

The EPIC-MOS Particle-Induced Background Spectrum¹

November 13, 2006

Source of Acquisition
NASA Goddard Space Flight Center

K. D. KUNTZ

*The Henry A. Rowland Department of Physics and Astronomy
The Johns Hopkins University, 3400 Charles Street, Baltimore MD, 21218*

and

*Astrophysics Science Division
Code 662, NASA/GSFC, Greenbelt, MD 20771
kuntz@milkyway.gsfc.nasa.gov*

S. L. SNOWDEN

*Astrophysics Science Division
Code 662, NASA/GSFC, Greenbelt, MD 20771
snowden@riva.gsfc.nasa.gov*

ABSTRACT

We have developed a method for constructing a spectrum of the particle-induced instrumental background of the *XMM-Newton* EPIC MOS detectors that can be used for observations of the diffuse background and extended sources that fill a significant fraction of the instrument field of view. The strength and spectrum of the particle-induced background, that is, the background due to the interaction of particles with the detector and the detector surroundings, is temporally variable as well as spatially variable over individual chips. Our method uses a combination of the filter-wheel-closed data and a database of unexposed-region data to construct a spectrum of the “quiescent” background. We show that, using this method of background subtraction, the differences between independent observations of the same region of “blank sky” are consistent with the statistical uncertainties except when there is clear evidence of solar wind charge exchange emission. We use the blank sky observations to show that contamination by SWCX emission is a strong function of the solar wind proton flux, and that observations through the flanks of the magnetosheath appear to be contaminated only at much higher solar wind fluxes. We have also developed a spectral model of the residual soft proton flares, which allows their effects to be removed to a substantial degree during spectral fitting.

Subject headings: methods: data analysis — instrumentation: detectors — X-rays: general

1. Motivation

Although *XMM-Newton* has a large field of view (FOV), compared to other current X-ray observatories, there are often observations in which the object of interest (such as diffuse Galactic emission, nearby galaxies, or a large cluster of galaxies) fills the entire FOV. Determining the background in such a situation can be problematic. The background is due primarily to energetic particles interacting directly with the detector, or interacting with material around the detector and producing fluorescent X-rays that then strike

the detector. This “particle-induced background” has multiple components and each component is temporally variable, though on different scales. Since the particle background is temporally variable, using the particle background derived from another observation is likely to be unsatisfactory.

“Blank sky” observations are often used to subtract a spectrum combining the particle background and the Galactic emission to allow the measurement of the spectrum of clusters of galaxies. Given that the Galactic background varies strongly with Galactic coordinate, use of blank sky data to remove the particle background and Galactic foreground emission in order to study some other object requires that one evaluate each of the blank sky fields for its appropriateness based

¹Based on observations obtained with *XMM-Newton* an ESA science mission with instruments and contributions directly funded by ESA Member States and NASA

on two criteria; the Galactic spectrum and the particle background spectrum. This method works well at $E > 2$ keV where neither the bulk of the Galaxy nor the particle background varies so greatly, but is much more difficult to apply at $E < 2$ keV. Although there is a large amount of blank sky data, there may not be enough to provide a match for any given observation. Since the blank sky is the object for those who study Galactic emission, it was necessary to develop a method for measuring each of the particle backgrounds independently of the emission in the FOV. This paper describes a method to model the spectrum of the quiescent particle background without the use of data from the FOV.

The bulk of the quiescent particle background spectrum can be modeled and subtracted directly from the source spectrum. The remaining particle background components, the strong Al and Si instrumental lines, and the residual contamination by soft proton flares can be fitted simultaneously with the source spectrum. These multiple components lead to a multiplication of fit parameters, but we show that these components can be relatively simply parameterized. N.B.: these methods require summing large areas of the detector, so they are not useful for small extended objects. However, the traditional method of annular background subtraction is often appropriate in these cases.

The design of the MOS cameras provides a measure of the particle background for each observation, the “unexposed pixels”. The unexposed pixels are portions of the outer CCDs that are masked off to prevent the incidence of cosmic X-rays. They do, however, “see” the same background of energetic particles as do the portions of the chips within the FOV. There are, however, four difficulties in using the unexposed pixels to determine the instrumental background. First, a typical observation does not produce sufficient counts within the unexposed pixels to produce a good spectrum of the particle background. Since the particle background is temporally variable, summing background spectra from other observations can be done only with care. Second, the response of the unexposed pixels to the background is not the same as the response of the pixels within the field of view. This spatial variation of the response to the particle background can be studied using data obtained when the filter wheel is in the closed position (FWC data). Third, a portion of the background is due to energetic particles striking material around the detector (the detector housing, filter wheels, etc.) and producing fluorescent X-rays. Since these X-rays have a local source, they are not smoothly distributed across the detector; the unexposed pixels may see more or fewer of these X-rays than do pixels within the field of view. Fourth, and finally, the mask protecting the unexposed pixels from cosmic X-rays also protects them from the softest and most highly temporally variable portion of the particle background, the “soft proton

flares”.

In the following discussion of the particle background, §2 describes our method of data preparation, §3 characterizes the “quiescent particle background” (QPB) observed by the unexposed pixels. This background is “quiescent” in that its properties, strength, and spectral shape, change relatively slowly, on the scale of several months. We do not consider time periods with strongly enhanced particle background rates. Section 3.1 describes the variation in the QPB, including fluorescent X-rays, over the face of the detectors, while §3.2 and §3.3 describes the temporal variation. Section 3.4 provides a prescription for creating a QPB spectrum for any particular observation. Section 4 discusses the spectrum due to the low-amplitude soft proton flares; §4.1 describes our method of determining that spectrum, §4.2 describes the variation of the flare spectrum with flare strength, §4.3 describes the spatial variation of the flare spectrum, and §4.4 describes a method for removing the effect of small-scale residual soft proton flares from the data. Section 4.5 characterises the flares themselves in terms of temporal and orbital distributions, a basic first step to understanding the underlying source of the proton flares. Section 5 applies our background subtraction method to show that the uncertainties in the resulting source spectrum significantly smaller than the statistical uncertainties. However, we do see large variations between different observations of the same sky region which can be attributed to contamination by solar wind charge exchange emission.

2. Data Preparation

For the characterization of the MOS instrumental backgrounds we used every observation publically available as of 1 April 2006. This collection of data contained ~ 3500 observations, ~ 8100 observation segments, and ~ 76 Ms of exposure for each MOS camera.

For each observation segment, the 2.5-8.5 keV light curve was automatically created from the entire field of view (but not the unexposed pixels). The “base” level of the count rate and the r.m.s. of the base level was found iteratively by fitting a Gaussian to the histogram of the count rates. Time intervals with count rates greater than 3.0 times the base level r.m.s. were deemed to be affected by flares and were removed. We then checked to ensure that the automatic measure of the base level was reasonable by examining plots similar to that shown in Figure 1, and removed observation segments where the fit was either in error or poor due to an insufficient amount of flare-free data. After this cleaning, there were ~ 2500 remaining observation segments and ~ 44 Ms of exposure for each MOS camera. Our light-curve cleaning removed $\sim 36\%$ of the total exposure time. Thus $\sim 36\%$ of the total exposure time was clearly affected by soft proton flares, with the frac-

tion of an individual exposure affected by flares ranging from zero to unity. Some of the remaining time is likely to be affected by residual soft proton contamination, but at relatively low levels. Note that if the analysis of point sources were the goal, less time would need to be removed, depending on the intensity of the flaring and the brightness/spectral shape of the source.

To some extent, our cleaning of the unexposed-region data was overkill; rarely does a soft proton flare effect the unexposed-region data. However, entry of the spacecraft into the particle belts will cause the lightcurve to rise in both the FOV and the unexposed pixels, so our cleaning guarantees the removal of those time periods as well.

There are a number of different prescriptions for the energy band to be used in light-curve cleaning (Lumb et al. 2002; Nevalainen et al. 2005; Pradas & Kerp 2005) so it is particularly important to justify our choice. In §4.1 we derive the spectrum of the soft proton flares and show it to be well approximated by a broken power law with a break energy of 3.3 keV. From this spectrum it is clear that the best energy band for light-curve cleaning will be low enough to contain a large fraction of the

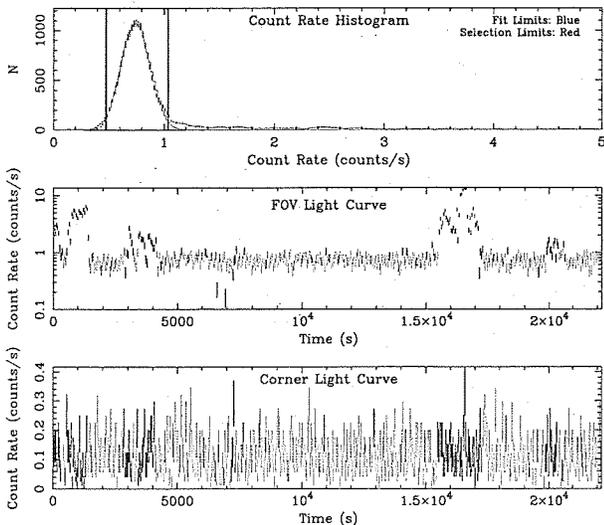


Fig. 1.— **Top:** A histogram of the count rate in the 2.5-8.5 keV band over the entire FOV. The blue curve is a Gaussian fitted to the histogram between the two vertical lines. The peak of the Gaussian defines the quiescent count rate, which will be composed of counts from the object and the quiescent particle background. Parts of the observation falling above 2.5σ above this mean quiescent rate are deemed to be contaminated by the soft proton flares. **Middle:** Light curve in the 2.5-8.5 keV band for the FOV smoothed with a 50 second boxcar function. The green parts of the curve are those accepted as “flare-free”. **Bottom:** Light curve in the same energy for the unexposed pixels. Note that the flares are not seen in the unexposed-region pixels, confirming their soft-proton origin.

flare emission, but high enough that the flare emission is greater than the source emission. The “best” energy band for light curve cleaning depends upon the spectrum of the dominant emission source. However, we have found that the 2.5-8.5 keV band is a good, general purpose energy band for light-curve cleaning. The lower bound is set to avoid the strong instrumental lines and the upper bound is set to exclude the highest energies where the flare spectrum has a very low count rate. The upper energy bound is thus rather arbitrary. We have not yet seen advantages to extending the energy band to lower energies.

3. Quiescent Particle Background (QPB)

Figure 2 shows the mean QPB spectra (as extracted from the unexposed pixel data) from all of the screened MOS1 and MOS2 data. The spectral shape is composed of two parts, lines and continuum. The lines are due primarily to the interaction of the particle background with the detector and the detector environment, and the subsequent fluorescence. As such, the line strengths vary with position on the detector, as the fluorescing surfaces are more or less visible to the detector (see Figure 3). The energy and width of the lines will also shift slightly with the residual instrument gain variations and with changes in the charge transfer inefficiency (CTI), thus direct subtraction of one spectrum from another is likely to produce P-Cygni-like profiles for the strongest lines. We will return to the lines below.

For reasons that will become apparent, we have chosen to characterize the shape of the continuum in two ways. The first is the $(2.5-5.0 \text{ keV}) / (0.4-0.8 \text{ keV})$ hardness ratio of the spectrum across the Al/Si/Au line complex. (Using 0.4-1.2 keV for the lower energy band does improve the signal-to-noise of the hardness ratio, but also decreases the dynamic range of the hardness ratio.) The other is the slope of a power law fitted in the 2.4-12.0 keV energy band exclusive of lines. As can be seen in the lower panel of Figure 2, a power law is not a perfect fit to this region; the continuum seems to fall away from the power law at energies above ~ 8 keV. However, the index of this power law is a useful measure with which to characterize the behavior of the QPB spectrum and the goodness of fit is not important.

3.1. Spatial Variation

When discussing the variation of the QPB spectrum, one should keep in mind that there are two different effects at work; the incident particle background varies with position (as seen in the variation in the line strengths over the detector) and the chips’ response to the particle background may vary. For our purposes it is generally not of interest to isolate which is responsible for the variation, but to characterize the variation as a whole. By QPB spectrum we mean that which is

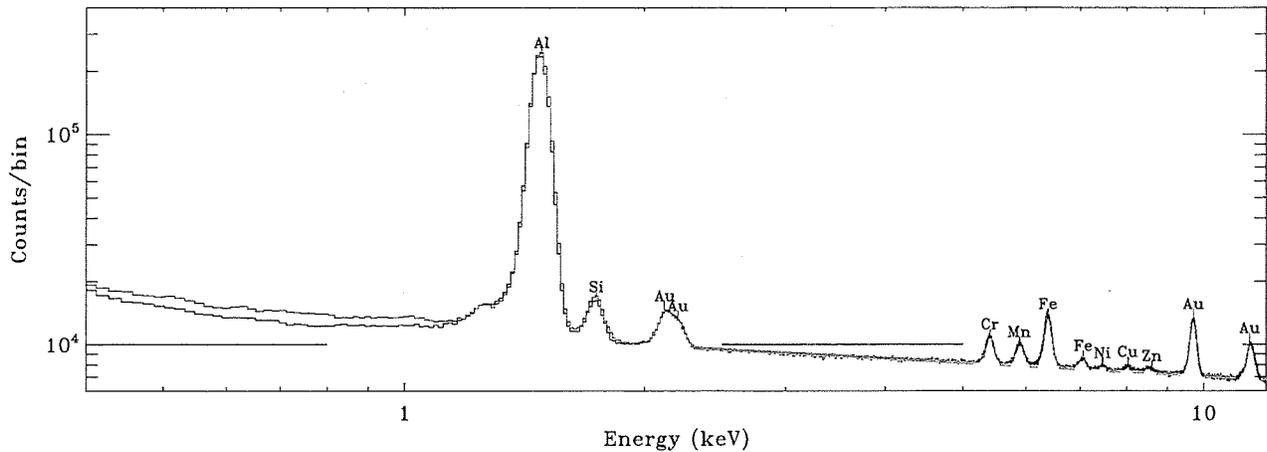


Fig. 2.— The mean QPB spectrum derived from the unexposed pixel data. The MOS1 spectrum is shown in black while the MOS2 spectrum is shown in blue. The heavy red lines indicate the two regions used to measure the hardness ratio. The heavy green line is the fitted power law above 2.4 keV. The prominent background lines are labelled.

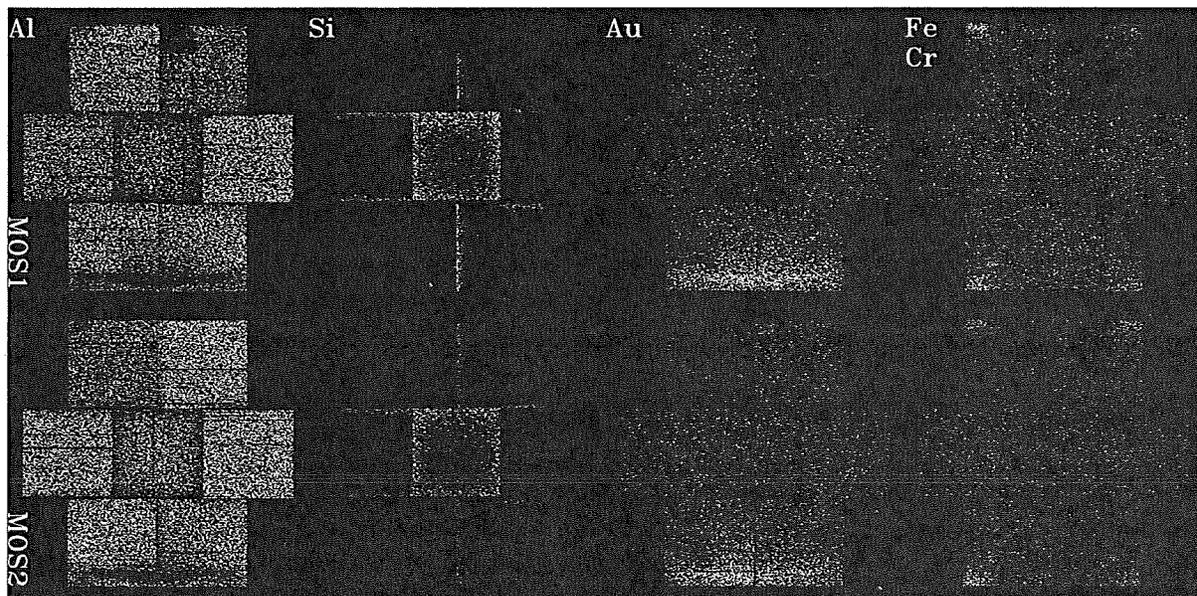


Fig. 3.— Filter-wheel-closed images in narrow energy bands including the principal lines. **Top:** MOS1, **Bottom:** MOS2. From left to right, Al 1.49 keV, Si 1.74 keV, Au 2.12, 9.71, and 11.4 keV, and Fe 6.4 keV + Cr 5.4 keV. The Fe and Cr distributions are similar so adding them together improves the visibility of the spatial distribution.

recorded; by variation we mean only the variation in what is recorded, and generally disregard the ultimate cause of its variation.

3.1.1. Chip-to-chip Variation

The solid lines in Figure 4 show the mean QPB spectra from different chips, as derived from the bulk of the unexposed-region data. Each of the chips has a somewhat different QPB rate. Above 2.5 keV, ignoring the fluorescent lines, the continuum shape shows little chip-to-chip variation. Below 2.5 keV, there are clear differences. Chip MOS1-5 shows an excess in the 0.4-

0.8 keV band while Chip MOS1-4 shows an excess in the 0.8-1.0 keV band. For the MOS2, there appear to be two groups of chips; 2, 5, and 7 in a group having a higher continuum, and 3, 4, and 6 in a group having a lower continuum. (Chip 1 has no unexposed regions.)

Thus, one can see that there are some chips whose continuum below the Al line are significantly different from the continua measured by the other chips. The implications for the analysis of extended sources and the choice of background regions can be significant.

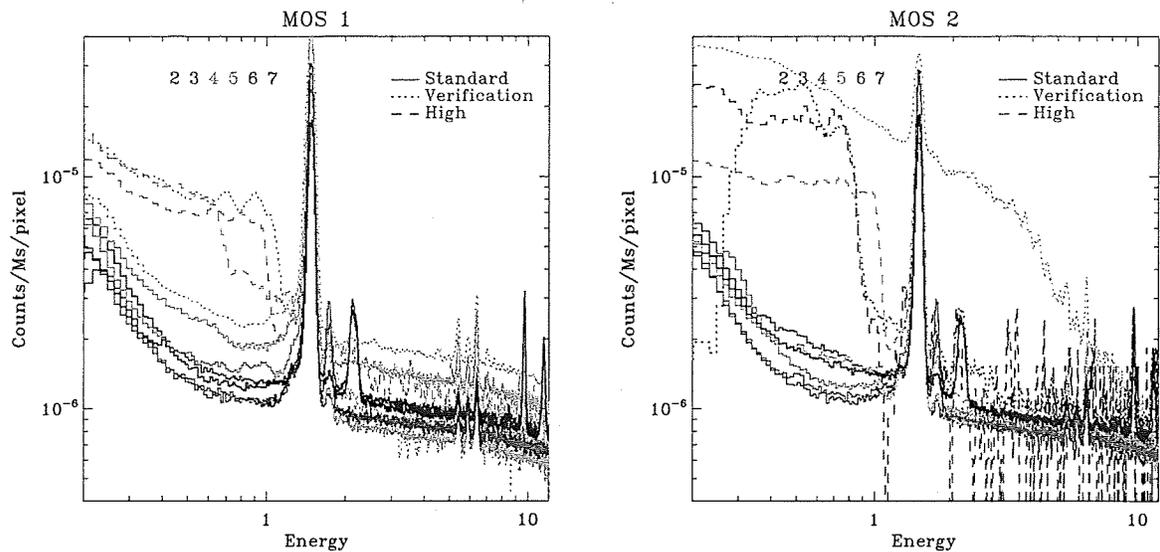


Fig. 4.— The mean QPB spectrum derived from the unexposed-region data for individual MOS chips for both standard and anomalous states. **Left:** MOS 1 **Right:** MOS 2. *Solid lines:* the spectra of the “standard state” unexposed-region spectra, *Dashed lines:* the spectra of the “high state” unexposed-region spectra, and *Dotted lines:* the spectra of the “verification state” unexposed-region spectra. The lines at 1.487 and 1.740 keV are due to Al and Si respectively. Note also that the Au line complexes at 2.2, 9.6, and 11.5keV are seen only in chips 2 and 7.

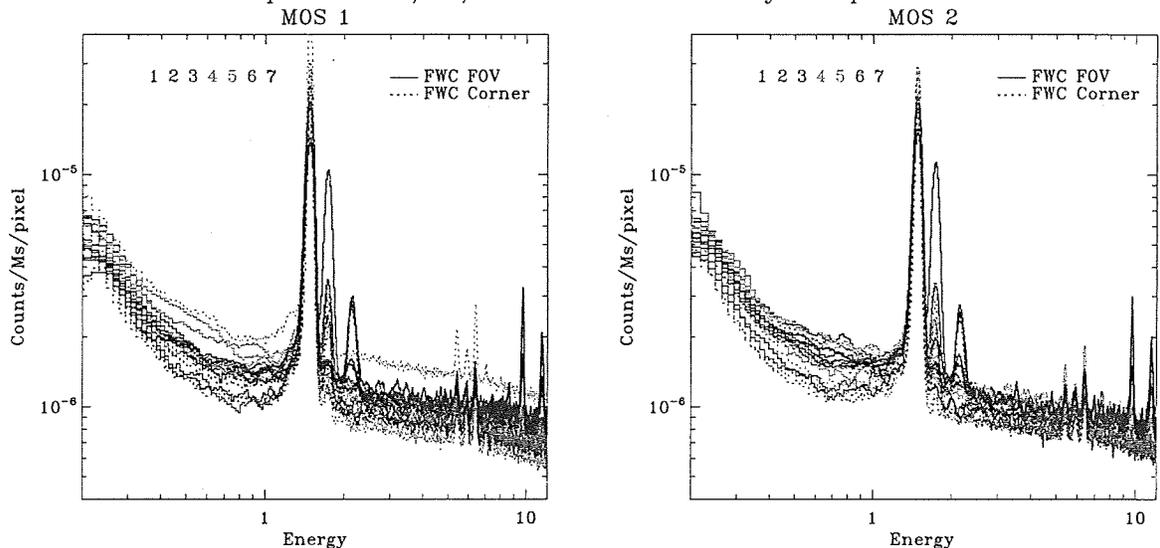


Fig. 5.— The mean QPB spectrum derived from the FWC data for individual MOS chips. **Left:** MOS 1 **Right:** MOS 2. *Solid lines:* the spectra of the FWC in the FOV, *Dotted lines:* the spectra of the FWC in the unexposed-region data.

3.1.2. Region-to-region Variation

The particle background of a given chip varies with the location on the chip. Given the limited amount of FWC data currently extant, we can characterize the variation only over relatively large regions. Of particular interest here is the difference in response between the unexposed region and the FOV region as this will significantly impact our ability to characterize the background. Figure 5 compares spectra extracted from the FOV region of FWC data with the spectra extracted

from the unexposed regions of the FWC data. Significant differences exist below the Al line. The FWC FOV data tend to have higher overall rates per pixel than the FWC unexposed data. Further, chip MOS1-1, which has no unexposed-region data, can be seen to have a continuum shape very similar to that of the MOS1 chips 2, 3, 6, and 7. Similarly, chip MOS2-1 has a continuum shape similar to that of the MOS2 chips 3, 4, and 6.

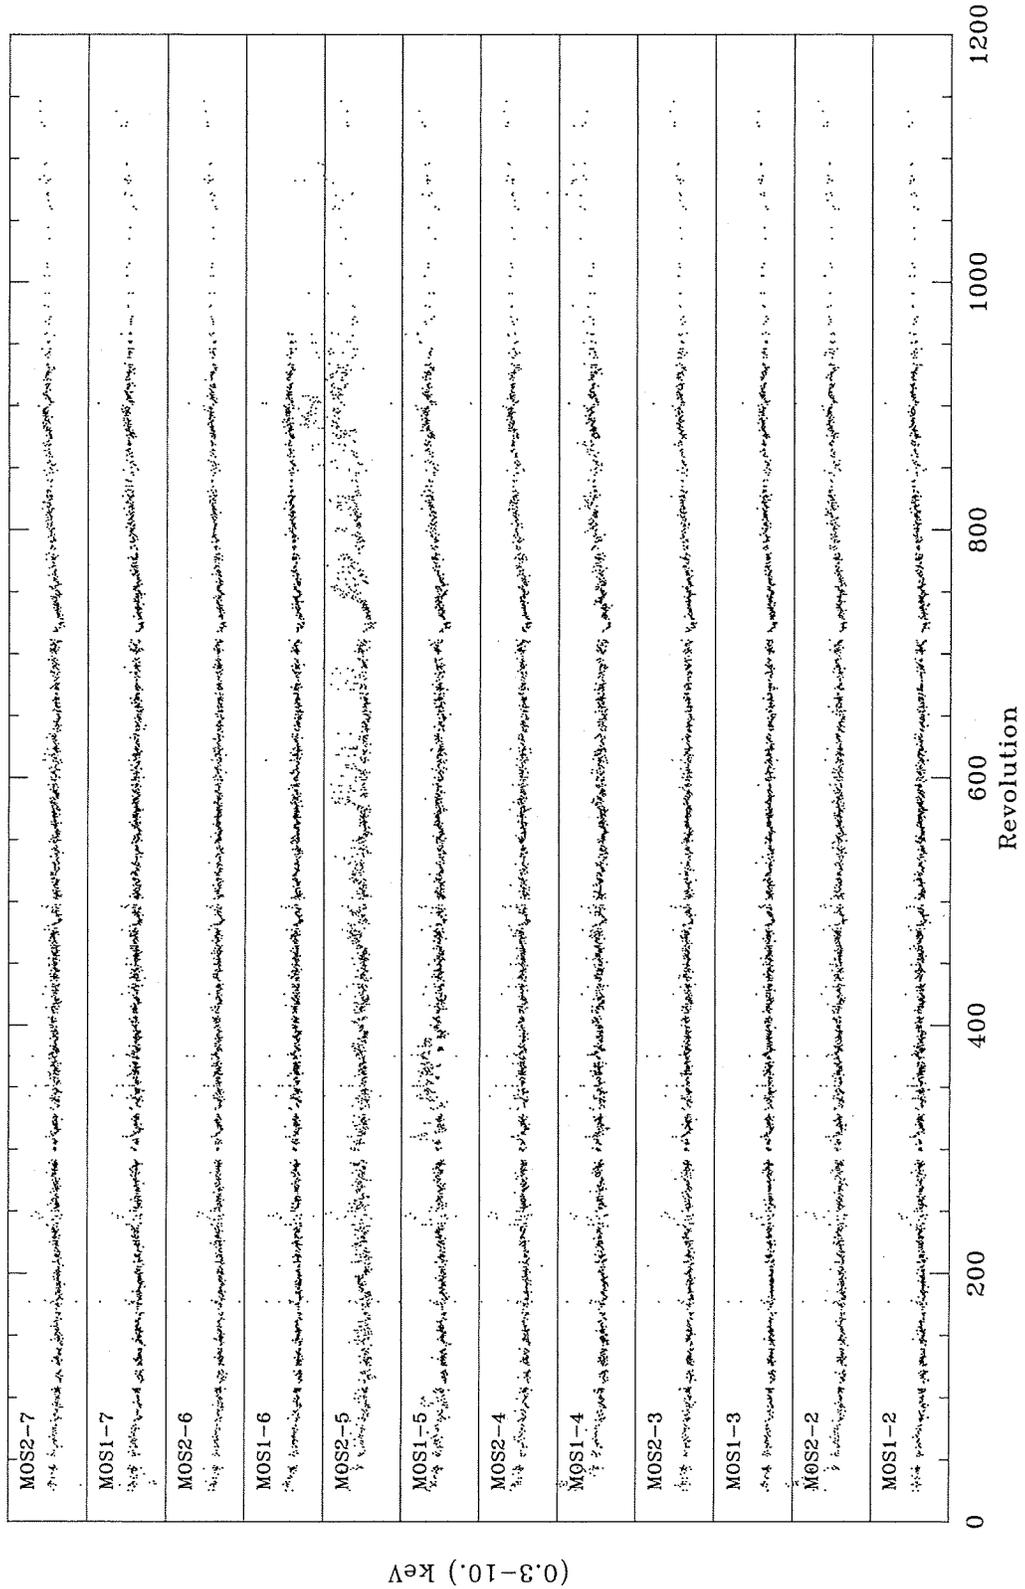


Fig. 6.— The 0.3-10.0 keV rate as a function observation date in revolutions, where one revolution is about two days.

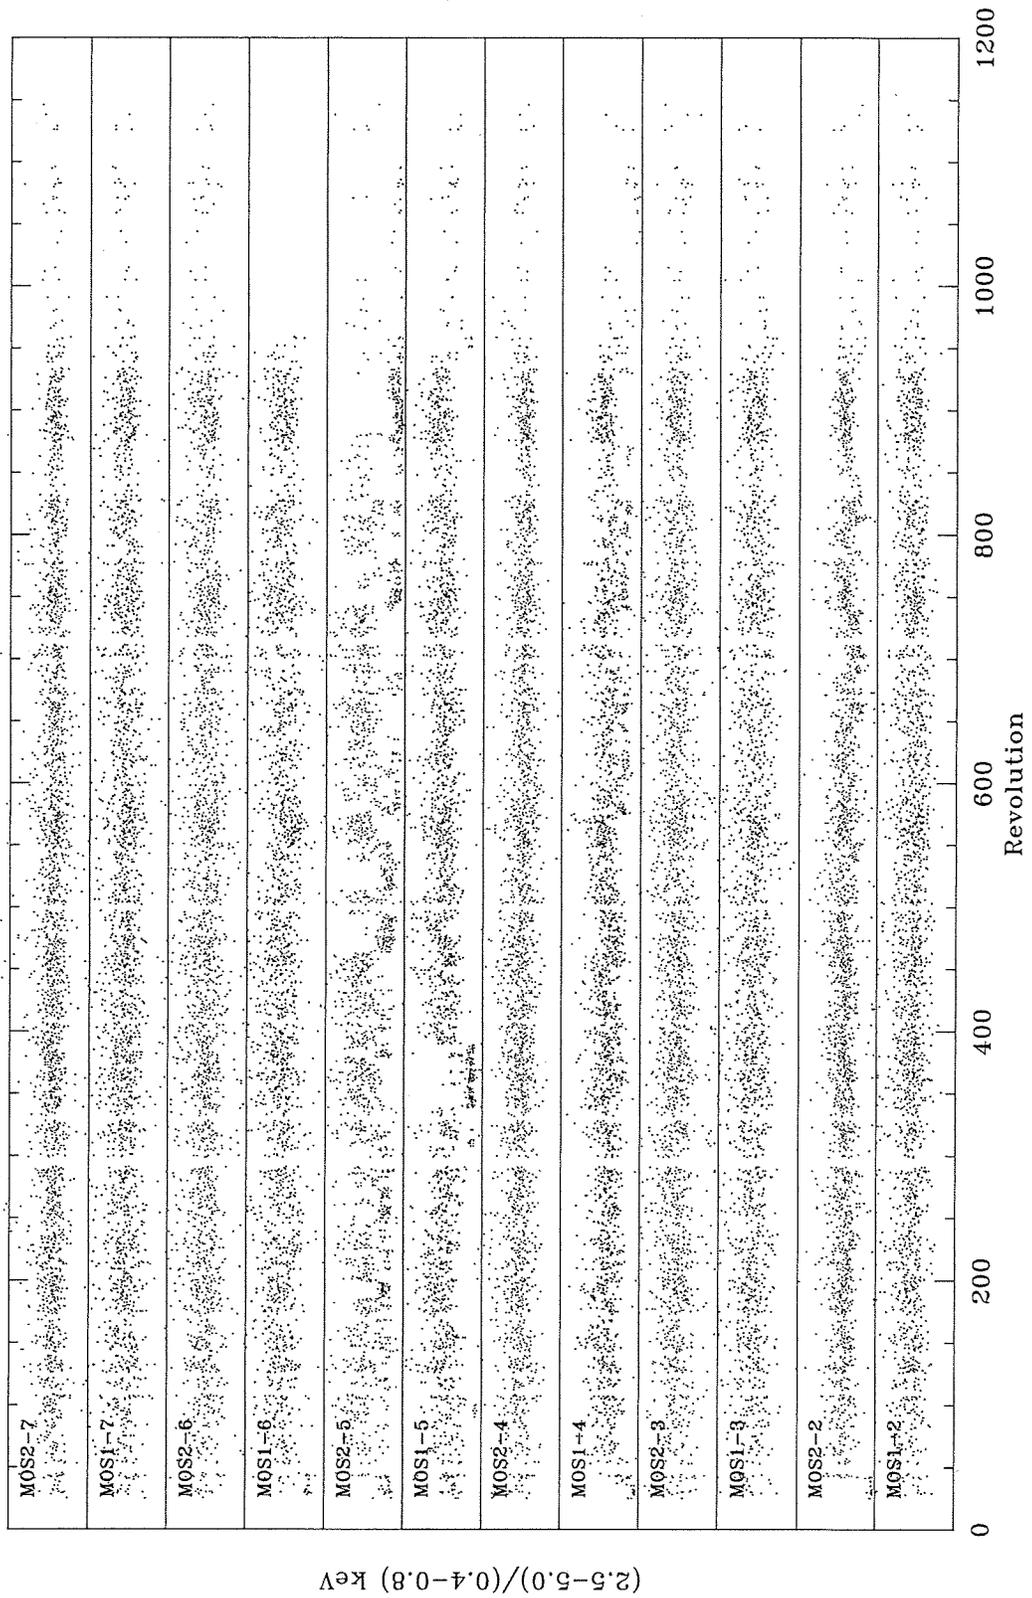


Fig. 7.— The $(2.5-5.0 \text{ keV})/(0.4-0.8 \text{ keV})$ hardness ratio as a function of revolution.

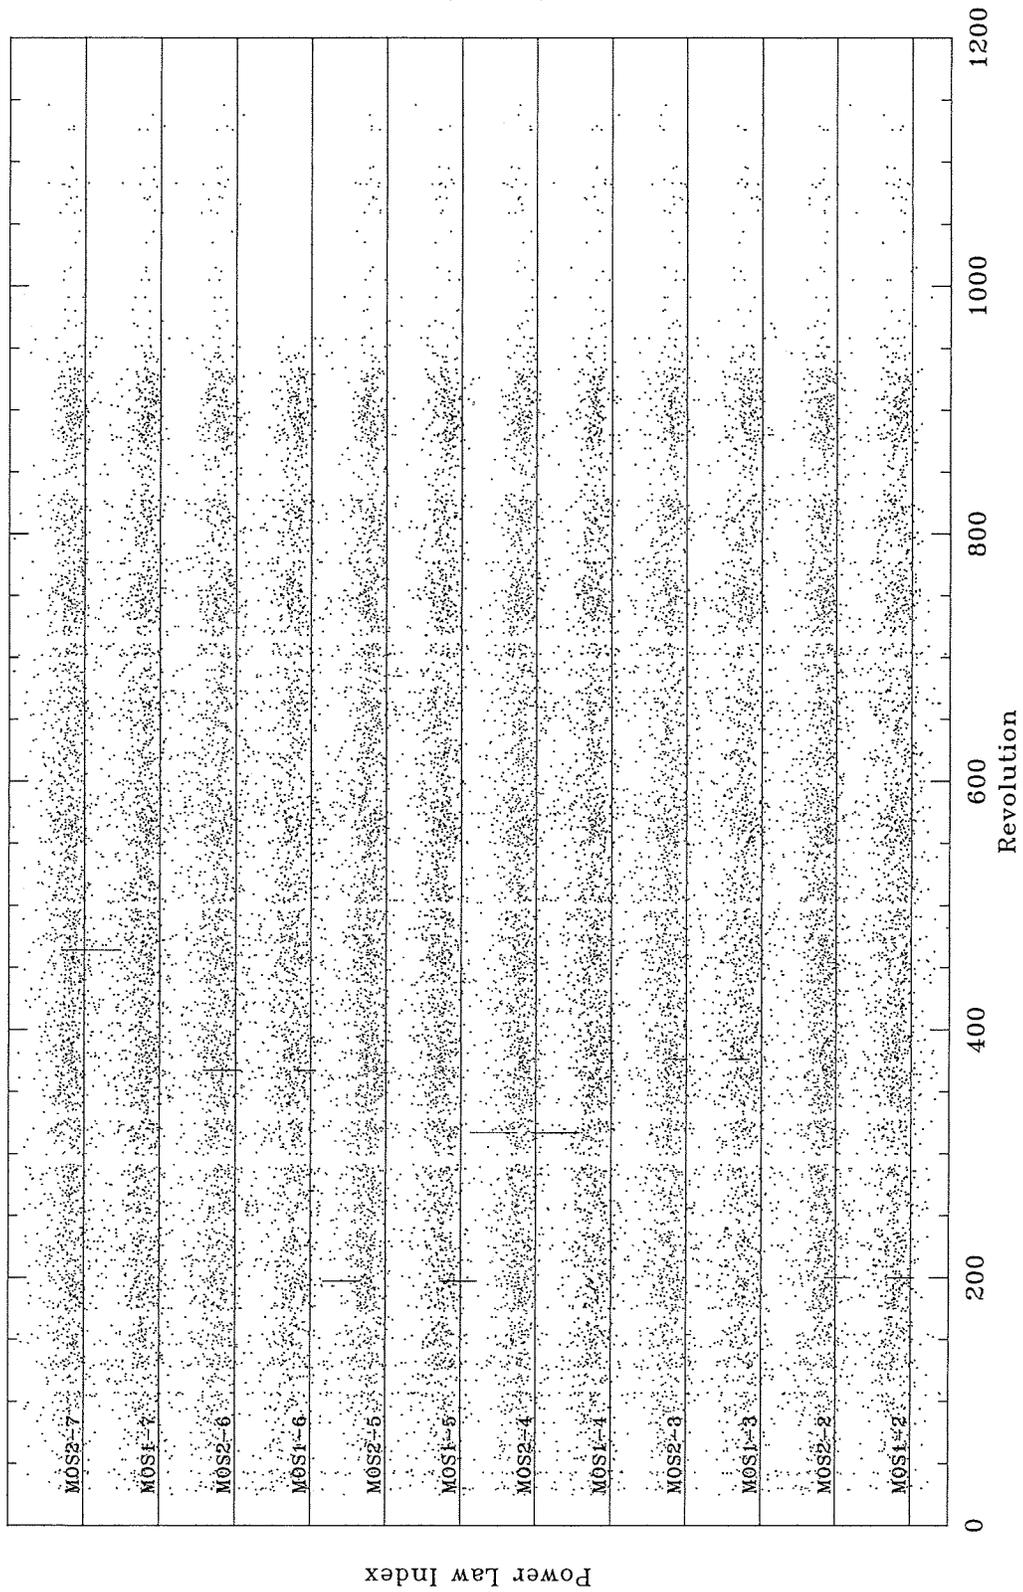


Fig. 8.— The power law index fitted in the 2.5-5.0 keV band as a function of revolution.

3.2. Temporal Variation

Figure 6 shows the temporal variation of the total QPB rate in the unexposed region for the cleaned data. Figure 7 shows the temporal variation in the (2.5-5.0 keV)/(0.4-0.8 keV) hardness ratio, and Figure 8 shows the power law index above 2.4 keV. The total QPB rate is falling before revolution 100, and gently rising thereafter, with an abrupt increase in the rate after revolution 717². Superimposed upon the long-term variation are some short (~ 10 revolution) upward excursions; those affecting all the chips are true variations in the particle background rates while those affecting only one chip are due to variations in the chip response to the particle background. For the bulk of the observations, the hardness ratio is not correlated with the total background rate. However, there are some limited periods for which a *single* chip will have an anomalously low hardness ratio and an anomalously high background rate. These observations are marked in red in Figure 6. As can be seen in Figure 9, these anomalous states occur only for chips MOS1-4, MOS1-5, MOS2-2, and MOS2-5, and, if not isolated in the rate-hardness plane, are at least distinguishable from ordinary observations in the rate-hardness plane.

Although there appears to be a certain amount of variation in the power law index, it is not correlated with variations in rate or hardness, and the power law index does not appear to be strongly affected by the anomalous states.

Figure 10 shows that, after the exclusion of observations in the anomalous states, the distribution of the (2.5-5.0 keV)/(0.4-0.8 keV) hardness ratio is wider than statistically expected. Using the mean QPB spectrum as a “true” spectrum, and the exposure times and count rates of the observed spectra we simulated a set of spectra with Poisson uncertainties. We then measured the (2.5-5.0 keV)/(0.4-0.8 keV) hardness ratio for each of these spectra in each of 1000 simulations. The result, in the left panel of Figure 10 is a narrower distribution than that observed; the statistical model must be convolved with a Gaussian with $\sigma \sim 1.5$ (in hardness ratio) to achieve the observed distribution. Conversely, the right panel of Figure 10 shows the result of the same simulation for the power law index above 2.4 keV. The observed distribution is very similar to the simulated distribution, leading us to conclude that the variation in the power law index is indistinguishable from statistical variation.

²There was no change in the instrument configuration at this time. However, one should note the gap in observations before revolution 717 which is due to the instruments having been put into their safe state due to a strong solar flare on 29 October 2003.

3.3. Anomalous States

All of the anomalous states identified to date are characterized by a low hardness ratio and a high total background rate. Almost all of the anomalous states are characterized by a background spectrum that is strongly elevated at $E \lesssim 1.0$, a sharp break, and a seemingly normal spectrum at $E \gtrsim 2.0$ keV. Mean unexposed-region spectra for chips in an anomalous state are shown in Figure 4 with dotted or dashed lines.

The anomalous periods can be divided roughly into two groups. The first of these is the “high state”, which has been seen in chips MOS1-4, MOS1-5, and MOS2-5, and is shown with dashed lines in Figure 4. The

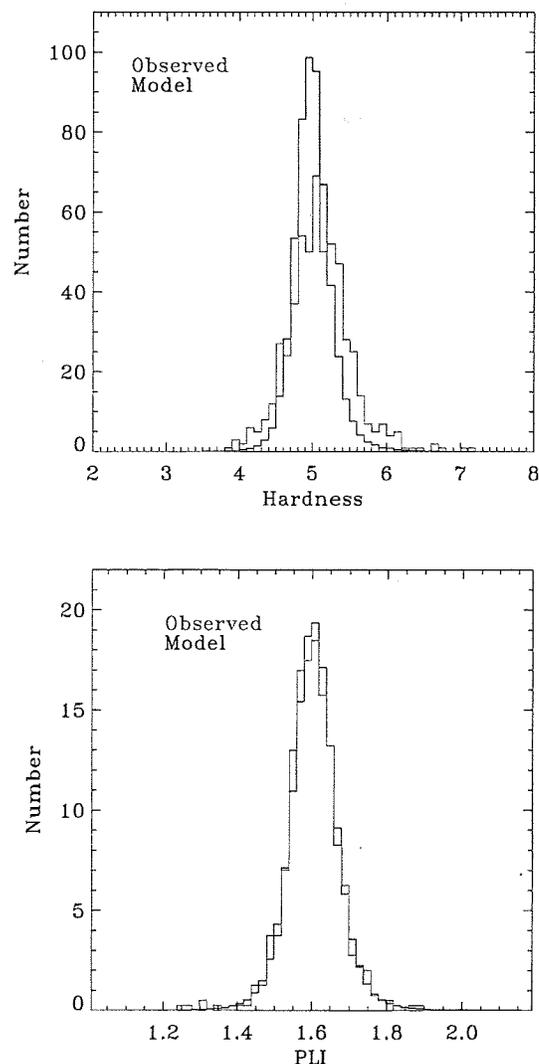


Fig. 10.— *Top*: Observed and statistically expected distribution of the (2.5-5.0 keV)/(0.4-0.8 keV) hardness ratio. The statistically expected distribution was calculated by Monte-Carlo simulation. *Bottom*: Observed and statistically expected distribution of the power law index above 2.4 keV.

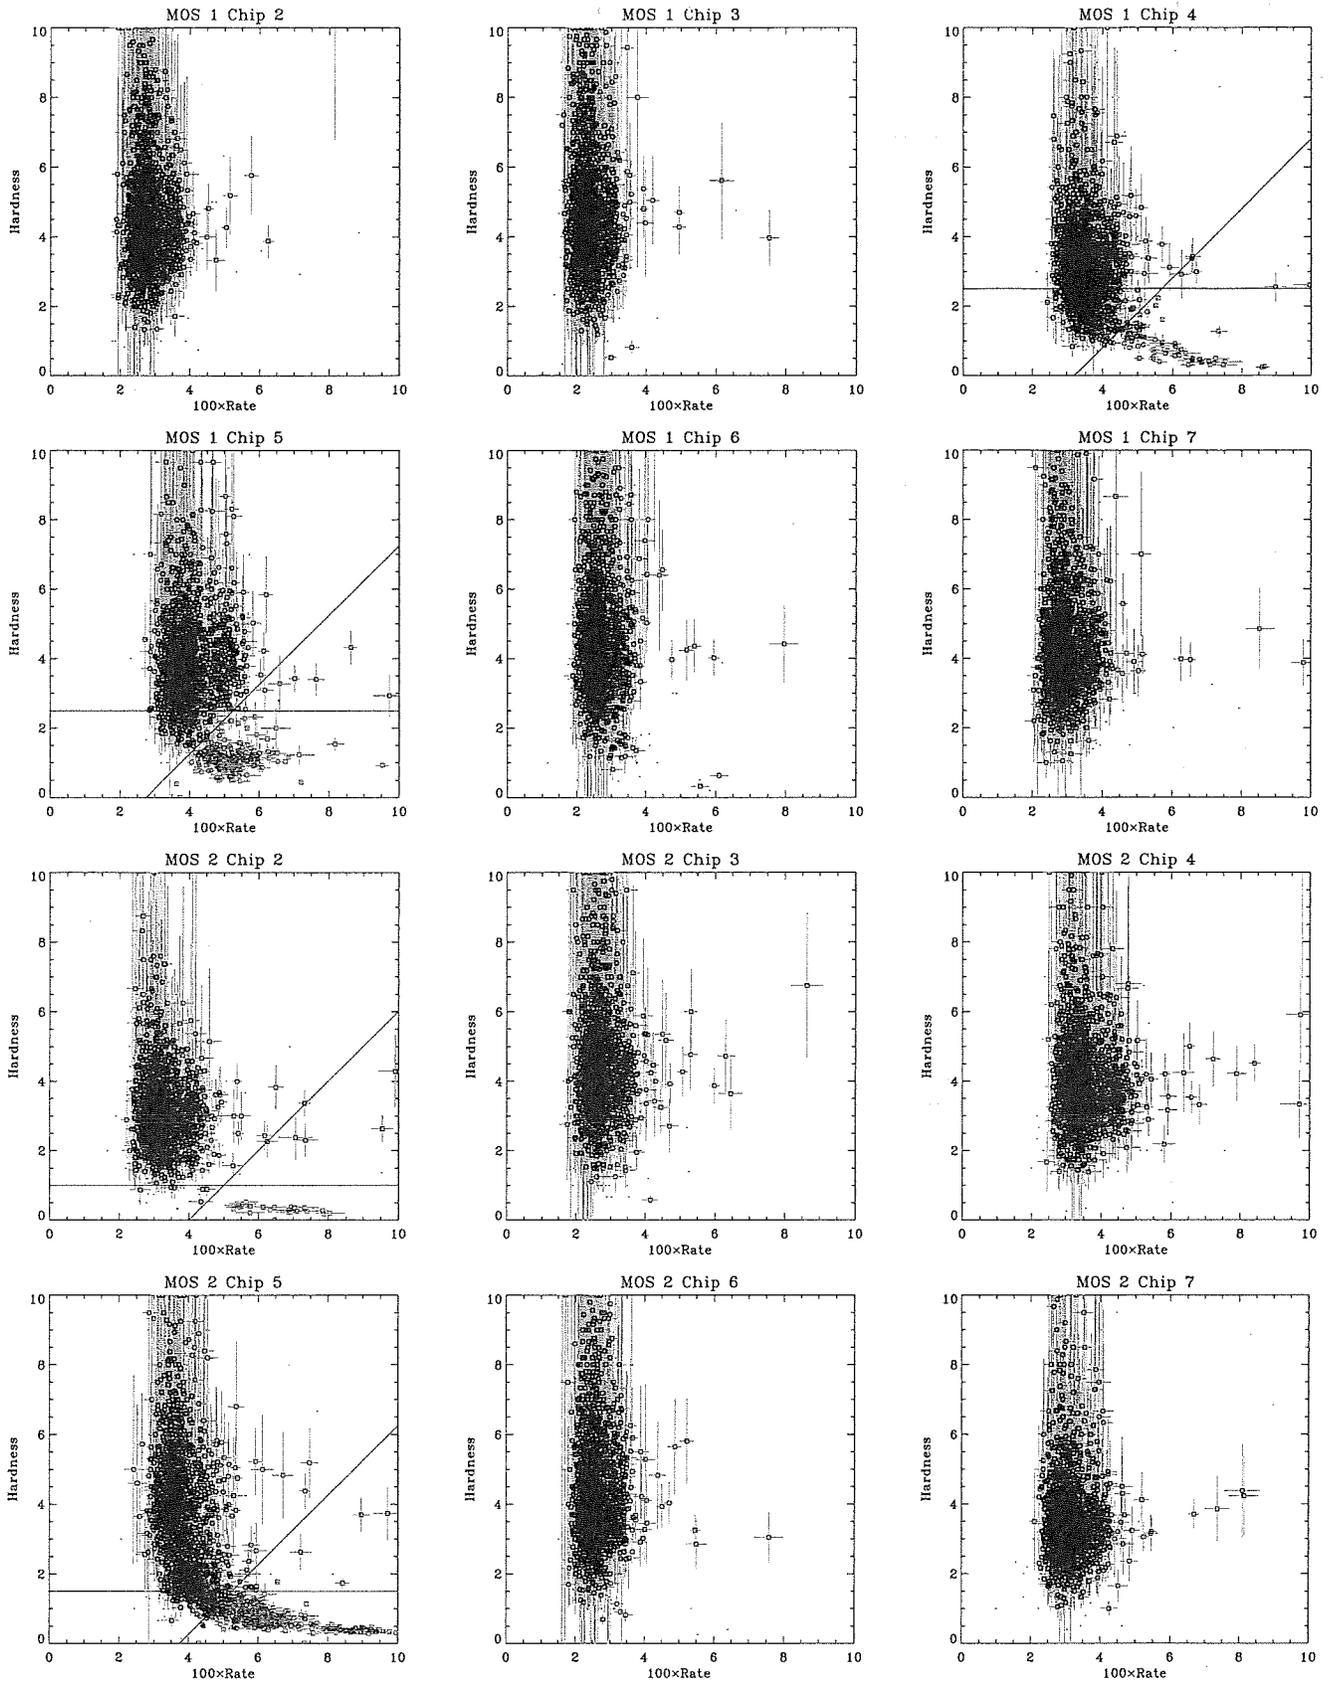


Fig. 9.— The $(2.5-5.0 \text{ keV})/(0.4-0.8 \text{ keV})$ hardness ratio plotted against the 0.3-10. keV rate for each chip. The red points match those in figure Figure 6. The blue lines show the criteria for separating anomalous states from normal ones.

breaks in the spectra typically occur at 0.8-1.0 keV. The periods when MOS2-5 is in a high state are clearly not the same as when MOS1-5 is in a high state. MOS1-5 was in a high state for most of revolutions 308-390, but rarely since then; MOS2-5 began suffering significantly extended high states only after revolution 750. MOS1-4 has had a few minor high states since revolution 800.

The second set of anomalous periods that can be identified occur before revolution 42 and have been named the “verification high state” since they occurred in the very early data. These high states occurred in chips MOS1-4, MOS2-2, and MOS2-5, and are shown with dotted lines in Figure 4. The spectra are noticeably different from the “high state” spectra, particularly for MOS2-5. Since few data were taken in this period, we have not attempted to characterize these states more fully.

All of the anomalous states discussed above were identifiable from the total rate and hardness ratios of the unexposed-region spectra. However, we have also identified a high state for chip MOS1-4 which is not readily identifiable from the unexposed-region spectra. This anomalous state, which we have dubbed the “anonymous state”, was noted while compiling the FWC data; a number of observations showed typical high state spectra in the FOV although the unexposed-region spectra did not. Figure 11 shows the FWC spectra for different portions of the MOS1-4 chip; the inset shows the extracted regions with respect to the chip and the location of the unexposed-region pixels. The unexposed-regions are dominated by those that have a “standard” background spectrum during high state anomalies. Similar problems have not yet been seen with any other chip despite careful checking of the FWC data.

With the exception of the unexposed-region anomaly with chip MOS1-4, the remaining anomalous states are easily identifiable from the rate-hardness diagrams formed from the unexposed-region data. The blue lines in Figure 9 mark the anomalous regions. For MOS1-4, MOS1-5, and MOS2-2 the anomalous states are well separated from the standard state. For MOS2-5, the anomalous points appear to blend smoothly with the normal points, and the strength of the low energy excess does change smoothly as the total rate decreases. We have set the division at the point where the amplitude of the low energy excess is comparable to the uncertainty in the mean spectra. The criteria defining the anomalous states is given in Table 1.

The root sources of the anomalous states are not yet understood, and independent methods of detecting them have not yet been fully developed. However, the anomalous states do appear to affect the pattern distributions. Figure 12 shows the pattern distribution for both the normal (solid lines) and anomalous states (symbols) for each of the CCDs as extracted from the

FWC data and normalized at the value of pattern 1. For the “high” states as well as the “anonymous” state of MOS1-4, patterns 2, 4, and sometimes 0 appear to be elevated. For the “verification” states, pattern 0 is strongly elevated (off of the plot) and patterns 2 and 4 may, or may not, be elevated. If the exposure is sufficiently long, the pattern distribution may be useful to detect anomalous states.

Our anomalous states are a more generalized version of the “bright CCD” problem discussed by Pradas & Kerp (2005) for chip MOS1-5. They also suggested that chip MOS1-2 was subject to a similar type of problem. Based on the rate-hardness diagrams in Figure 9 MOS1-2 does not display anomalous states. However, the MOS1-2 problems observed by Pradas & Kerp (2005) can be understood in terms of the soft proton flares discussed in §4.3.

3.4. QPB Removal Procedure

To summarize the preceding sections: 1) The continuum QPB spectrum is chip dependent. 2) The continuum QPB spectrum varies significantly across each chip. 3) The continuum QPB spectrum is temporally variable. This variation is best characterized as the variation of the (2.5-5.0 keV)/(0.4-0.8 keV) band ratio. The variation of the continuum at higher energies is consistent with statistical fluctuation. In the case of anomalous states, the variation of the spectral shape with chip position can be extreme, however, in most cases anomalous states can be detected with the unexposed-region spectra.

This analysis has ignored the strong Al lines at 1.486

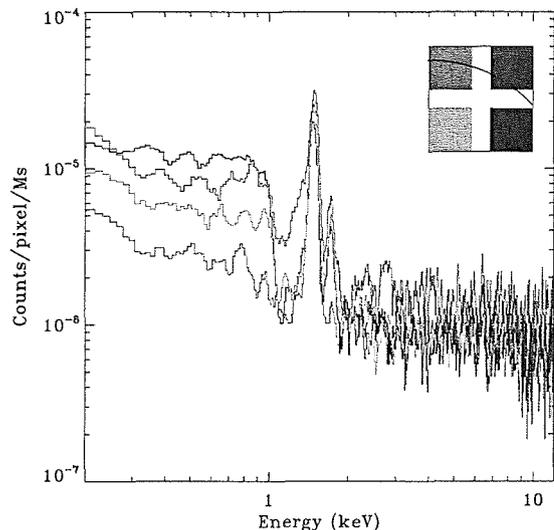


Fig. 11.— The spectrum from the four quadrants of chip MOS1-4 during the “anomalous” period. The inset shows the location of the quadrants with respect to the unexposed pixels (above the curved line).

TABLE 1
ANOMALOUS HARDNESS PERIODS

Detector-Chip	Criterion ^a	Code ^b (revs)	Interval
MOS1-4	H<2.5 and H<100R-3.20 and Rev< 42	Verification	Rev< 42
MOS1-4	H<2.5 and H<100R-3.20	High	
MOS1-4		Anonymous ^c	
MOS1-5	H<2.5 and H<100R-2.75	High	308-313
MOS2-2	Rev< 42	Verification	Rev< 42
MOS2-5	H<1.5 and R>0.1	Verification	Rev< 42
MOS2-5	H<1.5 and H<100R-3.75 and R<0.1	High	

^aH stands for the (2.5-5.0 keV)/(0.4-0.8 keV) hardness ratio, R stands for the 0.3-10. keV count rate, and Rev is the revolution number.

^bThe FWC files provided for background construction and described in §3.4 include the first letter of this code in their name.

^cThe A and H spectra for MOS1-4 are similar; the difference may be the extent to which the chip is affected. At this time A and H are combined for greater statistical significance. As more FWC data become available, the two states may be found to be truly different and the A and H files will be changed accordingly.

and 1.487 keV and the strong Si lines at 1.739 and 1.740 keV. These lines are sufficiently strong that small gain and strength changes between the object and background spectrum can cause large residuals. Further, the strengths of the Al and Si lines have a very strong spatial dependence (Figure 3 and Lumb et al. 2002). Our solution is to exclude these lines from the background analysis, and to interpolate the QPB spectrum over the region occupied by these lines (1.2-1.9 keV).

The Al and Si lines can then be fitted simultaneously with the source spectrum.

Given the variation in the QPB spectrum from chip to chip, the background spectrum used for any particular observation will depend upon the amount of object area to be extracted from each chip. Since the shape of the QPB spectrum is temporally variable, the QPB spectrum must be determined from the observation of interest, not from some temporally averaged spectrum. Since the active area outside of the FOV is rather small (see Table 2), and many exposures are rather short, the number of counts from which to derive the QPB spectrum may be quite small. The procedure outlined here uses the unexposed-region data from a given observation in conjunction with databases containing the unexposed-region data from “all” public observations, as well as the FWC data, to construct a QPB spectrum appropriate for that observation.

The first step is to derive the QPB spectra from the unexposed-region data for each chip. The selection statements in Table 3 should be used to extract the events that are outside of the FOV and are otherwise free of contamination by scattered X-rays. (The best method is to extract all of the unexposed-region events into a single event file using the “corner” criteria listed in Table 3 to exclude all pixels within the FOV or affected by light leaks, then use each of the chip criteria to extract spectra for the individual chips).

The second step is to improve the statistics of these spectra by augmenting them with other data. For each chip we measure the QPB rate (0.3-10.0 keV) and the (2.5-5.0 keV)/(0.4-0.8 keV) hardness ratio. From the database of public observations, we extract unexposed-region data with similar count rates and hardness ra-

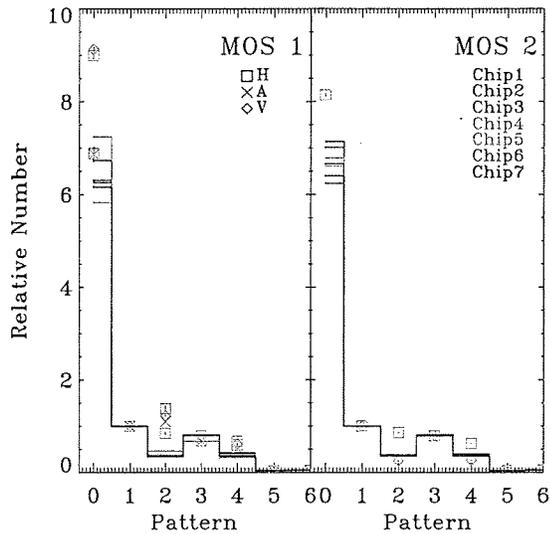


Fig. 12.— The distribution of pattern values for each state of each chip. All of the histograms have been normalized to unity for pattern 1. Normal states are shown by solid lines; anomalous states by symbols. The different chips are shown with different colors.

TABLE 2
RELATIVE CORNER AND FOV AREAS^a

Detector	Chip	FOV area (0.05'' pixels) ^b	Corner area	FOV area	Corner area (^{''} ²)
MOS 1	1	171522400	0	119	0
	2	106938197	33140127	74.3	23.0
	3	129397946	32245645	89.8	22.4
	4	107985152	51360274	75.0	35.7
	5	108333543	32537351	75.2	22.6
	6	127635245	32537351	88.6	22.6
	7	106611615	36668342	74.0	25.5
MOS 2	1	170966400	0	119	0
	2	108370216	39447693	75.2	27.4
	3	124263333	37076690	86.3	25.7
	4	106179861	45497553	73.7	31.6
	5	105525519	50773776	73.3	35.3
	6	127139313	34921667	88.3	24.3
	7	109727561	38608547	76.2	26.8

^aThe values are approximate and depend upon which bad pixel lists are used, and thus upon the date of the observation.

^bThe "BACKSCAL" parameter in the OGIP-compatible spectra produced by SAS are in these units.

tios. This creates a spectrum with approximately the same shape as that seen in the unexposed pixels during the observation, but with a much better signal-to-noise ratio.

The third step requires an assumption. Since there is not yet a sufficient amount of FWC data to study the temporal behavior of the region-to-region variation in any but the grossest manner, we assume that the ratios of spectra from different regions of a chip are temporally invariant. Clearly, anomalous states (particularly the anonymous state) need to be handled separately. From the observation data, we extract the object region. **Note:** because of SAS³ requirements, to support the following steps, one must create a definition of the object region in detector, not sky, coordinates. From the FWC data, extract spectra from the object region on a chip-by-chip basis. From the FWC data extract the unexposed-region spectra on a chip-by-chip basis, using the same criteria used to extract the unexposed-region spectra from the observation. For each chip, multiply the unexposed-region spectrum from the observation data by the ratio of the FWC object region spectrum to the FWC unexposed region spectrum on a spectral bin by bin basis. This "corrects" the unexposed-region spectra to the shape of the spectrum in the region of interest.

Given the limited amount of FWC data, taking the ratio of two spectra directly is unwise. Ideally, one would fit both spectra with a continuum-plus-line

model, and then take the ratio of the models in order to minimize the uncertainty in the ratio. We have tried such a method and found it very time consuming, with significant uncertainties for the weaker lines. Instead of fitting, we smooth both spectra by a small amount (9 channels, or 0.135 keV) and take the ratio of the smoothed spectra. Smoothing by a larger amount causes a significant spreading of the Al/Si lines into the surrounding spectral regions, as well as significant loss of the weaker lines.

The fourth step concerns the center CCDs (chips MOS1-1 and MOS2-1) which have no unexposed region, and for which this augmentation procedure is, clearly, not applicable. From Figure 5, one can see that the QPB spectrum derived for chip 1 from the FWC data is similar (but not statistically identical) to the spectra derived from some of the other chips. For chip MOS1-1, the similar chips are MOS1-2, 3, 6, and 7; for chip MOS2-1, the similar chips are MOS2-3, 4, and 6. We can then use the same method applied to the other chips, this time using the ratio of the FWC object region spectrum from chip 1 and the FWC unexposed-region spectra from the "similar" chips. We then multiply this ratio by the unexposed-region spectra from the "similar" chips derived from the observation data.

The fifth step is to combine the augmented and corrected unexposed-region spectra from the different chips, correctly weighted, to form a single background spectrum for the object region. This background spectrum is then subtracted from the source spectrum before spectral analysis.

The strong Al and Si lines remain in the spectrum

³SAS is the *XMM-Newton* Standard Analysis Software. We have used SAS v6.5.0 for this analysis.

TABLE 3
CORNER PIXEL SELECTION CRITERIA

Detector	Region	SAS Selection Expression	Comment	
MOS 1	!corner ^a	CIRCLE(100,-200,17700,DETX,DETY)	FOV	
		CIRCLE(834,135,17100,DETX,DETY)	FOV	
		CIRCLE(770,-803,17100,DETX,DETY)	FOV	
		BOX(-20,-17000,6500,500,0,DETX,DETY)	FOV extension	
		BOX(5880,-20500,7500,1500,10,DETX,DETY)	light leak	
		BOX(-5920,-20500,7500,1500,350,DETX,DETY)	light leak	
		BOX(-20,-20000,5500,500,0,DETX,DETY)	light leak	
		BOX(-12900,16000,250,4000)	scattered light(?) UL edge	
		BOX(80,18500,150,1300)	scattered light(?) UC edge	
		BOX(-10,-18800,125,1500)	scattered light(?) LC edge	
		chip 1	(CCDNR==1)&&BOX(4.5,4.5,6633,6633,0,DETX,DETY)	
		chip 2	(CCDNR==2)&&BOX(6529,-13572,6599,6599,0,DETX,DETY)	
		chip 3	(CCDNR==3)&&BOX(13295,-306,6599,6599,0,DETX,DETY)	
		chip 4	(CCDNR==4)&&BOX(6529,13027,6599,6599,0,DETX,DETY)	
chip 5	(CCDNR==5)&&BOX(-6435.5,13094,6633,6599,0,DETX,DETY)			
chip 6	(CCDNR==6)&&BOX(-13169,-105,6599,6599,0,DETX,DETY)			
chip 7	(CCDNR==7)&&BOX(-6502.5,-13438,6633,6599,0,DETX,DETY)			
MOS 2	!corner ^a	CIRCLE(435,1006,17100,DETX,DETY)	FOV	
		CIRCLE(-34,68,17700,DETX,DETY)	FOV	
		BOX(-20,-17000,6500,500,0,DETX,DETY)	FOV extension	
		BOX(5880,-20500,7500,1500,10,DETX,DETY)	light leak	
		BOX(-5920,-20500,7500,1500,350,DETX,DETY)	light leak	
		BOX(-20,-20000,5500,500,0,DETX,DETY)	light leak	
		chip 1	(CCDNR==1)&&BOX(5.6,6599,6633,0,DETX,DETY)	
		chip 2	(CCDNR==2)&&BOX(6673,-13427,6633,6599,0,DETX,DETY)	
		chip 3	(CCDNR==3)&&BOX(13372,-228,6633,6599,0,DETX,DETY)	
		chip 4	(CCDNR==4)&&BOX(6571,13104,6599,6599,0,DETX,DETY)	
		chip 5	(CCDNR==5)&&BOX(-6628,13172,6599,6599,0,DETX,DETY)	
		chip 6	(CCDNR==6)&&BOX(-13226,-61,6633,6633,0,DETX,DETY)	
		chip 7	(CCDNR==7)&&BOX(-6628,-13427,6599,6599,0,DETX,DETY)	

^aNote that this region is to be *excluded* in order to isolate the unexposed pixels and remove region affected by light leaks.

and must be fitted in the spectral analysis. We have found that the Al and Si lines are reasonably well removed by Gaussians of small width folded through the detector response. If there are a large number of fit parameters, it is best to fix the line energies to the expected values and fix the line widths to zero before fitting, allow the rest of the fit to approach its final value, and then allow the line energies and widths to float. To simplify the fitting further, the normalizations of the lines in the two detectors are relatively close for a single observation, so those parameters can be linked for the initial fitting. It is unwise to merely exclude the region containing these lines as the low energy wings are quite extensive and so will affect the lower energy data.

Table 4 shows the total exposure time of the currently available FWC observations. The FWC data sets were extracted from the entire public archive as of 1 April 1006. Using our understanding of anomalous states, we have divided that data into “standard” data sets, and the various anomalous states on a chip-by-chip basis.

The unexposed-region spectrum databases contain roughly 42 Ms of data from ~ 2200 observation segments, and provide good temporal coverage between revolutions 50 and 950 (April 2000-April 2006) and lesser coverage after that time. The unexposed-region spectrum databases will be updated from the public archive on a regular basis.

Software implementing the method described above is available in several different formats. The most current and most easily used versions are available through the *XMM-Newton* Guest Observer Facility website at Goddard (<http://heasarc.gsfc.nasa.gov>) and through our friends across the water at the *XMM-Newton* Science Operations Centre (<http://xmm.esac.esa.int>).

4. Soft Proton Flares

The light curve from the FOV can be represented as

$$\text{object} + \text{cosmic background} + QPB + \text{soft proton flares.} \quad (1)$$

The light-curve cleansing method described in §2 finds the minimum of the light curve, but does not guaran-

TABLE 4
AVAILABLE BACKGROUND DATA STATISTICS

Chip	Standard State			High State			Verification State			Anonymous State		
	Time (ks)	Counts	Number obs.	Time (ks)	Counts	Number obs.	Time (ks)	Counts	Number obs.	Time (ks)	Counts	Number obs.
MOS1-1	1030.1	213052	77	0.0	0	0	0.0	0	0	0.0	0	0
MOS1-2	1044.5	239881	79	0.0	0	0	0.0	0	0	0.0	0	0
MOS1-3	1044.7	213999	79	0.0	0	0	0.0	0	0	0.0	0	0
MOS1-4	776.5	150202	59	12.4	3545	2	106.4	30034	10	148.1	31325	8
MOS1-5	983.3	204621	70	60.4	16826	9	0.0	0	0	0.0	0	0
MOS1-6	914.8	202275	67	0.0	0	0	0.0	0	0	0.0	0	0
MOS1-7	1042.7	216287	79	0.0	0	0	0.0	0	0	0.0	0	0
MOS2-1	1018.8	230266	84	0.0	0	0	0.0	0	0	0.0	0	0
MOS2-2	921.9	221291	75	0.0	0	0	0.0	0	0	0.0	0	0
MOS2-3	1028.4	206075	85	0.0	0	0	0.0	0	0	0.0	0	0
MOS2-4	1029.1	220184	85	0.0	0	0	0.0	0	0	0.0	0	0
MOS2-5	695.2	155079	62	228.2	62164	14	93.7	94387	9	0.0	0	0
MOS2-6	1028.6	209553	85	0.0	0	0	0.0	0	0	0.0	0	0
MOS2-7	1028.1	233308	84	0.0	0	0	0.0	0	0	0.0	0	0
Corner Data												
MOS1-2	42375.8	1576523	2460	0.0	0	0	0.0	0	0	0.0	0	0
MOS1-3	42384.7	1225807	2460	0.0	0	0	0.0	0	0	0.0	0	0
MOS1-4	41047.8	1937938	2373	177.5	20185	21	0.0	0	0	0.0	0	0
MOS1-5	38697.0	1995104	2241	3641.4	250700	214	0.0	0	0	0.0	0	0
MOS1-6	41616.1	1424851	2460	0.0	0	0	0.0	0	0	0.0	0	0
MOS1-7	42306.4	1588072	2460	0.0	0	0	0.0	0	0	0.0	0	0
MOS2-2	44233.5	1959701	2497	203.1	15251	24	0.0	0	0	0.0	0	0
MOS2-3	44458.7	1534372	2527	0.0	0	0	0.0	0	0	0.0	0	0
MOS2-4	44483.7	1939511	2527	0.0	0	0	0.0	0	0	0.0	0	0
MOS2-5	37666.9	1834848	2151	6523.7	565654	341	185.2	41432	15	0.0	0	0
MOS2-6	44457.4	1445426	2527	0.0	0	0	0.0	0	0	0.0	0	0
MOS2-7	44446.0	1787256	2527	0.0	0	0	0.0	0	0	0.0	0	0

tee that the soft proton contamination is completely removed. The soft proton flares typically have strong variability on scales shorter than a few ks, and thus are readily identified. However, the soft proton contamination can occur with lower amplitude over longer time scales. The light curves in Figure 13 demonstrate that a short (10 ks) observation with only low amplitude fluctuations in the soft proton contamination could produce a very misleading QPB level. Worse, the user seeing a flat light curve, might assume the absence of flares and thus obtain a spurious emission component in spectral analysis.

The lower panel of Figure 13 shows two spectra of the same portion of sky. Flare intervals were removed from both observations. However, there was significant flare contamination remaining in the second observation. If the source does not have significant emission at $E \sim 3$ keV, then the flare contamination is readily seen as a strong excess above the expected extragalactic background.

The following section describes the spectrum of the soft proton flares, how it varies with flare strength, and how it varies with detector position. Although the spectrum can be modeled, this prescription should not be

taken as an indication that diffuse emission spectra can be adequately analyzed in the presence of flares; inclusion of flares introduces a significant uncertainty in both count rate and spectral shape.

4.1. Spectrum

We determined the shape of the spectrum of the soft proton flares by differencing spectra extracted from intervals with flares and spectra extracted from intervals without flares. For each observation segment in the public archive, spectra were extracted from the entire field of view from flareless intervals as well as intervals where the flare strength was 1.0-2.0 counts s^{-1} above the quiescent level. For each observation, the spectrum from the flareless interval was scaled by the exposure time and subtracted from the spectrum accumulated from the flare interval.

This method assumes that the underlying QPB does not change significantly during background flares, that is, that the background flares are an additional component rather than a modification of the QPB. This assumption seems to be well supported by observation; a comparison of the spectra extracted from the unexposed region during flares with those extracted from

the same region in flare-free periods do not show a significant difference. This method also assumes that the underlying X-ray sources do not vary significantly over the course of an observation. The Poisson variation of the brightest cosmic point-sources will be substantially greater than the flare signal. Thus, there may be a strong residual spectrum of the point-source in the flare spectrum, and that residual may be positive or negative (that is, the point-source may be over- or under-subtracted from the flare spectrum). The total number of counts in the flare spectrum measured for any given observation will be relatively small, so one must sum over many observations, which will minimize

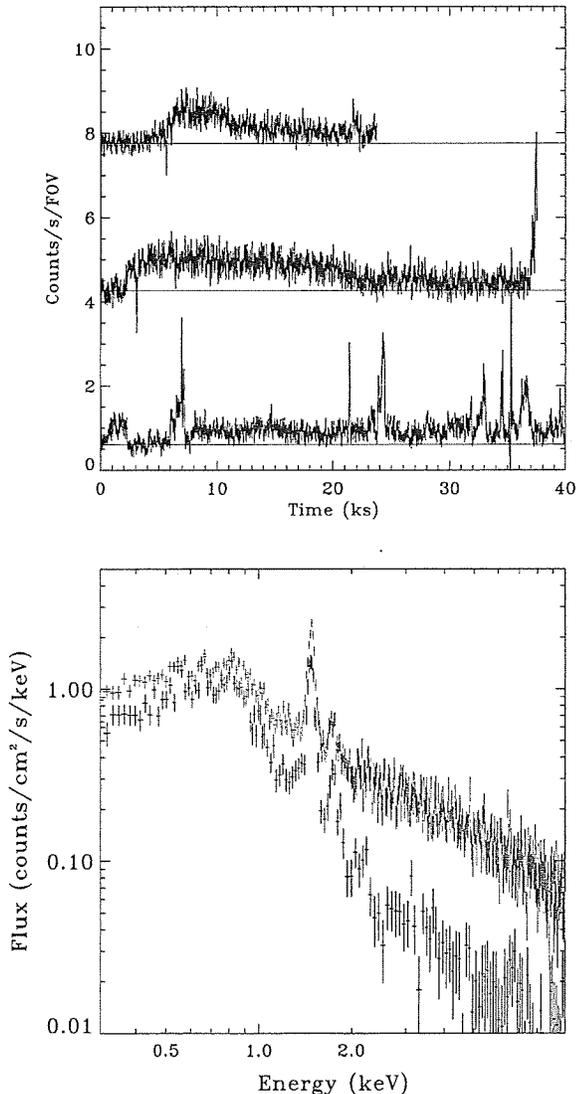


Fig. 13.— **Top:** Three light curves which, were the exposures much shorter, could have led to measurements of a significant overestimate of the QPB rate. **Bottom:** Two spectra of the same region of “blank” sky, both of were cleaned of flare intervals using the method described in §2.

the effects of any single bright source.

The flare level was chosen to maximize counts, minimize contamination by point sources, and provide the best spectral match to the residual flare contamination expected for a typical observation. Although stronger flares produce higher count rates, they also have a lower frequency. The flare spectrum is harder during stronger flares so, to make the studied flare spectrum as similar as possible to the typical residual flare spectrum, one would like to study the weakest identifiable flares. Conversely, the weaker flares are more susceptible to contamination by point sources.

Since the flare spectrum measured for many observations is too small for spectra fitting, we began our analysis by measuring the $(8.0-10.0 \text{ keV})/(0.35-1.35 \text{ keV})$ hardness ratio for observations through the medium filter, the filter for which there is the most data. The distribution of the hardness ratio is roughly Gaussian with a tail to higher values. It is not yet clear if the distribution is due to the inherent statistics of the spectra or whether is a true dispersion in the hardness ratio, but for our purposes the question is moot. We created mean spectra for ten hardness bins and fitted those mean spectra with a variety of functions. The spectra for the Medium filter of MOS1 are shown in Figure 14; they are relatively featureless.

Because the soft protons producing the flares are not X-rays, neither the redistribution matrix file (RMF) nor the ancillary response file (ARF) are appropriate. We initially fit the spectra using neither the RMF nor ARF. However, although *XSPEC* (v11) can fit spectra without applying an ARF, it is incapable of fitting spectra without applying the RMF. Thus we also fitted the spectra using the RMF.

Our goal was to create a model of the flare spectrum that required the fewest free parameters. We anticipated the need to simultaneously fit the object spectrum, the instrumental lines, and the flare spectrum so it was crucial to limit the number of free parameters. The best fit was the function

$$A_0 \exp(-B_0 E) + A_1 \exp(-B_1 E) \quad (2)$$

where

$$A_1 = a_0 + a_1 A_0 + a_1^2 A_0 \quad (3)$$

$$B_1 = b_0 + b_1 B_0 + b_1^2 B_0 \quad (4)$$

which requires only two fit parameters, once the lower case parameters are determined for the data in the public archive. However, this model is not implementable in *XSPEC*.

The disadvantages of fitting the spectra using the RMF are that the model spectra have two features not seen in the measured spectra, a small-scale feature at 1.75 keV, which will be hidden by the instrumental lines, and a decline below 0.45 keV. The latter causes

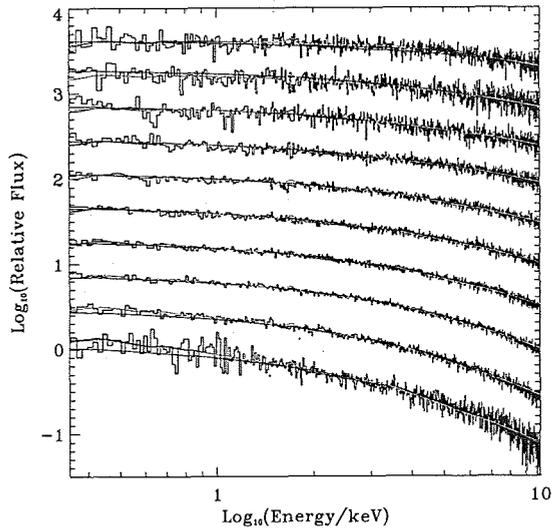


Fig. 14.— The mean flare spectra as a function of the (8.0-10.0 keV)/(0.35-1.35 keV) hardness ratio for MOS1 observing through the Medium filter. **Blue:** The fitting function was the sum of two exponentials where the parameters for the second exponential are quadratic functions of the parameters of first. No RMF or ARF was used in this fit. **Red:** The fitting function was a broken power law where all of the parameters were allowed to vary freely for all of the spectra. Only the RMF was used. **Green:** The fitting function was a broken power law where the break energy was set to a single value for all of the spectra. Only the RMF was used.

severe problems when attempting to fit the soft continuum. We fitted the spectra with a broken power law,

$$A_0 E^{-B_0} \text{ for } E < E_b \quad (5)$$

$$A_0 E_b^{B_1 - B_0} E^{-B_0} \text{ for } E > E_b \quad (6)$$

where B_1 and E_b were functions of B_0 . The only constraint we could place was to fix E_b ; further constraints produced significantly worse fits. For MOS1 $E_b = 3.318$, for MOS2 $E_b = 3.114$; in practice a break energy of 3.2 keV will work adequately for both detectors. The sum of two exponentials, when fitted using the RMF produces strong residuals for $E < 1.0$ keV.

4.2. Variation of the Spectrum with Flare Strength

We have extracted flare spectra for flares with different strengths. We found that the spectrum becomes harder with greater mean flare strength. We had hoped that one might be able to model the flare contamination for a given observation by making a histogram of the flare strengths and, using that histogram to weight model spectra for each flare strength, reconstruct the mean flare spectrum for the observation. This method

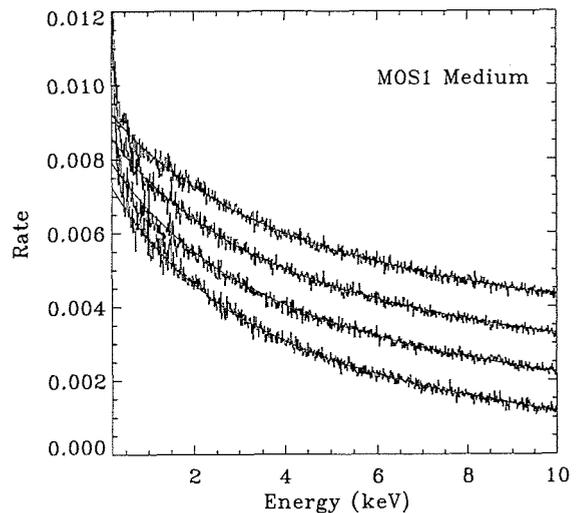


Fig. 15.— The flare spectrum for four different levels of flare strength. From the bottom, 1-2, 2-3, 3-4, and 4-5 counts/second above the quiescent level. As the flare strength increases, the mean spectrum becomes harder.

did not work; individual flares of a particular strength do not have spectra sufficiently similar to the mean.

The change in spectral shape with flare strength should give one pause; the spectrum measured here is for rates of 1-2 counts/second above the quiescent level while typical residual levels are likely to be lower. However, the above prescription is sufficiently general that it can compensate for a reasonable change of spectral shape with flare strength.

4.3. Spatial Variation of the Spectrum

In order to understand the spatial variation of the flare spectrum, we divided each detector into six annuli, as shown in the top panel of Figure 16. The inner region has a radius of 1/45 and each annulus has a width of 2/5. The middle panel shows the spectrum from all the annuli. The inner annuli appear to have flatter spectra at energies < 1 keV, but the spectral shapes at higher energies all seem to be consistent. Most bright point sources are placed in the innermost region for observation. As a result, the spectrum of the innermost region (not shown in the figure) is very strongly contaminated by bright sources. We constructed a spectrum of the innermost region from all those observations for which point sources were not obviously over- or under-subtracted. Although the signal-to-noise ratio was poor, it was consistent with the spectra obtained from the inner two annuli (A and B); thus the spectrum from the inner annuli may be safely used as a model for the spectrum at the very center of the FOV. The outer annuli have lower overall flare rates. The bottom panel shows the spectrum from the outer two annuli for each

chip. Except for MOS1-2, there do not appear to be significant chip-to-chip variations in the shape of the flare spectrum. Chip MOS1-2 has a spectrum that is more steeply rising below 0.7 keV which is due to a “hot” edge that can be seen in the vignetting maps.

Thus, any analysis that compares spectra from different regions of the FOV (such as measuring the temperature profile of a cluster) must take into account the radial dependence of the residual flare contamination. Fitting the same spectral shape at all radii using the ARF will fail because the soft protons do not have the same vignetting as the X-ray photons; thus fitted continuum values for the source will be systematically depressed with greater radii. This problem is addressed more completely in Snowden et al. (2005).

In order to determine a vignetting function for the soft proton flares, we constructed images of the flare counts in the FOV in six energy bands: 0.3-0.75 keV, 0.75-1.25 keV, 1.25-2.0 keV, 2.0-4.0 keV, 4.0-8.0 keV, and 8.0-12.0 keV. For each energy band, we constructed images from the events in flare time intervals and non-flare time intervals. The non-flare images were scaled by the relative exposure times and subtracted from the flare images, a method similar to that used for the spectra. However, before the subtraction, the non-flare image was binned into $25''$ pixels. Binned pixels falling 6σ or more above the median were presumed to contain point sources and those regions were removed from the unbinned image. This method reduces the contamination of the flare images by bright point sources, but does not entirely remove it; the center $1'$ often shows evidence of severe under- or over-subtraction.

The radial vignetting function for the flares is shown in Figure 17 where it is compared to the vignetting function for X-rays propagating through the telescope optics. Both functions were created for/from observations made through the Medium filter. The flare vignetting is significantly flatter than the X-ray vignetting function. The inner $1'$ is clearly effected by over- or under-subtracted point sources. The gap between the inner chip and the outer chips can be seen at $R \sim 5.25'$. For the inner $\sim 5'$ the vignetting functions above 2 keV are very similar to those below 2 keV. At larger radii the vignetting functions at the higher energies are very different from those at the lower energies.

The vignetting functions for Thin1 and Thick filters at a given energy have the same shape as the vignetting function for the Medium filter at the same energy. The amplitude of the vignetting function for the Thin1 filter is similar to that for the Medium filter, but the amplitude for the Thick filter is significantly smaller, indicating that the soft proton flares are attenuated by the additional filter material.

The flare vignetting maps are shown in Figure 18. We have only shown the maps for the MOS1 detector (the MOS2 maps are similar) for the Medium filter.

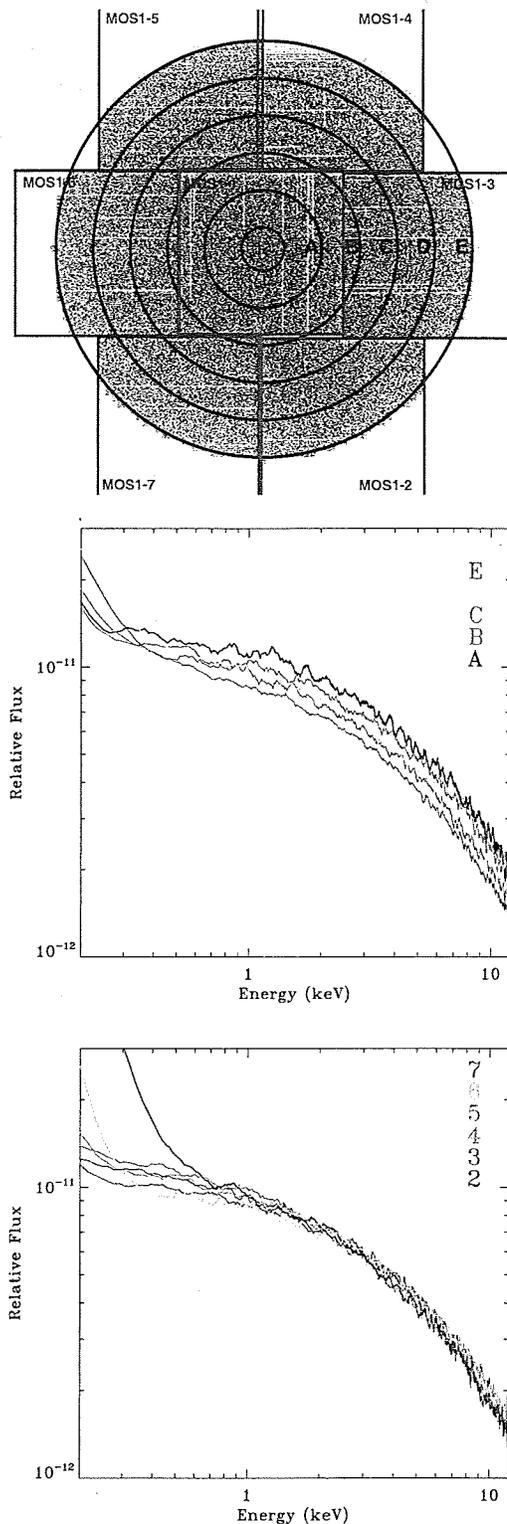


Fig. 16.— **Top:** Soft proton flare image for MOS1 showing the locations of the chips and the annuli used in this analysis. **Middle:** The flare spectrum as a function of radius. **Bottom:** The flare spectrum as a function of chip for the outer two annuli (D and E).

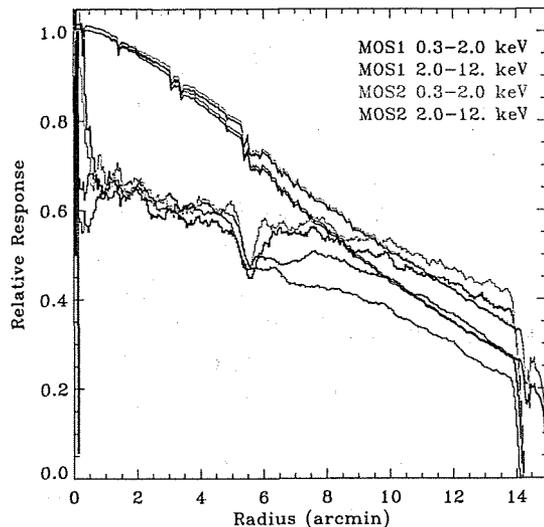


Fig. 17.— The vignetting function of the flares compared to the vignetting function of the X-rays for MOS1 observations made with the Medium filter. The upper set of curves are the vignetting function of the X-rays for both detectors and two energy bands. The X-ray vignetting functions are normalized to unity between $1'$ and $2'$. The lower set of curves are the vignetting function of the flare images. The flare vignetting functions were normalized to unity between $2'$ and $3'$, then shifted downwards for clarity. The flare vignetting functions show evidence of contamination by point sources for $R < 1'$.

The vignetting maps for the other filters have the same shape as that for the Medium filter, but with a different normalization. The vignetting maps reflect the behavior of the radial profiles; at lower energies the vignetting function is fairly flat while at higher energies the vignetting function is more centrally peaked. The vignetting maps show that there is little small-scale structure. The only small-scale structure is on chip MOS1-2, where the upper and left edges of chip MOS1-2 are “bright” at lower energies and dim at higher energies. The lower energy structure does affect the flare spectra shown in Figure 16, but the higher energy structure does not. There are no other features for the MOS1 detector, and there are no similar small-scale features for the MOS2 detector. This bright edge on MOS1-2 is probably the problem observed by Pradas & Kerp (2005).

4.4. Flare Removal Procedure

The method for determining the flare spectrum relied on fitting *mean* spectra. We found that the shape of the spectra was well fit by single parameter family of curves. However, that family of curves is difficult to implement in *XSPEC* and there is no guarantee that

any single flare spectrum will be well fit by that family of curves. Thus, we rely on a somewhat more flexible function, a broken power law where the break energy is fixed to ~ 3.2 keV. This function must be fitted without the ARF.

We recommend constructing one’s fit function as

$$source + e^{-\tau}(Galactic\ FG + Cosmic\ BG) + Al + Si + SPC \quad (7)$$

where SPC is the soft proton contamination, Al and Si are Gaussians to fit the instrumental lines, Galactic FG is the Galactic foreground emission (Local Hot Bubble, halo, etc.), and Cosmic BG is the unresolved AGN. We have been modeling the cosmic background with a power law of $\Gamma = 1.46$ (Chen et al. 1997) and a normalization of $10.5 \text{ keV cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1} \text{ keV}^{-1}$, which is reasonably successful at fitting blank sky spectra at $E > 3$ keV. (See also Luca & Molendi (2004).) After a preliminary fit where the normalization of the broken power law describing the soft proton contamination is set to zero, any soft proton contamination will stand out clearly at $E > 3$ keV. We have found that the normalization and power law indices for the soft proton contamination are similar but not identical for the two detectors. Thus, those parameters can be linked for the initial fits, and allowed to vary for more refined fits.

4.5. Flare Distribution

In order to produce the cleanest particle background, we removed periods of flaring. As a result, we have a record of the time and orbital location of the soft proton flares. Figure 19 shows the fraction of observing time that is not contaminated with flares as a function of orbital position; distance from the earth and orbital phase with respect to the sun. In this coordinate system the elliptical orbit precesses in one year; in June the apogee is near the sun-earth line and the spacecraft passes outside the nominal magnetosheath/solar wind bowshock while in December, the apogee is away from the sun and the spacecraft never passes outside of the nominal magnetosheath. The plot shows that the greatest flare-free time occurs when the spacecraft is furthest from the earth, particularly when the spacecraft is away from the sun. The part of the orbit that seems most susceptible to soft proton flares has radii near to the magnetopause for solar angles $< 90^\circ$. The source of the asymmetry about $Y = 0$ is not understood, but may be related to the way in which the spacecraft, traveling in a single direction, samples an environment that is strongly influenced by magnetic fields. There are also a number of possible asymmetries in the magnetosheath itself which need to be explored further. The second part of Figure 19 is a reprojection of the data showing the fraction of observing time not contaminated with flares as a function of month and distance from the earth.

On average, $\sim 36\%$ of observing time is contami-

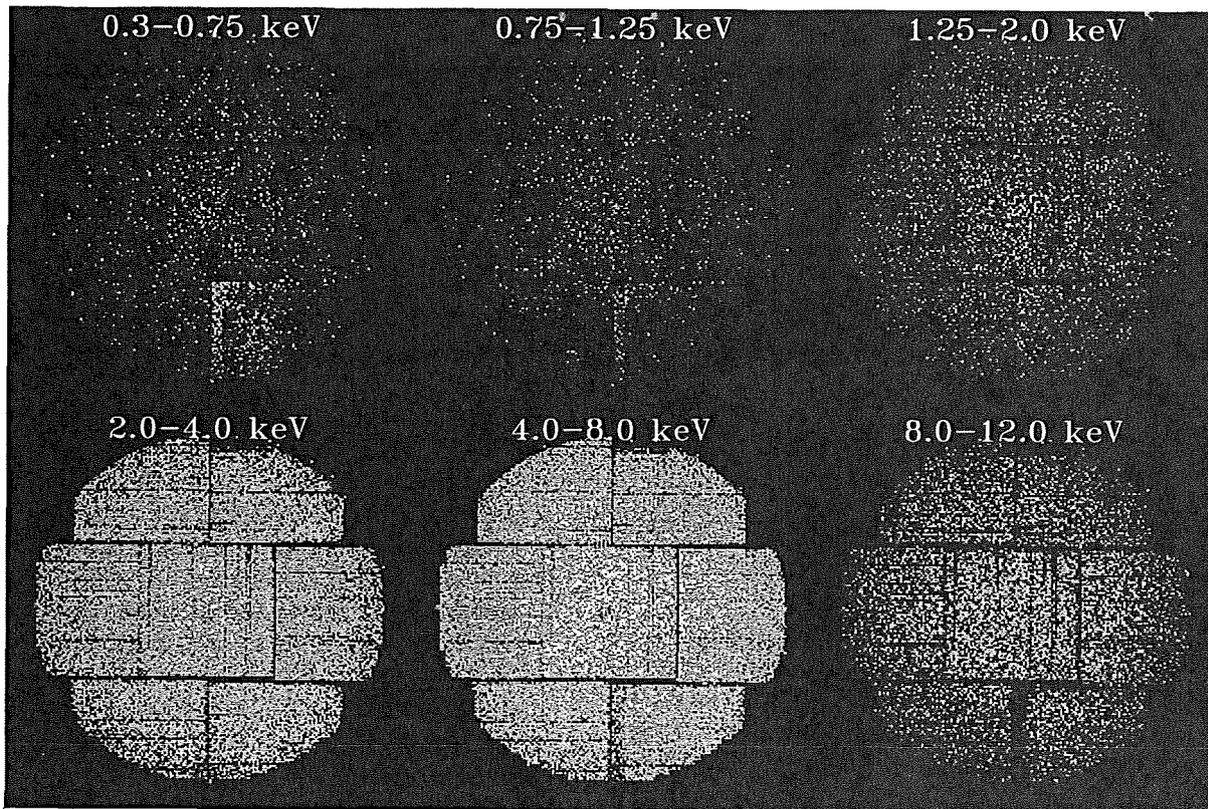


Fig. 18.— The flare vignetting maps for the MOS1 detector observing through the Medium filter.

nated by soft proton flares that can be detected from the light curve using the described method. The fraction of time affected by flares of a given strength is roughly

$$f(R) \sim \exp(-3.9 - 0.66R + 0.045R^2) \quad (8)$$

for a count rate R above the quiescent level (for $1 \text{ cnt s}^{-1} \text{ FOV}^{-1} < R < 5 \text{ cnt s}^{-1} \text{ FOV}^{-1}$).

5. Application

In order to study the repeatability of our particle backgrounds, we extracted from the archive all of the “blank sky” targets for which there were multiple observations. We chose “blank sky” targets such as the Lockman Hole or the AXAF Ultra-deep Field in order to avoid strong variable sources and bright sources with thermal spectra. Our intent was to measure the difference between spectra, which would place an upper limit on the variation introduced by the uncertainty in the particle background spectrum. Since most of the “blank sky” observations are at high Galactic latitudes, this sample is also a good sample of halo lines of sight useful for studying the variation in the temperature and emission measure of the Galactic halo.

One of our first tests was the Hubble Deep Field North observations where we discovered that one of the four spectra was very different from the others. The

difference was due to the first half of the discrepant observation; the latter part of the observation matched the other observations. The difference was due to increased emission in O VII 0.56 keV, O VIII 0.65 and 0.81 keV, C VI 0.37 and 0.46 keV, Ne IX 0.91 keV and Mg XI 1.34 keV. This line emission is expected from solar wind charge exchange (SWCX) emission which is due to the interaction of the solar wind with the earth’s magnetosheath, the region between the magnetopause and the bowshock (Snowden et al. 2004; Collier et al. 2005). The discrepant observation had a peculiar observation geometry; the spacecraft was observing tangentially through the subsolar portion of the earth’s bowshock and magnetosheath. In the near-earth environment, the magnetosheath has the highest density of neutral material, and thus the highest density of material with which to charge-exchange with the solar wind. Within the magnetosheath, the subsolar portion has the highest density of neutral material. The discrepant observation also happened to occur when the solar wind flux was particularly high, so there was a larger than usual number of ions for the charge-exchange reactions. From this observation, it was not clear whether the increase in SWCX emission compared to other observations in the same direction was due to the higher neutral density along the line of sight (due to the observation geometry), or the higher ion density in the solar wind (due to the particular time of the observation). Thus,

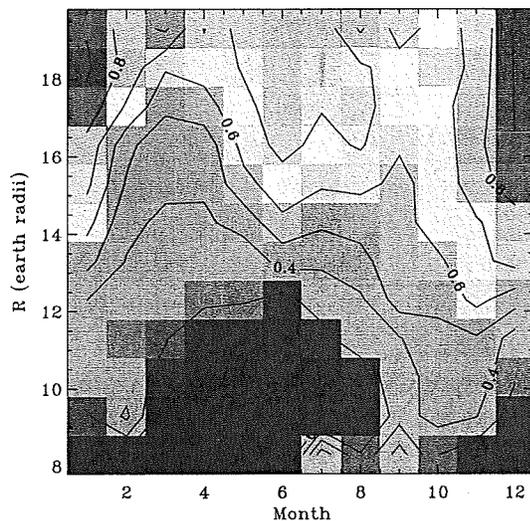
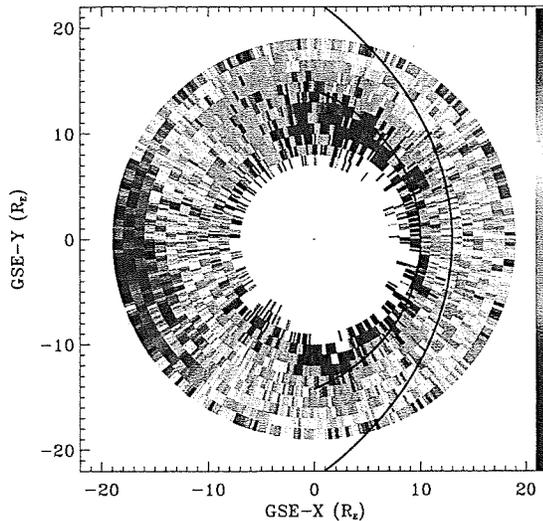


Fig. 19.— **Top:** The fraction of flare-free observation time as a function of distance from the earth (in earth radii) and the orbital phase with respect to the sun; the sun is to the right. Note that because the orbit has a significant inclination, this is not a projection into the XY plane. The color bar indicates the fraction of flare-free observation time where purple indicates severe losses and red indicates minor losses. The black lines are the intersection of the bowshock (exterior) and magnetopause (interior) with the XY plane for a nominal solar wind. We have not extended the model beyond solar angles $> 90^\circ$. **Bottom:** The fraction of flare-free observations time as a function of month and distance from the earth.

the collection of data useful for testing the stability of the particle background model also provides a small (but still useful) statistical sample with which to assess the probability of SWCX contamination in individual observations.

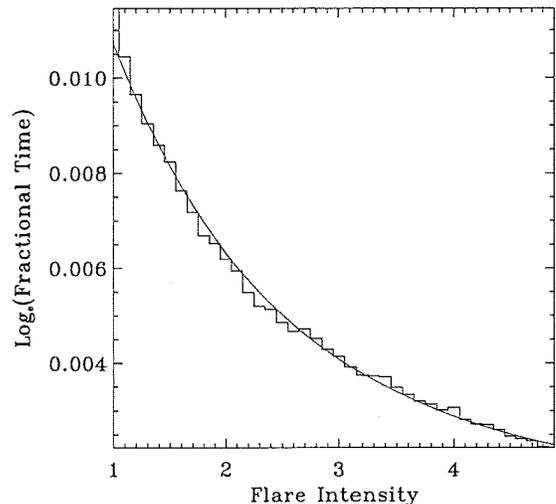


Fig. 20.— The fraction of the time experiencing flares of a given intensity in counts $s^{-1} \text{FOV}^{-1}$ above the quiescent level. The smooth curve is a rough fit.

Finally, it should be noted from §4 that there is an uncertainty in the removal of the soft proton contamination. Since the spectral signature of the soft proton contamination is easily distinguishable, this collection of data also allows one to assess the incidence of soft proton contamination after the removal of the obvious soft proton flares.

5.1. Data

We used a total of 39 obsids of eight “blank sky” targets; for each target all of the obsids have almost exactly the same pointing location. This data collection was supplemented by one “blank sky” target (RFT) with three obsids that do not have the same pointing but are sufficiently close to overlap one another. The data collection also includes one “blank sky” target (SOU) with 25 obsids which are scattered over a $\sim 2^\circ$ radius region; these data are not good for strict SWCX measurements but can be used to understand the Galactic halo and to augment the soft proton contamination statistics. Finally, we have included seven observations of as many different targets of “blank sky” or diffuse emission (such as clusters) to further augment the soft proton contamination statistics, as well as 38 cluster observations. The data used are listed in Table 5. Other obsids for these targets sometimes exist in the archive; they have not been included here due to flare contamination that was so strong that a quiescent level could not be determined.

In order to assess the stability of the particle background modelling we compared all of the spectra for each target with multiple observations. For this comparison we subtracted our model particle background

TABLE 5
DATA

Mnemonic	R.A	Dec.	ℓ	b
Blank Sky Targets with Multiple Observations				
AUD	53.08088	-27.79303	223.539	-54.46646
DFA	203.6549	37.88383	85.47977	75.93033
GWS	214.3077	52.37225	96.32891	60.09706
HDF	189.1987	62.18789	125.8984	54.85594
LHO	163.1901	57.50739	149.2398	53.13183
MAR	48.26479	-55.03955	270.1973	-52.12631
MBR	33.753	-73.98461	295.3369	-41.88856
PFL	165.4843	86.19036	124.9638	30.49656
Secondary Blank Sky Target with Multiple Observations				
RFT	187.805	20.77117	265.8689	82.14424
Blank Sky Targets Useful for Soft Proton Contamination Study				
SOU	35.87749	-4.46126	171.08078	-58.44255
Other Soft Proton Contamination Data				
CFRS	150.1631	25.26689	205.7925	51.8276
HALO	341.1985	-72.71964	314.9162	-41.3322
LHB3	347.413	61.6285	111.1486	1.1076
MB16	49.9786	11.25294	171.0738	-37.3755
ONNF	53.17258	-63.47839	278.6737	-45.3190
OFFF	49.98542	-62.45558	278.7296	-47.0916
SGP1	14.68725	-27.582	228.9665	-88.3144

from each observed spectrum, divided the MOS1 and MOS2 spectra by their respective responses (the probable reason for the upturn above 6 keV seen in Figure 21) and then added MOS1 and MOS2 spectra together. For each target we then plotted the spectra from all of the obsids in the plot for comparison (the solid lines in Figure 21). We chose the spectrum with the smallest value in the 0.5-0.7 keV band as that least likely to be contaminated by SWCX emission (plotted in black). We then subtracted that spectrum from each of the other spectra (the lower set of solid lines in Figure 21) and compared that difference with the expected 1σ uncertainty (the dotted lines Figure 21). We see two types of strong differences; either changes in the strengths of the O VII, O VIII, and other individual lines, or the addition of a smooth powerlaw-like spectrum. The former is due to the SWCX and the latter is due to the soft proton contamination.

We consider each of the targets in turn in the following paragraphs. For each observation we calculated the pathlength through a simple model magnetopause and bowshock of the forms

$$R_{MP} = \frac{14.21}{1 + 0.42 \cos \theta} \quad (9)$$

$$R_{BS} = \frac{22.74}{1 + 0.75 \cos \theta} \quad (10)$$

respectively, where R is in units of earth radii, and θ is the angle from the earth-sun line. These models assume a nominal solar wind pressure of 2.5 nPa; the distances scale as $P^{\frac{1}{2}}$ Petrinec & Russell (1996), so changes of a factor of two in the pressure cause changes of only $\sim 25\%$ in the distances. We also extracted the solar wind flux values from *ACE* and *WIND*; these are shown graphically in Figure 22.

HDF: observations of the Hubble deep field north. This is the set of observations discussed in Snowden et al. (2004) and Collier et al. (2005). The QPB value is well determined for all four observations, and is consistent among the observations. All four observations were made with the line of sight behind, through, or in front of the subsolar portion of the magnetosheath. Of the observations, HDF4 passes through the highest column densities. HDF4 also has the highest solar wind flux value of $1.7 \times 10^9 \text{ cm}^{-2} \text{ s}^{-1}$ (higher than 99.8% of the flux measurements in the *ACE* archive). HDF1 has an elevated solar wind flux of $5.5 \times 10^8 \text{ cm}^{-2} \text{ s}^{-1}$, higher than 91% of *ACE* observations) while HDF2 and HDF3 have nominal flux values. Of these observations, HDF2 had the smallest emission in the O VII and O VIII lines (see Figure 21). HDF3 and HDF1 have somewhat elevated values at O VII and O VIII; the difference between these spectra and that of HDF2 is about three times the uncertainty in that difference. HDF4 has very

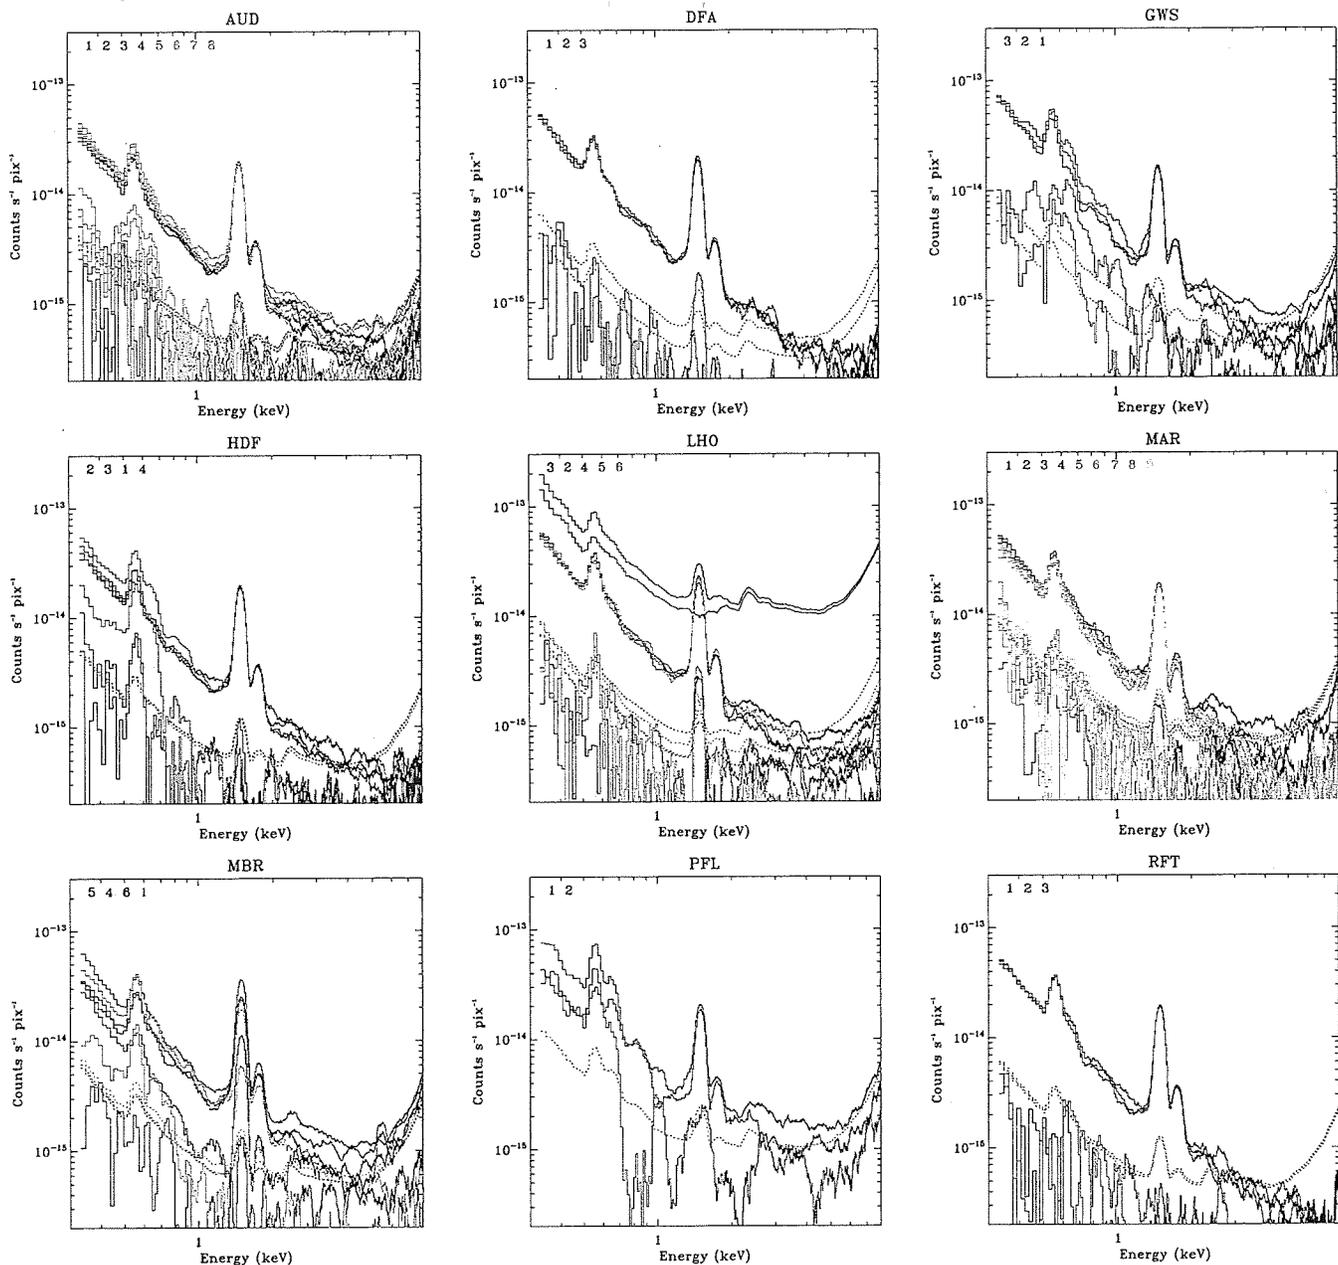


Fig. 21.— Each panel contains the all of the spectra for a single target. Of the upper set of plots (solid lines) the black spectrum is that which has the smallest O VII and O VIII flux. Each spectrum is the sum of the MOS1 and MOS2 spectrum after division by the response (arf). The lower set of solid lines shows the difference between the black spectrum and each of the other spectra. The dotted lines are the propagated uncertainties for the difference spectra.

strongly enhanced O VII and O VIII, as well as a number of other lines identified by Snowden et al. (2004). That the differences in the spectra are correlated with the solar wind strength and that those differences occur at strong lines argue that the differences are not due to the background subtraction.

GWS: observations of the Groth-Westfall strip. In all three observations the QPB level was well determined, though the QPB value for GWS1 is significantly higher than for the other two observations and shows

the greatest contamination. All three of these observations were made when the spacecraft was near the nominal bowshock and observing tangentially through or near the magnetosheath. The solar wind flux for GWS3 was nominal while the solar wind flux for the other two observations was significantly higher. There is some indication that the 0.5-0.75 keV lightcurve of GWS1 is correlated with the solar wind flux. The GWS1 and GWS2 do show enhanced O VII, O VIII, and 0.6-1.0 keV continuum with respect to GWS3. These observa-

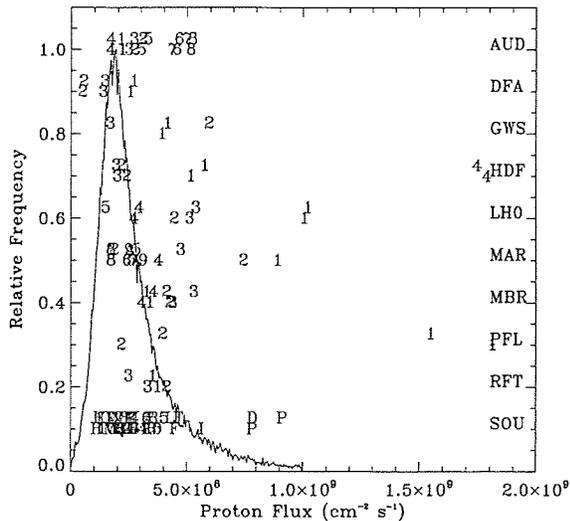


Fig. 22.— **Solid Line:** The relative frequency of the solar wind proton flux derived from the *ACE* archival data. A plot for the *WIND* archival data is very similar. The three letter mnemonics on the righthand side of the plot refer to the various sets of observations. Plotted at the same vertical level as the mnemonic are the values of the solar wind for each observation within that set. The set of values plotted in red slightly above the mnemonic are derived from the *ACE* data, while those plotted in blue slightly below the mnemonic are derived from the *WIND* data; not all observations have both *ACE* and *WIND* data.

tions have a geometry similar to that of the HDF observation discussed by Snowden et al. (2004) so there is little surprise that they are also contaminated by SWCX emission when the solar wind flux is high.

DFA: observations of a “blank” high latitude field. In all cases the QPB level was very well determined and although soft proton flares were frequent, they were well defined. All three of these observations were made when the spacecraft was near the nominal bowshock and observing tangentially through only a small portion of the magnetosheath, at times close to the subsolar point. The solar wind flux was nominal to significantly low, indicating that these observations had longer pathlengths through the magnetosheath than we have calculated. There is no significant variation at the oxygen lines.

RFT: observations across a high Galactic latitude absorption feature. For this set of three observations, the requirement of strict colocation was relaxed. These three observations almost overlap; each observation is offset by a half degree from the previous. The three observations span a feature with a $\Delta N_H = 1.3 \times 10^{20} \text{ cm}^{-2}$. The QPB levels were well determined for each of the observations. Each of the observations was taken

when the spacecraft was near or outside the nominal bowshock and pointed away from the earth, so that the pathlength through the magnetosheath is minimal. The solar wind was somewhat higher than nominal. No significant differences are seen at the O VII and O VIII lines.

PFL: observations of the Polaris flare. The QPB levels were reasonably well determined for both observations. If θ is the angle between the sun and the intersection of the line of sight and the magnetopause as measured from the earth, then both of these observations were taken with $\theta = 85^\circ\text{--}90^\circ$ and the line of sight passing through the flanks of the magnetosheath. The solar wind flux for PFL1 was somewhat higher than nominal ($\sim 4 \times 10^8 \text{ cm}^2 \text{ s}^{-1}$) while for PFL2 the solar wind flux is much higher ($\sim 15 \times 10^8 \text{ cm}^2 \text{ s}^{-1}$). Not surprisingly, PFL2 shows enhanced O VII, O VIII, and continuum below the O VII line. PFL1 shows some soft proton contamination compared to PFL2, though its source is not obvious in the lightcurve.

MAR: observations of the Marano field. There are nine observations of this target for which a QPB level could be determined. Most of the QPB values are around $0.35 \text{ counts s}^{-1}$, but there is one outlier at $0.45 \text{ counts s}^{-1}$ with no apparent reason for its excess. The observation with the higher QPB rate also shows an elevated spectrum at $E > 3 \text{ keV}$, suggesting soft proton contamination. For all of the observations the spacecraft is near or outside the nominal bowshock with little or no part of the line of sight being within the magnetosheath; where it is within the magnetosheath $\theta \sim 60$. The solar wind values are mostly nominal; those that are elevated are, ironically, those observations with the longest pathlength within the magnetosheath. However, the variation of the O VII line is $\lesssim 2\sigma$ and usually $< 1\sigma$; variation in the O VIII line is $< 1\sigma$.

AUD: observations of the AXAF deep field. There are nine observations of this target for which the QPB level was well determined. The QPB values varied from $0.36 \text{ counts s}^{-1}$ to $0.44 \text{ counts s}^{-1}$. The observation geometries are quite varied. For AUD1 and AUD2, the spacecraft was near the bowshock and the line of sight through the magnetosheath was short. For the remainder of the observations the spacecraft was near the magnetopause with a long pathlength within the magnetosheath. The solar wind fluxes for these observations range from nominal ($\sim 2 \times 10^8 \text{ cm}^2 \text{ s}^{-1}$) to elevated ($\sim 5 \times 10^8 \text{ cm}^2 \text{ s}^{-1}$). The strength of the O VII and O VIII lines in the spectra (Figure 21a) are also variable. The AUD1 and AUD2 spectra, which are extracted from the observations with the shortest pathlengths through the magnetosheath and nominal solar wind fluxes, have the smallest O VII and O VIII line fluxes. The AUD6 and AUD7 spectra, which are extracted from the observations with the highest solar wind fluxes (stronger than $\sim 80\%$ of solar wind fluxes

measured by *ACE*), have the highest O VII flux; the difference between these spectra and the AUD1 spectrum is three to four times the expected uncertainty in the difference. The AUD8 spectrum, which was also taken in a period of high solar wind flux, has a typical O VII flux, but appears to have enhanced O VIII; the difference between this spectrum and the AUD1 spectrum is twice the expected uncertainty in the difference. The AUD3, AUD4, and AUD5 spectra, which were taken during nominal solar wind fluxes, and have similarly long pathlengths through the magnetosheath have O VII levels are not significantly elevated; the difference between these spectra and the AUD1 spectrum is on the order of the uncertainty in the difference.

LHO: observations of the Lockman Hole. There are five observations of this target for which a QPB level could be determined. Four observations have well determined QPB levels which are self consistent, the fifth, LHO2, is poorly determined and significantly higher, showing strong soft proton contamination in the spectrum. The observation geometry is similar for all of these observations; the lines of sight pass through relatively long paths through the flanks of the magnetosheath. The solar wind fluxes range from nominal to elevated. However, the differences in the O VII strength are only $\sim 2\sigma$ and the variation at O VIII is less than 1σ .

MBR: observations of the Magellanic Bridge. There are four observations of this target where the QPB levels could be well determined, and two have strongly discrepant values for no apparent reason. The observation geometries are quite varied. For MBR1, the spacecraft is outside the nominal bowshock and is not observing through the magnetosheath, MBR2 and MBR3 have relatively long paths through the flanks ($\theta = 90^\circ - 120^\circ$), while MBR4 has a relatively short path closer to the subsolar point ($\theta = 66^\circ$). All observations have about the same somewhat elevated solar wind flux values. MBR1 and MBR2 have the strongest O VII and O VIII emission while MBR3 and MBR4 have the weakest; there seems to be no correlation of the line fluxes with the geometry or the solar wind.

5.2. SWCX

The DFA and RFT observations show that in the absence of SWCX or soft proton contamination, the differences between spectra are on the order of or smaller than the uncertainty except near the Al and Si instrumental lines. Many of the other sets of observations show that in periods of nominal solar wind fluxes and with similar observation geometries, the differences between spectra are also on the order of the uncertainty. Given that the differences between spectra can *usually* be attributed to changes in the solar wind flux or the observing geometry, and that the differences are either restricted to strong lines expected from the SWCX

or are smooth powerlaw-like additions over the entire spectral range, the differences are not likely to be due to the background modeling method.

The observations which have the strongest SWCX emission are those which pass tangentially through the magnetosheath near ($\sim 30^\circ$) of the subsolar point (HDF, GWS). The HDF observations show that the amount of SWCX emission depends strongly upon the solar wind flux. Although there is an HDF observation at nominal solar wind flux values, there are no HDF observations that do not have this tangential/subsolar geometry, so we are not able to place good limits on the amount of enhancement due to this geometry (compared, for example, to an observation passing through the flanks of the magnetosheath during a period of nominal solar wind flux). However, the DFA observations suggest that when the solar wind is low; even observations near the subsolar point of the magnetosheath may not be strongly contaminated with SWCX emission.

The PFL observations show that strong SWCX contamination can occur even when observing through the flanks of the magnetosheath if the solar wind is particularly strong; in this case the solar wind was stronger than 99.5% of the *ACE/WIND* measurements. The AUD observations suggest that, for nominal solar wind fluxes, lines of sight passing through the flanks of the magnetosheath have very slightly (1σ) elevated O VII values compared to lines of sight that do not pass through the magnetosheath. The LHO, MAR, and AUD observations suggest that when observing through the flanks of the magnetosheath, enhancements in the solar wind to the $\sim 80^{th}$ percentile can cause enhancements of $\sim 1\sigma$ in the strength of the O VII line with only occasional enhancement at the O VIII line.

Of course, the magnetosheath is not the only source of SWCX emission. The solar wind interacts with planetary neutral material along the entire line of sight. Since this emission is intergrated along the entire line of sight, the time variation due to the variation in the strength of the solar wind will be substantially smoothed over time scales longer than *XMM-Newton* observations. Variation from one observation to another could be due to emission from the interplanetary medium. However, there are a number of cases (HDF, GWS) where the SWCX strength can be correlated to the solar wind flux over the course of a single observation, implying that in those cases the SWCX is a local phenomenon.

5.3. Soft Proton Contamination

The rapidly declining response of the instrument to X-rays at $E > 8$ keV, the rapidly declining cosmic emission at the same energies, and the strength of the instrumental response to soft protons that appear as $E > 8$ keV photons is responsible for the success of the soft proton flare contamination criteria developed by

Luca & Molendi (2004). They compared the flux in the FOV with that of the unexposed region for $8 < E < 12$ keV where the instrument has a poor response to X-rays. If the ratio of surface brightnesses was > 1.3 they classified the observation as “strongly contaminated” by soft proton flares.

We have followed a slightly different method more closely tailored to the background removal method described above. We have used the FOV in the $8 < E < 12$ keV band and $11:58 < R < 14:08$ ($13900 < R < 16900$ detector coordinate pixels) after flare removal and after source elimination to compare to the unexposed-region data over the same energy range. Since we have already removed flares using the light-curve, this measure represents the remaining low-amplitude flares. To this value we compared a spectroscopic determination of the strength of the soft proton contamination.

For the data sets described in the following section we fit this model simultaneously with the spectrum of the cosmic background. To each of the spectra we fitted the function

$$N_L A_L(T) + \quad (11)$$

$$e^\tau (N_{HH} A_{HH}(T) + N_{HS} A_{HS}(T) + N_\Gamma \Gamma^E) + \quad (12)$$

$$Al + Si + SPC \quad (13)$$

where the $N_x A_x(T)$ represent thermal components due to the Local Hot Bubble (L), the soft Galactic halo (HS), and the hard Galactic halo (HH). The $N_\Gamma \Gamma^E$ represents the contribution of the unresolved AGN; $\Gamma = 1.46$ and $N_\Gamma = 10.5 \text{ keV cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1} \text{ keV}^{-1}$.

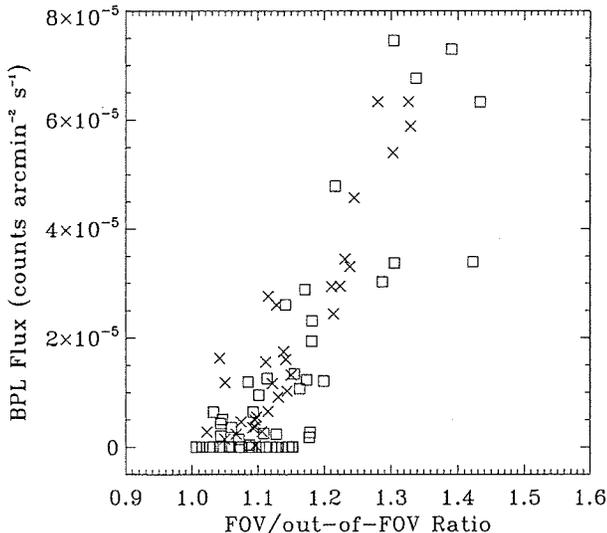


Fig. 24.— The FOV/out-of-FOV ratio in $8 < E < 12$ keV compared to the amount of soft proton contamination determined by fitting the spectrum. The boxes are 78 “blank-sky” images while the crosses are for 38 cluster observations.

The *Al* and *Si* represent the instrumental lines at 1.49 and 1.74 keV, which are modeled as Gaussians. The *SPC* represents the soft proton contamination and is modeled as broken power law with a break energy of 3.3 keV, as described in a previous section. For the blank sky observations the normalization and both of the indices of the broken power law were allowed to vary in this model. For the cluster data the indices were set to be the same. We then integrated the spectrum of the soft proton contamination over $8 < E < 12$. Figure 24 compares the FOV/out-of-FOV flux ratio with the integrated value of the soft proton spectral component. The two values track one another reasonably well. Using our selection criteria, it is clear that if the FOV/out-of-FOV flux ratio is greater than 1.2, there is significant contamination of the spectrum by soft protons. Below that value, the contribution may still be significant (depending upon one’s scientific aims) but is more likely to be negligible.

6. Summary

In order to analyse diffuse emission filling the FOV, it is necessary to model and remove the instrumental backgrounds: the quiescent particle background, and the background due to soft proton flares. Both of these components vary with time and show significant variation across the detector. Our method relies on the contemporaneous, though low signal-to-noise, measurement of the quiescent particle background to resolve the temporal variation, and we use the FWC data to resolve the spatial variation under the assumption that the spatial distribution of the response to the particle background is not strongly varying with time. As has been stressed elsewhere Lumb et al. (2002) using a background extracted from a region on one chip to model the background elsewhere may not be safe due to the strong patterning in the background lines. This is generally only a problem with large, extended sources, for which we developed this method. We do not have sufficient signal-to-noise in either the unexposed region spectra or in the FWC data to apply our method to regions much smaller than about a quarter of a chip, but for such regions the patterning of the background is generally not a significant problem.

The soft proton contamination is, in general, less of a problem. The ratio method developed by Luca & Molendi (2004) and modified here works well to determine whether the residual soft proton emission is significant. For low temperature emission, the spectral signature of the soft proton contamination can be clearly seen at energies higher than those occupied by the thermal emission, so the two can be fit simultaneously with little ambiguity. With higher temperature emission, there is significant overlap with the soft proton emission below the break energy of 3.3 keV, so determining the index of the softest part of the soft proton

emission is more difficult. It should also be noted that since the soft proton contamination has a much flatter radial distribution than the X-ray emission it poses a particular problem in the outer parts of the FOV where, for example, cluster emission is particularly interesting, and particularly weak. The specific application of our background removal method to clusters will be appearing shortly.

Characterization of the quiescent particle background and the soft proton contamination allows study of the next major contamination of the cosmic signal, the emission due to the SWCX. The HDF observations detailed by Snowden et al. (2004) showed that SWCX emission from the magnetosheath could be an important contribution to lines commonly used as plasma diagnostics, such as O VII and O VIII. The strength of the SWCX emission is dependent upon both the flux of the solar wind and the density of the neutrals in the magnetosheath with which they interact. The HDF observation that was most strongly affected by the SWCX emission happened to have a line of sight passing through the densest part of the magnetosheath during a period of strong solar wind, so it was ambiguous which, if either, density dominates. The collection observations detailed above show that significant SWCX contamination can occur when observing through the lower density flanks of the magnetosheath if the solar wind is particularly strong. Conversely, there are a number of observations through the densest part of the magnetosheath that do not show significant SWCX contamination, suggesting that the solar wind flux is a stronger factor than the density of the magnetosheath along the line of sight. However, for none of the observations through the densest part of the magnetosheath are there matching observations through the flanks, so we can not entirely rule out a significant role for the density variation within the magnetosheath when the solar wind flux is low. The combination of Figure 22 and our spectra shows that for observations near the densest part of the magnetosheath, a solar wind proton flux $\gtrsim 4 \times 10^8 \text{ cm}^{-2} \text{ s}^{-1}$ will cause significant SWCX contamination. Such a value for observations through the flanks does not appear to cause significant SWCX contamination, though much higher values ($1.5 \times 10^9 \text{ cm}^{-2} \text{ s}^{-1}$) will show SWCX contamination even observing 90° from the sun. With the current data we can not address the significance of SWCX emission due to the interplanetary medium.

We would like to thank Silvano Molendi, Andy Read, Andrea De Luca, and the XMM Background Working Group for their comments and encouragement. We would like to thank Michael Collier for many helpful discussions. The data for this work were provided through the HEASARC. This work was made possible in part by grants from the *XMM-Newton* Guest Observer Facility at NASA.

REFERENCES

- Chen, L.-W., Fabian, A. C., & Gendreau, K. C. 1997, *MNRAS*, 285, 449
- Collier, M. R., Moore, T. E., Snowden, S. L., & Kuntz, K. D. 2005, *Adv. Space Res.*, 35, 2157
- Luca, A. D., & Molendi, S. 2004, *A&A*, 419, 837
- Lumb, D. H., Warwick, R. S., Page, M., & Luca, A. D. 2002, *A&A*, 389, 93
- Nevalainen, J., Markevitch, M., & Lumb, D. 2005, *ApJ*, 629, 172
- Petrinec, S. M., & Russell, C. T. 1996, *J. Geophys. Res.*, 101, 137
- Pradas, J., & Kerp, J. 2005, *A&A*, 443, 721
- Snowden, S., Kuntz, K., & Mushotzky, R. 2005, in preparation
- Snowden, S. L., Collier, M., & Kuntz, K. D. 2004, *ApJ*, 610, 1182

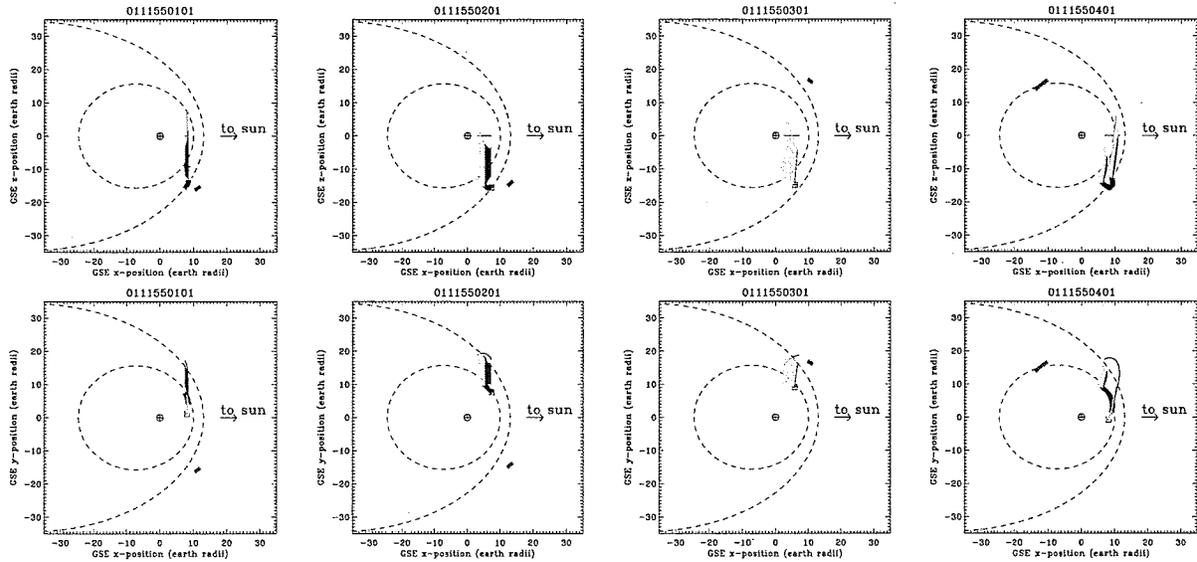


Fig. 23.— The observing geometry of the four different HDF observations. The top panel of each pair is plotted in GSE-X and GSE-Z, while the bottom panel of each pair is plotted in GSE-X and GSE-Y. In this coordinate system the earth is at the origin and the sun lies on the +X axis. The ellipsoids are the model magnetopause and bowshock described in Equations 9 and 10; these models are clearly not valid for angles $> 90^\circ$ from the sun and are only approximations for angles $< 90^\circ$ from the sun. The position of the spacecraft is plotted with a series of crosses. The intersection of the line of sight with the magnetosheath (from magnetopause to bowshock) is shown by the straight lines. Both the position of the spacecraft and the lines of sight are color coded; blue when the spacecraft is below the magnetopause, yellow when the spacecraft is above the magnetopause but below the bowshock, purple when the spacecraft is outside the bowshock but the line of sight intersects the the magnetosheath, and black when the spacecraft is outside the bowshock and the line of sight does not intersect the magnetosheath.