Correcting for X-ray Pileup in the EPIC pn-CCD camera using simulated Detector Response Matrices

Kjell Koch-Mehrin

SRE-S, ESAC, Madrid, Spain

October 2016

Abstract

The fraction of piled-up data in spectra as a function of source brightness for EPIC-pn observations is investigated, along with what effect this has on the source flux and spectral parameters. This is done for a large range of source count rates in order to quantify where the effect of pileup is still manageable and when it is most severe. The pileup correction method from the Science Analysis System, SAS, is updated with changes to the methodology and tested on simulated data using an IDL version of the correction code.

Contents

1 Introduction

The pileup correction method has been part of the SAS[1] since version 13.0.0, however is still considered experimental. The aim of this work is to evaluate how well the method works at recovering flux and the spectral parameters of a piled-up observation, and where its limits are. The results presented here are after much testing and some changes to the methodology. The method is therefore not analogous to the version in the latest SAS release (15.0.0)[2]. Furthermore, simulations are used to characterise the effect pileup has on a spectrum as a function of source event rate for the EPIC-pn camera on-board the XMM-Newton spacecraft, in the different observation modes.

Pileup occurs when two or more X-ray photons strike the same CCD pixel or adjacent pixels during the integration time of one read-out cycle (i.e. frame time). This can cause a number of problems. The on-board electronics register a single energy to one photon when it is in fact a sum of charges deposited by individual photons that have piled-up. This is reflected in the spectral data with counts migrating to higher energies; ultimately distorting a spectrum from its *true* shape to a harder one. X-ray photons are known to produce certain pixel patterns. Which patterns these are and the likelihood of them occurring, depend on the energy of the photon and the characteristics of the CCD (mainly the pixel size). Pileup changes these pattern fractions from the expected values, this affect is referred to as pattern migration. This can be, for example, two adjacent single pattern photons, forming a double pattern. Alternatively, two or more x-ray photons may produce a bad pattern (a pattern not known to be produced by x-rays), which are removed for spectral analysis. The two migration effects suppress the true event rate of the source and lead to flux loss. The type of pileup can be categorized into two groups:

1. Energy pileup: During a single read-out cycle the charge from more than one photon is deposited in the same pixel. The on-board electronics cannot distinguish if one, or multiple photons, are responsible for the charge found in the pixel. Therefore, the energy deposited in the pixel is attributed to a single event.

2. Pattern pileup: During a single read-out cycle charge deposited in two or more adjacent pixels is due to individual photons instead of a single multi-pixel event. Event reconstruction treats multi-pixel patterns as one event. Therefore, the separate events in adjacent pixels are erroneously combined and read-out as one event of higher pattern type and energy.

The severity of the affects due to pileup described above depend on the event rate of the source and the timing resolution of the observation mode being employed. When serve enough, the consequences of pileup prevent the characterisation of the true nature of the source. Both Strüder[1] and Jewtha[2] have given count rate limits for the brightness of a point source above which pileup will occur, for the different observation modes. Currently, the most common method used to counter pileup, for observations above the count rate limit, is to excise the core of the PSF of the source (where most of the photons fall and as such most of the pileup is located). Although this technique works well at reducing the amount of pileup in the data, many photons, valuable for accuracy and statistics, are lost. The pileup correction method is an alternative approach to deal with piled-up observations.

[1] An optional feature of the 'rmfgen' task

[2] The pileup correction method in the developmental SAS-build 11 has been updated with the changes to the version used here and will be included in SAS release 16.0.0

2 Methodology

I begin by defining the difference when referring to an event or a count. An event will be the charge deposited onto the CCD detector by a single photon of good pattern (pattern type <= 4 and known to be produced by X-rays). Therefore, when talking about the event rate, it is always in the absence of pileup (it is the true count rate of a source). The event spectrum *is* therefore the pileup free spectrum and determining its true characteristics from piled-up data is the ultimate goal of the pileup correction method, explained in section 2.2. A count is then the charge of a good pattern as actually read-out by the CCD, which, may be equivalent to that of the event or correspond to the sum of multiple events that have piled-up. The count rate is the actual measurable quantity from the output spectrum (i.e. the spectrum from the piled-up data). If there is no pileup in an observation, the count rate and event rate are equal.

2.1 PN eventlist simulator

The simulator is written in IDL and produces the eventlist for a point source. The simulator is given an input model from which it will simulate the data of a point source corresponding to that model. Additionally, the required program inputs are the number of expected events arriving at the detector per observation frame and the number of such frames that will be observed for (along with some calibration and instrument files)[3]. A Poisson distribution around the user supplied expected event rate is taken to obtain the exact number of events that are simulated for each frame. An energy is given to each event in every frame. The allocated energies are based on the probability of having that energy in the provided spectral model after convolution with the instrument's effective area. The charge(s) for the energy given to each event are simulated by randomizing over the detector response, the point spread function and the pattern distribution files. Each energy has now been reconstructed into a charge, or charges (depending on the pattern type), with a position on the detector. This information comprises the eventlist. By neglecting the positions of each event, there can be no pileup. I use this simplification to produce the event spectrum. The output spectrum, is produced by reconstructing the events in each frame to counts. This is achieved by merging the charges, which occupy the same position on the detector and/or form higher pattern types, into a single count. Bad patterns, if present, are removed. The desired region, from which the output spectrum is created, is spatially selected and the infinite-extent effective area is corrected according to this region.

All simulations are run for a minimum of 100,000 frames. Low event rates are typically run for many more (between 1 and 5 million frames). Essentially, the target is to have at least 100,000 counts in the spectrum at each event rate to obtain reasonable statistics. Albeit, this is not always possible for the very lowest event rates. Having a healthy number of frames per observation becomes very important when using the correction method. For this reason, even at very high event rates the minimum of 100,000 frames is enforced.

2.2 PN pileup correction method

The pileup correction method is run using IDL code, a testing platform to the official version in the SAS. Using the original instrument response file, the correction method simulates a new detector response matrix on the event list of an observation. Such, the corrected response can incorporate the presence of pileup in the observation.

The inputs required are the eventlist (containing all observed charges) and the same calibration files as mentioned for the PN simulator. The detector response matrix is rebuilt at each energy from, and only from, simulated test events that are added to the eventlist. Starting from the first input energy of the original PN detector response, for each frame in the eventlist which contains a charge, a **single** simulated event of known position, pattern and charge is

[3] these include the instrument response matrix, effective area, point spread function and pattern distribution

obtained by randomizing over the calibration files as above, and added to the frame. The simulated events are labelled with event "IDs" to distinguish them from the observed charges. The charges in each frame are reconstructed so that only good pattern counts remain. During this step, the simulated events are merged with the observed events which may cause them to pileup, changing their energy and/or pattern. The simulated events which do not form bad patterns during this process are extracted, assigned to the appropriate PI channel based on their charge and used to build the corrected response at the current input energy. Furthermore, the corrected response, at the given input energy, is normalised by the total number of simulated events added before the reconstruction process. The entire procedure described up to now is repeated for each energy in the detector response matrix. This is the foundation of how the correction method works. There are additional details and processes which depend, in part, on the user's needs. Most importantly, simulated events should only be given a position on the detector that is within the source extraction region used to make the spectrum (this must be specified by the user). The user must also indicate which patterns of the simulated events are to be included to build the corrected response. This should agree with the patterns of the counts in the spectrum for which the response will be used.

When adding a simulated event to each frame, the event rate of the observation is intrinsically increased. This must be taken into consideration. This part of the method is subject to some discussion and various approaches were looked into. The technique which appears to work best and therefore pursued in this work is removing one count (i.e. after reconstruction of the charges from the events, into counts) from each frame of the observation before adding the simulated event. This is a good approximation at low event rates where a single count is likely to also be a single event, and not multiple events that have piled-up. The original count rate of the observation is therefore conserved. The results of this method are presented and discussed in section 4.

2.3 Fitting the spectra

Spectral fitting is achieved using Xspec 2.9.0 with the use of PyXspec; an object orientated Python interface to the Xspec fitting program allowing the user to run Xspec commands from Python scripts.

Both the event and output spectrum are fit with the input model. The model is folded with the appropriate effective area and detector response. The event spectrum includes events across the entire detector of pattern types 0-12. The output data is fit, using the original response matrix, to the singles only and singles and doubles spectrum (pattern 0 and 0-4 respectively). When testing the correction method, the same fit is repeated, replacing the original detector response with the corrected one. A circular source extraction region of a 60 arc-second radius is used in all cases for the output data.

The spectra are all treated as Poisson distributed data and are therefore left unbinned. This is a good approximation since the only source of experimental noise is due to the number of events arriving at the detector and there is no background. The C-statistic is used to find the maximum likelihood during parameter estimation of the model with the Levenberg-Marquardt technique. To check whether the spectrum belongs to the model and its best-fit parameters, pearson's chi-squared test-statistic is employed. This test statistic is defined by:

$$x_p^2 = \sum_{i=1}^{N} \frac{(y_i - m_i)^2}{m_i}$$

where N is the total number of bins in the spectrum, y the observed data rate, and m the value of the predicted data rate based on the model.

3 A closer look at the pileup problem

3.1 Pileup fraction and event loss as a function of event rate

A hypothetical source yielding a delta function with a Gaussian width of 1 eV is used throughout. Simulations are carried out at four levels of hardness of the source; 1, 3, 5 and 8 keV. Simulating a delta function gives the distinct advantage of making counts due to exactly one photon and counts as a consequence of pileup very distinguishable when inspecting the spectrum (see figure 1). From this, the pileup fraction is determined within reasonable error.

$$\text{pileup} = \frac{\text{single and double pattern counts at wrong energies}}{\textit{total}\text{ number of single and double pattern counts}} \tag{1}$$

The number of photon events lost due to energy/pattern pileup, the removal of bad patterns and triple and quadruple pixel events, is given by:

$$\text{event-loss} = \frac{\text{number of single and double pattern counts in the output spectrum}}{\text{total number of events in the event spectrum}} \tag{2}$$

Figure 1 shows the event spectrum of the delta function at 1 keV and the corresponding output spectrum. At 4 counts per frame, we are observing a very bright source (7.5087e-11 ergs cm^{-2} s^{-1} if in the full-frame mode). For this reason a significant amount of pileup can be seen. The peaks succeeding the first one are purely due to pileup. The number of counts at 1 kev is less in the output spectrum due to multiple single photon events piling up and migrating to higher energies. This effect distorts the spectral shape. The difference in this peak is also due to the increased amounts of bad patterns which form when observing such a bright source.

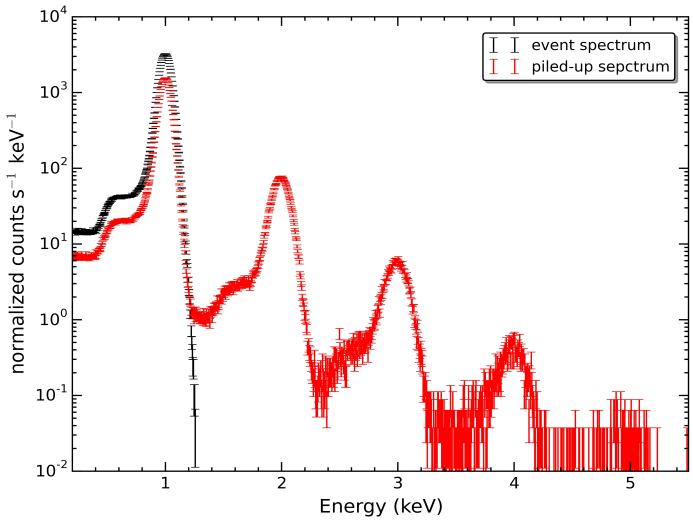

Fig. 1: Event spectrum compared with the piled-up one for the 1 keV spectral model at 4 counts/frame.

How the fraction of piled-up counts in a spectrum vary with the event rate is shown in figure 2. The event rate is quoted per frame; the events per second corresponding to these values depend on the time resolution of the observation mode. This information is given in figure 5. From figure 2 we see that doubles are approximately a factor 12 times more likely to be piled-up than singles. In all cases, the pileup fraction increases until 6 events per frame where it reaches a maximum. The pileup fraction then declines at a slightly slower rate to approximately 62% of it's peak value before increasing again from 22 events per frame upwards. The rate of the second increase is much slower than the first one at a factor of 70 times smaller. The behaviour of these curves as the event rate increases can be explained by looking at the PSF at different event rates. From figure 3 we see that the majority of the counts at 6 events/frame

are located at the center of the source. The PSFs corresponding to event rates greater than this, show the number of counts in the core fall off, which is due to an increasing amount of bad patterns. At 40 events/frame all counts in the core within a radius of 12 arcseconds are completely removed. The majority of the counts removed this way are also piled-up ones. The total pileup fraction therefore initially decreases, until the pileup contribution from the region around the core becomes greater than that from inside the core. The broken lines in figure 2 show the pileup fraction in an annulus of 15-60 arcseconds. From 50 events/frame upward these converge to the total pileup fraction, indicating that no piled-up counts are left in the core (0-15") of the PSF.

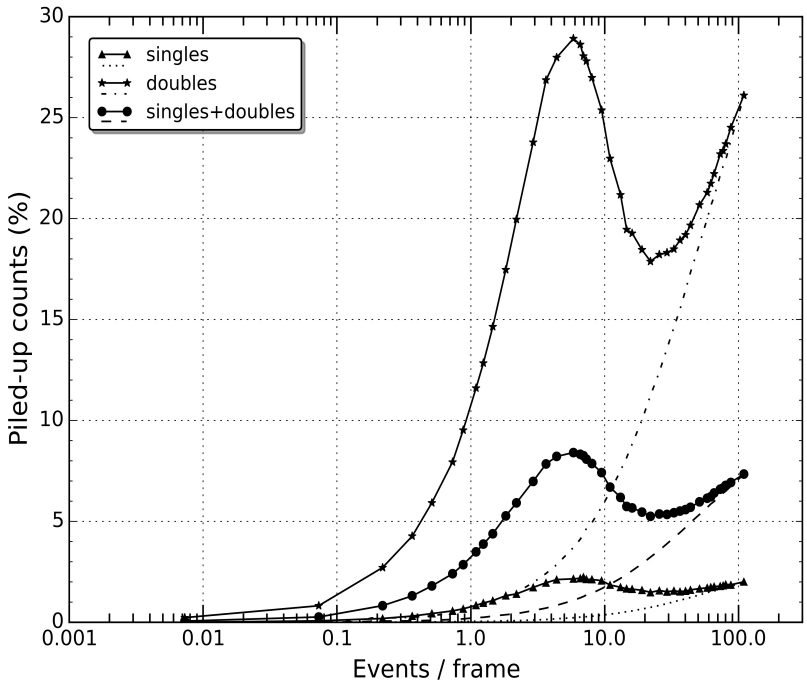

Fig. 2: Pileup fractions in spectra containing singles, doubles and both. The values are obtained from the delta source at 1 keV. The broken lines give the fraction in an annulus of 15-60 arcseconds. Each data point corresponds to a simulation run of one million frames.

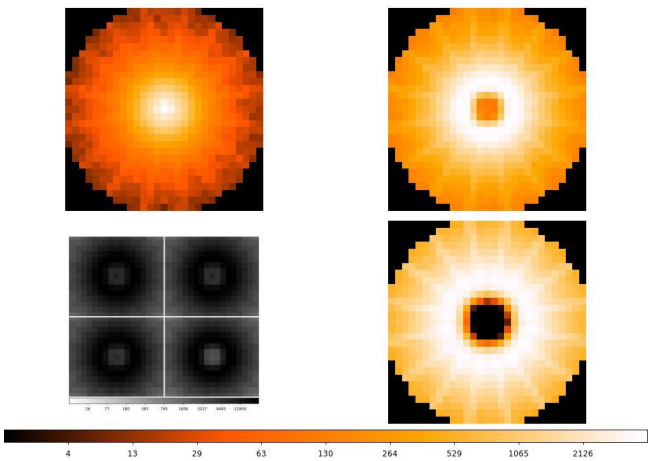

Fig. 3: PSFs from the 1 keV delta source at 6, 22 and 40 events per frame in clockwise direction. The region of the PSF with a radius of 60 arcseconds from the center is shown, including single and double pattern counts. The embedded image (bottom left) shows the PSFs for the different hardness levels all at 14 events/frame: 1, 3 keV from left to right at the top and likewise 5, 8 keV at the bottom.

As the pileup increases and more and more events either form bad patterns or erroneously combine into piled-up counts, the event rate of a source is lost. Figure 4 shows the relationship

between the event rate and the count rate for the different observation modes. The small window mode has the fastest timing resolution, decreasing the chance of pileup and preserving the event rate to a reasonable degree up to brighter sources. A source with a true brightness of 40 events/s is observed at 36 counts/s in the small-window mode but only 25 counts/s in the full-frame mode. For this particular plot, even in the absence of pileup the count rate is slightly smaller than the event rate. This is because the count rate is calculated from the number of counts inside the source extraction region, where the event rate comes from the total number of events regardless of where they strike on the CCD. Nevertheless, since the event rate is used as the independent variable in this work, this plot is particularly useful when the count rates corresponding to the event rates are desired.

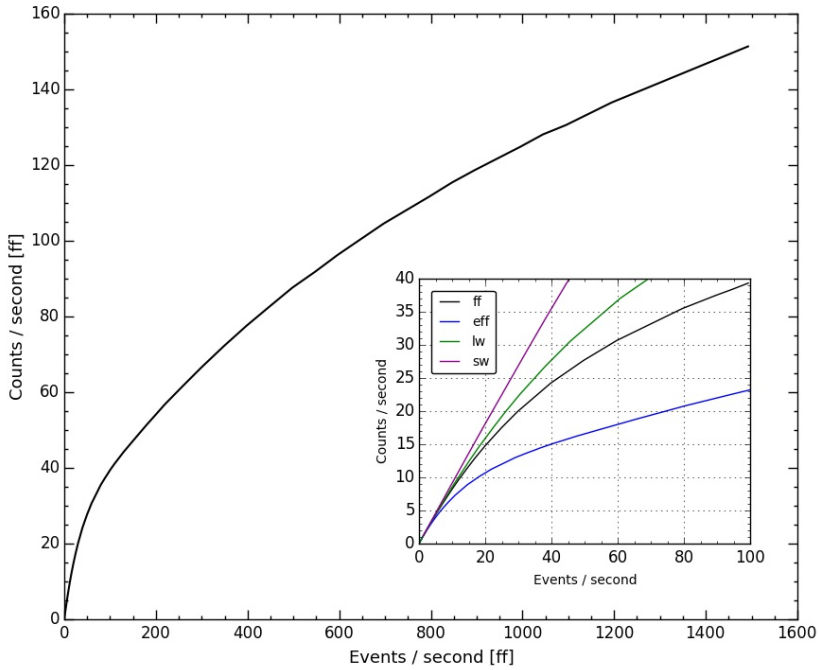

Fig. 4: The event rate in seconds against the count rate in seconds. The event rate includes all events on the detector space for all good pattern types. The count rate is the number of measurable single and double counts inside the extraction region of 60" radius. The values are taken from the delta function simulations at 1 keV. The main figure is for the full-frame only and the inner plot for all modes up to 100 events/second.

Figure 5 gives the pileup fraction and flux loss in the different PN observation modes. The frame times for the full-frame (FF), extended full-frame, large- (LW) and small-window (SW) modes are 73.3, 199.2, 47.7 and 5.7 ms [1] respectively. With the exception of the extended full-frame mode, the readout time is shortened by reducing the area to be read out. It follows, the longer the readout time of a frame is, the sooner pileup and therewith flux loss will occur as a function of source brightness. In the full-frame mode, the pileup peak is observed at 80 events/s compared with 1000 events/s in the small-window mode. Figure 6 examines how the hardness of a source affects the pileup fraction and event loss of an observation. As the hardness of the photons increase, the maximum pileup fraction decreases and shifts to a lower event rate. This is associated with harder photons having larger electron clouds. In terms of detector operation, this means that the fraction of multi-pixel to single pixel events increases with harder photons. As a result, pileup will occur more for fainter sources but also peak earlier. Accordingly, with a larger fraction of multi-pixel events, bad patterns are more likely to begin forming at lower count rates, leading to a greater loss of events. The bottom left image in figure 3 compares the PSFs for the different hardness levels at the same event rate. Notice, for the harder sources more of the core has already been removed.

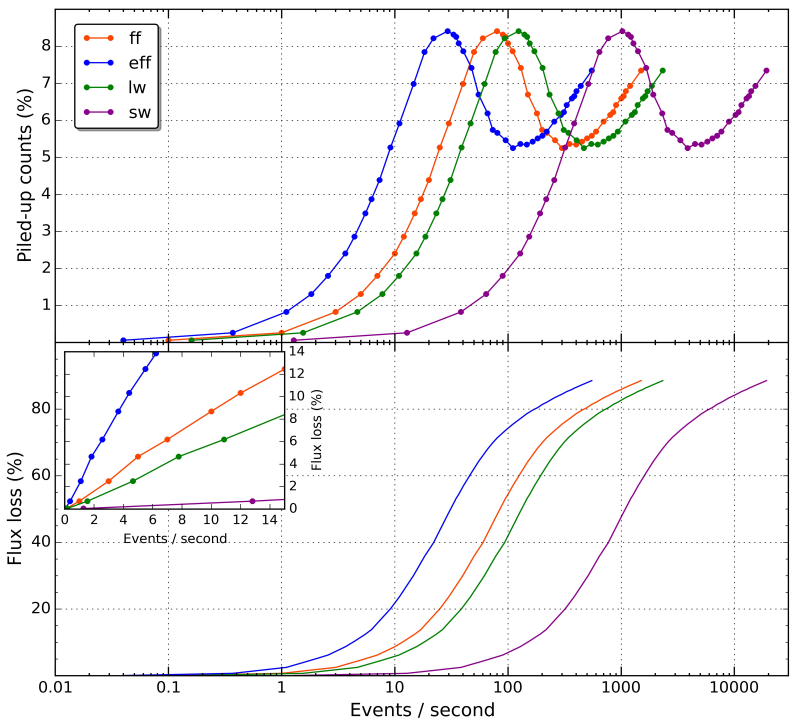

Fig. 5: The pileup fraction and flux loss as a function of event rate for the different EPIC-pn modes. The pileup fractions are from the singles+doubles spectra. The flux is calculated across the 0.2-10 keV energy range. The values are obtained from the delta function source at 1 keV hardness.

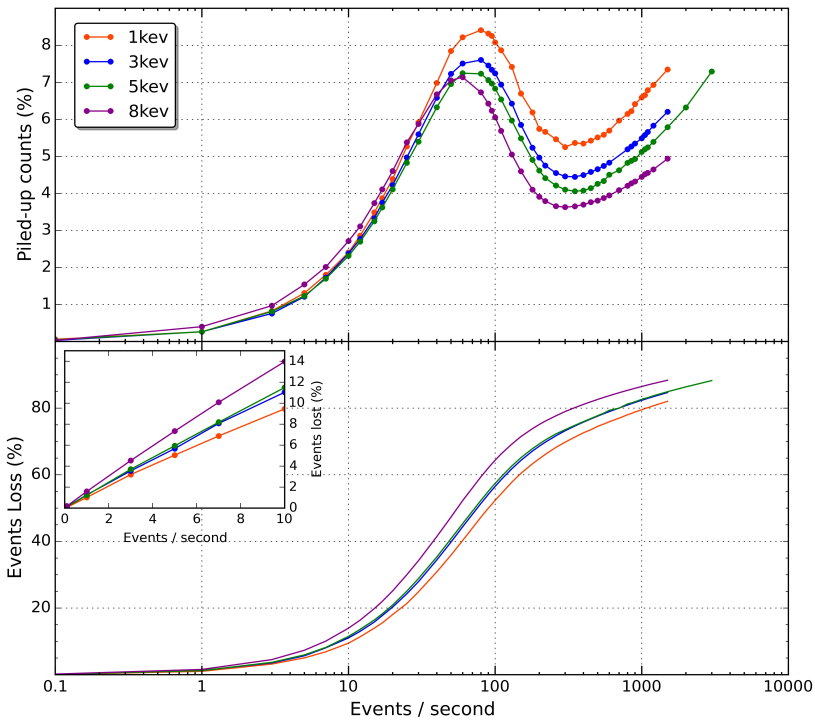

Fig. 6: The pileup fraction and event loss as a function of event rate for different levels of source hardness in the full-frame mode. The pileup fraction and event loss are from the singles + doubles spectra.

3.2 Fitting models to piled-up data

In this section I discuss how the model parameters and fit statistics are affected by pileup. We look at two simple models. A power law with a slope of $\Gamma = 1.7$ and a low temperature black-body at $kT = 0.04keV$, both with a hydrogen column density absorption fixed at $nH = 2 * 10^{20}$ atoms cm^{-2}. A fit is performed on both the output and the event spectrum between 0.2-10.0 keV. Figure 7 shows the fit results for the power law model spectra and figure 8 those for the black-body.

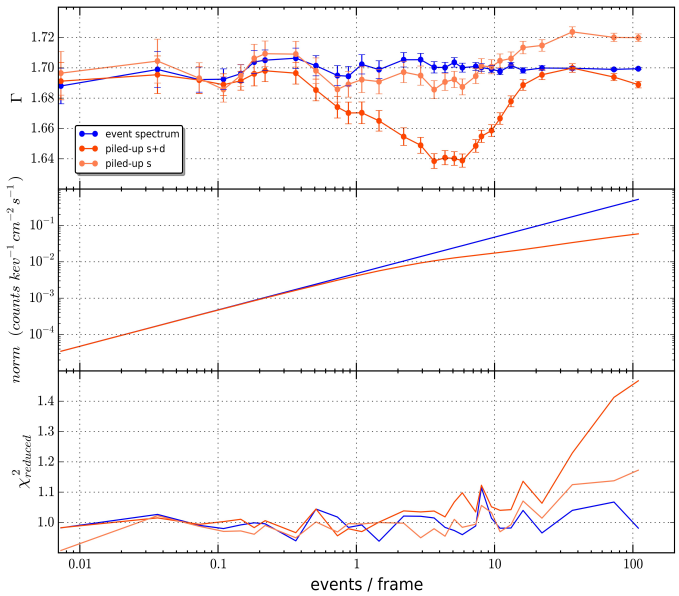

Fig. 7: Power law model fit to the event, singles and doubles output, and singles only output spectra at different brightness levels, leading to different levels of pileup. The fit results to the event spectrum are the true model parameters given no pileup (and thus equal the parameters of the input model given to the simulator). The filled circles represent the event rates at which a simulation is run. The errorbars give a 90% confidence level.

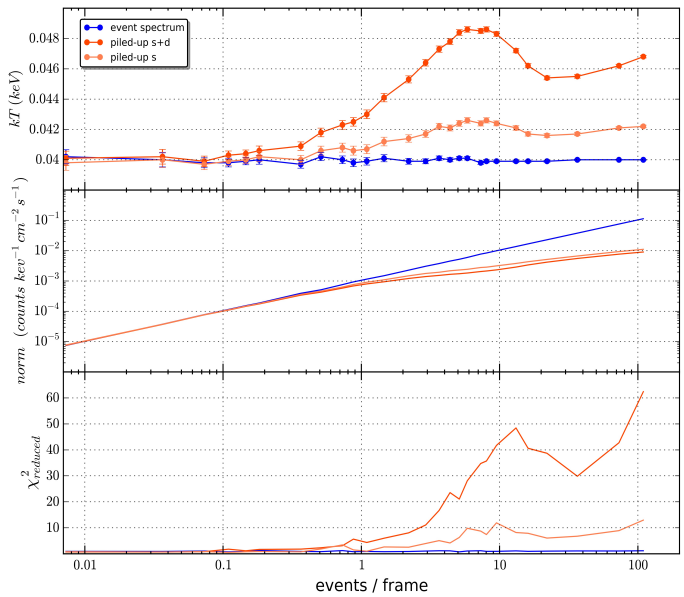

Fig. 8: Black-body model fit to the event, singles and doubles output, and singles only output spectra at different brightness levels.

One can see how the best-fit model parameters (normalisation, kT and power law slope) vary as a function of source flux when fitting the output spectra. At first, as the number of events per second increase, the spectra gain counts at higher energies from soft photons whose

charge combines. In consequence, a higher kT is needed for the black-body model to best fit the output spectra and a flatter slope for the power law model. Figure 9 compares the event spectrum with the output spectrum of the black-body at 4.38 events/frame (60 events per second in the full-frame mode). This plot clearly reveals the migration of photons to higher energies in the piled-up spectrum, compared to what a black-body spectrum should actually look like. At this count rate the spectrum has been distorted so much that it no longer resembles a black-body, thus the model cannot reach a minimum better than $\chi^2_{reduced} = 23.5$ for the fit. For the singles and doubles output power law spectra, even as the source brightness increases, the best-fitting model continues to produce good fits which remain under $\chi^2_{reduced} = 1.1$ up until roughly 8.5 events/frame. Nevertheless, the slope changes throughout to account for the photon energy migration. A power law spectrum, which usually spans over a relatively large range of energies and acts as the continuum, will be more robust and affected less by pileup than a low energy black-body component which only emits photons across a small range of energies. The normalisation parameter for both models decreases as a function of event rate due to increased amounts of bad patterns forming. This ultimately leads to a measured flux which is lower than the true flux of the source (as measured from the event spectrum). Both model spectra have a maximum at ≈ 6 events/frame where the best-fit spectral parameter is furthest from its true value. This corresponds with the pileup fraction curve introduced in section 3. At the maximum pileup fraction, we observe the worst best-fit spectral parameter. When the overall pileup fraction begins to decrease, the model spectral parameter starts re-approaching its true value. Analogously to the pileup fraction, the spectral parameter recovers instead of further worsening because much of the pileup in the core of the PSF is removed in the form of bad patterns. In essence, the removal of these counts has a similar effect as excising the core of an observation. At increasing event rates, however, the pileup fraction increases again, causing the spectral parameters to worsen once more.

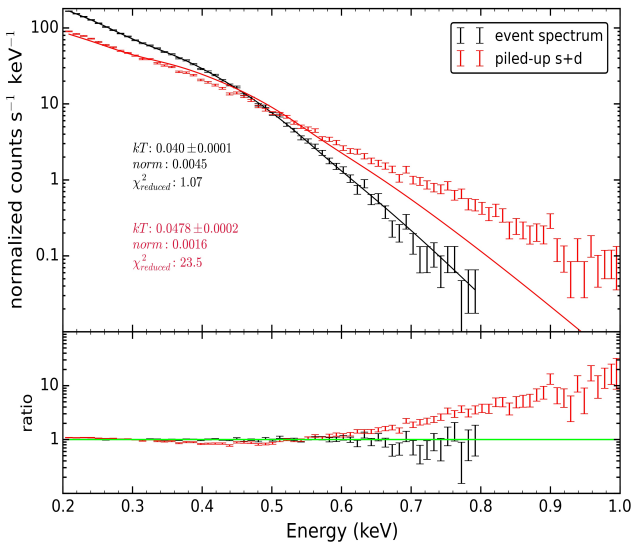

Fig. 9: Top: black-body spectrum and best-fit model at 4.38 events per frame for the event spectrum and output spectrum containing single and double counts. Bottom: corresponding ratio between the spectrum values and best-fit model values.

4 Testing the pileup correction method

In this section I discuss the pileup correction method which can be used as an alternative to excising the core of a piled-up point source. The method is tested on various simulated model spectra, some which closely resemble those from real sources. The ultimate goal is to recover the true flux and spectral parameters of the source. This is achieved when the true model (one which fits the event spectrum), convolved by the rebuilt corrected detector response, provides

the best-fit model for the output spectrum without the need for further parameter adjustment. Figure 10 illustrates the concept of this method, by showing the effect the corrected response has on the model, for a piled-up observation. The response ideally alters the model to fit the contribution from pileup in the spectrum (for the gaussian this is seen well by how the corrected model follows the peaks of the output spectrum which are purely due to pileup). The corrected black-body model provides a good fit at $\chi^2_{reduced} = 1.42$ while keeping the energy, kT, and normalisation close to the true values (0.04 keV and 0.0052 counts $kev^{-1}cm^{-2}s^{-1}$).

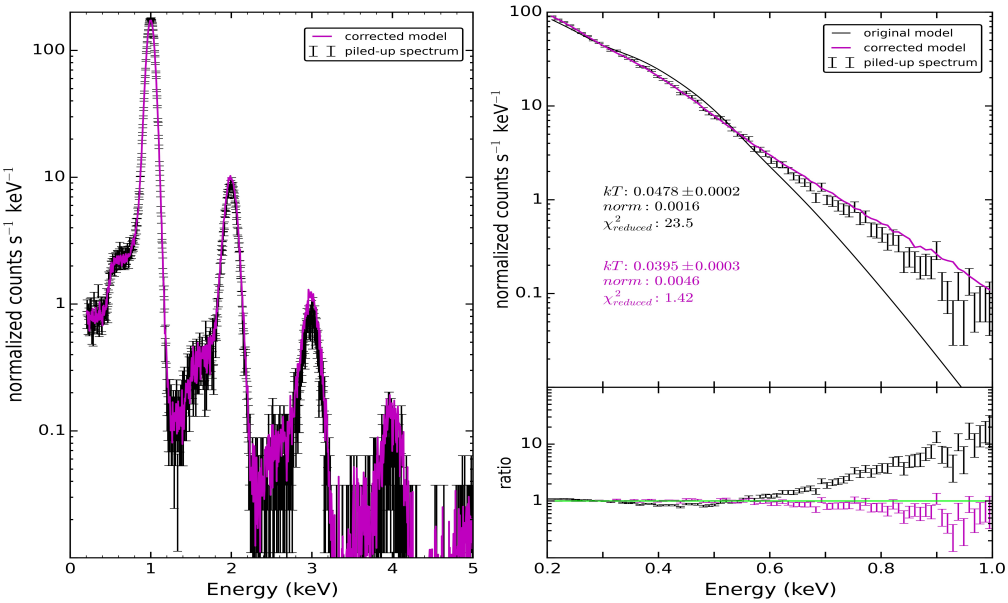

Fig. 10: Left: output spectrum for the 1 keV gaussian model at 5.11 events/frame fit with the input model folded by the corrected response. Right: the output black-body spectrum at 4.38 events/frame and best-fit model; one folded with the original response (black), the other with the corrected response (purple). The ratio between the spectrum and best-fit model are shown at the bottom. Both spectra include single and double counts.

The five models shown in table 1 are tested with the IDL version of the pileup correction method. During fitting, nH is always frozen and an energy range of 0.2-10.0 keV is used throughout. Errors are given to a confidence level of 90%. For the laor component which adds an accretion disk, black hole emission line, all parameters except for the line energy and normalisation are fixed and the Xspec default values are used: an index of 3.0, inner and outer radius of 1.235 to 400.0 GM/c^2 and an inclination of 30.0 degrees.

Input model	nH	Γ	kT	laor
wabs*(bbody)	0.02	0	0.04	0
wabs*(powerlaw)	0.02	1.7	0	0
wabs*(powerlaw)	0.02	2.4	0	0
wabs*(powerlaw+bbody)	0.02	1.7	0.4	0
wabs*(powerlaw+laor)	0.02	1.7	0	6.4

Table 1: Input models used to simulate the event and output spectra, and their given parameter values. nH in $*10^{22}$ atoms cm^{-2}, kT and laor line energy in keV.

4.1 wabs*black-body

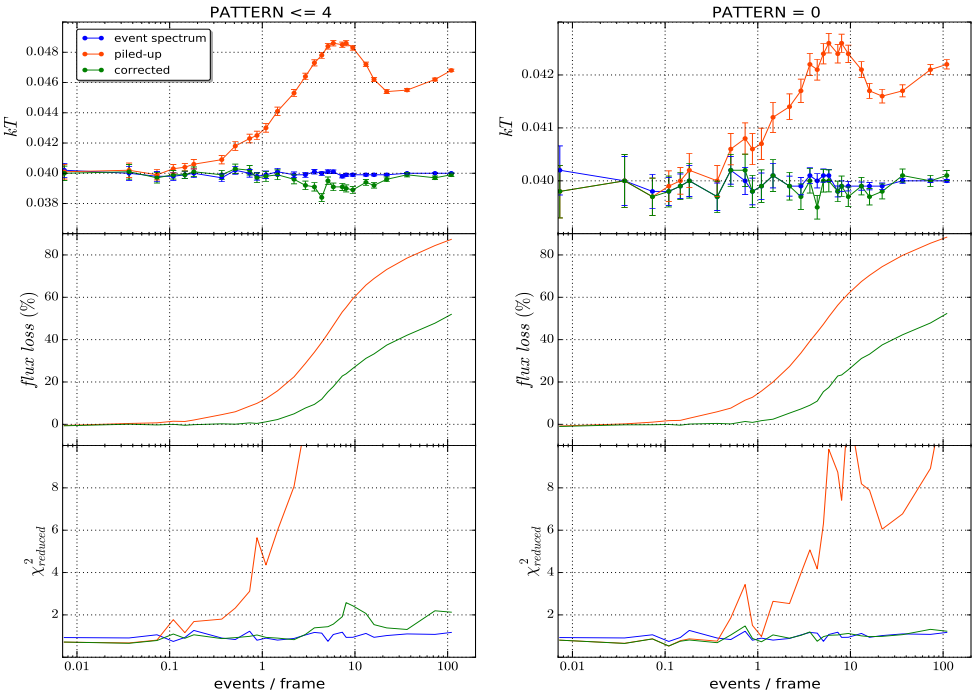

Fig. 11: Left column plots are for the singles and doubles spectra, right for singles only. For both columns, top to bottom respectively is: black-body energy kT, flux loss (in the range of 0.2-10.0 keV) and reduced χ^2, as a function of the events per frame.

4.2 wabs*powerlaw

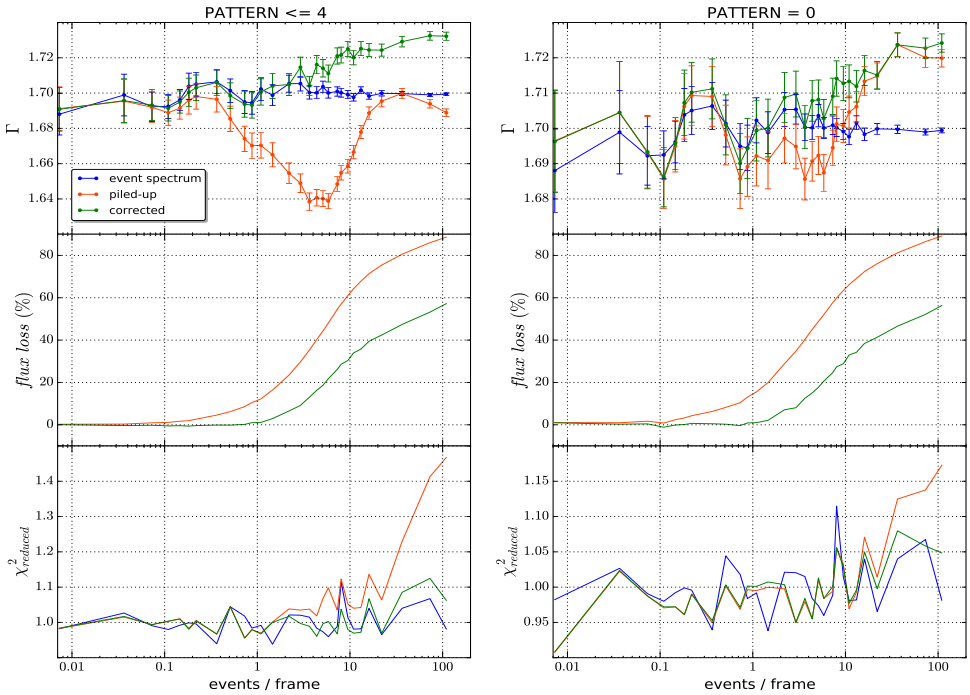

Fig. 12: Left column plots are for the singles and doubles spectra, right for singles only. For both columns, top to bottom respectively is: gamma index Γ, flux loss (in the range of 0.2-10.0 keV) and reduced χ^2, as a function of the events per frame.

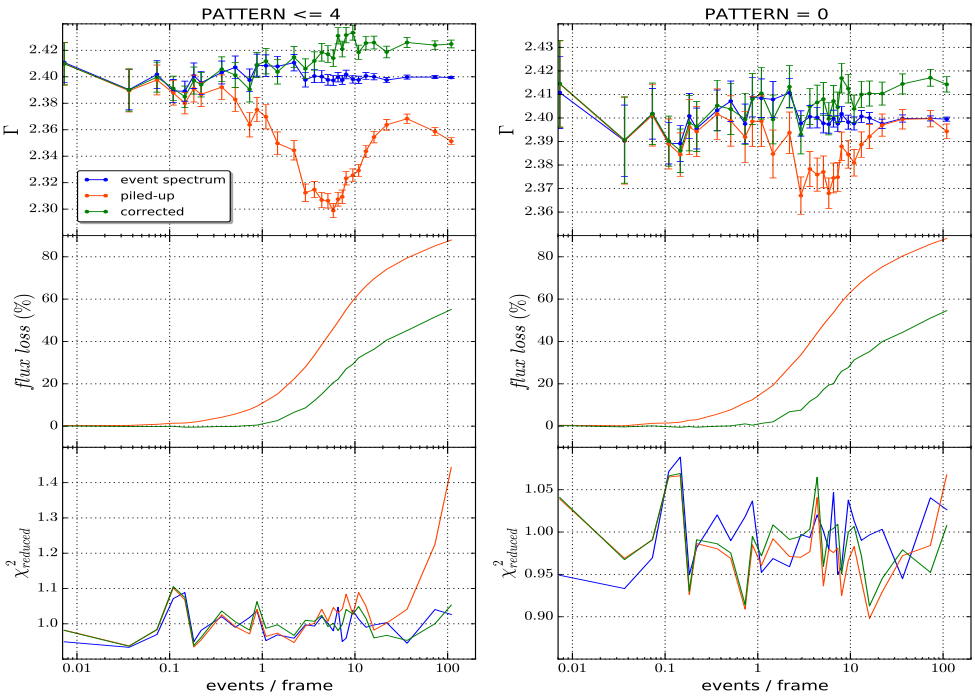

Fig. 13: Left column plots are for the singles and doubles spectra, right for singles only. For both columns, top to bottom respectively is: gamma index Γ, flux loss (in the range of 0.2-10.0 keV) and reduced χ^2, as a function of the events per frame.

4.3 wabs*(powerlaw+bbody)

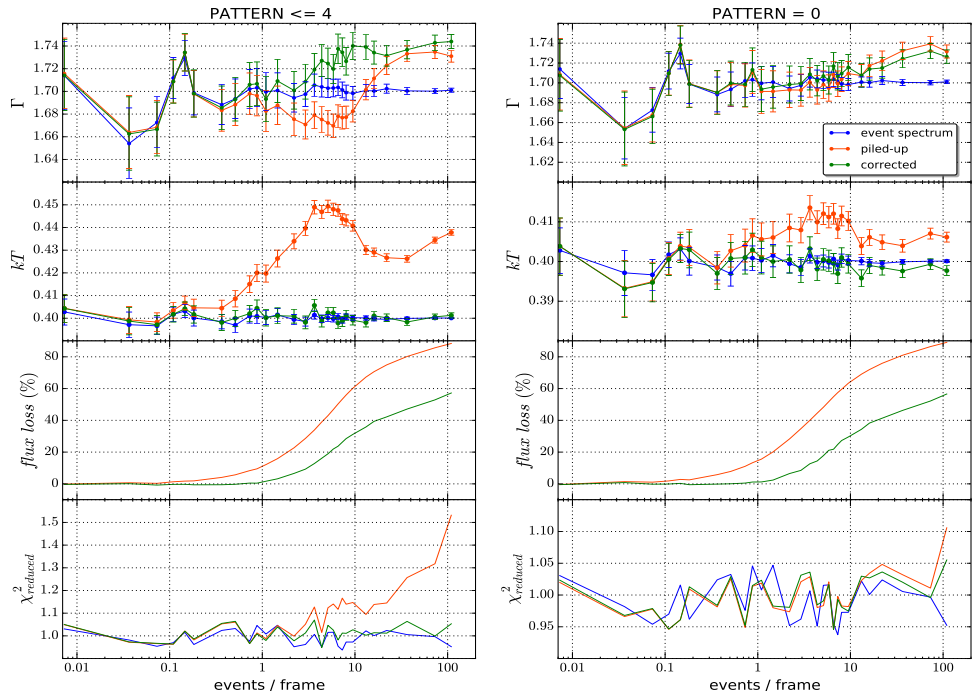

Fig. 14: Left column plots are for the singles and doubles spectra, right for singles only. For both columns, top to bottom respectively is: the gamma index Γ, black-body energy kT, flux loss (in the range of 0.2-10.0 keV) and reduced χ^2, as a function of the events per frame.

4.4 wabs*(powerlaw+laor)

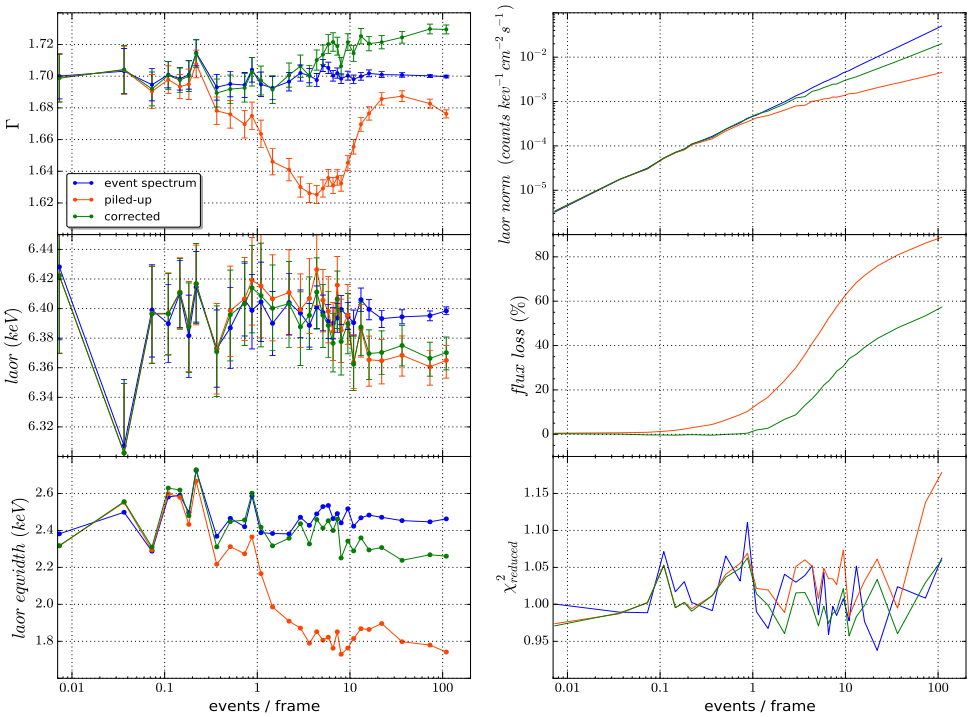

Fig. 15: All plots are for the singles and doubles spectra. In the left column, top to bottom is: gamma index Γ, the energy of the laor emission line and the equivalent width of the laor line. In the right column, top to bottom is: the normalisation of the laor line, flux loss (in the range of 0.2-10.0 keV) and reduced χ^2. All plots are given as a function of the events per frame.

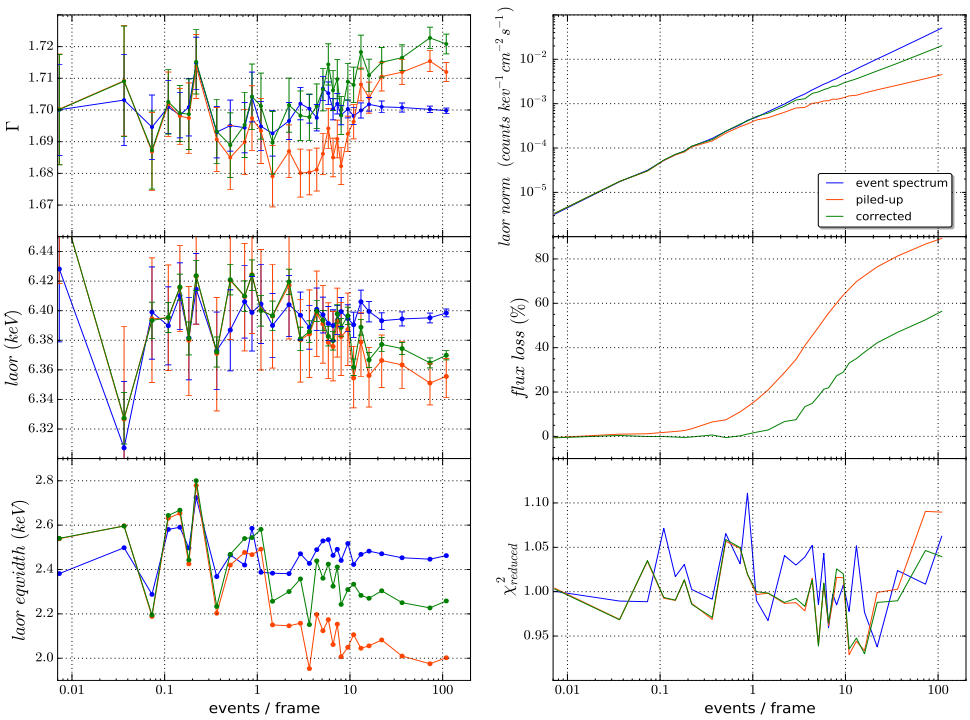

Fig. 16: All plots are for singles only spectra. Otherwise identical to figure 15.

4.5 Summary of the simulated results

From the figures presented in subsections 4.1-.4 I conclude that the flux recovery of the correction method works well (i.e. no flux is lost when using the corrected response matrix) up to **1.5 events per frame**. This converts to 20 events/second in the full-frame, 32 in the large-window and 263 in the small-window, all of which are significantly greater than the proposed count rate limits of the respective modes [1][2]. This flux recovery limit remains true for all of the simulated source models. Above this limit, using the corrected response will still decrease the amount of flux lost, but the true value will not be fully recovered. The magnitude of flux loss after using the correction gets worse as the event rate increases.

How well the spectral parameter(s) are recovered is more dependent on the model spectrum. The black-body spectral parameter, kT, is recovered extremely well at all event rates. We see that the corrected kT parameter is brought back to the original input value. This is true for the black-body both as a separate model and as an additive component to a power law. For both of the power law model spectra, the corrected gamma index, Γ, agrees with the true parameter value up to an incident flux of ≈ 3 events per frame. For observations simulated at a brightness greater than this, the corrected response matrix tends to be over-corrected. This results in a best-fit Γ index that is steeper than the true value. This effect is most severe when the pileup fraction is at its maximum. For the $\Gamma = 2.4$ power law model, the over-correction effect shows a similar behaviour as to that of the pileup fraction curve in section 3.1. The addition of a laor emission line at 6.4 keV does not affect the recovery of the Γ index. Fitting the input model to the output spectrum shows that the line energy of the laor component varies very little as the count rate, and therewith pileup, increases. However, the strength of the spectral line, shown by the equivalent width, quickly falls-off after 1 event/frame, reaching a maximum magnitude loss of 25%. The equivalent width is returned close to its original strength when fitting with the corrected model. In general, for all models, an improved reduced χ^2 is obtained when using the corrected response. Albeit, this is more evident for the singles and doubles spectra than for the singles only spectra.

Ultimately, based on the simulation results, I set two event rate limits up to which the pileup correction method works best. This is shown in table 2.

Recovery limit of	event rate [frame^{-1}]	mode	event rate [s^{-1}]
Flux	1.5	FF	20
		LW	32
		SW	263
Spectral Parameter	2.9	FF	40
		LW	62
		SW	508

Table 2: Event rate limits for the pileup correction method.

4.6 The quality of the corrected detector response

The reliability of the method depends, in part, on the quality of the corrected detector response. Section 2.2 describes how the response is built, energy by energy, from the single simulated events that are added into each frame of the observation. Therefore, if the duration of an observation is very short and thus does not have a large amount of frames, the quality of the corrected detector response will suffer. Figure 17 shows the difference between a clean response and one generated from an eventlist with too few frames. The corrected detector response created from an observation with 15,000 frames is on the left and one with 100,000 frames on the right, both have been convolved with a narrow gaussian model at 1 keV. The 'tail' following the 1 keV

peak is due to pileup. It is apparent that the response on the left is of poor quality and will be inadequate in estimating the pileup contribution to the spectrum. It should be insured that at least 100,000 frames are in an observation's eventlist when using the pileup correction. For very high count rates, it is best to have more, since many of the simulated events are lost due to forming bad patterns and therefore do not build the detector response.

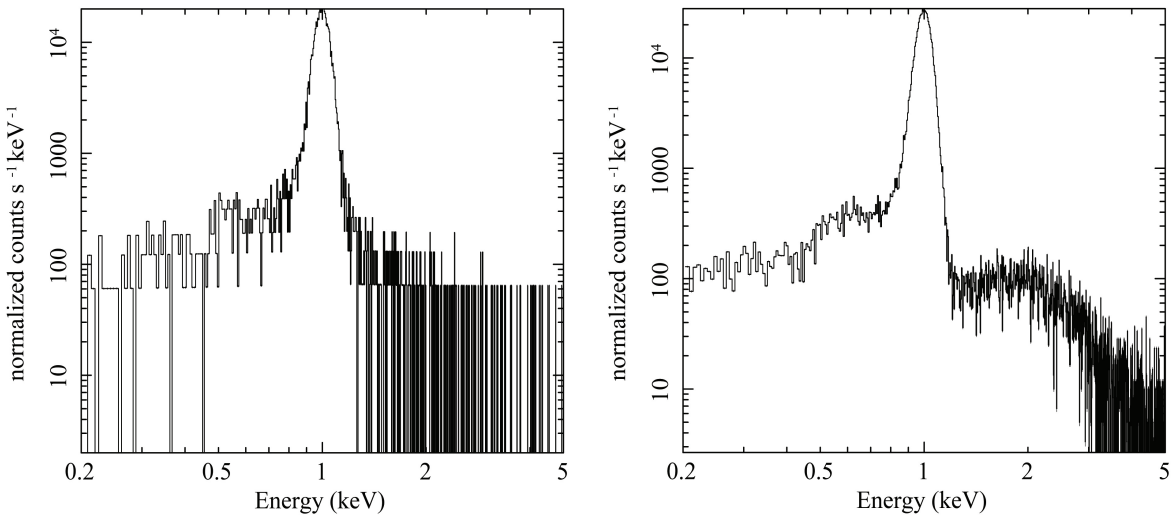

Fig. 17: Detector response corrected for pileup at 5 counts per frame folded with a narrow gaussian at 1 keV. Left: detector response generated on an eventlist with 15,000 frames. Right: response generated on an eventlist with 100,000 frames.

References

[1] L.Strüder. The European Photon Imaging Camera on XMM-Newton: The pn-CCD camera. *A&A*, 365:18–20, 2001.

[2] P.Jethwa. Pile-up thresholds for the EPIC cameras. pages 1–4, 2012.