Jan. 16, 2009–corrected right axis label in Fig. 24

A New Gain Map and Pulse-Height Filter for the Chandra LETG/HRC-S Spectrometer

B. J. Wargelin, P. W. Ratzlaff, and M. Juda

Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138

ABSTRACT

We describe the development of a time-dependent gain map for the Chandra X-ray Observatory's High Resolution Camera for Spectroscopy (HRC-S), a microchannel plate detector that is used with the Low Energy Transmission Grating (LETG). We also describe a wavelength dependent pulse-height filter for LETG/HRC-S data that reduces the detector background by 50–70% (relative to standard Level 2 processing) at wavelengths longer than ∼25 Å, with a loss of ∼1.25% of valid X-ray events.

Subject headings:

1. INTRODUCTION

The High Resolution Camera onboard the *Chandra* X-ray Observatory is a microchannel plate instrument comprising two detectors, one optimized for imaging (the HRC-I) and the other for spectroscopy (the HRC-S). Both detectors use microchannel plates in a chevron configuration with a CsI photocathode deposited on the top plate. Instrument operation and capabilities are described in detail in the *Chandra Proposers' Observatory Guide*[1] and in Kraft, et al. (2000).

The HRC-S is the primary detector for the Low Energy Transmission Grating (LETG; Brinkman, et al. 2000), which can also be used with the Advanced Camera for Imaging and Spectroscopy (ACIS) spectroscopy array (ACIS-S). The ACIS-S is a CCD detector with an intrinsic energy resolution of roughly 100 eV at LETG energies, and when used with the LETG covers wavelengths from 1.2 to more than 90 Å, but its very low quantum efficiency (QE) at low energies limits the effective LETG/ACIS-S coverage to ∼1.2–60 Å (200–10000 eV). The HRC-S is physically twice as long as ACIS-S and therefore covers wavelengths out to ∼180 Å (∼70 eV), but has almost no energy resolution and a much higher background than ACIS-S. The full width at half maximum (FWHM) of the HRC-S pulse-height distribution (PHD) for monochromatic X-rays at a given location is typically 40% of the PHD mean, and the mean changes by a little less than the FWHM for a factor

---

[1] http://cxc.harvard.edu/proposer/POG/

of ten change in photon energy. Nevertheless, this slight energy resolution is sufficient to provide significant discrimination between X-ray and background events.

In §2 we briefly describe the HRC-S and explain how it is used with the LETG to obtain high-resolution X-ray spectra. In §3 we describe the creation of a baseline gain map for the HRC-S based on pre-flight laboratory data. Temporal variations in the gain are calibrated by analysis of flight data as explained in §4, and §5 describes the characterization of PHDs and the development of a filter for use with LETG/HRC-S spectral data. We conclude in §6 with example applications of the filter and a discussion of its effectiveness.

2. THE HRC-S AND LETG

The HRC and LETG and their combined operation as the LETG Spectrometer are described in detail in Brinkman, et al. (2000). The HRC-S comprises three segments, also referred to as chips or plates. These flat segments lie nearly end to end with the outer segments tilted at $\sim 1.3^\circ$ to approximately follow the Rowland circle for dispersed LETG spectra. Event positions are determined with a resolution of $\sim 18\mu$m by collecting the charge emitted from the microchannel plates with a cross-grid charge detector (CGCD) and calculating the centroid of the charge cloud. The centroid calculation uses signals from the six "taps" nearest the event, three per axis, where a tap is an output from a group of wires along an axis of the CGCD. Event positions along an axis are determined using the number of the central tap of the triplet (coarse position) and the centroid of the tap signal values to obtain a fine position relative to the coarse position.

Each HRC segment spans 16 taps in the cross-dispersion (short) direction and 64 taps along the long axis, although the "active region" of the detector only covers approximately 12×59 taps on each segment. Because the dispersed spectrum is quite narrow, and to reduce the total telemetry rate (which is dominated by background events), only a 6-tap-wide region is normally used for LETG/HRC observations.

Taps are identified according to their "coarse position" numbers, with coarse U (CRSU) running from 0–15 across the short axis and CRSV from 0–191 on the long axis. Each tap is subdivided into 256 resolution elements, with square pixels 6.43 μm on each side. Raw position coordinates (RAWX, RAWY) are determined using the 3-tap algorithm for each axis as mentioned above. Because of the extent of the typical charge cloud, not all of an event's charge is collected by the 6 position-determination taps, resulting in a bias in the position toward the center tap; the charge deficit is largest when an event occurs at the mid-point between two taps. As a result, images of events plotted in RAW coordinates leave gaps between each tap. A degapping procedure applies corrections to the RAW coordinates to account for the bias. These corrected positions are then projected into several detector-, observatory- and sky-centric coordinate systems.

Figure 1 displays the various coordinate systems and detector regions. In order to avoid having the gaps between segments fall at the same wavelengths for both the positive and negative orders

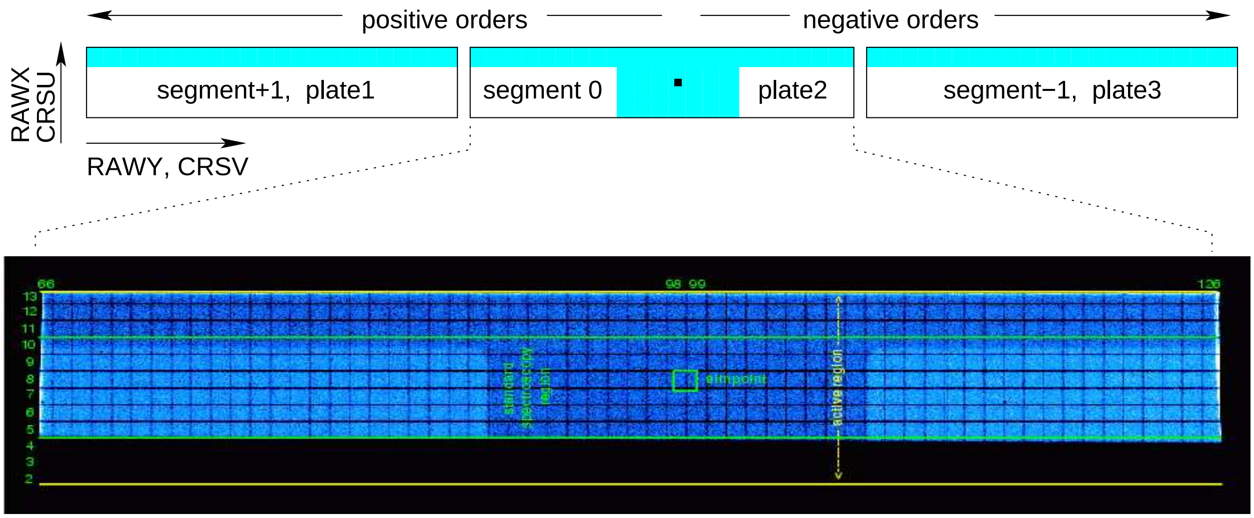

Fig. 1.— Top: Schematic of HRC-S layout, showing all three detector segments. Darker areas are the "T region" of the uv/ion shield; the small black square marks the optical axis (where 0th order falls for on-axis observations). Bottom: Intensity plot in RAW coordinates for the C-K flat field data on the HRC-S center segment. Numbers along the left and top edges denote tap numbers (CRSU, CRSV). Yellow lines mark the boundaries of the detector "active region" and green lines mark the standard LETG/HRC spectroscopy region (CRSU=5–10). The darker T region has fewer counts because of extra absorption by thicker Al on the uv/ion shield.

of the LETG spectrum, the center of the HRC-S is displaced from the telescope optical axis by a few taps along the dispersion axis. For on-axis observations, 0th order falls around CRSV=99. Along the short axis, the aimpoint is roughly centered on CRSU=8, near the middle of the standard LETG spectroscopy window (CRSU = 5–10). Over the course of the *Chandra* mission, which began in 1999, the aimpoint has drifted by a few tens of arc seconds (roughly 1/3 tap) toward smaller RAWX and larger RAWY values. The spacecraft also dithers around the aimpoint to avoid burn-in of the detector, with a peak-to-peak amplitude of 40" (304 pixels) for each axis. The effects of dither are removed via aspect reconstruction during standard data processing, with sky coordinates computed for every event.

The HRC has a 3-piece uv/ion shield, one for each segment, made of thin polyimide films ($\sim$2100 Å for the + and − segments and 2700 Å for the center segment) with Al coatings. Most of the shield has 300 Å of Al, but a wide "T" pattern has 800 Å of Al on the center segment and 2000 Å on the outer segments (see Fig. 1). The stem of the T extends from approximately CRSV=89 to 107 ($|\lambda| \lesssim 18$ Å) on the center segment, and the top of the T begins in the middle of the CRSU=10 tap and extends across all three segments. The thicker Al in the stem of the T mostly serves to reduce optical contamination in 0th order. The top of the T forms the LETG Low Energy Suppression Filter (LESF) but is almost never used and mostly lies outside the 6-tap spectroscopy region.

3. BASELINE GAIN MAP

The purpose of a gain map is to enable normalization of an event's signal so that a photon of a given energy will, on average, yield the same event amplitude everywhere on the detector. Unlike the HRC-I, which has a fairly uniform gain that peaks in the center, the HRC-S gain varies by more than a factor of 2.5 and has a complicated spatial behavior. Below, we describe our analysis of pre-flight data in order to create a baseline gain map. From this baseline we then derive time dependent gain maps to be used with flight data.

3.1. Laboratory Data

In order to calibrate quantum efficiency uniformity and gain on the HRC-S, flat field measurements were conducted prior to launch (Kraft, et al. 1997) using X-rays from an electron impact source at 8 energies from 183 eV (B-K) to 6.4 keV (Fe-K; see Table 1). The Ni-L, Ag-L, Ti-K, and Fe-K measurements were only made on the center segment, and the latter three are at X-ray wavelengths of little interest for LETG observations, so our analysis uses only the B-K, C-K, O-K, and Al-K lab data, which cover the entire detector. Source spectra were filtered with typically 5 or 6 mean free paths of absorption to reduce continuum contamination and the contribution from secondary lines. Although the intensity-weighted average X-ray energies are not exactly equal to the primary line energies listed in Table 1, the differences are negligible given the HRC's minimal energy resolution.

Only one HRC-S segment could be exposed to X-rays at a time, so during every exposure the other two segments collected background data. One ks of background data for the entire detector were also collected with each of the eight sets of X-ray data. Background data collected during the Ni-L, Ag-L, Ti-K, and Fe-K measurements showed anomalous emission near one edge of the − segment (along CRSV=128) and were not used in our analysis, yielding a net background exposure of 66 ks for each segment.

Lab data were processed to match the treatment of flight data as closely as possible. The data were collected in a special laboratory-only format and then converted to the equivalent of standard *Chandra* Level 0 event data files. The CIAO (Chandra Interactive Analysis of Observations; (Fruscione, et al. 2006)) tool `hrc_process_events` was used to convert the level 0 event files to Level 1 event files applying the relevant calibrations. All Level 1 data columns were carried over to Level 2 when generating L2 columns such as `tg_lam` (**t**ransmission **g**rating **lam**bda = wavelength), and a new data column called SAMP was created; SAMP is explained in §3.2. Filtering equivalent to that used in standard Level 2 processing was applied to generate the final event lists, except that events with PHA=255 were retained.

Table 1. Laboratory Flat Field Measurements

Emission	Energy[a] (eV)	Wavelength[a] (Å)	Exposure per segment (ks)
B-K	183	68	18
C-K	277	45	2
O-K	525	24	3
Ni-L	852	14.6	2
Al-K	1487	8.3	2
Ag-L	2984	4.15	2
Ti-K	4511	2.75	2
Fe-K	6404	1.94	2
Backgd.	···	···	66[b]

[a] Energy and wavelength for primary line; emission includes some continuum and often several lines.

[b] Total for combined measurements.

A very small fraction of the HRC-S active area showed enhanced background emission during the lab measurements. Seventeen of these "hot spots," usually caused by dust particles or by electric field effects along plate edges, were identified and removed from all the data files. Most hot spots contain at most a few hundred pixels (out of $\sim 1.4 \times 10^8$), but there are two areas on the + segment around CRSV=17, CRSU=11 that are more than 100 pixels across in both directions. These areas required some data interpolation, as described in §3.3. Note that hot spots seen in lab data are usually different from those seen in flight data, except along some plate edges.

Our analysis bins data on scales of 1/3 of a tap, which typically provides ∼500 counts per subtap. Bin-to-bin gain variations of a few percent are typical on this scale, often with the central subtap having the highest gain within a tap. Coarser binning would not adequately capture this spatial variation, while smaller bins would not provide the statistical accuracy we desire.

3.2. Gain Metrics

To determine the relative gain within a subtap we compute the mean of the event signal strength. Medians were also studied but provided no advantage. The Pulse Height Amplitude (PHA) value is the standard measure of HRC event signal strength, but our work uses a slightly different measure based on the scaled SUMAMPS amplitude, or SAMP. PHA is measured from the sum of all the detector amplifier signals, while SUMAMPS is the sum of only the six amplifier signals nearest the detected event, with three amplifiers for each axis (AU1–AU3 and AV1–AV3). SAMP is defined as SUMAMPS $\times\ 2^{\mathrm{AMP\_SF}-1}/128$ where AMP_SF is a scaling factor equal to 1, 2, or 3, and the division by 128 is somewhat arbitary but convenient. The value of AMP_SF varies from event to event and captures the selected setting of the autonomously switched amplifier range used to extend the dynamic range used to read the tap signals.

SAMP pulse height distributions for the various lab-calibration photon energies are shown in Fig. 2. PHA distributions are very similar. The main difference is that the mean values in SAMP PHDs show much less spatial variation on small scales than PHA means, particularly near tap centers, where PHA means can vary by up to ∼30% from one subtap to the next (see Fig. 3). Fig. 4 compares PHA and SAMP in more detail along CRSU=8, i.e., along the center of dispersed LETG spectra.

Because of the sensitivity of the mean to outliers, particularly at high channels where background events are most common, we use truncated means, ignoring the lowest 5% and highest 5% events. For the same reason, we estimate the uncertainty in SAMP means as the interquartile range (IQR; difference in SAMP between the 25th and 75th count percentiles) divided by the square root of the number of counts per subtap. With 1/3-tap binning, relative gains are measured with a typical accuracy of ∼1.5%. For the rest of this paper, we use S_i to represent the truncated mean of the SAMP PHD in subtap i.

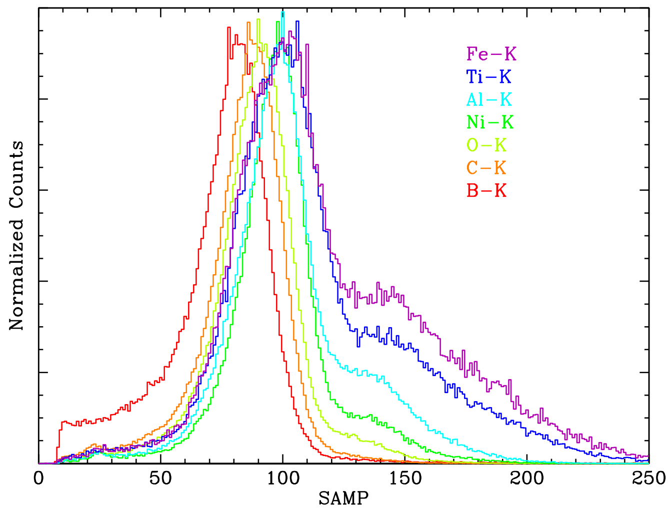

Fig. 2.— SAMP distributions at various photon energies. The lab data shown are taken from detector regions having similar gain.

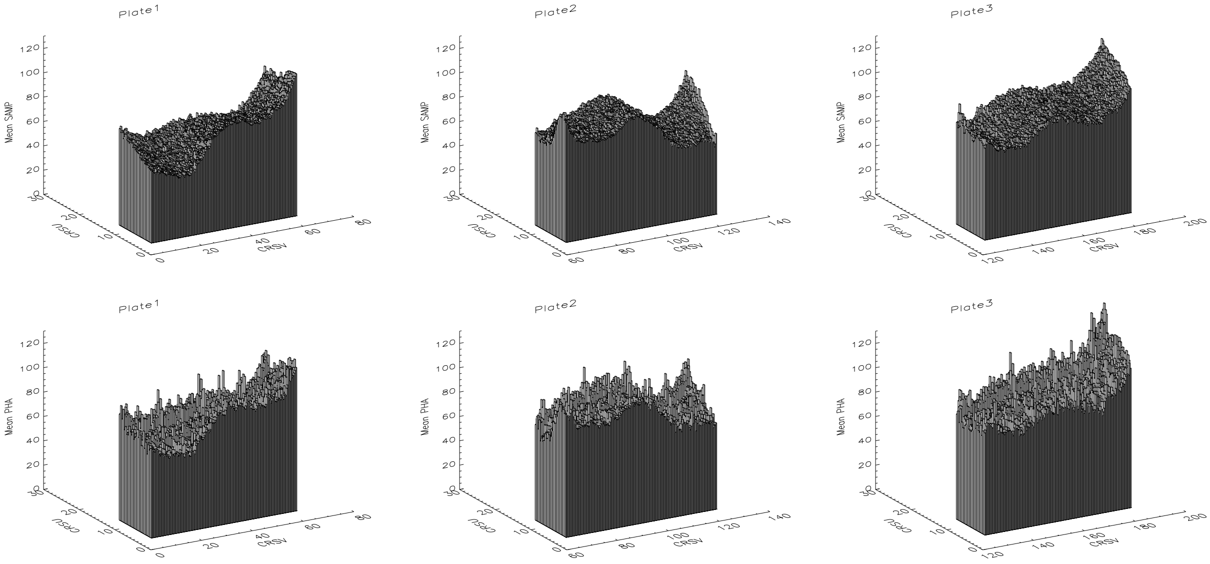

Fig. 3.— Surface plots of S_i^{C}, the SAMP mean value for C-K lab data with a 1/3-tap grid. Bottom panels plot the corresponding PHA means. Statistical uncertainties are typically 1 channel. SAMP means vary more smoothly than PHA means and provide a better metric for detector gain.

3.3. Low-Count Subtaps and Background Subtraction

Not all of the HRC-S active region could be illuminated during the lab measurements. Baffles were used to avoid reflections off the High Energy Suppresion Filter (HESF; never used in flight because of unexpectedly high background), but this also blocked the bottom edge of the HRC-S.

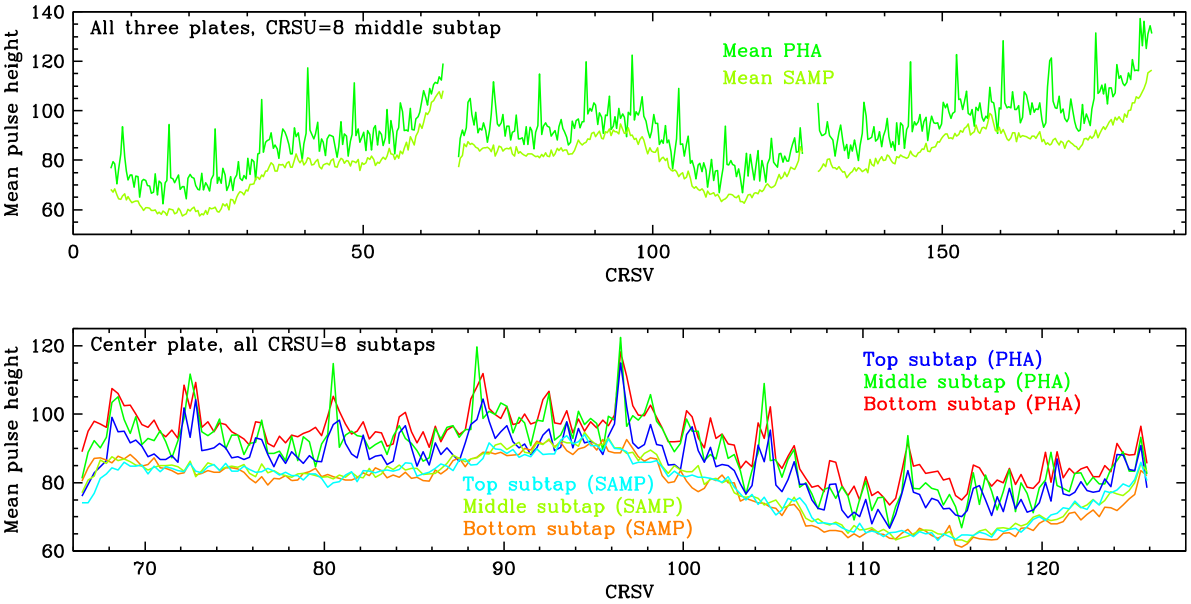

Fig. 4.— Comparison of mean PHA and SAMP values from C-K lab data along a single tap of the HRC-S. Top panel is for the middle subtap of CRSU=8 across all 3 detector segments. PHA means exhibit much more variation, particularly in subtap centers. Bottom panel focuses on the center segment, showing traces for all 3 subtaps of CRSU=8.

Taps with CRSU=2 or 3 are not exposed at all and coverage of taps with CRSU=4 or 5 is partial (see Fig. 1). A few subtaps have few or no counts because of excised hotspots. In unexposed regions we calibrate the gain based on background data (see §3.5), and in low-count subtaps we use larger data-sampling regions. When there are 100–150 counts in a subtap we use a 3×1 sampling region (along constant CRSU), and a 3×3 region when there are fewer than 100 counts. Uncertainty in S_i is computed as described above, but multiplied by $\sqrt{3}$ for 3×1 sampling and by 3 for 3×3 sampling for statistical compatibility with single-subtap sampling.

The B-K dataset requires special handling because of the low intensity of the X-ray source when using a B anode and the low detection efficiency of the HRC for B-K X-rays; note the much longer exposure times for the B measurements in Table 1. In the T region on the outer segments, each subtap has only ∼30 counts, with ∼25% of those counts from background. The T region on the center segment has about 120 counts per subtap, with 6% background. The rest of the detector (thin filter regions) typically has 300 counts per subtap, with 2% background. In comparison, the background fraction in data collected at higher energies is never more than 0.5%. For each sampling region, we therefore subtract the corresponding exposure-weighted background PHD and then calculate the truncated mean SAMP and IQR of the net PHD, grouping SAMP channels when necessary to avoid bins with negative counts. This results in a typical reduction of 1% in S_i for the center-segment T region and ∼5% for the outer-segment T regions, relative to uncorrected results. All further mentions of B-K lab data in this paper refer to background-subtracted results.

3.4. Gain Map Parameterization

To first order, the HRC-S gain map can be created using a normalization constant for each subtap. However, as shown in Fig. 5, which plots the ratio of S_i for Al-K and C-K on a tap-by-tap basis, the detector gain 'steepness' can vary with position, thus requiring a second parameter to characterize the local gain. Figure 6 shows SAMP means from a representative sample of taps over a range of gains and with varying degrees of steepness. As can be seen, there is a nearly linear relationship between S_i and the log of photon energy.

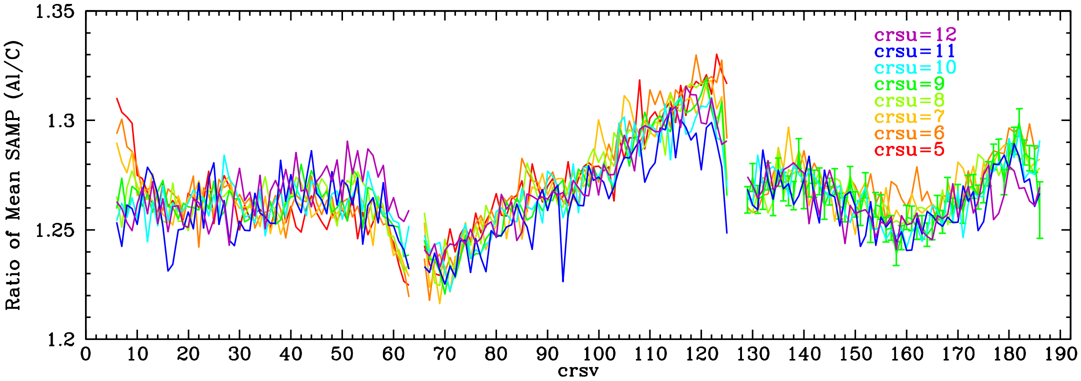

Fig. 5.— Ratio of SAMP means from Al-K and C-K lab data. Traces would be flat and coincident if the HRC-S gain could be characterized by a single normalization factor.

For the first order normalization, we divide S_i by $\mathcal{S}_i^{\rm norm}$, where $\mathcal{S}_i^{\rm norm}$ is the average of the means for B, C, O, and Al within each subtap, i.e., $\mathcal{S}_i^{\rm norm} = (S_i^{\rm B} + S_i^{\rm C} + S_i^{\rm O} + S_i^{\rm Al})/4$. After increasing the normalized means for B and O by 2.68% and 1.39%, respectively[2], we do linear fits to the data for each subtap i (see right panel of Fig. 6). The fit model has the form

$$\frac{S_i(E)}{\mathcal{S}_i^{\rm norm}} = m_i \log E + b_i \tag{1}$$

where $S_i(E)$ is the SAMP mean at E eV, m_i is the slope, and b_i is the y-intercept. Note that Eq. 1 can be rewritten as $S_i = M_i \log E + B_i$, i.e., the gain can be described using only two independent parameters; we explicitly list the unnecessary third parameter, $\mathcal{S}_i^{\rm norm}$, only for transparency.

Eq. 1 is then compared with a reference gain function to determine a pulse-height adjustment factor for any event in a given subtap, yielding a normalized SAMP value that we define as the SAMP Pulse-Height Invariant (SPI) value. Mathematically, this means taking an event's SAMP value, solving for $\log E$ in Eq. 1 (using SAMP in place of S_i), and then using that value in the

---

[2] To simplify the linear fitting process, adjustment factors were chosen to make the average normalized SAMP means for B and O (taken from CRSU subtaps illuminated by LETG spectra on orbit, i.e., CRSU=7b–8c) fall along the same line as the averages for C and Al.

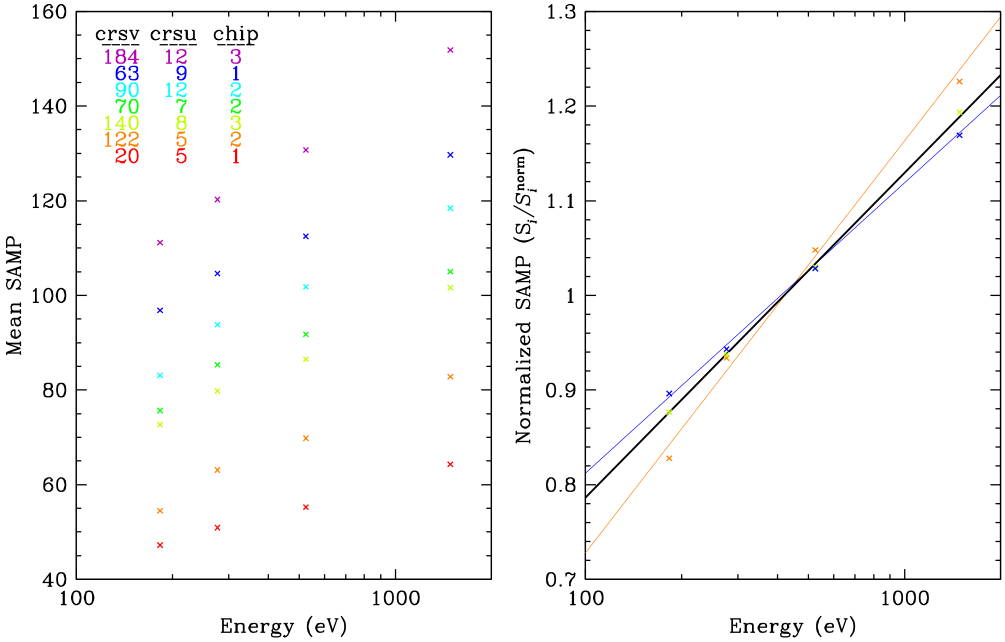

Fig. 6.— Left: Mean SAMP for 4 energies at 7 detector locations. Right: Means are normalized to $S_i^{\rm norm}$, the local average of S_i for B, C, O, and Al X-rays. B and O data points have been scaled up by factors of 1.027 and 1.014, respectively, to better illustrate the near-linear relationship of mean vs. $\log E$. Straight lines are drawn to guide the eye and emphasize differences in gain steepness at different detector locations. The thicker black line traces the reference gain function.

reference gain function to compute the normalized SAMP value, or

$$SPI = m_{ref} \log E + b_{ref} = \frac{m_{ref}}{m_i} \left(\frac{SAMP}{\mathcal{S}_i^{\rm norm}} - b_i \right) + b_{ref}. \tag{2}$$

The reference gain function follows the thick black line in the right panel of Fig. 6, which goes through the data points representing the average of S_i values for CRSV=7b–8c (see footnote 2). The specific values for the reference gain function parameters are $m_{ref} = 0.3431$ and $b_{ref} = 0.1003$. Lastly, we choose a global scaling factor of 113.32 such that $SPI = 128$ when $E = 1000$ eV, yielding

$$SPI = \frac{38.88}{m_i} \left(\frac{SAMP}{\mathcal{S}_i^{\rm norm}} - b_i \right) + 11.37. \tag{3}$$

Note that unlike SAMP, SPI values can be negative, although the fraction of such events is usually extremely small.[3]

---

[3] When the new gain map is incorporated into CIAO/standard processing, negative values of SPI will be set to zero. SAMP will become the new PHA, and SPI will become the new PI. Both will have integer values ranging from 0 to 511, instead of 0–255 for the old PHA.

3.5. Calibration of Unilluminated Region

As noted earlier, part of the HRC-S active region (roughly CRSU=2–4) could not be illuminated with X-rays during the lab calibration. To extend the gain map into unexposed areas we compare mean SAMP results from X-ray and background data, using 3 × 3-subtap sampling to obtain adequate statistics from the background data, which has only ∼7 counts per subtap. First, $\mathcal{S}_i^{\rm norm}$ is computed for each subtap, except that the sampling is over 3 × 3 subtaps for consistency with the background data analysis. Those results are then binned by 1 channel in SAMP, and the mean SAMP for background data from the corresponding subtaps ($S_i^{\rm BG}$) is computed.

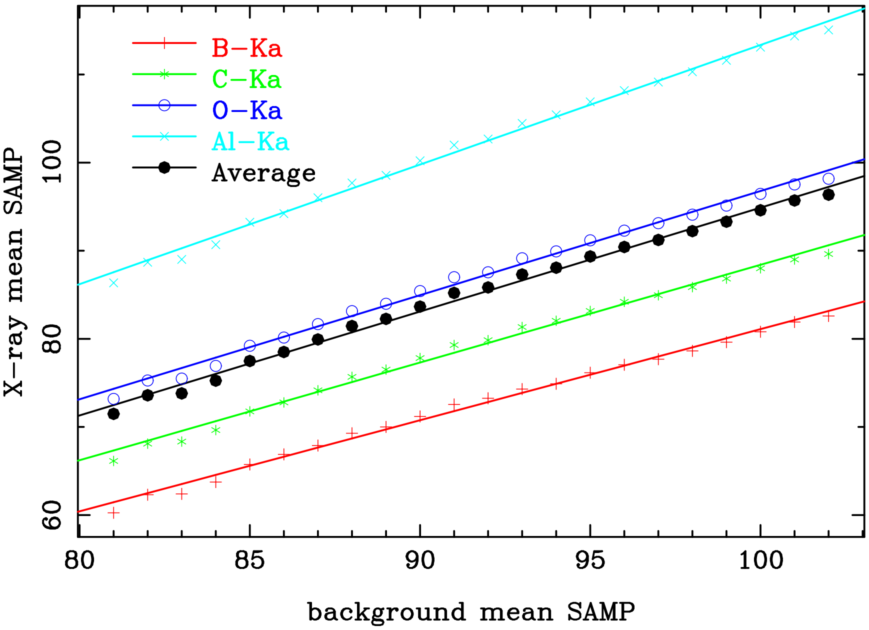

Fig. 7.— Mean SAMP from illuminated subtaps (lab X-ray measurements) versus mean SAMP of background data from corresponding subtaps. The fit to the averaged data is used to determine the value of $\mathcal{S}_i^{\rm norm}$ from background data in unilluminated regions.

As seen in Fig. 7, there is a tight relationship between $\mathcal{S}_i^{\rm norm}$ and $S_i^{\rm BG}$. Based on the results from a least-squares fit to those data, we use the relationship

$$S_i^{\rm norm} = 1.179 S_i^{\rm BG} - 23.029 \tag{4}$$

to determine values of $\mathcal{S}_i^{\rm norm}$ for unilluminated subtaps. Values of m_i are estimated by extrapolating linear fits of m_i versus CRSU subtap from the illuminated subtaps (note in Fig. 5 how m_i varies depending on CRSU, especially near some plate ends), and b_i is then calculated for each subtap by assuming that Eq. 1 goes through the point $E = 500$ eV, $S_i/\mathcal{S}_i^{\rm norm} = 1.028$ (see Fig. 6). The baseline gain map for the HRC-S comprises these three parameters ($\mathcal{S}_i^{\rm norm}$, m_i, b_i), one set for each subtap, to be used with Eq. 3.

4. TEMPORAL CALIBRATION

The baseline gain map described above is derived from laboratory data collected before launch. The gain measured when *Chandra* was in orbit, however, started out somewhat higher and has since gradually decreased over time. The next step in our analysis is to calibrate that decrease, which varies with position on the detector. Unlike with the lab calibration, only a small portion of the HRC-S can be practically calibrated on-orbit, specifically the part illuminated by LETG spectra. The dispersed spectrum is only ∼10 pixels wide (widening to ∼60 pixels on the outer plates) in sky coordinates (i.e., after aspect reconstruction) but over 300 pixels wide in detector coordinates because of spacecraft dither, which has an amplitude of 40" (304 pixels, 1.19 taps) in each direction. Only about 5 CSRU subtaps are therefore illuminated during on-orbit observations.

4.1. Analysis of Flight Data

We use two astrophysical sources for the temporal calibration, both of them bright continuum sources, so that a continuous range of wavelengths and taps can be calibrated with little background contamination. For the outer segments we use HZ43, a white dwarf with constant luminosity and strong emission longward of ∼60 Å (Fig. 8). The 2nd order grating efficiency for the LETG is a factor of ∼20 lower than 1st order, and 3rd order a factor of ∼11 lower, so the observed spectrum has negligible higher-order contamination out to the ∼180-Å limit of the LETG/HRC-S range. HZ43 is the most commonly used calibration source for the LETG and has been observed roughly twice a year since the launch of *Chandra* (Table 2).

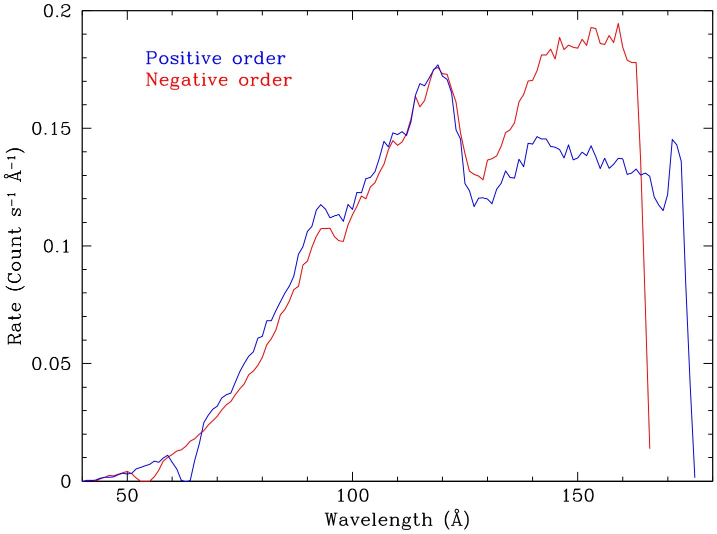

Fig. 8.— Background-subtracted positive- and negative-order spectra of HZ43 (ObsID 0059). Dips around 53 and 63 Å are from HRC-S plate gaps.

Table 2. *Chandra* Observations Used in Temporal Analysis

Source	ObsID	Total Exp. (s)	Net Exp. (s)	Obs. Date
HZ43	0059	40226	38757	Nov 12 1999
HZ43	1011	18762	18658	Mar 18 2001
HZ43	1012	20159	18488	Aug 18 2001
HZ43	2584	19189	19010	Jan 01 2002
HZ43	2585	20179	19990	Jul 23 2002
HZ43	3676	19981	19848	Dec 04 2002
HZ43	3677	20179	20010	Jul 24 2003
HZ43	5042	19853	19846	Dec 20 2003
HZ43	5044	19853	19666	Jul 19 2004
HZ43	5957	20186	20036	Feb 02 2005
HZ43	5959	18183	17921	Jul 29 2005
HZ43	6473	20181	20173	Mar 13 2006
HZ43	6475	20747	13640	Aug 07 2006
HZ43	8274	20188	19961	Mar 16 2007
PKS2155	0331	63166	62162	Dec 25 1999
PKS2155	1704	26005	25863	May 31 2000
PKS2155	1013	26833	24163	Apr 06 2001
PKS2155	3166	29983	29785	Nov 30 2001
PKS2155	3709	13832	12196	Nov 30 2002
PKS2155	4406	13854	12670	Nov 30 2002
PKS2155	5172	27166	25235	Nov 22 2004
PKS2155	6923	30179	29928	May 01 2006
PKS2155	6922	10083	10038	Jul 06 2006
PKS2155	7294	10104	9860	Jul 07 2006
PKS2155	7295	10151	10062	Jul 07 2006
PKS2155	8379	30140	29201	Apr 22 2007
PKS2155	9707	10184	10078	Jul 02 2008
PKS2155	9709	10188	10080	Jul 02 2008
PKS2155	9711	10189	10084	Jul 03 2008

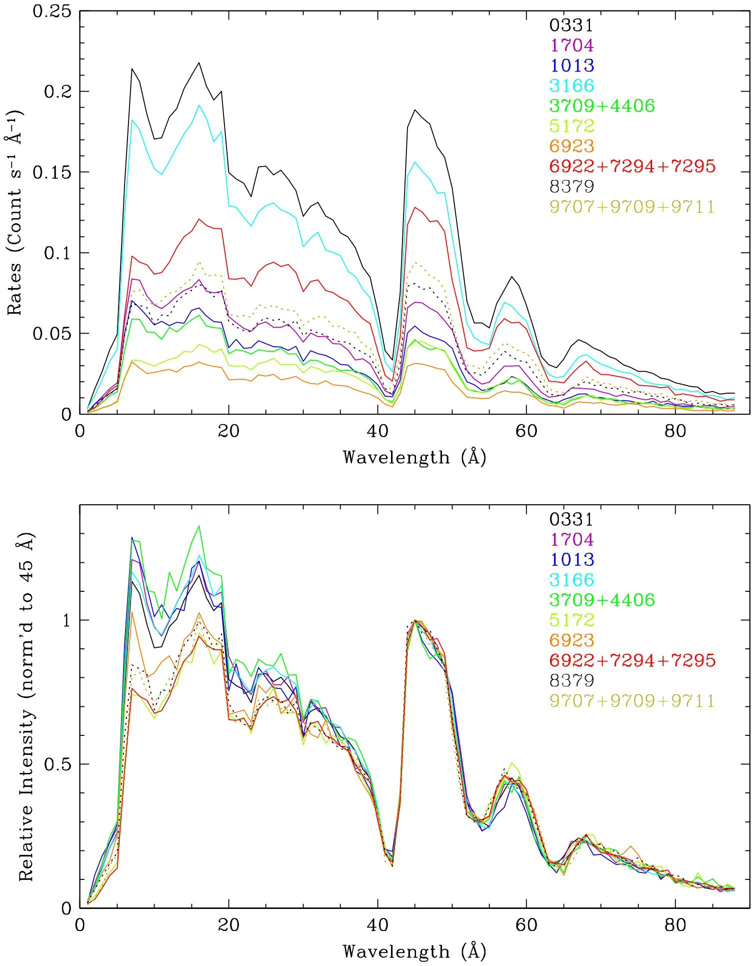

Fig. 9.— Background-subtracted spectra of PKS 2155-304 used in the temporal calibration, summing − and + orders. Dips around 57 and 67 Å are from HRC-S plate gaps. Top: Source intensity has varied by a factor of 7. Bottom: Spectra are normalized to intensity at 45 Å, showing that source hardness also varies, which affects the degree of higher-order contamination.

For shorter wavelengths, primarily on the center segment but with some overlap with HZ43 on the outer segments, we use PKS 2155-304, a blazar with strong emission out to ∼60 Å (see Fig. 9). Because of its strong emission at short wavelengths, PKS2155 has significant higher-order contamination at some wavelengths, particularly around the C-K absorption edge just shortward of 44 Å and at wavelengths beyond 60 Å. The effects of higher-order contamination are discussed in §4.2.

Light curves for each ObsID were inspected to remove periods of high background. Most ObsIDs required little or no filtering. The worst case was ObsID 6475 (HZ43) which lost 1/3 of its exposure. We also excluded roughly one-quarter tap (in detector coordinates) on both ends of plate 2 and on the inner end of plate 3 where event rates and/or PHDs are abnormal.

For the PKS2155 data, the extraction region for each dispersed X-ray spectrum is a long box 12 pixels wide in sky coordinates, centered on 0th order. Background regions of the same size are offset by 50 pixels to either side of the X-ray region, with the displacement chosen to minimize inclusion of cross-dispersion events diffracted by the LETG fine-support structure, which lie in a pattern that radiates from 0th order at integer multiples of 2.235°, relative to the primary dispersed spectrum. This extraction-region prescription is also used for HZ43 spectra on the center plate. On the outer plates, the HZ43 extraction regions for each ObsID are drawn by eye using 4-sided polygons, with a typical width of 15 pixels around 60 Å and 60 pixels around 180 Å. Background regions on the outer plates are offset from the X-ray region by 100 pixels.

X-ray extraction regions are narrower than used in standard *Chandra* data processing to reduce background at the expense of a small additional loss of diffracted X-ray events. PHDs from X-ray and background events are quite different, so we use background-subtracted PHDs in all our flight-data analysis. Even with background subtraction, it is desirable to minimize background contamination because the gain in the background region is not exactly the same as in the X-ray region; SAMP mean values typically vary by a few percent on 100-pixel scales, and so the PHDs of background events in the X-ray and background regions will be slightly different. Note that background rates on-orbit are roughly 100 times higher than in the lab, and background subtraction is required even for spectra of bright objects such as HZ43 and PKS2155.

Some X-ray events, dispersed by the LETG fine- and coarse-support structures, will be collected in the background regions, usually at a level of $\lesssim 2\%$ of the primary X-ray spectrum. Since we are only calibrating the spectral response of the HRC-S, X-ray-event contamination of the background data (which are subtracted from the X-ray region data) merely leads to a minuscule decrease in the statistical quality of the net X-ray PHD without altering its shape. Prior to background subtraction, contamination of X-ray spectra is usually at the few percent level, except for PKS2155 data at long wavelengths ($\gtrsim$60 Å) when the source is in a dim state, when background may exceed 10%. The net systematic error in determinations of mean SAMP after background subtraction is therefore never more than a few tenths of one percent, which is negligible compared to statistical uncertainties.

Another subtlety arises from the gradual drift of the *Chandra* aimpoint. Since late 1999, the position of 0th order has slowly moved by roughly 100 pixels in both RAWY (to larger values) and in RAWX (to smaller values; see Fig. 10). In order to study the change in detector gain over time, we analyze only data from a consistent detector region that is illuminated in all observations of a given target. For PKS2155, this is CRSU=7c–8c (4 subtaps), and for HZ43, CRSU=7c–8b (3 subtaps), where a,b,c denote thirds of a tap.

In the CRSV direction (along the dispersion axis), we divide source spectra into slices containing 2 to 7 taps (typically 3 or 4), yielding several thousand counts per slice. Because of aimpoint drift, the wavelength range for a given tap range will differ slightly from one observation to another, but the effect on determination of SAMP mean is usually negligible. An exception occurs when

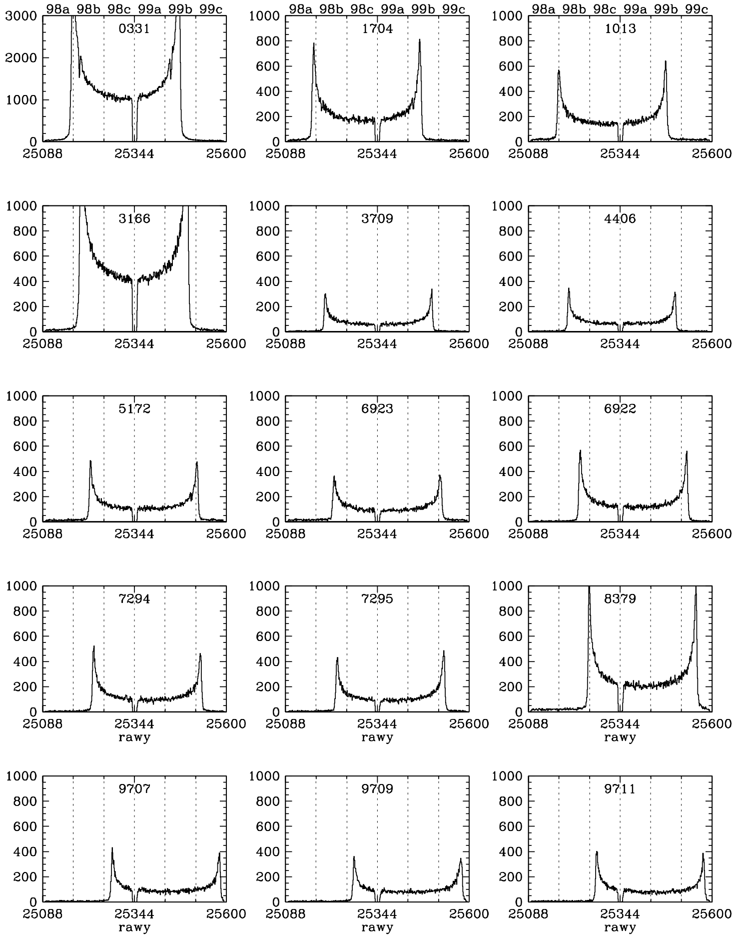

Fig. 10.— Distribution of 0th-order events in RAWY from PKS2155-304 observations. Dotted vertical lines mark 1/3-tap boundaries, with CRSV subtaps labeled across the top. A similar drift occurs in RAWX (CRSU).

the tap range encompasses a large change in spectral intensity, such as in the PKS2155 spectra across the C-K absorption edge (43.7 Å; from the HRC polyimide uv/ion shield) and at the I-$\mathrm{M_V}$ edge (20.0 Å; from the CsI photocathode). In such a case, aimpoint differences will lead to substantial differences in the relative exposure of different subtaps. We attempt to minimize these differences by isolating taps containing the the C-K edge (CRSV=76–77 and 119–121) and I-M edge (CRSV=87–89 and 108–110) and viewing their results more critically.

The last analysis subtlety is also related to aimpoint drift. As seen in Fig. 10, dithering exposes the outer edges of the dither pattern more than the center. To avoid giving data from particular U subtaps more weight in one observation than in another (because of the aimpoint drift), analysis results from each CRSU subtap are assigned equal weight, regardless of the number of counts; the statistical uncertainty for the combined result for a given spectral slice is based on the total number of counts. The effects of uneven dither-related weighting in the CRSV direction are not a concern (apart from the two absorption-edge exceptions noted above) because the spectrum lies along the CRSV axis and dither tends to smooth out the exposure.

To summarize the above, trimmed mean SAMP values and IQRs are calculated for background-subtracted data in specified CRSV tap ranges, for each of 3 or 4 CRSU subtaps, and then the means from the CRSU subtaps for each CRSV range are averaged with equal weights. Uncertainties in the means are estimated as the IQR divided by the square root of the total number of counts. PKS2155 observations that occurred within a day of each other (see Table 2) were combined for analysis. Results, with minor higher-order corrections described in §4.2, are shown in the top panel of Fig. 11. The bottom panel results are normalized to better illustrate the general trend of the gain decrease as a function of wavelength.

There is a gap around 0th order because of the lack of sufficient data short of ∼6 Å. HZ43 0th order data can be used to calibrate gain behavior around the aimpoint, but aimpoint drift and an unusually rapid gain decrease in this region early in the mission make this a challenging project, requiring study on a finer spatial scale. An approximate calibration of the aimpoint gain based on background data is described in §4.4. This short-wavelength range is, however, relatively unimportant for LETG observations and for now we interpolate the gain's temporal behavior from adjoining regions (see §4.3).

4.2. Corrections for Higher-Order Contamination and Intermediate Results

As noted earlier, the PKS2155 spectrum varies not only in intensity but also in shape, and has significant contamination from higher-order diffraction. In order to determine the effect of the higher orders on SAMP, we fit the positive and negative order spectra from each PKS2155 observation using LETG responses (ARFs and RMFs) through 12th order. The contribution of each order for ObsID0331 is shown in Fig. 12. As can be seen, nearly half of the observed events just short of the C-K absorption edge come from higher orders, and more than half at wavelengths beyond ∼80 Å.

Also shown, at the top of the bottom panel of Fig. 12, is the measured mean SAMP divided by the estimated mean SAMP in the absence of higher orders. The latter was determined by summing the product of the expected mean SAMP for each order m (determined from Eq. 1 by setting E to m times the 1st order energy) and the relative contribution of that order to the measured intensity at wavelength $m\lambda$. As seen in the figure, the net mean SAMP can be shifted by more than 10% at wavelengths longer than 80 Å, although the relative shifts when comparing results from one ObsID to another are never more than about 1/4 of the total shift. These corrections were applied to the data plotted in Fig. 11.

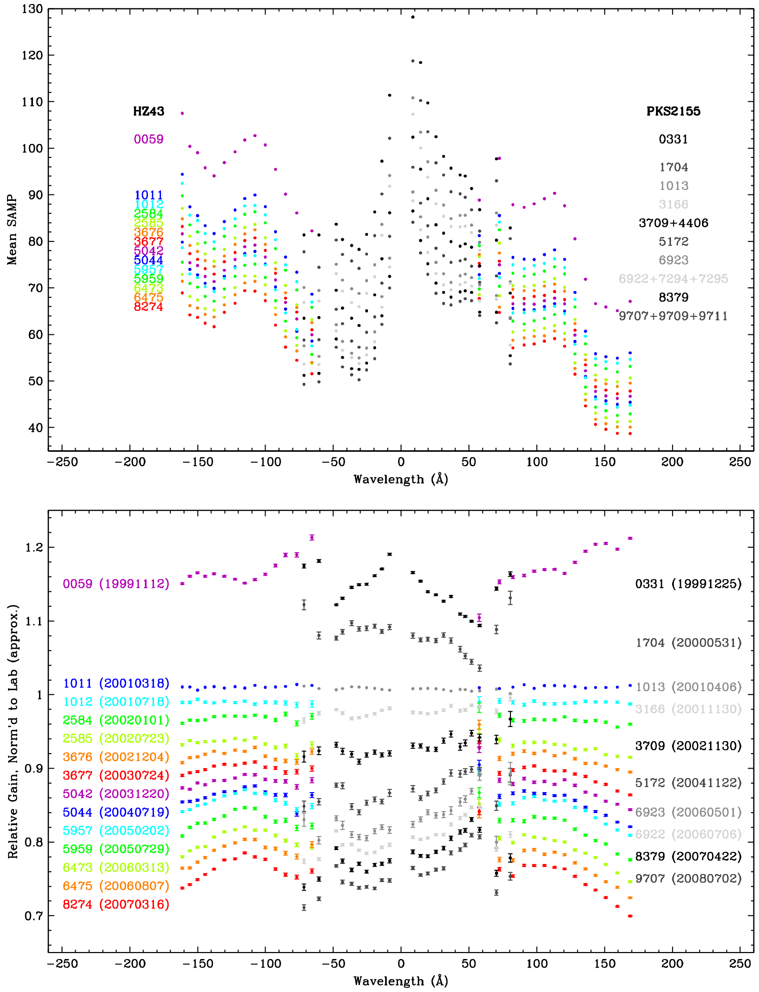

Fig. 11.— Top: Mean SAMP within slices of spectra from HZ43 (colors) and PKS2155 (gray shades), with some overlap on HRC-S plate edges. The peak toward short wavelengths is mostly due to the SAMP energy dependence (see Fig. 6). Bottom: Mean SAMP values are normalized to the average of ObsIDs 1011 and 1012 for HZ43, and a 3:1 weighted average of ObsIDs 1013 and 3166 for PKS2155. As discussed in §4.3, the gains during those ObsIDs were close to that during the lab measurements.

4.3. Temporal and Spatial Interpolation

Another way to display the data in Fig. 11 is mean SAMP versus time, as shown in Fig. 13 for the positive-order data. Least-squares fits to the data using exponential plus linear decay produced slightly better results than a double-exponential model, particularly on the outer plates (HZ43 data). Where PKS2155 and HZ43 data overlap near the ends of the plates, the fits are consistent within statistical uncertainties.

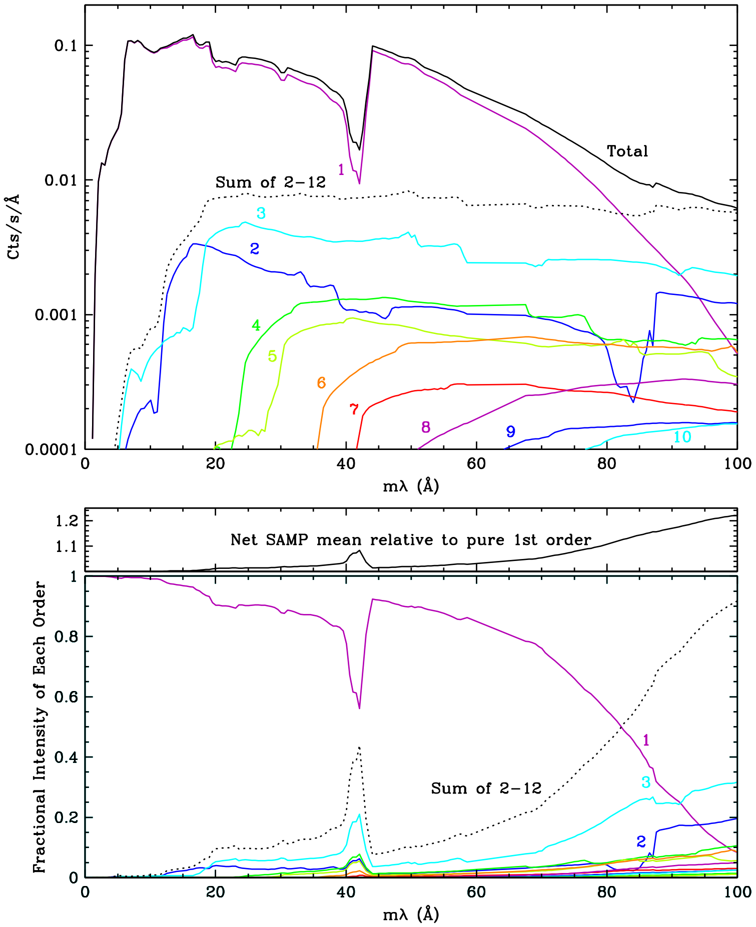

Fig. 12.— Top: Fits with 12 grating orders to the positive order spectrum of PKS2155 ObsID 0331. Bottom: Relative contribution of each order to the observed emission at $m\lambda$. Middle: Ratio of measured mean SAMP to mean SAMP without any higher-order contamination.

In order to apply time-dependent gain corrections we normalize the flight gains to the lab-based baseline gain map, implicitly assuming that when the lab and flight gains match at one detector location they match at all locations (but see §4.4). To do this we compare mean SAMP from flight data at positive and negative LETG wavelengths corresponding to B-K, C-K, and O-K emission with mean SAMP from lab data at corresponding detector locations. Taking into account systematic uncertainties (e.g., ignoring the +order B-K, which is very near a plate end with steeply varying gain, and giving the most weight to +order C-K, which falls in a region of slowly varying gain), we find that the lab gain best matches the flight gain very near the middle of 2001.

With the normalization thus established, we plot the fitted gain results for time intervals corresponding to roughly 3% changes in Fig. 14 and use spline fits for spatial interpolation. These curves are then used to determine the gain corrections for any observations as a function of time and V subtap, with temporal interpolation. Note that the corrections are not a function of U

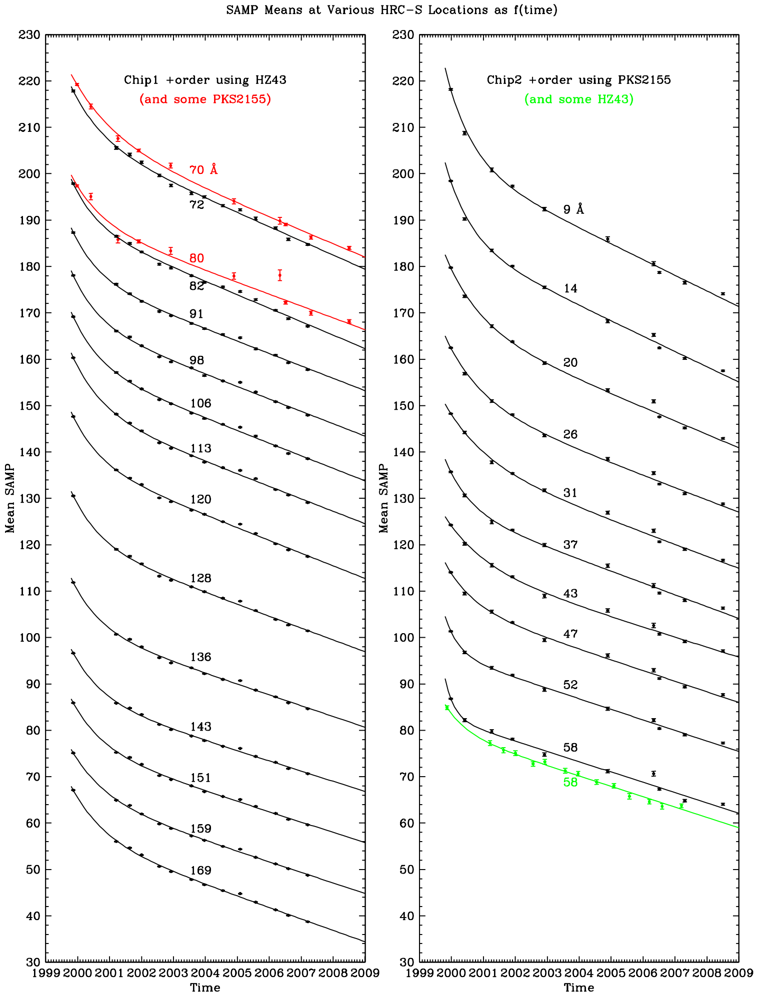

Fig. 13.— Mean SAMP as a function of time, with data from different wavelengths offset vertically in steps of 10. Smooth curves are from least-squares fits of exponential plus linear decay.

subtap, as our temporal gain calibration is derived from dispersed spectra, which only illuminate taps with CRSU∼8.

These temporal corrections are based on relative SAMP values so if we denote the relative gain for subtap i as $C_i(t)$, then Eq. 3 becomes

$$SPI = \frac{38.88}{m_i}\left(\frac{SAMP}{C_i(t)\mathcal{S}_i^{\rm norm}} - b_i\right) + 11.37. \qquad (5)$$

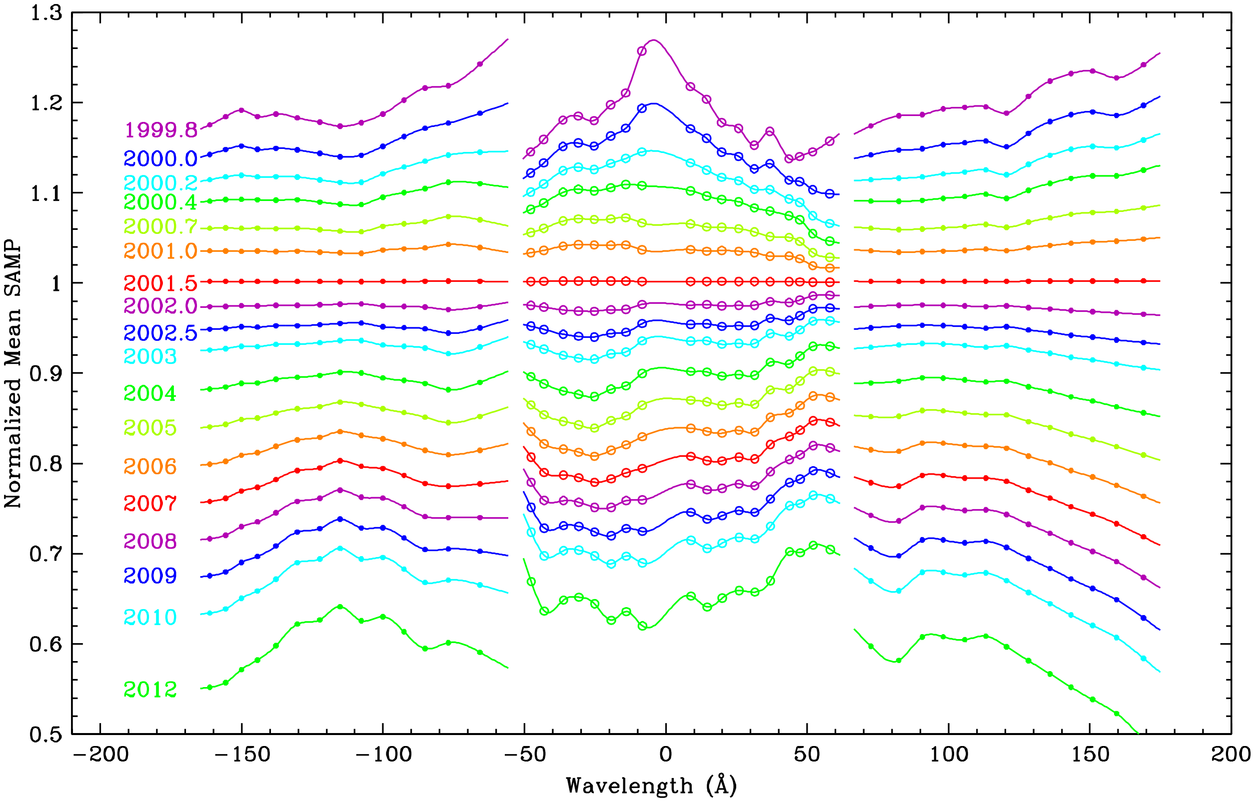

Fig. 14.— Relative gain curves derived from fits in Fig. 13. Gains are normalized to 2001.5, when flight gain most nearly matches the lab gain.

4.4. Spatial Gain Recalibration

After applying the temporal gain corrections above we confirmed that SPI pulse height distributions at a given wavelength from many different observations all look the same (e.g., Fig. 15). Data from multiple observations of the same source were then combined and the mean SPI of background-subtracted spectra was computed as a function of wavelength. Figure 16 shows the results from HZ43 and PKS2155 data. Merged PKS2155 data were adjusted for the effects of higher order contamination by applying the corrections described in §4.2. Also shown are results from an observation of HZ43 made 10' off-axis (ObsID 1170, Nov 1999), which shifts the spectrum by +34 Å.

As can be seen, the SPI means from the on- and off-axis HZ43 observations often disagree. Note in particular the trough in the off-axis curve near +95 Å, which corresponds to the gap between segments 0 and +1. Clearly the assumption in §4.3 that the detector gain in mid-2001 matches the lab gain everywhere is incorrect. The inconsistencies between on- and off-axis results must be the result of extra gain losses near the ends of each detector segment. This extra gain loss does not invalidate the temporal gain analysis, which was a calibration of relative changes at a given detector

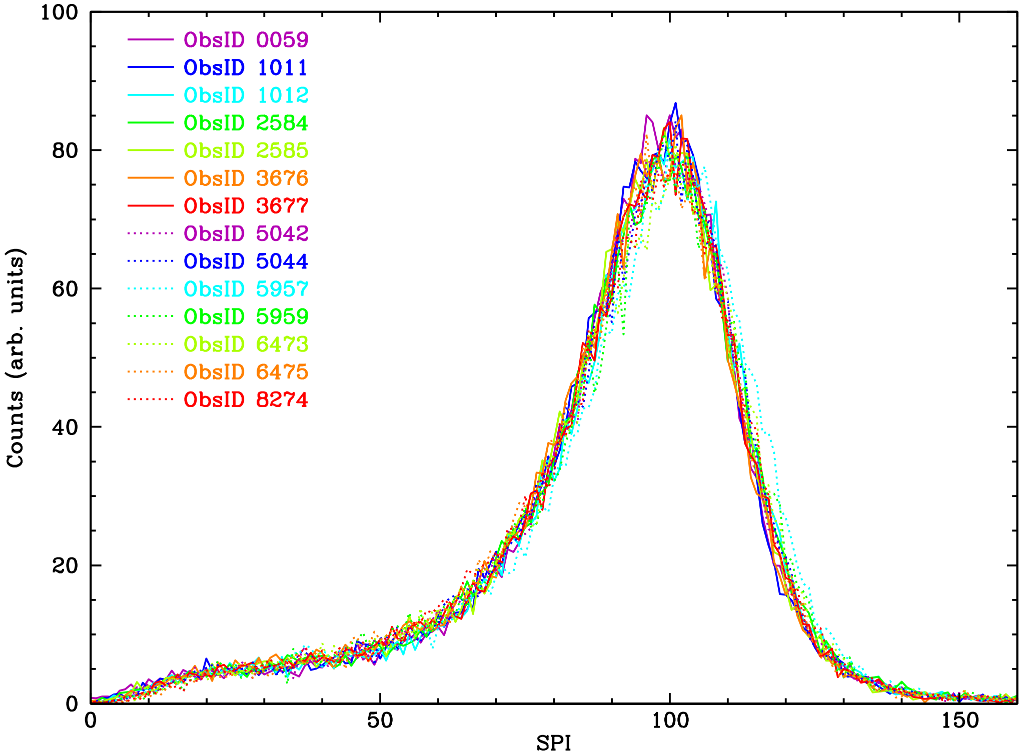

Fig. 15.— SPI distributions at +160 Å for all 14 HZ43 observations, normalized by total number of counts. The mean for ObsID 5957 (Feb. 2005) is slightly higher than the others, consistent with data in Fig. 13.

location once the HRC-S was in orbit. It does mean, however, that monochromatic photons will not produce the same pulse height everywhere on the detector. Off-axis observations, which shift the position of dispersed spectra on the detector, would therefore require different pulse-height filtering. In order to use a single pulse-height filter for all observations we must therefore determine the spatial dependence of the extra gain loss that mostly occurred soon after launch.

This is difficult to investigate using flight spectra, since there are very few off-axis observations. The next best thing is to use flight background data after excluding all X-ray source events, thus providing a sort of flat field dataset akin to the lab X-ray data. To cover all three plates, we must use data from grating observations, preferably long ones with weak sources.

The best set of background data comes from RXJ1856-3754, an isolated neutron star with 500 ks of exposure (ObsID 0331 in Mar 2000, ObsIDs 3380–3382 in Oct 2001). To exclude X-ray contamination, we remove a circle with radius 150 pixels (20”) around 0th order, a 100-pixel-wide box along the primary dispersed spectrum, and narrower boxes along the fine-support-structure cross-dispersion orders (plus and minus 1–3) and the three coarse-support-structure patterns around 0th order. Those exclusions remove a significant fraction of the detector area, particularly regions

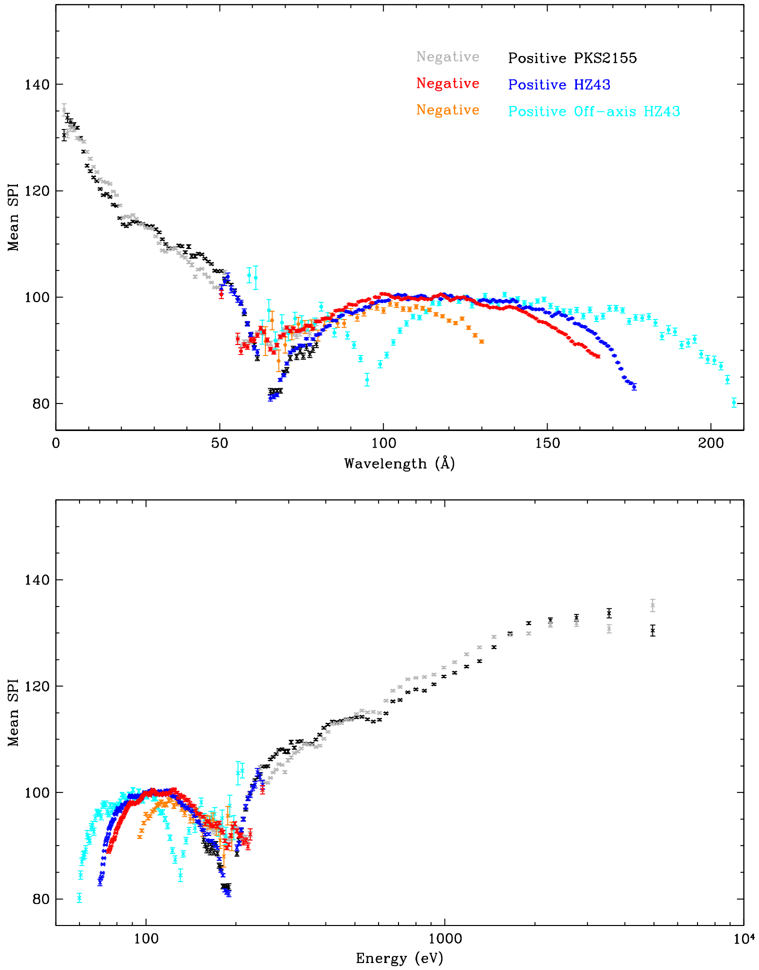

Fig. 16.— Mean SPI versus wavelength (top) and energy (bottom) for gain-corrected merged HZ43 and merged PKS2155 data, in 1-Å bins, prior to gain recalibration near detector segment ends. PKS2155 data have been adjusted for higher-order contamination. Also shown are results from an HZ43 observation made 10' off-axis (2-Å bins), which shifts the spectrum by 34 Å.

with CRSU=8 and near the aimpoint, so we also include data from a 50-ks observation of Sirius B (ObsID 9617, Jan 2008) that was displaced in the cross-dispersion direction by about 1' (1.8 taps toward higher CRSU), which helps fill in the low-count areas. The analysis region includes CRSU=7 and 8 and half of tap 9.

Background data from RXJ1856 and Sirius B were then analyzed in the same manner as for X-ray data, but binning spectral data by V subtaps instead of by wavelength. A plot of median SPI versus CRSV is shown in Fig. 17; medians were used because they were better behaved than means, particularly on plate 1. The medians from the central third of each tap (green) were systematically 3% lower than means from the other two subtaps, and have been adjusted in the figure. This effect

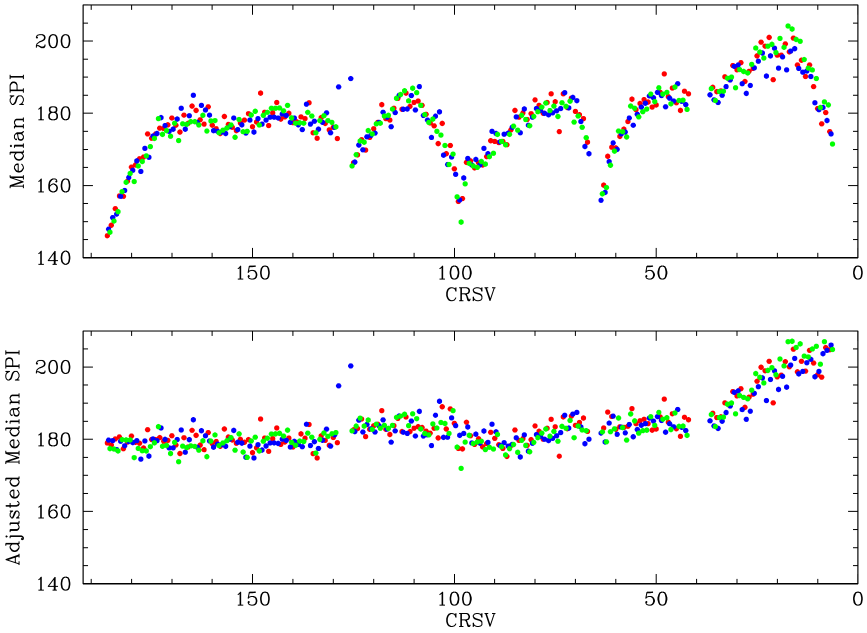

Fig. 17.— Median SPI versus CRSV for RXJ 1856 background data, in 1/3-tap bins. Medians from central V subtaps (green) have been increased by 3%. Top: before gain corrections. Bottom: after applying gain corrections on plate ends and near the aimpoint. In both cases, the gap around CRSV=40 is from exclusion of an off-axis X-ray source.

is not observed in the X-ray data.

As seen in the top panel of Fig. 17, the gain is not only low on the plate ends but also around the aimpoint. In both cases the extra gain decrease is presumably from extra charge extraction; there are many X-ray events at the aimpoint, and background emission is high along plate boundaries because of electric field edge effects. (Note that many events near the edges are removed during data processing via the 'bad pixel map.')

We fit piecewise smooth curves, one for each segment end and two around the aimpoint, to the data in the top panel of the form $S^{\rm BG} = S^{\rm BG}_{\rm Flat}(1 - A|({\rm RAWY} - {\rm RAWY}_0)|^p)$, where $S^{\rm BG}$ is the median SPI for the background data, $S^{\rm BG}_{\rm Flat}$ is the median SPI in regions without excess gain loss, ${\rm RAWY}_0$ is the point at which the gain starts to fall off, A is a constant, and p ranges from ∼1.4 (near the aimpoint) to ∼6.0 on the ends of segment+1. The gain around CRSV=155 is also adjusted manually to remove the small dip there. On the other side of the detector, the general rise in background gain across segment+1 toward small CRSU is not seen in X-ray data, and so is not adjusted.

SPI distributions from X-ray and background events differ substantially in shape—indeed, this is the basis of the pulse-height filtering described in §5—so we do not expect the background mean-SPI behavior to perfectly track X-ray-event gain. Differences are also seen between results from

Sirius B (observed during Solar Minimum) and RXJ1856 (Solar Maximum), mostly due to different charged particle composition and/or energy distributions, although the qualitative behavior of the gain fall-off remains the same.

The fits derived above are used to adjust mean SPI values for X-ray data according to $SPI^{\rm new} = SPI^{\rm old}/(1 - fA|({\rm RAWY} - {\rm RAWY}_0)|^p)$, where f is a factor used to optimize agreement between the on- and off-axis HZ43 results seen in Fig. 16. The values of f vary from 0.45 (on the $-$ side of the center plate) to 1.80 (on the + side of the same plate). Having no convenient X-ray data to guide gain adjustments near the aimpoint, we assume $f = 1$.

New gain maps were created by adjusting the temporal correction terms in Eq. 5 such that $C_i^{\rm new}(t) = C_i^{\rm old}(t)/(1 - fA|({\rm RAWY} - {\rm RAWY}_0)|^p)$, and all calibration data were reprocessed using the new gain corrections. Reanalysis of the HZ43 and PKS2155 data yields the results shown in Fig. 18. Results are also shown from RS Oph (ObsIDs 7296 and 7297, Mar and Apr 2006). That spectrum has almost all its emission between 15 and 30 Å and therefore provides a good medium-wavelength check for the higher-order corrections applied to the PKS2155 data.

4.5. Discussion

As seen in Fig. 18, results from on- and off-axis pointings now agree very well, and agreement between different sources and between negative and positive orders is also improved compared to Fig. 16. There remains, however, an apparent mismatch between results from the center segment and the outer segments near the plate gaps around 60 Å, as well as a somewhat suspicious dip in mean SPI at slightly longer wavelengths. Data points for negative and positive orders agree quite well, however, and changing the gain on any of the plate ends ruins the agreement among those results and between the on- and off-axis HZ43 data. This is a difficult issue to study because of the paucity of counts in this wavelength range.

If mean SPI versus wavelength follows the simple relationship described by the reference gain curve (shown as a dotted line in Fig. 18) at energies below those investigated in the lab (B-K is 183 eV, or 68 Å), perhaps with some additional flattening, then it would appear that the gain corrections need to be larger around 65–90 Å. However, the electron yield from the HRC-S CsI photocathode[4] (see Fig. 19) has a large dip at roughly the same wavelengths. There also appear to be slight jumps in mean SPI around 17 Å (∼730 eV) and 20 Å (∼620 eV) corresponding to the Cs-M and I-M edges, respectively. This suggests that the CsI electron yield may have a small effect on event pulse height, so that the dip in mean SPI near the plate gaps is a 'real' effect and not due to incorrect gain calibration.

In practice, the only potentially significant problem with this final gain map is the ∼10% gain

---

[4] M. Juda, 1997, `http://hea-www.cfa.harvard.edu/~juda/memos/mcad_QE_model.ps`, "A Simple Model for the QE of CsI Coated MCPs."

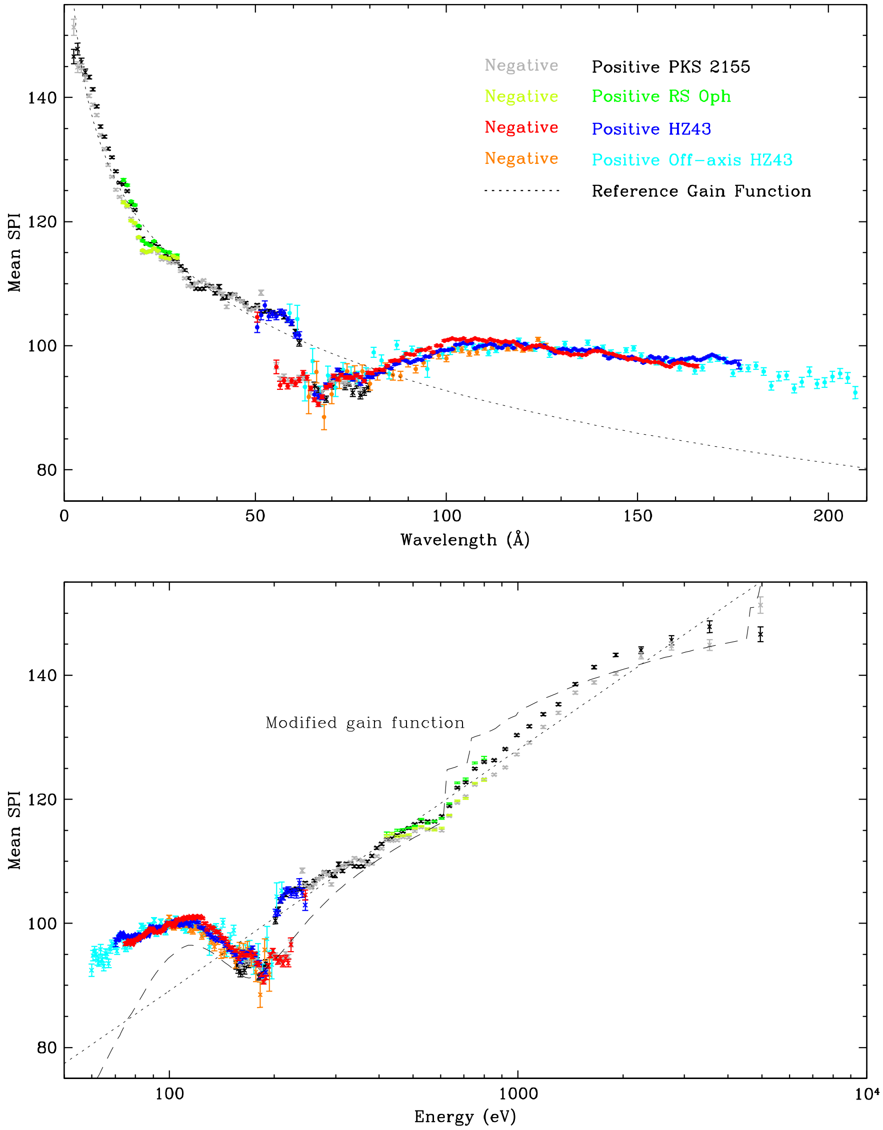

Fig. 18.— Mean SPI versus wavelength after final gain adjustments (c.f. Fig. 16). Also shown is the reference gain function derived from lab data (see Fig. 6) and a modified gain function (bottom panel) that roughly illustrates the possible effects of energy dependent electron yield from the CsI photocathode (see Fig. 19).

mismatch on either side of the central/outer plate gaps. The wavelength region around 60 Å, however, is not very important for most astrophysical sources, having very few emission lines. In any case, the only impact is that pulse-height filtering around the plate gaps must err on the high-SPI side to avoid removing too many X-ray events, thus slightly reducing the filter's effectiveness at removing background events, as described in the next Section.

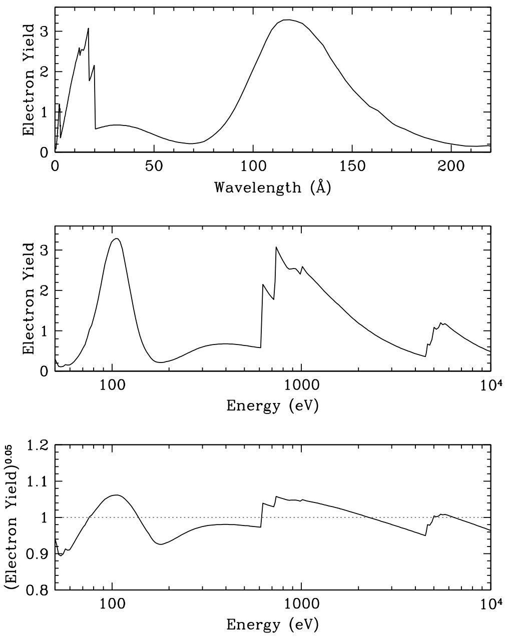

Fig. 19.— Model electron yield from the HRC-S CsI photocathode, from work by M. Juda (1997)[3]. The bottom panel shows the adjustment factor applied to create the modified gain function in Fig. 18. Some features in the model electron yield seem to correspond to those in the plot of mean SPI versus wavelength.

5. PULSE HEIGHT FILTER

Figure 20 shows SPI pulse height distributions for background and X-ray events. The background PHD varies somewhat with the solar cycle, and the total background rate roughly doubles going from Solar Maximum to Solar Minimum [5] because the solar wind weakens and more higher energy cosmic rays reach the inner solar system, but the main point is that the background event distribution extends to much higher channels than for X-rays. By rejecting events with SPI above a carefully chosen wavelength-dependent threshold, we can therefore remove a large fraction of background events with little loss of X-ray events.

Figure 21 shows the SPI values required to remove various fractions of X-ray events (upper rainbow-color traces). The bottom traces plot mean SPI for background subtracted HZ43 and PKS2155 data. After considering the scatter in each trace, gain residuals in some observations after applying the gain map, and the fraction of background versus X-ray events that would be removed for each curve, we decided that a filter that removes between 1.0 and 1.5% of X-ray events would be a good compromise.

---

[5] T. Isobe and M. Juda, `http://cxc.harvard.edu/contrib/cxchrc/Stowed_study/hrc_stowed_position_study.html`, "High Resolution Camera Stowed Background Study."

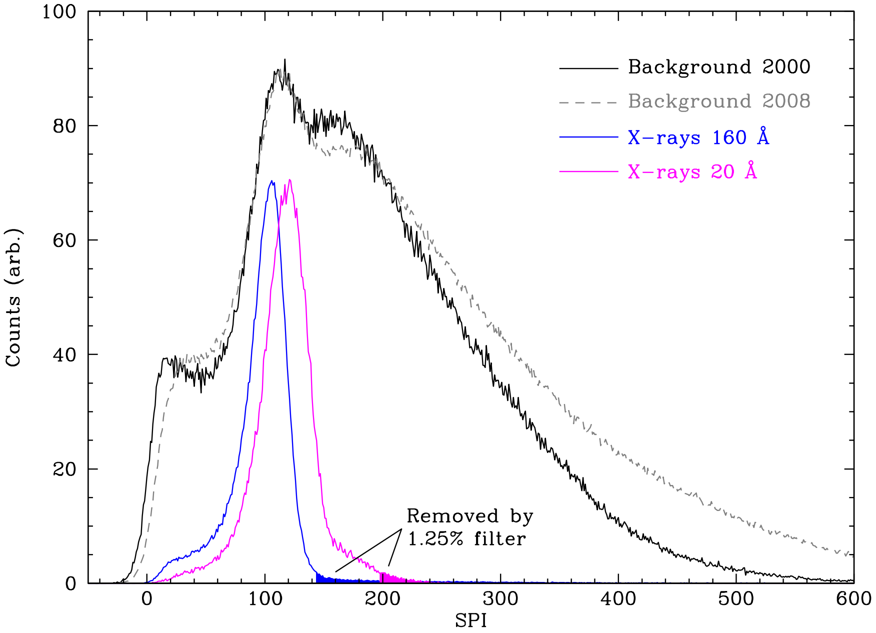

Fig. 20.— SPI pulse height distributions for flight background in 2000 and 2008, and for X-rays at 160 and 20 Å. Data are taken from RXJ1856 (ObsID 0331), Sirius B (Obs 9617), HZ43 (many ObsIDs), and RS Oph (ObsIDs 7296 and 7297). The difference between background and X-ray PHDs is the basis of the pulse-height filter.

As described in §4.2, we can correct SPI means for higher order contamination quite accurately, but the effects of HOs on the high-channel tails of SPI PHDs are harder to model. By comparing PKS2155 results with those from RS Oph and HZ43, which have negligible HO contamination, we empirically determined that using HO offsets for the SPI-filter values that are 3.9 (4.45) times larger than those used to correct SPI means works well at bringing the 1.5% (1.0%) filter limits from PKS2155 into agreement with the other sources. One exception is around the C-K edge (∼45 Å), where HO contamination effects are especially large; we interpolate results across this range.

Results are shown in Fig. 22, where the red curve marks the filter limit. Statistics near the plate gaps are relatively poor and, as noted in §4.5, there are some inconsistencies from one plate to another. We conservatively interpolate across these regions, erring on the side of removing too little background when in doubt.

6. RESULTS AND DISCUSSION

Results from application of the SPI filter are shown in Fig. 23. For comparison, results using an older gain map and the PI ‘light’ filter, which were based on PHA rather than SAMP and which did not take into account temporal gain loss, are also shown. Results are plotted relative to Level

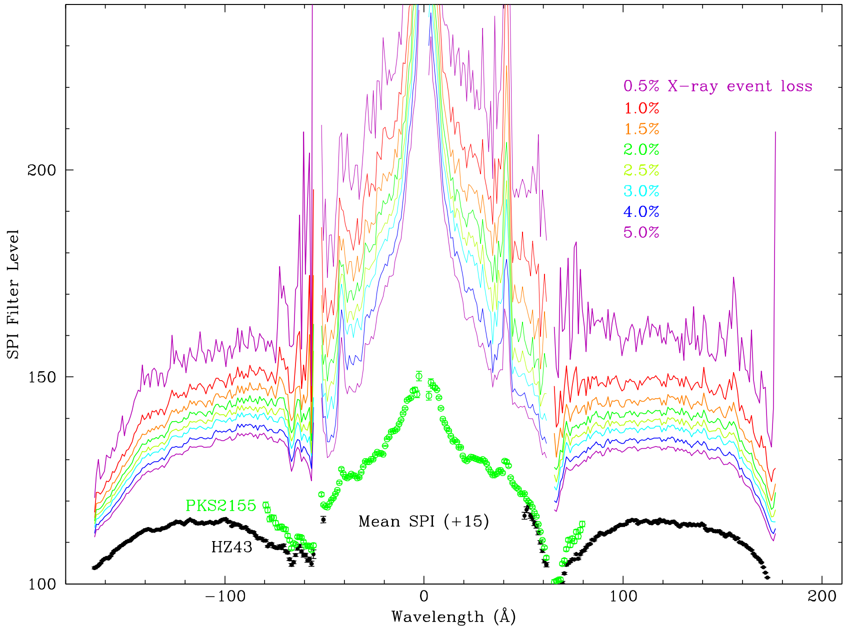

Fig. 21.— SPI limits for various levels of filtering, using data from HZ43 and PKS2155 (without higher order corrections). For clarity, overlapping data points near the plate gaps are not shown. SPI means (shifted upward) are shown in the bottom traces for comparison.

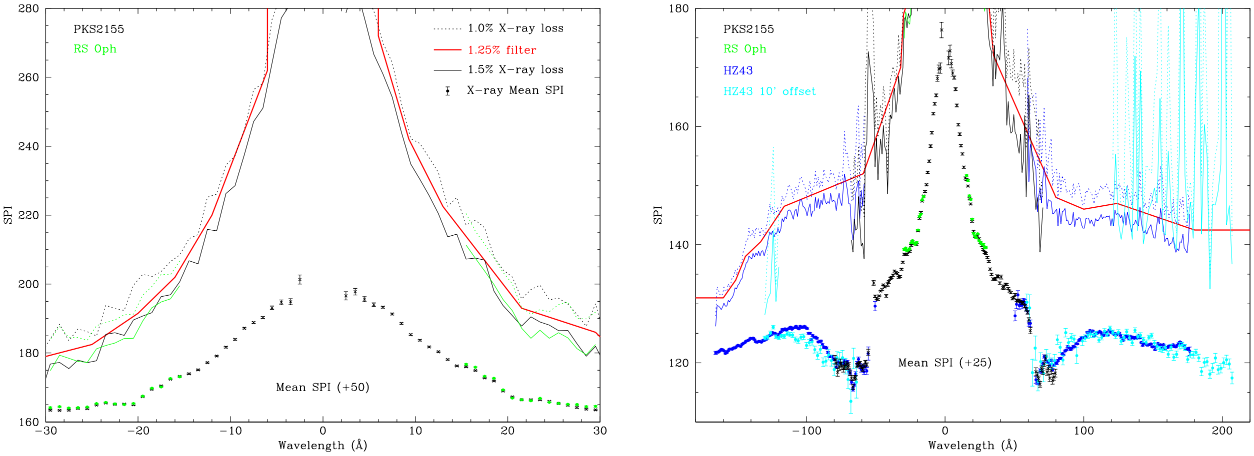

Fig. 22.— SPI limits for 1.0% and 1.5% filtering levels, emphasizing short wavelength (left panel) and long wavelength (right panel) ranges. The red line marks the 1.25% filtering specification. SPI means (shifted upward) are shown in the bottom traces for comparison. PKS2155 results have been corrected for higher-order contamination.

2 data, without any PHA filtering. Current standard processing excludes events with PHA=255, which (for 2008 data) removes about 8% of background events, though rather nonuniformly because of gain variations across the detector.

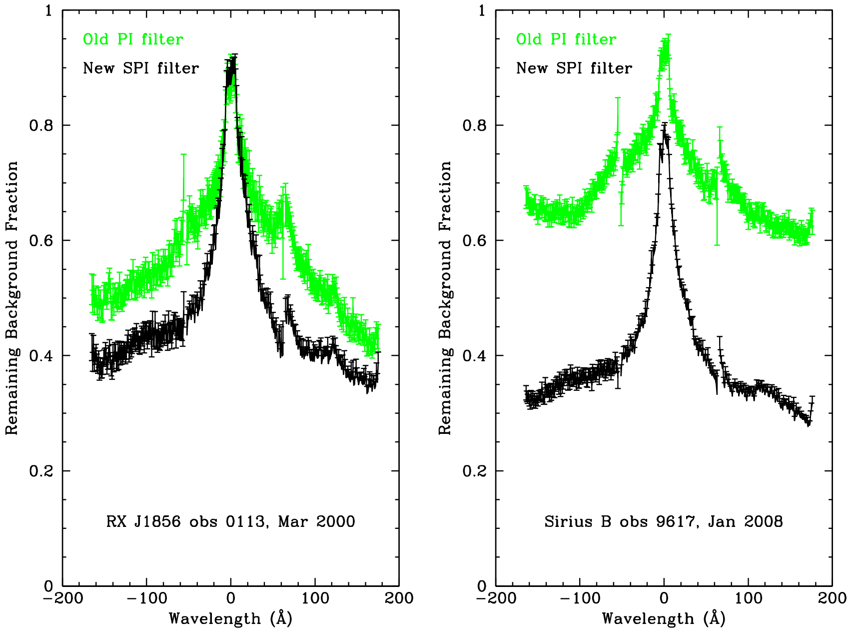

Fig. 23.— Comparison of old and new filters, as applied to data from 2000 and 2008. Background fractions are relative to Level 2 data that include events with PHA=255. Less background is removed at short wavelengths because the filtering threshold must be raised to avoid excluding X-ray events.

In addition to showing that the new SPI filter is more effective than the old PI filter, especially for recent observations, Fig. 23 shows that the SPI filter removes a higher fraction of the background in 2008 data than in 2000. This is because of solar-cycle variations in the energy distribution of the charged particles that cause background events (see Fig. 20 and §4.4). The improved filtering near Solar Minimum partially compensates for the higher raw background rate at those times, which is more than twice as high as during Solar Maximum. Raw and filtered background rates (for 2008 observations) when using the standard ‘bowtie’ LETG spectral extraction region are shown in Fig. 24.

7. Effect on Higher Orders

While the SPI filter removes $1.25 \pm 0.25\%$ of 1st order X-ray events, the fraction for higher orders is larger. Depending somewhat on photon energy, a factor of two difference in energy corresponds to a shift in mean SPI of ∼8%, and the mean of 8th order will therefore be about 25% higher than 1st order (see Fig. 20). As an example, about 11% of the 8th order of 20 Å ($m\lambda = 160$ Å) is removed by the filter. Extra filtering of higher orders will have negligible effect for nearly all analyses, but should be considered if deliberately studying wavelength ranges with heavy higher order contamination.

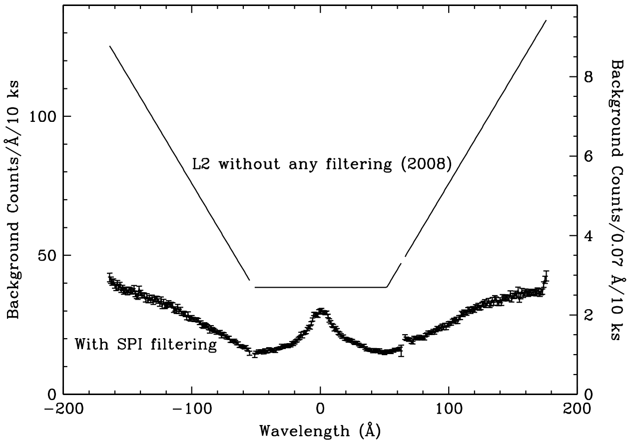

Fig. 24.— The LETG/HRC-S background rate with and without filtering, using the standard ‘bowtie’ spectral extraction region. The right-hand axis units refer to the ∼0.07-Å FWHM resolution of the LETGS.

8. FUTURE WORK

Analysis of the lab data can not be meaningfully improved any further, although the reference gain function can be adjusted to better fit lower energy flight data (see Fig. 18), which may slightly tighten the PHDs. Improvements to the reference function would be helped by gratingless off-axis observations of HZ43 around the plate gaps to calibrate the plate-end gain fall-off and resolve the discrepancies between the center and outer plates around 60 Å. Using the new gain function, the effects of higher order contamination on mean pulse height could be better estimated for PKS2155. Corrections for higher order contamination, though very small, could also be applied to RS Oph and HZ43 data.

An improved baseline gain map and plate-end gain recalibration could then be applied to HZ43 and PKS2155 data for the temporal gain calibration analysis. The gain map from our current work can be used to flatten the gains so that errors arising from variable subtap weights, particularly in the CRSU direction, would be negligible; the full width of the dispersed spectra could be analyzed, increasing the number of counts by roughly 20%. When binning spectra by CRSV, data could also be separated by wavelength, e.g., on either side of the Cs and I absorption edges, where small but relatively abrupt gain changes are expected.

The temporal gain decay can be well predicted out to c. 2012, but the fits in Fig. 13 will eventually need to be redone including newer data. So far there has been no indication of a decrease in HRC-S QE caused by gain loss (as low-channel events are lost below the signal-processing threshold), but it may be necessary before the end of the *Chandra* mission to raise the voltage on the HRC-S microchannel plates and recalibrate the gain.

The authors were supported during the course of this research by the Chandra X-ray Observatory Center, which is operated by the Smithsonian Astrophysical Observatory for and on behalf of the National Aeronautics Space Administration under contract NAS8-03060.

REFERENCES

Kraft, R. P., et al. 2000, Proc. SPIE, 4012, 493

Brinkman, A. C., et al. 2000, ApJ, 530, 111

Brinkman, A. C., et al. 2000, Proc. SPIE, 4012, 81

Kraft, R. P., Chappell, J., Kenter, A. T., Kobayashi, K., Meehan, G. R., Murray, S. S., & Zombeck, M. V. 1997, Proc. SPIE, 3114, 53

Fruscione, A., et al. 2006, Proc. SPIE, 6270, 62701V

This preprint was prepared with the AAS LaTeX macros v5.2.