

Long-term properties of accretion discs in X-ray binaries – III. A search for spin–superorbital correlation in SMC X-1

Kristen C. Dage^{1,2★}, William I. Clarkson¹, Philip A. Charles^{3,4†}, Silas G. T. Laycock⁵ and I-Chun Shih⁶

¹Department of Natural Sciences, University of Michigan-Dearborn, 4901 Evergreen Rd. Dearborn, MI 48128, USA

²Department of Physics and Astronomy, Michigan State University, East Lansing, MI 48824, USA

³Astrophysics, Department of Physics, Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, UK

⁴Physics & Astronomy, University of Southampton, SO17 1BJ, UK

⁵Department of Physics & Applied Physics, University of Massachusetts, Lowell, MA 01854, USA

⁶GEPI, Observatoire de Paris, Université PSL, CNRS, 5 Place Jules Janssen, 92190 Meudon, France

Accepted 2018 September 17. Received 2018 September 17; in original form 2018 June 27

ABSTRACT

Due to long-term X-ray monitoring, a number of interacting binaries are now known to show X-ray periodicities on time-scales of tens to hundreds of binary orbits. In some systems, precession of a warped accretion disc is the leading model to explain the superorbital modulation. The High-Mass X-ray Binary SMC X-1 showed two excursions in superorbital period (from ~ 60 d to ~ 45 d) during the 1996–2011 interval, suggesting that some characteristic of the accretion disc is varying on a time-scale of years. Because its behaviour as an X-ray pulsar has also been intensively monitored, SMC X-1 offers the rare chance to track changes in both the accretion disc and pulsar behaviours over the same interval. We have used archival X-ray observations of SMC X-1 to investigate whether the evolution of its superorbital variation and pulse period are correlated. We use the 16-year data set afforded by the *RXTE* All-Sky Monitor to trace the behaviour of the warped accretion disc in this system, and use published pulse-period histories to trace the behaviour of the pulsar. While we cannot claim a strong detection of correlation, the first superorbital period excursion near MJD 50 800 does coincide with structure in SMC X-1's pulse-period history. Our preferred interpretation is that the superorbital period excursion coincides with a change in the long-term spin-up rate of the SMC X-1 pulsar. In this scenario, the pulsar and the accretion disc are both responding to a change in the accretion flow, which the disc itself may regulate.

Key words: accretion, accretion discs – pulsars: individual: SMC X-1 – X-rays: binaries.

1 INTRODUCTION

With orbital periods typically a few days or shorter (e.g. Liu, van Paradijs & van den Heuvel 2006), X-ray binaries allow behaviours to be charted over hundreds of orbital time-scales through multi-year X-ray monitoring, offering a window into the accretion/outflow process which is important to astrophysical objects on all length-scales (e.g. Frank, King & Raine 2002; Charles & Coe 2006). As suggested by X-ray monitoring data sets on time-scales of years-decades, a number of X-ray binaries are now suspected to undergo accretion disc warping, in which the accretion disc warps out of the plane and may precess, giving rise to observed high/low cycles

in X-ray brightness on a time-scale of tens of days (e.g. Charles, Kotze & Rajoelimanana 2010).

The leading candidate mechanism to produce the warp is *radiation-driven warping* (Pringle 1996), in which the interception and re-emission by the accretion disc of radiation from the accretor leads to a perturbative force out of the plane. For the more luminous accretors at the centre of sufficiently large accretion discs, a warped configuration becomes stable (Wijers & Pringle 1999) and can precess. Under this model, the configuration of the warp, and thus the superorbital modulation, should depend on the binary separation (a proxy for the disc radius), with more widely separated binaries producing more complicated warp configurations, or possibly a competition of modes (Ogilvie & Dubus 2001, hereafter OD01). Comparison of superorbital X-ray modulations of a sample of systems chosen to trace the warp modes of OD01 suggested that the X-ray binaries as a class might indeed follow this scheme (Clarkson et al. 2003a, Clarkson et al. 2003b, hereafter Paper I & Paper II, respectively). Smoothed-particle hydrodynamics simula-

* E-mail: kcdage@msu.edu

† Leverhulme Emeritus Fellow, The Leverhulme Trust, London EC4A 3BG, UK

tions of accretion discs in a variety of realistic configurations added further support to the radiation-driven warping model (Foulkes, Haswell & Murray 2006, 2010); when the central radiation-source is switched on in the simulation, the accretion disc does indeed warp and precess.

Because it is the interception of X-ray luminosity that drives the warp in the first place, the configuration of the warp should be sensitive to the mass transfer rate onto the central accretor (Ogilvie & Dubus 2001; Foulkes et al. 2010). An X-ray binary in which X-ray superorbital modulation is observed to vary, *and* in which variations in \dot{M} can be observed, would therefore provide a direct test of the disc/accretor link for warped discs.

The well-studied X-ray binary SMC X-1 is the perfect system in which to probe this connection. A superorbital X-ray modulation at ~ 60 d has long been observed (Bonnet-Bidaud & van der Klis 1981 provides an early discussion of superorbital modulation of this source). X-ray spectroscopic evidence suggests this variation is likely due to repeated obscuration, just as would be expected from a precessing disc-warp (see e.g. Wojdowski et al. 1998; Vrtilek et al. 2005, as well as Blondin & Woo 1995). Optical monitoring supports this general model (Coe et al. 2013).

The superorbital modulation from SMC X-1 was suspected to vary at least as early as 1998 thanks to the first year of X-ray monitoring by *RXTE* (Wojdowski et al. 1998). Further monitoring has led to improved characterization of this variation (e.g. Ribó et al. 2001, Paper I; Trowbridge, Nowak & Wilms 2007; Kotze & Charles 2012; Hu et al. 2013). Whether the detected period of modulation varies smoothly or discretely, on at least two occasions the superorbital modulation has been recorded to be as short as 45 d. We identify, in particular, two superorbital period excursions: the first superorbital excursion takes place between approximately MJD 50 500 and MJD 51 500, with period minimum taking place near MJD 50 800. The second, slightly shorter superorbital excursion takes place between approximately MJD 53 500 and MJD 54 500, with period minimum taking place near MJD 54 000 (e.g. Trowbridge et al. 2007; Kotze & Charles 2012). SMC X-1 lies in the region of (binary radius, mass ratio) space in which a competition of modes might be expected (OD01). It is the best current candidate for a system in which an accretion disc warp shows strong variation in its configuration.

Some indication of a link between superorbital modulation and \dot{M} has been claimed from long-term X-ray monitoring alone. For example, SMC X-1 shows a weak correlation between the superorbital period and superorbital modulation amplitude (in the sense that longer modulation period would suggest greater modulation amplitude), although the link between rms amplitude and X-ray flux (tracing \dot{M}) is unclear (Hu et al. 2011). If this amplitude indeed relates to \dot{M} directly, then this claimed correlation would be a point of similarity with Her X-1, the prototypical X-ray binary exhibiting stable accretion disc warping (varying in the range 33–37 d, e.g. Jurua et al. 2011), for which a correlation has also been measured between its superorbital modulation period and the average X-ray flux over the superorbital cycle (Leahy & Iğna 2010).

SMC X-1¹ is also a very well-measured X-ray pulsar, with the evolution of the X-ray pulse profile offering further support for a warped accretion disc scenario. Its X-ray pulse profile is energy-dependent; the ‘hard’ component (at energies > 2.0 keV) is double peaked and does not show strong variation in profile shape with

superorbital phase (e.g. Neilsen, Hickox & Vrtilek 2004 and references therein). However, a ‘soft’ (energies < 1.0 keV) component has also been detected, whose profile does vary with superorbital phase in a manner consistent with reprocessing of the hard component by the inner regions of a warped accretion disc (Hickox & Vrtilek 2005).

The pulse-period history of SMC X-1 suggests its neutron star accretor has consistently been spun up over all five decades in which X-ray pulse-period measurements have been taken (Naik & Paul 2004), which suggests a persistent accretion disc despite its HMXB nature (the companion is a $17 M_{\odot}$ B0 supergiant; van der Meer et al. 2007).²

SMC X-1 therefore offers a crucial test-case in which changes in *both* the superorbital period and the accretion torque on the pulsar can be measured. We have therefore undertaken a side-by-side comparison of changes in superorbital period P_{sup} (from publicly-available *RXTE*/*ASM* data) and pulse period P_{spin} (from published pulse periods). Our aim is to learn whether changes in the putative accretion disc-warping and precession are indeed accompanied by changes in the accretion-induced torque on the neutron star.

This paper is the third in a series examining long-term X-ray periodicities in X-ray binaries and comparing them to the radiation-driven warping framework. Paper I demonstrated the variation in P_{sup} from SMC X-1, building on indications from Wojdowski et al. (1998) and Ribó et al. (2001). In Paper II, the analysis was extended to include Her X-1, LMC X-4, and Cyg X-2, covering a range of locations in the (binary radius – mass ratio) diagram of OD01 and suggesting that the radiation-driving framework of OD01 is consistent with measurements (Boyd & Smale 2004 provide a different but complementary interpretation). Both papers used *RXTE* All Sky Monitor data over the interval 1996 December – 2002 August.

This paper is organized as follows: the data sets used are introduced in Section 2, and their characteristics relevant to this study summarized; Section 3 communicates the techniques used to isolate changes in the superorbital and spin periods, respectively, with Appendices A and B presenting detailed methodology on the determination of superorbital period and particularly its measurement uncertainty. The side-by-side comparison of the superorbital- and spin-periods are then presented in Section 4, with implications for future work discussed in Section 5.

2 OBSERVATIONAL DATA

Most of the data used for this study were taken with the Rossi X-ray Timing Explorer (*RXTE*) satellite, using All-Sky Monitor data (Section 2.1) to chart the superorbital modulation and its variation, and the Inam, Baykal & Beklen (2010, hereafter I10) compilation of pointed observations with the Proportional Counter Array (PCA) to trace evolution in the pulse period. However, the pulse periods are not exclusively from *RXTE*; in the *RXTE* era, pulse-period measurements from *RXTE*, *BeppoSAX*, *Chandra*, *INTEGRAL*, and *XMM* are considered in the analysis, while ultimately the published pulse period history of SMC X-1 includes measurements from 16 missions (I10, Neilsen et al. 2004 and Section 2.2).

¹SMC X-1 is object SXP 0.72 in the Coe et al. (2005) catalogue of X-ray pulsars in the SMC, as well as in subsequent catalogues using the same nomenclature (e.g. Yang et al. 2017).

²SMC X-1 is not the only HMXB pulsar to show a near-monotonic spin-up trend over many years; at least a dozen additional systems show mostly negative period derivative within the SMC alone (e.g. Klus et al. 2014; Yang et al. 2017).

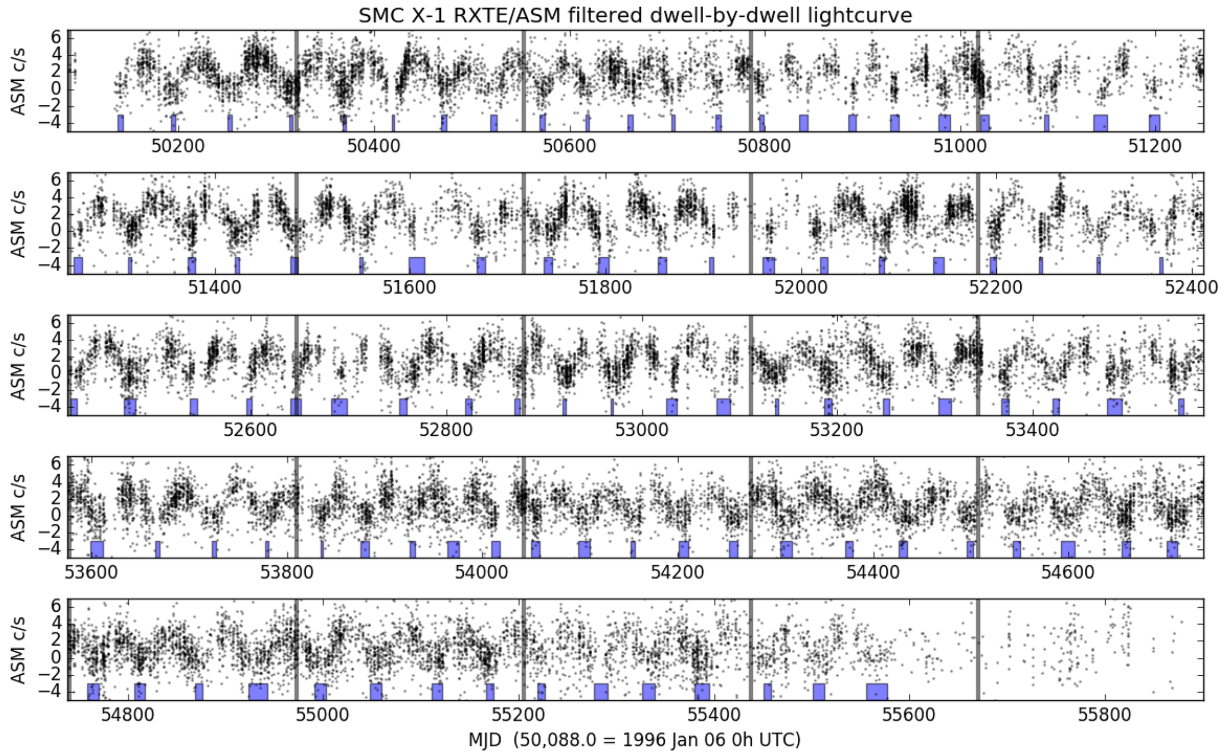


Figure 1. The *RXTE*/*ASM* dwell-by-dwell lightcurve used here. Blue-shaded regions show estimated cycle minima estimated using methods similar to those in Trowbridge et al. (2007). This plot shows the interval from the first measurement (at MJD 50 088.0 = 1996 January 06) to MJD 55 901 = 2011 December 05. Vertical grey lines indicate an example division of the lightcurve into 25 equal-length strips. See Section 3.1.

2.1 Superorbital modulation: *RXTE*/*ASM*

RXTE carried an All-Sky Monitor (*ASM*), which provided regular coverage of the entire sky at 1.3–12.1 keV energies. It consisted of three scanning detectors whose pointing and flux history were resolved into roughly $(0.20) \times 90$ s measurements of individual sources per day, with timing accurate to within a thousandth of a day (Levine et al. 1996). Crude spectral information is also available through three energy channels (1.3–3.0 keV, 3.0–5.0 keV, and 5.0–12.1 keV), although we do not use this channel-by-channel information in the present analysis. In total, *RXTE*/*ASM* accumulated just under 16 years’ monitoring on SMC X-1 (or about 1,494 binary orbits).

We considered the entire *RXTE*/*ASM* time baseline for SMC X-1 (MJD 50 088.35–55 902.96, or 1996 Jan 06–2011 Dec 07) in our analysis, although the coverage degrades somewhat in the final year of operation so that (for example) superorbital minima become difficult to distinguish; see e.g. Fig. 1. Following previous practice (e.g. Levine et al. 1996, Paper I, Trowbridge et al. 2007), we filtered the *ASM* data for background (< 10.0 counts s^{-1}) and on the quality of the least-squares fit to the flux (keeping points with reduced chi-square statistic $\chi^2_\nu < 1.5$). This filtering removed about 7 per cent of the data points (cutting the sample from 60 246 down to 56 263 measurements).

2.2 Pulse periods: *RXTE*/*PCA* and other missions

SMC X-1 has been observed many times with pointed X-ray observations (> 160 times with *RXTE*/*PCA* alone; I10), allowing its ≈ 0.71 second pulse period to be tracked on a timescale of decades. Wojdowski et al. (1998) presented a pulse period history of SMC X-1

over the interval 1971 January – 1996 January. Since then, considerably more pointed X-ray observations have taken place (e.g. I10, Raichur & Paul 2010 and Falanga et al. 2015). A complete pulse-period history of SMC X-1 including all recent observations from all missions does not yet appear to be available.

Instead, we rely heavily on the I10 compendium of pulse period measurements, which considers ~ 130 *RXTE*/*PCA* pointings totaling 250 ks over the time interval 1996 January – 2003 December, as well as a handful of pointed observations taken with *ROSAT*, *BeppoSAX*, *Chandra*, and *INTEGRAL*. This provides sufficient coverage to track the pulse period history of SMC X-1 over its first superorbital excursion near MJD 50 800 and its subsequent behaviour (e.g. Paper I), and to make predictions for the behaviour of SMC X-1’s spin period during the second (at MJD ≈ 54 000, e.g. Trowbridge et al. 2007) and further superorbital period excursions.

Full details of the extraction of the pulse periods themselves can be found in I10. Briefly, the authors first produced an initial estimate for the observed pulse period from the periodogram peak in the barycentre-corrected lightcurve (at energies 3–20 keV, the I10 pulse periods correspond to the ‘hard’ component of the pulse profile; Hickox & Vrtilek 2005). Folding on this initial estimate then produced an empirical pulse profile template, against which 200 s-long intervals of the lightcurve were cross-correlated to measure pulse arrival times. Where near-contiguous portions of the data covering several orbits were available, I10 deduced orbital elements, allowing an orbit-corrected history of SMC X-1 pulse periods. We extracted the spin frequencies and uncertainties from table 3 of I10, converting them into pulse periods and uncertainties.

Seven additional X-ray pulse-period measurements tabulated in Neilsen et al. (2004, hereafter N04) were also included in our analysis; as this second data set includes five measurements during the

two years after the first X-ray superorbital excursion event (three with *Chandra*, two with *XMM*, with two more *ROSAT* pulse period measurements filling out the coverage before the *RXTE* era), this substantially improves X-ray pulse-period coverage after the first superorbital period excursion.

The union of the **I10** and the **N04** pulse period measurements completes the set of archival X-ray pulse periods for SMC X-1. Not all archival measurements of otherwise sufficient duration yielded published pulse periods; for example, *Chandra* observations 1027–1030 at MJD 51 854, 51 856, 52 024, and 52 025 (all P.I. Vrtilek) were taken during the superorbital low state, while *RXTE* observations in 2004 Jan 27–30 inclusive (MJD 53 031–53 035, from program P80078, P.I. Eikenberry) apparently failed to detect the pulsar, despite the source not being at superorbital minimum at the time of these latter observations (Section 5.2).

3 METHODS

Here, we describe the methods we used to isolate changes in P_{sup} and P_{spin} .

3.1 Changes in superorbital variation

Fig. 1 shows the filtered *RXTE/ASM* lightcurve used here. To track the evolution of superorbital X-ray flux variations from SMC X-1 (whether in P_{sup} , or in the shape of the lightcurve, or both), we employed the Dynamical Power Spectrum (DPS), in which the Lomb–Scargle (LS) periodogram (Scargle 1982) is charted for sub-intervals within the total time interval, allowing the variation in power spectrum behaviour to be charted over time (e.g. Kotze & Charles 2012). Other methods that have been used to probe SMC X-1’s long-term variability include profile fitting in the time domain to estimate duration of individual variation cycles (Wojdowski et al. 1998; Trowbridge et al. 2007), wavelet analysis (Ribó et al. 2001), and the Hilbert–Huang transform (Hu et al. 2011, 2013); these methods find largely similar behaviour to power spectrum methods, with differences in sensitivity to instantaneous changes in superorbital variation time-scale.

The DPS approach suffers from two difficulties that complicate its interpretation, both related to the breaking of the total time baseline into intervals over which to compute each LS power spectrum. First, as pointed out by Trowbridge et al. (2007), if the time intervals overlap, there is a danger of artificially smoothing out variations in P_{sup} which may intrinsically be discontinuous. Secondly, the choice of the length τ of the time-intervals is not immediately obvious. At least a few cycles per interval are required to detect cyclic variation, but if τ is too great then the sensitivity to periodicity variation is lost.

We therefore characterized the DPS of SMC X-1 at two levels. The first investigation broke the total time baseline into 25 non-overlapping intervals, allowing the uncertainty in period detection to be estimated for a particular strip-length τ (Section 3.1.1). For the second investigation, the behaviour of the DPS was charted for a wide range of τ values (Section 3.1.2).

3.1.1 Period uncertainty for a given strip-length

The detection of stable periodicities in *RXTE/ASM* data is not always straightforward, as the periodogram shows structure due to windowing from the uneven sampling leading to structures in the periodogram which might mimic periodicities or shift the position

of detected peaks (e.g. Scargle 1982). Additionally, the underlying noise structure can often include time-scale-dependent ‘red noise’ that increases the false alarm probability for a periodic signal at a given LS power (e.g. Homer et al. 2001).

We therefore conducted a suite of Monte Carlo simulations to assess the likely period detection uncertainty associated with the DPS approach, tuning the parameters of the underlying noise model and injected periodicities to reproduce the observed peaks, the fine structure, and the underlying slope and plateaux of the LS power spectrum of the true data. The total time baseline was broken into 25 non-overlapping strips of equal duration (at $\tau = 230$ d per strip, this allows 3–4 superorbital cycles per strip; e.g. Fig. 1).

Full details about the determination of superorbital period measurement uncertainty are presented in Appendix A. Briefly, the synthetic lightcurves were generated by adding a smoothed representation of the detected superorbital modulation to a signal-free time-series with time-scale-dependent red-noise following the methodology of Timmer & Koenig (1995) as implemented in the Python module `DELCgen.py` (Connolly 2015).³ In tests, we found strong variation in execution time between the strips; a procedure to remove this variation, which afforded a factor ~ 20 reduction in total simulation time, is communicated in Appendix B.

The resulting synthetic lightcurves resemble the observed ASM data set quite well (Fig. 2 shows an example lightcurve and power spectrum for one of the synthetic lightcurves compared to that measured for SMC X-1). We were unable to reproduce the observed periodogram structure using a pure-noise model (with no superorbital periodicity added); a pure-noise model with break frequency at or near $1/P_{\text{sup}}$ and amplitude tuned to reproduce observed superorbital modulation, produces too much power (by a factor 10 or more) at timescales $> P_{\text{sup}}$.

The top panel of Fig. 3 shows the detected superorbital periods and the estimated measurement uncertainties (since these uncertainties are small, see also Fig. A4 in Appendix A). All but the final, poorly sampled strip show random uncertainty in recovered superorbital period in the range $0.22 \leq \text{rms}(P_{\text{sup}}) \leq 0.44$ d, with a high outlier at 0.90 d. The detections also show a small bias, in the sense $-1.0 \leq (\text{measured} - \text{simulated}) \leq +0.5$ d, although for the majority of strips this bias is in the range ± 0.25 d.

From our Monte Carlo trials, we therefore conclude the following: first, random uncertainty due to the Lomb–Scargle measurement method is not a significant contributor to the uncertainty in detected P_{sup} ; secondly, it is highly unlikely that the superorbital excursions near MJD 50 800 and 54 000 are spurious.

3.1.2 Allowing the strip-length to vary

Use of the DPS forces a choice as to the strip length τ . A particular choice of τ might cause intrinsic variation to be missed or artificially enhanced, depending on where the strip boundary lies with respect to intrinsic variation in either the profile or the period of superorbital modulation (see also Trowbridge et al. 2007).

As a second method to estimate uncertainty in P_{sup} , we therefore computed the DPS for a wide range of strip-lengths τ and examined the ensemble of DPS estimates thus produced for commonalities in structure. The uncertainty estimate thus produced contains contributions from intrinsic variation in light-curve structure, from statistical uncertainty due to both measurement uncertainty and temporal sampling variations, and from the range of choices of strip-length.

³<https://github.com/samconnolly/DELlightcurveSimulation>

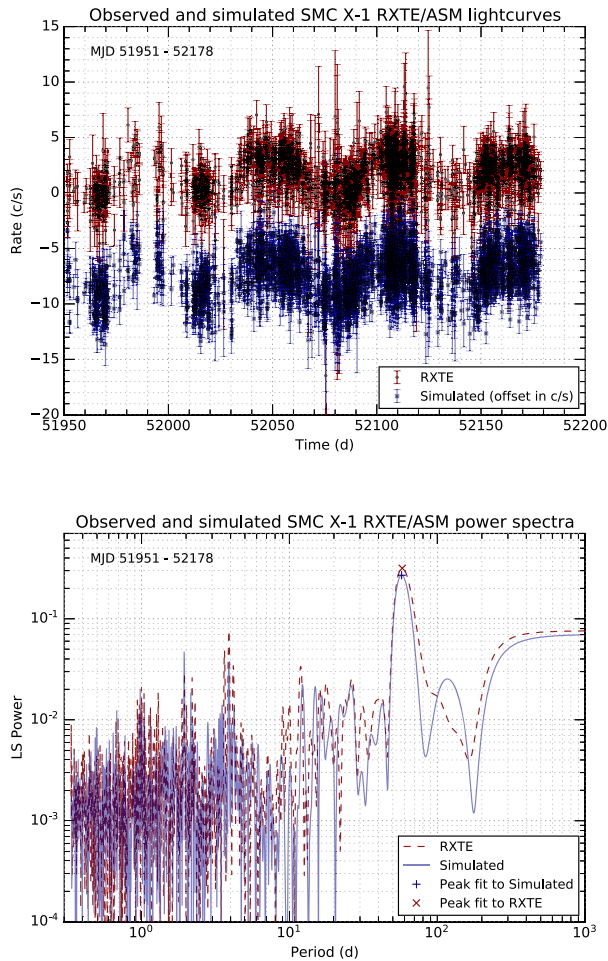


Figure 2. *Top:* Lightcurve for a data strip of length $\tau \approx 230$ d (MJD $\approx 51\,950$ – $52\,180$); the true RXTE/ASM is shown in red, the simulated lightcurve in blue and offset for clarity. *Bottom:* Log–log power spectrum of both RXTE data and simulated data showing the superorbital variation, as well as the 3.89-day orbit, over a time-scale-dependent noise continuum. Both the structure due to red-noise and the uneven sampling are clear. See Section 3.1.1 and Appendix A.

Where the detected period near a given date shows a small range of values over this ensemble, we interpret this as suggesting variation that is intrinsically reasonably stable. Conversely, where the detected period shows a large range for a particular date, we interpret this as indicating intrinsic changes on a short-enough time-scale that the choice of τ does impact the recovered value of P_{sup} .

The bottom panel of Fig. 3 shows the evolution of the superorbital peak in the DPS as computed separately for 3-day steps in $180 \leq \tau \leq 553$ d. The first and second superorbital excursions near MJD 50 800 and MJD 54,000 are visible with all strip-lengths, although the dates of the period minima can change by a few months depending on the value of τ . (Smoothness in the P_{sup} curves may be illusory, since each value of τ represents averaging over a different number of superorbital cycles.)

The scatter in recovered P_{sup} for all dates is larger (by about a factor 5–10) than that suggested by the Monte Carlo trials with constant P_{sup} (Section 3.1.1). We therefore adopt the rms scatter over the ensemble of strip-lengths τ as our running estimate of superorbital period uncertainty for the rest of this work.

3.2 Changes in spin period

Fig. 4 shows the pulse period history of SMC X-1. As shown by I10, SMC X-1 exhibits a long-term spin-up trend, consistent in sense (if not magnitude) over the almost five decades for which this source was observed.⁴ As pointed out by I10, the magnitude of this spin-up changes on a time-scale of years. I10 divide their four-decade data set into five segments, each separately fit with a linear trend (some of the segments have fewer than four samples in each). For the RXTE era the pulse history becomes substantially better-sampled, with >100 pulse-period measurements in the MJD 50 000–53 000 interval. We followed I10 and fit a piecewise linear trend to the pulse-period history, adopting the same partitioning by time interval. When fitting the spin-period trends to estimate P_{spin} , the interval MJD $>52\,500$ was explicitly avoided in the fitting to avoid artificially pulling this time interval back to the baseline behaviour. (For completeness, we did attempt to fit a single low-order polynomial to the entire pulse-period history of SMC X-1 in the manner of Wojdowski et al. (1998), but found that this process invariably left discontinuities in the pulse-period history when the complete time range of I10 was considered.)

3.2.1 Estimating the smoothed pulse-period derivative

Under our working hypothesis, if the spin- and superorbital variations really are related to each other, then we should see changes in P_{spin} on the same range of time-scales as those in P_{sup} . We are therefore interested in variations in P_{spin} of the order of weeks–years. That I10 & N04 do not report the *instantaneous* P_{spin} derivatives is therefore of no real consequence to the present investigation.

To isolate changes in spin-period derivative on time-scales of interest, the detrended values of P_{spin} were binned to 50 d (approximately one superorbital cycle) and the period derivative estimated as the derivative of a spline fit to the resulting smoothed periods. Use of this quantity, $\langle \dot{P}_{\text{spin}} \rangle$, is analogous to the application of a low-pass filter tuned to the average superorbital cycle length.

4 RESULTS

The comparison of superorbital period P_{sup} with the pulse-period P_{spin} and the smoothed pulse-period gradient $\langle \dot{P}_{\text{spin}} \rangle$ is shown in Figs 4 and 5. If there is a strict correlation between P_{sup} and either P_{spin} or $\langle \dot{P}_{\text{spin}} \rangle$, it is not simple or stable.

4.1 Variation of P_{spin} with P_{sup}

The first superorbital excursion (at MJD $\approx 50\,500$ – $51\,500$) coincides with an apparent shortening in P_{spin} over about the same time-scale. However, the P_{spin} behaviour after this interval does not closely resemble the changes in P_{sup} . Although coverage is sparse near MJD 53 000, the P_{spin} measurements in that region show no sign of returning to a baseline level, as P_{sup} appears to do (Fig. 4).

4.2 Variation of $\langle \dot{P}_{\text{spin}} \rangle$ with P_{sup}

Near the first superorbital excursion (MJD $\approx 50\,800$), the smoothed $\langle \dot{P}_{\text{spin}} \rangle$ history shows a striking similarity to the behaviour of P_{sup} ,

⁴SMC X-1’s pulse spin-up is not entirely monotonic; for example, a short spin-down event is visible in the pulse periods in the interval MJD 51 100–51 125 – itself during an interval with unusually intensive time-sampling.




Figure 3. Two estimates of the evolving superorbital period P_{sup} from RXTE/ASM data, using forms of the Dynamical Power Spectrum (DPS). *Top:* The time-series is broken into 25 non-overlapping strips of equal length, and the Lomb–Scargle periodogram computed separately for each. The grey-scale shows the DPS, with each strip renormalized to a maximum LS power of 1.0. Yellow errorbars are centred on the detected peak from each strip, and represent the 1σ uncertainties from Monte Carlo trials (see Section 3.1.1 for details). The Cyan triangles show the peak periods, offset by the median bias (in the sense detected minus simulated) suggested by Monte Carlo trials. *Bottom:* Maxima in the dynamical power spectrum, where the total ~ 5600 -day time interval is broken into non-overlapping strips of duration τ ; the complete run of DPS for $(180 \text{ d} \leq \tau \leq 553 \text{ d})$ is shown, in steps of τ of three days. See Section 3.1.

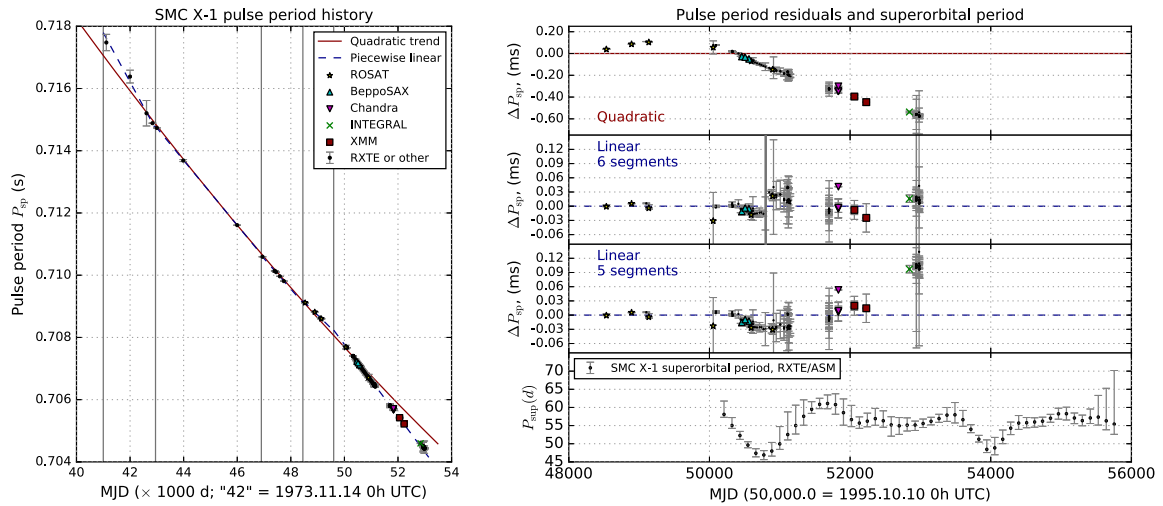


Figure 4. Pulse-period history of SMC X-1, from the compilation of I10, supplemented by pulse periods reported in N04. *Left-hand panel:* the pulse-period history. *Right-hand panels:* residuals after a quadratic (right-top) and piecewise linear trends (right upper- and lower-middle panels.) The piecewise linear fit with five segments uses the same time partitioning as I10, which assigns the entire RXTE mission to a single partition. For the six-segment scheme, the RXTE era is partitioned at MJD 50 800 d (the vertical grey line in the right upper-middle panel). Plot symbols and colours indicate different missions: after MJD 48 000, black points with errorbars indicate RXTE observations, while before this date the same symbols indicate 10 missions (*Uhuru*, *Aerobee*, *Apollo-Soyuz*, *SAS 3*, *Ariel V*, *Einstein*, *EXOSAT*, *Ginga*, *HEXE*, or *ASCA*). In all panels, the dark red solid line indicates the quadratic trend fit to the entire pulse-period history, while the dark blue dashed line shows the piecewise linear fits. Finally, the *right-hand-bottom* panel shows the RXTE/ASM superorbital period determination; medians are interpolated from the run over light-curve strip-lengths τ presented in the bottom panel of Fig. 3, with asymmetric 1σ errorbars indicating the root mean square of the variation of superorbital period above and below the median detected values. See Section 3.2.




Figure 5. Smoothed pulse-period derivative history overlaid on the superorbital variations. The superorbital period variation is represented by the median values of P_{sup} over the range of τ values, with 1σ upper and lower bounds indicated as grey-scale. See Section 4.

with changes in $\langle \dot{P}_{\text{spin}} \rangle$ taking place about 200 d before the changes in P_{sup} (Fig. 5). As with the detrended spin period P_{spin} , however, the intervals at $\text{MJD} > 51\,500$ do not obviously confirm the trend – although we caution again that the coverage in pulse period measurements is relatively sparse in this interval.

5 DISCUSSION

Even with the huge P_{spin} data set amassed by I10, complemented by the measurements of N04, we cannot claim a detection of a strict correlation between P_{sup} and either P_{spin} or $\langle \dot{P}_{\text{spin}} \rangle$. Both these quantities do show variations that appear to *coincide* with variations in P_{sup} , but only for the interval near the first superorbital excursion at $\text{MJD} \approx 50\,800$.

5.1 X-ray modulation scenarios

Without coverage of the *second* superorbital period excursion near $\text{MJD} 54\,000$, a number of interpretations are still consistent with the publicly-available RXTE measurements. We consider the alternatives in turn.

5.1.1 Correlated excursions in both P_{spin} and P_{sup}

One possibility is that P_{sup} and P_{spin} both experience an excursion to shorter periods nearly simultaneously. This might indicate a scenario in which the accretion disc and inner accretion flow both change simultaneously. However, behaviour after the first P_{sup} excursion is markedly different, and, under this scenario, it is not clear why the cluster of P_{spin} measurements near $\text{MJD} 53\,000$ does not follow the same behaviour.

5.1.2 A change in slope of P_{spin} , coincident with an excursion in P_{sup}

As shown by I10, and references therein, and as we confirm for this study, SMC X-1’s pulse-period history is best characterized as a series of linear segments of differing gradient but all decreasing in period, rather than a single low-order polynomial over its entire spin period history (Fig. 4). Discontinuous changes in spin-period derivative are not unknown in X-ray pulsars; examples include GX 1+4 (Chakrabarty et al. 1997), 4U 1626-67 (which, like SMC X-1,

is Roche-lobe fed Camero-Arranz et al. 2010) and GX 301-2 (e.g. Kaper, van der Meer & Najarro 2006), although in those cases the changes in spin period are much more dramatic than in SMC X-1, which never undergoes an extended interval of torque reversal. Models to explain these changes in spin-period gradient generally suggest a change in accretion rate onto the pulsar (e.g. Yi & Vishniac 1999), possibly caused by changes in the inner accretion disc (e.g. Wijers & Pringle 1999). For the case of SMC X-1, the discontinuous changes in \dot{P}_{spin} have generally taken place between intervals of P_{spin} coverage (I10, also the left-hand panel of Fig. 4).

However, it may be that what we interpreted in Section 5.1.1 as an excursion in P_{spin} near $\text{MJD} 50\,800$, may more accurately be described as another of the discontinuous changes in the spin-period derivative \dot{P}_{spin} . Under this interpretation, the spin period behaviour near $\text{MJD} 50\,800$ may represent a transition to a slightly slower spin-up rate from the SMC X-1 pulsar. The right-hand panel of Fig. 4 shows pulse-period residuals over the RXTE era, divided at $\text{MJD} 50\,800$, approximately the epoch of the first superorbital excursion. Under this model, the spin-period gradient changes from $dP_{\text{spin}}/dt = -0.432(2) \text{ ms yr}^{-1}$ before $\text{MJD} 50\,800$ to $-0.396(6) \text{ ms yr}^{-1}$ after it (with the figures in parentheses giving the formal uncertainty in the last decimal place; these gradients are consistent with the smoothed period gradients shown in Fig. 5). The reduced chi-square statistic χ^2_{ν} is unlikely to show values close to unity when linear trends are used to fit structure as complicated as that in SMC X-1’s pulse-period history (particularly in the presence of a small number of measurements with extremely small uncertainty estimate; I10), but we do see an improvement in the fit when the RXTE era is partitioned in this way: representing the RXTE era with a single linear trend produces $\chi^2_{\nu} = 122,738/135 = 909.2$ while a two-segment fit produces $\chi^2_{\nu} = 45,244/133 = 340.2$. Thus, while these χ^2_{ν} are far from formally acceptable,⁵ they do suggest improvement in the fit to the pulse period history when a two-segment fit is used for the RXTE era.

We thus suspect that the compilation of I10 may have caught a torque-change event from SMC X-1 near $\text{MJD} \approx 50\,800$. If so, then the P_{sup} excursion near $\text{MJD} 50\,800$ may relate to the putative torque change near the same date. With these data, we are unable to match the events more accurately than a few viscous time-scales.

5.1.3 A link between $\langle \dot{P}_{\text{spin}} \rangle$ -and P_{sup} ?

If indeed $\langle \dot{P}_{\text{spin}} \rangle$ and P_{sup} are varying together, a time-lag of $\Delta t \sim 200 \text{ d}$ between variations in P_{sup} and $\langle \dot{P}_{\text{spin}} \rangle$ might be expected. If the smoothed $\langle \dot{P}_{\text{spin}} \rangle$ traces changes in the accretion rate onto the neutron star, then the interception of accretion luminosity by the warped disc might lead to adjustment in the disc configuration. With disc adjustments taking place on one to a few viscous time-scales (OD01), this might then lead to changes in measured P_{sup} that lag those in $\langle \dot{P}_{\text{spin}} \rangle$ by $\sim 200 \text{ d}$.

5.1.4 No relation between the spin- and superorbital behaviours

A final possibility is that there is no intrinsic link between the variation in superorbital and pulse period. We note that the data considered in this paper together do *not* rule out this possibility. In

⁵Under the hypothesis that residual variation after subtraction of the best-fit model is due to measurement uncertainty, producing $\chi^2_{\nu} \approx 1$ would require multiplying the pulse period uncertainties by a factor ~ 18 , which seems unreasonably large.

this scenario, the accretion disc and the pulsar both do respond to local changes in accretion rate, but changes in the accretion flow are so complex that variations in the accretion disc configuration are observationally decoupled from variations in the pulse period. We point out that this scenario would not invalidate the radiation-driven warping scenario (in which the accretion luminosity from the accretor drives the disc into its warped configuration; Wijers & Pringle 1999, OD01), but instead the changes in the pulsar and the disc are too intrinsically noisy to correlate.

5.2 Pointed X-ray coverage of the second P_{sup} excursion

The second superorbital excursion, at approximately MJD 53 500–54 500 (Trowbridge et al. 2007; Kotze & Charles 2012, roughly the 2005 May – 2008 Feb interval) is not covered in the I10 compendium of pulse periods. We examine here the archival coverage of pointed observations which may allow P_{spin} measurement during and after the second superorbital excursion.

To our knowledge, the most recent pointed RXTE/PCA measurements sensitive to the pulse period, are those obtained on MJD 52 981–52 987 (2003 Dec 08–14) and published in Raichur & Paul (2010) and Falanga et al. (2015). A recent examination of type-II bursts from SMC X-1 in its high state also uses RXTE/PCA data no more recent than 2003 December; Rai, Pradhan & Paul (2018).

The HEASARC archive shows only a single set of RXTE/PCA pointed observations with coverage beyond 2003 December. This is RXTE program 80078 (PI Eikenberry), which observed SMC X-1 several times in the 2003 October – 2004 January interval. A subset of these observations were used by I10 and are tabulated in their table 3. Curiously, the 2004 January observations within program 80078 do not appear in I10, and a subsequent processing using the Southampton RXTE/PCA pipeline (see Yang et al. 2017)⁶ confirms that the pulsar is not detected in those observations. These observations did not take place during mid-eclipse (using the ephemeris of Falanga et al. 2015), nor precisely at superorbital minimum, and the RXTE/ASM did detect SMC X-1 over the same interval, so it remains unclear why the SMC X-1 pulsar has not been detected in the 2004 January observations. Regardless, for dates more recent than 2004 January, the HEASARC archive does not contain any RXTE/PCA pointings within one degree of SMC X-1 (roughly the distance off-axis within which RXTE/PCA has sensitivity).

While SMC X-1 has been observed several times by other missions during and after the second superorbital excursion (roughly 2005 May – 2008 Feb), it is not immediately obvious that any of the measurements will allow useful measurements of SMC X-1's pulse period. For example, SMC X-1 was within the *ACIS-I* field of view for *Chandra* observations 5493 & 5494 (both P.I. Coe, taken during 2008 Feb and March), but the 3.2 s time resolution of *ACIS-I* is too poor to permit pulse period measurement of this source. Of the five *Swift* observations made of SMC X-1 during the 2006–2009 interval (obsIDs 00035216001, 00035216002, 00035216003, 00037160002, & 00037160003), only the first was taken in a mode with sufficient time resolution to measure SMC X-1's spin period. However, as these data were taken in Windowed-Timing mode (by which the image is collapsed along one spatial axis) and SMC X-1 itself was not at the aim-point of the observation, it is not clear that SMC X-1's spin period can be measured from these data. Finally, *XMM* observed the Nova SMC 2005 field, including SMC X-1, during 2006 March 13th (MJD 53 807.64 – 53 807.77, ObsID

0311590601, PI Scharfel). Using the ephemeris presented in table 3 of Falanga et al. (2015), this observation took place at orbital phase 0.90–0.94. With eclipse commencing at phase ≈ 0.90 (Falanga et al. 2015), it is likely that extracting a pulse period measurement from this latter observation may also be problematic.

Extraction of pulse periods from this heterogeneous and patchy archival data set near the second superorbital excursion is beyond the scope of the present investigation; at this juncture, we expect perhaps 0–2 pulse-period measurements will be extractable from this interval. However, the *MAXI* satellite continues to furnish long-term X-ray monitoring coverage of bright X-ray binaries in a manner similar to the *RXTE/ASM* (e.g. Sugimoto et al. 2016), including SMC X-1 itself (Sugimoto et al. 2014). Therefore, we anticipate X-ray monitoring of further superorbital period excursions of SMC X-1 will allow triggering of observational campaigns to measure P_{spin} evolution coincident with further large changes in P_{sup} .

5.3 Predictions for future P_{sup} excursions

While comprehensive mining of archival data for MJD $\geq 53,000$ is beyond the scope of this paper, we can make the following predictions for what analysis of the full archive will reveal, depending on the scenario:

- (i) If the P_{spin} and P_{sup} excursions are strictly correlated, then P_{spin} will decrease during the second and future superorbital period excursions, returning to a baseline level after the end of the excursion (Section 5.1.1);
- (ii) If instead the P_{sup} excursion heralds a sudden change in torque, then the spin-period gradient \dot{P}_{spin} will change near the middle of the excursion, and remain roughly constant until any future superorbital excursions take place (Section 5.1.2);
- (iii) If the smoothed spin-period gradient $\langle \dot{P}_{\text{spin}} \rangle$ varies with P_{sup} , then the variation of this quantity will resemble that of the P_{sup} excursion itself, just shifted earlier by ~ 200 d (Section 5.1.3);
- (iv) If instead there is no link between P_{spin} and P_{sup} , then the P_{spin} behaviour during superorbital excursion will show no relation to that of P_{sup} (Section 5.1.4).

5.4 Comparison to Her X-1

The eclipsing X-ray binary Her X-1 (P_{spin} 1.24s, orbital period 1.70d; donor mass $2.2 M_{\odot}$, e.g. Jurua et al. 2011) provides a second example of an X-ray binary system in which both the X-ray pulse behaviour and superorbital modulation may be tracked together, providing important context for the present study. This source shows a 35-day superorbital modulation, with intrinsic scatter of about ± 2 d (e.g. fig. 5 of Leahy & Igna 2010), or about ± 7 per cent variation in P_{sup} compared to ± 20 per cent for SMC X-1. The relative stability of the superorbital period exhibited by Her X-1 is broadly consistent with its system properties under the radiation-driven warping framework (e.g. OD01, Paper II). Under this framework, the figure of the accretion disc warp in Her X-1 thus likely changes less severely than is the case for SMC X-1, but is still characterized by a range of superorbital periods and associated disc configurations against which other behaviours can be compared. A correlation has been measured between P_{sup} and the average X-ray flux during a superorbital cycle (Leahy & Igna 2010), which might suggest a link between P_{sup} and \dot{M} .

The relatively stable superorbital modulation of Her X-1 is complicated by occasional ‘Anomalous Low States’ (ALS), in which the observed X-ray flux drops to very faint levels for intervals of a

⁶<http://www.xraypulsars.space/projects>

few months, coincident with and possibly also preceded by intervals of rapid spin-down (e.g. Parmar et al. 1999, Staubert et al. 2006) and most likely corresponding to an increase in X-ray obscuration rather than a cessation of \dot{M} (see Still et al. 2001 and references therein). The currently favored model is that the ALS state indicates a small change in the form of the accretion disc warp, leading to both an increase in X-ray obscuration and altering the instantaneous \dot{M} through the inner disc (e.g. Staubert et al. 2006). That the superorbital modulation at optical wavelengths shows very little change between the ALS and normal high-state, suggests that the disc warp need not change strongly in its form to produce the large increase in X-ray obscuration characterizing the ALS (Jurua et al. 2011, and references therein).

Outside ALS events, Her X-1 shows strong correlation between the pulse period and the phasing of the main-on in the superorbital cycle, with the superorbital main-on variations lagging behind the P_{spin} variations by a few tens of days (Staubert et al. 2006). Additionally, variations in the pulse profile from Her X-1 are precisely synchronized with superorbital main-on phasing (Staubert et al. 2013). The two leading candidates to account for the pulse-profile variations from Her X-1 are free precession of the neutron star (e.g. Postnov et al. 2013 and references therein), and varying occultation of the neutron star by the inner accretion disc (Scott, Leahy & Wilson 2000), although both models have difficulties (see Staubert et al. 2013 for discussion).

The two sources differ in the energy dependence of pulse-profile variations. From Her X-1, the pulse profile at 9–13 keV is observed to vary (e.g. Staubert et al. 2013), while from SMC X-1 it is only the soft (blackbody, at $\lesssim 1.0$ keV) component of the pulse profile that varies with superorbital phase (Hickox & Vrtilek 2005). While the Scott et al. (2000) mechanism invokes progressive occultation by the inner accretion disc to explain pulse-profile variations seen from Her X-1, the Hickox & Vrtilek (2005) mechanism invokes varying reprocessing at the inner disc to explain pulse-profile variations seen from SMC X-1.

While we have not performed precisely the same tests as Staubert et al. (2006) – we characterize the period of the superorbital modulation from SMC X-1 over several cycles rather than its phasing from individual cycles – our results from SMC X-1 are probably intermediate between the lockstep correlation observed from Her X-1 and a case with no correlation at all. We see an excursion in P_{spin} that appears to coincide with an excursion in P_{sup} (and, in the case of the smoothed P_{spin} derivative, also possibly leading the superorbital modulation by tens of days) but spin–superorbital correlation does not appear to recur later in the *RXTE* interval.

The observed behaviour from SMC X-1 is probably more consistent with ALS-type behaviour (a small change in accretion disc-warp figure accompanied by an episode of pulsar spin-down) than with the correlated changes in (presumed) accretion warp figure and pulse profile seen from Her X-1. While the physics underlying the (presumed) response of the accretion disc to changes in the pulsar behaviour may be expected to be similar for SMC X-1 as for Her X-1 (with the main difference being the stability of the accretion disc warp configuration and the complexity of the form of the warp), more detailed cross-comparison of the two systems probably awaits X-ray pulse-period measurements during another superorbital excursion from SMC X-1.

6 CONCLUSIONS

With X-ray variability measurements that trace both superorbital and pulsar spin-period time-scales, and with large-amplitude

(~ 20 per cent) variation in superorbital period, SMC X-1 is perhaps the best test-case for probing the interplay between the behaviour of a large, persistent accretion disc, and variations in the accretion flow onto the compact object. We have therefore performed an investigation to determine whether variations in the pulsar-spin-period derivative are correlated with long-term variations in the superorbital X-ray variability of this source. Our conclusions are:

- (i) We find an intriguing possible correlation between the behaviour of P_{sup} and the pulsar spin period, particularly in the interval around the first P_{sup} excursion at MJD $\approx 50\,800$;
- (ii) The measurements do not yet allow us to discriminate between several possible scenarios. The measurements are most consistent with either a model in which the accretion disc and pulsar are both responding to a discontinuous change in accretion flow, possibly regulated by the disc (Section 5.1.2), or one in which the variations in pulsar torque and accretion disc configuration are too complicated to be observationally coupled (Section 5.1.4);
- (iii) The second P_{sup} excursion near MJD 54 000 offers an opportunity to distinguish between these scenarios, at least in principle, but the coverage of observations *sufficient to measure the pulse period* is likely much more sparse, and we defer detailed analysis of these observations to further work;
- (iv) With ongoing long-term X-ray monitoring from *MAXI*, future superorbital-period excursions may allow renewed observational probing of the variation of SMC X-1’s spin-period evolution as the superorbital period changes.

We therefore believe that a follow-up to I10 is now required, to assess the full pulse-period history of SMC X-1. By comparing the pulse-period history during the complete *RXTE* interval to the superorbital modulation, including the second superorbital excursion near MJD 54,000 and subsequent recovery, it may be possible to reveal if and how the accretion disc and pulsar really are responding to the same underlying processes. If, however, archival observations are indeed not sufficient to probe the spin-period evolution during this second superorbital excursion, then further probing of the spin–superorbital link in SMC X-1 may have to await future observations, triggered by long-term monitoring by *MAXI* or similar future facilities. This would allow new tests of the radiation-driven warping framework (OD01), and finally allow SMC X-1 to yield its secrets into the accretion/outflow process that is so central to astrophysics at all scales.

ACKNOWLEDGEMENTS

This work was only possible because of the efforts of the ASM/*RXTE* team at MIT and NASA/GSFC. This research has made use of data and/or software provided by the High Energy Astrophysics Science Archive Research Center (HEASARC), which is a service of the Astrophysics Science Division at NASA/GSFC and the High Energy Astrophysics Division of the Smithsonian Astrophysical Observatory. This research made use of *Astropy*, a community-developed core Python package for Astronomy (*Astropy Collaboration 2013*). This project made use of the *sympy* symbolic algebra Python library. PAC gratefully acknowledges receipt of a grant from the Leverhulme Trust. SL acknowledges support of NASA Astrophysics Data Analysis program grant 80NSSC18K0430. We thank Malcolm Coe, of the University of Southampton, for insightful discussion. We thank the anonymous referee for a thorough and careful reading of the manuscript, and for thoughtful comments and suggestions which greatly improved the paper.

REFERENCES

- Astropy Collaboration, 2013, *A&A*, 558, A33
- Blondin J. M., Woo J. W., 1995, *ApJ*, 445, 889
- Bonnet-Bidaud J. M., van der Klis M., 1981, *A&A*, 97, 134
- Boyd P. T., Smale A. P., 2004, *ApJ*, 612, 1006
- Camero-Arranz A., Finger M. H., Ikhsanov N. R., Wilson-Hodge C. A., Beklen E., 2010, *ApJ*, 708, 1500
- Chakrabarty D. et al., 1997, *ApJ*, 481, L101
- Charles P., Kotze M., Rajoelimanana A., 2010, *AIPC*, 1314, 303
- Charles P. A., Coe M. J., 2006, in Lewin W. H. G., van der Klis M., eds, *Compact Stellar X-Ray Sources*. Cambridge Univ. Press, Cambridge, UK, p. 215
- Clarkson W. I., Charles P. A., Coe M. J., Laycock S., 2003b, *MNRAS*, 343, 1213 (Paper II)
- Clarkson W. I., Charles P. A., Coe M. J., Laycock S., Tout M. D., Wilson C. A., 2003a, *MNRAS*, 339, 447 (Paper I)
- Coe M. J., Angus R., Orosz J. A., Udalski A., 2013, *MNRAS*, 433, 746
- Coe M. J., Edge W. R. T., Galache J. L., McBride V. A., 2005, *MNRAS*, 356, 502
- Connolly S., 2015, preprint ([arXiv:1503.06676](https://arxiv.org/abs/1503.06676))
- Falanga M., Bozzo E., Lutovinov A., Bonnet-Bidaud J. M., Fetisova Y., Puls J., 2015, *A&A*, 577, A130
- Foulkes S. B., Haswell C. A., Murray J. R., 2006, *MNRAS*, 366, 1399
- Foulkes S. B., Haswell C. A., Murray J. R., 2010, *MNRAS*, 401, 1275
- Frank J., King A., Raine D. J., 2002, *Accretion Power in Astrophysics*, 3rd edn. Cambridge Univ. Press, Cambridge, UK
- Hickox R. C., Vrtilek S. D., 2005, *ApJ*, 633, 1064
- Homer L., Charles P. A., Hakala P., Muhli P., Shih I.-C., Smale A. P., Ramsay G., 2001, *MNRAS*, 322, 827
- Horne J. H., Baliunas S. L., 1986, *ApJ*, 302, 757
- Hu C.-P., Chou Y., Wu M.-C., Yang T.-C., Su Y.-H., 2011, *ApJ*, 740, 67
- Hu C.-P., Chou Y., Yang T.-C., Su Y.-H., 2013, *ApJ*, 773, 58
- Inam S. Ç., Baykal A., Beklen E., 2010, *MNRAS*, 403, 378 (I10)
- Jurua E., Charles P. A., Still M., Meintjes P. J., 2011, *MNRAS*, 418, 437
- Kaper L., van der Meer A., Najarro F., 2006, *A&A*, 457, 595
- Klus H., Ho W. C. G., Coe M. J., Corbet R. H. D., Townsend L. J., 2014, *MNRAS*, 437, 3863
- Kotze M. M., Charles P. A., 2012, *MNRAS*, 420, 1575
- Kovacs G., 1981, *Ap&SS*, 78, 175
- Leahy D. A., Igna C. D., 2010, *ApJ*, 713, 318
- Levine A. M., Bradt H., Cui W., Jernigan J. G., Morgan E. H., Remillard R., Shirey R. E., Smith D. A., 1996, *ApJ*, 469, L33
- Levine A. M., Bradt H. V., Chakrabarty D., Corbet R. H. D., Harris R. J., 2011, *ApJS*, 196, 6
- Liu Q. Z., van Paradijs J., van den Heuvel E. P. J., 2006, *A&A*, 455, 1165
- Naik S., Paul B., 2004, *A&A*, 418, 655
- Neilsen J., Hickox R. C., Vrtilek S. D., 2004, *ApJ*, 616, L135 (N04)
- Ogilvie G. I., Dubus G., 2001, *MNRAS*, 320, 485 (OD01)
- Parmar A. N., Oosterbroek T., dal Fiume D., Orlandini M., Santangelo A., Segreto A., del Sordo S., 1999, *A&A*, 350, L5
- Postnov K., Shakura N., Staubert R., Kochetkova A., Klochkov D., Wilms J., 2013, *MNRAS*, 435, 1147
- Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 2002, *Numerical Recipes in C++ : the Art of Scientific Computing*. Cambridge Univ. Press, Cambridge, UK
- Pringle J. E., 1996, *MNRAS*, 281, 357
- Rai B., Pradhan P., Paul B. C., 2018, *RAA*, preprint ([arXiv:1806.03244](https://arxiv.org/abs/1806.03244))
- Raichur H., Paul B., 2010, *MNRAS*, 401, 1532
- Ribó M., Peracaula M., Paredes J. M., Núñez J., Otazu X., 2001, in Gimenez A., Reglero V., Winkler C., eds, *ESA Special Publication Vol. 459, Exploring the Gamma-Ray Universe*. ESA Publications Division, Noordwijk, The Netherlands, p. 333
- Scargle J. D., 1982, *ApJ*, 263, 835
- Scott D. M., Leahy D. A., Wilson R. B., 2000, *ApJ*, 539, 392
- Staubert R., Klochkov D., Vasco D., Postnov K., Shakura N., Wilms J., Rothschild R. E., 2013, *A&A*, 550, A110
- Staubert R., Schandl S., Klochkov D., Wilms J., Postnov K., Shakura N., 2006, in D'Amico F., Braga J., Rothschild R. E., eds, *American Institute of Physics Conference Series Vol. 840, The Transient Milky Way: A Perspective for MIRAX*. AIP Publishing, Melville, New York, USA, p. 65
- Still M. et al., 2001, *ApJ*, 554, 352
- Sugimoto J., Mihara T., Kitamoto S., Matsuoka M., Sugizaki M., Negoro H., Nakahira S., Makishima K., 2016, *PASJ*, 68, S17
- Sugimoto J., Mihara T., Sugizaki M., Serino M., Kitamoto S., Sato R., Ueda Y., Ueno S., 2014, *Proceedings of the 12th Asia Pacific Physics Conference (APPC12), JPS Conference Proceedings, Vol. 2*. The Physical Society of Japan, Tokyo, Japan, p. 1
- Timmer J., Koenig M., 1995, *A&A*, 300, 707
- Trowbridge S., Nowak M. A., Wilms J., 2007, *ApJ*, 670, 624
- van der Meer A., Kaper L., van Kerkwijk M. H., Heemskerk M. H. M., van den Heuvel E. P. J., 2007, *A&A*, 473, 523
- Vrtilek S. D., Raymond J. C., Boroson B., McCray R., 2005, *ApJ*, 626, 307
- Wijers R. A. M. J., Pringle J. E., 1999, *MNRAS*, 308, 207
- Wojdowski P., Clark G. W., Levine A. M., Woo J. W., Zhang S. N., 1998, *ApJ*, 502, 253
- Yang J., Laycock S. G. T., Christodoulou D. M., Fingerman S., Coe M. J., Drake J. J., 2017, *ApJ*, 839, 119
- Yi I., Vishniac E. T., 1999, *ApJ*, 516, L87

APPENDIX A: LOMB–SCARGLE PERIOD DETECTION UNCERTAINTY FOR LONG-PERIOD VARIABILITY

Several effects complicate the effort to estimate the measurement uncertainty on the superorbital periods shown by SMC X-1, as we demonstrate below. With recurrence interval in the range 45–60 d, the superorbital variation from SMC X-1 takes place in a period regime at which time-scale-dependent red-noise can be important (e.g. Homer et al. 2001). Intrinsic variability at these time-scales can cause substantial power at other periods, and a period-finding by Gaussian fitting to power spectrum peaks at long time-scales can be biased (usually to higher periods) since the peaks themselves are not Gaussian. Pre-treatment of the power spectrum to remove underlying red-noise can circumvent these difficulties to some degree (e.g. Levine et al. 2011), but to produce superorbital periods and their uncertainties with as few assumptions as possible, we have opted to work with the power spectrum without pre-treatment. This Appendix communicates the techniques used to estimate the superorbital period, the underlying noise model, and thus the measurement uncertainty of the superorbital period for each of the time strips in the *RXTE/ASM* lightcurve.

The superorbital period measurement uncertainty was estimated by generating a large number of synthetic *RXTE/ASM* lightcurves tuned to match the observed lightcurve and power spectrum of each time strip, and the random and systematic components of the period uncertainty estimated by characterizing the distribution of the difference ΔP between the recovered and simulated superorbital periods. In the process, we accounted for bias due to asymmetry of the power spectrum peak (Appendix A1) and paid particular attention to the shape of the underlying noise continuum in the periodogram (Appendix A2).

For each simulation, the synthetic lightcurve was generated using the Timmer & Koenig (1995) method (for a time-scale-dependent red-noise continuum) and from unit Gaussian random variates scaled by *RXTE/ASM* measurement uncertainties (testing the white-




Figure A1. Illustration of bias in a Gaussian fit (dashed line) to the strongest Lomb–Scargle power spectrum peak from SMC X-1 (solid line). Due to asymmetry in the true profile, the Gaussian peak is artificially shifted to higher values. The thick grey bar shows the selection region bracketing the true peak; instead of the peak of the Gaussian approximation, we use the maximum profile value within this selection region to determine the power spectrum peak period. See Section A1.

noise case). Sinusoids at the orbital period and its first harmonic were added to simulate orbital modulation, and the superorbital modulation injected by adding the fold to the observed data set on the detected superorbital period. In this way, the average superorbital modulation profile for a given strip is preserved, allowing spectral leakage from the superorbital modulation to be accurately included in the simulations. Fig. 2 in the main paper shows an example of the quality of the match obtained to the measured data through these methods.

A1 Power spectrum peak asymmetry bias

At long periods, peaks in the Lomb–Scargle power spectrum from SMC X-1 become sufficiently asymmetric that a Gaussian fit to the peak is biased high by up to about 2.5 d – much larger than the typical random uncertainty we estimate in the period detection. Fig. A1 shows an example. Instead of using a model fit to the power spectrum peak, we use the Gaussian approximation to determine a region of evaluation that brackets the true peak, then determine the maximum value of the true peak using a highly oversampled evaluation of the power spectrum within this region.

A2 Continuum model

To characterize the power spectrum continuum, the observed power spectrum was coarsely binned in log–log space (ignoring the period interval surrounding the detected superorbital peak, as well as intervals bracketing the orbital period and its first harmonic). To remove spectral leakage from the superorbital modulation and thus isolate the underlying continuum, a fold on the superorbital modulation was subtracted from the lightcurve, followed by subtraction of a fold on the orbital period. Fig. A2 shows an example; the subtraction of measured orbital and superorbital variability reveals an underlying continuum that depends on modulation time-scale.

The strips show substantial variability in the indicated continuum (Fig. A3). In some cases, subtraction of the superorbital modulation leaves behind a time-scale-dependent ‘red-noise’ continuum with weaker time-scale dependence than before. In some others, removal




Figure A2. Estimating the power spectrum continuum for a single light-curve strip. The red curve and step-plot show the observed and binned Lomb–Scargle power spectrum, respectively. The grey curve and dashed black step-plot show the observed and binned power spectrum, after folds on the superorbital and orbital periods are subtracted from the observed lightcurves. The grey bars along the top horizontal axis indicate regions ignored from the continuum fitting due to proximity to the superorbital and orbital power spectrum peaks. See Section A2.

of the superorbital modulation almost removes the time-scale dependence of the remaining power spectrum altogether. This might indicate that the remnant noise continuum may be substantially due to intrinsic period and/or profile variation – which would manifest as apparent ‘red-noise’ when a constant-profile, constant-period modulation were subtracted.

To characterize period uncertainty, the median over all the strips of the periodicity-subtracted power spectrum was used. This median continuum power spectrum is reasonably well-fit by a Bending Power-Law (BPL), of the form (Connolly 2015)

$$f(\nu) = A \nu^{-a_{lo}} \left[1 + \left(\frac{\nu}{\nu_{bend}} \right)^{a_{hi}-a_{lo}} \right]^{-1} + c \quad (\text{A1})$$

where $\nu = 1/P$. Based on the fit to the fold-subtracted continuum power spectrum (Fig. A3, bottom panel), for the model parameters $\{A, a_{lo}, a_{hi}, \nu_{bend}, c\}$ we adopt the values listed in the right-most column of Table A1 to produce the synthetic data sets. (In practice, while we find $a_{lo} \approx 0$, we retain the full expression as implemented in the DELCgen.py simulation methods of Connolly 2015.)

A3 Comparison of period uncertainties

Here, we present the comparison between the superorbital period uncertainty in the full BPL-noise case and the simpler case assuming timescale-independent white-noise, as well as comparing our simulation-based uncertainties with two standard prescriptions in the literature.

Fig. A4 shows the Monte Carlo estimates of random and median offset $\Delta P = P(\text{measured}) - P(\text{injected})$ thus obtained, for the BPL continuum model and for a white-noise model. For both noise models, the typical random uncertainty in the superorbital period is close to $\sigma(\Delta P) \approx 0.35$ d for the BPL noise model, and close to $\sigma(\Delta P) \approx 0.20$ d for the white-noise model. This is qualitatively as expected, since under a BPL noise model, the continuum surrounding the main power spectrum peak is stronger. While period uncertainties are small in the presence of both types of noise, the




Figure A3. Continuum power spectra for the first 24 light-curve strips (the 25th and final strip shows such sporadic time coverage that its utility for period determination is low). In each panel, the thin lines show the continua for individual strips, while the symbols and thick line show the median continuum power spectrum over the strips, and a fit to the median continuum power spectrum, respectively. The *top panel* shows the binned power spectrum, while the *lower panel* shows the result after superorbital and orbital modulation is subtracted. See Section A2.

Table A1. Bending power-law noise model (equation A1) fit to the *RXTE/ASM* binned power spectrum. The rightmost column represents the power spectrum after subtraction of the fold to the superorbital and orbital modulations, and was used to generate synthetic data sets for the estimation of superorbital period measurement uncertainty. See Appendix A2.

| Parameter | Raw | Fold-subtracted |
|--------------|-----------------------|------------------------|
| A | 6.84×10^{-3} | 5.01×10^{-3} |
| a_{lo} | 6.32×10^{-2} | -2.39×10^{-2} |
| a_{hi} | 1.56 | 2.21 |
| ν_{bend} | 4.33×10^{-2} | 1.94×10^{-2} |
| c | 1.52×10^{-3} | 1.04×10^{-3} |

impact of the underlying continuum on the detection precision is substantial – detection of a ~ 50 -day signal over a BPL continuum shows about 1.75 times the period uncertainty as it does over a white-noise continuum.

Under both noise models, the detected superorbital period is also biased, typically by $\lesssim 0.25$ d but for some strips by as much as -1.0 d

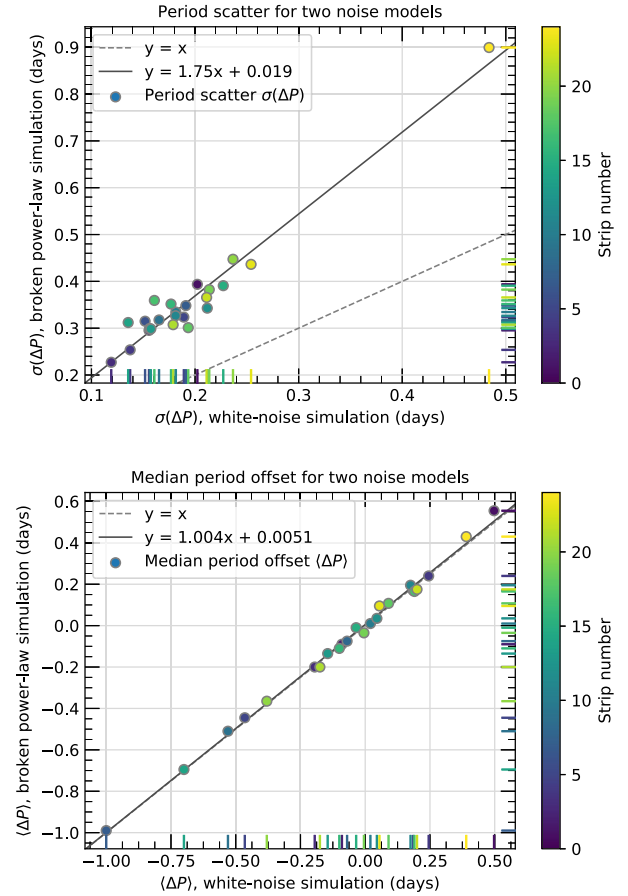


Figure A4. Comparison of the rms (*top panel*) and median (*bottom panel*) of the offset $\Delta P = P(\text{measured}) - P(\text{injected})$, for simulations employing two models of the underlying continuum: a broken power-law (vertical axes) and time-scale-independent white-noise (horizontal axes). The plot symbols are colour-coded by strip number (strip zero being the earliest). In the top panel, the line corresponding to identical rms(ΔP) is indicated with a dashed line. See Section A3.

(see the bottom panel of Fig. A4; note that this is *not* the same bias as the Gaussian-peak bias presented in Appendix A1). This bias is indistinguishable between the BPL and white-noise cases and in most cases is substantially larger than the random uncertainty. The bias probably represents the impact of the window function due to the uneven sampling (e.g. Scargle 1982) and is reproducible, allowing us to correct for it when estimating the superorbital period for each light-curve strip (Section 3.1.1).

We may also compare the ~ 0.3 -day estimate for the superorbital period uncertainty with commonly used formal estimates (e.g. Horne & Baliunas 1986). The conservative upper-limit for frequency uncertainty from a signal that is measured over total time interval T , i.e.

$$\delta f_2 = \frac{1}{2T} \quad (\text{A2})$$

translates to period uncertainty $\sigma(P) \approx 7$ d, much larger than our simulation-based estimate, as expected.

For completeness, we also consider the widely-used period uncertainty estimate of Horne & Baliunas (1986, itself derived from Kovacs 1981), which predicts the frequency uncertainty under a model of a single sinusoidal signal over Gaussian white noise. As




Figure A5. Evolution of the superorbital period measurement uncertainty as the strip-length τ is varied. Dark red points show the rms of ΔP (i.e. measured minus simulated period) for each strip with end date prior to MJD 55,600 (at which point *RXTE/ASM* coverage becomes sparse; Fig. 1). The grey-shaded areas visualize the distribution of rms(ΔP) for each value of τ , with the median and extrema indicated with black horizontal lines in each case. The black vertical notches on the horizontal axis indicate the strip-length τ in each case. The blue dashed line indicates a power-law fit to the median rms(ΔP) for $\tau > 175$ d, while the green dot-dashed line indicates a polynomial fit to $\ln(\tau)$ for all strip-lengths in the Figure. See Appendix A4.

presented in Levine et al. (2011), this prescription takes the form

$$\delta f_1 = \frac{3}{8} \frac{1}{T\sqrt{P_r}} = \frac{3\delta f_2}{4} \sqrt{\frac{\langle P \rangle}{P_{\text{raw}}}} \equiv K\delta f_2 \quad (\text{A3})$$

In the first form of δf_1 , P_r is the peak power in the power spectrum, and the power spectrum has first been normalized so that the average power is unity (Levine et al. 2011). In the second, P_{raw} is the peak period in the power spectrum before normalization, and $\langle P \rangle$ is the appropriate average power in the periodogram.

Expression (A3) was originally derived for the white-noise case (Kovacs 1981), and the presence of time-scale-dependent red-noise complicates the determination of the normalizing factor $\langle P \rangle$. For a given strip, the equivalent normalization is likely in the range $10^{-3} \lesssim \langle P \rangle \lesssim 2 \times 10^{-2}$ d (Fig. A3). With raw periodogram peaks in the range $0.1 \lesssim P_{\text{raw}} \lesssim 0.35$ d, this suggests $(0.04 \lesssim K \lesssim 0.33)$, or, finally, $(0.3 \lesssim \delta f_1 \lesssim 2.4)$ d. While the lower bound of this range is consistent with the full simulations, it corresponds to normalization by the short-timescale end of the underlying noise-continuum. Strict application of the estimate (A3) to the *RXTE/ASM* power spectrum of SMC X-1 is not well-justified and we caution against its use in red-noise cases. We prefer to quote uncertainties directly from the synthetic lightcurves that we have tailored to match the observed behaviour of SMC X-1.

A4 Dependency of measurement uncertainty on strip-length τ

Fig. A5 presents the variation of superorbital period measurement uncertainty as the strip-length τ is varied. For each strip,

all the methods of Appendices A1 and A2 were applied and the uncertainty in ΔP characterized in the same manner. There is a weak dependence on the strip-length, with the longest strip-length affording the smallest measurement uncertainty, on the assumption that the superorbital period is constant throughout the strip. For $(175 \lesssim \tau \lesssim 325)$ days, the median value of the measurement uncertainty for each value of τ seems to be well-represented by a power-law of the form $\langle \text{rms}(\Delta P) \rangle \approx 234 \tau^{-1.2}$, while shorter strips suffer from larger uncertainty than predicted by this trend. The strip-length $\tau = 230$ d adopted in the main paper (i.e. breaking the *RXTE/ASM* data set into 25 equal strips) presents a good compromise between minimizing measurement uncertainty and retaining sensitivity to variation in superorbital period.

APPENDIX B: INCREASING THE SPEED OF RED-NOISE LIGHT-CURVE SIMULATION

When conducting red-noise Monte Carlo tests (Section 3.1.1) with the *DELIGHTcurve* suite, we found large reproducible variation (by up to a factor 90) between the strips in the amount of time required to generate and analyse the synthetic lightcurves. We show here that this is most likely due to the execution time for the inverse Fast Fourier Transform (iFFT), and its dependence on the factorization properties of the length of a particular intermediate array in the Connolly (2015) implementation of the Timmer & Koenig (1995) method. Since a single time interval τ per strip is unlikely to allow convenient array-lengths for all strips when the data are unevenly sampled, we present a simple workaround to remove the strong non-uniformity in execution time, which removes a bottleneck to practical generation of sufficient numbers of synthetic lightcurves for Monte Carlo analysis.

In tests, we traced the bottleneck to the *TimmerKoenig* method within the *DELTCgen.py* module (Connolly 2015).⁷ A step within this method takes the iFFT of an array of length $L \sim 100$ times the input data length (extending the light-curve length by a large factor to account for spectral leakage). Since in general the length of this large-array will not be a power of two, we might expect the time elapsed in taking the iFFT to scale with the square of the number of elements in the longest partition of the length- L array (usually the largest prime factor of L , e.g. Press et al. 2002).

Within the module *DELTCgen.py*, the length L of the finely-sampled array is constructed from the number of input data points N and a length-parameter *RedNoiseL*, via the relation $L = (N \times \text{RedNoiseL}) + 1$. As shown in Fig. B1, for a uniform value of *RedNoiseL* for all strips, the execution time is dominated at the long end by an approximate $\propto x^2$ scaling, where x is the largest prime factor of the array-length L . This supports the hypothesis that it is the partitioning of data for the iFFT computation that dominates the total execution time.

Since the parameter *RedNoiseL* is user-controlled (with default value 100), it allows the large-array length L to be tuned for performance. We follow a simple heuristic, looping through trial values in the range $90 \leq \text{RedNoiseL} \leq 105$ and picking the value that minimizes the largest prime factor x , using the Python package *sympy*⁸ to perform the factorization.

⁷f Note that this bottleneck is likely a feature of the algorithm itself and not the *DELTCgen.py* implementation.

⁸<http://www.sympy.org/en/index.html>




Figure B1. Average elapsed time per synthetic light-curve (generation plus measurement and recovery of superorbital period) with periodicities over red-noise (Section 3.1.1), for each of 25 equal-length strips using the RXTE/ASM sample times of SMC X-1 and the parameter `RedNoiseL` fixed at its default value of 100. Durations are plotted as a function of the largest prime factor x of the length L of the finely-sampled array appropriate for each strip, and the colour scale shows the number N of data points in each strip. For $x \gtrsim 10^4$ elements, the elapsed time scales approximately as $\propto x^2$ (we find $\propto x^{2.16}$). These timings were measured on a 1.8 GHz laptop with 8GB RAM, using `DELCgen.py` as last modified 2016 December 01 (the latest version as of 2017 December). See Appendix B for discussion.

Fig. B2 shows the improvement in execution time and its uniformity when `RedNoiseL` is dynamically varied rather than set equally for all simulations. No $\propto x^2$ dependence was found, and instead the duration now depends most strongly on N , the number of points in the original data set that lie within the strip. For the most problematic strips, the execution time was reduced from ~ 500 s to $\sim 5 - 6$ s, now likely dominated by the measurement process applied to each simulation and scaling with N .




Figure B2. Average elapsed time to simulate and analyse a single synthetic lightcurve (similarly to Fig. B1), but plotted against the number N of data points in each strip and with the colour scale showing the largest prime factor of L for each strip. This run differs from that presented in Fig. B1 in that for each strip the oversampled array-length L was dynamically tuned (via the parameter `RedNoiseL`) to minimize the largest prime factor of L and thus the elapsed time. See Appendix B for discussion.

In our tests, this parameter-tuning procedure reduces the total time required to conduct our Monte Carlo simulations by about a factor 20, at a cost of slight non-uniformity in the length-multiplication factor `RedNoiseL` between the strips. If it is particularly important that `RedNoiseL` be as uniform as possible over the strips, then more sophisticated approaches might be considered, including walking out from the default value of 100 until x drops below an acceptable threshold determined from a prior run of the simulations. For the present, however, we consider the simple heuristic used here to be an acceptable compromise between uniformity and execution time.

This paper has been typeset from a \LaTeX file prepared by the author.