An Introduction to Exposure Maps

John Houck

Center for Space Research, Massachusetts Institute of Technology, Cambridge MA 02139

1. Introduction

The counts image of an extended X-ray source invariably contains instrumental artifacts which complicate surface brightness measurements. These artifacts are caused by imperfections in the mirror and the detector and are both energy and position dependent. Often, one would like to generate an image in physical units (e.g. photon sec^{-1} cm^{-2} arcsec^{-2}), with these artifacts removed. A common approach is to divide the counts image by an *exposure map* to re-scale all parts of the image to the same relative exposure.

Unfortunately, in the general case, no exposure map with this property exists. The reason is that, because the effective area is energy dependent, the number of counts produced in a given broad pulse-height band is not proportional to the incident flux in that band. Instead, an integral equation relates the total counts produced in a given pulse-height bin to the incident source spectrum; solution of this integral equation for point-source observations is the primary function of XSPEC and SHERPA.

The imaging case is further complicated by the need to account for the position-dependence of the effective area. As with point-sources, an integral equation describes the relationship between the counts image and the source surface-brightness distribution in a given energy band. In general, simply dividing the counts image by an exposure map does not constitute a solution to this integral equation and does not yield the true source surface brightness vs. position in the chosen band. However, when the effective area is roughly constant within a given energy band, the relevant integral equations may be greatly simplified, allowing one to divide by an exposure map to effectively remove instrumental artifacts from a counts image within the specified band.

2. Definition of the Exposure Map

For an explicit first-principles derivation of the equations defining the ARF and the exposure map, see Davis (2000). Here, we merely state the relevant results and discuss how they might be used in the analysis of X-ray imaging data. After simplifying for the case when the mirror effective area and point-spread function shape are slowly varying, the integral equation relating the counts image to the source surface brightness distribution is

$$C(h, \hat{\mathbf{p}}) = \tau_{\mathrm{eff}} \int d\lambda \; E(h, \lambda, \hat{\mathbf{p}}) S_{\mathcal{F}}(\lambda, \hat{\mathbf{p}}) \tag{1}$$

where the integral extends over all wavelengths λ which contribute counts to pulse-height bin h. In this expression, $C(h, \hat{\mathbf{p}})$ is the expected number of aspect-corrected counts with pulse-height h attributed to sky position $\hat{\mathbf{p}}$, $S_{\mathcal{F}}(\lambda, \hat{\mathbf{p}})$ is the PSF-smeared source surface brightness at wavelength λ and sky position $\hat{\mathbf{p}}$, and τ_{eff} is the effective exposure time. The quantity $E(h, \lambda, \hat{\mathbf{p}})$ is the exposure map; from equation (1), it follows that E must have units of cm^2 counts photon^{-1} (given $S_{\mathcal{F}}\Delta\lambda$ in photon sec^{-1} cm^{-2}, τ_{eff} in seconds, and C in counts).

The exposure map for a given observation combines the effective area of the telescope and detector with a map of the dwell time vs. sky position, accumulated by following the telescope pointing motion during the observation. The effective area varies with position on the detector and is also energy dependent. At a given energy, the effective area as a function of position on the detector is called the instrument map. The map of dwell time vs. pointing direction, built up by the telescope pointing motion, is called the aspect histogram. Combining the instrument map with the aspect histogram, the exposure map is a map of the total exposure as a function of position on the sky. In particular, note that the exposure map has units of *Area* (ignoring for the moment the dimensionless ratio counts photon^{-1}). Therefore, to convert a counts image to flux units, the relevant quantity is the product (effective aperture *Area*)×(exposure *Time*) (see equation (5) below).

To obtain an image with a statistically significant number of counts per pixel in the region of interest, it is often necessary to sum over counts in a pulse-height interval (Δh). Summing equation (1) over such an interval and taking care to extend the integral over the wavelength range $(\Delta\lambda)$ which contributes to (Δh), we obtain

$$C(\Delta h, \hat{\mathbf{p}}) = \tau_{\text{eff}} \int_{\lambda \in \Delta\lambda} d\lambda \; E(\Delta h, \lambda, \hat{\mathbf{p}}) S_{\mathcal{F}}(\lambda, \hat{\mathbf{p}}) \tag{2}$$

with the understanding that the pulse-height range Δh is summed over, so that, e.g.

$$C(\Delta h, \hat{\mathbf{p}}) \equiv \sum_{h \in \Delta h} C(h, \hat{\mathbf{p}}). \tag{3}$$

Note that the wavelength range $\Delta\lambda$ does not correspond exactly to the band width Δh in PI units. In particular, because of the finite energy resolution of the detector, $\Delta\lambda$ must extend over a somewhat broader range than Δh so that it can cover the entire range of λ which significantly contributes to Δh.

If we choose a wavelength range $\Delta\lambda$ over which $E(\Delta h, \lambda, \hat{\mathbf{p}})$ is relatively constant (see Figure 1), we can remove E from the integrand in equation (2) to obtain

$$C(\Delta h, \hat{\mathbf{p}}) \approx \tau_{\text{eff}} E(\Delta h, \lambda_0, \hat{\mathbf{p}}) \int_{\lambda \in \Delta\lambda} d\lambda S_{\mathcal{F}}(\lambda, \hat{\mathbf{p}}). \tag{4}$$

Given a counts image for this band, $C(\Delta h, \hat{\mathbf{p}})$, we can divide by the corresponding exposure map, $E(\Delta h, \lambda_0, \hat{\mathbf{p}})$, to obtain an image, in physical units, of the PSF-smeared source integrated over the

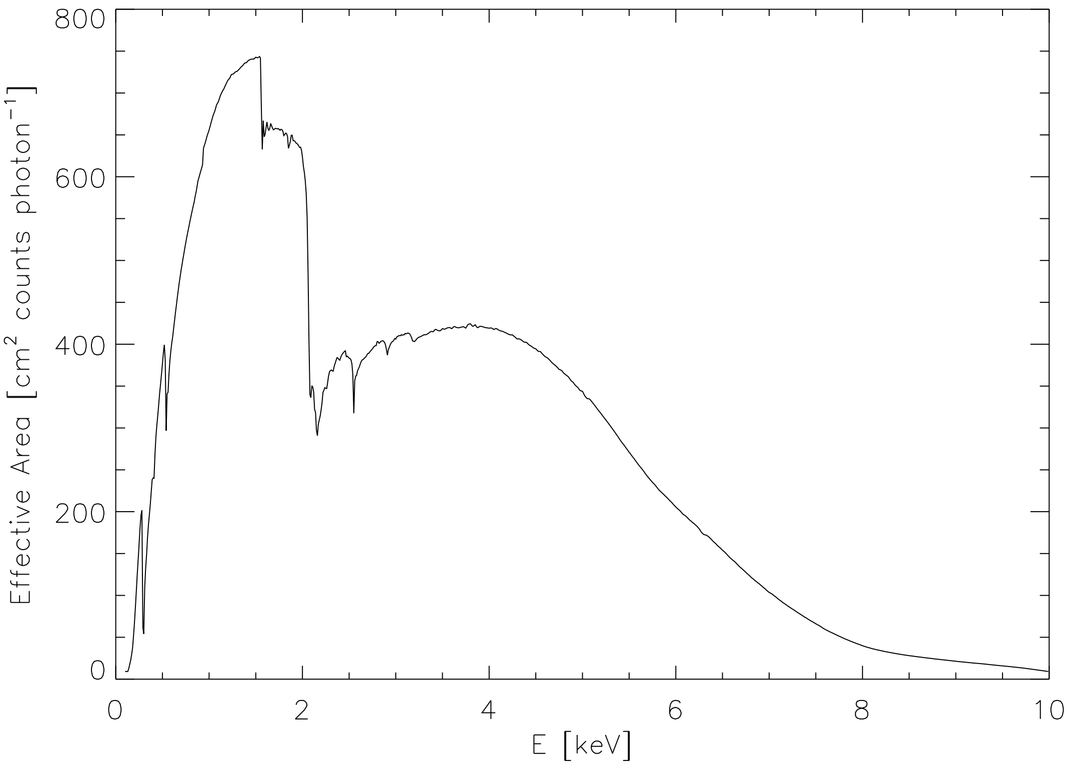

Fig. 1.— A typical effective area curve.

wavelength band $\Delta\lambda$:

$$\int_{\lambda \in \Delta\lambda} d\lambda S_{\mathcal{F}}(\lambda, \hat{\mathbf{p}}) = \frac{1}{\tau_{\rm eff}} \cdot \frac{C(\Delta h, \hat{\mathbf{p}})}{E(\Delta h, \lambda_0, \hat{\mathbf{p}})}. \tag{5}$$

We emphasize that equation (5) is valid only to the extent that the effective area is constant over the wavelength range of interest ($\Delta\lambda$).

Rather than requiring that the effective area be constant over the entire wavelength range of interest, we can account for some of the energy dependence by approximating the effective area function with a number of piecewise-constant regions. Neglecting the $\hat{\mathbf{p}}$ dependence, we have

$$C(\Delta h) = \tau_{\rm eff} \sum_i E_i \int_{\Delta\lambda_i} d\lambda S_{\mathcal{F}}(\lambda) \tag{6}$$

where $E_i \equiv E(\Delta h, \lambda_i)$ is constant within each $\Delta\lambda_i$. If we now define a set of weights

$$w_i \equiv \frac{\int_{\Delta\lambda_i} d\lambda S_{\mathcal{F}}(\lambda)}{\int_{\Delta\lambda} d\lambda S_{\mathcal{F}}(\lambda)} \tag{7}$$

where $\Delta\lambda = \sum_i \Delta\lambda_i$, we can define a *weighted exposure map* $\mathcal{W} \equiv \sum_i w_i E_i$ (MKINSTMAP allows such weighting to be included). Dividing by this weighted exposure map we obtain a better estimate of the PSF-smeared source integrated over the broad wavelength band $\Delta\lambda$ so that

$$\int_{\Delta\lambda} d\lambda S_{\mathcal{F}}(\lambda) \approx \frac{C(\Delta h)}{\tau_{\mathrm{eff}}\mathcal{W}}. \tag{8}$$

Although the incident source spectrum is usually unknown, one can use a model spectrum to compute approximate weights using equation (7).

To the extent that the weights w_i accurately represent the shape of the incident spectrum over the entire image and that the effective area function is piecewise constant over the wavelength range of interest, one can use equation (8) to compute a surface brightness image spanning a relatively broad wavelength range. However, we emphasize that by neglecting the $\hat{\mathbf{p}}$ dependence, we have assumed that a single spectral shape applies to the entire image. If the shape of the incident spectrum varies significantly across the image, e.g. so that a single set of weights does not represent the incident spectrum over the entire image with sufficient accuracy, then it is not possible to accurately compute the correct surface brightness at every point in the image by simply dividing the counts image by a single exposure map.

To make this more precise, we examine the extent to which spectral variations or, equivalently, errors in the spectrum model, produce errors in the derived flux. To simplify the notation, assume $\tau_{\mathrm{eff}} = 1$ and ignore the Δh and $\hat{\mathbf{p}}$ dependencies; similarly, ignore the dither motion, so that the exposure map is essentially just the effective area curve, hence $E(\lambda) \rightarrow A(\lambda)$. The incident source spectrum may be written in the form $s(\lambda) = s_0(\lambda) + s_1(\lambda)$, where $s_0(\lambda)$ is our approximation of the incident spectrum and $s_1(\lambda)$ is the error term. In the limit of arbitrarily fine wavelength subdivisions, the weighted exposure map has the form

$$\mathcal{W}[s_0] \equiv \frac{\int d\lambda\ A(\lambda)s_0(\lambda)}{\int d\lambda\ s_0(\lambda)}. \tag{9}$$

From equation (8), dividing the counts image, C, by this weighted exposure map, $\mathcal{W}[s_0]$, gives the correct flux if

$$\int d\lambda\ s(\lambda) = \frac{1}{\mathcal{W}[s_0]} \int d\lambda\ A(\lambda)s(\lambda) \tag{10}$$

By substituting $s(\lambda) = s_0(\lambda) + s_1(\lambda)$ in this expression and using the definition of $\mathcal{W}[s_0]$, it follows that we can obtain the correct flux if $\mathcal{W}[s_0] = \mathcal{W}[s_1] = \text{constant}$.

Figure 2 shows $\mathcal{W}[s]$ for two common types of smoothly varying spectra. Because $\mathcal{W}[s]$ is slowly varying for these spectra, an accurate model for the incident spectrum should lead to correspondingly small errors in the computed flux. Figures 2a and 2b can be used to examine the accuracy of source fluxes computed using weighted exposure maps for both thermal and power-law sources. For example, if an exposure map weighted for a thermal spectrum with kT = 3 keV is applied to an image for which the true incident spectrum is kT = 3.5 keV, Figure 2a indicates that the error in the weighted exposure map will be only a few percent, giving a similarly small error in the computed flux.

On the other hand, if the image contains significant spectral variations, the resulting systematic errors may range from a few percent to nearly a factor of two, depending on the chosen bandpass $\Delta\lambda$, even for a weighted exposure map created with arbitrarily small energy bins. In such a situation, no single exposure map can apply to the entire image. For example, suppose we have an image of a galaxy cluster where the outer parts correspond to kT = 6 keV but the core is at 2 keV. In the 0.5-10 keV band, if we compute an exposure map weighted for a temperature of kT = 6 keV, Figure 2a indicates that the flux in the core would be overestimated by about 40%. Instead, if we restrict the bandpass to the 2.5-5 keV band where the response is nearly flat, the systematic flux error should be quite small over most of the image. Note that because the cosmic X-ray background spectrum is quite different from the thermal X-ray spectrum of the cluster, the flux computed for the cluster-free areas of the image may have significant systematic errors. For a more detailed example based on a MARX simulation using thermal spectra, see §4.

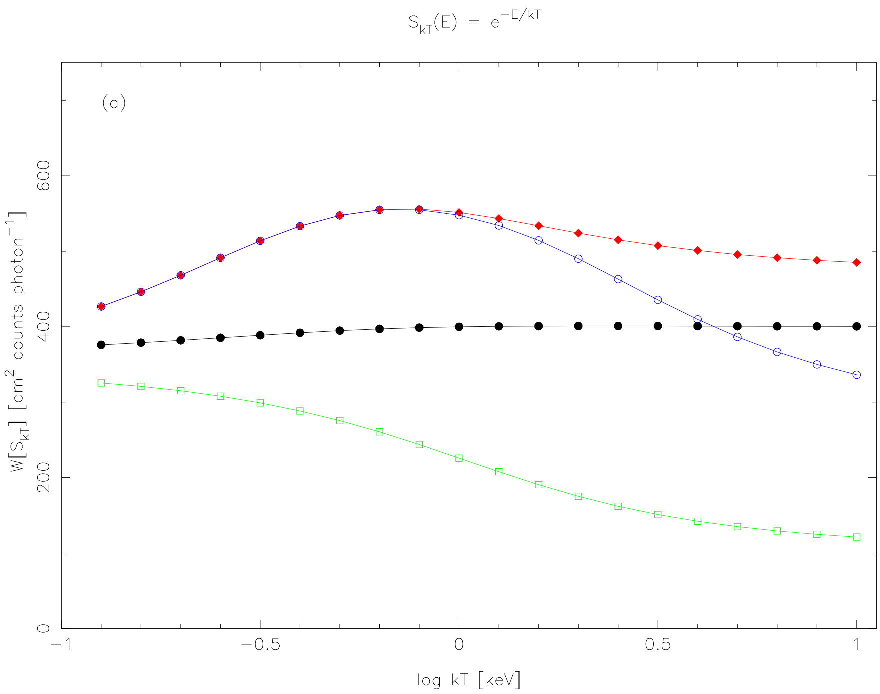

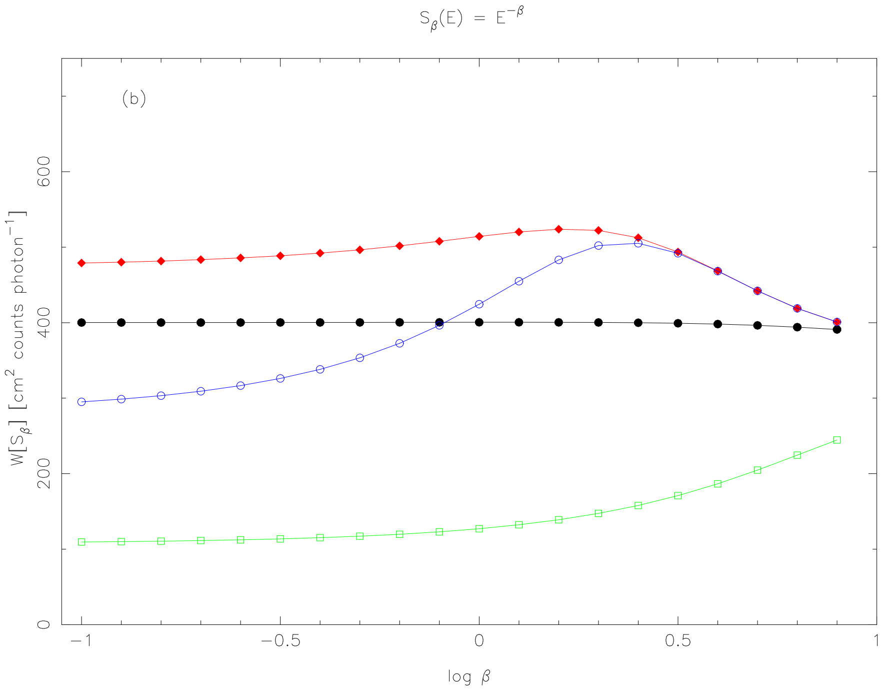

Fig. 2.— Dependence of weighted exposure map on spectral shape for (a) exponential and (b) power-law spectra. Each figure shows results for four bandpasses, $\Delta\lambda$; 0.5-10 keV (open circles), 0.5-5 keV (filled diamonds), 5-10 keV (open squares), 2.5-5 keV (filled circles).

3. Example 1: Monochromatic Sources

3.1. Source Properties

As an illustration of the problems which arise when the effective area varies significantly with energy, consider the problem of exposure correction within the relatively narrow energy range 1.4-2.7 keV which lies near the peak of the effective area curve and which also spans the Iridium edge at 2 keV. To simplify the interpretation of the image, consider an extended source which consists of two uniform surface brightness disks of 1 arcmin radius, with their centers spaced 1 arcmin apart such that the disks overlap. Each disk separately contributes a monochromatic flux of 0.1 photon $\mathrm{sec}^{-1}\mathrm{cm}^{-2}$, one at 1.47 keV (Magnesium), the other at 2.61 keV (Sulfur). It follows that, integrated over the energy range 1.4-2.7 keV, the surface brightness where the disks overlap should be exactly twice the surface brightness in the monochromatic regions.

To test the accuracy of the exposure correction, we will compare our results with the known surface brightness in each part of the source image. Each disk contributes a monochromatic flux $F(\lambda)$ which is distributed uniformly over an area of angular size A, so that the monochromatic surface brightness is $S(\lambda) = F(\lambda)/A$, with $A = \pi R^2$ and $R = 60$ arcsec. The surface brightness in each region is then $S = 8.84 \times 10^{-6}$ photon $\mathrm{sec}^{-1}\mathrm{cm}^{-2}\mathrm{arcsec}^{-2}$. and twice that, or 1.77×10^{-5} photon $\mathrm{sec}^{-1}\mathrm{cm}^{-2}\mathrm{arcsec}^{-2}$, in the overlap region.

3.2. Observation and Analysis

Using MARX 3.0, we generated a realistic simulation of a Chandra ACIS-S3 imaging observation of this idealized source. See the Appendix for details of the simulation and data analysis. To generate an exposure map for the 1.4-2.7 keV image of this source, we can use an exposure map for a single representative energy or we can use spectral weights chosen according to a model for the source spectrum. We examine four likely choices: three monochromatic exposure maps, computed at 1.47 keV, 2.61 keV and at the average energy, 2.0 keV and a weighted exposure map based on the known input spectrum, with equal weights at 1.47 keV and 2.61 keV.

Figure 3 shows the results of applying these four exposure maps to the counts image using equation (5). From the figure, it is clear that *none of the exposure maps gives the correct result for the entire image.* The exposure map computed for the intermediate photon energy produces significant surface brightness errors in all parts of the image. Where there is an exact match with the incident source spectrum, the monochromatic exposure maps work well, but in the overlap region, they produce large errors. Although the weighted map based on the *known input spectrum* produces the correct answer in the overlap region, it gives the wrong surface brightness in each of the monochromatic regions.

This simple example illustrates two important points. First, even with perfect knowledge of the input spectrum, it may not be possible to compute the incident surface brightness at every point

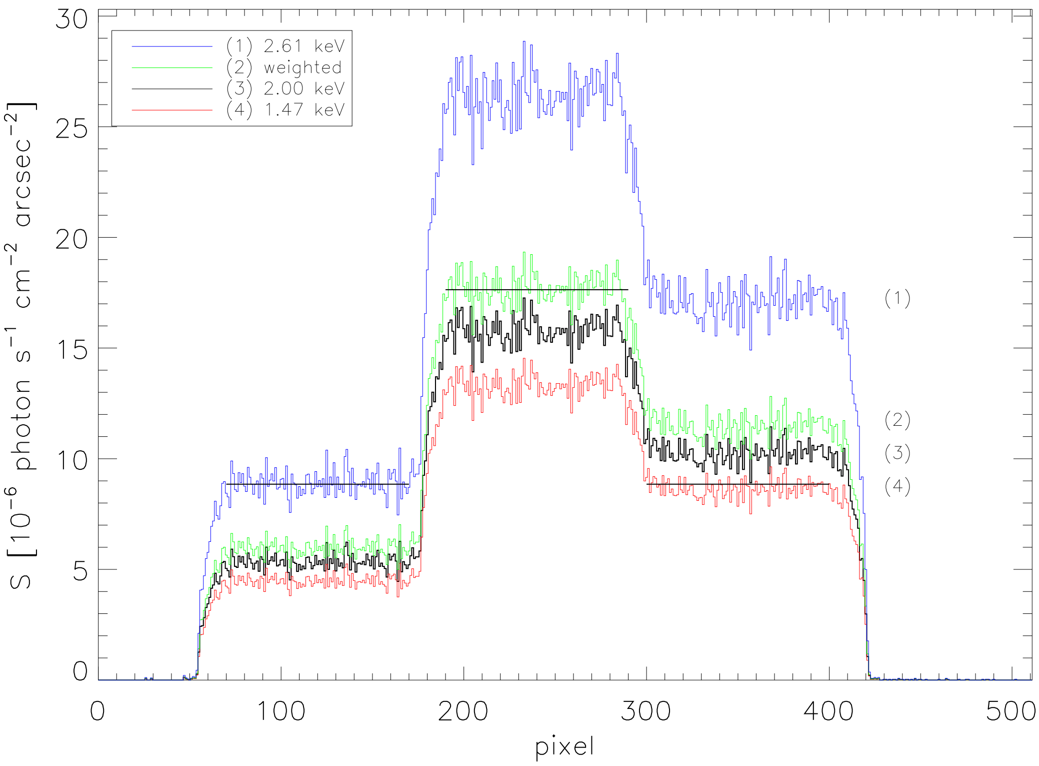

Fig. 3.— Surface brightness profiles generated using four different exposure maps. Straight lines (black) mark the correct surface brightness values in each source region.

in an image using a single exposure map when the incident source spectrum is position dependent. Second, even when the source spectrum is independent of position, an energy-dependent effective area can combine with errors in the presumed form of the incident spectrum to cause significant errors in the resulting surface brightness image. It follows that one should take care to apply exposure maps only to energy bands for which the effective area is relatively constant and that one should be especially careful when interpreting images in which the source spectrum may be position-dependent.

4. Example 2: Thermal Sources

4.1. Source Properties

To examine the more realistic case of sources with thermal spectra, we again consider a simulation consisting of two uniform surface brightness disks, but with each emitting an unabsorbed, solar-abundance, thermal X-ray spectrum. Each disk contributes an integrated flux of 0.1 photon $\mathrm{sec}^{-1}\mathrm{cm}^{-2}$ (0.18-10 keV); one disk has kT = 2 keV and the other has kT = 6 keV. To illustrate the effect of the energy-dependence of the effective area, we consider two energy bands: the 2.5-5.0 keV band for which the effective area curve is roughly constant, and the 1.0-5.0 keV band which includes the Iridium edge at 2 keV. The surface brightness values for the various source regions are given in Table 1.

4.2. Observation and Analysis

Because the effective area is only weakly energy dependent between 2.5-5.0 keV, we expect to obtain fairly accurate surface brightness values in this band using a monochromatic exposure map for E=3.75 keV. Indeed, the lower curve (# 4) in Figure 4 shows that this exposure map gives the correct result. However, because the effective area has rather strong energy dependence in the broader energy band between 1.0-5.0 keV, we do not expect a monochromatic exposure map to give the correct surface brightness in all parts of an image in this band. The upper curve (# 1) in Figure 4 shows that a monochromatic exposure map leads to significant errors in the computed surface brightness ($\sim$ 50%).

We also examine two weighted exposure maps for the 1.0-5.0 keV band, one for each component in the thermal spectrum. The weights computed using equation (7) are given in Table 2. Examination of the two middle curves (# 2 and # 3) of Figure 4 shows that each of these weighted exposure maps yields surface brightness values with errors $\lesssim$ 10%.

Table 1
Incident Surface Brightness

Region	S^a (2.5-5.0 keV)	S^a (1.0-5.0 keV)
"2 keV"	0.45	3.20
"6 keV"	0.94	3.17
"Overlap"	1.39	6.37

[a] Units 10^{-6} photon sec^{-1} cm^{-2} arcsec^{-2}

Table 2
Exposure Map Weights

E (keV)	kT=2 keV	kT=6 keV
1.3	0.626	0.418
1.8	0.123	0.151
2.1	0.056	0.062
3.6	0.195	0.368

Given the results of the earlier examples and the spectral variation present in this image, the comparative accuracy of this last result may seem somewhat surprising. However, because the incident spectra are relatively smooth and have similar shapes, the weighted exposure map and, therefore, the computed flux is relatively insensitive to the exact spectral shape within the chosen band (see equation 9 and e.g. the corresponding curve in Figure 2a).

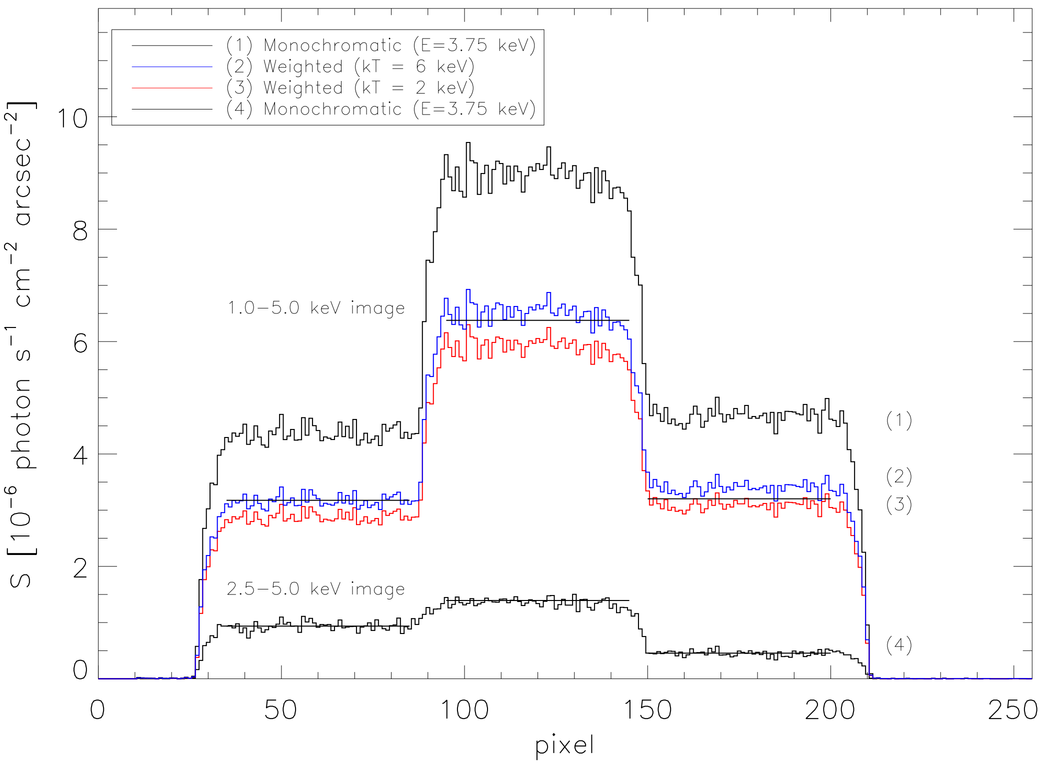

Fig. 4.— Surface brightness profiles for 2 different energy bands computed using monochromatic and weighted exposure maps. The top 3 curves give the surface brightness in the 1.0-5.0 keV band while the bottom curve gives the surface brightness in the 2.5-5.0 keV band where the effective area is roughly constant. Straight lines (black) mark the correct surface brightness values in each source region for each energy band.

5. Summary

Surface brightness images generated by dividing a counts image by an exposure map are likely to contain significant systematic errors unless

1. the effective area is approximately constant over the specified energy band,
2. the incident source spectrum varies slowly across the specified energy band (e.g. thermal, power-law),
3. the contribution to the exposure map from each piecewise-constant band is weighted using an accurate model for the incident spectrum (e.g. according to equation 7), and
4. spectral variations in the image are sufficiently small that a single spectrum model is applicable to the entire image.

6. Appendix

To generate the MARX simulation, we ran the following programs:

- MARX run twice, one for each disk (`SourceType="DISK"`).
- MARXCAT to combine the two disk simulations
- MARX2FITS to create the input FITS event list
- MARXASP to create the aspect solution

We then analyzed these simulated data using the standard Ciao tools:

- ASP_CALC_OFFSETS, ASP_APPLY_SIM, ASPHIST to generate the aspect histogram giving the pointing duration for each sky position.
- MKINSTMAP to generate the instrument maps
- MKEXPMAP to generate the exposure maps
- DMCOPY to create the counts image
- DMIMGCALC to divide the exposure map into the counts image

REFERENCES

Davis, J.E., 2000, ApJ, submitted.