Monte Carlo Processes for Including Chandra Instrument Response Uncertainties in Parameter Estimation Studies

Jeremy J. Drake[a], Peter Ratzlaff[a], Vinay Kashyap[a], Richard Edgar[a], Rima Izem[b], Diab Jerius[a], Aneta Siemiginowska[a] and Alexey Vikhlinin[a]

[a]Smithsonian Astrophysical Observatory, 60 Garden Street, Cambridge MA 02138, USA;
[b]Department of Statistics, Harvard University, 1 Oxford Street Cambridge, MA 02138, USA

ABSTRACT

Instrument response uncertainties are almost universally ignored in current astrophysical X-ray data analyses. Yet modern X-ray observatories, such as Chandra and XMM-Newton, frequently acquire data for which photon counting statistics are not the dominant source of error. Including allowance for performance uncertainties is, however, technically challenging in terms of both understanding and specifying the uncertainties themselves, and in employing them in data analysis. Here we describe Monte Carlo methods developed to include instrument performance uncertainties in typical model parameter estimation studies. These methods are used to estimate the limiting accuracy of Chandra for understanding typical X-ray source model parameters. The present study indicates that, for ACIS-S3 observations, the limiting accuracy is reached for $\sim 10^4$ counts.

Keywords: Chandra X-ray Observatory, X-ray Optics, Calibration, Uncertainties, Monte Carlo Methods

1. INTRODUCTION

The development of X-ray astronomy was pioneered primarily by rocket and satellite missions carrying gas-filled proportional counters. The response of this type of detector to X-ray photons is characterised by an electronic pulse with a magnitude that depends on the energy of the ionising photon and on stochastic aspects of the interaction of this photon with the detector and its gas. The signal resulting from a number of X-ray events of a given energy is a fairly broad distribution of "pulse heights" with a peak that is approximately proportional to the photon energy. The instrument response of such detectors provides a low degree of spectral energy resolution of $E/\Delta E \sim$"a few". For investigating the nature of the spectrum of an astrophysical source, this type of instrument response often precludes direct measurement of source parameters or spectroscopic diagnostics because signals from different energies are strongly overlapping and cannot be easily separated. Instead, the nature of this type of data prompted the use of parameterised source models which could be compared with observation through convolution with the instrument response and subsequent application of a "goodness of fit" (typically χ^2) statistic. Best-fit model parameters are then found through a minimisation scheme such as Levenberg-Marquardt. This remains the approach used today for the great majority of analyses of X-ray observations from missions past and present, using "fitting engines" such as XSPEC[1] and Sherpa.[2]

Incorporating independent uncertainties in both observational data and the parameterised models used for comparison is, at least in principle, quite straightforward for least-squares type minimisation. Though more problematic in the very low count limit that occurs frequently in X-ray observations of celestial sources, photon counting uncertainties are routinely incorporated in all modern astrophysical X-ray parameter estimation analyses. They are also the only uncertainties considered: instrument response uncertainties and uncertainties in the models being fit to the data are almost universally ignored. Such dramatic simplification of the problem is prompted mostly by analytical and computational expedient: uncertainties in the response of a typical X-ray telescope and detector system are not independent and easily incorporated but are *correlated* in complicated ways. Likewise, uncertainties in source models, such as those used to model hot, optically-thin astrophysical plasmas, might be subject to a very complex set of uncertainties involving an extensive assemblage of atomic data, all correlated through the plasma ionization balance, atomic level populations and element abundance. There

Further author information: Send correspondence to jdrake@cfa.harvard.edu

is no standard set of procedures for incorporating complicated correlated systematic uncertainties in non-linear parameter estimation: the approaches used for treating independent errors simply do not apply.

The complexity of the correlated uncertainties problem in X-ray astronomy is a daunting prospect for analytical solution and suggests instead the use of Monte Carlo techniques. Indeed, Monte Carlo methods were employed to derive the original calibration requirements for the *Chandra* mirrors.[3] Here we present methods to treat instrument response uncertainties in a reasonably realistic way. We apply these to the *Chandra* X-ray Observatory Advanced CCD Imaging Spectrometer (ACIS) and to assess the limiting accuracy of this system for typical parameter estimation analysis. We also discuss how this method can be adopted for public release

2. CORRELATED UNCERTAINTIES IN INSTRUMENT CALIBRATION

Comprehensive fully-empirical ground-based calibration of X-ray instrumentation to a desired accuracy prior to flight is often not realistically achievable: too many subassembly components might require sampling at more energies than can be reached within reasonable synchrotron or laboratory time. Instead, the response of flight instruments are generally characterised by analytical or semi-empirical performance models underpinned by a more limited set of fiducials from laboratory measurements. The uncertainties in the resulting response curves are prime examples of uncertainties that are correlated: response models, for example, provide cross-talk between individual laboratory measurements. Even the latter might be correlated to some extent through, for example, common calibration uncertainties of the laboratory instrumentation.

The uncertainties, $\sigma(E)$, in a vector such as a telescope effective area as as function of energy, $A(E)$, are then more properly described not by scalar quanties at different energies, but by a matrix of related uncertainties. For example,

$$\sigma(E) = \begin{pmatrix} \sigma_{E_1,E_1} & \sigma_{E_2,E_1} & \cdots & \sigma_{E(n-1),E_1} & \sigma_{E_n,E_1} \\ \sigma_{E_1,E_2} & \sigma_{E_2,E_2} & & & \\ \vdots & & \ddots & & \vdots \\ \sigma_{E_1,E(n-1)} & & & \sigma_{E(n-1),E(n-1)} & \sigma_{E_n,E(n-1)} \\ \sigma_{E_1,E_n} & & \cdots & \sigma_{E(n-1),E_n} & \sigma_{E_n,E_n} \end{pmatrix} \tag{1}$$

where

$$\sigma_{E_i,E_j} = \sqrt{\sigma^2_{E_i,E_i} + \sigma^2_{E_j,E_j} - 2\rho_{E_i,E_j}\sigma_{E_i,E_i}\sigma_{E_j,E_j}} \tag{2}$$

In the example error matrix in Eqn. 1, diagonal elements σ_{E_i,E_i} represent the absolute uncertainties at energies E_i. However, these uncertainties are not independent because physically the instrument does not behave discontinuously between adjacent energy bins, except perhaps at the absorption edges of its constituent materials. The off-diagonal elements then represent the *relative uncertainty* in the response between energies E_i and E_j. Physically, it is to be expected that σ_{E_i,E_j} increases with increasing $j - i$. The *correlations* between energy bins E_i and E_j are represented by a covariance matrix of values ρ_{E_i,E_j}; generally, the correlation ρ_{E_i,E_j} *decreases* with increasing $j - i$.

While an error matrix such as the above might be easy to write down in theory, there are as yet no simple "goodness-of-fit" formulae or well-accepted methods to employ such a specification in traditional parameter estimation studies. In practice, populating such an error matrix with physically meaningful quantities is also far from trivial and requires a comprehensive specification of the instrument and its uncertainties; moreover, for such a matrix to be of use, self-consistency among all the possible combinations of σ_{E_i,E_j} must be enforced.

3. MONTE CARLO APPROACH FOR CHANDRA ACIS-S

Monte Carlo "brute force" methods offer an alternative to the manifold complexity of developing a suitable error matrix and associated statistical methods to use it. In a Monte Carlo approach, the calibration of an instrument can be sampled or altered within existing uncertainty bounds and the effects of such perturbations on the instrument response assessed through the distributions of parameters obtained from least-squares fits employing the different response realisations. Such an approach would have been computationally prohibitive some years ago, when simply performing a single fit required significant computation time. With the processing power available today, thousands of fits can be undertaken in relatively short order.

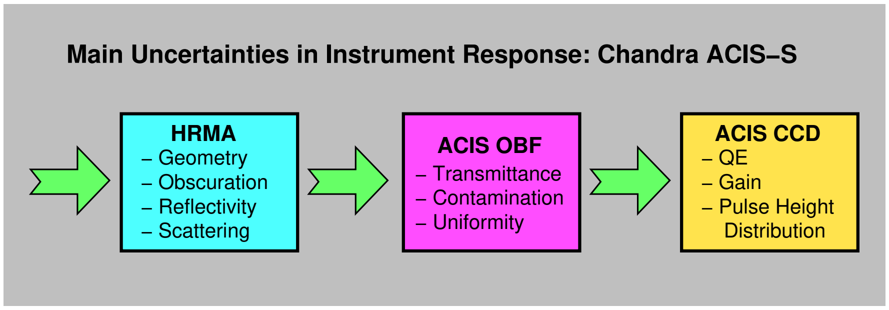

Figure 1. Illustration of the main contributions to the instrument response uncertainties in the *Chandra* ACIS-S photon path.

The proof-of-concept Monte Carlo method we describe here has been applied to the on-axis back-illuminated S3 chip of the *Chandra* ACIS-S detector; this currently being the most used configuration of the observatory.

3.1. Sources of Uncertainty in the Photon Response

The *Chandra* ACIS-S observing configuration includes along the photon path the High Resolution Mirror Assembly (HRMA) and the ACIS detector. The latter is comprised of an Optical Blocking Filter (OBF) and the CCDs themselves. The main uncertainties in the response of these components are shown schematically in Figure 1.

Our Monte Carlo approach treats each subassembly separately using two different methods: (a) use of subassembly model responses computed for different input parameters sampled within an uncertainty range (HRMA contamination overlayer, ACIS QE, ACIS gain and pulse height distribution as a function of energy); and (b) a more arbitrary *Perturbation Function* that is a smooth curve deviating from unity within a given energy range by an amount that lies within the uncertainty of the subassembly response. These are described in more detail below.

3.2. The Perturbation Vector

A flexible means (literally) of introducing variations into the subassembly responses is through a series of *perturbation vectors*. Each subassembly response is quite naturally divided up by the absorption edges of its constituent materials. Prominent examples are the C K edges in the OBF, the Si K edge in the ACIS QE, and the Ir M edge in the HRMA. At energies with significant edge structure, the subassembly responses are generally less certain than in regions with smooth variations as a function of energy; the "jump" discontinuity over an edge also introduces some latitude in relative uncertainty between the responses above and below the edge. We therefore divide each subassembly response into sections defined by prominent edges or discontinuities. Within each region we construct a perturbation vector, $\mu(E)$, defined by parameters that can be chosen at random such that $\mu(E)$ lies within a pre-defined uncertainty envelope. Vectors for neighbouring energy ranges are tied together by an "edge constraint" that specifies the maximum deviation permitted in the magnitude of the jump discontinuity. The perturbation vector is illustrated schematically in Figure 2a.

The perturbation vectors for the different energy ranges of a subassembly are concatenated to produce a single vector, $\mu(E)$. In the present method, we construct separate vectors for the HRMA effective area, the OBF transmittance, the OBF contamination layer, and the ACIS QE. In the cases of the HRMA and ACIS QE, the perturbation vectors supplement the variations induced by sampling the predictions of computer models (see Sect. 4 below). For any given set of vectors, the modified effective area of the system, $A'(E)$, is defined by the products of the vectors and nominal effective area, $A(E)$:

$$A'(E) = \mu_{HRMA}(E)\mu_{contam}(E)\mu_{OBF}(E)\mu_{QE}(E)A(E).$$

An arbitrary number of effective areas, $A'(E)$, all differing from the nominal area by different perturbations within the specified uncertainty limits, can be generated either on-the-fly or to build up a library of effective area files.

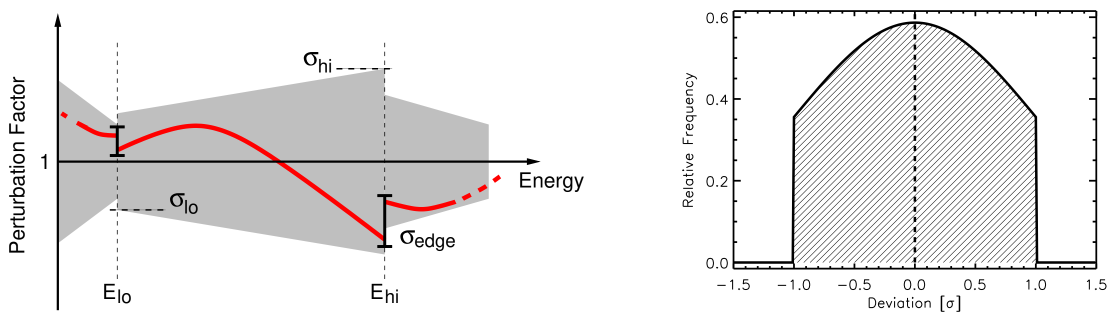

Figure 2. Left: Illustration of a segment of a *perturbation vector* used to apply deviations from a nominal subassembly response within a given energy range. Within each energy range, E_{lo}–E_{hi}, a smooth curve is generated that is constrained to lie within the grey shaded region defined by the uncertainties σ_{lo} and σ_{hi}, and also to join up with neighbouring segments within the edge constraints σ_{edge}. **Right:** The truncated normal distribution used to represent the distribution of calibration uncertainties that are used in the perturbation vector and computer model realisations.

3.3. Distribution of Uncertainties

In order to implement any Monte Carlo sampling of the instrument response within specified uncertainties, some knowledge of the distribution of the uncertainties is of course essential. Uncertainties are often specified in terms of "σ", with the underlying implicit assumption that they are distributed according to the normal distribution where σ takes its usual meaning in which there is a 68% probability that the true value lies within $\pm 1\sigma$ of the estimated one. The normal distribution naturally describes measurements of summary statistics of any quantity, especially when a large number of measurements are made. However, it fails to properly describe anything that may have a skewed, or otherwise systematically different distribution. It inevitably has extended wings, with, for example, about 32% probability that the true value deviates from the estimate by more than 1σ.

In the context of instrument calibration, the quantities being measured are not completely unknown and these extended wings often do not correspond to researchers' intuition as to the true limits of the uncertainties of a measurement. Determining the actual distributions of the real calibration uncertainties and their moments is in practice extremely difficult. Such measurements are often described with phrases like "the uncertainties are no larger than 10%". Rather than admitting a 1 in 3 probability that the true value lies outside a 10% error budget, in such a description the researcher is confident (rightly or wrongly!) that, based perhaps on experience or knowledge of the system being calibrated that is difficult to specify in rigorous statistical detail, the likelihood of the true error being larger than 10% is extremely small or negligible. Such a distribution is probably not well-described by either a uniform distribution between the limits ($\pm 10\%$) or by a normal distribution, but by something with much less extended wings, though perhaps still peaked at the center. For this study, we have adopted a truncated normal distribution to represent the distribution of uncertainties. In a Bayesian context, these are informative prior distributions on the parameters that describe the calibration, and are simply the product of a Gaussian with variance σ^2 and a rectangular Step function with unit density between $\pm\sigma$. These are then specified by σ in the traditional fashion, except that the normal distribution is truncated at 1σ. Our 1σ errors here then represent the sharp cut-off in probability that is intended to reflect the "gut feeling" of experience and prior knowledge that is otherwise very difficult to include. This truncated normal distribution is illustrated in Figure 2b.

4. ASSESSMENT OF SUBASSEMBLY RESPONSE UNCERTAINTIES

It is often (but not always) possible to design laboratory tests such that Poisson uncertainties from photon counting devices are negligible. Exceptions might include "flat field" measurements that can require a vast number of counts to cover a large filter or to populate a large detector with fine spatial resolution. When Poisson uncertainties are very small—in the case of *Chandra* calibration "very small" can be taken to mean less than 1%—the systematics of the measurement system dominate. The uncertainties in the calibration of nearly all aspects of the *Chandra* telescope and detector system are dominated by these systematics.

Chandra Calibration measurements were performed in different laboratories and synchroton facilities, and at the NASA X-ray Calibration Facility at the Marshall Space Flight Center.[4] It is well beyond the scope of this article to describe these calibration activities and results; the reader is instead referred elsewhere for detailed descriptions relevant to ACIS-S.[4–10]

Table 1. Adopted "1σ" fractional uncertainties and edge constraints (see text and Figure 2a) for the components of the HRMA and ACIS subassemblies

Component	E_{lo} Edge	E_{lo} [keV]	E_{hi} [keV]	σ_{lo}	σ_{hi}	σ_{edge}
HRMA Effective Area	...	0.05	2.156	0.07	0.06	0.02
	Ir M_4 $3d_{3/2}$	2.156	3.183	0.05	0.05	0.01
	Ir M_1 $3s$	3.183	6.400	0.05	0.05	0.005
	...	6.400	12.0	0.05	0.20	...
OBF Contamination	...	0.05	0.284	0.50	0.08	0.04
	C K $1s$	0.284	0.410	0.08	0.05	0.03
	N K $1s$	0.410	0.543	0.05	0.03	0.01
	O K $1s$	0.543	0.697	0.03	0.03	0.01
	F K $1s$	0.697	12.0	0.01	0.005	...
OBF Transmittance	...	0.05	0.284	0.15	0.07	0.03
	C K $1s$	0.284	0.543	0.05	0.04	0.02
	O K $1s$	0.543	1.560	0.03	0.03	0.02
	Al K $1s$	1.560	12.0	0.03	0.01	...
ACIS-S3 QE	...	0.05	0.543	0.10	0.05	0.03
	O K $1s$	0.543	1.839	0.03	0.03	0.03
	Si K $1s$	1.839	12.0	0.03	0.03	...

The adopted "1σ" uncertainties in the different energy ranges for the different subassemblies used in our truncated normal distribution Monte Carlo perturbation vectors, together with edge constraints (described above in Sect. 3.2) are listed in Table 1. These numbers were arrived at through study of the calibration data and reports and through the hands-on experience of the authors with *Chandra* data and observations. In cases of doubt (ie most of the numbers in Table 1), we admit to a mild conservative tendency, allowing perhaps slightly larger, rather than smaller, uncertainties than information sources might have implied.

In providing the numbers in Table 1, we emphasise that the goal of the present current study is not to provide a final, rigorous assessment of the uncertainties of each of the *Chandra* subassemblies, but to make the first attempt at reasonable estimates of these for use in assessing the limiting accuracy of *Chandra*, and for developing the methods to do so.

Some very brief explanatory notes on these individual components are included below.

4.1. High Resolution Mirror Assembly

The HRMA area is defined by meticulous measurements of the reflectivity of witness flats,[9] XRCF measurements of the mirror throughput[11] and a sophisticated ray trace computer model.[12]

The XRCF flow proportional counters (FPCs) and solid state silicon detectors (SSDs) were independently calibrated at synchrotron facilities to accuracies now believe to be of order 3%, including allowance for all systematic terms; this number then probably represents about the best absolute accuracy achievable through XRCF measurements at that time. Additional small sources of uncertainty include, e.g., incomplete modelling of the bowing of the FPC windows and the obscuration of their mesh supports,[13] and small non-uniformities and time-variability in the X-ray beam.

Uncertainties in calibration at XRCF were effected by systematic differences between FPC and SSD measurements which could not be reconciled within the Poisson errors of measurement. These differences amount to between 5 and 10% or so, with the true values believed to lie between these extremes.[11] The uncertainty envelope encompassing these errors is then about $\pm5\%$. Slightly larger uncertainties rising to 7% have been included at the lowest energies.

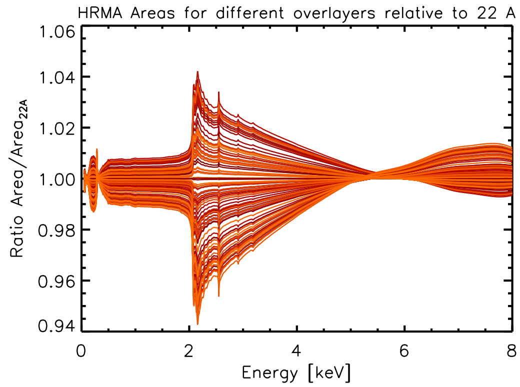

Figure 3. Illustration of the relative change in the HRMA effective area caused by different hydrocarbon contamination layers. The range shown corresponds to the nominal adopted 22 ± 6 Å layer thickness.

Problems with matching in-flight observations of strong continuum sources in the region of the Ir M_4 edge at 2.156 keV using the *Chandra* transmission grating spectrometers were uncovered after launch.[14] This has been attributed to a thin hydrocarbon contamination layer approximately 20 Å thick that is believed to have built up on the mirrors prior to launch. It modifies the HRMA area in the Ir M edge region by up to 10-15 %. Ray trace models for different overlayer thicknesses have been used to account for the different HRMA areas possible within the range of uncertainty of the layer; these areas were sampled using the curtailed normal distribution and an overlayer thickness and 1σ uncertainty of 22 ± 6 Å. The effect of this uncertainty on the HRMA area is illustrated in Figure 3.

4.2. Optical Blocking Filter and CCD Quantum Efficiency

The X-ray transmittance of a filter such as the ACIS OBF can be measured relatively easily at a synchrotron facility by comparing the intensity of a monochromatic beam seen through the filter with that when the filter is removed. The ACIS OBF was calibrated at the National Synchrotron Light Source[15] and is probably the best-calibrated component of the *Chandra* ACIS-S system. Sources of uncertainty are again dominated by systematics, such as out-of-band transmission of the monochromator used for calibration measurements, and filter non-uniformity; the latter is specified as being uniform in transmittance to "better than 2%".[15] We have assessed our 1σ uncertainties at 5% at the C K edge, tapering down to 3% at Al K and 1% at higher energies where the transmittance approaches unity. While edge structure is likely to be less accurately represented, discontinuities between different regions appears well-constrained and we assess a 2% uncertainty at C K and O K edges.

The ACIS CCD QE was calibrated in the laboratory and at XRCF; corrections have also been made prompted by in-flight observations, and in particular concerning the ratio of the QEs of front-illuminated and back-illuminated devices.[14, 16] Recent re-analysis of XRCF data[16] reveals statistical uncertainties in the measurement of the chip-averaged QEs of $\sim$ 1-3%; the final adopted model for the S3 chip agrees with the measurements to within about 5%. QE uniformity maps indicate residual relative uncertainties of 1-2%[17] We have levied 1σ uncertainties of 3% for the absolute QE and its principal edges (O K 0.543 keV and Si K 1.839 keV), with a larger uncertainty of 10% at the lowest energies, tapering to 5% below the O K edge.

In addition to the perturbation vector approach, we also used Monte Carlo samples of the ACIS QE computed using a semi-analytical "slab and stop" model of the ACIS CCD gate structure.[18] Parameters varied were the depletion depth and SiO_2 layer. At the time of writing of the pre-launch ACIS calibration report,[18] the uncertainty in the depletion depth was stated as $\pm 15\%$; we adopted a slightly lower $\pm 13\%$ as our perturbation vector allows for some additional uncertainty. For the SiO_2 layer, we adopted an uncertainty of $\pm 20\%$ from the nominal value. The effects of the uncertainties in depletion depth and SiO_2 layer are illustrated in Figure 4; the former affects primarily the higher energies, while the latter becomes insignificant for energies above 1 keV.

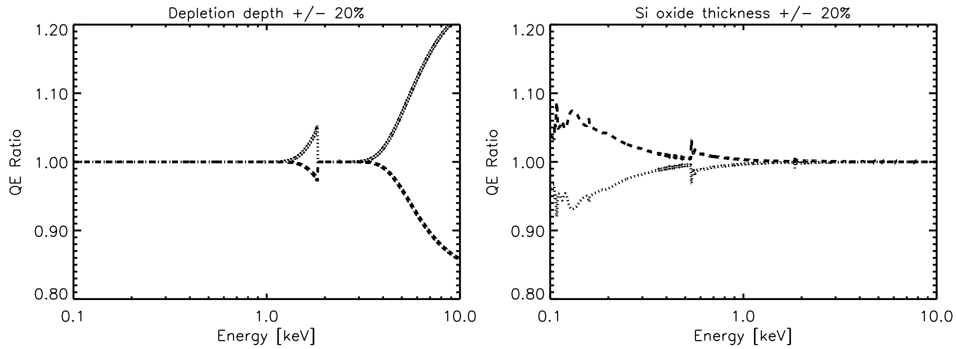

Figure 4. Relative changes in the model ACIS S3 QE caused by adoption of different parameters for the depletion depth (left) and SiO_2 thickness (right). The former shows the effect of a $\pm 20\%$ change; the 1σ range adopted here is ± 13 %, which corresponds to a range of about $\pm 10\%$ in the QE at 10 keV.

4.3. Contamination Layer

There are three aspects of uncertainty in the *Chandra* ACIS-S effective area resulting from the contamination layer that has been building up on the ACIS-S OBF: the contaminant chemical composition, its average thickness at any given time, and its spatial uniformity. The composition of the contamination layer has been assessed to accuracies of a few percent in the optical depths of the most important elements (C, O, F) using observations of strong celestial continuum sources.[10] The spatial uniformity and change in time of the contamination has been measured using both the on-board External Calibration Source and spatially-extended celestial sources.[17, 19]

The uncertainties in the derived effective area using the time-dependent contamination model[19] will also be time-dependent to some extent: uncertainties can be expected to be slightly larger in later observations when optical depths are greater. The nominal ACIS-S effective area upon which our Monte Carlo experiments are based corresponds to an observation date of 2002 September 9—well after the discovery of the contamination problem and typical of the dates of some of the observations used for characterising its composition.[10] Our adopted uncertainties in Table 1 taper from 8% at the C K edge down to 0.5% at high energies, with uncertainties at the C, N, O and F edges of 4, 3, 1,and 1%, respectively.

4.4. CCD Gain and Pulse Height Response

Unlike the quantities discussed so far, the ACIS CCD gain and pulse height response does not enter into the effective area of the instrument. These aspects of the ACIS response are accounted for in the *Chandra Interactive Analysis of Observations* (CIAO) software using a 2-D response matrix file (RMF) that encapsulates the mapping between the incoming photon energy and the resulting detector pulse heights. The gain is a function of position on the CCD and evolves with time as a result of gradual changes in the CCD charge transfer inefficiency (and also because of electronic drift in the I2 chip).[20] The mean secular drift is known to of order 0.2% from ECS observations and observations of extended celestial sources.[20] The gain itself is thought to have uncertainties of 1% near 0.7 keV, 0.5% near 1.5 keV, and of order 0.2% or less at 4 keV and above, with uncertainties in the pulse height FWHM being of order 1%.[21]

We modelled the gain and pulse height uncertainties using the code `calcrmf2`, which is a more general and flexible version of the CIAO program `mkacisrmf`. A Perl wrapper program samples the uncertainties in gain and FWHM and runs `calcrmf2`, creating a large library of RMF files (typically 1000).

5. THE LIMITING ACCURACY FOR PARAMETER ESTIMATION

One of the goals of this study is to determine how accurately *Chandra* can probe the characteristics of different categories of celestial X-ray source. Observations in which photon noise is smaller than uncertianties in the effective area are now routine; uncertainties in parameters obtained from such observations could be grossly underestimated if calibration uncertainties are ignored. Knowledge of the limiting accuracy of the system can

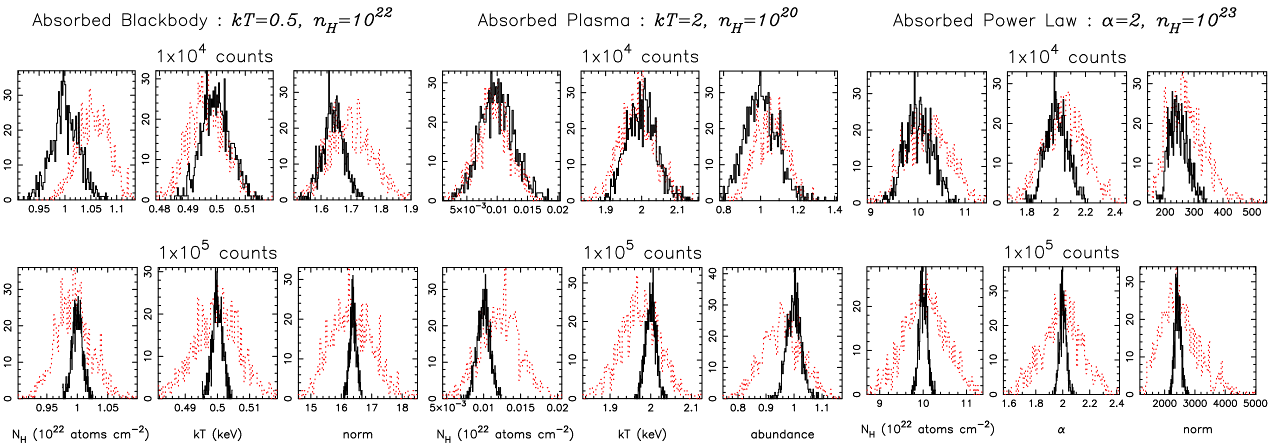

Figure 5. Frequency distributions of best-fit parameters obtained for typical blackbody, thermal plasma and powerlaw models from XSPEC for synthetic data sets containing 10^4 (upper panels) and 10^5 (lower panels). Black histograms are distributions resulting from 1000 Monte Carlo samplings of the synthetic data allowing Poisson noise variations alone. Red histograms are the distributions of parameters resulting from fits to a single synthetic data set using 1000 Monte Carlo-generated effective areas and response matrices.

also help define maximum observation integration times beyond which increased signal strength provides no further insights, and to assess the feasibility of observations that might prove of little scientific value when calibration uncertainties are accounted for.

Our method for assessing the effects of the Monte Carlo sampled effective areas on a given analysis consists of two steps: (i) generation of a synthetic data set using the spectral model of interest and a nominal effective area; (ii) repeated parameter estimation using a different effective area and response matrix each time. A sample of effective areas used is compared with the nominal one in Figure 6. The distributions of the model parameters found from step (ii) can then be compared with those resulting from photon noise alone. For consistency, we assess this by examining the distribution of best-fit parameters obtained for a large sample of synthetic data sets identical except for variations due to Poisson noise.

We carried out computations as described above using version 12.2.0 the XSPEC program[1] driven by a Perl module to control the Monte Carlo aspects and data I/O. Models investigated were: typical absorbed blackbodies with parameters temperature, neutral H column density N_H, and normalisation; optically-thin thermal plasmas with solar composition and parameters temperature, metallicity, N_H, and normalisation; and power laws with parameters slope α, N_H, and normalisation. The computations were repeated for synthetic ACIS-S3 spectra containing 10^4 and 10^5 counts; the former represents the signal in a typical observation of reasonable quality, while the latter probes the case in which photon noise is negligible. Typical distributions of parameters resulting from calibration uncertainties and Poisson noise for 10^4 and 10^5 count synthetic spectra are illustrated in Figure 5. In Figures 6 and 7, we illustrate the mode and highest posterior density $\pm 95\%$ confidence intervals for all the different models examined.

The distribution of best-fit parameters in the upper panels of Figure 5 corresponding to simulations with 10^4 counts show that the retrieved *parameter errors due to calibration uncertainties are at least as large as those due to Poisson noise.* Taken at face value, these results indicate that obtaining significantly more than 10^4 counts will lead to no further gain in accuracy. If we have overestimated the calibration uncertainties, then it is likely that this number should be slightly higher.

Comparison of the distributions for 10^4 and 10^5 counts show little change between those corresponding to calibration uncertainties, except in some cases a small change in centroid. This is not surprising since the centroid for 10^4 count simulation will be subject to some Poisson uncertainty indicated by the black histograms in the upper panels.

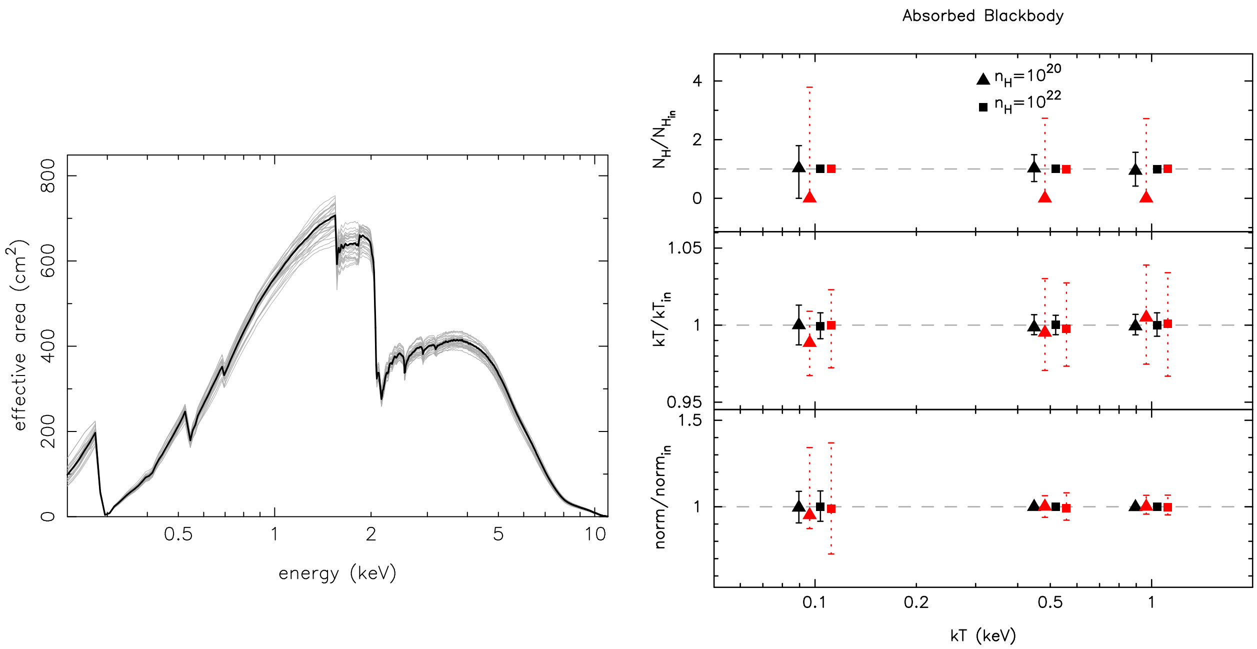

Figure 6. Left: The nominal "seed" *Chandra* ACIS-S effective area (black) compared with a sample of 30 effective areas generated using the Monte Carlo modification method described in the text (grey). **Right:** Modes and highest posterior density ±95% confidence intervals obtained for the blackbody models investigated using XSPEC for synthetic data sets containing 10^5 counts. The y-axes correspond in both cases to the ratio of the input parameter to that retrieved in the model fit. Black error bars correspond to 1000 Monte Carlo samplings of the synthetic data and show the effects of Poisson noise variations alone. Dashed error bars correspond to fits to a single synthetic data set using 1000 Monte Carlo-generated effective areas and response matrices.

We note that some parameters distributions are naturally skewed or shifted from the input values which are all well-represented by the sharper peaks of the black histograms corresponding to Poisson errors in the lower panels; particular examples are the parameters for the optically-thin thermal plasma with $kT = 2$ keV, $N_H = 10^{20}$ cm^{-2} and solar metallicity.

The summary statistics shown for all the models investigated here in Figures 6 and 7 indicate fairly large uncertainties in some parameters. Most conspicuous are those for the cool $kT = 0.1$ keV absorbed blackbody, and for low column densities. These are not surprising since they depend heavily on the low energy calibration for which the uncertainties in this current study are largest. Surprises include the retrieved temperature for a 5 keV optically-thin plasma with $N_H = 10^{22}$ cm^{-2}, and the power law models with a relatively high $N_H = 10^{23}$ cm^{-2}. The limiting accuracy for estimation of the metallicity of an optically-thin plasma appears to be about 10%, with more absorbed models showing largest uncertainty. It seems likely that when estimating individual abundances the uncertainties will be larger, though confirmation must await a more thorough study.

6. PRINCIPAL COMPONENT ANALYSIS DIMENSIONALITY REDUCTION AND EXTENSION TO PUBLICLY AVAILABLE SOFTWARE

The end goal of this study is to develop a system whereby users of the CIAO software can routinely incorporate calibration uncertainties in analyses of *Chandra* data. The Monte Carlo experiments described here were undertaken using a fairly cumbersome assemblage of specialised software and data. This method is computationally intensive and is difficult to implement as a portable software solution in the various extant fitting environments such as Sherpa and XSPEC. We have therefore investigated methods by which calibration uncertainties might be compressed into a much more compact format that could be used directly in a module of existing CIAO software. Principal Component Analysis appears to be well-suited to this problem.

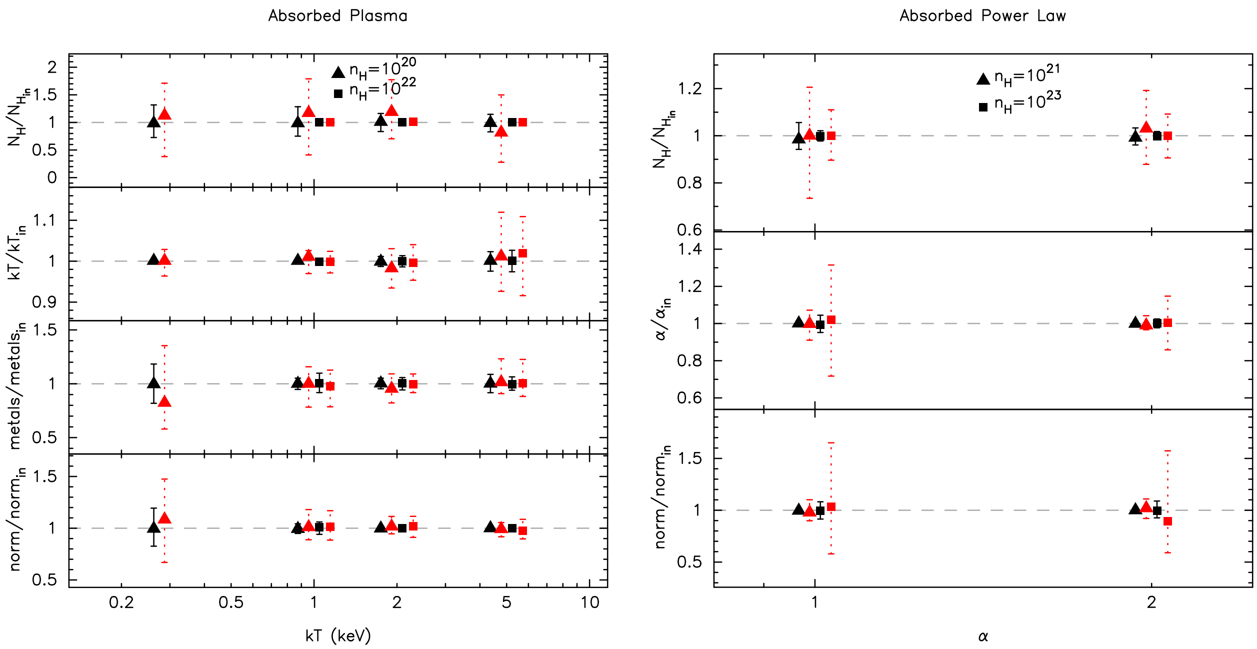

Figure 7. Same as Figure 6 right panel for the thermal plasma and powerlaw models investigated here.

Principal Component Analysis (PCA) is a linear transformation that chooses successive new coordinate systems for the data set such that the greatest variance by any projection comes to lie on the first axis (called the first principal component), the second greatest variance on the second axis, and so on. The low-order components generally contain the most important aspects of the variations in the data. By keeping lower-order principal components and ignoring higher-order ones, PCA can be used for reducing dimensionality in a dataset while retaining the characteristics that contribute most to its variance.

Experiments in applying PCA to our Monte Carlo Effective area library indicate that all variations in the data down to a level of a few percent can be contained in of order 10 principal components. A PCA decomposition of the effective areas represented in Figure 6 is shown in Figure 8.

We are developing the CIAO `sherpa` fitting engine to use the PCA decomposition of instrument response uncertainties through a Markov-Chain Monte Carlo (MCMC) technique, whereby a new effective area is realised at each iteration of the fit. Such a system should be highly portable, can be generalised to other aspects of calibration (such as the point spread function) and would be easily applied to other instruments and missions. The MCMC approach is computationally efficient, since full probability distributions of the model parameters are obtained as a matter of course, including the effects of both statistical and systematic errors, in a single MCMC fit.

7. SUMMARY AND FUTURE REFINEMENTS

We have described Monte Carlo methods to include a reasonably realistic description of the uncertainties in the calibration of the *Chandra* ACIS-S system in parameter estimation analyses. Application of this system to synthetic data provides insights into the limiting accuracy of the observatory for typical classes of celestial X-ray source. Based on our assessments of the current uncertainties in effective area, gain and pulse height distributions, it appears that the limiting accuracy of *Chandra* is reached in spectra containing about 10^4 counts. Beyond this, errors in in best-fit parameters due to calibration uncertainties completely dominate those due to photon noise.

Current work aims to extend this type of analysis to the other instrument combinations of *Chandra*. However, there are also several obvious ways that the current proof-of-concept system can be improved. These include:

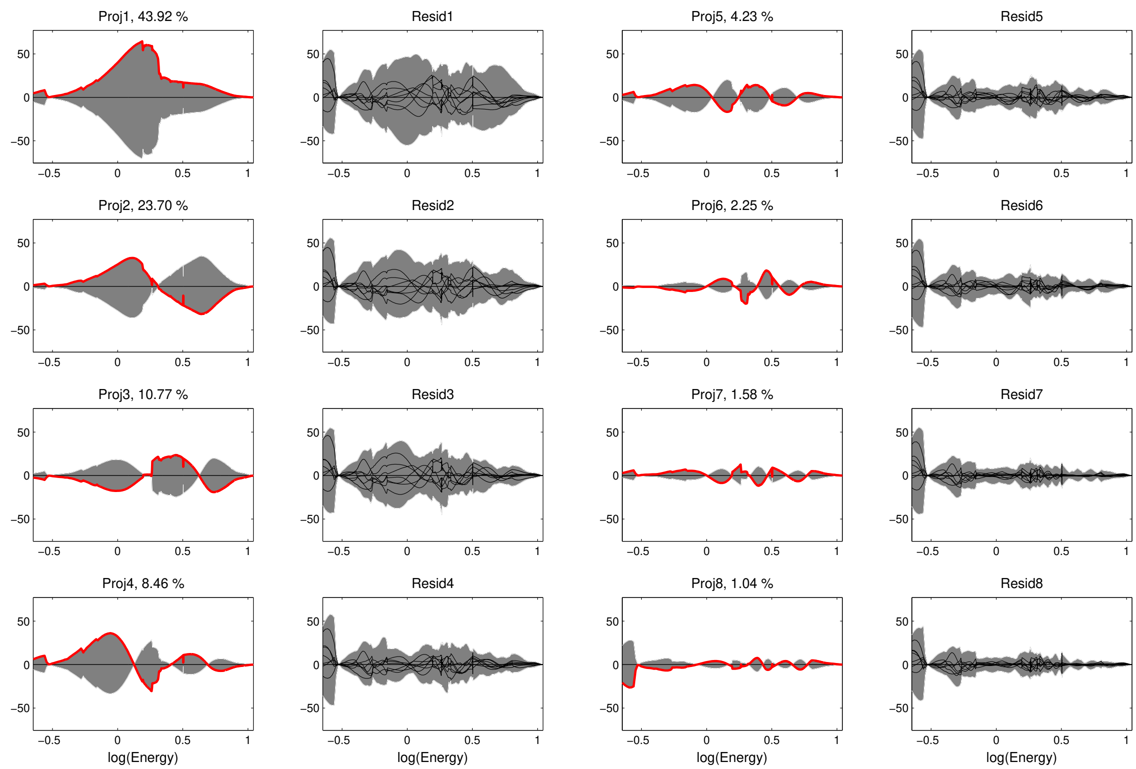

Figure 8. The first eight components from a Principal Component Decomposition of the effective areas represented in Figure 6. The left and right-center columns show the principal component directions superimposed on the range of variation of the projections of the curves onto each component (grey). The contribution of these projections to the total variation in the data is given in the title of each panel. The left-center and right columns show the cumulative curve residuals' range of variation (grey).

allowance for asymmetric uncertainties; more extensive use of Monte Carlo-driven model response predictions; and the imposition of further observational constraints, such as, for example, might be afforded on the overall "smoothness" of the effective area by high quality grating observations of continuum sources.

Improvements in model predictions are perhaps most readily achievable for the HRMA and its alignment on which the effective area also depends. The ACIS CCD model used here is also schematically simple and effects such as the escape photon fraction from fluorescence following events at energies above the Si K edge are currently not included. These improvements are currently under investigation and will be reported on in future publications.

ACKNOWLEDGMENTS

We would like to thank Catherine Grant for help with the ACIS QE model and Alanna Connors, Andreas Zezas and members of the *Chandra*-Harvard-Smithsonian-California Astrostatistics Collaboration (CHASC) for insightful discussions. This work was supported by NASA contract NAS8-39073 to the *Chandra X-ray Center* and by NASA grant AISRP NNG06GF17G.

REFERENCES

1. K. A. Arnaud, "XSPEC: The First Ten Years," in *ASP Conf. Ser. 101: Astronomical Data Analysis Software and Systems V*, G. H. Jacoby and J. Barnes, eds., pp. 17, 1996.
2. P. Freeman, S. Doe, and A. Siemiginowska, "Sherpa: a mission-independent data analysis application," in *Proc. SPIE Vol. 4477, Astronomical Data Analysis*, J.-L. Starck and F. D. Murtagh, eds., pp. 76–87, Nov. 2001.

3. J. P. Hughes, D. A. Schwartz, D. M. Worrall, and M. V. Zombeck, "Derivation of requirements for calibrating the Advanced X-ray Astrophysics Facility (AXAF) telescope," in *X-ray instrumentation in astronomy II; Proc. SPIE Vol. 902*, L. Golub, ed., pp. 380–386, 1988.
4. M. C. Weisskopf and S. L. O'dell, "Calibration of the AXAF observatory: overview," in *Proc. SPIE Vol. 3113, Grazing Incidence and Multilayer X-Ray Optical Systems*, R. B. Hoover and A. B. Walker, eds., pp. 2–17, July 1997.
5. C. L. Allen, P. P. Plucinsky, B. R. McNamara, R. J. Edgar, N. S. Schulz, and J. W. Woo, "Analysis of the AXAF HRMA+ACIS effective area measurements from the XRCF," in *Proc. SPIE Vol. 3444, X-Ray Optics, Instruments, and Missions*, R. B. Hoover and A. B. Walker, eds., pp. 198–209, Nov. 1998.
6. J. A. Nousek et al., "Joint AXAF high-resolution mirror assembly and AXAF CCD Imaging Spectrometer calibration at the MSFC X-Ray Calibration Facility," in *Proc. SPIE Vol. 3444, X-Ray Optics, Instruments, and Missions*, R. B. Hoover and A. B. Walker, eds., pp. 225–233, Nov. 1998.
7. D. A. Schwartz et al, "Absolute effective area of the Chandra high-resolution mirror assembly (HRMA)," in *Proc. SPIE Vol. 4012, X-Ray Optics, Instruments, and Missions III,*, J. E. Truemper and B. Aschenbach, eds., pp. 28–40, July 2000.
8. P. P. Plucinsky et al, "Flight spectral response of the ACIS instrument," in *X-Ray and Gamma-Ray Telescopes and Instruments for Astronomy. Edited by Joachim E. Truemper, Harvey D. Tananbaum. Proceedings of the SPIE, Volume 4851*, pp. 89–100, Mar. 2003.
9. D. E. Graessle, R. Soufli, A. L. Aquila, E. M. Gulliks on, R. L. Blake, and A. J. Burek, "Iridium optical constants for the Chandra X-ray Observatory from r eflectance measurements of 0.05-12 keV," in *X-Ray and Gamma-Ray Instrumentation for Astronomy XIII. Edited by Flanagan, Kathryn A.; Siegmund, Oswald H. W. Proceedings of the SPIE, Volume 5165*, pp. 469–481, Feb. 2004.
10. H. L. Marshall, A. Tennant, C. E. Grant, A. P. Hitchcock, S. L. O'Dell, and P. P. Plucinsky, "Composition of the Chandra ACIS contaminant," in *X-Ray and Gamma-Ray Instrumentation for Astronomy XIII. Proceedings of the SPIE, Volume 5165*, K. A. Flanagan and O. H. W. Siegmund, eds., pp. 497–508, Feb. 2004.
11. P. Zhao et al, "Chandra X-ray Observatory mirror effective area," in *X-Ray and Gamma-Ray Instrumentation for Astronomy XIII. Proceedings of the SPIE, Volume 5165*, K. A. Flanagan and O. H. W. Siegmund, eds., pp. 482–496, Feb. 2004.
12. D. H. Jerius et al, "The role of modeling in the calibration of the Chandra's optics," in *X-Ray and Gamma-Ray Instrumentation for Astronomy XIII. Proceedings of the SPIE, Volume 5165*, pp. 402–410, Feb. 2004.
13. B. J. Wargelin, P. W. Ratzlaff, D. O. Pease, V. L. Kashyap, and J. J. Drake, "Chandra LETG higher-order diffraction efficiencies," in *X-Ray and Gamma-Ray Instrumentation for Astronomy XIII. Volume 5165*, pp. 509–517, Feb. 2004.
14. H. Marshall, "Preliminary Results Concerning the Calibration of the LETG/ACIS Effective Area," in *Internal Memorandum*, pp. `http://space.mit.edu/ASC/calib/letg_acis/`, May 2002.
15. G. Chartas, G. P. Garmire, J. A. Nousek, L. K. Townsle y, F. R. Powell, R. L. Blake, and D. E. Graessle, "ACIS UV/optical blocking filter calibration at the National Synchr otron Light Source," in *Proc. SPIE Vol. 2805, Multilayer and Grazing Incidence X- Ray/EUV Optics III, Richard B. Hoover; Arthur B. Walker; Eds.*, pp. 44–54, July 1996.
16. R. Edgar and A. Vikhlinin, "Absolute QE of ACIS S1, S2 and S3 from XRCF data at selected energies," in *Internal Memorandum*, pp. `http://cxc.harvard.edu/cal/Acis/Cal_prods/qe/index.html`, Aug. 2004.
17. A. Vikhlinin, "High-resolution QEU maps for ACIS-I and S1,2,3 chips," in *Internal Memorandum*, pp. `http://cxc.harvard.edu/cal/Acis/Cal_prods/qeu/index.html`, Mar. 2004.
18. M. Bautz and J. Nousek, "Science Instrument (SI) Calibration Report for the AXAF CCD Imaging Spectrometer (ACIS)," in *Final Pre-launch Calibration Report*, Jan. 1999.
19. C. ACIS Team, "Analysis of ACIS Data Affected by the Low Energy QE Degradation," in `http://cxc.harvard.edu/cal/Acis/Cal_prods/qeDeg/2003_11_19/index.html`, Mar. 2004.
20. A. Vikhlinin, R. Edgar, and N. Schulz, "Corrections for Time-dependence of ACIS Gain," in *Internal Memorandum*, pp. `http://cxc.harvard.edu/cal/Acis/Cal_prods/tgain/index.html`, July 2004.
21. N. Schulz et al, "ACIS Energy Response Performance with the External Calibration Source," in *Chandra Calibration Workshop 2005*, Oct. 2005.