HRC-I Gain Correction

Jennifer Posson-Brown[a] and Vinay Kashyap[a]

[a]Harvard-Smithsonian CfA, 60 Garden St, Cambridge, MA, USA

ABSTRACT

We study the gain variations in the HRC-I over the duration of the *Chandra* mission. We analyze calibration observations of AR Lac obtained yearly at the nominal aimpoint and at 20 offset locations on the detector. We show that the gain is declining, and that the time dependence of the gain can be modeled generally as a linear decrease in PHAs. We describe the spatial and temporal characteristics of the gain decline and discuss the creation of time-dependent gain correction maps. These maps are used to convert PHAs to PI channels, thereby removing spatial and temporal dependence, and allowing source pulse-height distributions to be compared directly regardless of observation date or location on the detector.

Keywords: Chandra, HRC-I, gain correction

1. INTRODUCTION

Regular monitoring of both the HRC-I and the HRC-S have shown that there exists a steady decline of the gain since the launch of the *Chandra* observatory.[1,2] The gain decline takes the form of a shift in the pulse height amplitude (PHA) profiles of observed sources to lower detector channels, and is expected to occur as the cumulative dosage on the MCPs increases.[3]

Gain monitoring serves as a proxy for the health of the HRC instruments. It is therefore necessary to establish a baseline behavior of the gain drop so that departures from it can be detected quickly. Here we describe the characteristics of the spatial and temporal gain corrections on the HRC-I, and develop tools that can also be applied to the HRC-S (Posson-Brown et al, in preparation).

Furthermore, the low background, large field-of-view, and the comparatively high oversampling of the *Chandra* PSF provided by the HRC-I makes it an attractive instrument for large area surveys. The main hindrance to this type of usage is the poor spectral resolution of the micro-channel plate detectors. However, while poor, the detectors are capable of sufficient spectral discrimination to allow crude estimates of the spectral shape via hardness ratios and quantile-width diagrams.[4] In order to facilitate such usage, it is necessary to calibrate the PHA values such that any spatial and temporal variations are accounted for.

1.1 Pre-Flight Gain Calibration

The HRC-I gain response was measured during pre-flight ground calibration with a series of flat field maps at six energies spanning 183 – 6404 eV. Gain maps were created for each energy by calculating the median PHA value for the events over the detector in half-tap (128 x 128 pixels) bins. Each map was then normalized by the mean value in its central $\frac{1}{9}$ area, thus generating, in its reciprocal, a multiplicative gain correction map, designed to modify the off-axis PHAs to match the PHAs in the center of the detector and thus create "pulse invariant" (PI) pulse height units. Finally, these normalized correction maps were averaged to create a single "pre-flight gain map" for the detector.[5] This gain correction map is available in the *Chandra* Calibration Database as the file

`$CALDB/data/chandra/hrc/bcf/gain/hrciD1998-10-30gainN0002.fits`

On October 4, 1999, shortly after launch, the voltage on the HRC-I MCP was lowered. A series of calibration observations of AR Lac were carried out at the aimpoint and 20 offset locations at both the high and the new flight voltage settings to determine any changes to the gain.

Further author information: (Send correspondence to J.P.B.)
J.P.B.: E-mail: jpossonbrown@cfa.harvard.edu, Telephone: 1 617 495 7025
V.K.: E-mail: vkashyap@cfa.harvard.edu, Telephone: 1 617 495 7173

We computed a new gain correction map for the flight voltage settings based on these observations,[5] following a process which will be discussed in §5.1. This map is available in the *Chandra* CALDB as the gain correction map for flight voltage observations:*

`$CALDB/data/chandra/hrc/bcf/gain/hrciD1999-10-04gainN0003.fits`

Here, we derive gain correction maps at the flight voltage that correct for both spatial and temporal variations in the gain. We describe the observations used to derive the new gain correction maps in §2, and the data reduction steps in §3. Characteristics of the PHA temporal decline are discussed in §4, and the time dependent gain maps are computed in §5. Finally, we discuss our analysis in §6 and summarize our results in §7.

2. OBSERVATIONS

In order to monitor the gain response of the HRC-I, yearly calibration observations are taken of AR Lac (Table 1) at 21 locations on the detector ($(Y_{\rm sim}, Z_{\rm sim}) = (0', 0'), (0', \pm 2'), (\pm 2', 0'), (\pm 2', \pm 2'), (0', \pm 4'), (\pm 4', 0'), (0', \pm 6'), (\pm 6', 0')$, and $(\pm 10', \pm 10')$). The observations in each yearly set are done successively, or, if this is not possible (due to an interruption such as a radiation shut-down), they are done as closely together as possible. Each observation is nominally 1 ks long; however, effective observation times may be shorter because of background flares.

Table 1. AR Lac stellar parameters

Parameter	**Value**
Other Names	HR 8448 / HD 210334 / RX J2208.6+4544 / HIP 109303
$(\mathrm{RA,Dec})_{ICRS2000.0}$	(22:08:40.818, +45:44:32.12)
$m_{\rm V}, B-V$	6.13, 0.72
Distance	$42 - 47$ pc
Spectral Type	G2IV/K0IV (RS CVn)
$M_{\rm V}$	3.5/3.3
Masses	1.3/1.3 $M_\odot$
Radii	1.8/3.1 $R_\odot$
Ephemeris	1.98316^d ; conjunction @ 2445611.6290 HJD

After the operating voltage of the HRC-I was lowered on 4 October 1999 (see §1.1), we have obtained 8 sets of AR Lac observations at the current (low) voltage. These observations were carried out in October 1999 (AO 1), December 1999 (AO 1), December 2000 (AO 2), January 2002 (AO 3), February 2003 (AO 4), November 2004 (AO 5), October 2005 (AO 6), and September 2006 (AO 7).

3. DATA REDUCTION

A pulse height amplitude (PHA) is assigned to each event based on the total charge collected by the cross-grid charge detector (CGCD) of the HRC-I and is recorded in the Level 1 event lists. We monitor the gain response by tracking the median PHA over time at each of the 21 observation locations. In all cases, we first reduce the data with CIAO (v3.4; CALDB v3.3), following the HRC Data Preparation analysis guide† and analyze the data with pre-packaged and custom-built IDL routines (e.g., `PINTofALE`[6]).

We extract source events from an 800x800 box centered on the nominal observation location in chip coordinates. Background is estimated by collecting the events in the *same location*, but from the 20 other observations carried out in that cycle. The background counts thus accumulated are normalized by their appropriate exposure times prior to subtracting them from the counts accumulated at the source location. The median values of the source PHA spectrum are estimated using Monte Carlo simulations.‡

*Note that the date encoded in the filename, 4 October 1999, marks the beginning of the calibration program to obtain data at different voltage settings. Therefore, for any observations carried out at the high voltage setting after this date, the `gainfile` parameter in `hrc_process_events` must be manually and explicitly set to use the pre-flight gain correction map.

†http://cxc.harvard.edu/ciao/guides/hrc_data.html

‡For each detector channel, the posterior probability distribution of the source intensity is computed, taking into

4. CHARACTERISTICS OF PHA DECLINE

The median PHAs, calculated as described in §3, are shown in Figure 1 as a function of observation location. The gain response decreases steadily and monotonically at all the monitored locations, as evidenced by the consistent lowering of the median PHAs. (But note the large drop in median PHA at the aimpoint between December 1999 and December 2000.) The errors on the medians are typically ≤ 1 channel.

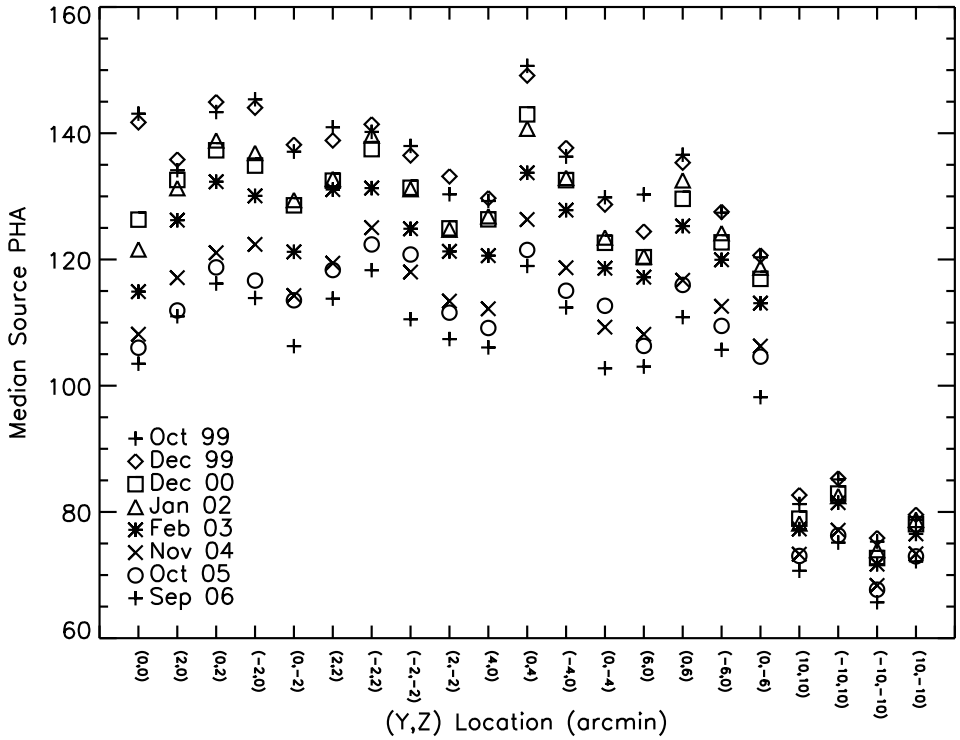

Figure 1. Median PHAs for HRC-I AR Lac observations. The background subtracted median PHAs are plotted in a vertical row for all observations at a given location, for data obtained during different cycles, all at the low voltage setting. Data from different cycles are marked with different point types (see legend). The statistical errors on the medians are typically ≤ 1 channel. Note that the median PHAs drop monotonically.

With the exception of the aimpoint, the time dependent decrease in the median PHAs at all the locations are all well-fit by straight lines. Even the data from the aimpoint can be well-fit by a straight line after December 2000. It is apparent that the aimpoint underwent an as yet unexplained stressing that resulted in a large gain drop soon after the voltage was lowered. With that exception, the gain drop is well-fit as a linear decrease at all locations and at all times. Furthermore, the best-fit slopes are roughly equal for all cases (excluding the four pointings at $14.14'$ off-axis), indicating that the rate of gain decline is relatively uniform across the detector (Figure 2).

5. TIME DEPENDENT GAIN CORRECTION

In order to calculate time-independent pulse invariant (PI) values, we carry out two stages of corrections to PHA($\vec{\mathbf{x}}$,t) . (Here $\vec{\mathbf{x}}$ denotes location on the detector and t denotes observation start time.) Since the gain decline is linear in time and similar in rate at all monitored locations (see Figure 2), we separate the variables and compute the spatial and temporal gain corrections independently. At each observation epoch, the raw

account the measured background.[7] Source counts realizations are repeatedly drawn from this distribution and the median is computed at each iteration. The average value of the medians obtained from such Monte Carlo simulations is reported as the median PHA and the standard deviation of the sample as the 1σ errors on it.

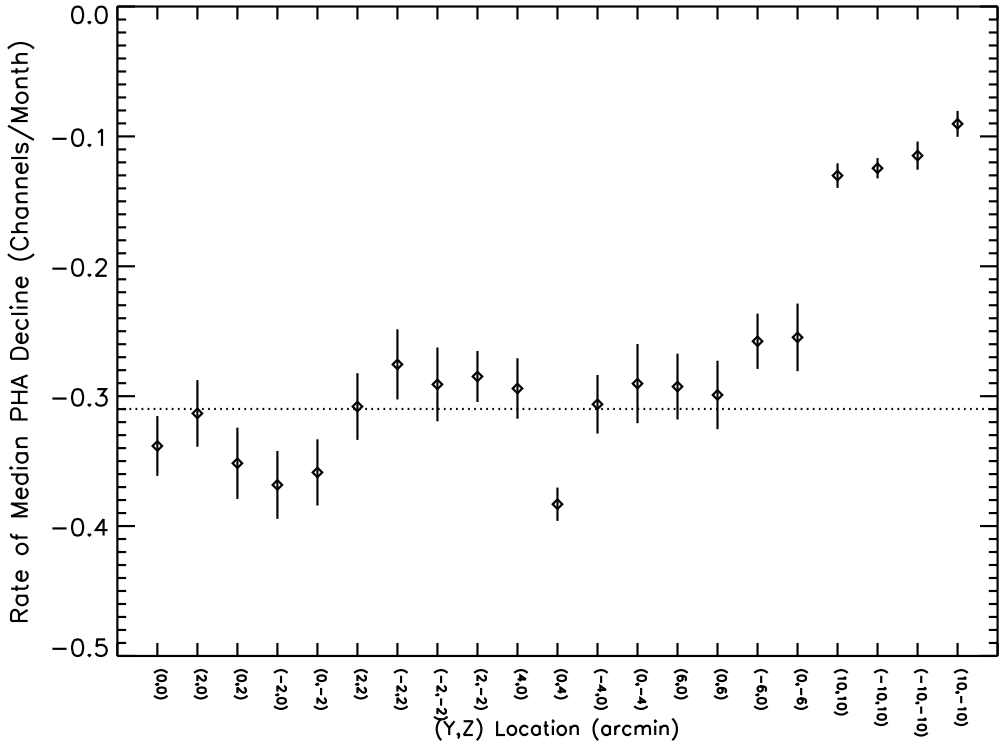

Figure 2. Uniformity of gain decline across the HRC-I. The best-fit slopes from the linear fits to the median PHAs are shown for each location. The average of the slopes (neglecting the four pointings at $14.14'$ offaxis) is shown as the dotted horizontal line. To a first approximation, excluding the far offaxis locations, the gain decline is uniform across the detector.

PHA($\vec{\mathbf{x}}$,t) are multiplied by a non-linear gain correction surface g($\vec{\mathbf{x}}$|t) that carries out a "flat-fielding" of the raw PHA values, i.e., at each epoch the PHA at location $\vec{\mathbf{x}}$ are transformed to what the PHA are at the aimpoint location. After this "flat-fielding", the PHA are no longer a function of $\vec{\mathbf{x}}$, and thus will be denoted PHA($\vec{\mathbf{0}}$|t). A time dependent correction, TC(t) is then applied to PHA($\vec{\mathbf{0}}$|t) to transform them to PI, which are effectively the same as the PI values as on October 1999, when the HRC-I flight voltages were reset. Thus,

$$\begin{aligned} \mathrm{PI} &= \mathrm{PHA}(\vec{\mathbf{x}},\mathrm{t}) \times \mathrm{g}(\vec{\mathbf{x}}|\mathrm{t}) \times TC(\mathrm{t}) \\ &\equiv \mathrm{PHA}(\vec{\mathbf{0}}|\mathrm{t}) \times TC(\mathrm{t}) \end{aligned} \tag{1}$$

The time-dependent gain correction is currently implemented via a series of gain correction maps, each of which include the effects of both g($\vec{\mathbf{x}}$|t) and TC(t) for specific epochs. This solution is the least disruptive of existing CALDB and CIAO programs and data structures. Work on producing a more elegant implementation, e.g., via a functional correction, is ongoing.

5.1 Spatial Corrections

We correct for the spatial variations in gain response by creating a series of 7 correction maps g($\vec{\mathbf{x}}$|t), one from each AO. (We used the Dec 99 dataset when creating the AO1 map.) We compute them as modifications of the high-resolution gain correction map $\mathrm{g_{LAB}}(\vec{\mathbf{x}})$ originally derived during ground calibration:[5] a corrective factor γ is determined at each pointing, a smooth surface is fit to these corrective factors, and the gain correction map at that epoch is derived as

$$\mathrm{g}(\vec{\mathbf{x}}|\mathrm{t}) = \mathrm{g_{LAB}}(\vec{\mathbf{x}}) \times \gamma(\vec{\mathbf{x}}|\mathrm{t}) \tag{2}$$

This procedure preserves the high spatial-frequency information present in the lab calibration data, while accounting for the gross changes that have occurred in the gain since launch.

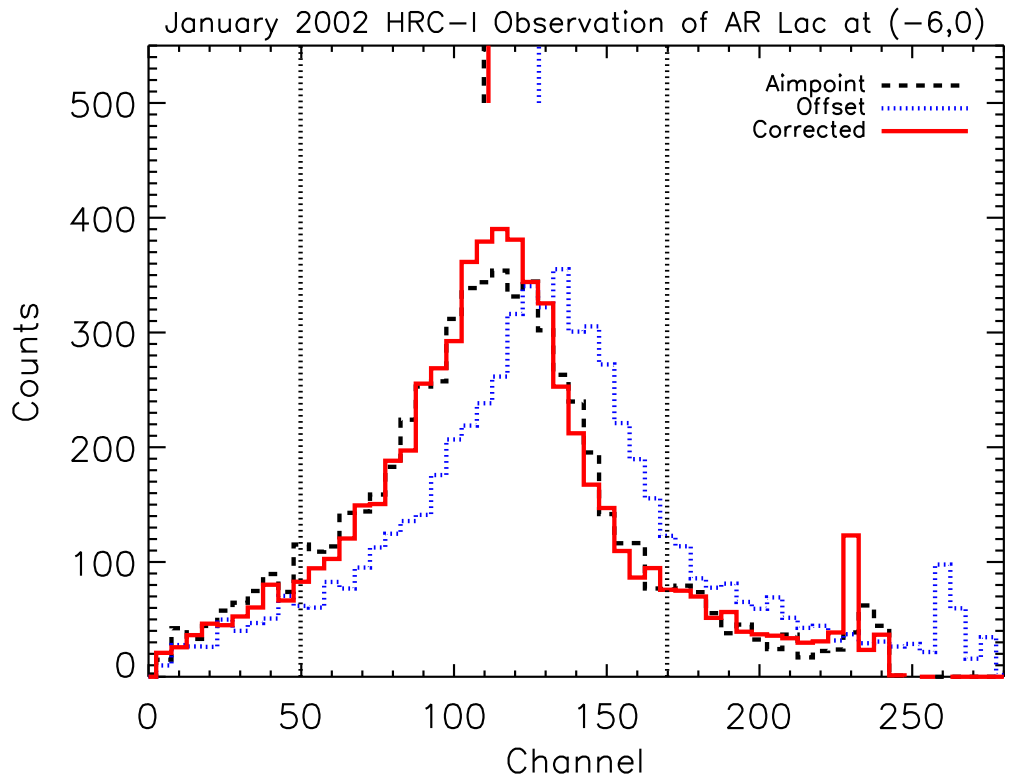

Figure 3. Matching the offset spectrum with the aimpoint spectrum to derive a corrective factor for the gain map. The figure shows the case of the spectrum from the January 2002 AR Lac observation at an offset pointing of $(-6', 0')$ (dotted blue histogram) compared to the aimpoint spectrum (dashed black histogram). The spectra are of the quantities $g_{LAB}(\vec{x}) \cdot \mathrm{PHA}(\vec{x}, t)$. A correction factor is determined by imposing a multiplicative corrective factor on the gain until the two spectra match each other (solid red histogram; $\mathrm{PHA}_{LAB}(\vec{x}, t)$, $\gamma = 0.875$). The matching is done over PHA values bracketed by the vertical dotted lines. Also shown on the plots as short vertical bars at the top are the locations of the medians of the spectra, in the same style as the corresponding histograms. Note that the off-axis spectra have been normalized to the same number of counts as the aimpoint spectrum.

The corrective factors γ are computed by a direct comparison of the spectra at different pointings to the aimpoint spectrum. First, the PHA($\vec{x}$,t) are randomized within each bin to avoid aliasing effects, and then the putative spatially gain-corrected PHAs are computed as

$$\mathrm{PHA}_{LAB}(\vec{x}, t) = \mathrm{PHA}(\vec{x}, t) \cdot g_{LAB}(\vec{x}) \,. \tag{3}$$

These modified PHAs are binned into spectra $f(PHA)$, and the best-fit value of γ that results in the best match between $f(\gamma \cdot \mathrm{PHA}_{LAB}(\vec{x}, t))$ and $f(\mathrm{PHA}_{LAB}(\mathrm{aim}, t))$ is determined via a grid-search algorithm ($\gamma = 0.7 : 1.2 : 0.004$) that minimizes the χ^2 value between the two functions. We limited the comparison to values $median(\mathrm{PHA}_{LAB}(0, t)) \pm 60$ to minimize the effects of any lingering background counts and outliers. We determine errors by Monte Carlo bootstrapping, generating different realizations of the counts in each bin as Poisson deviations for both spectra 1000 times and recomputing the best-fit in each case. An example of this process is shown in Figure 3.

For each epoch, the corrective factors γ at each of the 21 pointings are calculated[§] as described above. Next, we use them to interpolate a minimum curvature surface at all locations over the detector to obtain the corrective surface $\gamma(\vec{x}|t)$. This is multiplied by the high-resolution gain map $g_{LAB}(\vec{x})$ to obtain the gain correction map $g(\vec{x}|t)$ for the epoch (Equation 2).

[§]By definition, $\gamma = 1$ at the aimpoint.

We test the gain correction maps by independently applying $g(\vec{\mathbf{x}}|t)$ to the $PHA(\vec{\mathbf{x}},t)$ values and comparing the *median*$(PHA(\vec{\mathbf{0}}|t))$ for all the datasets. The results are shown in Figure 4. As expected, the medians for each epoch are uniform, i.e., the gain correction has removed the spatial dependence in the $PHA(\vec{\mathbf{x}},t)$. Note that these maps are intermediate products, and are *not* distributed within the calibration database.

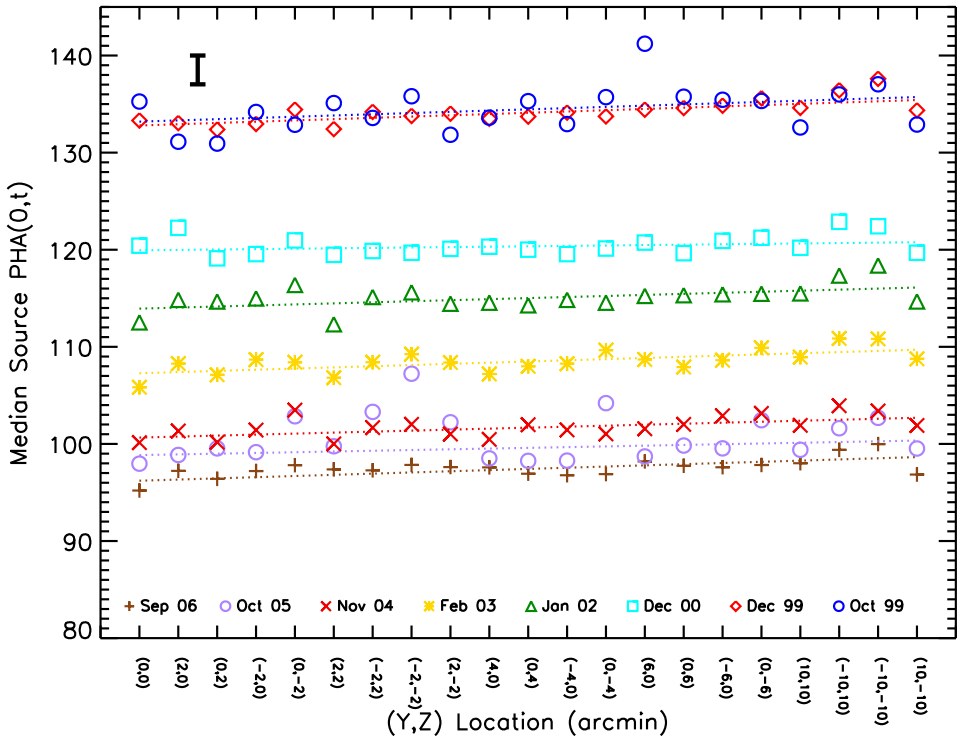

Figure 4. Median "flat-fielded" PHA values as a function of location on the detector, for all the AR Lac observations. The dashed lines show the best-fit line to each set. Since the spatial correction is relative to the aimpoint at each epoch, we expect these lines to be horizontal; the slopes of these lines are statistically indistinguishable from 0. The thick vertical line in the upper left corner shows the typical $\pm 1\sigma$ error on the medians. (Note that this is larger than the error on the raw PHA due to the uncertainty in γ, typically $\sim 1\%$.)

5.2 Temporal Correction

Having made correction maps for the spatial non-uniformity of the detector response, our next task is to correct for the time-dependence in the gain decline. As described above, we seek to calculate the correction as a function of time only, and then correct the gain correction maps from each epoch ($g(\vec{\mathbf{x}}|t)$, see §5.1) by multiplying with this factor (see Equation 1). That is, we want to find a temporal correction factor (TC) such that

$$PI = \mathrm{PHA}(\vec{\mathbf{0}}|t) \times TC(t)\,, \tag{4}$$

where t is the time since October 1999, and PI is the spatially and temporally invariant PHA (equivalent to the flat-field PHA observed at the aimpoint in October 1999), and $PHA(\vec{\mathbf{0}}|t)$ are the "flat-field" PHA values.

In order to determine $TC(t)$ we proceed as follows. We first fit lines to the medians of the spatially corrected PHA spectra,

$$median(\mathrm{PHA}(\vec{\mathbf{0}}|t)) = m \cdot t + b\,, \tag{5}$$

for each observation location separately. This results in a set of 21 slopes m and intercepts b. We exclude the October and December 1999 data from all the fits (not just for the aimpoint data) because the sharp non-linear drop in gain between December 1999 and December 2000 (see Figure 1) is now incorporated in $PHA(\vec{\mathbf{0}}|t)$

(compare Figures 1 and 4). The best-fit slopes are shown in Figure 5, and confirm[¶] that the temporal dependence can be modeled as being uniform across the detector, with

$$\bar{m} = mean(m) = -0.318 \pm 0.014 \text{ channel month}^{-1}. \quad (6)$$

The average y-intercept

$$\bar{b} = mean(b) = 123.0 \pm 1.6 \text{ channel}, \quad (7)$$

corresponds to the expected value of the median PI during the October 1999 observation if that dataset had followed the linear trend established by the succeeding observations. The observed difference,

$$\Delta \equiv median(PI|_{\text{Oct99}}) - \bar{b} = 12.29. \quad (8)$$

Including Equations 5 and 8 in Equation 4, we get the form for the time correction for the PHA,

$$TC(\text{t}) = \begin{cases} \frac{\bar{b}+\Delta}{\bar{m}\cdot t+\bar{b}} & t > 0 \\ 1 & t = 0 \end{cases} \quad (9)$$

For the gain correction maps from AO2 onwards, the fact that we are implementing the time correction as a step function (with t corresponding to the observation date of the AR Lac set) is not problematic, since the gain change is happening slowly. However, for AO1, when the gain was changing rapidly, a single time correction is not sufficient for the whole time period. This problem is illustrated in §6.1.

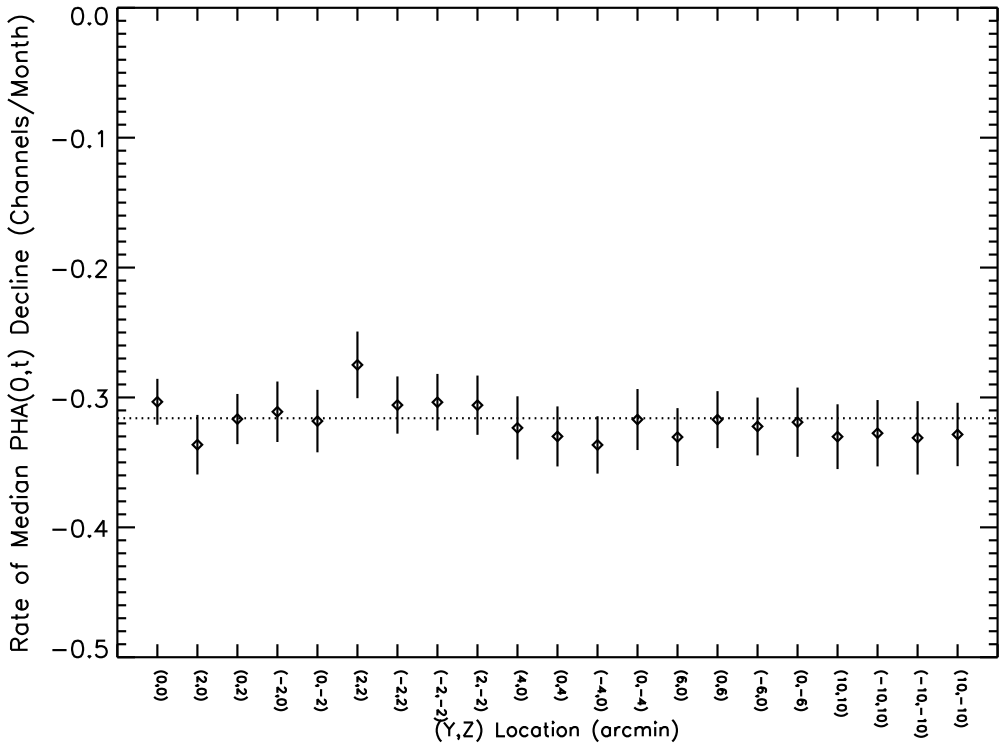

Figure 5. Best-fit slopes (with 1σ errors) from linear fits to the median spatially-corrected PHAs (excluding Oct and Dec 99) for each observation location. The horizontal dashed line shows the average rate of decline.

For each epoch of observation t, we obtain the corresponding correction factor $TC(\text{t})$, and multiply the previously derived "flat-fielded" gain maps $g(\vec{x}|\text{t})$ to obtain the gain correction map at each epoch. These maps, one for each of the 7 epochs, are the final product of our analysis (Figure 6). They were implemented in the CALDB with release version 3.4.

[¶]Note that Figure 5 differs from Figure 2 in that the medians of the spatially-corrected PHAs are used rather than the medians of the raw PHAs.

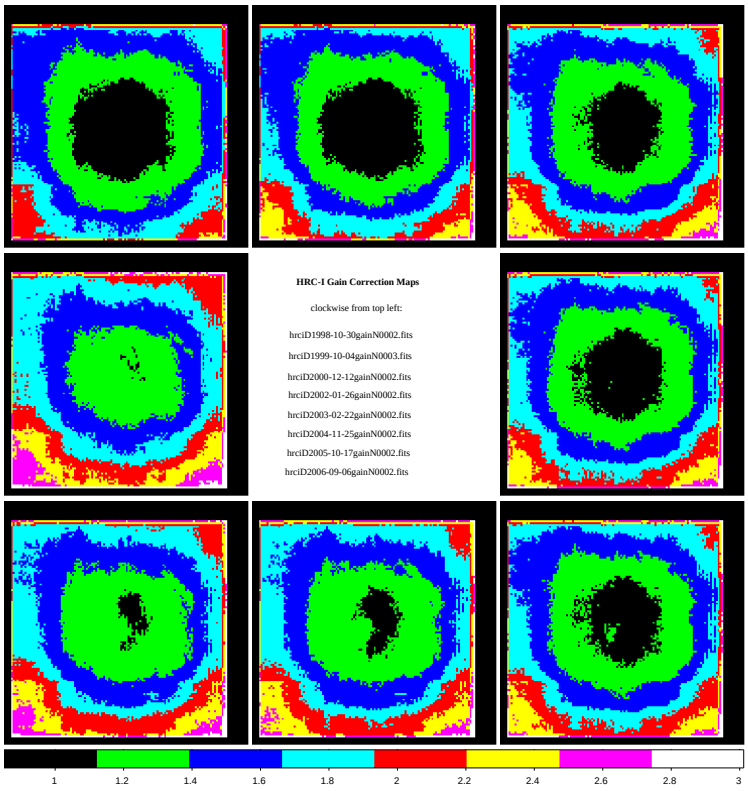

Figure 6. Gain correction maps, shown on a linear display scale from 0.85-3. The pre-flight gain correction map $g_{\rm LAB}(\vec{\mathbf{x}})$ is in the upper left corner. The subsequent time-dependent maps, with updates based on AR Lac observations, are shown chronologically in clockwise order.

6. DISCUSSION

6.1 Testing the New Maps

To test the new gain-correction maps, we again return to the raw source and background PHA values extracted from the AR Lac Level 1 event lists (§3). We convert the values to PI using the appropriate map, then find the median background subtracted PI. Figure 7 shows the median PI values for each AO as a function of observation location on the detector. Comparing this figure to raw PHAs versus location (Figure 1) and spatially-corrected PHAs versus location (Figure 4), we can see that the new gain correction maps have performed their task: the spatial and temporal dependencies from pulse-height values have been removed.

For a more independent test of the newly-created gain correction maps, we use HRC-I calibration observations of the white dwarf HZ 43 (RA = 13:16:21.853, Dec = +29:05:55.44 (FK5)) and the supernova remnant G21.5-0.9 (RA = 18:33:32, Dec = -10:33.6 (FK5)). These sources serve as useful test sets because they have been observed regularly since launch, and therefore we have the opportunity to use each of the time-dependent gain correction maps. In addition, since these sources emit at different energies than AR Lac (peak at ∼60 eV for HZ 43 and reaching till ∼2 keV for G21.5, compared to ∼3 keV for AR Lac), they provide an opportunity to explore the energy dependence (if any) of the gain correction maps. (Unfortunately, we do not have multiple observations of either of these sources at offset locations, so we cannot use them to test the performance of the new gain correction maps away from the aimpoint.) We apply the appropriate gain-correction map to the source PHA values, converting to PI, then find the median PI value. The observations taken before December 1999, to which we apply the AO1 map, are overcorrected by approximately 10%, due to the insufficiency of a constant time-correction over this period when the gain was changing rapidly. However, the gain correction maps work quite well for the later data. Fitting the post-December 1999 medians as a function of time gives best fit slopes that

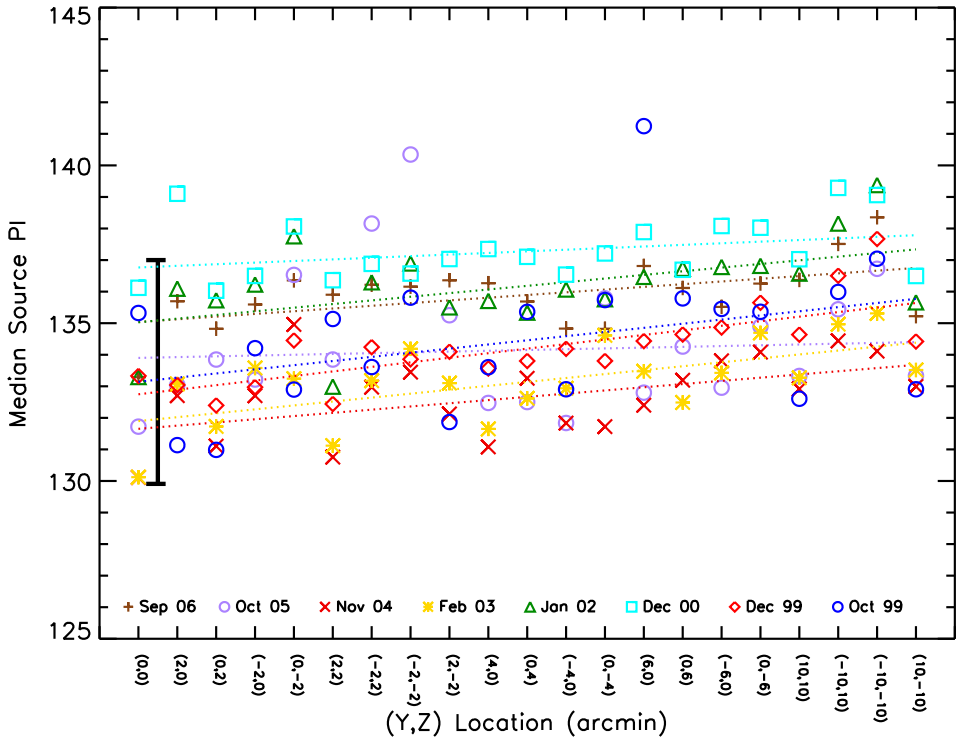

Figure 7. Median PI values of AR Lac datasets, calculated using the new gain correction maps, as a function of observation location. The bold vertical bar at the left shows the typical $\pm 1\sigma$ error on the medians, which includes the uncertainties in γ and $TC(\mathrm{t})$, typically $\sim 1\%$ and $\sim 2\%$, respectively. The horizontal dotted lines show linear least-squares fits to data for each AO. The slopes are statistically consistent with zero.

are nearly flat (-0.03 ± 0.04 channels per month for HZ 43 and -0.02 ± 0.08 channels per month for G21.5-0.9) indicating that the gain maps have successfully removed the temporal dependence from the PIs.

6.2 Correlation with Dosage

Cumulative dosage maps for the HRC-I are made each month by summing the Level 1 event lists from observations taken that month to the previous month's dosage map. Because the dosage is expected to affect the gain, we consider the possibility that the gain decline at a given location on the detector (as evidenced by the median PHA drop at that point between Oct 99 and a later date) is correlated with the cumulative dosage that location received. Surprisingly, we find that there is not a strong correlation between gain decline and dosage. For example, the points at $14.14'$ off-axis and the points at $4'$ off-axis have received roughly the same amount of dosage since launch, 10^6 counts, but the median PHAs in the first group have declined by about 12 channels, while the median PHAs in the second group have declined by ~ 30 channels. It seems that areas of the detector where the gain is higher (i.e. near the center) experience relatively larger gain drops than the low gain areas far offaxis. This pattern is seen in the HRC-S as well (Posson-Brown et al, in preparation).

However, the dosage maps include only the telemetered events. The total event rate for HRC-I, prior to on-board vetoing, averages to $\sim$250 cts s^{-1}.[8] This rate is dominated by the particle background which, unlike X-ray events, is not concentrated at the aimpoint. Rather, it is very roughly uniform (to $\pm$ 20%) across the detector.[9] This may explain why the gain decline does not correspond as strongly to telemetered dosage as we might expect.

Based on lab measurements, we expect a gain decline of 10% after a particle fluence of approximately $6 \times 10^7 \mathrm{cm}^{-2}$,[3] or 6×10^9 events (the detector area is $\sim 100 \mathrm{cm}^2$). Estimating the average HRC total event

rate to be $250s^{-18}$ and the total ONTIME of the detector thus far as 4.4×10^6 s gives $\sim 1.1 \times 10^9$ counts,[‖] well under the lab limit. Note though that this number is an underestimate, by a factor 2-3x, because the HRC high-voltage (HV) is occasionally on even outside of the observation ONTIME (M. Juda & T. Isobe, private communication). Nevertheless, the fact that we have already seen gain declines $\gg 10\%$ suggests that the correlation of gain decline with dosage may be more complex than expected. We continue to investigate this connection.

7. SUMMARY

By monitoring the background-subtracted median source PHAs from AR Lac observations taken regularly at 21 locations on the HRC-I, we have seen that the detector gain has declined since the beginning of the mission. At the aimpoint, the gain has dropped rapidly between October 1999 and December 2000, with the median PHA falling by 17 channels. Since then, the gain has declined more slowly and linearly, with the median PHA changing at an average rate of ~ -3.7 channels per year at all monitored locations (excluding the four outermost pointings at $14.41'$ off-axis, where the rate is ~ -1.4 channels per year).

To address this gain decline, we have made a series of time-dependent gain correction maps by using the AR Lac observations to update the pre-flight gain map. We first derive a set of spatial corrections for each year by comparing the offset to aimpoint AR Lac profiles for that year. We then derive a time correction as a step function of time, based on linear fits to the spatially-corrected PHA values. The final gain correction maps remove the spatial and temporal dependencies from the resulting PI spectra, allowing for direct comparison regardless of observation date or location on the detector. (Note, however, that caution should be used when dealing with observations taken between October 1999 and December 2000.)

Surprisingly, the gain decline is not strongly correlated with the telemetered HRC-I dosage. Even estimating the total HRC-I dosage to date does not explain the magnitude of gain decline seen. We conclude that correlation of the gain decline with dosage is not as expected, and continue to explore this relationship. We also continue to monitor the gain with regular observations of AR Lac, and investigate the possibility of a functional gain correction.

ACKNOWLEDGMENTS

The authors thank Michael Juda, Frank Primini, and Brad Wargelin for helpful discussion. This work was supported by NASA contract NAS8-39073 to the Chandra X-ray Center.

REFERENCES

1. Posson-Brown, J., Donnelly, H., 2003, *Evolution of PHA Response in the HRC* (CXC Memo, available at `http://asc.harvard.edu/cal/Hrc/gain.html`)
2. Pease, D., Drake, J., 2003, *Monitoring the HRC-S Gain with the LETG/HRC-S* (CXC Memo, available at `http://asc.harvard.edu/cal/Hrc/gain.html`)
3. Juda, M., 2001, *HRC Rates and High Solar Activity* (CXC Memo, available at `http://hea-www.harvard.edu/ juda/memos/hrc_bkg/high_solar.html`)
4. Kashyap, V., Posson-Brown, J., 2005, *HRC RMF* (CXC Memo, available at `http://asc.harvard.edu/cal/Hrc/RMF/`)
5. Wilton, C., Posson-Brown, J., Juda, M., Kashyap, V., 2005 *The HRC-I Gain Map*, Poster #9.10 at the 2005 Chandra Calibration Workshop, available at `http://asc.harvard.edu/ccw/proceedings/05_proc/presentations/wilton/`
6. Kashyap, V. & Drake, J.J., 2000, BASI 28, 475
7. van Dyk, D.A., Connors, A., Kashyap, V.L., & Siemiginowska, A., 2001, ApJ, 548, 224
8. Weisskopf, M.C., Brinkman, B., Canizares, C., Garmire, G., Murray, S., Van Speybroeck, L.P., 2002, PASP, 114, 1
9. Juda, M. et al, 2002, Proceedings of SPIE, 4851, 112

[‖]This is larger than the numbers derived from the dosage maps, because the latter includes only events in the Level 1 event lists, whereas the former is the total event rate prior to on-board vetoing.