(Re)processing and Analysis Issues for ACIS Data

Glenn E. Allen
for the CXC

CIAO Workshop
October 20, 2008

Outline

- Reprocessing
  - General
  - Specific
  - Aside: custom bad-pixel files
  - Aside: custom biases
- Filtering
  - General
  - Aside: afterglows
  - Aside: horizontal streaks
  - Specific
- Analyses
  - Spatial
  - Spectral
  - Timing

Reprocessing: general

In general, users are encouraged to reprocess ACIS data in order to

- use the latest CIAO tools,
- use the latest CALDB files (e.g. an appropriate t_gain), and
- help insure that a matched set of CALDB files are used (e.g. to compute the pulse heights, ARFs and RMFs)

(see `http://asc.harvard.edu/ciao/threads/createL2/`).

Since Level 2 event-data files do not include all of the events and all of the information about each event, Level 1 files should be used for reprocessing.

Reprocessing: specific

Depending on the type of analysis being performed, there are a variety of other, specific reasons to reprocess data.

- *Pixel randomization:*
  The chip coordinates of an event are randomized within a pixel before calculating the sky coordinates of the pixel. It is possible to disable this feature (see `http://asc.harvard.edu/ciao/why/acispixrand.html`).

- *Pulse-height randomization:*
  The pulse heights are randomized within an adu. It is possible to disable this feature (see `http://asc.harvard.edu/ciao/why/acispharand.html`).

- *VFAINT background:*
  There is an algorithm that may help remove some of the background events for VFAINT mode observations (see `http://asc.harvard.edu/ciao/threads/aciscleanvf/`). However, be cautious. This algorithm can remove a significant fraction of real X-ray events.

- *Bad-pixel files:*
  The default choice to exclude data in columns adjacent to bad columns and in the columns along the quarter-chip node boundaries may adversely affect certain types of analyses. It is possible to reprocess the data using a custom bad-pixel file (see `http://asc.harvard.edu/ciao/threads/acishotpixels/` and the following aside about custom bad-pixel files).

- *Continuous-clocking mode:*
  The times in continuous-clocking mode observations may be inaccurate if the values of `RA_TARG` and `DEC_TARG` are not accurate to better than $0.5''$ (see
  `http://asc.harvard.edu/ciao/why/ccmode.html`).

- *Biases:*
  Some biases have anomalous features that may adversely affect data analysis. It is possible for the CXC, not users, to reprocess using custom bias files. See the following aside about custom bias files.

Aside: custom bad-pixel files

Presently, a suite of tools (see `acis_run_hotpix`) produces an observation-specific list of all of the bad pixels for the observation. There are several different kinds of bad pixels, including

- known bad pixels†,
- known bad columns (or portions of columns)†,
- the quarter- and mid-chip node boundaries,
- pixels with suspicious bias values†,
- pixels with bias-parity errors†,
- pixels affected by the "FEP0" problem† (see `http://asc.harvard.edu/ciao/caveats/acis_caveats.html`),
- temporarily hot pixels†, and
- pixels affected by cosmic-ray afterglows†.

By default, all of these types of bad pixels and the pixels adjacent to the ones marked with a dagger are identified as bad. Figures 1 and 2 depict nominal sets of bad pixels for the ACIS-I3 and ACIS-S3 CCDs, respectively.

While possible, it is not easy to customize a bad-pixel file. The tool `acis_build_badpix` is being modified to make it easy to do so.

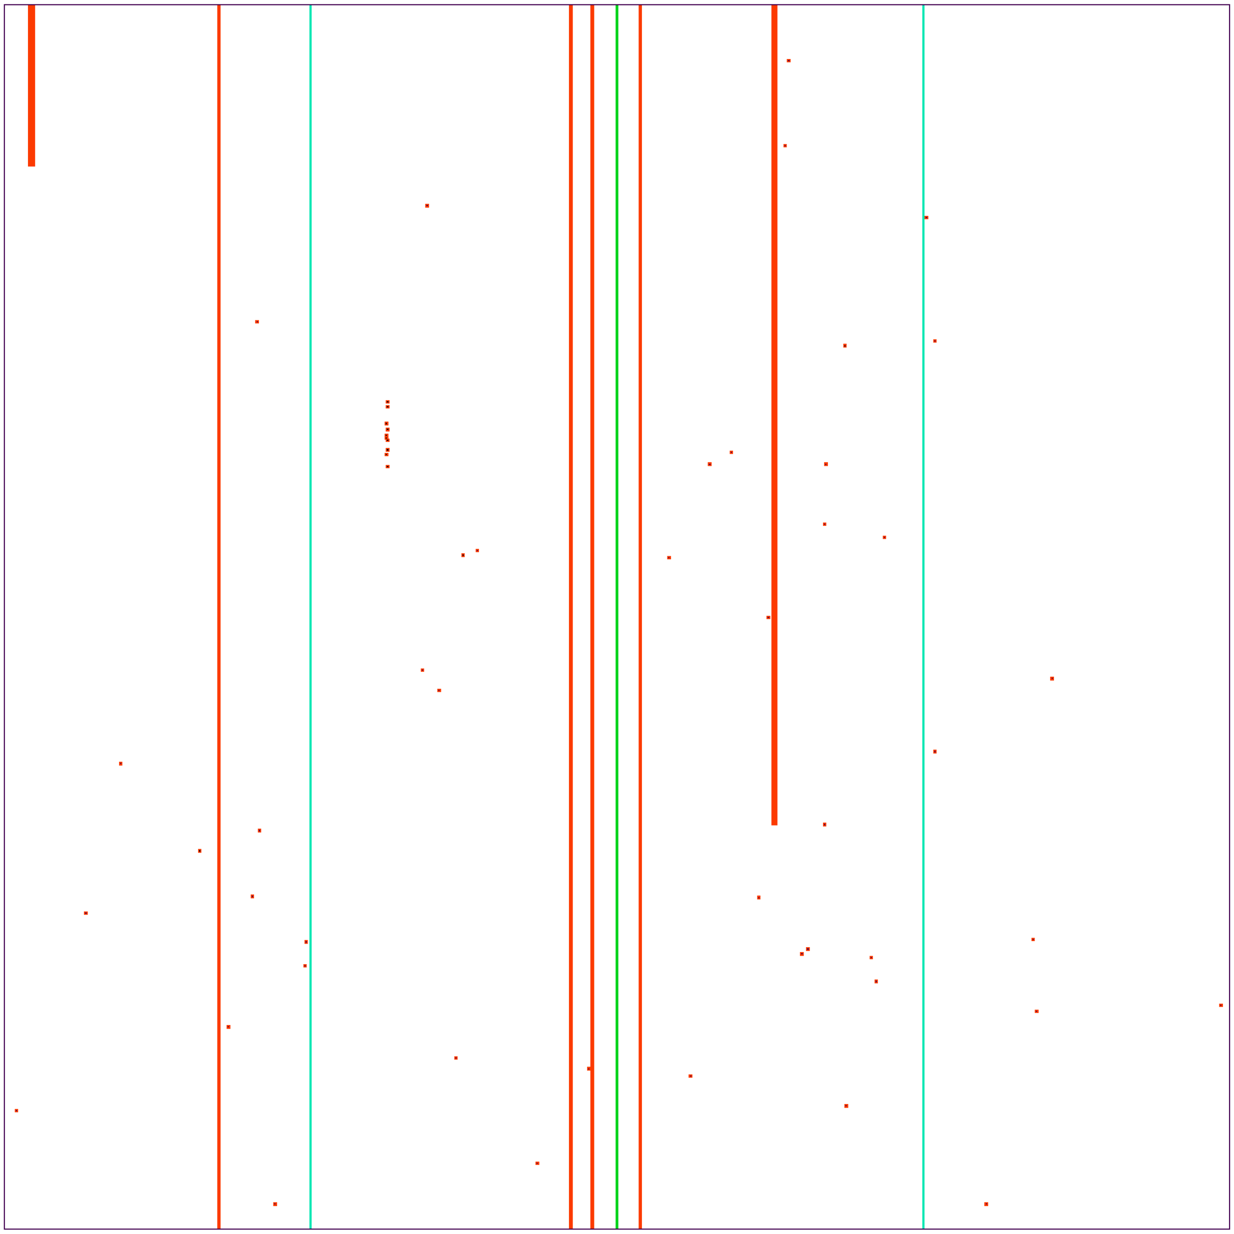

Figure 1. A nominal set of bad pixels for the CCD ACIS-I3. Different types of bad pixels are color coded. However, not every type of bad pixel has a unique color.

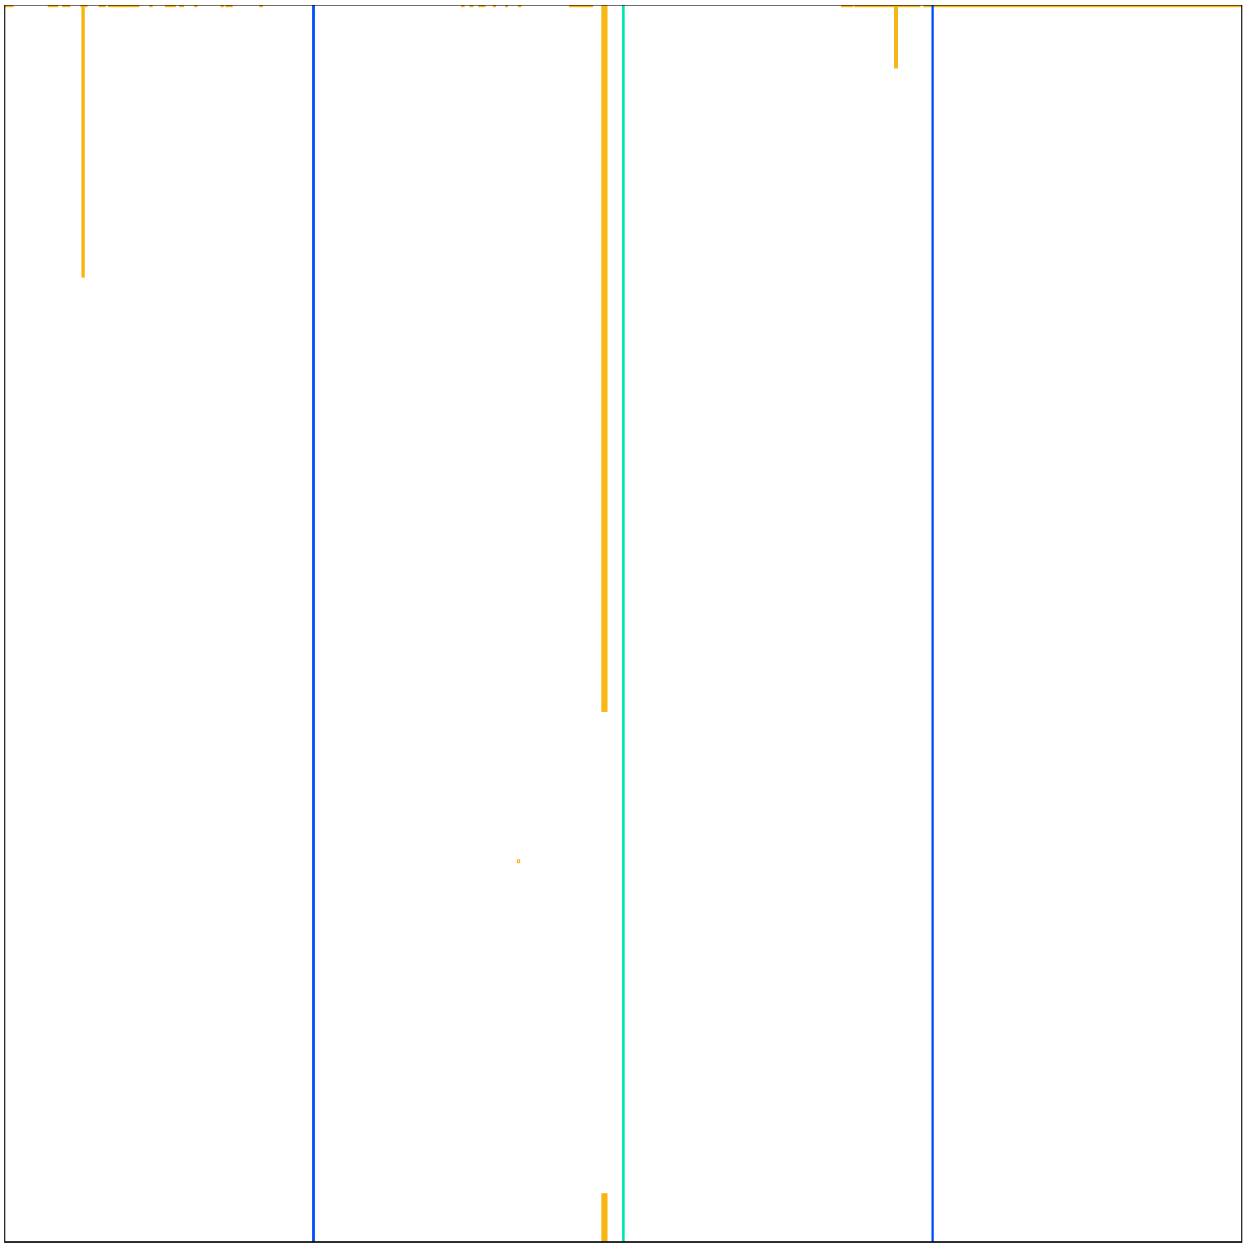

Figure 2. A nominal set of bad pixels for the CCD ACIS-S3. Different types of bad pixels are color coded. However, not every type of bad pixel has a unique color.

Aside: custom biases

Bias images represent the amount of charge on each pixel of a CCD in the absence of an X-ray event. The bias is subtracted from the charge distribution of a potential X-ray event to determine the pulse height of the event. There are several potential problems that are directly or indirectly related to bias images. These problems include

- bias-parity errors,
- the "FEP0" problem,
- excess charge from charged particle interactions,
- excess charge near the top of a CCD, and
- light leaks from fields with bright optical sources (such as Jupiter).

The standard pipeline processing automatically handles the problems associated with bias-parity errors, FEP0, excess charge near the top of a CCD and some charged-particle interactions. Furthermore, the biases for each observation are visually inspected to look for other problems. If necessary, timed-exposure mode biases are modified to remove problems that would be missed in the pipeline. This process is performed by hand. There is no CIAO tool which serves this purpose. Continous-clocking mode biases, which can have artifacts from charged-particle interactions, are not modified. If you think your data are adversely affected by a bad bias, please contact the Helpdesk.

Figures 3 to 9 show examples of nominal and unusual biases for both timed-exposure and continuous-clocking mode observations. Note that continuous-clocking mode biases are effectively one dimensional.

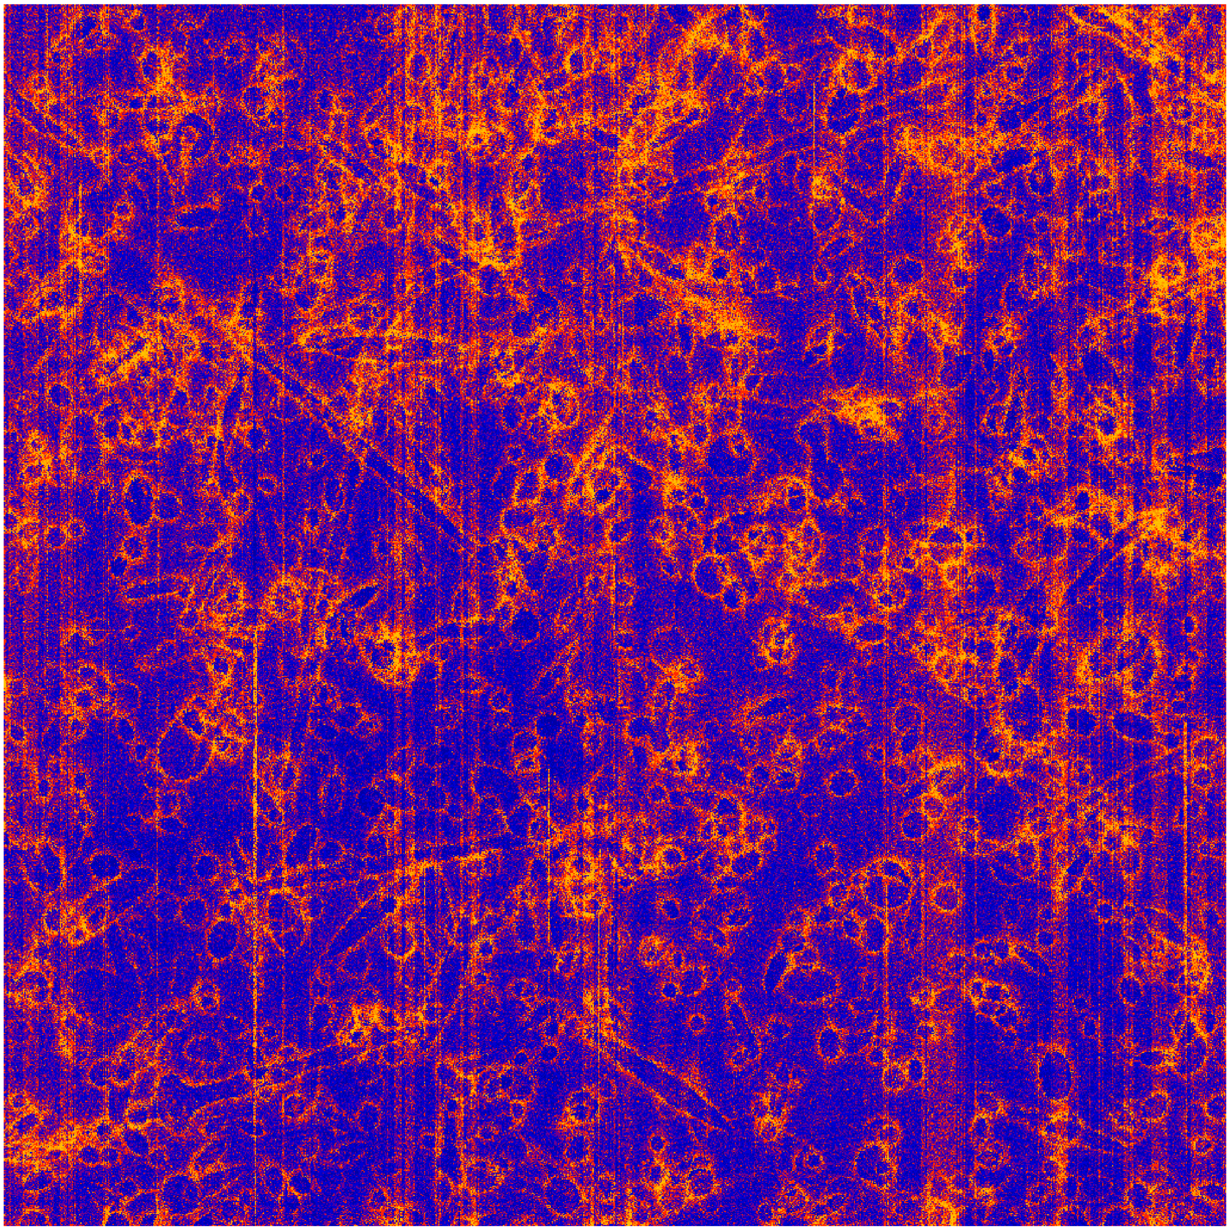

Figure 3. A nominal timed-exposure mode bias for a front-illuminated CCD (ACIS-I2). Large node-to-node offsets have been removed.

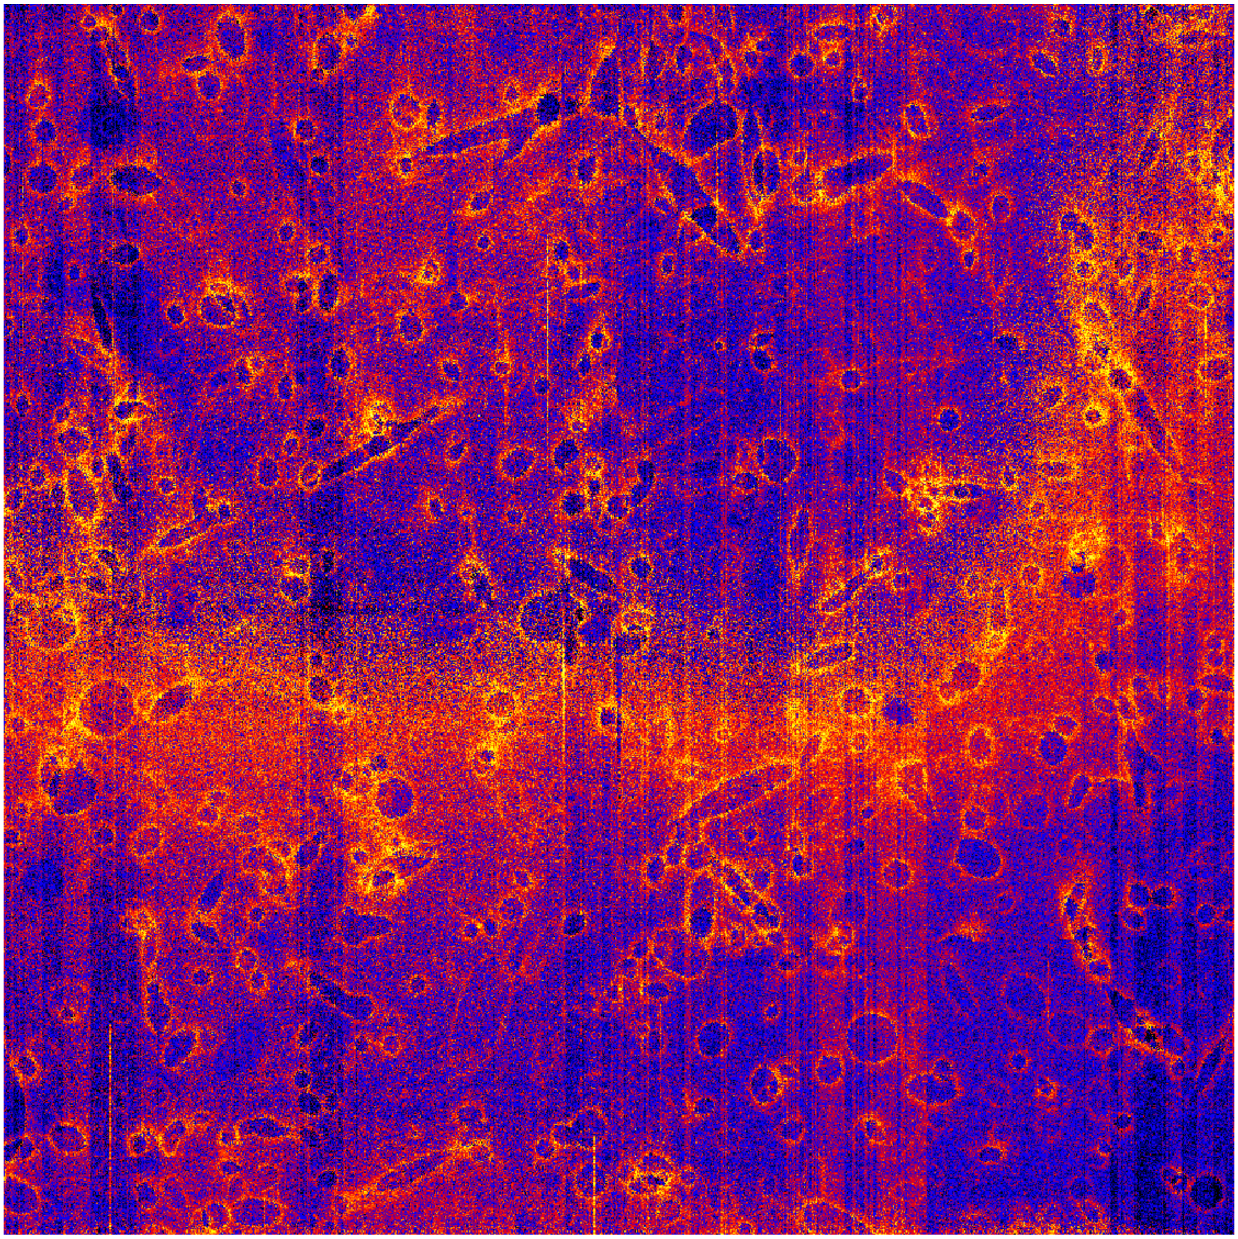

Figure 4. An unusual timed-exposure mode bias for a front-illuminated CCD (ACIS-I2). Large node-to-node offsets have been removed. An anomalous hook like feature is evident through the middle of the CCD from left to right and up the right-hand side.

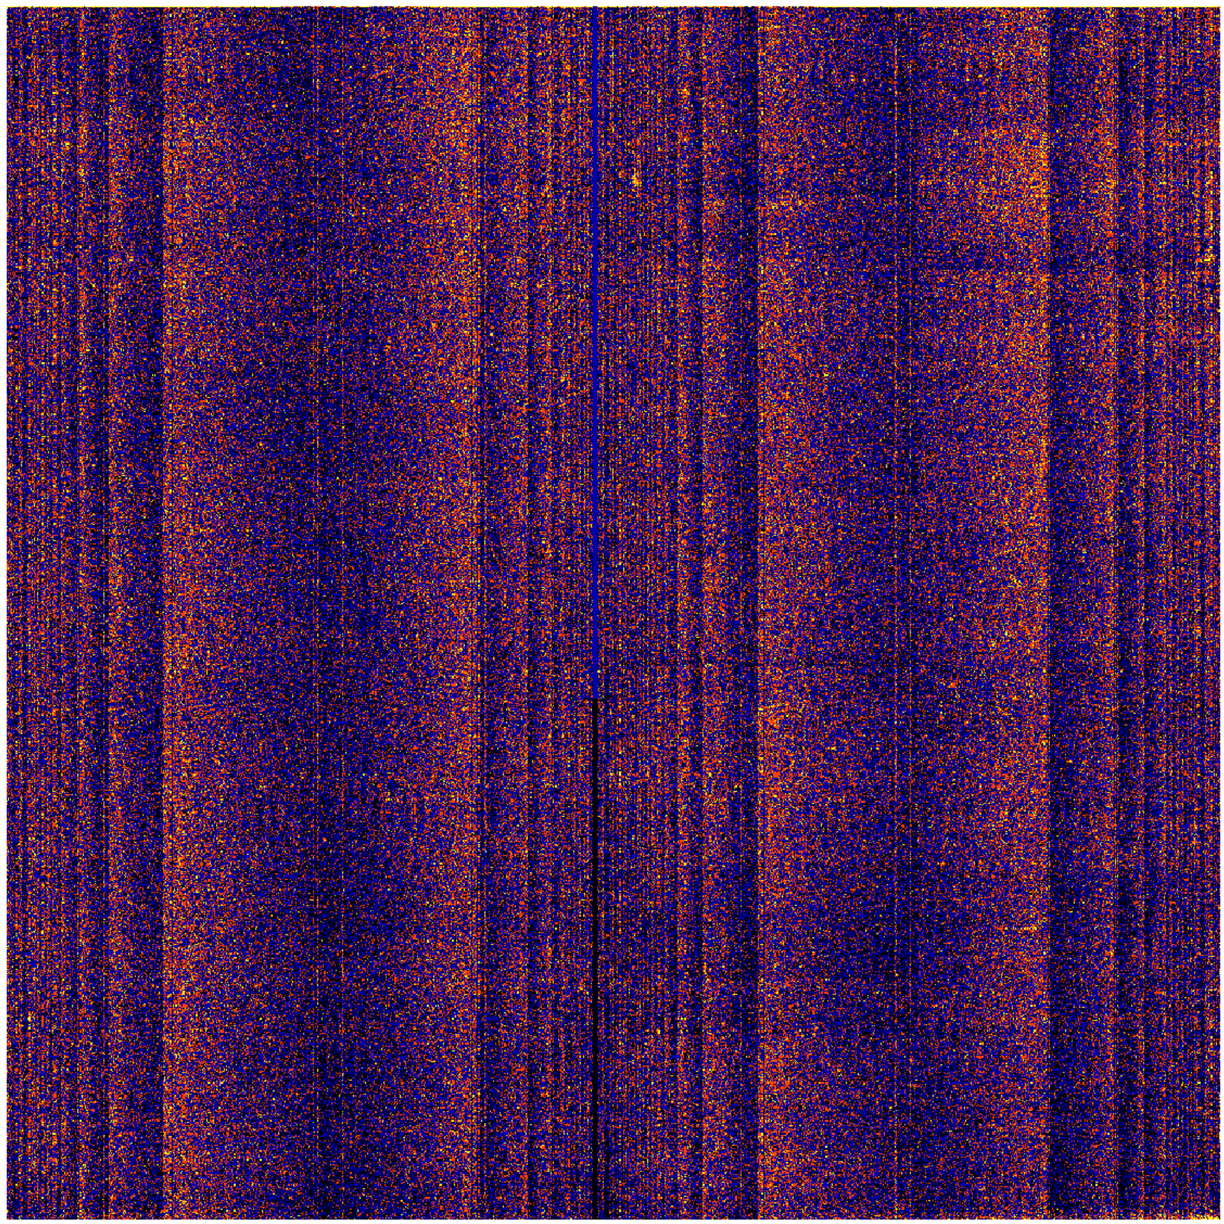

Figure 5. A nominal timed-exposure mode bias for a back-illuminated CCD (ACIS-S3). Large node-to-node offsets have been removed.

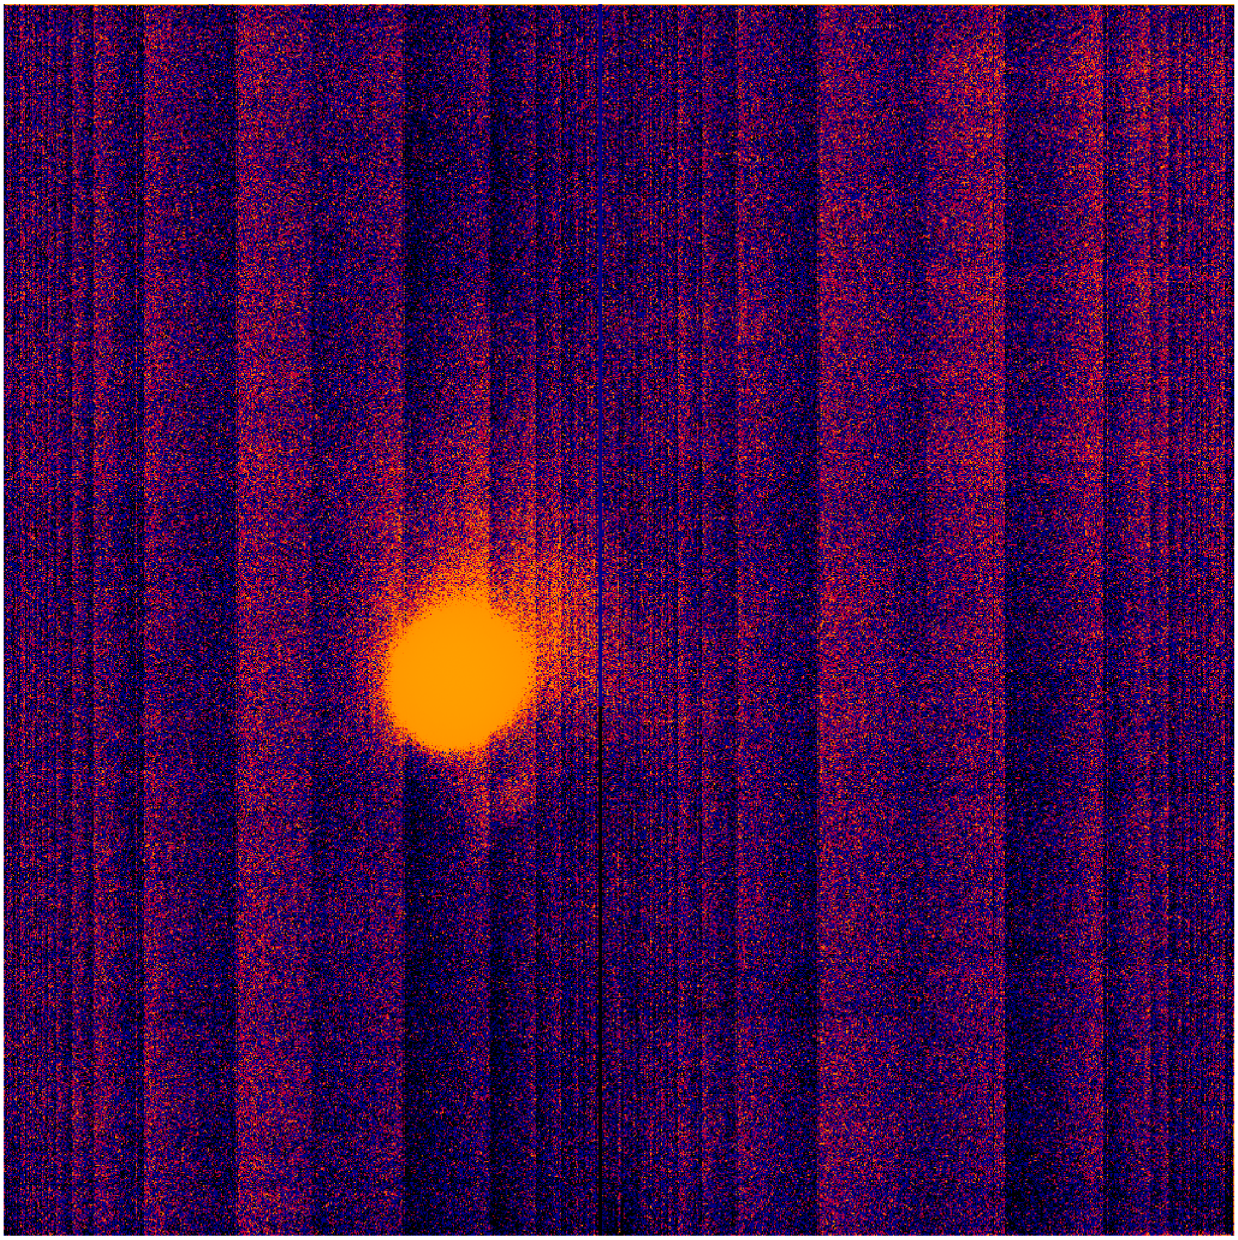

Figure 6. An unusual timed-exposure mode bias for a back-illuminated CCD (ACIS-S3). Large node-to-node offsets have been removed. The anomalous bright spot is source (Jupiter).

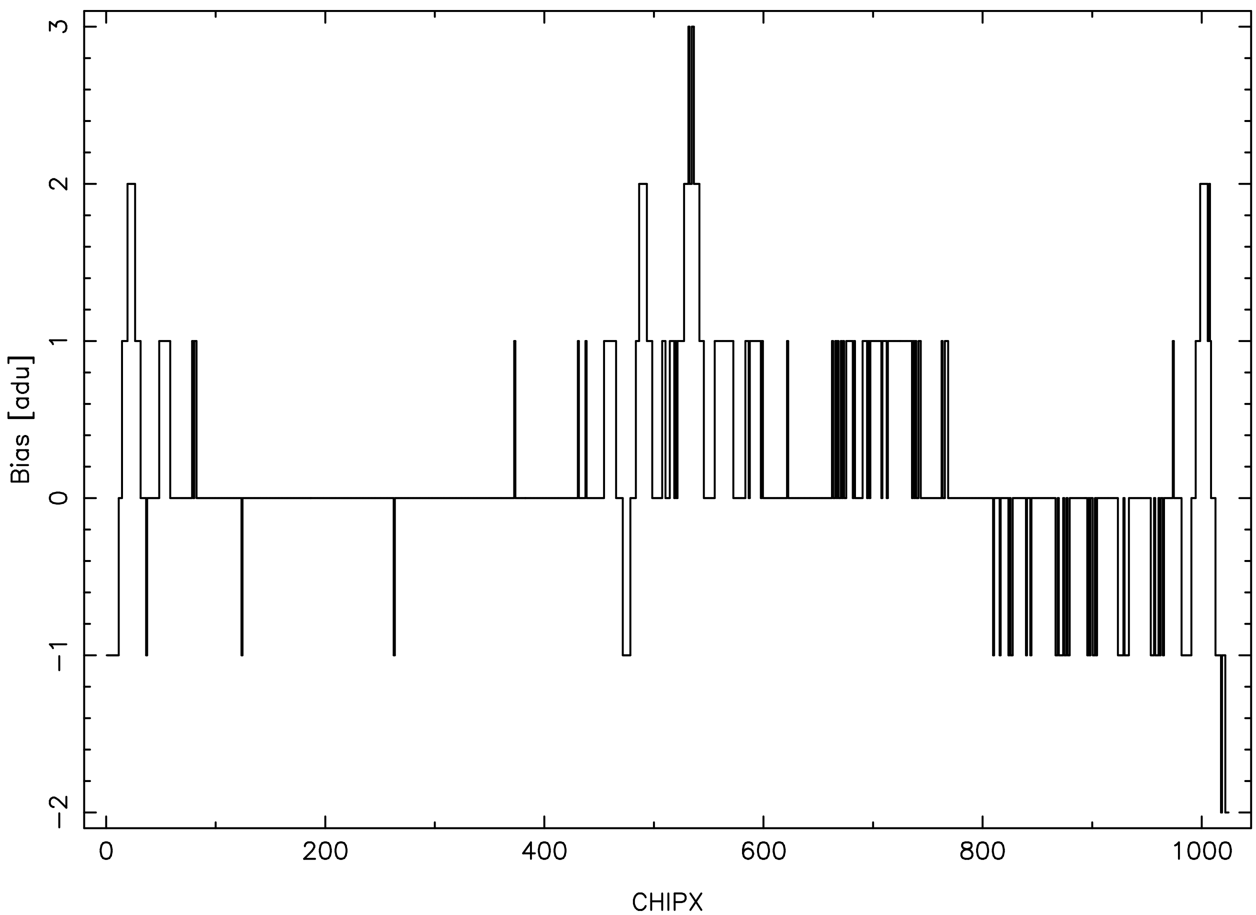

Figure 7. A nominal continuous-clocking mode bias for a front-illuminated CCD (ACIS-S2). Large node-to-node offsets have been removed.

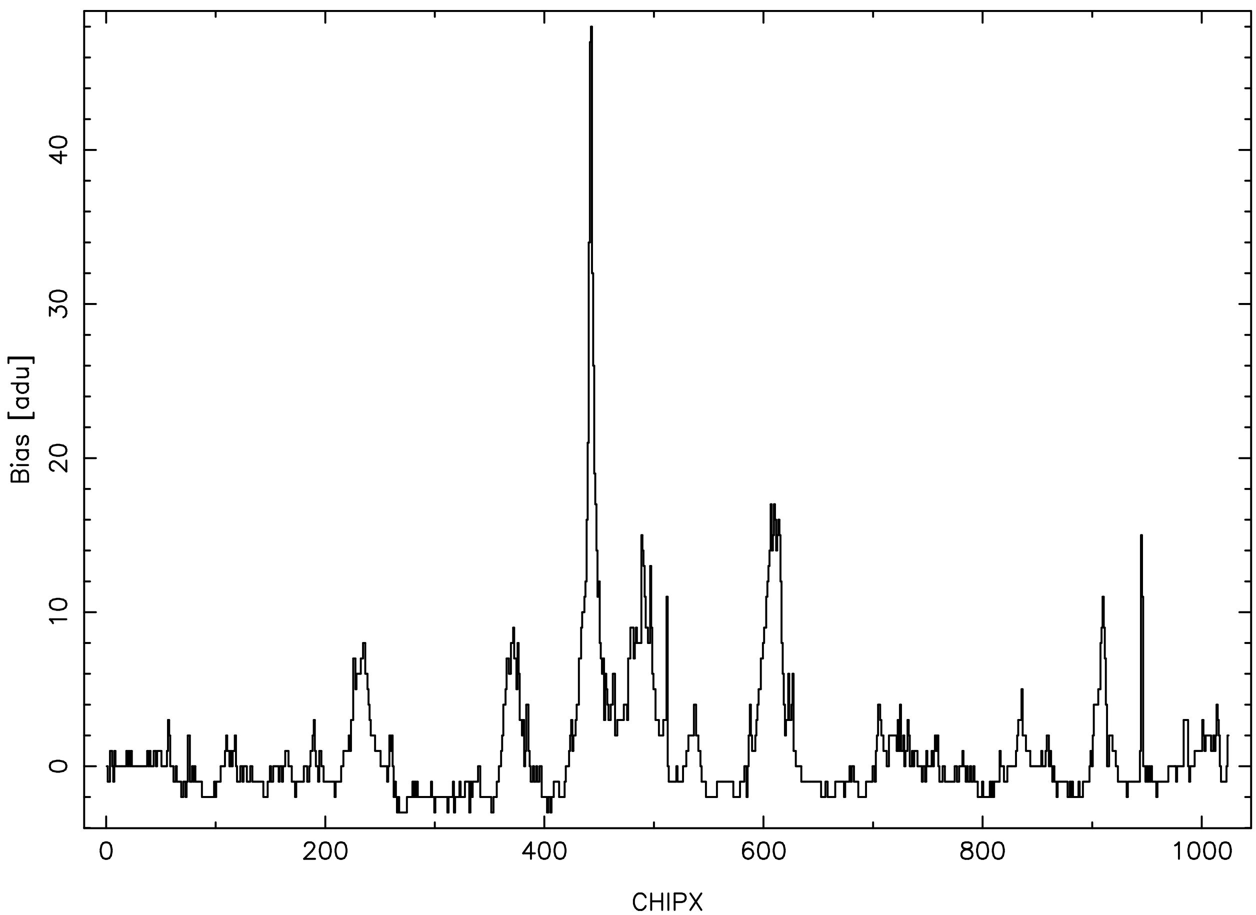

Figure 8. An “unusual” continuous-clocking mode bias for a front-illuminated CCD (ACIS-S2). Large node-to-node offsets have been removed. The anomalous peaks are caused charged particles.

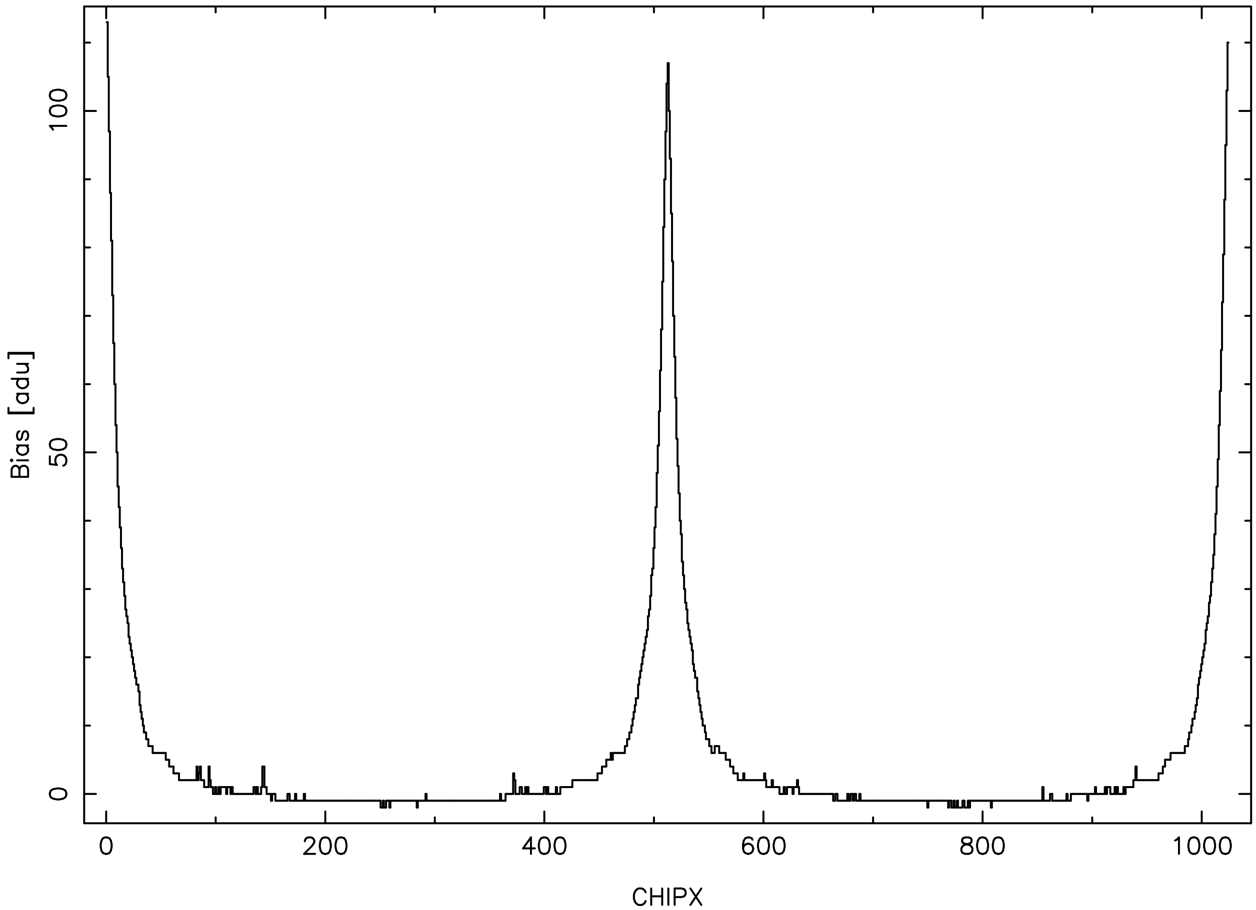

Figure 9. A nominal continuous-clocking mode bias for a back-illuminated CCD (ACIS-S3). Large node-to-node offsets have been removed. The peaks near the node boundaries at CHIP = 1, 512, 513, and 1024 are not unusual. They are due to a drop in the voltage.

Filtering: general

After reprocessing ACIS event data, the data should be filtered to exclude events that may adversely affect subsequent analyses. In general, users are encouraged to exclude events that

- are outside the good-time intervals,
- have a GRADE = 1, 5 or 7, and
- have a STATUS > 0

(see `http://asc.harvard.edu/ciao/threads/createL2/`). One or more STATUS bits are set to one if there are potential problems with an event. For example, a bit is set if an event

- is on a bad pixel or has invalid coordinates,
- has invalid pulse-height information,
- is part of a cosmic-ray afterglow (see the following aside), or
- is part of a horizontal streak (see the following aside).

For a complete list of STATUS bits in event-data files, see `http://cxc.harvard.edu/ciao/dictionary/status.html`.

Aside: afterglows

A cosmic-ray "afterglow" is produced when a large amount of charge is deposited on a CCD by a cosmic ray. Most of the charge is clocked off of the CCD in a single frame. However, a small amount can be captured in charge traps, which release the charge relatively slowly.

Events in an afterglow:

- occur on the same pixel,
- occur over a short time interval, and
- have pulse heights that usually decline with time.

The pulse height as a function of time is plotted for a selected sample of afterglows in Figure 10.

In general, afterglows do not adversely affect analyses of Chandra ACIS data. However, if you are searching for faint sources in the field of view, then some of the potential sources identified by a source detection algorithm may be associated with afterglows instead of astrophysical objects. Figure 11 illustrates this problem for a typical observation. Here, the tool `acis_run_hotpix` was used to identify potential afterglows, which were subsequently filtered from the data. The filtered data were searched for potential sources using `wavdetect`. Nearly one quarter of the potential sources are associated with afterglows that were missed by `acis_run_hotpix`. The reason the afterglows were missed is that they have too few events to be considered significant by `acis_run_hotpix`. As shown in Figure 12, the detection efficiency (i.e. the fraction of the potential afterglows that are found) declines quickly as the number of events in an afterglow drops below 10.

If your analysis would be adversely affected by afterglows, then you may examine each potential source individually to determine if it is a real astrophysical source. Alternatively, the tool `acis_detect_afterglow` can be used to remove some of the afterglows missed by `acis_run_hotpix`. However, be cautious because `acis_detect_afterglow` can remove a significant fraction of real X-ray events from bright sources. Ultimately, a new algorithm will be implemented that eliminates the problems associated with `acis_run_hotpix` and `acis_detect_afterglow`.

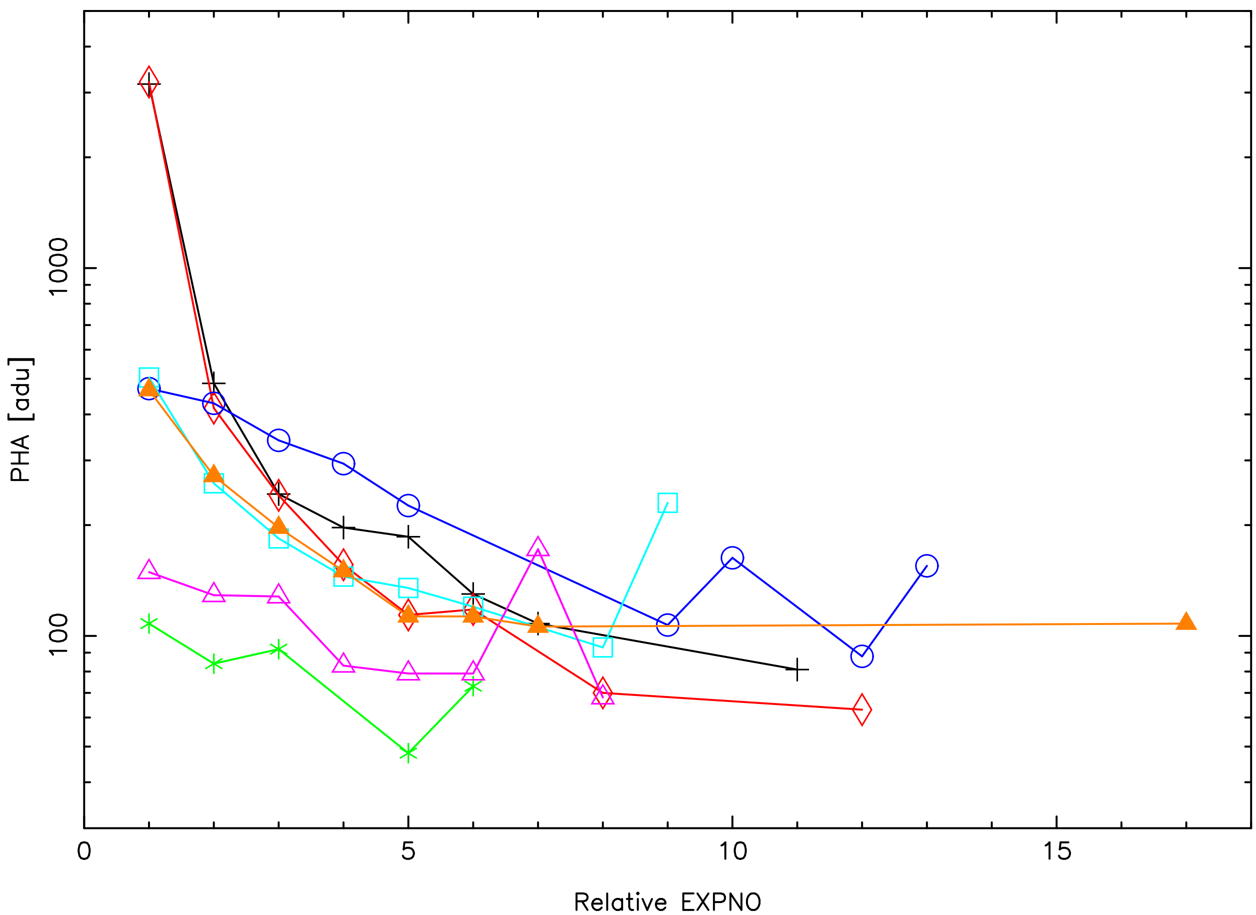

Figure 10. The pulse height (**PHA**) as a function of time (**EXPNO** or frame number) for a selected sample of afterglows. While the pulse height typically decreases from one event to the next, this trend is not universal. The pulse height can increase when the amount of charge "bottoms out" near the end of an afterglow. Also note that events do not have to occur in consecutive frames. There can be gaps.

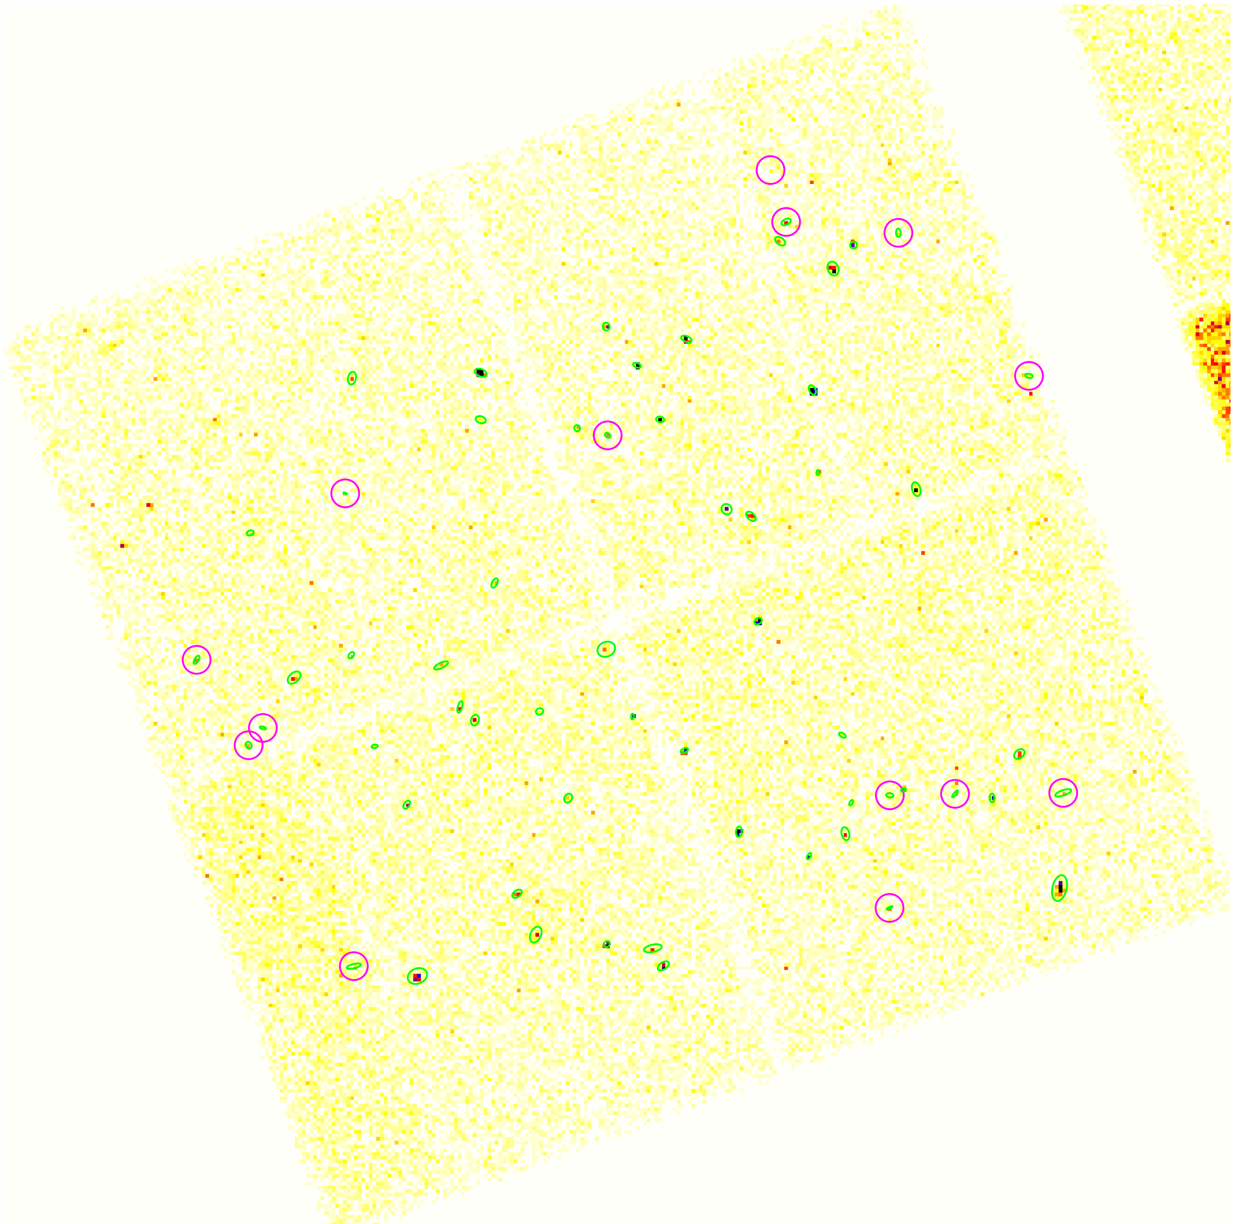

Figure 11. An ACIS-I image in which 60 sources (green ellipses) were identified as potential sources by `wavdetect`. Of the 60, 14 (or 23%) are associated with afterglows instead of astrophysical objects. The 14 are surrounded by magenta circles.

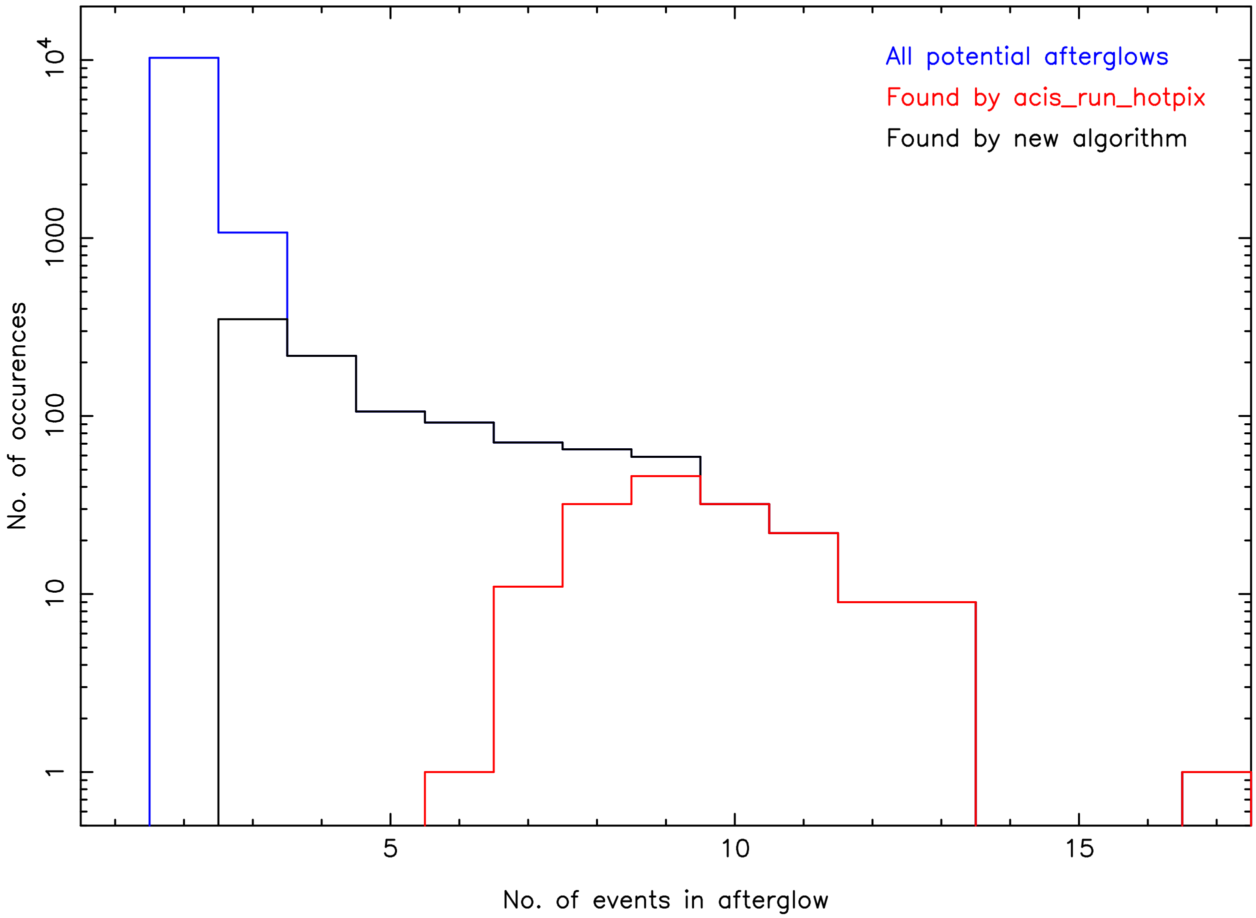

Figure 12. The afterglow-detection efficiency as a function of the number of events in an afterglow. The blue histogram is number of potential afterglows. While some of the potential afterglows with only 2 or 3 events may not be real, the afterglows with 4 or more events appear to be legitimate. A comparison of the red and blue histograms shows that the fraction of afterglows identified by **acis_run_hotpix** declines sharply as the number of events in the afterglows drops below 10. A newer, as yet unreleased, algorithm effectively finds all afterglows that have 4 or more events. (Compare the black and blue histograms.)

Aside: horizontal streaks

During the serial read-out process, electrical noise can produce spurious events that appear to be a horizontal streak. Figure 13 displays an example of such streaks, which are horizontal in chip, not sky, coordinates. Events in a horizontal streak

- occur in the same frame,
- occur on the same node, and
- occur in the same row.

The only CCD that exhibits a substantial number of horizontal streaks is ACIS-S4.

Unlike bad pixels, the pixels affected by horizontal streaks vary from one streak to another. Streaks seem to be uniformly distributed in the CHIPY direction, but in the CHIPX direction are concentrated closer to the read-out locations at CHIPX = 1, 512, 513 and 1024 (Fig. 14). For more information about streaks, see `http://asc.harvard.edu/ciao/why/destreak.html`.

By default, the tool `destreak` is used in the pipeline to remove horizontal streak events.

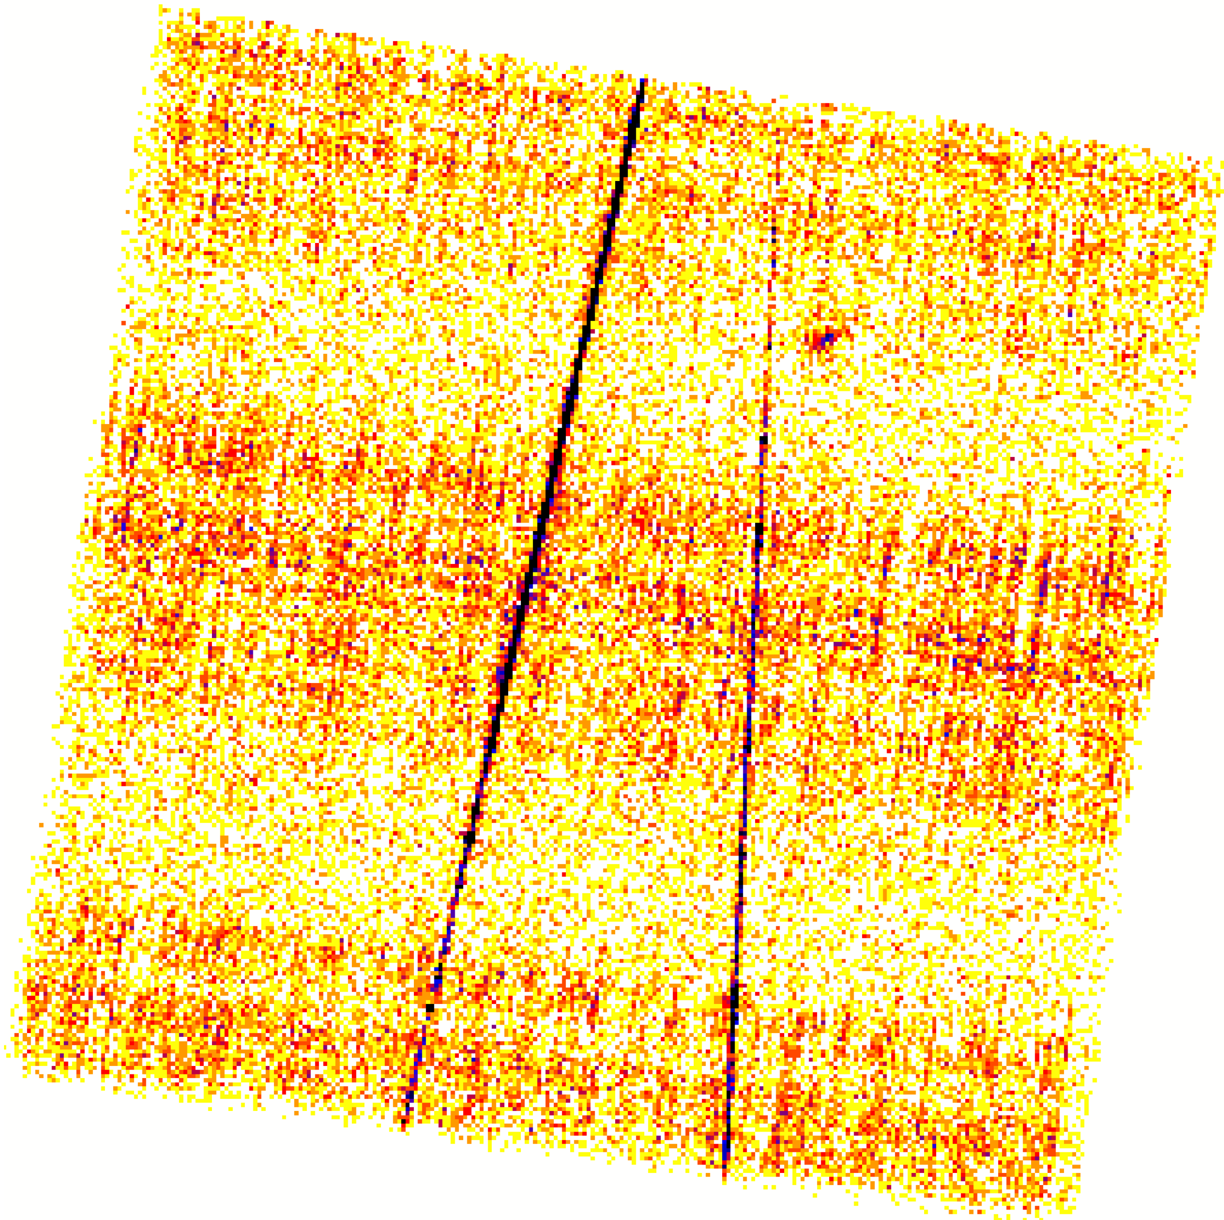

Figure 13. An image of the CCD ACIS-S4 before the horizontal streaks have been removed. The streaks are horizontal in chip, not sky, coordinates. The image also includes a source near the upper right-hand corner and portions of the HETG grating arms running through the middle of the CCD from the top to the bottom.

The CHIPX distribution

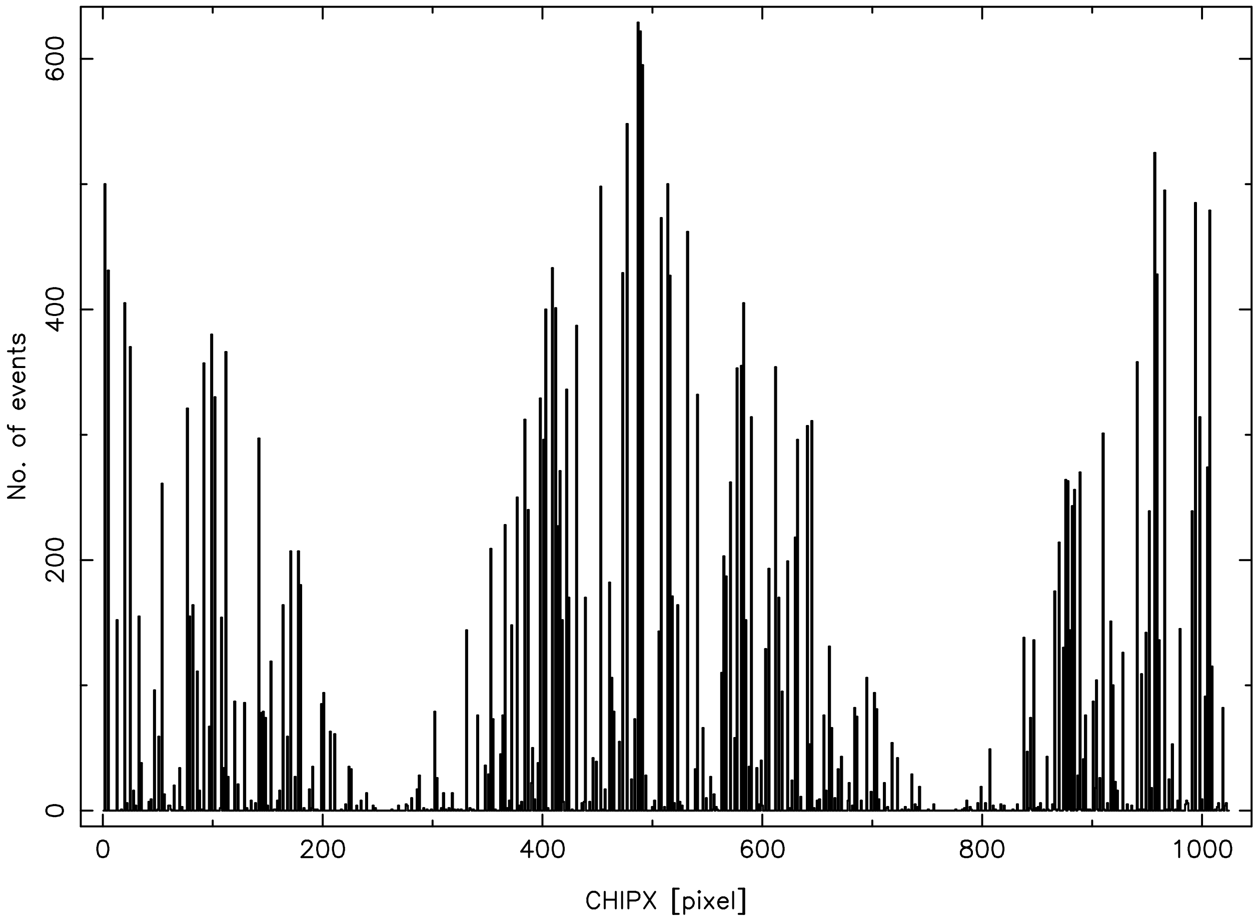

Figure 14. The CHIPX distribution of only the streak events. These events are concentrated closer to the read-out locations at CHIPX = 1, 512, 513, and 1024, than to the opposite ends of the nodes at CHIPX = 256, 257, 768, and 769, respectively.

Filtering: specific

Depending on the type of analysis being performed, there are a variety of other, specific filters that can be used.

- *Background flares:*
  It may be desirable to eliminate intervals with relatively large rates of charged-particle background events (see `http://asc.harvard.edu/ciao/threads/filter/index.html#filterback`).

- *VFAINT background:*
  If the data were obtained using the VFAINT mode, then they can be reprocessed with `acis_process_events` using `check_vf_pha=yes` (see `http://asc.harvard.edu/ciao/threads/aciscleanvf/`). However, be cautious. This algorithm can remove a significant fraction of real X-ray events.

- *Miscellaneous:*
  Other filters on TIME, location, and ENERGY may be advantageous.

Analyses

Once the data have been (re)processed and filtered, a variety of analyses can be performed. The three basic types of analyses include

- spatial,
- spectral, and
- timing.

An analysis may also involve combinations of these basic types. The following pages describe some of the issues and concerns associated with each type of analysis.

Analyses: spatial

The spatial distribution of X-ray events can be affected in several ways by the hardware and software.

- *Spatially-dependent detection efficiency:*
  The detection efficiency is not uniform across the field of view. Aside from vignetting, normal pixel-to-pixel variations, gaps between CCDs, bad pixels, user-specified windows, and the frame-store shadow influence the detection efficiency. Except for the frame-store shadow, which affects the bottom 8 rows or so of a CCD, these effects are handled automatically when an exposure-map is created.

- *Vertical streaks:*
  Since ACIS is a shutterless camera, vertical streaks are produced for bright sources during the frame-transfer process. The streaks are vertical in chip, not sky, coordinates. An image of such a streak is shown in Figure 15. The tool `acisreadcorr` can be used to reposition streak events at the location of the source. However, the tool is not designed to accurately reproduce the point-spread function.

- *Pile up:*
  If the event rate for a source is sufficiently large, the charge clouds for two or more events can overlap and be treated as a single event. As a result, the number of events identified is smaller than the number that occurred. This effect is particularly acute at the core of the point-spread function. Therefore, a point source may appear to be slightly extended since more the events are missing from the center than the edges. If the total amount

of charge in the composite event is large enough, then the event will not be telemetered. Therefore, as shown in Figure 15, a bright source may have few or no events at the core of the point-spread function. Furthermore, the combination of pile up and charge-transfer inefficiency (CTI), can cause sources to have dark "spokes" in addition to dark cores. The effects of pile up on the point-spread function are uncalibrated. For more information about pile up see `http://asc.harvard.edu/ciao/why/pileup_intro.html`.

- *Pixel randomization:*
  By default, the chip coordinates of an event are randomized within a pixel before the sky coordinates are calculated. This randomization, which can be disabled (see `http://asc.harvard.edu/ciao/why/acispixrand.html`), broadens the point-spread function.

- *Sub-pixel localization:*
  A handful of sub-pixel positioning algorithms have been developed to improve the localization of events whose charge clouds are split across 2, 3, or 4 pixels. The most sophisticated algorithm is the one published by Li, J. et al. (2004, ApJ, 610, 1204). An IDL implementation of this algorithm is available from `http://www.cis.rit.edu/people/faculty/kastner/SER/ser.html`. The optimum algorithm will be added to the tool `acis_process_events` in the future.

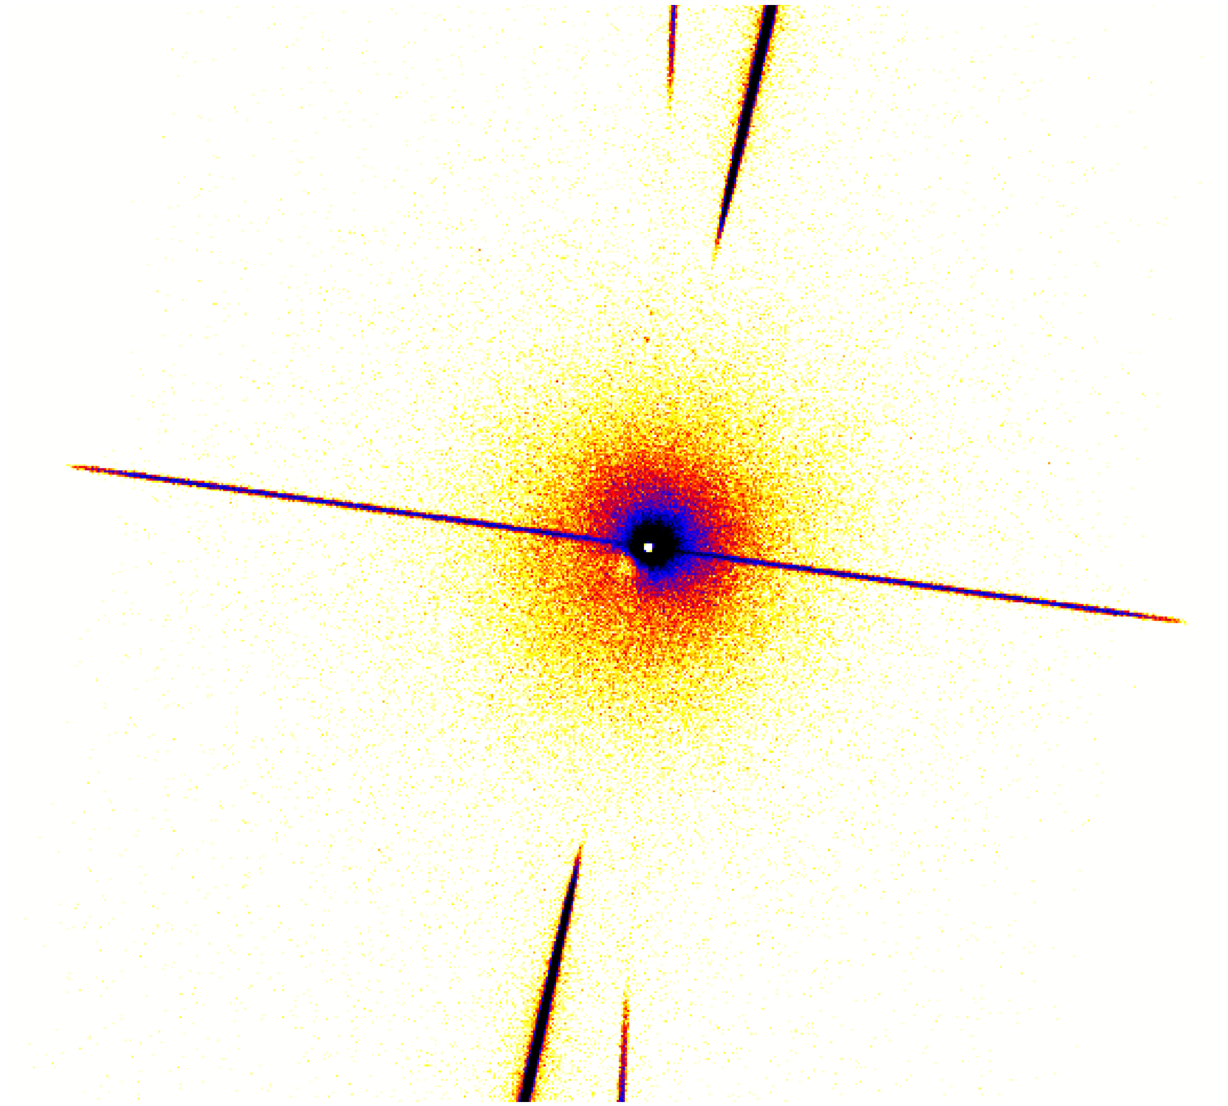

Figure 15. An example of a vertical frame-transfer streak. The streak is vertical in chip, not sky coordinates. In this image, the streak is nearly horizontal. The linear features at the top and bottom are portions of the HETG grating arms.

Analyses: spectral

The hardware and software also affect the spectral distribution of X-ray events.

- *CTI:*

  Due to charged-particle radiation, the number of charge traps in the ACIS CCDs has increased over time. As a result, the charge-transfer inefficiency (CTI) during the read-out process has increased. For an X-ray of a given energy, the mean amount of charge that reaches the read out has diminished. This apparent gain shift is calibrated and corrected using a combination of CTI and time-dependent gain files. The time-dependent gain file used in the pipeline cannot be up-to-date. Therefore, it may be necessary to reprocess the data using the latest time-dependent gain file (see `http://cxc.harvard.edu/ciao/threads/createL2/`).

  The effects of CTI also cause the distribution of the amount of charge that reaches the read out to broaden for a monoenergetic source. In other words, the energy resolution is worse. Some, but not all, of this degradation is corrected using the CTI-adjustment algorithm in `acis_process_events`. The effect of CTI on the energy resolution has been calibrated.

- *Continous-clocking mode:*

  The effective gain and spectral response of the detectors are not the same for the timed-exposure and continous-clocking modes. While the timed-exposure mode is relatively well calibrated, spectral analyses of continous-clocking mode data would benefit from an improved calibration of this mode. The calibration team is working on such a calibration.

Since the pulse-height calibration varies from one location to another on a CCD and since continous-clocking mode event data contains only one spatial coordinate (`CHIPX`), the values of `RA_TARG` and `DEC_TARG` are used to calculate the other coordinate (`CHIPY`). Therefore, these values should be as accurate as possible, preferably to better than $0.5''$. If necessary, they can be changed using **`dmhedit`** and the data can be reprocessed using **`acis_process_events`**.

- *Contamination:*
  A contaminant is being deposited on the ACIS optical-blocking filters or detectors. As a result, the detection efficiency at low energies is declining with time. This effect has been calibrated and the response tools automatically produce products that are adjusted to compensate for this effect. For more information see `http://asc.harvard.edu/ciao/why/acisqedeg.html`.

- *Dead area:*
  Some charged particles deposit a lot of charge on the front-illuminated CCDs creating "charge blooms." Since it is not possible to detect X-ray events that land within such a bloom, these regions represent temporary dead areas of the detector. As a result the detection efficiency (i.e. effective area) of the detector is diminished by a few percent. This effect has been calibrated and the response tools automatically produce products that are adjusted to compensate for this effect. For more information see `http://asc.harvard.edu/ciao/why/acisdeadarea.html`.

- *"Blank-sky" background:*
  Obtaining a representative background spectrum can be challenging for extended sources, particularly if the source covers an

entire CCD. A set of "blank-sky" data files are available for such cases (see `http://asc.harvard.edu/ciao/threads/acisbackground/`). However, be cautious. The background data may not have been obtained in the same charged-particle radiation environment as your observation and the background data may not have been processed in the same manner as your data.

For point sources, take care to avoid the wings of the point-spread function and the vertical read-out streak when selecting representative background regions.

- *Pile up:*
  The effect of pile up can significantly alter the spectrum of a source. Fewer events are detected and the spectrum will be skewed toward higher energies. For more information about pile up see `http://asc.harvard.edu/ciao/why/pileup_intro.html`.

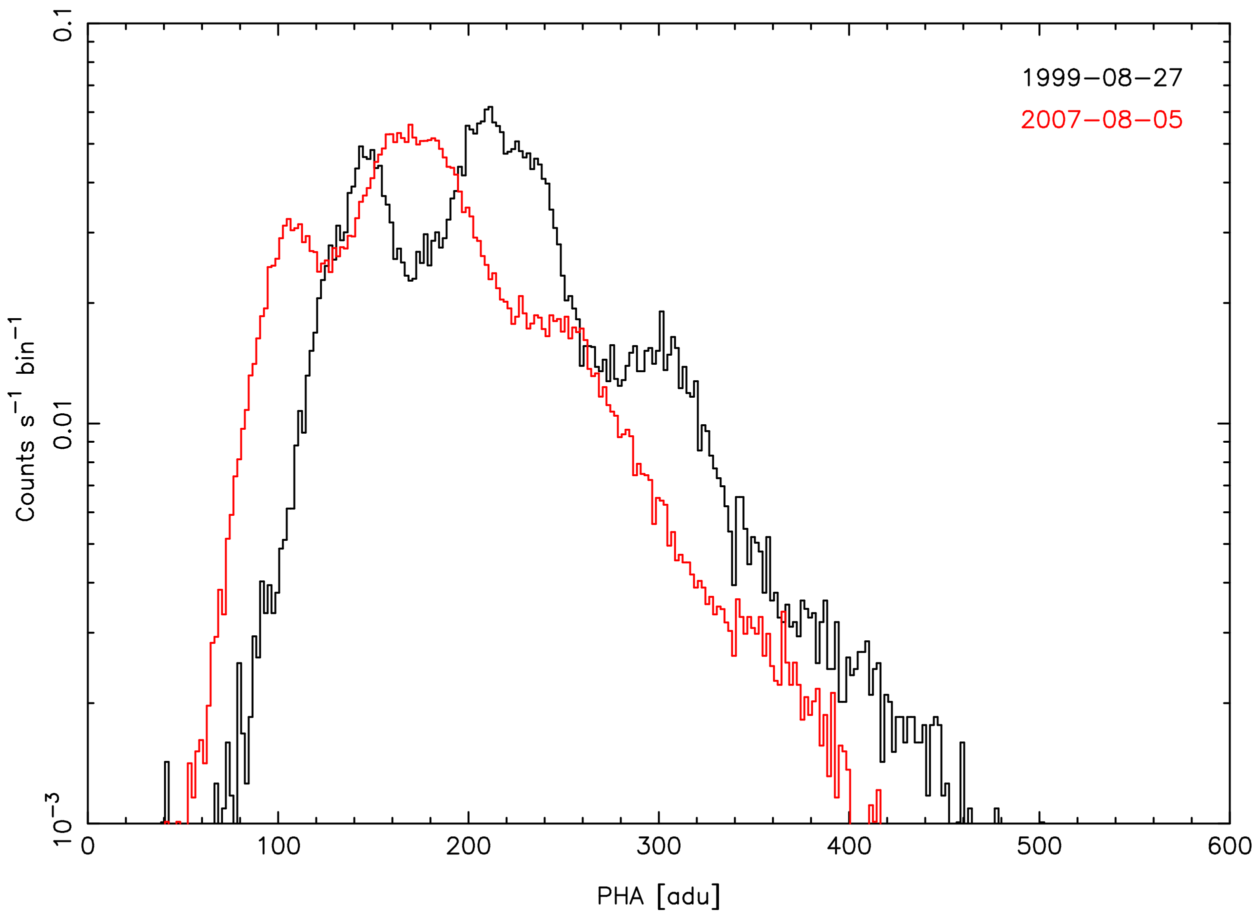

Figure 16. An example of the effects of CTI on the the spectrum of 1E0102.2-7219. Spectra from two separate observations are shown. One is relatively free from the affects of CTI (black histogram). The other has been altered significantly by these effects (red histogram). Note that the red histogram is shifted to lower pulse heights and that it has broader line features.

Analyses: timing

The hardware and software affect the timing data as follows.

- *Timed-exposure mode:*
  For timed-exposure mode observations, the relative times of arrival are accurate to within one frame time. However, in situations were there are a significant number of events in the vertical read-out streak, the times of these events can be calculated to a relative accuracy of about 40 μs (see `acisreadcorr`).

- *Continous-clocking:*
  The relative accuracy of the times of arrival of events in continous-clocking mode event data is about 0.0285 s. The absolute accuracy of the times depends on accuracy of `RA_TARG` and `DEC_TARG`. To insure that the absolute times are accurate to within about 0.00285 s, `RA_TARG` and `DEC_TARG` should be accurate to within $0.5''$. If necessary, they can be changed by using `dmhedit` and the data can be reprocessed using `acis_process_events` (see `http://asc.harvard.edu/ciao/why/ccmode.html`).

- *Barycentering:*
  The times associated with events are the times of arrival at the detector. The tool `axbary` can be used to calculate the times of arrival at the barycenter.