ultraluminous state. These intermediate states are also associated with the presence of transient jet ejections. Type C QPOs are generally also accompanied by broad red-noise (or flat-top) components, usually also modeled with Lorentzians whose characteristic frequencies (Equation (4)) evolve in parallel with that of the QPO, and one or more harmonics. Models for these QPOs often involve Lense–Thirring precession around a rotating BH, and put their frequency in relation with other oscillatory components such as broadband noise or HFQPOs. These models use different approaches from single-particle motion (the original relativistic precession model; L. Stella & M. Vietri 1999), to precessing rings (e.g., D. Psaltis & C. Norman 2000), to entire regions of the disk that precess like a solid body (P. C. Fragile et al. 2007; A. Ingram et al. 2011; see also A. R. Ingram & S. E. Motta 2019 for a review).

Interpreting M82 X-1’s QPO as a Type C QPO, one can also be tempted to go one step further: all timescales around a gravitating body scale with mass, and notably this includes

orbital frequencies. The fact that this QPO is about an order of magnitude slower than typical Type C QPOs leads to a mass estimate of an order of magnitude above stellar-mass BHs, in the regime of small IMBHs. This kind of scaling is often attempted, using different variability components, and in the case of M82 X-1 this has often led to claims of an IMBH origin for this source (e.g., T. E. Strohmayer & R. F. Mushotzky 2003; D. R. Pasham et al. 2014).

However, similar claims were made for other sources. One such source is M82 X-2, for which H. Feng et al. (2010) estimated a mass of 12,000–43,000 M_{\odot} by rescaling the Chandra-detected millihertz QPOs to Type C QPOs. However, this famously turned out to be a pulsar. Another pulsating ULX that recently showed a QPO in the millihertz range is M51 ULX-7 (M. Imbrogno et al. 2024), further stressing that care is needed when estimating the mass of the (ULX) accretor from the QPO frequency. It is generally difficult to compare QPO phenomena from different objects, and there is a wealth of QPO phenomena on many timescales in stellar-mass compact objects, including at even lower frequencies. One example is the millihertz QPOs observed in low-mass NS and BH X-ray binaries (e.g., M. Revnivtsev et al. 2001; H. Xiao et al. 2025).

Moreover, we note that the spectrum of M82 X-1 does not depart significantly from the bulk of ULXs (M. Brightman et al. 2020), which are increasingly identified as super-Eddington accreting stellar-mass objects. In some cases the detection of pulsations unequivocally identify them as NSs, but in general, their spectral shapes do not resemble the bulk of sub-Eddington accreting objects (J. C. Gladstone et al. 2009), and they frequently show signatures of the strong winds expected from super-Eddington accretion (M. J. Middleton et al. 2014, 2015b; C. Pinto et al. 2016, 2017; P. Kosec et al. 2021; C. Pinto & D. J. Walton 2023). In particular, NGC 5907 X-1 has a comparable flux to M82 X-1 despite being an NS (G. L. Israel et al. 2017; F. Fürst et al. 2023).

In addition, the decrease of the coherence of the M82 X-1 QPO with frequency observed here (Figure 3) is puzzling, and unlike what is observed in Type C QPOs where the coherence tends to generally increase with frequency, at least up to the onset of the soft state where these oscillations disappear.




Figure 6. Simultaneous detections of the QPO in NuSTAR and Chandra data, showing a clear association with M82 X-1.




Figure 7. Image of Chandra ObsID 18072, simultaneous to NuSTAR ObsID 90202038004. Blue circles indicate the seven ULXs in the catalog by J.-F. Liu & J. N. Bregman (2005). The QPO is significantly detected in NuSTAR data and not in Chandra data of any source, but the Chandra image shows that the emission is dominated by M82 X-1, while M82 X-2 is undetected.

Admittedly, this might be an observational bias: the integration time required to detect the QPO might just be longer than the variation time for the frequency. We need thousands of oscillations of the QPO before being able to detect it, and if the frequency is changing rapidly this would artificially increase the measured width, and so its quality factor. A clear example of this are ObsIDs 80002092006 and 80002092007, where using the shift-and-add technique clearly improves the quality factor (Section 3.5). In an effort to determine if this played a role, we used the shift-and-add technique in all observations with a strong QPO and at least ~ 90 ks exposure. The technique always improved the Q factor throughout the frequency range, but without significant advantages for the higher frequencies.

Interestingly, bright accreting pulsars also show LFQPOs, with a phenomenology and characteristic frequencies similar to what we observe here. For these sources, LFQPOs are often hypothesized to arise from the Keplerian frequency at some important radius, or from a beat between the Keplerian frequency at the truncation radius and the spin frequency (beat frequency model; M. A. Alpar & J. Shaham 1985). H. Manikantan et al. (2024) provide a table of the energy-dependent QPO parameters for a number of sources over multiple observations. V0332+53 is

the only one with sufficient observations to see the evolution of the quality factor over a wide range of QPO frequencies, and Q is interestingly seen to decrease between 10 and two with increasing frequency, as we observe for M82 X-1. As a general rule, we argue against using QPOs alone to infer the mass of accreting objects (see M. J. Middleton et al. 2011). Alternative explanations to geometric or relativistic origins for LFQPOs have also been proposed. In particular, the ν_{kompth} model (K. Karpouzas et al. 2020; L. Bellavita et al. 2022) interprets LFQPOs as a result of oscillatory Comptonization in a compact corona, driven by feedback between heating and soft-photon cooling. This model does not require disk precession or general relativistic effects and instead focuses on radiative instabilities within the Comptonizing region. Although our analysis does not reveal clear energy-dependent phase lags or spectral pivoting that would be diagnostic of such a mechanism, the observed trends in coherence and the lack of harmonics make it a viable alternative that warrants consideration in the broader context of ULX QPO phenomenology. The QPOs studied in this work are compatible with phenomenology observed in accreting sources of different kinds, including NSs.

Our timing analysis yields a fractional-rms amplitude of $22\% \pm 2\%$ for the 50–300 mHz QPO in the 10–20 keV band (Figure 4). This already exceeds the $\lesssim 20\%$ typically reported for Type A/B/C LFQPOs in Galactic BHs below 10 keV (e.g., P. Casella et al. 2005; T. M. Belloni 2009).¹⁶ Given that M82 X-2 contributes $\simeq 30\%$ – 40% of the 3–20 keV continuum flux (Table 1) and shows no significant power at the QPO frequency, the dilution-corrected (intrinsic) fractional rms of the M82 X-1 QPO must be $\gtrsim 30\%$.

A similar behavior is also observed in M51 ULX-7, a pulsating ultraluminous X-ray source (M. Imbrogno et al. 2024), with similarly high values of rms. In both sources the QPO fractional rms steadily increases with photon energy, reaching $\sim 50\%$ at about 8 keV (Q.-C. Shui et al. 2025, Figure 3).

The close match might imply that geometry, rather than the nature of the accretor, dominates the observed variability. It should be noted, however, that a comparison with M51 ULX-7 is only possible in the 0.3–10 keV band, where the millihertz QPO has been detected (M. Imbrogno et al. 2024; Q.-C. Shui et al. 2025). Future observations of M51 ULX-7 at energies > 10 keV

¹⁶ In Galactic systems, fractional-rms amplitudes below 10 keV usually fall in the 5%–15% range and only rarely reach 20%.

with facilities such as NuSTAR could help us understand whether the energy dependence of the QPO is really similar to M82 X-1 or differences start to arise at higher energies due to the distinct nature of the accretor. Finally, it is worth noting that the QPO frequency in M82 X-1 varies widely in the 20–300 mHz range. This is not the case for M51 ULX-7, in which the frequency of the QPO is always found to be ≈ 0.5 mHz. The lack of variability in the latter could be a sign of the different nature of the accretor, but it could also be the consequence of an observational bias, since there are fewer detections of the QPO from M51 ULX-7 (five Chandra observations and three XMM-Newton observations; M. Imbrogno et al. 2024) than from M82 X-1 (26 NuSTAR observations, see Table 1).

The nondetection of the twin 3–5 Hz QPOs reported by D. R. Pasham et al. (2014) can be explained in various ways. First of all, the oscillations might be transient, and have disappeared over time: the result by D. R. Pasham et al. (2014) is based on the integration of many years of data. It is also possible that the filtering of “flaring” observations by D. R. Pasham et al. (2014), which reduced considerably the time intervals analyzed to obtain their average periodogram, was more aggressive than ours, and that we need more observations without flaring or high variability to detect those QPOs with NuSTAR. Another possibility might even be that the source of QPOs is *not* in M82 after all: the FOV of RXTE is 1° , which includes three ULX host galaxies (M82, M81, and Holmberg IX) with at least nine ULXs (J.-F. Liu & J. N. Bregman 2005).

5. Conclusions

We made an extensive, 10 yr study of the QPOs from M82 X-1 using NuSTAR. Using archival Chandra data, we made a robust identification of the 20–300 mHz QPO from M82 X-1, which confirms previous evidence. Thanks to the sensitive response at and above 10 keV, we were able to detect the QPO over $\sim 70\%$ of the 3 Ms of existing M82 observations, characterizing its behavior. We used a novel approach to cleaning the periodogram in order to better fit the low-frequency component and get reliable fit parameters for the QPO as well. The QPO tends to decrease its coherence as its frequency increases, but it is not clear whether this is due to fast variations of the frequency that we are not able to follow due to the long exposure required by the detection at these count rates. We note that LFQPOs in this range of frequencies and with similar behavior are observed in many accreting systems, including NSs, so that any inference on the mass of the compact object based on the frequency of the QPOs should be taken with a grain of salt.

Acknowledgments

The authors wish to thank Phil Uttley for discussions on the spectral cleaning procedure. M.B. was supported in part by the Italian Research Center on High Performance Computing Big Data and Quantum Computing (ICSC), project funded by European Union—NextGenerationEU—and National Recovery and Resilience Plan (NRRP)—Mission 4 Component 2 within the activities of Spoke 3 (Astrophysics and Cosmos Observations). D.J.W. acknowledges support from the Science and Technology Facilities Council (STFC; grant code ST/Y001060/1).

Software: astropy (Astropy Collaboration et al. 2013, 2018, 2022), Stingray (D. Huppenkothen et al. 2019), CIAO (A. Fruscione et al. 2006), Astroquery (A. Ginsburg et al. 2019).

Appendix A Periodogram Production and Filtering

Our source light curve is produced by sampling photons from the source region falling into equispaced time intervals with resolution t_{samp} . Alongside the source light curve, we create synthetic visibility light curves with the same sampling interval that are equal to the mean counts per bin of the source light curve during GTIs and zero outside. We take a periodogram of this synthetic light curve, which shows strong features corresponding to the missing data while going rapidly to zero at high frequencies, because it does not contain noise. In this periodogram, we can set a threshold, and single out the strongest peaks to be used for notch filtering. The idea is that we will eliminate these frequencies from the final periodogram, and then make a geometrical rebinning (as it is commonly done) that will average the remaining nearby bins, maintaining the statistical properties of the periodogram.

However, if we take the source light curve at face value, using zero outside GTIs, we can see that if one applies the notch filter above straight away, a large number of powers do not follow the expected distribution (Figure 8, middle panels). This is because some noise is still leaking at frequencies near the blacklisted ones.

But we can also fill the BTIs of the light curve not with zeros, but with the mean of the data. This produces a much smaller effect of the visibility windows on the final periodogram, creating an almost flat distribution of powers, and the notch filter will just be an additional cautionary measure (Figure 8, bottom panels, black data). In a real-life situation, when there is some long-term source variability, filling bad intervals with the mean will not be as clean as this example, so notch filtering will be useful. There is an additional measure to take, however. The periodogram is a measure of the variance of the data, and it is calculated from an FFT which contains a division by total number of data points. However, the filled data points do not contribute to the variance, and this means that the calculated power will be lower than the expected value. To reinstate the correct normalization of the white noise, we need to multiply this periodogram by $n_{\text{tot}}/n_{\text{gti}}$, where n_{tot} is the total number of bins and n_{gti} is the number of bins in GTIs (Figure 8, bottom panels, magenta data). Additionally, the measured rms of signal will still be underestimated by another factor $n_{\text{tot}}/n_{\text{gti}}$.

To verify that the method does not alter the response of the periodogram in a frequency-dependent way, we performed the following test: We generated synthetic light curves with a sinusoidal modulation at a 10% fractional amplitude, using the same GTIs and mean flux as ObsID 80002092006. The light curve had no Poisson noise, only the smooth sinusoidal modulation. Outside GTIs, we filled the time intervals with the mean value of the flux as we did for the data. We repeated the experiment for 1000 frequencies distributed log uniformly over the range 10^{-4} –1 Hz. We always used frequencies extracted from the grid of the periodogram, in order to avoid the expected sinc² response degradation when moving away from the center of the frequency bin (M. van der Klis et al. 1989). We applied the same notch filters as the real data, and we measured the decrease of rms (hereafter, the damping factor) in the remaining frequencies compared to an uninterrupted pulsation over the whole observation. The decrease of rms was the same at all frequencies and exactly what expected from the fact of having missing data (i.e., outside GTIs there is no variability and the total variance decreases).

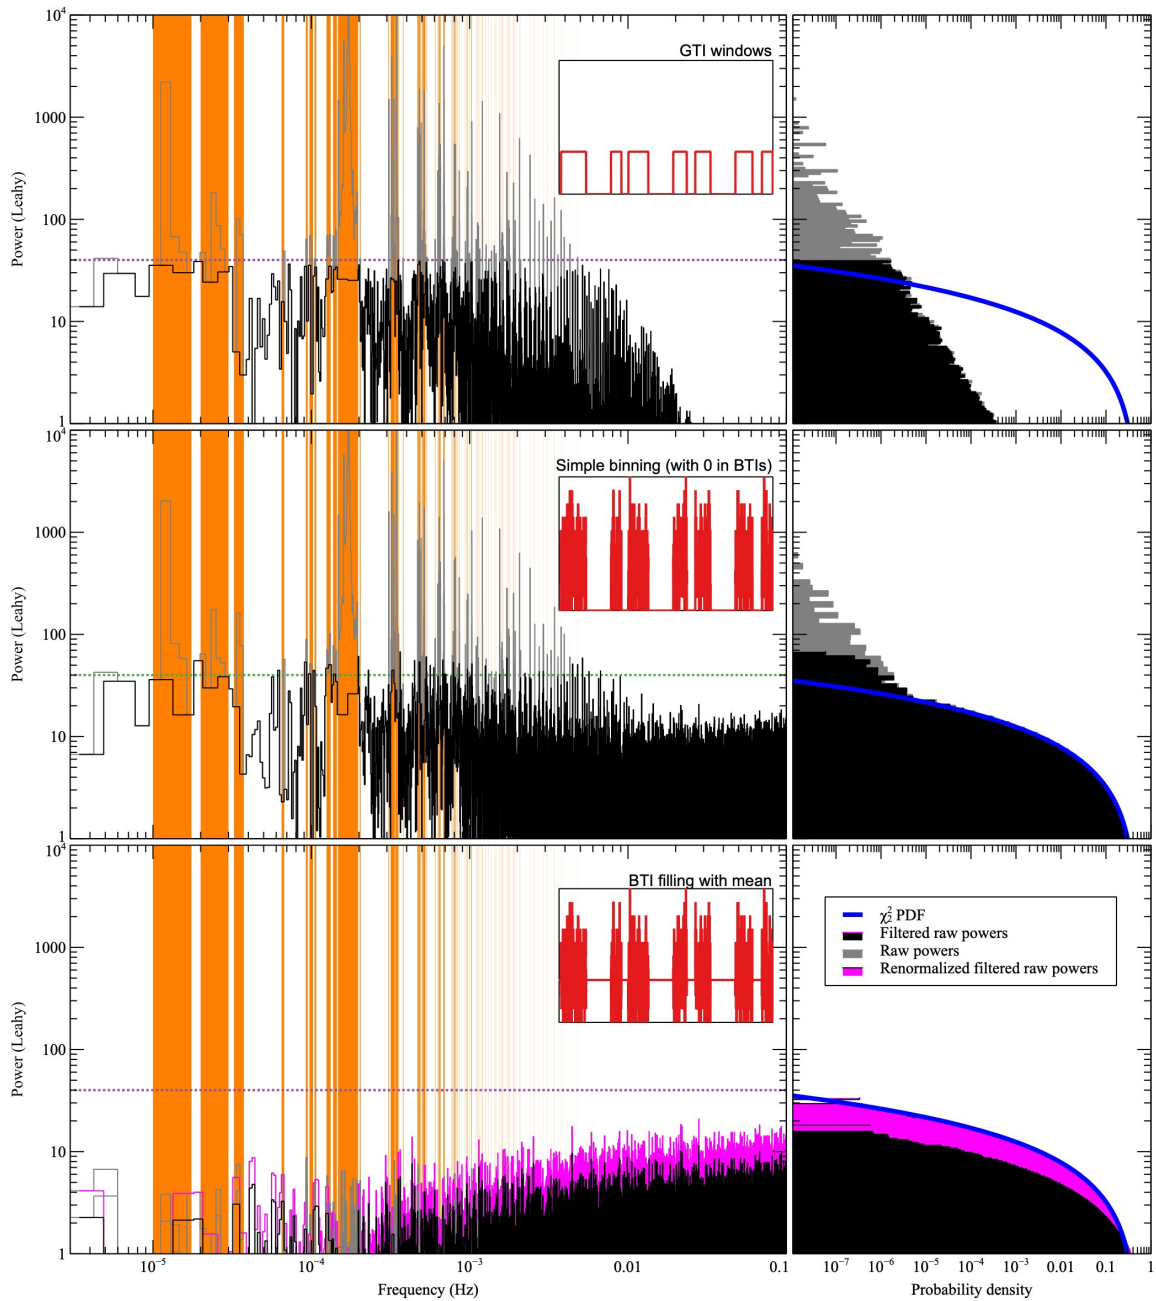


Figure 8. Procedure to filter periodograms from the effect of visibility windows described in Appendix A, using simulated Poissonian data (so, no source variability) with the same mean count rate and GTIs of ObsID 80002092006. Top: periodogram of the visibility light curve, showing the features corresponding to the missing data. We set a threshold and select a number of bad frequency intervals (orange) to be blacklisted. Middle: periodogram of the binned light curve with zero outside GTIs; it contains many of the the same features, with similar powers, plus the expected white noise from the data. After notch filtering, some powers clearly exceed the expected χ^2_2 distribution. Bottom: periodogram of the binned light curve with the mean counts per bin used as a filler outside GTIs instead of zero. Most of the features disappear from the periodogram even before notch filtering, and the powers follow the correct distribution, but with the wrong normalization, which is corrected as described in the text

Only very close to the notch-filtered frequencies do we noticed some “wiggles” in the damping factor, by 10% at most.

Appendix B

Alternative Periodogram Analysis: Bartlett + Lomb–Scargle

Another widely used tool for detecting and characterizing periodic signals is the Lomb–Scargle periodogram (N. R. Lomb 1976; J. D. Scargle 1982), which has proven to be more effective in detecting periodic patterns even when observations are

unevenly spaced, providing a reliable solution for analyzing time series data in astronomy and diverse scientific fields as well. The Lomb–Scargle periodogram allows for covering a wide range of low frequencies in comparison to the Bartlett periodogram, allowing for a more comprehensive analysis of low-frequency signals.
















The only downside of this periodogram is that its powers are not guaranteed to be uncorrelated, and the assumption of χ^2_2 -distributed powers is not as solid as for the periodogram. For a comprehensive view on modern algorithms to compute

the Lomb–Scargle periodogram and their limitations, see J. T. VanderPlas (2018).

Nonetheless, the Lomb–Scargle periodogram is very useful for the analysis of unevenly sampled data, including light curves with missing data. Our original approach to the analysis of the data sets in this paper was a hybrid approach, using the Bartlett periodogram and the Lomb–Scargle periodogram at the same time. We write it here because it could be an interesting inspiration for future works and, in any case, it represents an alternative (even if not completely independent) approach to the analysis presented in this paper. We calculated the Lomb–Scargle periodogram simply discarding light-curve bins outside GTIs, and avoiding oversampling (i.e., using the same spectral resolution of the FFT), which limited the correlations between powers. Our sample time was 0.1 s, giving 5 Hz as the Nyquist frequency, but we found that the periodogram departed significantly from the expected noise level of two when reaching about half the Nyquist frequency. Otherwise, the Lomb–Scargle and the Bartlett periodograms have very good overlap in the common frequency ranges, and we decided to eliminate the frequencies above 1 Hz from the Lomb–Scargle periodogram for extra caution. We checked that the assumption of χ^2_2 -distributed powers is justified for our case by plotting the distribution of the Lomb–Scargle powers of the simulated and real data similarly to Figure 8. Moreover, even though in theory it should not have been needed, we blacklisted the frequencies corresponding to the orbital occultations similarly to the description in Appendix A. The raw periodograms have very different frequency resolutions, but we rebinned both periodograms with geometrically increasing frequency bin sizes, and defined a threshold frequency ν_{thr} where the frequency resolution of the rebinned Lomb–Scargle periodogram reached the frequency resolution of the Bartlett periodogram. From that point on, we used a hybrid periodogram containing the Lomb–Scargle powers below ν_{thr} and the Bartlett ones above. The periodogram was characterized by the power values and the number of averaged powers, either from rebinning or—in the case of the Bartlett one—averaging of multiple periodograms.

We proceeded to fit the multi-Lorentzian model of Section 3.3, with the same methods, to this hybrid periodogram. Indeed, the results using this alternative method are compatible with those obtained for the single filtered periodogram in Section 3.3.

ORCID iDs

Hamza El Byad  <https://orcid.org/0009-0001-5473-4953>
 Matteo Bachetti  <https://orcid.org/0000-0002-4576-9337>
 Silvia Columbu  <https://orcid.org/0000-0001-5420-0719>
 Giuseppe Rodriguez  <https://orcid.org/0000-0001-9054-8712>
 Maura Pilia  <https://orcid.org/0000-0001-7397-8091>
 Matthew J. Middleton  <https://orcid.org/0000-0002-8183-2970>
 Dominic J. Walton  <https://orcid.org/0000-0001-5819-3552>
 Murray Brightman  <https://orcid.org/0000-0002-8147-2602>
 Hannah Earnshaw  <https://orcid.org/0000-0001-5857-5622>
 Karl Forster  <https://orcid.org/0000-0001-5800-5531>
 Brian Grefenstette  <https://orcid.org/0000-0002-1984-2932>
 Marianne Heida  <https://orcid.org/0000-0002-1082-7496>
 Matteo Imbrogno  <https://orcid.org/0000-0001-8688-9784>
 Eleonora Veronica Lai  <https://orcid.org/0000-0002-6421-2198>
 Thomas Maccarone  <https://orcid.org/0000-0003-0976-4755>

References

- Abramowicz, M. A., & Kluźniak, W. 2001, *A&A*, 374, L19
 Akaike, H. 1974, *ITAC*, 19, 716
 Alpar, M. A., & Shaham, J. 1985, *Natur*, 316, 239
 Altamirano, D., & Belloni, T. 2012, *ApJL*, 747, L4
 Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *ApJ*, 935, 167
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, 156, 123
 Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33
 Atapin, K., Fabrika, S., & Caballero-García, M. D. 2019, *MNRAS*, 486, 2766
 Bachetti, M., Harrison, F., Walton, D. J., et al. 2014, *Natur*, 514, 202
 Bachetti, M., Heida, M., Maccarone, T., et al. 2022, *ApJ*, 937, 1
 Barret, D., Kluźniak, W., Olive, J. F., Paltani, S., & Skinner, G. K. 2005, *MNRAS*, 357, 1288
 Barret, D., & Vaughan, S. 2012, *ApJ*, 746, 131
 Bartlett, M. S. 1950, *Biometrika*, 37, 1
 Bellavita, L., Méndez, M., Karpouzas, K., et al. 2022, *MNRAS*, 515, 2099
 Belloni, T., & Hasinger, G. 1990, *A&A*, 230, A103
 Belloni, T., Homan, J., Casella, P., et al. 2005, *A&A*, 440, 207
 Belloni, T., Psaltis, D., & van der Klis, M. 2002, *ApJ*, 572, 392
 Belloni, T. M. 2009, *Lecture Notes in Physics*, Vol. 794. *The Jet Paradigm: From Microquasars to Quasars* (Berlin: Springer), 53
 Belloni, T. M., & Altamirano, D. 2013, *MNRAS*, 432, 10
 Belloni, T. M., & Stella, L. 2012, *SSRv*, 183, 43
 Bradt, H. V., Rothschild, R. E., & Swank, J. H. 1993, *A&AS*, 97, 355
 Brightman, M., Harrison, F. A., Bachetti, M., et al. 2019, *ApJ*, 873, 115
 Brightman, M., Walton, D. J., Xu, Y., et al. 2020, *ApJ*, 889, 71
 Byrd, R. H., Lu, P., Nocedal, J., & Zhu, C. 1995, *SJSC*, 16, 1190
 Casella, P., Belloni, T., & Stella, L. 2005, *ApJ*, 629, 403
 Cooley, J. W., & Tukey, J. W. 1965, *Math. Comp.*, 19, 297
 Dewangan, G. C., Titarchuk, L., & Griffiths, R. E. 2006, *ApJL*, 637, L21
 Efron, B. 1979, *AnSta*, 7, 1
 Efron, B., & Tibshirani, R. J. 1994, *An Introduction to the Bootstrap* (London: Chapman and Hall/CRC)
 Feng, H., & Kaaret, P. 2007, *ApJ*, 668, 941
 Feng, H., Rao, F., & Kaaret, P. 2010, *ApJL*, 710, L137
 Fragile, P. C., Blaes, O. M., Anninos, P., & Salmonson, J. D. 2007, *ApJ*, 668, 417
 Freedman, W. L., Hughes, S. M., Madore, B. F., et al. 1994, *ApJ*, 427, 628
 Fruscione, A., McDowell, J. C., Allen, G. E., et al. 2006, *Proc. SPIE*, 6270, 62701V
 Fürst, F., Walton, D. J., Israel, G. L., et al. 2023, *A&A*, 672, A140
 Ginsburg, A., Sipőcz, B. M., Brasseur, C. E., et al. 2019, *AJ*, 157, 98
 Gladstone, J. C., Roberts, T. P., & Done, C. 2009, *MNRAS*, 397, 1836
 Harrison, F. A., Craig, W. W., Christensen, F. E., et al. 2013, *ApJ*, 770, 103
 Huppenkothen, D., Bachetti, M., Stevens, A. L., et al. 2019, *ApJ*, 881, 39
 Imbrogno, M., Motta, S. E., Amato, R., et al. 2024, *A&A*, 689, A284
 Ingram, A., Done, C., & Fragile, P. C. 2011, *MNRAS*, 419, 2369
 Ingram, A. R., & Motta, S. E. 2019, *NewAR*, 85, 101524
 Israel, G. L., Belfiore, A., Stella, L., et al. 2017, *Sci*, 355, 817
 Iwasawa, K. 2021, *A&A*, 652, A18
 Kaaret, P., Feng, H., & Roberts, T. P. 2017, *ARA&A*, 55, 303
 Kaaret, P., Simet, M. G., & Lang, C. C. 2006, *ApJ*, 646, 174
 Karpouzas, K., Méndez, M., Altamirano, D., et al. 2020, *MNRAS*, 492, 1399
 King, A., Lasota, J.-P., & Middleton, M. 2023, *NewAR*, 96, 101672
 Kluźniak, W. 2006, *AcPPB*, 37, 1361
 Kluźniak, W., & Abramowicz, M. A. 2001, *AcPPB*, 32, 3605
 Kosec, P., Pinto, C., Reynolds, C. S., et al. 2021, *MNRAS*, 508, 3569
 Leahy, D., Darbro, W., Elsner, R., et al. 1983, *ApJ*, 266, 160
 Liu, J.-F., & Bregman, J. N. 2005, *ApJS*, 157, 59
 Lomb, N. R. 1976, *Ap&SS*, 39, 447
 Manikantan, H., Paul, B., Sharma, R., Pradhan, P., & Rana, V. 2024, *MNRAS*, 531, 530
 Matsumoto, H., & Tsuru, T. G. 2001, *ApJL*, 549, L97
 Méndez, M., van der Klis, M., Van Paradijs, J., et al. 1998, *ApJL*, 494, L65
 Middleton, M. J., Heil, L., Pintore, F., Walton, D. J., & Roberts, T. P. 2015a, *MNRAS*, 447, 3243
 Middleton, M. J., Roberts, T. P., Done, C., & Jackson, F. E. 2011, *MNRAS*, 411, 644
 Middleton, M. J., Walton, D. J., Fabian, A., et al. 2015b, *MNRAS*, 454, 3134
 Middleton, M. J., Walton, D. J., Roberts, T. P., & Heil, L. 2014, *MNRAS*, 438, L51
 Motta, S. E. 2016, *AN*, 337, 398

- Mucciarelli, P., Casella, P., Belloni, T., Zampieri, L., & Ranalli, P. 2006, *MNRAS*, **365**, 1123
- Pasham, D. R., Strohmayer, T. E., & Mushotzky, R. F. 2014, *Natur*, **513**, 74
- Pinto, C., Alston, W., Soria, R., et al. 2017, *MNRAS*, **468**, 2865
- Pinto, C., Middleton, M. J., & Fabian, A. C. 2016, *Natur*, **533**, 64
- Pinto, C., & Walton, D. J. 2023, High-Resolution X-ray Spectroscopy: Instrumentation, Data Analysis, and Science (Singapore: Springer), 345
- Psaltis, D., & Norman, C. 2000, arXiv:[astro-ph/0001391](https://arxiv.org/abs/astro-ph/0001391)
- Ransom, S. M., Eikenberry, S. S., & Middleditch, J. 2002, *AJ*, **124**, 1788
- Remillard, R. A., & McClintock, J. E. 2006, *ARA&A*, **44**, 49
- Revnitsev, M., Churazov, E., Gilfanov, M., & Sunyaev, R. 2001, *A&A*, **372**, 138
- Scargle, J. D. 1982, *ApJ*, **263**, 835
- Shaposhnikov, N. 2012, *ApJL*, **752**, L25
- Shui, Q.-C., Zhang, S., Feng, H., et al. 2025, *ApJ*, **984**, 130
- Stella, L., & Vietri, M. 1998, *ApJL*, **492**, L59
- Stella, L., & Vietri, M. 1999, *PhRvL*, **82**, 17
- Strohmayer, T. E., & Mushotzky, R. F. 2003, *ApJL*, **586**, L61
- van der Klis, M., Ögelman, H., & Van Den Heuvel, E. P. J. 1989, in NATO ASI Ser. 262, Timing Neutron Stars (New York: Kluwer), 27
- van der Klis, M. 2006, in Compact Stellar X-ray Sources, ed. W. Lewin, Walter, & M. van der Klis (Cambridge: Cambridge Univ. Press), 39
- VanderPlas, J. T. 2018, *ApJS*, **236**, 16
- Walton, D. J., Mackenzie, A. D. A., Gully, H., et al. 2022, *MNRAS*, **509**, 1587
- Wijnands, R., & van der Klis, M. 1999, *ApJ*, **514**, 939
- Xiao, H., Ji, L., Tsygankov, S., et al. 2025, *ApJ*, **982**, 180
- Zhang, C. M., Yin, H. X., Zhao, Y. H., Zhang, F., & Song, L. M. 2006, *MNRAS*, **366**, 1373