1.1 Basic CCD Model Components and Measurement Strategy

The fundamental strategy for calibration of the ACIS instrument is to make measurements which constrain free parameters of mathematical models of the system. This approach was followed in calibration of the CCD detectors themselves, and the reader must therefore understand the detector models in order to understand the detector calibration.

In this section we describe some principles of X-ray CCD detector operation and describe the parameters of the (simplified) CCD model which we have fit to the data. In section 1.2, we present an overview of the measurements made to constrain those parameters.

1.1.1 A Brief Description of ACIS CCDs

The ACIS focal plane contains ten MIT Lincoln Laboratory CCID17 charge-coupled detectors. The CCID17 ([24]) is a three-phase, frame transfer imager with 1026 rows of 1024 columns in each of the imaging and framestore areas . The 2 rows nearest the framestore section are ignored by the ACIS digital electronics. The pixels in the imaging area are 24 microns square. Each device is served by four output nodes which are (usually) operated in parallel. The architecture is illustrated in Figure 1.1.

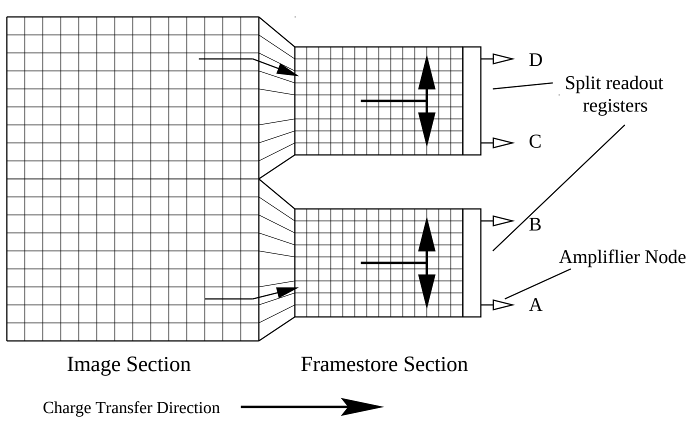

Figure 1.1: Layout of the ACIS CCID17 detector, showing framestore architecture and four output nodes. Ten such devices comprise the ACIS focal plane.

ACIS flight detectors are fabricated of high-resistivity silicon ($\rho > 6500\Omega$-cm) to maximize depletion depth and, therefore, high-energy X-ray detection efficiency. The devices are three-side-abuttable to minimize inter-chip gaps in the flight focal plane. Flight CCDs have been fabricated in both front-illuminated and back-illuminated configurations. Other characteristics of the detectors are listed in Table 1.1.

Parameter	Value
Architecture	3-phase, frame-transfer, 3-side-abuttable
Format	1024 columns by 1026 rows, image array
	2x 512 columns by 1026 rows, framestore
Pixel Size	24 x 24 μm, image array
	21(H) x 13.5(V), μm, framestore
Output Nodes	4
Readout Rate	100 kpix s^{-1} at each of output four nodes
	400 kpix s^{-1} total per CCD
Operating Temperature	-120C (nominal)
Readout Noise:	
CCD	$< 2e^{-}$ RMS, typical
System	2 - 3 e^{-} RMS
Charge Transfer Inefficiency:	
(Low-illumination, 5.9 keV)	
FI Devices	$< 3 \times 10^{-6}$ per pixel transfer
BI Devices	$< 1 - 3 \times 10^{-5}$ per pixel transfer
Mean Dark Current (-120C)	$< 5 \times 10^{-2} e^{-} s^{-1}$ pixel^{-1}

Table 1.1: Some Characteristics of the ACIS CCD (MIT Lincoln Laboratory CCID17)

1.1.2 X-ray CCD Detection Efficiency: Processes, Models and Parameters

A cross-section of a generic, front-illuminated CCD is shown in Figure 1.2. In this figure, X-rays incident from the top must pass through a number of "dead" layers which make up the so-called gate structure before entering the photosensitive volume of the device. The function of the gate structure is to allow charge collected in the device to be moved (transferred) from the vicinity of the photon interaction site to an output amplifier which converts the charge to a measurable electrical signal. The direction of charge transfer is indicated in the figure. The gate structure consists of overlapping strips of highly-doped, conductive polysilicon, interlarded with insulating silicon dioxide, and undergirded by relative thin, uniform insulating layers of silicon nitride and silicon dioxide. Although the gate structure is not completely insensitive to X-rays, (see section 1.14 below) for present purposes we assume it to be so.

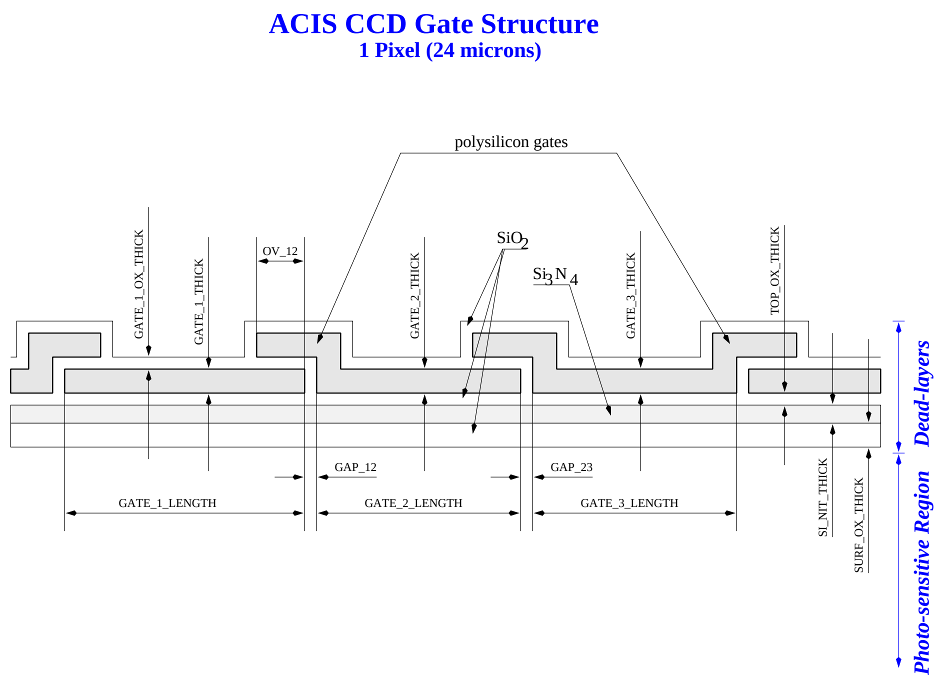

Figure 1.2: Cross-section through a front-illuminated CCD. X-rays are incident from the top. The gate structure poses a deadlayer; only the depletion region region is usefully photosenstive. Calibration data must constrain the physical dimensions of these and other regions shown in the figure.

A section of the CCD in a plane orthogonal to that shown in figure 1.2, would reveal a set of structures called the channel stops. (The channel stops are illustrated and discussed in great detail in section 1.5 below). The channel stops function to isolate adjacent CCD columns from one another, and present additional dead and partially dead layers which affect CCD detection efficiency.

Back-illuminated CCDs differ from front-illuminated devices in two crucial ways. First, back-illuminated devices are thinned; essentially all of the material shown "below" the depletion region in Figure 1.2 is removed during device fabrication. In fact, to ensure a sufficiently strong electric field at the back surface, some of the putative depletion region is removed during the thinning. Therefore, a back-illuminated device is thinner than the depleted region of a front-illuminated device. Second, the device is mounted so that the X-rays are incident from the bottom in Figure 1.2. As a result, the gate structure does not function as a deadlayer in a back-illuminated device. A very thin deadlayer of silicon dioxide (of order $\sim$ 500 Angstroms thick) does exist on the back (illuminated) surface of the back-illuminated devices.

The total thickness of the gate structure is of order 0.5 microns. Deadlayers associated with the channel stops are somewhat thicker, but cover a relatively small fraction of the device (see section 1.5, below). In any event, the deadlayer thicknessses (rather than, say, electronic noise) largely determines the low-energy limit of the response of the front-illuminated CCD. By construction the gate structures and channel stops are not deadlayers for the back-illuminated CCD, and the "low-energy limit" of these devices is much lower than that of the front-illuminated detectors.

X-rays penetrating the gate structure may be absorbed photoelectrically by the silicon below. If the X-ray is absorbed, the resulting primary photoelectron produces secondary ionization within a volume which is very small compared to the scale of 1 CCD pixel (generally much less than 1 micron in radius for X-rays in the 0.1- 10 keV bandpass of interest). Auger relaxation of the absorbing silicon atom will also contribute to the charge produced. The volume immediately below the gates (for the first 60 - 75 μm) is depleted of free charge carriers. The residual negative fixed charge provided by the ionized acceptor impurities (the substrate is very lightly doped p-type), together with the charge induced on the gates when they are biased, produces a relatively strong electric field in the depletion region. This field sweeps photo-liberated electrons up to the charge transfer channel very rapidly (on timescales of nsec or less). Once in the transfer channel, the charge may be transferred to the CCD output node by application of proper voltage waveforms to the various gates, and then measured.

The depletion region of a typical front-illuminated ACIS flight device is between 65 and 75 microns thick; the back-illuminated devices have photosensitive volumes which are 30 to 40 microns thick. These thicknesses essentially determine the high-energy detection efficiency limit.

In simplest terms, then, the CCD detection efficiency is the probability than an incident X-ray will interact (photoelectrically) in the depletion region, rather than being absorbed in the gate structure or passing entirely through the depletion region without interaction. This is in principle a straightforward problem requring that one find the spatially averaged transmission of the gate structure, and the transmission of the (assumed uniformly thick) depletion region as functions of incident X-ray energy. If the relevant mass absorption coefficients are known, (except near edges, where, as discussed in section 1.6.4, below, we have measured the mass absorbtion coefficients, we adopt the values published by Henke in 1993) the problem reduces to one of determining dimensions.

In practice, even this simple model is rather more complex than the one we have fit to the data. Specifically, all analysis presented in this report will assume that each gate structure layer is uniform along the charge transfer direction shown in Figure 1.2, and piecewise uniform in the direction normal to the page of the figure. The quantum efficiency model then has only the seven gate structure parameters listed in Table 1.2; the thickness of the depletion region is the eighth and final model parameter.

Parameter/Description	Typical Value (μm)
Pixel Width	24.0
Silicon Gate Thickness	0.25 - 0.30
SiO_2 Insulator Thickness	0.20 - 0.35
Si_3N_4 Insulator Thickness	0.02 - 0.04
Channel Stop Width	4.1
Channel Stop Silicon Implant Thickness	0.35
Channel Stop SiO_2 Thickness	0.45

Table 1.2: Parameters for "Slab and Stop" Model of CCD Gate Structure

While this model oversimplifies the three-dimensional gate structure, it is worth noting that it is exactly correct to first order in the gate structure optical depth. Thus it is least accurate at energies where the gate optical depth is large (e.g., at energies immediately above those of the K absorption edges of oxygen and, to a lesser extent, of nitrogen and silicon). Our detection efficiency model also neglects the influence of charge collection efficiency (described below) on detection efficiency; fortunately, we have found that by judicious choice of event selection criteria, that the latter effects are independent of energy. We discuss this point further in the following paragraphs. Finally, it is worth stressing that the depletion approximation is just that; in adopting this approximation by fitting for the "depletion depth" parameter (see section 1.6.2), we have ignored the fact that the electric field does not drop linearly to zero in the substrate. The remarkable accuracy of this approximation is a consequence of the fact that the depletion region is quite large compared to the amount the initial charge cloud diameter changes over the energy band.

Although to date we have used only the simplified detection efficiency model described in this section to predict ACIS detector quantum efficiency, we have developed a more detailed model for use in future analysis. This model is described in some detail in section 1.14. We must rely on the sub-pixel structure measurements described in section 1.5, together with (destructive) electron micrographs of sibling devices produced by Lincoln Laboratory, to estimate typical values for some of the parameters (gate overlaps, phase-to-phase gate thickness variations, and channel stop dimensions) of the detailed gate structure model.

1.1.3 Spectral Redistribution: Processes, Models and Parameters

A spectral redistribution function describes the probability of an instrument response in each pulse-height channel to photons of any particular energy. In this section we discuss the spectral redistribution function and its models.

To guide this discussion, a typical spectral redistribution function for a front-illuminated CCD (the response of a flight-like detector to a radioactive ^{55}Fe source) is shown in Figure 1.3. The main components of the redistribution function are indicated: in addition to the two primary photopeaks at 5.9 and 6.4 keV (from the Kα and Kβ lines of the manganese daughter of the radioactive decay), one sees Si-K escape and flourescence peaks, an assymetrical shoulder onn the low-energy side of the photopeaks, a low-energy continuum, and an upturn (tail) at very low energies. (The Manganese L lines from the source are also visible). The off-nominal features (everything except the primary photopeaks) together contribute about 2% of the integrated area under the response function at this energy. The task of modelling the spectral response function is to predict this curve in sufficient detail to meet AXAF calibration requirements.

We divide our discussion into three parts: the energy scale (or "gain"); the spectral resolution (essentially the width of the photopeak); and the off-nominal features (everything else).

Consider first the first moment of the spectral redistribution function, that is, the variation of the mean response (essentially the centroid of the primary photopeak) with incident energy. To high accuracy, the output of an ACIS front-illuminated detector (and its associated detector electronics) is a highly linear function of the *charge* generated in the detector. In the absence of pileup, the photo-charge is, in turn, a highly linear function of the input photon energy. This linear relationship between input energy and the output pulseheight is sometimes referred to as the energy scale or "gain" relation. While the linearity of the energy scale rests on the well-understood physics discussed briefly below, (not to mention nearly-ideal detectors and meticulous electrical engineering), the program of the energy scale calibration is the empirical one of determining the slope and intercept of this relationship for each of the forty CCD output nodes as a function of focal plane and detector electronics temperatures. Results of this determination are presented in sections

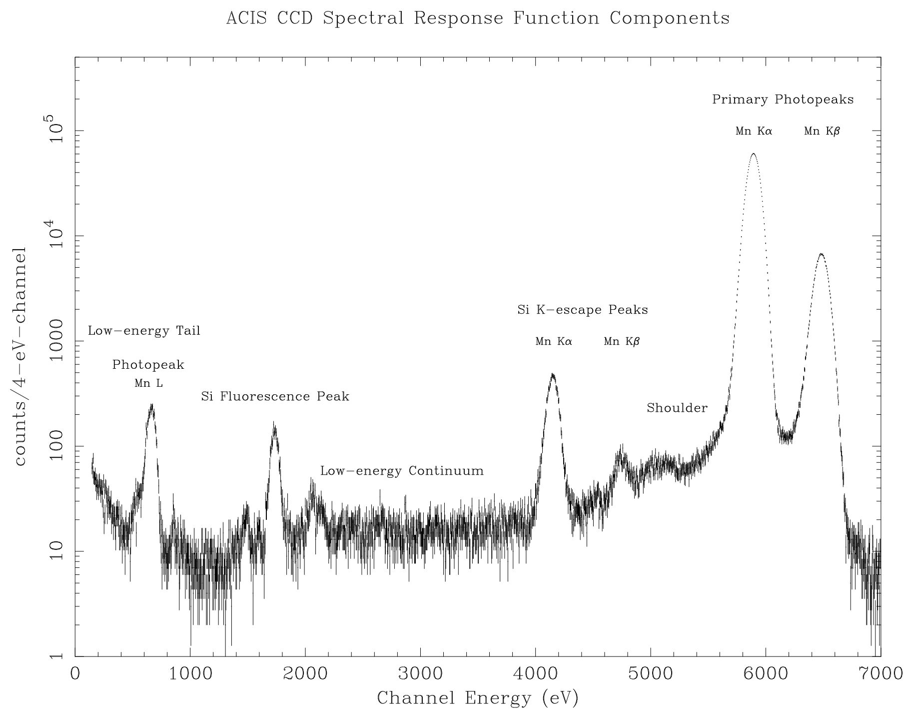

Figure 1.3: Principal components of the CCD spectral response function. The (asymmetric) primary photopeaks, silicon K-escape and fluorescence features, and the extended low-energy tail are indicated. Note that the upturn in the tail at very low energies (E< 400 eV) is not due to electronic noise, but is a characteristic of the detector itself. Note that the ordinate is logarithmic.

1.3.1 and 1.9.2. It is worth noting, that, as is shown in section 1.3.1, the back-illuminated detectors exhibit significant non-linearity that we have not yet modelled in physical terms.

The second moment of the redistribution function (the detector's "spectral resolution", as measured by the width of the photopeak) is, in the case of front-illuminated detectors, equally amenable to simple modelling. In the mean, the quantity of charge liberated by an X-ray interacting in the CCD is proportional to the energy of the incident X-ray:

$$N_e = \frac{E}{w} \tag{1.1}$$

where N_e is the number electrons liberated, E is the photon energy and $w \approx 3.7eV/e^-$, the mean ionization energy per electron-hole pair, is a function of the temperature of the silicon.

The two most important factors affecting the spectral resolution of the CCD are readout noise and the stochastic nature of the ionization process. The readout noise characterizes the accuracy with which one can measure the quantity of charge deposited in any given pixel. For ACIS detectors, the readout noise is found to be a normally distributed, zero mean random deviate (see, e.g., [1]) of standard deviation $\sigma_r \sim$ 2-3 electrons, RMS, per read. The variance σ_N^2 on the charge liberated N_e, is

$$\sigma_N^2 = F \times N_e = F \times \frac{E}{w} \tag{1.2}$$

where F, the Fano factor, has the value $F = 0.135$, and is characteristic of crystalline silicon. Using the room-temperature measurement of Scholze et al (1996) [2], viz, $w = 3.64 \pm 0.03eV/e^-$ and the temperature dependence reported by Canali et al (1972) [3], who find that w increases linearly with temperature between 77K to 300 K, and we expect that $w = 3.71eV/e^-$ at the CCD operating temperature of $T = 153K$.

The effect of these two processes alone, then, is that the primary photopeak produced in response to a beam of monochromatic incident photons of energy E will be a Gaussian distribution with mean proportional to E and standard deviation, expressed on the input energy scale, given by

$$\sigma_E = \sqrt{(w\sigma_r)^2 + (F \times w \times E)} \tag{1.3}$$

The form of equation 1.3 is an accurate model of the spectral resolution of front-illumianted ACIS detectors, as is shown in section 1.3.1. This accuracy is the more remarkable in light of the fact that equation 1.3 neglects the summation of multiple pixels to compute the event amplitude (see below). The insensitivity of the spectral resolution to multi-pixel summation arises because the readout noise σ_r is low enough, and the depletion depth large enough, that $\sigma_N^2 \gg \sigma_r^2$ at all energies where a significant fraction of events must be summed. Note also that equation 1.3 implies that the the spectral resolution can be computed from i) a measurement of the readout noise and ii) knowledge of w and F. As discussed in section 1.3.1, the best fit model parameters differ significantly from the values expected on this basis. Finally, the resolution of the back-illuminated devices deviates considerably from the predictions of equation 1.3. This deviation has yet to be modelled in detail.

As Figure 1.3 shows, a number of processes complicate the CCD spectral redistribution function. In addition to the K-escape and fluorescence phenomena common to all photoelectric X-ray detectors, the CCD response function is influenced by incomplete charge collection and, to a more limited extent, by photoelectric interactions in the gates and channel stops.

The amplitude of the K-escape peak may be modelled reasonably accurately through straightforward simulation of photon transport in the detector. With somewhat less accuracy, the (lower) amplitude of the silicon flourescence line can also be modelled; see section 1.3.2, below. The locations and widths of these peaks are also modelled in the obvious way.

In practice, then our model of the portion of the spectral redistribution function discussed to this point consists of a sum, for each incident energy, of a set of Gaussian distributions of fixed relative centroid locations, with absolute locations, relative amplitudes and widths constrained by fits to measurements. Together, these components account for all but about 1% of the spectral redistribution function for the front-illuminated devices.

Finally, we come to the remaining off-nominal features: the shoulder, low-energy continuum and low-energy tail. The physical origin of these features is believed to be well-understood, but their amplitude is not easily desribed analytically.

In spite of the relatively small volume occupied by initial the photoelectric charge cloud, there is a non-zero probability that charge can be shared among multiple pixels, because events can occur near pixel boundaries. Although the event amplitude computation algorithm sums together pixels in which signficant charge is detected, a pixel is only included in the sum if it exceeds a so-called split-event threshold, typically set at 15 electrons (equivalent to about 55 eV). The split-event threshold excludes empty (but noisy) pixels from the event amplitude summation, but has the undesirable consequence of excluding some bonafide photo-ionization from the sum. As a result, the spectral redistribution function can contain a low-energy shoulder.

Moreover, some charge from events which interact below the depletion region (in front-illuminated devices) will diffuse to the back of the device and will never be collected in the charge transfer channel. Events of this sort can generally be identified because the charge that is collected must diffuse large distances, and hence can occupy several pixels (more than 3). Charge can also be lost when photons interact in the channel stops and gate structure. These processes produce a "low-energy continuum" tail on the spectral response function. A more detailed account of this model, developed at MIT by Prigozhin et al. (1998), is given in section 1.3.2.

1.1.4 De-coupling Detection Efficiency and Spectral Resolution by Choice of Event Selection Criteria

The charge diffusion and charge loss processes just described also affect the effective detection efficiency of the CCD. This coupling operates mainly by way of the "event-shape-based" event selection criteria which, as discussed above, are used to reject events for which charge collection is so poor that the event amplitude is a poor measure of the incident photon energy. If each and every event rejected on grounds of shape resulted from interactions outside the depletion region, then the effect of the selection criteria could be modelled in a straightforward "geometrical" way. In fact, the selection criteria are not perfectly efficient;

they accept a fraction of events occuring outside the depletion region, and reject a fraction of events occuring in the depletion region.

While the ultimate CCD model will predict this coupling by accurate simulation of the details of the charge diffusion process, for present purposes we have attempted to avoid the problem by choosing event selection criteria that minimize the influence of charge collection effects on detection efficiency. Specifically, we elect to compute the response of the CCD for ASCA event grades 0, 2, 3, 4 and 6, with a split event threshold of 15 electrons. (This criterion accepts all 1- and 2-pixel events, as well as "L-shaped" 3-pixel events and square or "Quad-shaped" 4-pixel events. More precisely, our standard grade selection criterion accepts events in accordance with the following ACIS grade map: 0x1d05 0x0047 0x0004 0x0047 0x1133 0x0003 0x1100 0x0000 0x1d04 0x0000 0x0004 0x0000 0x0000 0x0003 0x0000 0x0000. Readers without decoder rings may see the ACIS Software Requirements Specification for an interpretation.

This criterion also turns out to produce a response function which is relatively insensitive to the precise value of the split-event threshold.

In principle, the event selection criteria we've adopted represent the high-detection efficiency extreme of detection efficiency-*vs.*-spectral resolution tradeoff. However, in practice, the readout noise of the ACIS detectors is so low, and the depletion depth so large, that, at least for front-illumianted detectors, there is little sacrifice in spectral resolution entailed by this choice. As is discussed elsewhere in this report, the situation is not so clear for the back-illuminated detectors.