


Chandra X-Ray Center



CIAO Workshop Introduction to X-Ray Data Analysis

David Huenemoerder (MIT)
Randall Smith (CfA)

Abstract

We describe in general terms how flux incident on an X-Ray telescope observatory is transformed into the detected signal via the response of the instruments, and then describe the specific kinds of files and operations we use in data analysis. This will include an overview of data products, calibration files, and analysis products needed for imaging, imaging spectroscopy, dispersive spectroscopy and timing investigations.

Fundamental Response Equation

$$C_D(\sigma, h, t) = T(\sigma, t) \int d\lambda \int d\hat{p} R(\sigma, h, \lambda, \mathcal{R}(t) \cdot \hat{p}, t) S(\lambda, \hat{p})$$

Counts
[Counts]

Good Time
[s]

Response
[cm²counts photon⁻¹]

Truth
[phot cm⁻²s⁻¹Å⁻¹]

σ = Detected position (x, y pixel on the detector)

h = Detected “wavelength” (a.k.a., “Pulse Height Amplitude” or PHA)

t = Time

λ = Incident wavelength (or energy)

\hat{p} = Incident angle (or x, y on the sky)

$C_D(\sigma, h, t)$ comes from The fundamental file of X-ray astronomy, the **event list**. Each **event** has a measured

(x, y) : position on detector;

h : detector pulse height (also, PHA)

t : time of readout

Each of these is imprecisely known; uncertainties on each value are frequently correlated. Each is related to a desired physical quantity of the source.

The event's detector position is converted to a sky position using the telescope's *aspect solution*. The position is smeared by the telescope optical aberrations.

The event's pulse height (in a CCD/proportional counter/microchannel plate) is related to the incident X-ray's energy via the detector's *spectral redistribution* – another smearing.

Unless the response is near *diagonal* (or approximately a delta function), we cannot invert the fundamental response equation. Hence, we address these smearings by convolving an *assumed* model of the source through the instrument response (so-called *forward-folding*) and comparing the measured and modeled results. *Or*, we build an instrument with a more nearly diagonal response (e.g., diffraction gratings).

The Event File is your Friend!

Become intimate with it. List it. Plot it. Filter it. Slice & dice it.

```
unix% dmlist acis_evt2.fits cols
```

Columns for Table Block EVENTS					
Col	Name	Unit	Type	Range	Description
1	time	s	Real8	7.7e7:7.8e7	S/C TT corresponding to mid-exposure
2	ccd_id		Int2	0:9	CCD reporting event
3	node_id		Int2	0:3	CCD serial readout amplifier node
4	expno		Int4	0:2147483647	Exposure # of CCD frame
5	chip(chipx,chipy)	pixel	Int2	1:1024	Chip coords
6	tdet(tdety,tdety)	pixel	Int2	1:8192	ACIS tiled detector coordinates
7	det(detx,dety)	pixel	Real4	0.50:8192.50	ACIS detector coordinates
8	sky(x,y)	pixel	Real4	0.50:8192.50	sky coordinates
9	pha	adu	Int4	0:36855	total pulse height of event
10	energy	eV	Real4	0:1.e6	nominal energy of event (eV)
11	pi	chan	Int4	1:1024	pulse invariant energy of event
12	fltgrade		Int2	0:255	event grade, flight system
13	grade		Int2	0:7	binned event grade
14	status[4]		Bit(4)		event status bits

```
unix% dmlist hrc_evt2.fits cols
```

Columns for Table Block EVENTS					
Col	Name	Unit	Type	Range	Description
1	time	s	Real8	6.9e7:7.0e7	time tag of data record
2	rd(tg_r,tg_d)	deg	Real4	-2.0: 2.0	Grating angular coords
3	chip(chipx,chipy)	pixel	Int2	1:4096	Chip coords
4	tdet(tdety,tdety)	pixel	Int4	1:49368	Tdet coords
5	det(detx,dety)	pixel	Real4	0.50:65536.50	Det coords
6	sky(x,y)	pixel	Real4	0.50:65536.50	Sky coords
7	chip_id		Int2	1:3	
8	pha		Int2	0:255	
9	pi		Int2	0:255	
10	tg_m		Int2	-62:62	Diffraction order (m)
11	tg_lam	angstrom	Real4	0: 400.0	wavelength (lambda)
12	tg_mlam	angstrom	Real4	-400.0:400.0	Order times wavelength (m * lambda)
13	tg_srcid		Int2	0:32767	source ID, index from detect t
14	tg_part		Int2	0:99	HEG, MEG, LEG, HESF regions
15	tg_smap		Int2	0:32767	source map; flags for up to 10
16	status[4]		Bit(4)		event status bits

Stages of Data Products

Telemetry — (Level 0)

Very raw

(voltages, frames, encoder counts)

Physical — (Level 1)

scaled, aspect corrected

(RA,Dec.; PHA; PI; Energy)

Source dependent — (Level 1.5)

Transmission grating

(wavelengths, diffraction angles)

Filtered, concatenated, binned — (Level 2)

Good detections, good aspect, good pixels;

Statistical filters (cosmic rays, streaks, glowing pixels);

Allows construction of user-level data, from which you can make: **images, spectra, light curves, models.**

File Formats

Most X-ray astronomical data files come in “FITS” format (Flexible Image Transport System). These consist of one or more logical blocks, each of which has a “header” and a “data” component. The header is a list of strings which both describes the format of the data in the block and documents what the data contains:

```
unix> dmlist acisf00614_002N002_evt2.fits[EVENTS] header
```

```
-----
Header keys for block EVENTS
-----
```

```
...
0007 DATE          2000-12-10T17:05:29      String  Date and time of file creation
0008 DATE-OBS     2000-03-27T04:40:34      String  Date and time of observation start
0009 DATE-END     2000-03-27T16:46:34      String  Date and time of observation stop
0020 TSTART       70519234.6478440017      Real8   Observation start time
0021 TSTOP        70562794.5869369954      Real8   Observation end time
0024 MISSION      AXAF                      String  Mission
0025 TELESCOP     CHANDRA                      String  Telescope
0026 INSTRUME     ACIS                          String  Instrument
0027 DETNAM       ACIS-456789                    String  Detector
0028 GRATING      HETG                          String  Grating
0029 OBJECT       FK COM/HD 117555              String  Source name
0030 TITLE        FK COM: SHROUDED IN MYSTERY  String  Proposal title
0031 OBSERVER     Dr. Derek Buzasi              String  Principal investigator
0032 OBS_ID       614                           String  Observation id
```

The data block can either be an image (an array of number) or a table (like a spreadsheet) with a named columns (such as **time**, **energy**, **ccd_id**) and a number of rows.

```
unix> dmlist acisf00614_002N002_evt2.fits blocks
```

```
-----
Dataset: acisf00614_002N002_evt2.fits
-----
```

	Block Name	Type	Dimensions	
Block	1: PRIMARY	Null		
Block	2: EVENTS	Table	21 cols x 422577	rows
Block	3: GTI7	Table	2 cols x 1	rows
Block	4: GTI4	Table	2 cols x 1	rows
Block	5: GTI5	Table	2 cols x 1	rows
Block	6: GTI6	Table	2 cols x 2	rows
Block	7: GTI8	Table	2 cols x 2	rows
Block	8: GTI9	Table	2 cols x 1	rows
Block	9: REGION	Table	8 cols x 3	rows

```
unix> dmlist acisf00614_002N002_evt2.fits[EVENTS] cols
```

```
-----
Columns for Table Block EVENTS
-----
```

ColNo	Name	Unit	Type	Range	Null	
1	time	s	Real8	70519234.6:70562794.5	-	time tag of data reco
2	expno		Int4	0:2147483647	-	
3	rd(tg_r,tg_d)	deg	Real4	-2.0: 2.0	-	Grating angular coo
4	chip(chipx,chipy)	pixel	Int2	1:1024	-	Chip coords
5	tdet(tdetx,tdety)	pixel	Int2	1:8192	-	Tdet coords
6	det(detx,dety)	pixel	Real4	0.50: 8192.50	-	Det coords
7	sky(x,y)	pixel	Real4	0.50: 8192.50	-	Sky coords
8	ccd_id		Int2	0:9	-	
9	pha		Int4	0:36855	-	
10	pi		Int2	1:1024	-	
11	energy		Real4	0: 1000000.0	-	
12	grade		Int2	0:7	-	
...						

To Get What You Want...

All that stands between you and Truth, $S(\lambda, \hat{p})$, is the *response*, which depends upon the three critical items:

1. Calibration,
2. **Calibration,**
3. *Calibration!*

(The rest is “just” software, organization, and analysis.)

Response is comprised of:

Mirror area (Effective Area)

Mirror point-spread-function (PSF)

Grating diffraction efficiency

Grating line-spread-function (LSF)

Detector quantum efficiency (QE)

Detector redistribution function (RMF)

Time-dependent coordinate transformations ($\mathcal{R} \cdot \hat{p}$)
(Aspect solution, “dither”)

The Calibration Database (CALDB)

The *CALDB* is a directory tree which contains Basic Calibration Files (BCF) and Calibration Product Files (CPF; derived from BCF) in FITS format with descriptive headers. These are used by the analysis software in order to create response matrices, telescope vignetting functions, and the like. *The CALDB is updated as calibrations are improved or revised!*

```
/usr/local/data/caldb/data:
  chandra
  rosat
  xmm
/usr/local/data/caldb/data/chandra:
  acis
  ephin
  hrc
  pcad
  sim
  tel
/usr/local/data/caldb/data/chandra/acis:
  bcf
  cpf
  index
/usr/local/data/caldb/data/chandra/acis/cpf:
  2dpsf
  fefs
  osip
  rmf
/usr/local/data/caldb/data/chandra/acis/bcf:
  badpix
  evtsplt
  gain
  grade
  gtilim
  qe
/usr/local/data/caldb/data/chandra/acis/bcf/qe:
  acisD1996-11-01qeuN0001.fits
  acisD1999-09-16qeuN0001.fits
  acisD1999-09-16qeuN0002.fits
  acisD2000-01-29qeuN0001.fits
  acisD1997-04-17qeN0003.fits
```

Analysis Preparation

Know your data —

Were “automatic” calibrations accurate?

View sky and detector coordinate images;

Examine PHA vs. time; Examine aspect, Good Time Intervals (GTI), lightcurve, aspect solution ...

Check the V&V report for issues;

Check the Chandra web pages for caveats or bugs.

Check calibration files —


Know your CALDB

Version, date, revision; do they match data products? Is re-processing needed?

Prepare Binned Observations —

What type of analysis is to be done?

Primary considerations:



Secondary considerations:

Spectral resolution;

Spatial resolution;

Source extent;

Temporal resolution;

Signal quality.

Bin the Data According to Analysis Type

Spatial: Use `dmcopy` to make images;

- Either point or extended source;
- Usually on the **ACIS-I**, **ACIS-S3**, or **HRC-I**.

Spectral: Use `dmextract` or `tgextract` to make spectra;

- Low-resolution ($E/\Delta E \sim 10$): **ACIS-I**, **ACIS-S3**
- High-resolution ($\lambda/\Delta\lambda \sim 1000$): **HETGS**, **LETGS**


Temporal: Use `lightcurve` to bin events;

- Low-resolution: **ACIS-TE**, **HRC-I**
- High-resolution: **ACIS-CC**, **HRC-S**


Type of analysis \implies Types of **RESPONSES** required.

Data Screening Examples

In this (early) dataset not all the bad columns were removed from the Level 2 event (`*_evt2.fits`) file:



$sky(x, y)$




$detector(x, y)$

In this case, it is much easier to see the problem by viewing the file in detector (chip) coordinates, rather than sky coordinates. (Note: the command is:

```
ds9 obsid_evt2.fits -bin cols detx dety.)
```


For a start, good files to examine are the `*_evt2.fits`, `*_pha2.fits`, `*_asol1.fits`. This can be easily done with the `firstlook` or `prism` tools, along with the `oif.fits` file for an observation. (Generally, the column names are recognizable quantities.)

Cas A, ACIS



This image looks pretty good; no bad pixels can be seen. There may be a bad column or two, but the faint streaks could also be due to events exposed during the frame-shift.

Checking the light curve for your observation is vital, since this is not in general checked by either the pipeline or in V&V. Solar flares can easily cause huge temporary flares in count rates. For point sources, this may not be important, but for extended sources the data taken during the flare should probably be removed.



The [analyze_ltcrv.sl](#) script attempts to identify flares in an automatic fashion, as shown in this light curve for a 3C 273 observation.

Chandra dithers about the pointing direction, in order to:

Fill chip gaps

Smooth over pixel-to-pixel variations

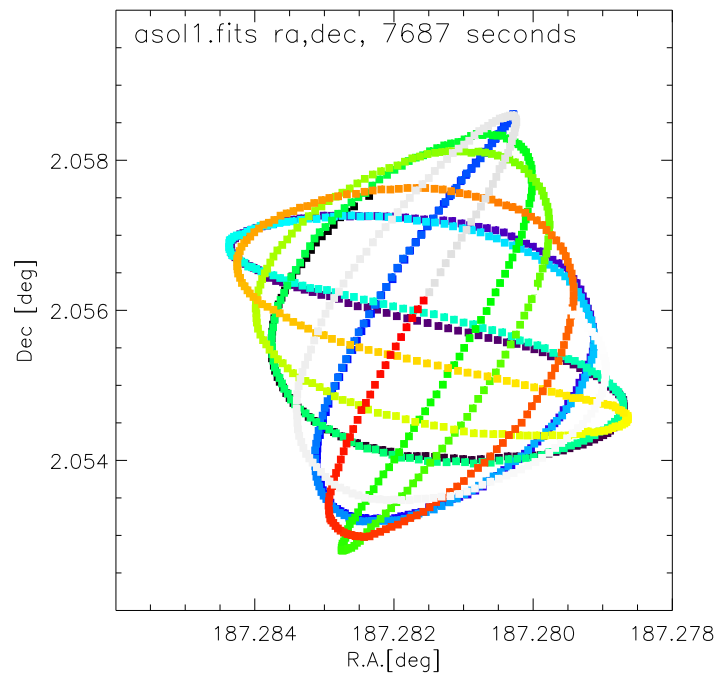
Limit radiation damage to pixels from ultra-bright sources

Some effects of dither:

Responses become observation-specific;

Calibration at the pixel scale is not needed.

Smears bad pixels' affects over a “dither patterns” in sky projection



The aspect information is in the `pcad*_asol1.fits` file (in this case, `pcadf063902942N002_asol1.fits`.) In this example, no wild excursions are present, so no bad-aspect times need to be excluded.

To reprocess, or not to reprocess. . .

Data centers routinely re-process as the instrument calibration is improved and as software is improved. Each Chandra dataset contains all the datafiles from Levels 1, 1.5 (if applicable), and 2. When deciding whether or not to reprocess your data, ask the following questions:

- Do the Level 2 images, lightcurves, spectra, and aspect solution appear reasonable? Are there obvious outliers that are not excluded? If so, reprocess from the Level 1 data.
- Has there been a new **CALDB** released with relevant updates since these data were processed? (This can be checked on the CXC main page, <http://asc.harvard.edu/>.)

If you decide reprocessing is required, check out the Chandra threads at

http://cxc.harvard.edu/ciao/documents_threads.html

to learn how. In many cases, reprocessing from Level 1 is not required; instead, the level 2 can be filtered as necessary to “destreak” the data, to detect sources, remove noisier pixels, and so on.

Most Wanted Files

If reprocessing from Level 1, you mainly need:

`*_evt1.fits`: Level 1 event list;

`pcad*_asol1.fits`: Aspect solution;

`*bpix1.fits`: Bad pixel list.

...from which you make a Level 2 event list.

If as delivered Level 2 files are adequate, you mainly need:

`*_evt2.fits`: Level 2 event list;

`pcad*_asol1.fits`: Aspect solution;

`*pha2.fits`: Binned spectrum (if a grating observation)


Spatial Response: EXPOSURE MAP

The *Exposure Map*, $E(\Delta h, \lambda, \hat{p})$, retains spatial information at the expense of spectral. It has units of $[\text{cm}^2 \text{s counts photon}^{-1}]$.

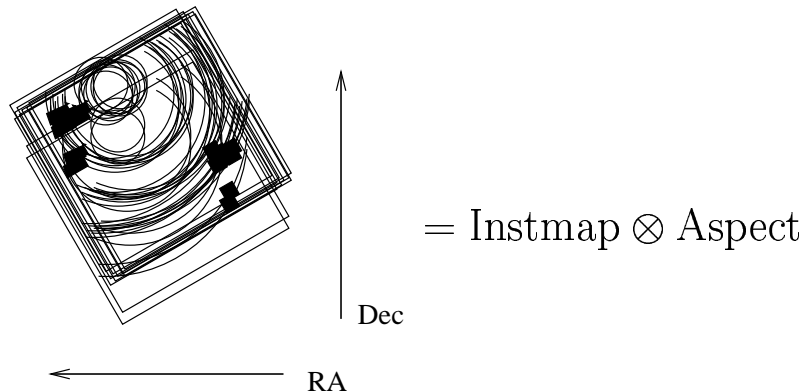
$$\int d\lambda S(\lambda, \hat{p}) \approx \frac{C(\Delta h, \hat{p})}{E(\Delta h, \lambda, \hat{p})}$$

C is the observed counts per spatial bin in a pulse-height bin. S is the source flux, with units of $[\text{photons cm}^{-2} \text{s}^{-1} \text{\AA}^{-1}]$.

Instrument Map — efficiency calibration information, band integrated. (create with `mkinstmap`)



Exposure Map — applies telescope aspect history and coordinate transformations (= area \times time). (create with `mkexpmap`).



Spectral Response, Low Resolution

The spectral response retains spectral information at the expense of spatial.

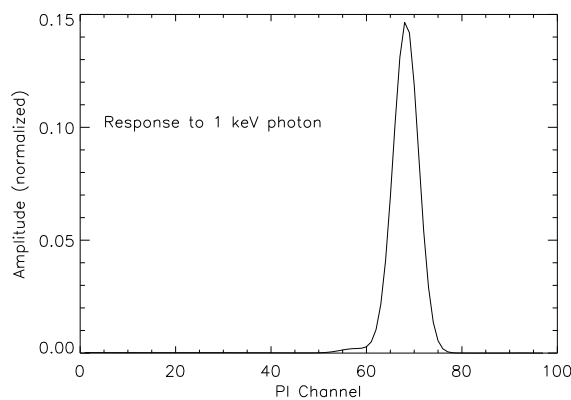
$$C_{\Omega}(h) = \tau_{\text{eff}} \int d\lambda D_R(h, \lambda) A_{\Omega}(\lambda, \hat{q}) S(\lambda)$$

for a point source at location \hat{q} for region Ω .


$D_R(h, \lambda)$ = Redistribution Matrix Function (or File), also known as the *RMF*; a probability. (create with `mkrmf`)

$A_{\Omega}(\lambda, \hat{q})$ = Auxiliary Response Function (or File), also known as the *ARF*; an area times efficiency, or *effective area*. This includes the detector quantum efficiency and mirror area. (create with `mkarf`)

τ_{eff} = effective exposure time.



$D(PI, \lambda_0)$



$A_{\Omega}(E, \hat{q})$

Imaging Spectroscopy: some details

PHA file Pulse-height spectra (of either **PI** or **PHA**) can be made with **dmextract**.

Spectral Modeling Once you have a **PHA** file, **RMF**, and **ARF** for your source, you can use Sherpa (or **XSPEC**, **SPEX**, OR **ISIS**, . . .) to create a spectral model, fold it through the **RMF/ARF**, and compare the result to the **PHA** file.

Backgrounds: Especially for faint sources, a “background” region should be selected and a background **PHA** spectrum created. This can then be subtracted from the source **PHA**, or modeled separately and summed with the source model.

Extended Sources: For **ACIS**, both the **RMF** and **ARF** are somewhat position-dependent. Need and methods for creating “averaged” **ARF** and **RMF** are in progress.

Pile-Up CCDs are subject to pile-up: more than one photon strikes the same pixel in the same frame. A 6 keV and 1 keV photons in the same pixel and frame, for example, may appear to be a single 7 keV photon. Pile-up percentages can be calculated from the source’s flux and spectrum; if more than 2-3% of photons are “piled-up,” care should be used. **ISIS** (<http://space.mit.edu/CXC/ISIS>) and Sherpa have a pile-up model for these cases. (It is a non-linear model in the source spectrum.)


PHA, Energy & PI: There is a non-linear relation between *ENERGY* and PHA, and a linear relationship between PI and *ENERGY*. The event file has an **ENERGY** column (units of eV) that can be used for filtering (remember that this is a *smear*ed energy!)

Grouping: Events must be binned before the spectrum is fit. This binning does **not** have to have constant spacing, but bins can be grouped together in order to increase the number of counts per bin.

Fitting: — χ^2 statistics systematically underestimate flux for low-count data, because errors are underestimated.

— Use χ^2 if errors are Gaussian distributed ($\gtrsim 10$ counts/bin).

— For low-count data, either (1) bin the data, (2) use an approximate error estimate, such as χ -Gehrels, or (3) use a Bayesian method, such as the Cash statistic.



A 20,000 count spectrum from the SNR G21.5

Spectral Response, Dispersive


- Spatial & spectral information is tightly coupled;
- Imaging spectral (“PHA”) & spatial information are tightly coupled.

$$C_{\Omega_i}^m(h) = \tau_{\text{eff}} \int d\lambda G_{\Omega_i}^m(\lambda, \hat{q}) A_m(h, \lambda) S(\lambda)$$


$G_{\Omega_i}^m(\lambda, \hat{q})$ = Grating RMF, diffraction order m , which describes the line-spread-function (LSF) (a redistribution from λ into spatial bin Ω_i) (make with [mkgrmf](#), or get off-the-shelf from CALDB.)

$A_m(h, \lambda)$ = Grating ARF, for order m . This includes grating efficiencies as well as the mirror’s and detector’s. (make with [mkgarf](#))

LETG/HRC-S Observation of NGC6624



Capella with HETG/ACIS-S



More about High Resolution Spectra

Large wavelength range: 1.2-170Å, over two orders of magnitude. Wavelength as the natural unit, since the gratings disperse linearly in that unit.

Coordinates: The zero-order centroid is the origin of the diffraction coordinate system. Hence, we find the zero-order image, measuring each event's position relative to it, and use the dispersion relation to calculate $m\lambda$, the order \times wavelength. If ACIS is the detector, the CCD resolution can be used to distinguish between different orders; on the HRC, this must be modeled.

Spatial-spectral coupling: The spatial and spectral elements are tightly coupled. If the zero-order image is slightly displaced (as can easily happen with heavily piled-up sources), the \pm order wavelengths will be offset from each other, and the order-sorting will also be compromised. (If this occurs, you may wish to reprocess with a new position measured by hand.)

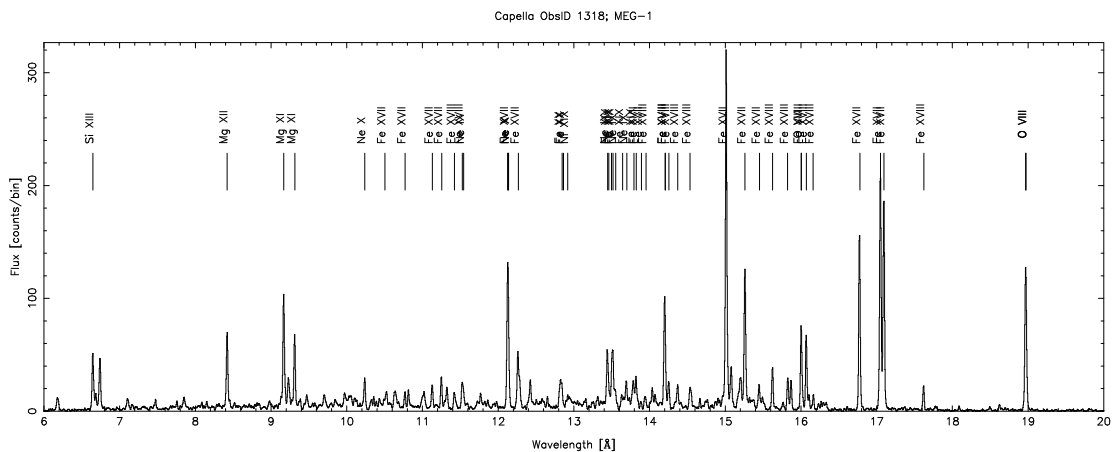

Grating PHA For ACIS grating data, 12 spectra are created ($\pm 1, 2, 3$ orders for both the HEG and MEG). For HRC data, 2 spectra are created ($\pm \sum_n i$ orders).

Grating ARF This can be created for each order using `mkgarf`; it gives the effective area for each order.

Grating RMF Optional; this is the Line Spread Function (LSF) for grating data, but this can be modeled using gaussians. A custom grating RMF (for non-standard grids or orders) can be made with [mkgrmf](#).

Fitting: In ISIS or Sherpa, single or multiple orders can be fit, with or without ARFs and RMFs. (They should both be used if there are instrumental spectral edges in a source feature.)

Capella, ObsID 1318



Response, Temporal

— spectrum, source counts, or image vs time bin.

Need exposure map or ARF vs time, or assurance that the mean is appropriate. If necessary, integrate:

$$C(\Delta h, \Delta t) = \int dt \int dh C(h, t) = \dots$$



Chandra's dither will smooth out most variations in the effective area, and so a time-resolved ARF will likely not be needed (unless the source is on a chip gap and has variability on the scale of the dither period (~ 1 ksec)).

Currently, CIAO has few analysis tools for temporal data. The main tool is **lightcurve**, which can be used to bin event data by exposure time per bin, counts per bin, or rate. **lightcurve**'s output is a FITS file, which can be plotted with **chips**. The lightcurve "rate" can then be divided by the effective area (ARF) to get a flux.

ACIS-TE Fundamentally limited by the 3.2 s frame time.

ACIS-CC Much higher time resolution (3 msec) at the expense of one spatial dimension.

HRC-S Can achieve 16 μ sec timing accuracy if only the central plate is on and telemetry not saturated.



A light curve of the active coronal binary, II Pegasi, showing a flare. With data-model filters, you can easily make light curves for narrow bands or specific features.

Response, Spatial PSF

- The PSF is a complicated redistribution (especially for X-ray optics!). The ARF must include the PSF fraction appropriate to the ARF region and source position. (For most regions about well isolated point sources, that fraction is very close to 1.0.)
- The PSF is generally a function of energy and source position relative to the optical axis (two vectors):
 $PSF \rightarrow PSF(\lambda, \hat{p}, \hat{q})$.
The PSF is often treated empirically (via a library of observed or ray-traced images), and is not formalized as a “response” function.

Some Practical Matters and Issues

Much analysis can be done with the three basic response functions:

ARF (spectral response, or effective area; low or high resolution)

RMF (spectral redistribution; low or high resolution)

Exposure Map (imaging mode)

But some cannot, such as ray-trace simulations or extended source analysis.

Forward-folding — is a necessary technique:

- Guess model parameters
- Predict observed signal (counts vs channel, or position)
- Repeat, until you minimize residuals.

E.g., **Sherpa**, **ISIS**, **XSPEC** spectral fitting;

Or, source spatial model times exposure map.

“Inversion” —

- Division of counts image by exposure map;
- Division of high-resolution spectrum by ARF;

WARNING: inversion can be dangerous! Ignoring source-model dependence in maps or redistribution in spectra can lead to inaccurate fluxes.

Simulation —

Ray-trace simulations necessarily use calibration data at the lowest, most detailed level. For instance,

rays require 3D geometry, and realistic offsets between filters and detectors. Analysis does not, since the signal can usually be defined adequately by the product of efficiencies.

Signal-to-Noise ratio —

The S/N can affect what calibration is appropriate for a given goal. For example, perhaps the signal is not large enough to warrant spectral fitting, and exposure map “inversion” is adequate for a survey of faint sources.

Some rules of thumb for low resolution spectra:

100 counts: enough for a hardness ratio.

1000 counts: A poor spectrum that may be able to distinguish between a power-law and thermal source.

10,000 counts: Some strong individual lines can be measured.

Computation speed —

For a survey, one must balance the speed and accuracy of making exposure maps against that of making ARFs. Analysis is slightly different: dividing counts by map, or folding a spectrum by the ARF.

Calibration Quality —

Calibration is not perfect! However, little analysis takes the uncertainty of the calibration into account. This is indeed difficult because it is typically a systematic term (e.g., 3%). With high-signal data, the systematic can dominate residuals.

Furthermore, there are always regimes where calibration is poor (10%, 20%, or more).

Don't over interpret data without considering calibration quality!

Note: systematics cannot be “added in quadrature” to statistical uncertainties. Responses typically have some rigidity to the function (i.e., 10% means over some correlation distance within which the function can only be perturbed in a smooth fashion).

Add data, or fit jointly? – It may be tempting to add spectra or images, and also add the responses, in order to have fewer objects to analyze, and to improve counting statistics. However, you may compromise knowledge of systematics or calibration quality dependent on spatial or spectral region. (e.g., you can't as easily ignore the bad parts in favor of the good.) You may also make certain coordinates ambiguous (e.g., position of the optical axis).

High-Resolution issues: For grating spectra, you should consider the following:

Continuum: Many thermal sources have far more lines than our models include. Be careful to fit the continuum in line-free regions only.

Line Shapes: Lines can be modeled as gaussians, or as delta functions if an RMF is used. If line shape (broadening) is important, a RMF should be used.

Adding Orders: Will increase counts, but with a cost in resolution. Joint fitting will likely give the best results, after a summed (and possible smoothed) qualitative analysis is done to guide the eye.