



Open Research Online

Citation

Barnard, R.; Trudolyubov, S.; Kolb, U. C.; Haswell, C. A.; Osborne, J. P. and Priedhorsky, W. C. (2007). On the artificial nature of aperiodic variability in XMM-Newton observations of M 31 X-ray sources and the ultraluminous X-ray source NGC 4559 ULX-7. *Astronomy & Astrophysics*, 469(3) pp. 875–879.

URL

<https://oro.open.ac.uk/10856/>

License

None Specified

Policy

This document has been downloaded from Open Research Online, The Open University's repository of research publications. This version is being made available in accordance with Open Research Online policies available from [Open Research Online \(ORO\) Policies](#)

Versions

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding

On the artificial nature of aperiodic variability in XMM-Newton observations of M 31 X-ray sources and the ultraluminous X-ray source NGC 4559 ULX-7

R. Barnard¹, S. Trudolyubov², U. C. Kolb¹, C. A. Haswell¹, J. P. Osborne³, and W. C. Priedhorsky⁴

¹ The Department of Physics and Astronomy, The Open University, Walton Hall, Milton Keynes, MK7 6BT, UK
e-mail: r.barnard@open.ac.uk

² Institute of Geophysics and Planetary Physics, University of California, Riverside, CA 92521, USA

³ The Department of Physics and Astronomy, The University of Leicester, Leicester, LE1 7RH, UK

⁴ Los Alamos National Laboratory, Los Alamos, NM 87545, USA

Received 1 September 2006 / Accepted 5 March 2007

ABSTRACT

Context. Power density spectra (PDS) that are characteristic of low mass X-ray binaries (LMXBs) have been previously reported for M 31 X-ray sources, observed by XMM-Newton. However, we have recently discovered that these PDS result from the improper addition/subtraction of non-simultaneous lightcurves.

Aims. To understand the properties and origins of the artefact.

Methods. We re-analysed our XMM-Newton observations of M 31 with non-simultaneous and simultaneous lightcurves, then combined simulated lightcurves at various intensities with various offsets and found that the artefact is more dependent on the offset than the intensity.

Results. The lightcurves produced by the XMM-Newton Science Analysis Software (SAS) are non-synchronised by default. This affects not only the combination of lightcurves from the three EPIC detectors (MOS1, MOS2 and pn), but also background subtraction in the same CCD. It is therefore imperative that all SAS-generated lightcurves are synchronised by time filtering, even if the whole observation is to be used. We also find that the reported timing behaviour for NGC 4559 ULX-7 was also contaminated by the artefact; there is no significant variability in the correctly-combined lightcurves of NGC 4559 ULX-7. Hence, the classification of this source as an intermediate-mass black hole is no longer justified.

Conclusions. While previous timing results from M 31 have been proven wrong, and also the broken power law PDS in NGC 4559 ULX-7, XMM-Newton was able to detect aperiodic variability in just 3 ks of observations of NGC 5408 ULX1. Hence XMM-Newton remains a viable tool for analysing variability in extra-galactic X-ray sources.

Key words. X-rays: binaries – galaxies: individual: M 31 – galaxies: individual: NGC 4559 – galaxies: individual: NGC 5408 – X-rays: general

1. Introduction

The variability and spectral properties of Galactic low mass X-ray binaries (LMXBs) are well known to depend more on the accretion rate than on whether the primary is a neutron star or black hole (van der Klis 1994). At low accretion rates, the power density spectra (PDS) may be characterised by broken power laws, with the spectral index, γ , changing from ~ 0 to ~ 1 at some break frequency in the range 0.01–1 Hz (van der Klis 1994); such variability has a rms power of ~ 10 –40% (e.g. van der Klis 1995). We describe such PDS as type A (Barnard et al. 2004). At higher accretion rates, the PDS is described by a simple power law with $\gamma \sim 1$ and rms variability $< 10\%$ (van der Klis 1994, 1995). We refer to these PDS as type B (Barnard et al. 2004).

Type A variability is characteristic of disc-accreting X-ray binaries; any X-ray source that exhibits such variability may be identified as an X-ray binary, rather than a foreground star or background active galaxy (see e.g. Barnard et al. 2003b). Such variability has been reported in XMM-Newton observations of M 31 by some of the authors (see e.g. Barnard et al. 2003b, 2004; Williams et al. 2005). Type A PDS have also been reported

for ultraluminous X-ray sources (ULXs) in NGC 4559 (Cropper et al. 2004) and NGC 5408 (Soria et al. 2004).

However, we have now discovered an artefact that produces false type A PDS, caused by the improper addition of lightcurves from the three EPIC instruments on board XMM-Newton (MOS1, MOS2 and pn), where the lightcurves are not synchronised (Barnard et al., this issue). The data were analysed using SAS 6.0.0 and FTOOLS 5.3.1. In the Sect. 2 we discuss how the error occurred. We then describe a new, corrected analysis of the M 31 data in Sect. 3. Section 4 describes our investigation into the artefact using simulated lightcurves. Then, Sect. 5 details the causes of the artefact. Section 6 re-analyses XMM-Newton observations of two ultraluminous X-ray sources with published type A PDS. We draw our conclusions in Sect. 7.

2. Manipulating lightcurves generated by SAS

XMM-Newton has observed M 31 extensively, and Barnard et al. (2003b,a, 2004, 2006) conducted a survey of bright sources in the central region, over four observations. For each observation of every source, 0.3–10 keV lightcurves were obtained from




Fig. 1. Power density spectra of 0.3–10 keV lightcurves of r3-60, using Method I (*left*) and Method II (*right*). We present best fits to the Method I PDS of power law and broken power law models; a broken power law is clearly required. The PDS is hence reminiscent of type A variability in Galactic LMXBs. However, the Method II PDS is flat, revealing that this variability is artificial.

source and background regions in the MOS1, MOS2 and pn images, using 2.6 s binning.

We have recently discovered that the results of manipulating lightcurves produced with the SAS tool `evselect` using the `lcmath` FTOOL produces different results, depending on how the time selection in `evselect` is phrased. Originally, lightcurves were produced from events files that were filtered in `evselect` with the expression “`(PI in [300:10000]) && (TIME in [tstart:tstop])`”; additional filtering was instrument dependent: MOS lightcurves were additionally filtered using “`#XMMEA_EM&&(PATTERN<=12)`”, while the pn equivalent was “`#XMMEA_EP&&(PATTERN<=4)`”. We shall refer to this filtering as Method I.

However, different lightcurves are obtained when `evselect` is used in either of the following ways. Firstly, one may assign `tstart` and `tstop` to the `TLMIN1` and `TLMAX1` keywords in the pn, MOS1 and MOS2 events files, using FTOOLS such as `fmodhead` or `fparkey`; we refer to this method as Method II. Method III involves filtering each lightcurve with the `evselect` tool, using the parameters “`timemin = tstart timemax = tstop`”. Methods II and III are equivalent. Vitaly, they are not equivalent to Method I, although this is not mentioned in any documentation. Consequently, we analysed lightcurves obtained with each method to determine which, if any, were correct.

3. Re-analysing the data

3.1. Combined EPIC lightcurves

We obtained MOS1, MOS2 and pn lightcurves from the source extraction regions of several sources in the M 31 observations, using Methods I and II. For Method I, we found that uncertainties of combined EPIC (MOS1+MOS2+pn) lightcurves were underestimated by a factor of ~ 10 –40%. This ratio varied within each lightcurve, with the standard deviation in the ratio for a given lightcurve decreasing with increasing luminosity. The reason for the small uncertainties is discussed in Sect. 5.1.

Closer inspection of the lightcurves has revealed the difference between Method I and Method II lightcurves: the Method II lightcurves are synchronised, while the Method I lightcurves are not. Of course, for the Method II or III lightcurves to be synchronised, `tstart` must be the same for all detectors, and must be present in all the events lists.

Many PDS from lightcurves combined using Method I exhibit type A variability, while identical data combined with Method II do not. To investigate the cause of the type A variability in the observed PDS, we examined the Method I and II lightcurves of XMMU J004208.9+412048 (r3-60 Kong et al. 2002). The Method I and Method II PDS are shown in the left and right panels respectively of Fig. 1; they were obtained using the FTOOL `powspec`. The Method I PDS is well modelled ($\chi^2/\text{d.o.f.} = 6/8$) by a broken power law where the spectral index changes from 0 to 0.33 ± 0.05 at a frequency of 29 ± 4 mHz. Such a PDS is characteristic of Galactic LMXBs. However, the PDS of the Method II lightcurve is flat, with a Leahy power of 2, consistent with Poisson noise. Hence the variability in the Method I PDS is artificial.

3.2. Background subtraction

Background subtraction requires two lightcurves: source and background. We produced background-subtracted pn lightcurves of various X-ray sources from the 2002 January XMM-Newton observation of the M 31 core, to test whether the lightcurves from the same detector were synchronised. We present the Method I PDS from source, background and background subtracted pn lightcurves of r3-125 in Fig. 2, panels (a), (b) and (c) respectively. The source and background PDS are flat, while the background-subtracted PDS is of type A. Hence, background subtraction causes artefacts when non-synchronised source and background lightcurves are used.

4. Investigating the artificial variability with simulated lightcurves

To determine the cause of artificial variability we combined artificial lightcurves with varying intensities and offsets. We defined 100 evenly spaced intensity intervals from 0.01 to 1.00 count s^{-1} , and for each interval generated a single “source” lightcurve and 100 “background” lightcurves (using the terminology of `lcmath`: the “background” lightcurve may be added to or subtracted from the “source” lightcurve). This intensity range reproduces well the observed intensities from the bright X-ray sources in the M 31 Core. Each lightcurve had 60 ks lightcurve duration, with 2.6 s binning.




Fig. 2. PDS from source+background **a)**, background **b)** and background-subtracted **c)** pn Method I lightcurves of r3-125. The background-subtracted PDS is of type A, although the component PDS are flat.

For each lightcurve, we generated a series of photons with random arrival times, using the Park & Miller Minimal Standard random number generator, shuffled by the Bays & Durum method (Knuth 1981); the random number seed for the n th “background” lightcurve at intensity i was defined as $-(10\,000i + n)$. The source lightcurves were generated with seeds of $-10\,000i$; hence, each lightcurve was reproducible and unique. We then conducted a series of tests for artificial power in the PDS of combined lightcurves.

4.1. Adding two lightcurves with different start times

Firstly we added each “background” lightcurve for every intensity interval to the corresponding “source” lightcurve at various offsets with `lcmath`; the offset was defined as the delay in seconds as a fraction of the 2.6 s time binning, and ranged from 0.05 to 0.95 at intervals of 0.05. We then produced geometrically grouped, 64-bin PDS of the resulting lightcurve with the `FTOOL` `powspec`, averaged over ~ 180 time intervals. Initially, we examined the PDS for excess variability by fitting them with a constant Leahy power of 2; for our PDS, a $\chi^2 \gtrsim 25 / 10$ d.o.f. signified a 3σ detection of power above the Poisson noise. We find that the offset is more important than intensity, with all combinations of lightcurves with offsets 0.25–0.80 producing significant excess variability; 12% of lightcurves with an offset of 0.05 or 0.95 produced artificially variable PDS. However, not all variable PDS are considered type A.

To be classed as type A, a power law fit to the PDS must be rejected, and a broken power law fit must be reasonable. A broken power law fit may be described by spectral index α that changes to spectral index β at frequencies higher than some break frequency ν_b ; the final parameter is the power at ν_b . When fitting the PDS, we rejected power law fits with good fit probability < 0.005 ($\chi^2 > 24$) and accepted broken power law fits with good fit probability > 0.05 ($\chi^2 < 16$); for the broken power law fit, we set $\alpha = 0$, as observed in Galactic LMXBs. Less than 10% of lightcurves with offsets in the ranges 0.05–0.25 and 0.75–0.95 exhibited type A variability, while 30% of lightcurves with $f = 0.50$ were of type A.

Table 1. Summary of variability observed when three artificial lightcurves are combined, for Cases A and B. For each case, the fraction of combined lightcurves exhibiting type A were shown, for offsets $f = 0.05$, $f = 0.50$ and $f = 0.95$.

Case	$f = 0.05$	$f = 0.50$	$f = 0.95$
A	0.003	0.13	0.011
B	0.12	0.75	0.12

4.2. Adding three lightcurves with different start times

Since Barnard et al. combined source lightcurves from the three EPIC detectors in the original M 31 survey work (e.g. Barnard et al. 2003b,a, 2004) we added a third lightcurve to the combined artificial lightcurves discussed in the previous section. For Case A, this lightcurve was the 50th background lightcurve with an offset of 0.1, and for Case B the 50th background lightcurve with an offset of 0.5 was added. The resulting lightcurves had intensities of 0.03–3 count s^{-1} , with offsets of 0.05–0.95 as before.

We tested the resulting PDS for type A variability. The results are summarised in Table 1. Again, we found that the fraction of type A lightcurves dependent on offset rather than intensity, with the maximum at $f = 0.50$. However, even an offset of 0.05 produced artificial type A variability in 12% of combinations for Case B. Combining three non-synchronous lightcurves (in Case B) results in a factor of > 2 more false type A variability than just combining two lightcurves. Hence the artefact is strengthened as more non-synchronous lightcurves are added (combining detectors) or subtracted (background subtraction).

5. The causes of the artefact

Three factors contribute to the artificial type A variability: underestimated uncertainties, excess power in the PDS, and suppression of high-frequency variability. We now discuss these in turn.

5.1. Underestimating Method I uncertainties

The smaller uncertainties of Method I are due to the workings of `lcmath`. Whether adding or subtracting lightcurves, `lcmath` treats the first lightcurve as the source and the second lightcurve as the background; it calculates the mean “background” intensity for each time interval in the “source” lightcurve. If the lightcurves are offset by some non-integral number of time bins, then the intensities of two “background” intervals are weighted by their overlap with the “source” interval. The intensity of the n th interval of the summed lightcurve, $S(n)$, is then given by

$$S(n) = A(n) + f \times B(n-1) + (1-f) \times B(n) \quad (1)$$

where A is the “source” lightcurve, B is the “background” lightcurve and f is the fractional offset. The latter is given by

$$f = \frac{t_0^B - t_0^A}{\Delta T} \quad (2)$$

where t_0^A and t_0^B are the start times for the source and background lightcurves, respectively, and ΔT is the width of the time interval.

By default, the lightcurves begin at the arrival time for their first photons. In practise, however, t_0^A is determined by the arrival of the first photon in the “source” lightcurve but `lcmath` sets t_0^B to the start of the first “background” interval to overlap the “source”

lightcurve. In addition, `lcmath` ignores all source intervals with no “background” interval. Hence, $0 \leq f \leq 1$. For synchronised lightcurves, t_0^A and t_0^B are both set to the same user-defined value. Hence $f = 0$ and $S(n) = A(n) + B(n)$, as it should be.

The underestimated uncertainties are caused by an error in the `lcmath` code `lcmathrdbkg.f`. When two or more background bins coincide with a source interval, with count rate C_i and fractional overlap O_i , then `lcmathrdbkg.f` assigns an uncertainty of $O_i \times C_i^{0.5}$, rather than the Gaussian uncertainty $(O_i \times C_i)^{0.5}$ if no uncertainty is provided. Similarly, when an uncertainty σ_i is provided, then `lcmath` calculates the total uncertainty as $[(O_1 \times \sigma_1)^2 + (O_2 \times \sigma_2)^2 + \dots]^{0.5}$, when it should be $(O_1 \times \sigma_1^2 + O_2 \times \sigma_2^2 + \dots)^{0.5}$. Since $O_i < 1$, $O_i/O_i^{0.5} < 1$, and the uncertainties are underestimated.

The underestimated Method I uncertainties are also responsible for the excess power in the PDS. This is because the Leahy PDS produced by `powspec` are normalised such that the white noise level expected from the data errors corresponds to a power of 2 (see `powspec` documentation). To do this, Gaussian errors are assumed and the normalisation, A_{Leahy} , must be

$$A_{\text{Leahy}} = 2 \frac{\Delta t}{\sigma^2 N} \quad (3)$$

where Δt is the sampling time, N is the number of points and $\overline{\sigma^2}$ is the mean squared error (see e.g. Appendix A of Vaughan et al. 2003). I.e., $A_{\text{Leahy}} \propto 1/\overline{\sigma^2}$. Therefore, underestimating the Method I uncertainties by a factor of 1.1–1.4 will overestimate the PDS power by a factor of ~ 1.2 –2.

5.2. Suppression of high-frequency power

After correcting for the power excess described above, we find that the Method I PDS actually dips below the Poisson noise level at high frequencies. I.e., the high frequency variability is suppressed, causing the break in the PDS at frequency ν_b .

For those M 31 sources identified with broken power law PDS, $\nu_b = 20 \pm 4$ – 57.8 ± 0.5 mHz, independent of the luminosity. We have studied the break frequencies of the 190 000 lightcurves that were constructed from three non-simultaneous, randomly-generated lightcurves and discussed in Sect. 4.2. We found ν_b to vary over ~ 3 –70 mHz, with no strong correlation between ν_b and offset or intensity. While the break frequency is necessarily related to the bin size, this relation is very weak also because ν_b varies by a factor of ~ 30 while the bin size is constant.

6. Ultra-luminous X-ray sources in NGC 4559 and NGC 5408

We conducted a brief literature search for published PDS of combined EPIC lightcurves that may also be affected by this data-analysis artefact. We found a published broken power law PDS for NGC 4559 ULX7 (Cropper et al. 2004) and also for a ULX in NGC 5408 (Soria et al. 2004). We re-analysed XMM-Newton observations of these sources, first following the methods described in these papers, using non-synchronised lightcurves, then comparing their PDS with those from synchronised lightcurves. However, it is beyond the scope of this paper to re-analyse all published XMM-Newton variability data!

Cropper et al. (2004) reported a broken power law PDS from an ultra-luminous X-ray source (ULX) in NGC 4559 with




Fig. 3. Method II PDS of NGC 4559 ULX-7. The power is normalised to give fractional rms² variability. The expected noise is subtracted, following Cropper et al. (2004). This PDS is consistent with zero power ($\chi^2/\text{d.o.f.} = 20/16$, with a null hypothesis probability of 0.22). Hence there is no evidence for the variability reported by Cropper et al. (2004).

a UV/X-ray luminosity exceeding 2.1×10^{40} erg s⁻¹. We extracted MOS1, MOS2 and pn lightcurves from a source region with 30'' radius using Method II. We constructed a PDS that was normalised to give fractional rms² variability, and the expected noise was subtracted. The expected noise had a power of 4.3. The frequency bins were geometrically grouped. The resulting PDS is shown in Fig. 3. This PDS is consistent with zero power (good fit probability = 0.22); hence, the variability reported by Cropper et al. (2004) is artificial.

We note that Feng & Kaaret (2005) also re-analysed the XMM-Newton observation of NGC 4559 ULX-7, and were able to reproduce the variability reported by Cropper et al. (2004) only when the intervals of background flaring were included; hence, they mistakenly attributed the variability to the background flares, unaware that Cropper et al. (2004) also excluded the background flaring from their analysis. Feng & Kaaret (2005) produced their PDS directly from the events lists that describe each photon (Kaaret 2006, private communication). Hence they did not observe the artefact after the intervals of high background were excluded.

Soria et al. (2004) studied a series of XMM-Newton observations of a ULX in the dwarf galaxy NGC 5408 over an 18 month period between 2001, July and 2003, January. Its 0.2–12 keV spectrum requires two components, a power law with photon index $\Gamma = 2.6$ –2.9 and a thermal component with blackbody temperature of 0.12–0.14 keV; such a spectrum is typical of a black hole X-ray binary in the very high/steep power law state (McClintock & Remillard 2004). The inferred 0.2–12 keV luminosity is $\sim 10^{40}$ erg s⁻¹ (Soria et al. 2004). They present a broken power law PDS from ~ 3 ks of the 2003, January XMM-Newton observation of NGC 5408, where γ increases from ~ 0 to 1.3 ± 0.2 at 2.5 mHz. Such a PDS is also consistent with a black hole X-ray binary in the very high state.

We extracted EPIC lightcurves from a circular source region with 30'', following Soria et al. (2004). A Method II PDS is presented in Fig. 4; the PDS is rms² normalised, and averaged over 3 intervals of 256 bins, with 5 s time bins. Fitting this PDS with zero power results in a $\chi^2/\text{d.o.f.}$ of 92/14, hence the variability is real. Figure 4 also shows the best power law




Fig. 4. Method II PDS for the NGC 5408 ULX, averaged over 3 intervals of 256 bins with 0.1 s resolution. The PDS is normalised to give the fractional rms² variability, and the white noise is subtracted. Power law and broken power law fits are shown, with best fit χ^2 /d.o.f. of 77/12 and 24/11 respectively.

and broken power law fits to the PDS, with χ^2 /d.o.f. = 77/12 and 24/11 respectively. The broken power law fit is clearly preferred (f-testing yields a 99.95% chance that the improvement is significant), but we caution that it is not an acceptable fit: the goodness of fit probability is just 1.2%.

7. Conclusions

Some of the authors have previously reported variability in XMM-Newton observations of M 31 X-ray sources that are characterised by broken power law PDS, which are signatures of disc-accreting X-ray binaries (Barnard et al. 2003b, 2004; Williams et al. 2005). However, this “type A” variability observed in XMM-Newton observations of M 31 is an artefact and may be entirely attributed to errors introduced by lcmath when combining non-synchronised lightcurves (Barnard et al. 2007). However, extragalactic analogues of the Galactic X-ray binaries should exhibit type A and type B behaviour; type A variability should be observed in sufficiently bright X-ray sources.

In the survey of XMM-Newton observations of X-ray sources in M 31, we found that the type A variability was

observed at lower luminosities than type B variability, as expected for LMXBs. However, the most likely explanation for this observed behaviour is that the offset was smaller for the brighter sources; the start time for an XMM-Newton lightcurve is defined by the arrival time of the first event, so brighter sources will naturally have smaller offsets between CCDs.

The published broken power law PDS of NGC 4559 ULX7 is also artificial. The PDS of the NGC 5408 ULX shows genuine intrinsic variability. Timing analysis of XMM-Newton observations of extra-galactic X-ray sources is indeed viable. M 31 X-ray sources have exhibited pulsations (Osborne et al. 2001; Trudolyubov et al. 2005); bursts (Pietsch & Haberl 2005); periodic dipping due to photoelectric absorption (Trudolyubov et al. 2002; Mangano et al. 2004), from a precessing disc in one case (Barnard et al. 2006), finally branch movement in a trimodal colour-intensity diagram reminiscent of Galactic Z-sources (Barnard et al. 2003a).

Acknowledgements. Power density spectra were fitted using fitpowspec, provided by P.J. Humphrey. R.B is funded by PPARC, while S.T. acknowledges support from NASA grant NAG5-12390. We thank the anonymous referee for their hard work and constructive comments.

References

- Barnard, R., Kolb, U., & Osborne, J. P. 2003a, *A&A*, 411, 553
- Barnard, R., Osborne, J. P., Kolb, U., & Borozdin, K. N. 2003b, *A&A*, 405, 505
- Barnard, R., Kolb, U., & Osborne, J. P. 2004, *A&A*, 423, 147
- Barnard, R., Foulkes, S. B., Haswell, C. A., et al. 2006, *MNRAS*, 366, 287
- Barnard, R., Kolb, U. C., & Osborne, J. P. 2007, *A&A*, 469, 873
- Cropper, M., Soria, R., Mushotzky, R. F., et al. 2004, *MNRAS*, 349, 39
- Feng, H., & Kaaret, P. 2005, *ApJ*, 633, 1052
- Gehrels, N. 1986, *ApJ*, 303, 336
- Kong, A. H. K., Garcia, M. R., Primini, F. A., et al. 2002, *ApJ*, 577, 738
- Mangano, V., Israel, G. L., & Stella, L. 2004, *A&A*, 419, 1045
- McClintock, J. E., & Remillard, R. A. 2004, *Compact Stellar X-ray Sources* (Cambridge University Press), in press [arXiv:astro-ph/0306213]
- Osborne, J. P., Borozdin, K. N., Trudolyubov, S. P., et al. 2001, *A&A*, 378, 800
- Pietsch, W., & Haberl, F. 2005, *A&A*, 430, L45
- Soria, R., Motch, C., Read, A. M., & Stevens, I. R. 2004, *A&A*, 423, 955
- Trudolyubov, S., Kotov, O., Priedhorsky, W., Cordova, F., & Mason, K. 2005, *ApJ*, 634, 314
- Trudolyubov, S. P., Borozdin, K. N., Priedhorsky, W. C., et al. 2002, *ApJ*, 581, L27
- van der Klis, M. 1994, *ApJS*, 92, 511
- van der Klis, M. 1995, *X-ray Binaries* (Cambridge University Press), 256
- Vaughan, S., Edelson, R., Warwick, R. S., & Uttley, P. 2003, *MNRAS*, 345, 1271
- Williams, B. F., Barnard, R., Garcia, M. R., et al. 2005, *ApJ*, 634, 365