epreject

August 15, 2005

Abstract

Task **epreject** (1) corrects shifts in the energy scale of specific pixels which are due to high-energy particles hitting the EPIC PN detector during the calculation of the offset map and (2) suppresses the detector noise at low energies by statistically flagging events according to the known noise properties of the lowest energy channels. In the case of timing mode data, **epreject** also permits the flagging of soft flare events.

1 Instruments/Modes

Instrument	Mode
EPIC PN	IMAGING

2 Use

pipeline processing	yes (?)
interactive analysis	yes

3 Description

3.1 Correction of the energy scale in specific pixels

The usual mode of operating the EPIC pn camera consists of computing an offset map immediately before the beginning of an exposure. Ideally, this map contains the energy offset for each pixel (expressed in analog–to–digital units, adu). During the exposure, these offsets are subtracted onboard from the measured signals, and only events where the difference exceeds a lower threshold (usually 20 adu, which formally corresponds to 100 eV) are transmitted to Earth.

High–energy particles hitting the EPIC pn CCD during the offset map calculation may cause the affected pixels to get offset values which are incorrect by a few adu. As a consequence, the energies of all events in these pixels appear to be shifted by the same amount. Due to the specific method of deriving the offset map onboard, the affected areas occur often in blocks of four consecutive pixels along readout direction.

Depending on the orientation of the trail caused by the high–energy particle with respect to the CCD, these areas may also extend over several consecutive pixels perpendicularly to the readout direction. The affected pixels usually get an offset which is too small. If the affected area extends over several pixels along a CCD row (perpendicularly to the readout direction), then the remaining pixels within this CCD row may get an offset which is too high.

If the offset to be subtracted is too small, then the adu values which are assigned onboard to events in such pixels become too high. Thus, events which have adu values below the lower threshold and which would normally be rejected, may show up in the data set. As most of such events are due to detector noise, which is steeply increasing towards lower energies (Fig. 1), any reduction of the lower energy threshold leads to a considerable increase in the number of events. An immediately apparent consequence of this effect is the occurence of bright patches in EPIC pn images which are accumulated at low energies (e.g. Fig. 3). A less obvious consequence is a shift in the energy scale over the whole spectral bandwidth. This shift degrades the energy resolution for extended objects. For point sources, the X–ray spectrum may be shifted by some 10 eV, in most cases towards higher energies, if the position of the source happens to coincide with one of these patches.

Task **epreject** provides two methods to correct the energy scale:

1. If offset maps are available these can be used directly to estimate the offset errors: Since the effected chip areas are limited to regions where the chip was hit by high–energy particles during the offset map calculation, the offset error can be determined by subtracting the value of the offset map in the effected areas from the value in the unaffected areas. This is achieved by in turn subtracting its median value from each column and each row of the offset maps. The values in the remaining residual offset map can then be regarded as the offset errors which need to be subtracted from each event falling into the respective pixels.

2. If no offset maps are available the offset errors can be corrected using a method based on the count images accumulated in the lowest energy channel. As the detector noise is monotonically increasing towards lower energies, a correlation is expected between the brightness of such pixels at low energies and the amount of offset shift which they have received. Evidence for such a correlation was indeed found, in particular when only the lowest transmitted adu value (usually 20) is used for determining the pixel brightness. However, this correlation is disturbed by the fact that the brightness of a pixel at 20 adu is also influenced by other factors, in particular by its individual noise properties. In order to separate offset–induced changes of the 20 adu pixel brightness from other brightness variations at 20 adu across the detector, a reference image is subtracted from the 20 adu image. This reference image contains for each pixel the nominal, i.e., temporally constant, value of its 20 adu brightness (Fig. 2). The reference image was derived by accumulating images at 20 adu from long FF exposures with no bright X–ray sources in the field, and computing the median value for each pixel.

   The intensity in the subtracted, normalized 20 adu images is then used to reconstruct the value of the offset shift, which was incorrectly applied onboard, and the raw amplitudes of all events in the corresponding pixels are shifted back by this amount to their nominal value (Fig. 4). The reconstruction of the offset shift is done by using calibration data which were derived from exposures where offset maps were available.

Although this task readjusts the energy scale, there are some effects left which cannot be corrected for:

- If the energy scale has to be corrected towards lower energies, then events may get raw amplitudes which are below the low energy threshold applied onboard. In order to re–establish a common lower energy threshold, such events should be removed in a subsequent step.

- If the energy scale has to be corrected towards higher energies, then events with very low raw amplitudes are missing (because they were not included in the telemetry). In order to re–establish a common lower energy threshold, this threshold could be increased by the maximum shift which was applied towards higher energies. However, while a shift towards higher energies is usually required only for very few pixels, information about the lowest energies would then be lost for all pixels. Furthermore, the area around the object of interest might not be affected at all. Therefore, it should be checked in each specific case whether a general increase of the lower threshold is appropriate.

Caveats:

In option (2), above, the task attempts to reconstruct the offset shifts from the brightness of pixels at 20 adu. While it is guaranteed that the offset shifts can only occur at discrete adu steps, the correspondence between the 20 adu brightness and the value of the offset shift is not always unique. The presence of Poissonian noise in the 20 adu images, in particular for short exposures, limits the sensitivity for spotting the bright patches and deriving the appropriate energy correction. The parameter which specifies the minimum significance which a block of four consecutive pixels along readout direction must have in order to trigger the offset shift correction task for this block, can be set by the user. Tests indicate that setting this parameter to $\sim 4\,\sigma$ is a good choice for short ($\sim$ 5 ks) exposures; for longer exposures this parameter can be increased (to $\sim 5-6\,\sigma$ for more than 20 ks). It is recommended to control the results by accumulating an image below 20 adu after this task: this image shows the pixels where an offset shift was applied (Fig. 5).

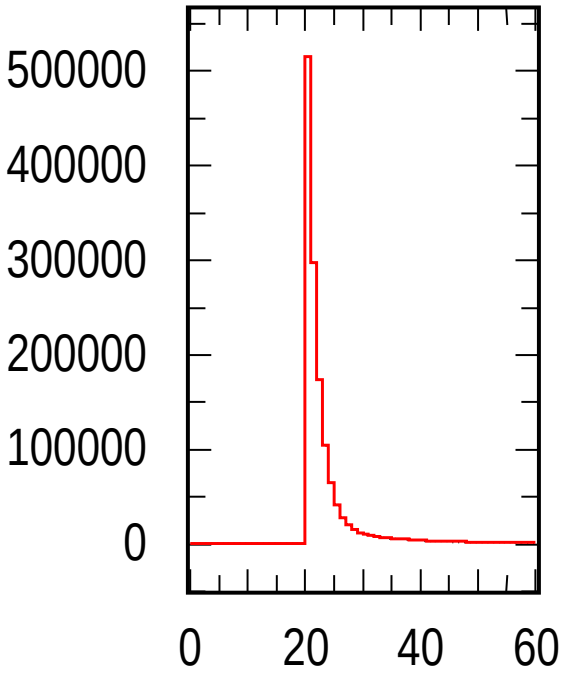

Figure 1: Number of events as a function of PHA [adu] for quadrant 0, obtained during 50.7 ks in a closed FF exposure (`0059_0122320701_PNS003`). The peak at 20 adu corresponds to an event rate of 0.75 events/frame/quadrant, or $2 \cdot 10^{-5}$ events/frame/pixel.

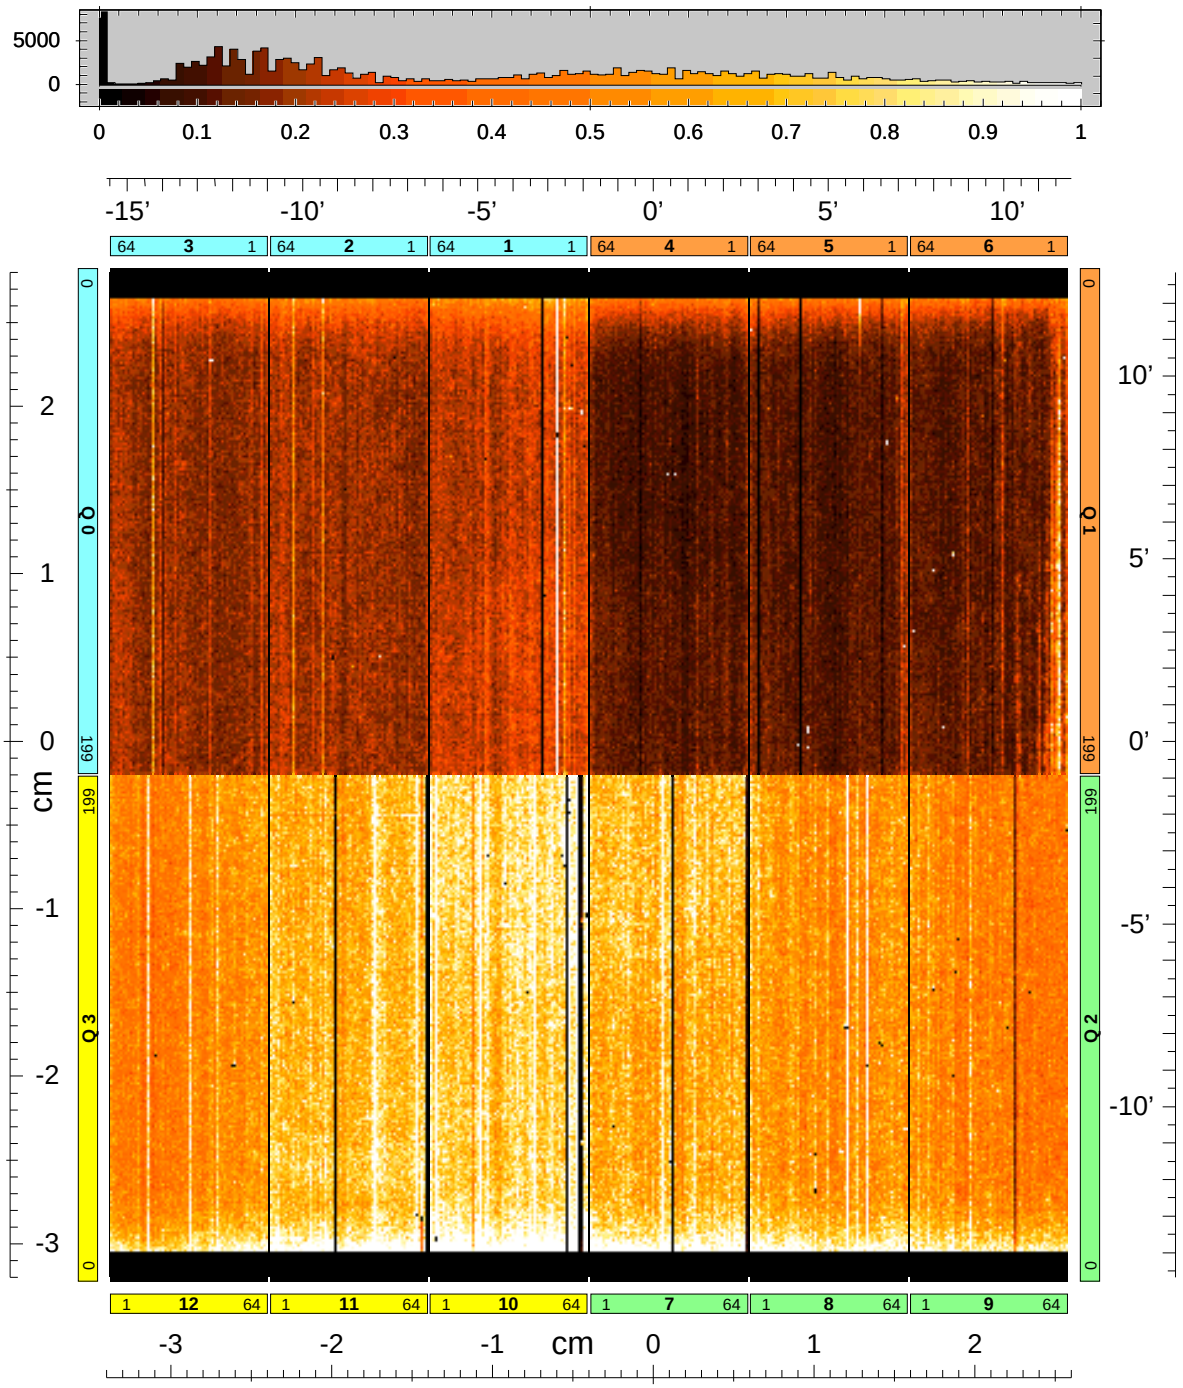

Figure 2: Reference image for the 20 adu intensities per pixel, normalized to an exposure of 1 ks. It was obtained by accumulating images at 20 adu from long FF exposures with no bright X–ray sources in the field, and computing the median value for each pixel. This image contains for each pixel the temporally constant value of its 20 adu brightness.

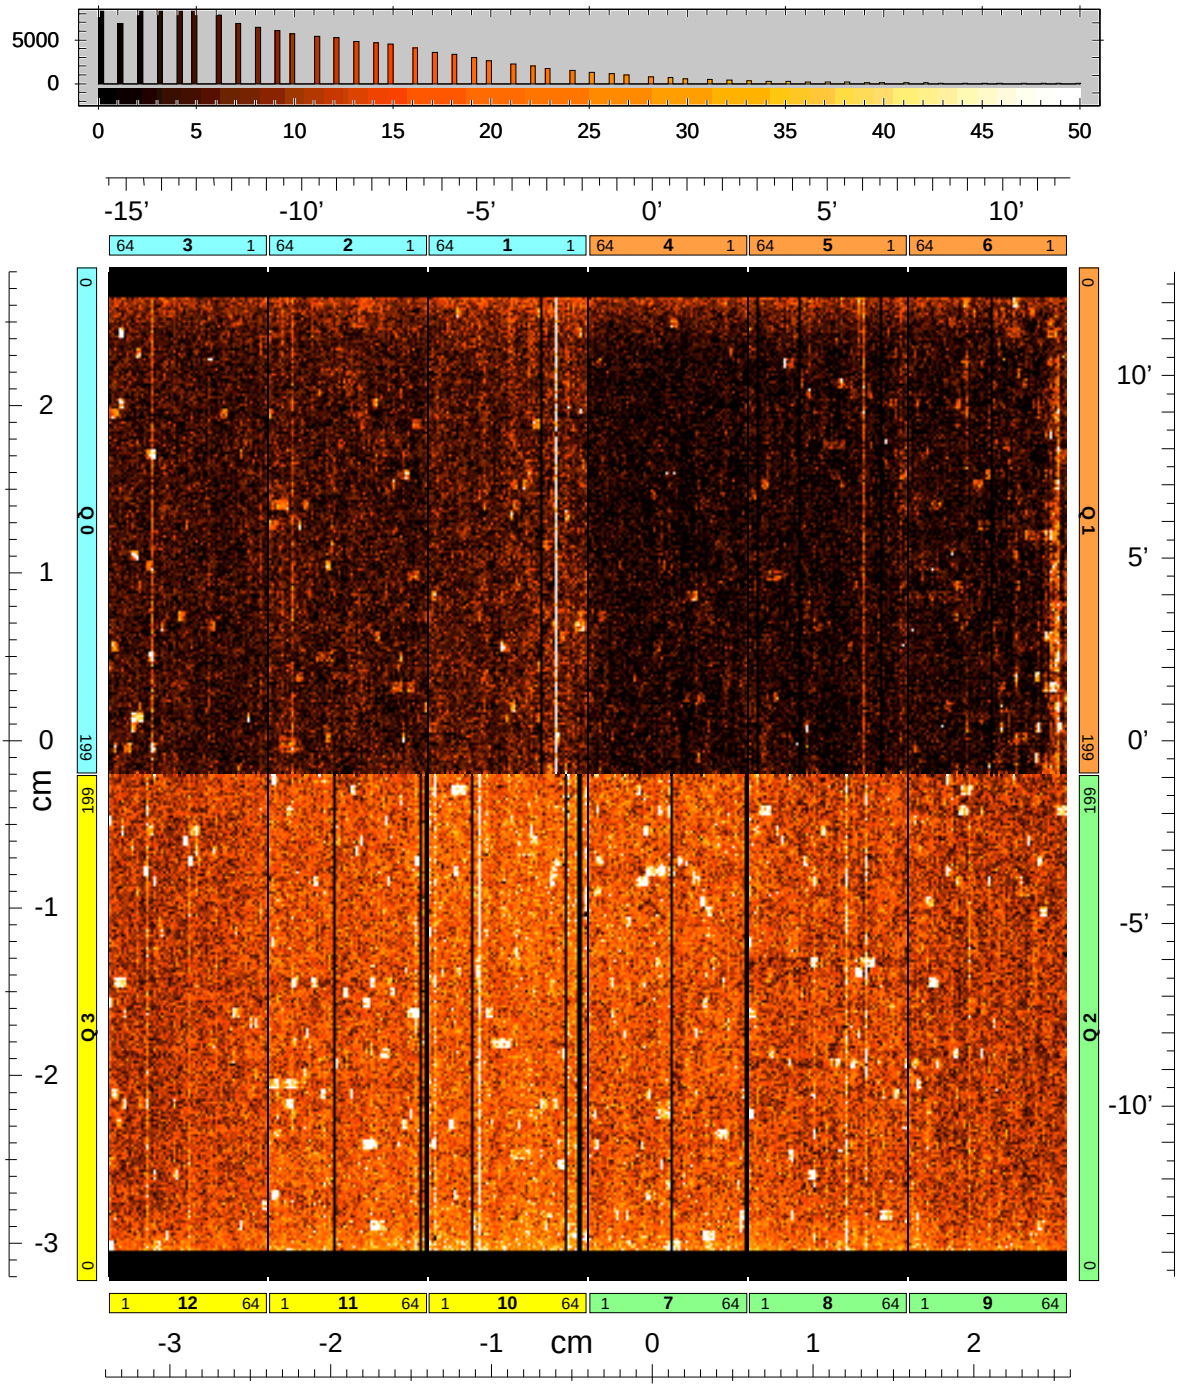

Figure 3: Events with raw amplitudes of 20 adu in the closed 23.2 ks FF exposure 0462_0134521601_PNS005. The number of events per pixel is color coded according to the color bar at top, ranging from zero (black) to 50 (white).

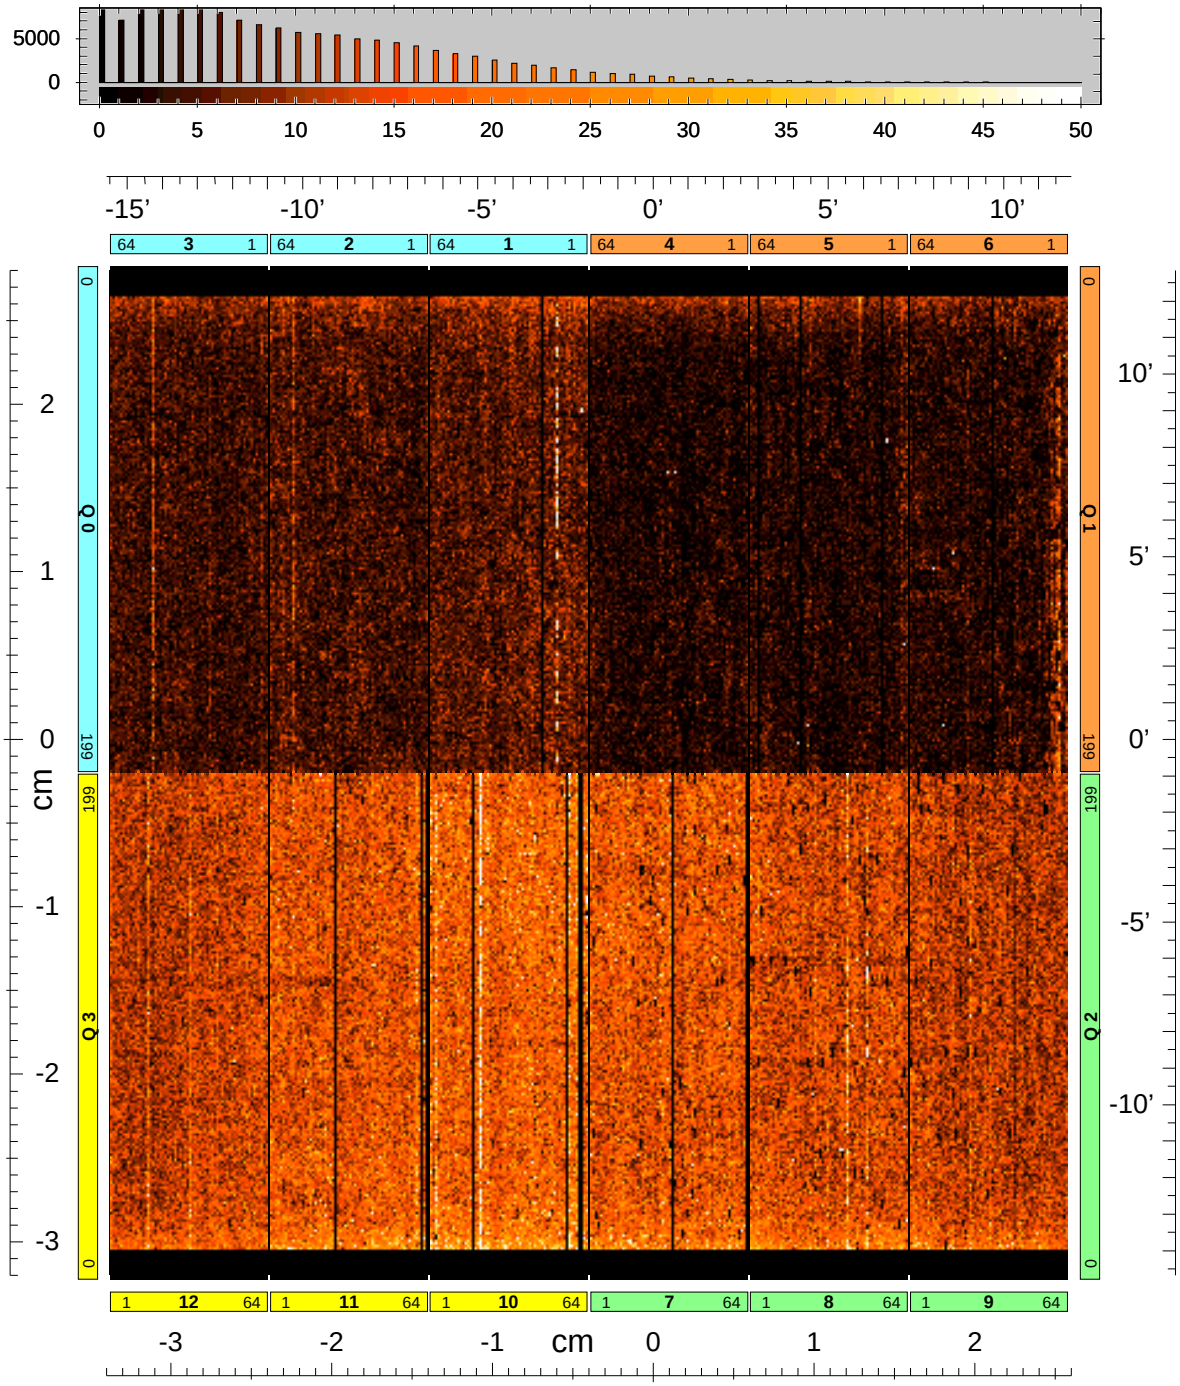

Figure 4: Same as Fig. 3, but after applying the task `epreject` for correcting the energy scale in specific pixels.

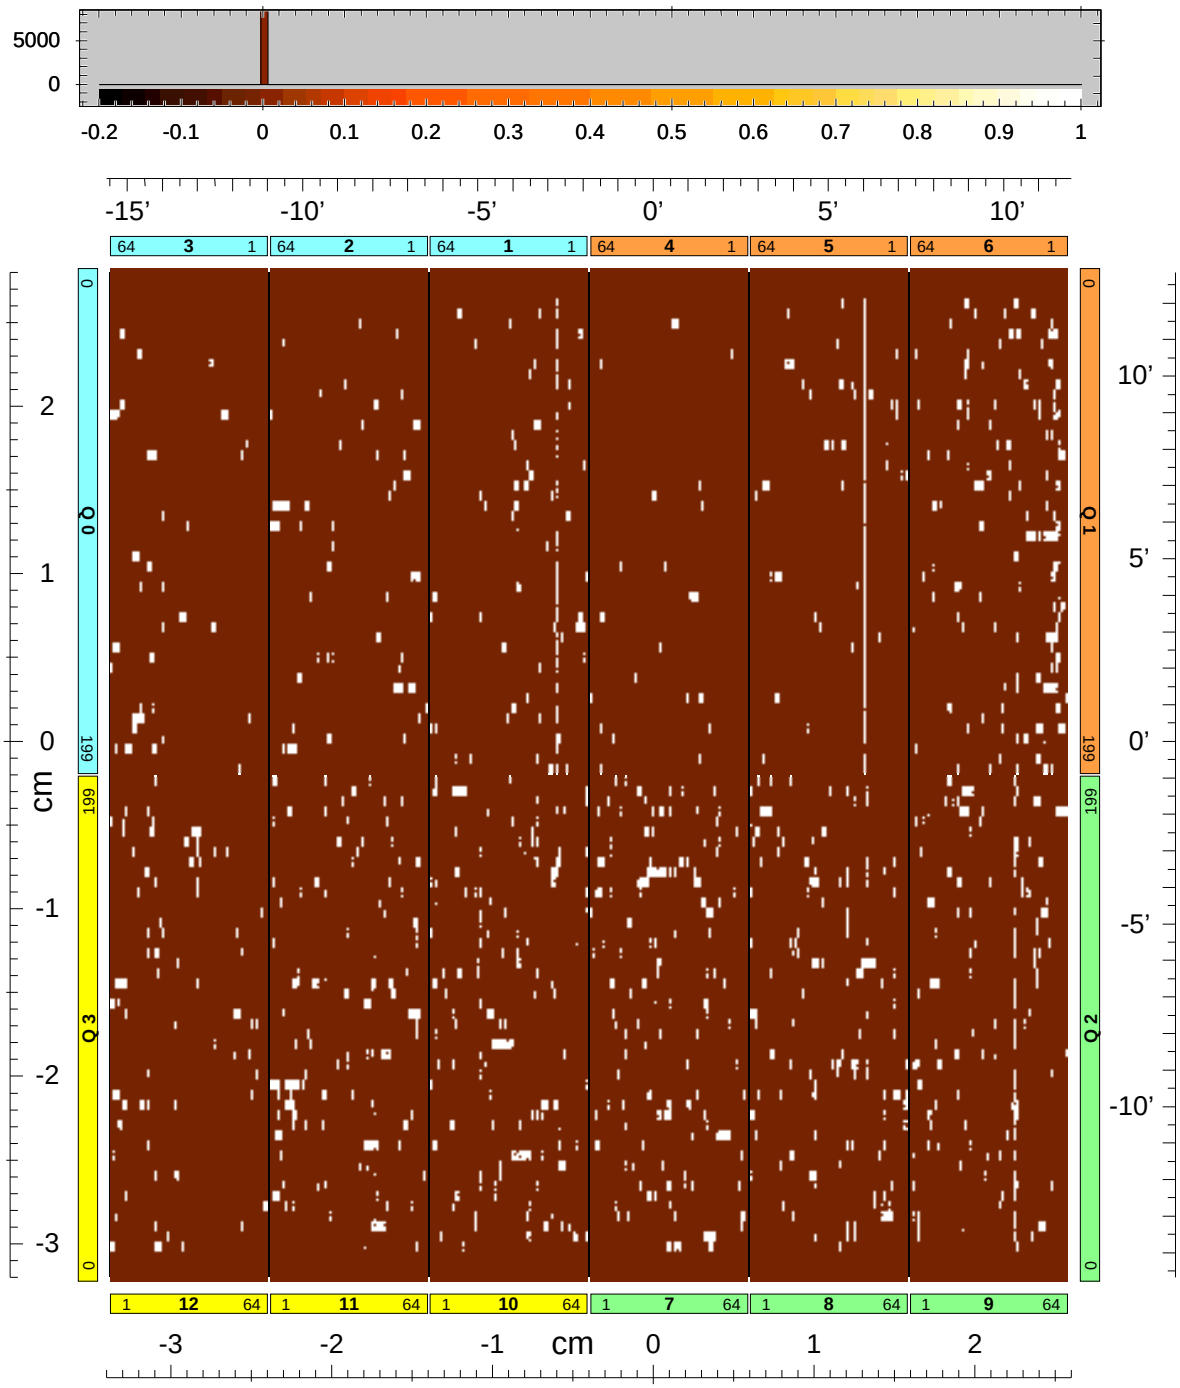

Figure 5: Pixels in the exposure 0462_0134521601_PNS005 where offset shifts were applied. This image was accumulated from events with PHA < 20 adu, after running the task `epreject`. A $4\,\sigma$ threshold was applied for the identification of bright patches.

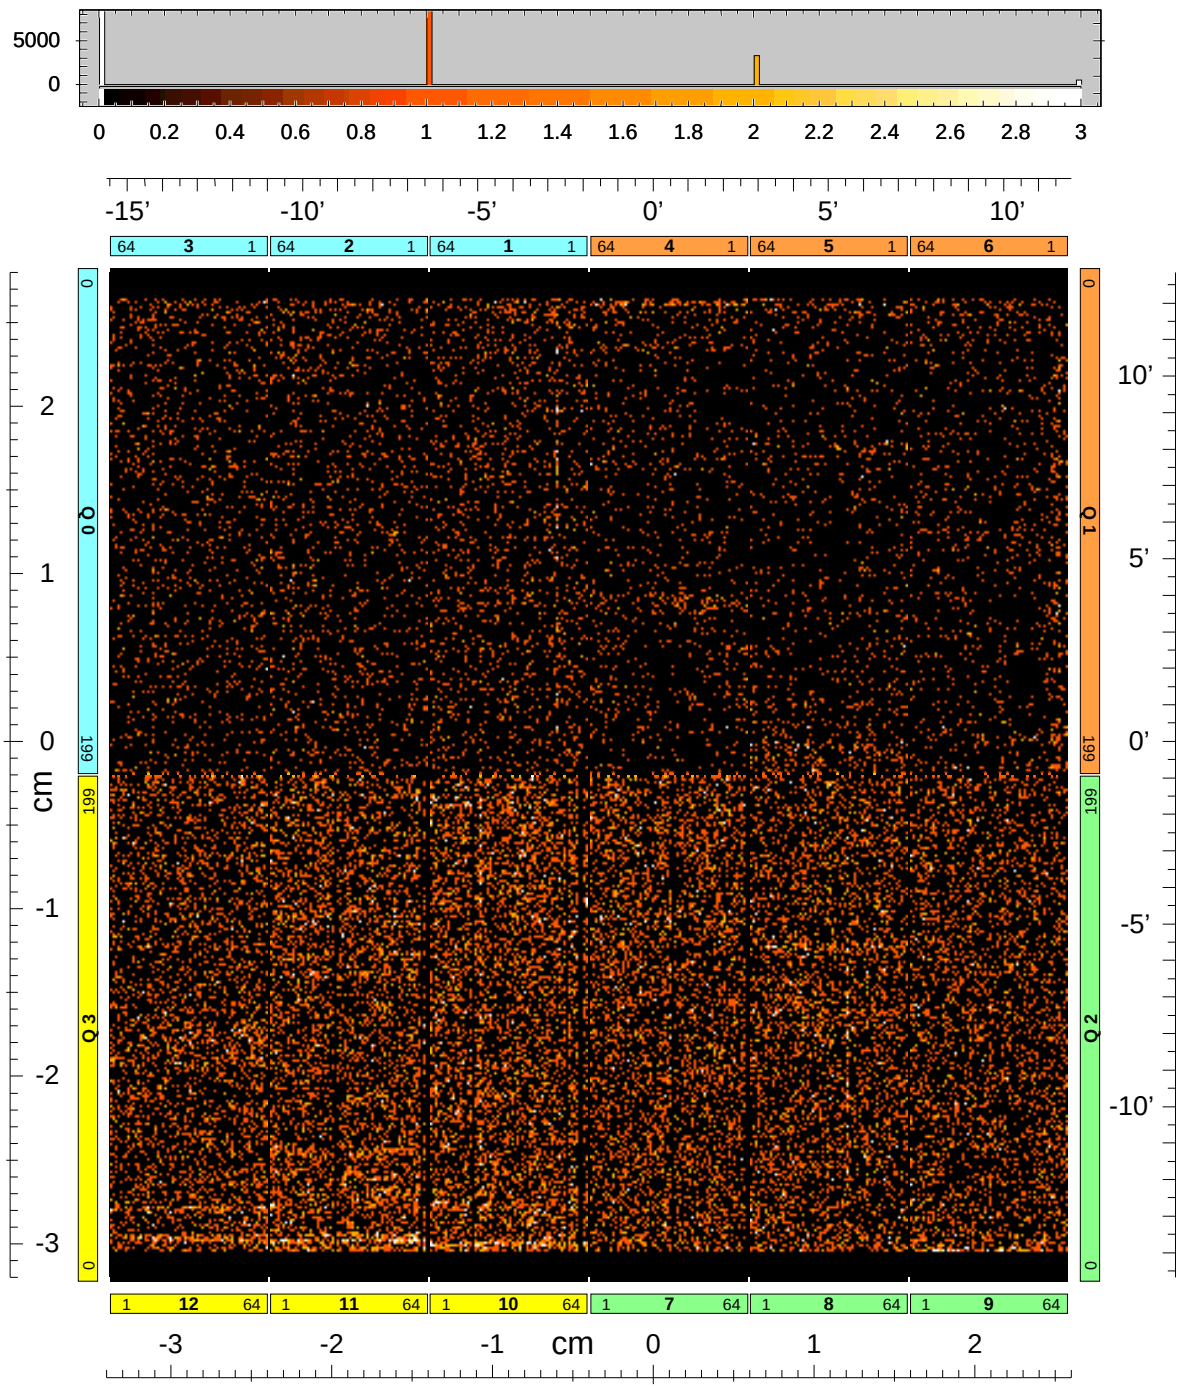

Figure 6: Events with raw amplitudes of 20 adu in the closed 23.2 ks FF exposure 0462_0134521601_PNS005, after correcting the offset shifts and after suppressing the noise. Note that the intensity scale extends now from 0 to 3, while it covered the range from 0 to 50 in Figs. 3 and 4.

3.2 Suppression of detector noise at low energies

While there is practically no detector noise present at energies above 250 eV, X-ray data below ~ 200 eV are considerably contaminated by noise events, which become more and more dominant towards the lowest transmitted energies (Fig. 1). Investigations of 40 hours of in–orbit calibration data with the filter wheel closed, taken over a period of more than two years, showed that the noise properties vary with position and energy, but are fairly stable in time for most areas on the detector. This property enables a statistical approach for suppressing the detector noise. The task uses the information about the spatial and spectral dependence of the detector noise in order to flag and optionally remove, on a statistical basis, the amount of events which correspond to the expected noise. Apart from a dramatic reduction of the event file size (if the noise events are physically removed) this procedure thus has the additional advantage to correctly treat the spacial and spectral properties of the detector noise in a more straightforward manner than would be possible by conventional background subtraction.

The noise properties were derived from a total of 10 exposures with closed filter wheel between XMM–Newton revs. 129 and 532, yielding a total exposure time of 144 ks. These measurements were first corrected for offset shifts (see above) and then used for accumulating PHA spectra below 65 adu, individually for each CCD row. The fine spacing in RAWY was chosen because the noise properties change considerably with distance from the readout node, in particular close to the CAMEX. In order to get a sufficient number of events, no subdivision was made along the CCD rows (i.e., along RAWX). This approach was motivated by the fact that the noise properties do not show a pronounced dependence along this direction (unless there is a bright column). As the resulting spectra suffered somewhat from low count rate statistics, they were smoothed along RAWY with a running median filter extending over ± 5 rows. This smoothing was not applied to rows with RAWY < 20, where the spectra contained more counts and where the dependence of the spectra on RAWY was high.

For the flagging of potential noise events in a specific observation, the number of events in each (CCDNR, RAWY, PHA) bin is compared with the corresponding value in the noise data (scaled to the same exposure). According to the ratio between the actual number of events and the expected noise contribution in each (CCDNR, RAWY, PHA) bin, individual events are then randomly flagged. In order to improve the statistics somewhat, the spectra from the observation to be corrected are internally smoothed by a running median filter along RAWY, for rows with RAWY ≥ 20, in a similar way as the noise spectra, before the noise contribution is computed. The data set itself, however, remains unchanged by this task, except for receiving the additional flags.

The data from exposures where the filter wheel was closed, show that the noise properites may exhibit some temporal variations, which are usually similar in all CCDs of a quadrant. In order to allow some fine tuning to an individual observation, twelve parameters can be controlled by the user. These parameters specify the relative amount of noise in each CCD. An additional parameter controls the maximum percentage of events in each (CCDNR, RAWY, PHA) bin, which may be considered as noise. This parameter should be set to a value which is slightly below 100% (default: 98%), to take the fact into account that even in the exposures with the filter wheel closed, which were used for deriving the noise properties, not all events are due to detector noise. There is, e.g., some additional flux present from fluorescence at the filter wheel itself, triggered by energetic particles. This component would change with the position of the filter wheel. The appropriate setting of these parameters should be checked by accumulating an image of all events with low PHA values (below ~ 25 adu) which were not flagged as noise events. If there was no bright, extended X–ray source in the field of view, then the events should be distributed homogeneously across the image.

If a satisfactory suppression of noise events can be achieved, then this task allows to extend the useful energy range to an instrumental energy of ~ 120 eV. Below this energy, parts of the detector become essentially insensitive due to the low energy threshold applied onboard and the combined effect of charge transfer loss and gain variations within the 768 amplifiers in the pn camera. In addition to the improvement of the data quality, removing the events which were flagged as noise events will also make the files

considerably smaller and easier to handle. Tests have shown that up to $\sim 90\%$ of all the events can be removed from the original data set by this method (e.g. Fig.6).

Caveats:

Tests with exposures taken between XMM–Newton revs. 129 and 532 indicate that the noise properties did not change much during this period. However, there may be cases where the noise properties deviate significantly from the behaviour expected by this task. This may have happened in particular early in the mission when the offset map was computed with a different method

This task was verified to yield satisfactory results down to exposures of ~ 5 ks. For considerably shorter exposures, however, there is a limitation due to count rate statistics (how should one remove, e.g., 20% of the events if there is only one event present ?). It is therefore always recommended to check the results with the method described above.

3.3 Soft flare rejection for fast timing mode data

3.3.1 Background: motivation for developing this procedure

The EPIC-pn fast timing mode was designed so that a time resolution of 0.029 milliseconds can be achieved. This mode is important for observing bright variable sources with a very high time resolution. Up to now the it has only been possible to use the spectra down to 300 eV in this mode. Below this energy the data is affected by so-called soft flares which are caused by stack overflows generated by high energy particles.

These so-called soft flares have sharp rise (from 0 counts/sec to several thousand counts/sec in $\sim$0.10-0.15 sec) then decaying back 0 counts/sec slightly slower in about 0.3-0.4 sec). In some cases (extremly bright flares) they can lead to a FIFO overflow thus causing a small gap in the data of about $\sim$0.08 sec. These gaps are not yet handled by the SAS.

In this task we provide a method to mitigate the effect these flares have on the data for flagging such soft flare events so that they can later be screened from the data by the epreject command in the epchain.

With this method it at last becomes possible to use the spectral information in the data down to the lowest energies detectable in this mode i.e. 200 eV.

This is particularly interesting for timing studies of isolated neutron stars, X-ray Binaries and other variable objects, such as magnetic CVs, with very soft spectra.

3.3.2 The implemented version for the screening:

In this version of epreject only the boxcar method for screening the soft flares has been implemented. This method has been proven to be very robust. In order to perform the soft flare event selection the program generates a lightcurve in 0.1 sec bins in the specified PHA range using the rawevents obtained from the epframes procedure. This lightcurve is then smoothed with a boxcar function. The width of the boxcar can be given as parameter by the user. This smoothed lightcurve is then used to help perform the soft flare event selection. Here each photon is checked to see whether it has a PHA value in the user given PHA range (default: 40-50 ADU) and the lightcurve is above the user given softflarethreshold1 (this absolute count per 0.1 sec bin value has to be optimised for each observation as the contribution of the source photons in the given PHA range varies a lot from observation to observation). The criteria above keep events above the upper PHA threshhold during the softflare events as good events. A flag

column labeled "Flare" is added to the rawevents file. A "1" in the column denotes that is has been flagged as a soft flare event. This column can the be used to screen the data.

3.3.3 Caveats: things to take into account when using the screened data

Screening the data on the Flare column requires the user to keep several things in mind when interpreting the data down to 200 eV.

- **Spectra:** They are problematic as all photons in the given PHA range during soft flares are removed from the eventfile (both source and softflare events). Spectral fits to the screened data can be performed for PHA values greater than the upper PHA limit given by the user. Below that the total exposure time is reduced as a consequence of the flare removal. A correct handling for this would require energy dependant GTIs. In the future a method for statiscally separating source and softflare events could correct this but could be problematic for stronly variable sources.
- **Timing:** There is no problem above the upper PHA limit. But below the upper PHA limit the same applies here as for the spectra. Here the also as a consequence of the flare removal total exposure time is reduced.

3.4 Usage

Task **epreject** is intended to be run immediately after task **epframes**, either in the PPS or alternatively as part of the offline analysis using the **epchain** script. Noise events flagged by **epreject** may be physically removed from the event list, considerably reducing the file size of the event set. The **epchain** task parameters `runepreject`, `sigma`, `badcolumnset`, `noiseparameters`,and `screenrejected` are available to control the behaviour of **epreject** and the subsequent event filtering. Care should be taken that parameter `screenlowthresh` is set to a sufficiently low value (120 eV) to preserve the noise-screened low energy events in the final calibrated event set. A typical **epchain** call to run **epreject** would look like this (see **epchain** task description for other epchain parameters):

```
epchain runepreject=yes screenrejected=yes screenlowthresh=120 ...
```

After running epchain, a noise screened low energy image may, e.g., be created by the following call

```
evselect expression='PHA>20 && PI in (120:200)' ...
```

Note, that events with PHA values below 20 need to be removed to create a clean image.

4 Parameters

This section documents the parameters recognized by this task (if any).

Parameter	Mand	Type	Default	Constraints
eventset	yes	dataset	rawevents.dat	

Name of input EPIC PN raw event file

set	no	dataset	odfset	

input EPIC OB ODF datset name (actual or symbolic); see description of task **epframes** for details

badcolumnset	no	dataset	badcolumns.tab	

Name of optional ascii file containing pairs of ¡ccd nr.¿ ¡bad column nr.¿ (one per line), to be omitted from the offset correction

sigma	no	real	4.0	

Sigma threshold for offset correction

noiseparameters	no	real	0.98 12 × 1.0	

Noise fraction parameters (cutoff parameter and 12 chip specific correction factors; only for expert use)

withoffsetmap	no	boolean	yes	

enables use of offset map to calculate energy shifts

withsoftflarescreening	no	boolean	no	

enables soft flare screening

softflarethreshold1	no	real	10.0	

threshold for flare screening in units of counts/0.1 s

softflarethreshold2	no	real	1.0	

threshold for flare screening

softflaresmooth	no	string	BOX	

smoothing method for flare screening

softflareenergyrange	no	list of int	40 50	

energy range for flare screening (in ADU units)

softflaresmoothparams	no	list of real	2 0 0	

smoothing parameters

5 Errors

This section documents warnings and errors generated by this task (if any). Note that warnings and errors can also be generated in the SAS infrastructure libraries, in which case they would not be documented here. Refer to the index of all errors and warnings available in the HTML version of the SAS documentation.

noEvents *(error)*
No EVENTS extension in input data set

wrongEventSet *(error)*
Single CCD event file required

wrongSubMode *(warning)*
Could not read SUBMODE keyword; FF mode assumed
corrective action: check eventset SUBMODE keyword

wrongRawCoords *(warning)*
Invalid RawX, RawY
corrective action: check event set

6 Input Files

EPIC PN raw event list (output of task **epframes**); if present the OFFSET and BADPIX extensions will be read.

7 Output Files

EPIC PN raw event list; two new columns, `OFF_COR`, containing the offset correction, and `NOISE`, which flags events as noise events are added to the event table. Column `PHA` is updated by task **epreject** (the original PHA value can be restored by adding the value in the `OFF_COR` column). If soft flare screening is enabled, an additional flag column, `FLARE`, will be created.

8 Algorithm

```
Energy scale correction:

** bin 20 ADU and 30-40 ADU images from input event file

** determine chip mask area on which to perform offset correction
     o remove offset columns (read from offset extension) from mask
     o remove 10 sigma excesses (3x3 pixel sliding box) in 30-40 ADU image
       due to soft sources or optical loading from mask

** inside mask area,
```

```
      if offset map is available and -withoffsetmask=true

        o create residual offset map by subtracting median
          from each column and row

      else

        o subtract 20 ADU calibration reference image from 20 ADU image
        o if statistically significant excesses are found on 4x1 pixel grid
          look up offset correction value corresponding to the
            observed excess in each pixel

      end if

  ** loop over input event file and

      o apply offset correction value to PHA value of each event
      o copy correction value into newly created OFF\_COR column

  Supression of detector noise:

  ** create averaged low energy spectra for each chip row

  ** determine median spectrum for each +/- 10 rows range

  ** for each row and each spectral bin determine the ratio of the
       calibration noise spectra divided by the accumulated median spectra

  ** apply cutoff value and chip specific correction factors to noise ratios

  ** loop over the event file and randomly flag the events according
       to the noise ratio of the corresponding chip row and energy channel
```

9 Comments

References