AXAF Exposure Maps: *Modus Operandi*

David Huenemoerder
Revision 1.0

15 January 1998

Contents

1 Purpose of This Document

This document is intended to define the exposure map, some intended usage scenarios, and to discuss some issues related to use or implementation. It is *not* a user's guide or reference manual to the software, which will be provided in future documents.

2 Exposure Map Definition

The **Exposure Map** is defined as

$$E(\lambda_D,\, \lambda,\, \hat{p}) \tag{1}$$

which has units of [$\mathrm{cm}^2\,\mathrm{s}$]. λ_D is the detected "wavelength" (proportional to the inverse pulse-height), λ is the photon wavelength, and $\hat{p}$ specifies the sky coordinate — a vector, representing the two spatial coordinates, (x, y). E is essentially a position-dependent redistribution matrix, scaled by the exposure time and the effective area. It most generally includes the mirror efficiency (and vignetting), detector efficiencies, and detector geometry. Alternatively, it can be calculated as exposure time vs. sky position. It *does not* contain any PSF information.

The relationship between the source spectrum and the counts is given by

$$S_{\hat{p}}(\lambda_D) = \int d\lambda\, E(\lambda_D, \lambda, \hat{p})\, f_{\hat{p}}(\lambda), \tag{2}$$

which, without the spatial dependence, is the "response matrix" multiplication solved by tools such as `XSPEC`. Here, $f_{\hat{p}}(\lambda)$ is a spatially-dependent source spectrum (PSF effects ignored), and $S_{\hat{p}}(\lambda_D)$ is the pulse-height counts-spectrum per sky bin, $\hat{p}$.

Issue: I didn't include a background term.

At the full spatial and spectral resolutions of AXAF (and conceivably with some temporal binning), this object would be very time-consuming to compute and unwieldy in size. Instead, we will allow its computation over specifiable regions, coarse spatial resolution, or coarse pulse-height (λ_D) resolution, specifically tailored to the user's analysis. **Hence, the type of analysis to be done will determine what is kind of exposure map is computed.**

E could also be determined as a function of time; that dependence will be implicit unless made explicit for time-dependent applications.

3 Modes of map calculation

3.1 Mean exposure in a sky region

$$E(\lambda_D,\, \lambda) = \int_{\Delta\hat{p}} d\hat{p}\, E(\lambda_D,\, \lambda,\, \hat{p}) \tag{3}$$

Here, the integral over area can be as simple as a circle around a point source, or it could be an annular region, or an arbitrarily shaped region. The resulting map

The user must be aware and correct for PSF affects, if necessary. For example, for an isolated point source, the region should be large enough that the same fractional flux is enclosed at any energy for the presumed spectrum, or an Auxiliary Response File (ARF) generated which contains that information if the assumption isn't valid. For extended sources, regions should be large enough so that contributions from neighboring bins are small.

is effectively the average response in the region of interest. It is analogous to the $ARF \times RMF$ (Auxiliary Response File times Response Matrix File).

The region, in sky coordinates, may contain events from devices of very different response, such as back- or front-illuminated CCDs of ACIS-S. It is assumed that *either* of the following is true:

- the response is uniform over the region, or
- the photons are uniformly distributed over the region.

This is expected to be the default mode for point-source analysis in which the spatial distribution of the counts within the region is uniform.

3.2 Wide-field, narrow-band images

For wide-field, diffuse-source imaging, it can be useful to correct an image by providing a flat-field image in a narrow wavelength band, $\Delta\lambda$: $E_{\Delta\lambda}(\hat{p}) = E_{\Delta\lambda}(x, y)$. This requires some crude assumptions (depending upon the bandpass and redistribution function), but can be useful for creating approximately calibrated diagnostic images or for comparison with images from other observatories. This case is also an example for general usage of the map for other binning options to follow.

The source counts per pixel can be written as:

$$S_{\hat{p}}(\lambda_D) = \int d\lambda \, E_{\hat{p}}(\lambda_D, \lambda) \, f_{\hat{p}}(\lambda). \tag{4}$$

In a narrow pulse-height band, λ_D, we have:

$$\begin{aligned} S_{\hat{p}}(\Delta\lambda_D) &= \int_{\Delta\lambda_D} d\lambda_D \, S_{\hat{p}}(\lambda_D) && (5) \\ &= \int d\lambda \, f_{\hat{p}}(\lambda) \int_{\Delta\lambda_D} d\lambda_D E_{\hat{p}}(\lambda_D, \lambda) && (6) \end{aligned}$$

The last integral (over λ_D) is a known quantity, given the calibration data and the aspect solution. For any model spectrum, $f_{\hat{p}}(\lambda)$, we could predict the counts-image in λ_D bins. However, we often don't know the spectrum, but wish to examine its spatial variations in a narrow band, either over a line-complex or small continuum region, without instrumental artifacts. To approximate this, we can limit our integrals to a narrow $\Delta\lambda$ about a specific wavelength, λ_0, which in turn defines a range in λ_D. We ***assume*** that the spectral shape in $\Delta\lambda$ is constant with position, and replace the flux with the mean value:

$$f_{\hat{p}}(\lambda) \to \overline{f_{\hat{p}}}. \tag{7}$$

Integrating over this wavelength bin we derive

$$S_{\hat{p}}(\Delta\lambda_D) \approx \overline{f_{\hat{p}}} \int_{\Delta\lambda} d\lambda \int_{\Delta\lambda_D} d\lambda_D \, E_{\hat{p}}(\lambda_D, \lambda). \tag{8}$$

This can be solved for the "fluxed" image:

$$\overline{f_{\hat{p}}} \approx \frac{S_{\hat{p}}(\Delta\lambda_D)}{\int_{\Delta\lambda} d\lambda \int_{\Delta\lambda_D} d\lambda_D \, E_{\hat{p}}(\lambda_D, \lambda)} \equiv \frac{S_{\hat{p}}(\Delta\lambda_D)}{E_{\Delta\lambda}(x, y)} \tag{9}$$

The right-hand-side is comprised of known quantities.

3.3 Response for an assumed model spectrum

The narrow-band case is practical for a detector with modest spectral resolution, such as achieved by the CCDs of ACIS. For micro-channel plate detectors of HRC, however, narrow band images are not meaningful since the PHA distribution is broad and relatively independent of photon energy. In this case, the response to an assumed model spectral energy distribution can be computed by doing the integrals in Equations 5 and 6, and integrating over the entire range of λ_D. (This mode is, of course, also useful for ACIS.) The resulting product is thus a sensitivity map: the expected total counts vs. position for the assumed spectrum.

Since we need to integrate over all PHA bins, we can write

$$\begin{aligned} S_{\hat{p}} &= \int_0^\infty d\lambda_D \, S_{\hat{p}}(\lambda_D) && (10) \\ &= \int d\lambda \, f_{\hat{p}}(\lambda) \int_0^\infty d\lambda_D E_{\hat{p}}(\lambda_D, \lambda) && (11) \\ &= \int d\lambda \, f_{\hat{p}}(\lambda) A_{\hat{p}}(\lambda), && (12) \end{aligned}$$

where $A_{\hat{p}}(\lambda)$ is the response integrated over all $\Delta\lambda_D$, and is like the effective area function (or ARF, but scaled by time). We can take this one step further by assuming an analytic form for the spectrum to compute a "generic-source" definition of an exposure map. If we let the spectrum be a power-law, $f_{\hat{p}}(\lambda) = \mathcal{F}_{\hat{p}} \lambda^\alpha$ (in which $\mathcal{F}_{\hat{p}}$ is a position-dependent normalization factor for the spectrum), then we can produce an exposure map as defined by

$$E_{\hat{p}} = \int d\lambda \, \lambda^\alpha \, A_{\hat{p}}(\lambda). \tag{13}$$

This can be computed for a couple values of α which characterize the gross shape of spectra.

3.4 Diagonal redistribution matrix

The exposure map may be factored into two terms, the effective area function ("ARF") and the redistribution function ("RMF") (each is in general spatially dependent):

$$E_{\hat{p}}(\lambda_D, \lambda) = A_{\hat{p}}(\lambda) R_{\hat{p}}(\lambda_D, \lambda). \tag{14}$$

If we assume that the redistribution matrix is strictly diagonal (a useful approximation for ACIS, but risky business for HRC), we have $R_{\hat{p}}(\lambda_D, \lambda) = \delta(\lambda_D - \lambda)$, which allows us to trivially integrate for the counts:

$$S_{\hat{p}}(\Delta\lambda_D) = A_{\hat{p}}(\lambda_D) f_{\hat{p}}(\lambda_D) \tag{15}$$

or

$$f_{\hat{p}}(\lambda_D) = S_{\hat{p}}(\Delta\lambda_D) / A_{\hat{p}}(\lambda_D), \tag{16}$$

which is a "quick-and-dirty inversion" by division of the PHA histogram by the mean effective area function.

3.5 Exposure in coarse spatial bins

To reduce spatial resolution, we can integrate over large bins:

$$E(\lambda_D,\, \lambda,\, \hat{p}_i) = \int_{(\Delta\hat{p})_i} dA\, E(\lambda_D,\, \lambda, \hat{p}) \tag{17}$$

where the notation, $(\Delta\hat{p})_i$ indicates integration over the limits, $\hat{p}_i - \Delta\hat{p}_i/2$ to $\hat{p}_i + \Delta\hat{p}_i/2$ (remember that $\hat{p}$ has two spatial coordinates, so $\Delta\hat{p}$ represents both Δx and Δy).

This is similar to Section 3.1, except that we have allowed a grid of rectangular regions as opposed to a single arbitrarily shaped region.

3.6 Exposure in coarse pulse-invariant bins

Similarly to spatial binning, we can integrate over pulse-height bins:

$$E(\lambda_{D_i},\, \lambda,\, \hat{p}) = \int_{(\Delta\lambda_D)_i} d\lambda_D\, E(\lambda_D,\, \lambda,\, \hat{p}) \tag{18}$$

This is similar to the PI integration of Section 3.2, where we subsequently factored out the flux and integrated over wavelength. With HRC's low spectral resolution, only one to three PI bins may be practical (such as "soft", "hard", and total).

If broad PI bins are used, application requires a model spectrum, since the source flux term cannot be removed from the integral (see Section 3.3).

3.7 Binning in both space and λ_D

Binning over multiple dimensions may be done simultaneously and may be appropriate for quick-look or photon-limited observations.

3.8 The map may be calculated for specified time-intervals

$$E(\lambda_D,\, \lambda,\, \hat{p}, t_i) = \int_{(\Delta t)_i} dt\, E(\lambda_D,\, \lambda,\, \hat{p},\, t) \tag{19}$$

This mode is necessary for light-curve investigations, where the response is needed in temporal bins commensurate with the light-curve bins. It would most likely be used with the mode of Section 3.1 to produce a series of response matrices as a function of time (and optionally combined with spatial or spectral binning).

Issue: need a compressed mode for this. If redistribution function does not change with time, or $QE(\lambda)$, then we need a single exposure map for the region and a time-dependent normalization.

3.9 Maps which maintain detector identity

To preserve correspondence of photons with detector elements on may filter events on detector element or region and make appropriate maps. For example, if a source wanders (via aspect motion) from front- to back-illuminated CCDs, separate PI spectra or images can be made for by selecting only photons from the relevant detector element or region and the appropriate map generated. This mode may be of use if one detector element's calibration is more uncertain than the other and has adjustable parameters (TBR).

3.10 Grating Mode

For grating spectra, an extraction mask is specified by a zero-order position, cross-dispersion width, and PI fraction. These must be consistent with the counts spectrum binned from the Level 1.5 event file. It is assumed that order-sorting has been done via the grating response and CCD response matrix for ACIS.

The spectroscopic map is specified by

$$E_{\hat{p}_0}(\lambda_G, m), \tag{20}$$

in which $\hat{p}_0$ is the zero-order position, λ_G is the detected grating wavelength, and m is the spectral diffraction order. The detected grating wavelength, λ_G, is close to λ, the source wavelength, but is not identical due to the instrumental line spread function.

This map is a two-dimensional object of size $n(\lambda)$ by $n(m)$ (on the order of 3000×20). It is different from the imaging-mode map because for any given detector position, its value depends on the zero-order location. Hence, it can only be made on a per-source basis.

The predicted counts are approximately

$$S_{\hat{p}_0}(\lambda_G, m) = f_{\hat{p}_0}(\lambda_G) E_{\hat{p}_0}(\lambda_G, m) \tag{21}$$

(having equated λ with λ_G in f).

For HRC, since there is no order resolution, the counts-spectrum is given by a sum over orders at each diffraction angle, $\beta \approx m\lambda/P$, where P is the mean grating period:

$$\begin{aligned} S_{\hat{p}_0}(\beta) &= \int_{m=1}^{m=\infty} dm\, E_{\hat{p}_0}(\lambda_G(\beta), m)\, f_{\hat{p}_0}(\lambda(\beta)) \qquad (22) \\ &= \int_{m=1}^{m=\infty} dm\, E_{\hat{p}_0}((P\beta/m), m)\, f_{\hat{p}_0}(P\beta/m) \qquad (23) \end{aligned}$$

(where $\infty \approx 30$).

3.11 Exposure map gradient utility

To provide the user with a quick assessment of whether particular region of the sky is likely to be affected by strong gradients in the response (either QE or redistribution), a utility will provide a map of regions exposed by different array or filter elements. For example, if dither exposes a sky region on two CCDs of ACIS-I, then the utility will outline that sky region. The user should then take particular care analyzing any sources in that region.

This utility will function by saving the extrema in the projection of detector element boundaries onto the sky and performing the boolean "more-than-one" operation on these regions to show the overlaps.

For example, "more-than-one" for 3 detector elements, A, B, and C, means
{ (A AND B) OR (A AND C) OR (B AND C)}

3.12 Background Map applications:

—*TBD*—

—Get input from A.D.—

4 Miscellaneous issues: Speed, accuracy, implementation.

This section new 980115; un-reviewed

1. Aspect solution: histogram or time-series?

   The AXAF aspect solution provides a pointing and detector position every 0.256 s. Since the exposure map computation time scales directly with the number of aspect pointings to be processed, it is important to compress the aspect solution into a set of unique pointings.

   The nominally programmed dither pattern is large compared to the on-axis point-spread-function (40 arcsec vs 1 arcsec). Hence, it is not clear that any cell in a 6-dimensional aspect solution will be re-visited, which means that an aspect histogram does not economize over a sequential time integration. Actual performance (and better performance models) will help answer this. For example, SIM motion is expected to be very slow so 3 of the 6 dimensions may compress into a few bins. If roll variation is as small as predicted, it also fits into a small number of bins. Hence, an aspect histogram as used in past missions may be efficient.

   Alternatively, stable aspect intervals may be defined which compress the sequential aspect points into the times during which the mean pointing does not change by more than some fraction of a sky bin. This also has dependencies on the actual aspect performance. For instance, if roll is significant, these intervals depend on the off-axis angle of the detector region.

   Both aspect histogram (via a 64-tree, the 6-dimensional counterpart of the oct-tree of a 3D aspect solution) and sequential options will be available.

   The time to process TBD aspect bins for map resolutions of TBD is TBD.

   TBD Aspect time scaling.

2. Pipeline vs. User applications:

   Rigorous detailed analysis of an imaging or spectroscopic observation involves many choices regarding spatial region definition, spatial resolution desired, spectral region of interest, and source model characteristics. These cannot be known *a priori* for purposes of autonomous standard pipeline processing. Hence, exposure map tools must be flexible under user control.

   On the other hand, some source-model-independent products can be derived in certain cases:

   - For well-defined point-sources or for a small sky region of interest, a mean response can be calculated ($E(\lambda_D, \lambda)$, Equation 3 of Section 3.1).
   - For ACIS imaging observations, a moderately source-model-independent map can be made in several PI bins: $E_{\Delta\lambda}(x, y)$ (Equation 9, of Section 3.2).

   For HRC imaging observations, a compromise must be made, such as simple spectral functions (such as described by Equation 13 in Section 3.3).

3. Model spectrum mode (see Section 3.3): Given a model spectral energy distribution, one can pre-calculate the integrated (over $\Delta\lambda_D$ and wavelength) response in every detector pixel, including the mirror response. Then the aspect solution can be applied to project this scalar image onto the sky. This is perhaps the most efficient mode, but requires the most knowledge of the source.

4. Sky grid resolution directly affects the computation time by the number of bins whose response must be computed. A related term is the degree of accuracy to which the mirror vignetting with off-axis angle must be interpolated.

   Optimization trade-offs are TBD.

5. Pixel sub-sampling (also known as anti-aliasing) can be applied to produce a smoother map, but costs computationally. This may be important since the PSF is comparable to a pixel, but may be sufficiently smoothed by the dither. Use is TBD.

6. Detector uniformity resolution and detector size affects computation time. Regions of uniform response are projected to the sky, and sky pixels filled with a constant value. More regions require more time. Optimization trades are TBD.

7. For grating spectroscopy, the regions over which the exposure map is calculated and over which the 1D spectrum is extracted must be consistent. Whether a 2D (wavelength and cross-dispersion coordinate) maps is required is TBD.

8. PSF issues: far off axis, less resolution is required. Spatial grids could be coarser.

9. Multi-resolution issues: For all types of maps, multi-resolution calculations are needed to handle both the low spatial resolution of the detector quantum-efficiency uniformity maps and the high spatial resolution of the edges and bad-pixels/regions. For small-scale maps, these can be combined into one map at the desired resolution. For large-scale maps, size constraints may demand maintaining multiple-resolution files (TBR); if multiple resolution files are not kept, the affect on exposure of the fine structure will be included into the coarse maps.

10. For small-scale maps the redistribution function probably will not change over many bins (while the normalization may). Whether it is practical to factor these "ARF", "RMF", and normalization terms is TBD.

11. Spatially binned maps may be associated as a "stack" with the ASC Data Model (TBD).

5 The Bottom Line: Recommended Products

980116 New section; un-reviewed

5.1 Pipeline Processing

Pipeline (autonomous) processing should provide a small number of useful diagnostic exposure maps, but not necessarily products for rigorous and detailed source modeling and analysis, since they depend heavily on subjective scientific decisions.

5.1.1 ACIS Imaging

ACIS imaging observations should produce wide-field, narrow-band exposure maps as described in Section 3.2 (see Eq. 9).

In addition, if there are point-sources detected, the target and few brightest in the field will have the mean exposure map calculated in the source region, as described in Section 3.1 (see Eq. 3).

5.1.2 ACIS, continuous clocking mode imaging

Continuous clocking mode processing must assume a source position in the parallel transfer direction. This can be estimated from the pointing and known field sources. Given that, 2D or source region maps can be calculated and collapsed to 1D.

5.1.3 HRC Imaging

HRC imaging observations require assumption of a model spectrum. An exposure map as described in Section 3.3 for several model assumptions will be computed. Assuming a unit flat spectrum (power-law index $\alpha = 0$) will produce a map of detector-mirror exposure. One or two other prudent choices for α (generally characterizing hard or soft thermal and power-law sources) will also be made and maps computed.

5.1.4 HETGS

HETGS observations (that is, HETG with ACIS-S) will have exposure maps (as defined in Section 3.10) for the target source and default extraction region for first orders of MEG and HEG. In addition, an imaging-mode response will be calculated for the zero-order region. For background/confusing source estimation, one or more narrow-band wide-field maps will be computed.

In continuous-clocking mode, the source position in the parallel-transfer direction must be assumed or known *a priori*. Given that, spectroscopic mode maps can be derived. HEG odd order response will have to be combined with MEG even-order response.

5.1.5 LETGS

LETGS observations (LETG with HRC-S) will have exposure maps generated similarly to HETGS.

HETGS/HESF (Drake Flat) maps are TBD.

Drake Flat maps needs definition.

5.2 Interactive Processing

The user will be able to recreate any of the standard products for different filtering criteria (spatial, spectral, detector element, time, grade-set, etc). Some parameter choices may be prohibitively expensive to compute, however. Hence, there will be some diagnostic tools and re-usable pre-computable objects, and tabulated benchmarks to guide the analysis. The user should be able to compute exposure maps of the forms described in Sections 3.1–3.12.

5.3 Diagnostic Tools and Pre-computable Objects

There are several things which can be done to reduce the scale of the computation or to diagnose problem areas. (These will grow with experience.)

1. Aspect Histogram: For some applications, the aspect solution can be binned into a histogram. For AXAF, this is a 6-dimensional object, since there are 3 directions for the optical axis (right-ascension, declination, and roll) and another three for the detector (SIM x, y, and θ). This only needs to be calculated once for any set of good time intervals (GTI), aspect quality, or aspect grid parameters. It can then be re-used for different map parameters (such as spectroscopic vs. imaging, or different model spectra or PHA bands).

2. Aspect Stable Intervals: Another way to compress the aspect histogram is to determine times when the pointing is within a specifiable sky bin, and then save this list of mean pointings. This can then be used in different exposure map parameter cases as described for the histogram. (This mode is appropriate for light-curve analysis, for instance, for which the histogram is not.)

3. Regions of steep response gradient: The detector arrays are heterogeneous. ACIS-I is comprised of 4 CCDs, each with a unique response. ACIS-S has very different elements: front- and back-illuminated CCDs. HRC-S has several different filter regions, with fuzzy boundaries (they are shadows of the converging cone). If the aspect dither shuffles an interesting source across boundaries, analysis could be difficult. A pre-processing tool can give a map on the sky of the maximal extent of each detector element.

4. Merged Response: For any selection which affects the instrumental calibration products (e.g., ACIS grade-set), or for any choice of detector regions and resolution, or sky-grid resolution, the detector and mirror responses can be multiplied before applying the aspect solution. Furthermore, if a form for the source spectrum is known, that can also be convolved with the response and possibly integrated. The final map can then be made with this merged response.

6 Relationship to previous missions

980115 new section, unreviewed

1. ROSAT HRI used bright-earth data to make a map of size 512×512, blocked 30×30 pixels (to 5.0 arcsec). Vignetting was only done for 0.53 keV. The `FTOOL` is `hriexpmap`, and is based on Snowden et al. 1992 (ApJ, 393 819). `FTOOLS` documentation says that the task can take a "long time" (without further elaboration).

2. ROSAT PSPC (also based on Snowden et al. 1992; `FTOOL pcexpmap`) used a composite detector-source response map derived from all-sky survey data, with known sources removed. This is an empirically determined flat field *for some mean incident background spectrum.* As such, it is an approximation to the instrument response. The output is a 512×512 image with 15 arcsec pixels, which is a blocking by 16 of the detector pixels. It may be calculated in any of several pre-determined PHA ranges.

   The input map dimensions are TBD by TBD spatially, and TBD PHA bins. To create a sky map of size TBD by TBD for TBD PHA bins takes TBD seconds on a TBD computer.

   ROSAT TBD, TBD, TBD benchmark.

3. ASCA: Exposure maps for ASCA (`FTOOL ASCAEXPO`) can be made for just exposure time, exposure time and mirror response, or time with both mirror and detector responses. The detectors are assumed to be uniform. Attitude records may be decimated or a minimum significant step-size set to improve speed, but possibly sacrifice accuracy. A time-independent map may also be made (no attitude file used). Input and output maps may be rebinned (but by the same factor). The `FTOOL` doesn't include GIS dead-time or SIS frame saturation.

   The input instrument map determines PHA regions. This can be created via the `FTOOL`, `ascaeffmap` for either the SIS or GIS. `FTOOLS` documentation indicates that an "extremely long time" can be required for making an instrument map over the entire energy band.

   A timing benchmark is TBD.

   ASCA benchmark TBD.

7 Document History

971230 dph, revision 0.9, small distribution for comments.

reviewed by: D.Dewey, D-W.Kim

980115 dph revision 1.0.

8 References

1. Examples of extended source analysis in narrow PI bands for ROSAT and ASCA.

   (a) Hwang, U., and Gotthelf, E.V., (1997), Ap.J., 475, 665.

   (b) Snowden, S.L., McCammon, D., Burrows, D.N., and Mendenhall, J.A., (1994), Ap.J., 424, 714.

2. HEASARC resources:

   (a) ROSAT GOF FAQ has some exposure map tips:
   (`http://heasarc.gsfc.nasa.gov/docs/rosat/faqs/pspc_faq_list.html`)
   (b) ROSAT FTOOLS: `HRIEXPMAP`, `PCEXPMAP`
   (c) ASCA Exposure Map and tutorial:
   `http://heasarc.gsfc.nasa.gov/docs/asca/abc/node11.html`
   `#SECTION00115000000000000000`
   (d) `ASCAEFFMAP`: ASCA effective area map `FTOOL`.
   (e) `ASCAEXPO`: ASCA exposure map `FTOOL`.

3. SAO Historical resources:

   (a) SAO Special Report #393 (1982), L. Van Speybroeck, "Exposure Map Definitions", p.255.