Advanced CCD imaging spectrometer (ACIS) instrument on the *Chandra X-ray Observatory*

G. P. Garmire[*a], M. W. Bautz[**b], P. G. Ford[b], J. A. Nousek[a], and G. R. Ricker Jr[b]
[a]The Pennsylvania State University; [b]Massachusetts Institute of Technology

Abstract

The ACIS instrument has been operating for three years in orbit, producing high quality scientific data on a wide variety of X-ray emitting astronomical objects. Except for a brief period at the very beginning of the mission when the CCDs were exposed to the radiation environment of the Outer van Allen Belts which resulted in substantial radiation damage to the front illuminated CCDs, the instrument has operated nearly flawlessly. The following report presents a description of the instrument, the current status of the instrument calibration and a few highlights of the scientific results obtained from the Guaranteed Observer Time.

1. Introduction

The Advanced CCD Imaging Spectrometer (ACIS) on board the *Chandra X-ray Observatory(CXO)* [1] is a powerful tool for conducting imaging, spectroscopic and temporal studies of celestial X-ray sources. The instrument consists of ten Charge Coupled Devices (CCDs) especially designed for efficient X-ray detection and spectroscopy [2]. Four of the front illuminated (FI) CCDs are arranged in a square array with each CCD tipped slightly to better approximate the curved focal surface of the Chandra Wolter type I mirror assembly. The remaining six CCDs are set in a linear array, tipped to approximate the Rowland circle of the objective gratings that can be inserted behind the mirrors (see Figure 2.1). One CCD next to the center-line of the grating array is essentially flat and is a back illuminated (BI) CCD that is useful for imaging soft X-ray objects. Each CCD subtends an 8.4 arc minute by 8.4 arc minute square on the sky. The individual pixels of the CCDs subtend 0.492 arc seconds on the sky. The on-axis performance of the telescope is better than 0.5 arc second. The spacecraft normally is commanded to conduct an observation utilizing a dither motion in the form of a Lissajous pattern over a 16 arc second square area of sky. This motion slowly moves (typically, 0.1 arc second per CCD exposure) any X-ray object across a 32 by 32 pixel region of the CCDs. It is possible to deconvolve the images obtained from the CXO/ACIS to better than 0.5 arc second resolution. In Figure 1 the image of SN1987A, the most recent supernova in the LMC, is resolved even though it subtends only about 1.3 arc seconds (see Burrows, [3] for more details).
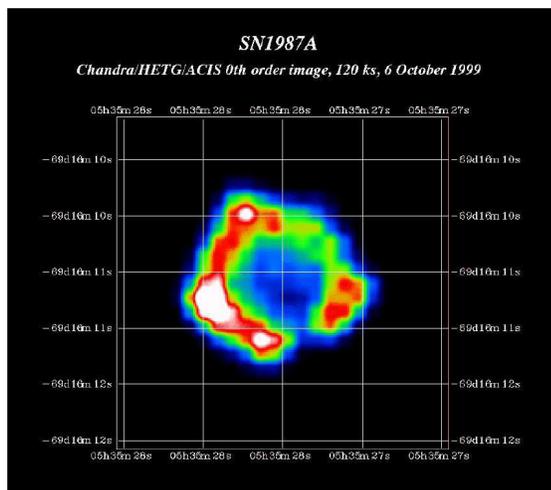

Figure 1

The CCDs utilize a framestore design, such that no shutter is required, the image being shifted to the framestore area in 41 ms, very short compared to a typical exposure of 3.24 seconds. The frame transfer causes streaking for very bright sources, but more typical exposures have no bright source in the field. Shorter exposures are possible by using only a portion of the CCD area, but with a loss of field on the sky and observing efficiency. The CCDs may also be used in a continuous clocking mode to achieve a line readout time of 2.85 ms (equivalent to a time resolution of ~6 ms for a point source) but at the loss of one spatial dimension. The background on orbit is very low (6 x 10^{-8} count/pix/sec in the 0.5 – 2.0 keV band and 1.3 x 10^{-7} counts/pix/s in the 2.0 – 8.0 keV band) as obtained from the dark moon observations [4]. Even in a 2Ms exposure [5] 40% of the pixels are free of backgr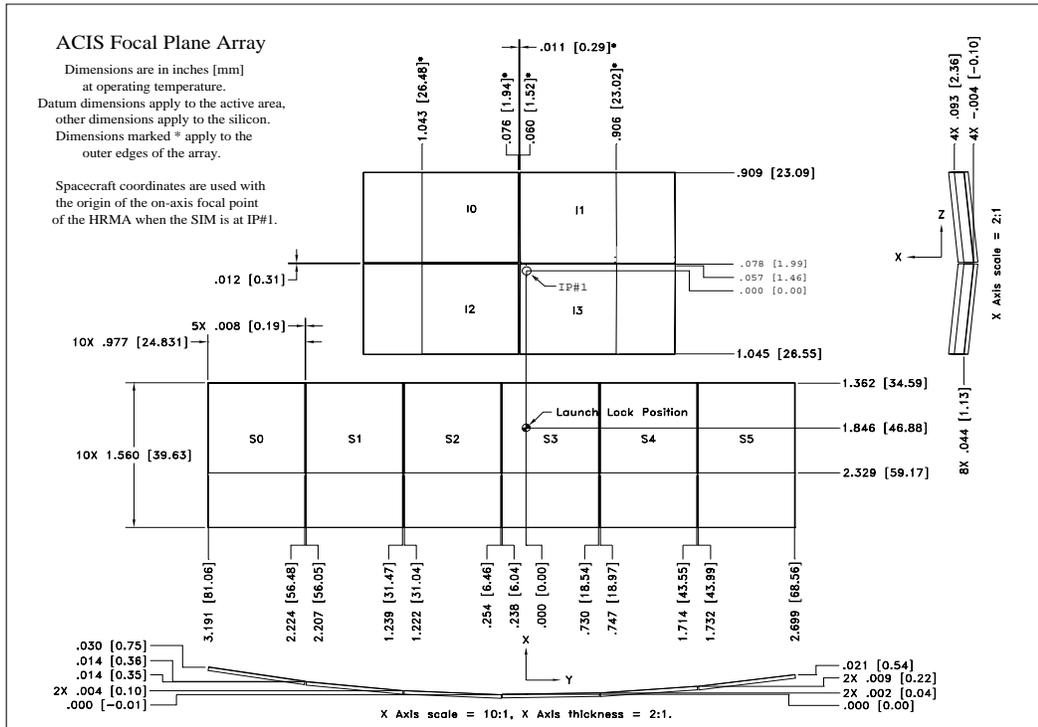ound events. A much more complete description of the instrument is available at the web site listed at the end of this article. This report is intended to give an over view of the instrument with sufficient detail to aid the potential observer in understanding how the instrument works and what its capabilities are for observations.

2. Focal Plane Design

The ACIS focal plane (Figure 2.1) is composed of two arrays with one array designed to optimize imaging and the other array for spectroscopy utilizing the objective gratings that are part of the Chandra X-ray Observatory (CXO). In the case of the imaging array, the CCDs are tipped to lie tangent to the optimum focal surface of the Wolter Type I mirrors employed for imaging. This nearly spherical focal surface lies 10.04 meters from the joint separating the paraboloid and hyperboloid mirrors and has a radius of curvature of 85 mm. By tilting the CCDs in the imaging array more pixels encompass regions of high angular resolution (see Figure 2.3 for a comparison between a flat array and a tilted array. The dashed curve in (a) starting at 7 arc seconds is for the S-array portion of the image). The spectroscopic array is also tilted, but to a lesser degree, since the Rowland circle for the gratings has a radius of curvature of 4.2 m. A detailed layout of the focal plane is shown in Figure 2.1. The position of the detected X-ray in the spectroscopic array is directly proportional to the wavelength of the photon. The intrinsic energy resolution of the CCDs is used to separate the overlapping orders of the dispersed spectrum. The spectroscopic array is composed of two different CCD designs, four (FI) three phase CCDs and two (BI) three phase CCDs , S1 and S3 in Figure 2.1. The (BI) device labeled S3 is essentially flat (perpendicular to the optical axis of the telescope) and provides an additional imaging capability, especially for sources with mainly low energy X-ray emission (below 1 keV) where the BI CCDs have superior quantum efficiency. Fabrication limitations require small gaps to exist between the CCDs (~11 arcsec for the I-array and 8.8 arcsec for the S-array). Dithering of the spacecraft in a Lissajous figure over a 16 arc second square pattern during an observation fills in these gaps to some extent. The actual flight focal plane is shown in Figure 2.2. The four posts a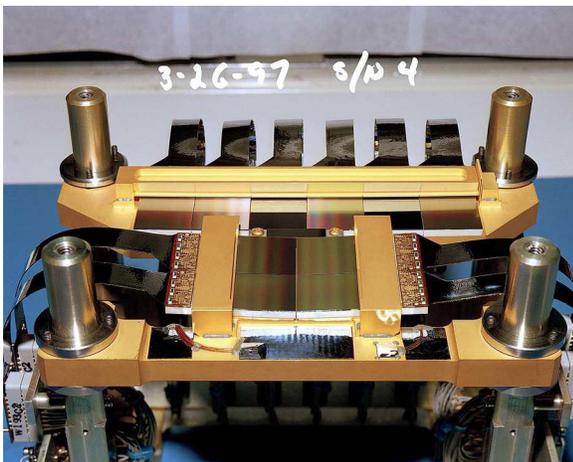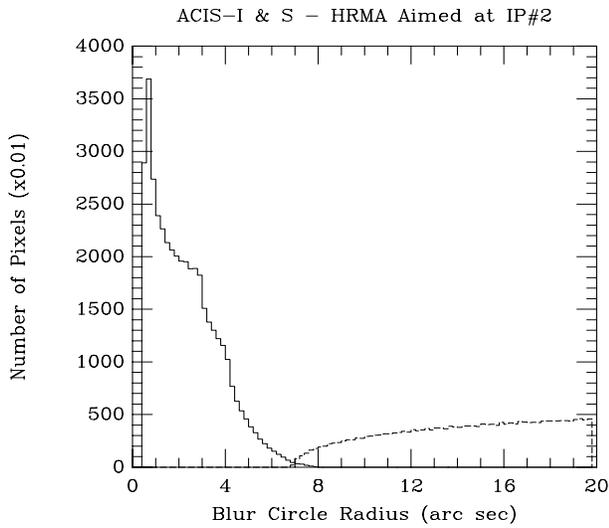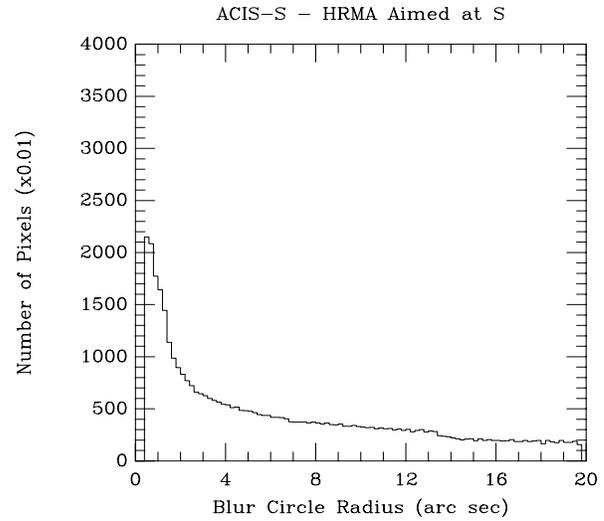re made from a plastic Torlon to insulate the focal plane from the rest of the camera body. Gold-coated aluminum bars cover all of the framestore areas to shield them from focused X-rays. The curved straps attached to each CCD are flexprints which carry the electrical signals to and from each CCD.

Figure 2.1

Figure 2.2

3561   03/31/95   16:46:32

Figure 2.3a

6784   04/03/95   17:00:59

Figure 2.3b

3. CCD Characteristics

Figure 3.1

The ACIS CCDs were developed at the MIT Lincoln Laboratory on high-purity p-type float-zone wafers of silicon with resistivities of about 7000 ohm-cm. A complete description can be found in [2]. Each CCD is a 1024x1026-pixel frame-transfer imager which is divided into four sectors. As shown in Figure 3.1, the framestore is split, and the framestore pixels are smaller than those in the imaging area. This arrangement allows for four independent output amplifiers, and facilitates 3-side abutment of the detectors. Each amplifier is a floating-diffusion output circuit with responsivity of 20 microvolts/e⁻ and a noise of about 2 e⁻ RMS at the operating rate of 100 kpix/s All four amplifiers are operated in parallel to minimize the frame readout time, which is typically 3.24 s. The room temperature dark currents were usually found to be about 500 pA/cm^2 with about 50 pA/cm^2 coming from the bulk emission and the remainder from surface-state dark current. By operating the device below -100 C and with frame times less than 10 s, the dark current per pixel is less than 1 e⁻/pixel/frame. Each pixel is 24 by 24 microns which at the focal surface of the CXO corresponds to 0.492 arc seconds. A cross section of the FI and BI CCDs is shown in Figure 3.2. The FI CCD design results in the charge from cosmic rays being generated in a nearly field free region of the CCD which spreads over a large number of pixels by the time it reaches the buried channel where it is stored for readout.
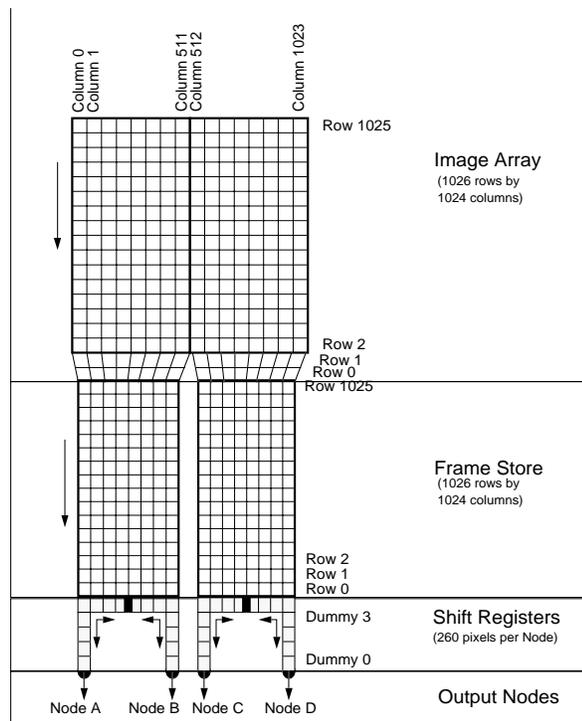
Also, high energy X-rays interact deeper in the silicon on average thus giving them more opportunity to diffuse into a large charge cloud as illustrated in the Figure 3.2. The BI CCDs, on the other hand, are just the reverse for X-rays, where the low energy X-rays interact far from the buried channel and diffuse into a larger charge cloud. Cosmic rays generate much smaller charge clouds in the BI architecture, since there is a much thinner field free region near the back surface. Nevertheless, the background rejection efficiency is much higher in the FI CCDs, with the result that the background is higher in the BI CCDs by a factor of 2 –3, depending on the energy.

The CCD design used for the ACIS includes a trough implant in the buried channel that increases the ability of the device to tolerate charge particle radiation encountered in the space environment. Two types of radiation damage are

.

Front-Illuminated X-ray CCD Structure
(not to scale)

Back-Illuminated X-ray CCD Structure
(not to scale)

Figure 3.2

typical of silicon based CCDs; displacement damage and ionizing damage. The former is induced by protons and other heavy particles in the cosmic ray flux, while the latter is induced by electrons or UV or X-ray photons that can deposit a charge in the insulating layer separating the gates from the buried channel. This has the effect of changing the potential well created by the clock voltages which in tern can affect the charge transfer efficiency. Over a five year baseline life of the *Chandra* Mission the accumulated dose of the displacement damage was estimated to be 1-3 kRad, while the ionizing dose is computed to be less than 10 k Rad (however, see section 3.1 concerning On-Orbit anomaly). These levels of displacement damage will produce a measurable increase in Charge Transfer Inefficiency (CTI) over a five-year period.

The CCDs were carefully calibrated in the laboratory at MIT using ten X-ray lines on 32 by 32 pixel regions of theCCDs to match the dither pattern size. This process smoothes out pixel to pixel variations and reduces the calibration time by a factor of one thousand. The calibrations were compared to a reference CCD that was calibrated at the synchrotron light source at BESSY in Berlin, Germany. The overall calibration was deemed to be accurate to a few percent ([6], ACIS Calibration.

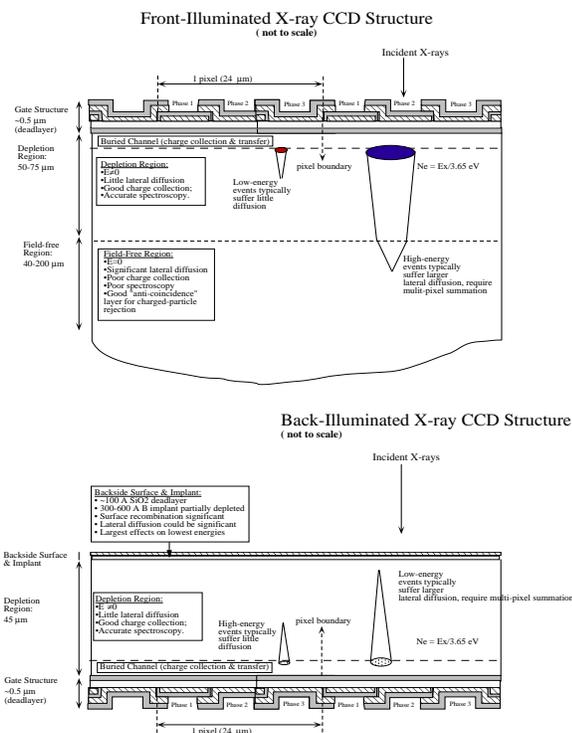Report).   An Optical Blocking Filter (OBF) was placed about 2 centimeters above the CCDs to limit the amount of light that could reach the CCDs from stars and from scattered light in the spacecraft. The filter was supplied by the Luxel Corporation using a 200nm substrate of free standing polyimide and coated with 160 nm of aluminum over the imaging portion of the array and 130 nm of aluminum over the spectroscopy array where less visible light was anticipated since much of it would be dispersed by the gratings. The X-ray transmission was calibrated over the same 32 by 32 pixel equivalent areas at the synchrotron light source at the University of Wisconsin. The transmission over the range 200 to 1200nm range was measured at Penn State, at the Denton Vacuum Corp using a Perkin Elmer UV/VIS/Near IR spectrometer and at Brookhaven National Laboratory. The
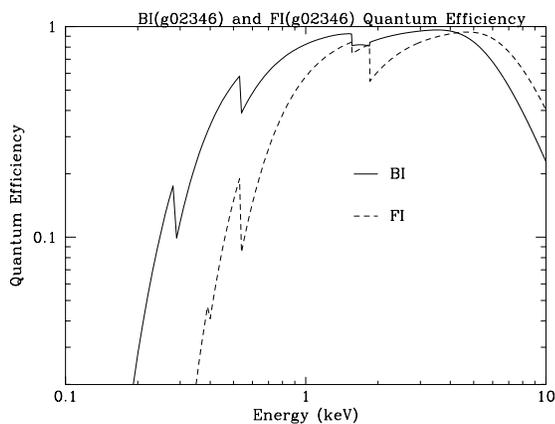

Figure 3.3

results of these calibrations are included in the ACIS Calibration Report referenced above. The observational consequences of the transmission of visible light are discussed below.

The overall, average quantum detection efficiency of the two types of CCDs used on ACIS is shown in Figure 3.3. The fine structure near the absorption edges (EXAFS) have been omitted in this figure. The range of variations from CCD to CCD is about 10% and within a FI CCD the variations are less than 3% from pixel to pixel. The variations within the BI CCDs are greater and reach about 15% for extreme cases. The QE of the CCDs is dependent upon what combination of pixels is used to create each X-ray event. It is frequently the case, depending upon X-ray energy, that some of the charge of the X-ray induced ionization reaches more than one pixel. The onboard processing electronics can be programmed to send several different pixel combinations to the ground through the telemetry of the spacecraft. Normally, a 3x3 pixel island (faint mode) is transmitted, with the central pixel containing the peak charge. An event threshold is set in the front end processor to only select events above the threshold. A second threshold is used by the processor to detect the charge from the adjacent pixels, the "split event threshold' which is usually lower than the event threshold. Each event is generally "graded" as determined by the pattern of split events. The QE will depend upon which grades are selected. For example, if only unsplit events are selected, the QE will be less than events where splits are allowed. See the above referenced ACIS Calibration Report for details.

A detailed calculation was made to determine the sensitivity of the CXO with ACIS to stray light. First, it was important to protect the CCDs from receiving light scattered in the telescope structure and from the background sky, since the CCDs are extremely sensitive to visible light. Secondly, light from the object under study must also be rejected. In the case of a bright star or planet this is a challenging requirement because the thickness of the filter for visible and near IR light also affects the transmission of the filter for low energy X-rays. The filter choices are given above. The aluminum was coated on both sides of the polyimide for electrical grounding, with 30nm on one side and the remainder on the other. The measured transmission curves are given in the ACIS Calibration Report. The limiting magnitudes for various configurations of ACIS and are given in Table 3.1. It is worth noting that the effects of light from bright objects can be circumvented, somewhat, by using a 5x5 island (very faint mode) for each event and then using the outer 16 pixels of the island to recompute a bias for the event. In addition, the event threshold must be increased so that any bias offset created by the visible light does not trigger the event recognition algorithm. As an example, for Jupiter at opposition, the event threshold should be increase from 20 ADU to 53 ADU and the split threshold from 13 ADU to 46 ADU to prevent every pixel over the disk of the planet from triggering the threshold for the BI CCD. In the initial observations of Jupiter, the event threshold was not raised, and every pixel illuminated by the disk of Jupiter produced an event. Since the on board processor was set to reject all events in which all of the pixels in the 3x3 island was above either the event or split thresholds, most of the events, including the X-ray events, were discarded on orbit. By inserting the MEG/HEG an optical attenuation of about 2.6 magnitudes is introduced.

Table 3.1

The magnitude for which one electron is generated per 3.24 s exposure from a point source. It takes 13 electrons to exceed the split event threshold and 20 to exceed the event threshold for the BI CCD.

Stellar Temperature (K)	BI Chip in S-Array (V-magnitude)	FI Chip in I-Array (V-magnitude)
3000	8.5	3.5
4000	8.3	3.2
5000	7.7	2.6
6500	7.3	2.2
10000	6.8	1.6
20000	6.7	1.4

3.1 On-Orbit anomaly

The ACIS was launched with the S3 CCD at the aimpoint of the telescope in the launch-lock position. Initial performance of the detectors determined by the Fe[55] in the ACIS door, was found to be consistent with the ground calibration. After the ACIS door was opened on day 220 of 1999, and the mirror covers were opened on day 225, the ACIS S3 CCD was used to obtain the focus of the SIM and boresight information on the relative alignment of the SIM and the star tracker. Then a number of observations were conducted using both the ACIS and the HRC until day 252, when the ACIS was placed under the calibration source to record calibration data. The data showed that there had been a *very* large increase in the FI CCD charge transfer inefficiency (CTI)! Such a change was totally unexpected, based on all earlier analyses of radiation damage. A number of emergency meetings were held to determine what the cause of such a large increase might be. At a meeting held at Lincoln Laboratory on 09/23/99, a nuclear physicist suggested that the most likely cause of the damage was 100 keV protons from the outer trapped radiation belt forward scattering (Rutherford scattering) off of the very smooth iridium coating on the X-ray mirrors. These protons would not be deflected significantly by the magnets and would reach the CCDs with nearly their original energies. Given the high flux of these protons, a significant dose to each CCD could be expected. By using the Science Instrument Module (SIM) translation to place the ACIS behind a shield for each radiation belt passage, the CCDs could be protected. Subsequent analyses have confirmed this suggestion and the protection procedure has worked very well since it was implemented.
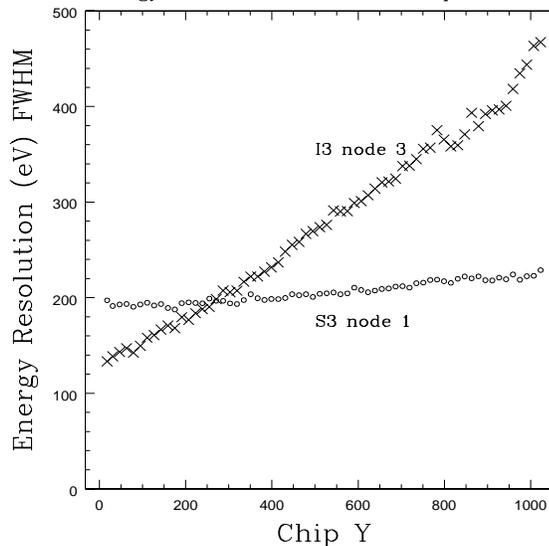
.
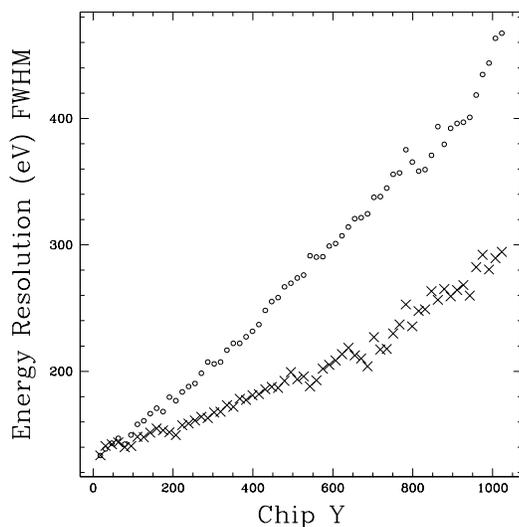
Figure 3.1.1

Figure 3.1.2

The radiation damage was confined to the FI CCDs, since the BI CCDs have about 40 microns of silicon over the buried channels to protect them from the low energy protons. The effect of the radiation damage was to increase the CTI from ~$<10^{-6}$ e/transfer to about 2 x 10^{-4} e /transfer. The impact of this very large increases was to make the energy output and energy resolution of the CCD highly position dependent. As an example, in Figure 3.1.1 the output of I3, node 3 (containing the aimpoint for the I-array) is shown versus the chip y-position for 5.9 keV X-rays. For the best spectroscopy if imaging is not of top priority, place the target near chip y =0. The aimpoint on I3 is at y=964 and at y=510 on S3. The energy resolution of the two amplifiers and chip nodes containing the aimpoints are compared in Figure 3.1.2 for 5.9 keV X-rays. All of these data were obtained at a focal plane temperature of -120C. The effects of the CTI increase for the FI CCDs and for the CTI of the BI chips can be reduced by ground processing and calibrations of the variations in the CCD properties (see [3] and [4] and Figure 3.1.1).

3.2 Stability of ACIS on-orbit

The properties of the ACIS camera have been relatively stable during the 2.5 years since the focal plane was cooled to -120C. There has been a gradual decrease in the gain, and an increase in the energy resolution, as shown in Figures 3.2.1 and 3.2.2. The degree of change is about what was expected prior to launch. Because the focal plane is operated at -120C and the optical blocking filter is at about -60C it was anticipated that there might be a slow accumulation of condensation on either or both components, although there is only a very small path from the main cavity of the telescope to the focal plane. A provision was made to be able to turn on heaters that can raise either the focal plane temperature to +30C or the camera body and filter to a temperature of +25C, or both simultaneously. An observation of a pulsar showed that the flux was less than expected (Pavlov, private communication). A study by Catherine Grant of the onboard calibration source revealed that the Mn and Fe L lines were decreasing in intensity relative to the Mn K alpha line at a steady rate ever since launch (http://space.mit.edu/~cgrant/lkratio_s3.ps). The rate of buildup required to reproduce the observed rate of change of the attenuation was about an order of magnitude more than expected from models of the spacecraft out-gassing. An empirical model of the deposition on the filter has been created by George Chartas and incorporated into XSPEC by Keith Arnaud as *acisabs* (http://asc.harvard.edu/cgi-gen/cont-soft/soft-list.cgi).
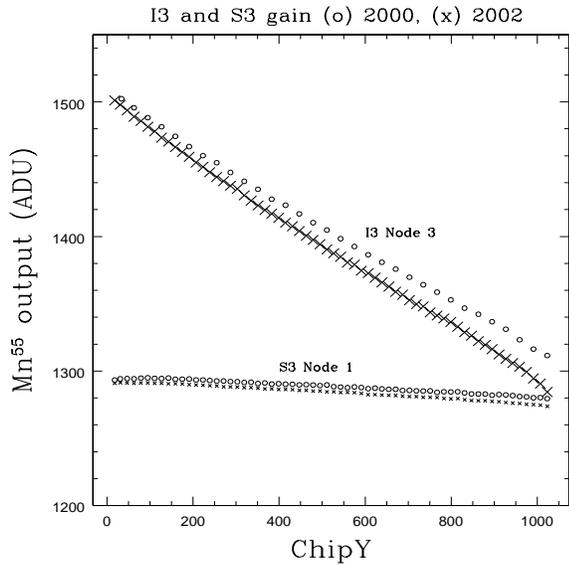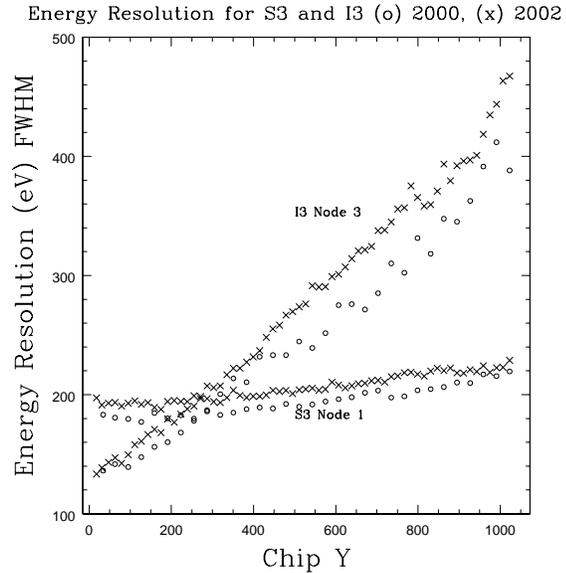
Figure 3.2.1

Figure 3.2.2

3.3 The ACIS Camera
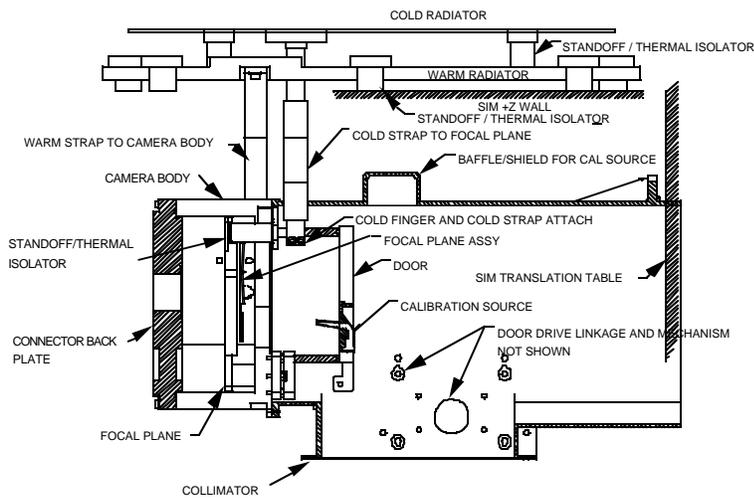
Figure 3.3.1

The ACIS camera is shown schematically in Figure 3.3.1. The collimator of the camera is made from titanium for its strength and thermal properties. It is coated with gold to control its emissivity in the infrared as well as to reduce its contributions to the X-ray background generated on orbit. Aluminum was used in the body of the camera for shielding. The close proximity of the electronics boxes to the camera also provides shielding of the CCDs against radiation damage from cosmic rays and particles trapped in the Earth's magnetic field. The cooling of the focal plane and camera body is accomplished by two large radiators that were protected from any direct sunlight (Figure 3.3.2). The camera body is held at -60C, while the focal plane is held at -120C.

A thermostat can change the focal plane temperature and hold it at an assigned point. A heater is provided to protect the camera and electronics from cooling excessively in eclipse conditions, or to heat the camera and focal plane to bake out any condensations that might build up over time.

.
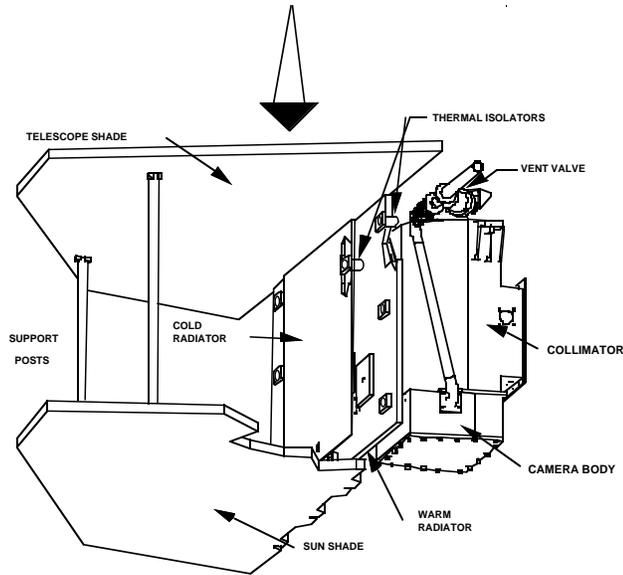
Figure 3.3.2

A vacuum door is provided to protect the filter and focal plane during ground handling and during launch into orbit. Built into the door is an Fe[55] calibration source and a light emitting diode to monitor the CCDs, optical blocking filter and electronic system during ground testing. The door opening and closing is controlled by a wax actuator. During the observatory level thermal vacuum test at TRW the door failed to open. A very thorough investigation by the group at Lockheed Martin who designed and built the door as well as MSFC and several outside experts revealed no flaw in the design or operation of the door. A potentiometer was added to track the angle of opening of the door, and an opening procedure was devised to protect the actuator from failure, should the door stick during the orbital opening procedure. The door was tested at TRW in the instrument module prior to shipment to the Kennedy Space Center, with no anomalous behavior. Needless to say the door opened as planned in orbit

In addition to the door calibration source there are three sources mounted in the Science Instrument Module (SIM) that are directed toward the ACIS camera when the camera is in an offline position (the High Resolution Camera Spectroscopic array is at the aimpoint in this position). The calibration sources utilize the radioactive isotope Fe[55]. The characteristics are given in table 3.3.1. When ACIS is at the aimpoint, no radiation from these sources can reach any detector. Normally, ACIS will transmit the calibration data through the housekeeping telemetry, but since the radiation damage event, the data have been transmitted for each orbit with the full telemetry for typically 30 ks.

Table 3.3.1

Serial No.	55FE Activity mCi (date measured)	Target Material	Principal Line (Energy)	Flux ct s-1 CCD-1 (August 1998)	Other lines (keV)
SN301	0.055 (8/96)	none	Mn K (5.9, 6.4)	45	Mn L (.64-.67), Au M (2.2)
SN802	11 (5/97)	Ti	Ti K (4.5, 4.9)	25 (est)	Mn K (6.4)
SN702	93 (5/97)	Al	Al K (1.5)	22 (est)	Mn K (6.4)

3.4 ACIS Electronics

Figure 3.4.1

The ACIS electronics system, illustrated in figure 3.4.1, is designed to readout the CCDs under programmable computer control, to identify X-ray events in the pattern of signals generated, and to send the information to the spacecraft telemetry. The *Detector Electronics Assembly* (DEA) consists of ten independent, identical subsystems, each controlling a single CCD, and a pair of redundant 'common' sections. The common section is responsible for controlling focal plane functions and for passing commands and housekeeping data between the other DEA subsystems and the Back End Processor (BEP) of the *Digital Processor Assembly* (DPA). There are a total of 40 analog processing chains within the DEA, one for each of the 4 output nodes of each CCD, but no more than 24 chains can be in use during an observing run. The chains are multiplexed to six front-end processing chains, one for each DPA Front End Processor (FEP). Each DEA subsystem comprises a 'driver' section and a 'video' section. To process data from a CCD requires the interaction of three DEA sections — driver, video, and common — in conjunction with a FEP and a BEP from the DPA.
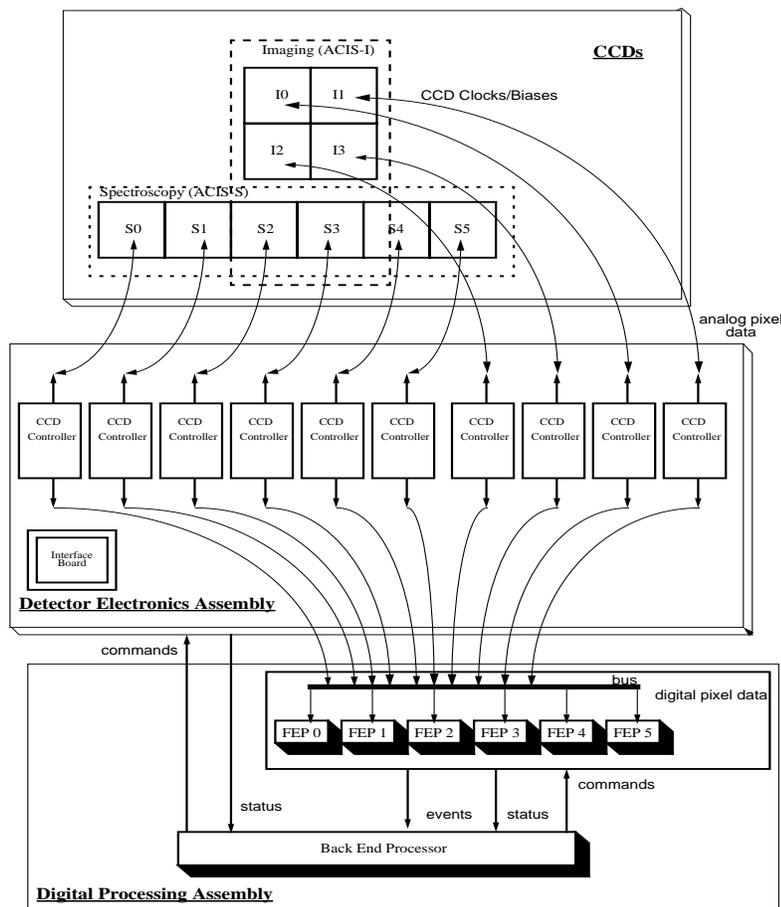
The *Driver Section* provides signal levels, conditions the clock signals delivered to the CCD, and controls the CCD bias voltages. Within the driver section, a Field Programmable Gate Array (FPGA) chip receives commands from the common section. After they are decoded, control signals are generated and distributed throughout the driver section to execute various functions necessary for proper CCD Driver operation. The FPGA also controls the Digital to Analog (DAC) banks, which determine the CCD clock and bias voltages. The Driver section can also be commanded to report several housekeeping channels—voltages and temperatures that characterize the overall functional health of the Section.

The *Video Section* is responsible for amplifying, sampling and converting analog video data from the CCD into a 12-bit-per-pixel digital stream that is sent to a selected Front End Processor (FEP) within the DPA. This section contains four independent video chains that process the four nodes of its CCD simultaneously. Each chain is given its own DC voltage offset, as determined by a set of DAC banks which are set by the FPGA chip under BEP control. The video pre-amplifiers have two gain settings: 1 ADU/e⁻ or 0.25 ADU/e⁻. The conditioning network that powers the video board has the ability to shut down its rail potentials when excessive current is being drawn. This provides a safety feature to prevent overcurrent and failure resulting from a component "latch-up". Depending upon the readout mode specified upon initialization, the power on the video chains A and C or B and D can be independently shut down to conserve power.

The *Common Section* is responsible for receiving commands and clock-sequencing parameters from the BEP, dispatching them to the various DEA subsystems, and collecting and reporting housekeeping data to the BEP.
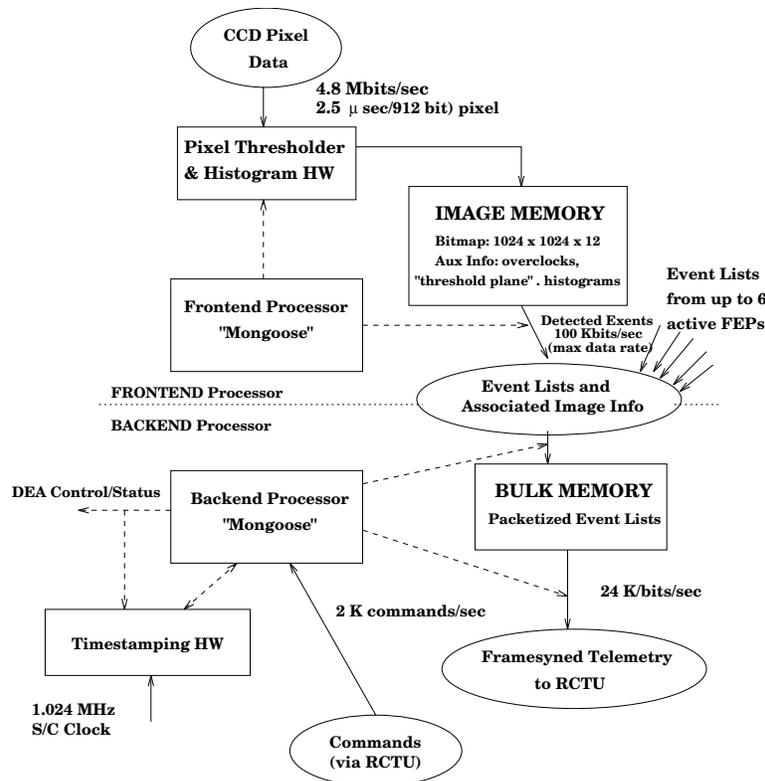
.

Figure 3.4.2

The *Digital Processor Assembly* (DPA) is a computer-based on-board data system which receives data from the DEA and commands from the Remote Command and Telemetry Unit (RCTU), processes data according to the commanded mode, and submits the processed data back to the RCTU for eventual telemetry to the ground. The DPA is customized to meet the unique requirements of the CCD data acquisition and reduction tasks, and to achieve satisfactory compression of the incident data stream to meet the limitations of the telemetry capacity. The DPA is capable of receiving 28 Mbits/s of raw pixel data, corresponding to an exposure rate of 3.24 seconds per frame for six simultaneously active CCDs, and, after event detection and filtering, reduce the average data rate to the allocated telemetry bandwidth of 24 kb/s.

The DPA is composed of two types of processing units; a *Back End Processor* (BEP) which oversees supervisory tasks including uplink, downlink and CCD control, and also performs the X-ray event filtering and packetizing; and a set of *Front End Processors* (FEP) which ingest CCD data on a pixel-by-pixel basis, detect candidate X-ray events, and pass them on to the BEP. There are two redundant BEPs in the DPA, each cross-strapped to the other, but only one is powered up at any one time. There are six FEPs, each listening to a CCD through its DEA subsystem. Should an FEP processor fail, ACIS could only process data from five CCDs, but they can be *any* of the ten CCDs in the flight array, as selected by ground command. The BEPs and FEPs have identical 32-bit processors: the MIPS3000 RISC CPU as implemented by LSI, commonly called the 'Mongoose'. This is a radiation resistant (rad-hard) version of LSI's Logic LR33000 (Self-Embedding Processor) adapted to the special requirements of space-systems electronics. It includes an interrupt controller, a DMA controller, timer/counters, a serial port/debug interface, and cache control. The BEP has a 512 kByte RAM instruction cache and 256 kByte data cache, each with 35 ns access time to support the Mongoose running at 10 MHz. The RAM is extremely rad-hard and immune to single event upsets. Each BEP is booted from a ROM which contains non-volatile flight software for BEP and FEPs. Along with a Mongoose processor, each FEP also contains a pair of FPGAs that assist in the task of pixel processing and X-ray event location.

3.5 Operations

The CCD is a powerful tool for both simultaneously recording images and spectral information. The large number of pixels in each CCD preclude the transmission of all of the pixel information to the ground for processing. X-ray event recognition must be carried out on board and only valid X-ray events passed on to the telemetry. Cosmic rays are rejected on board at this time as well, by discarding certain patterns of charge associated with each event, known as its

grade. The observer is permitted to craft the observation using several options for exposure time and event recognition, as well as which CCDs are included in the observation. From one to six CCDs may be chosen from any of the 10 CCDs comprising the focal plane array. The typical 3.24 second exposure time for a full frame readout has the consequence that a bright source will place several events in a single pixel per readout — a process called 'pileup'. Pileup creates a situation in which the energy of the event becomes unknown and the grades of the events are lost. For extreme cases, the events look like cosmic ray events and are rejected on board, producing an image that has a 'hole' in its center. It is best to keep the pileup below 5% to prevent distortions in the spectrum.

ACIS operates in one of two standard modes. In *timed-exposure mode*, it integrates over a fixed time interval, the *frame time*, and then transfers charge from the image portion of the CCD to the framestore. This transfer takes 41 ms, after which the framestore is read out to determine where each X-ray has interacted in the CCD, while the next exposure is being recorded on the imaging portion of the CCD. Since it takes approximately 3.2 seconds to read the entire framestore, typical full-frame exposures last 3.2 4seconds and the frame-to-frame time is 3.241 seconds. Shortened exposure times may be achieved by discarding part of the framestore, *i.e.,* processing only a *sub-frame*, or by intentionally *pre-flushing* and discarding the contents of the image area after reading the framestore. Of the two choices, *sub-frame* readout is usually to be preferred to *pre-flushing* since the former optimizes detector efficiency at the expense of a smaller detector area. Exposure times can be commanded from 0.1 to 10.0 seconds, in units of 0.1 seconds.

In *continuous-clocking mode*, single rows of pixels are read from the framestore, after which the charge in the image and framestore is moved down by one row. The entire process takes 2.85 ms, which is therefore the effective exposure time in this mode for an on-axis point source. This method of clocking the CCD has the effect of collapsing the 1024x1024 image into a 1x1024 image. All of the events in a 1024 column are added together, which greatly increases the background, especially in the case where the point source is surrounded by a supernova remnant or if another source should fall on the same column as the target.

An observing run begins when the instrument is commanded to power-up a set of DEA sub-systems and FEPs. If necessary, the BEP will load flight software into the FEPs. Then a parameter block is loaded into the BEP and the BEP is told to start the run. It begins by configuring the DEA hardware and by generating and loading the DEA microcode that will be responsible for sequencing the CCDs. The BEP then determines whether the FEP bias maps must be re-calculated, which is necessary if a FEP is being powered up, and very advisable after perigee passages and if the existing bias map was calculated for a different instrument mode.

Each FEP contains a block of RAM addressed as a 1024x1024x12 bit array known as a *bias map*. Each pixel of each CCD generates a certain charge when it is read out, even when no X-ray has generated any charge in the pixel. In order to determine whether an X-ray has contributed any charge to a pixel, it is necessary to know what its 'dark' level is. The bias map is this array of dark levels and must be subtracted from the incoming data pixels prior to any further processing. Normally, a new bias must be created for each CCD after a change of observing mode, perigee passage and radiation induced shutdown. In addition to the stored bias maps, a bad pixel and bad column map is also stored to remove these artifacts from the data that is telemetered. The FEP software contains several choices of algorithm for computing bias maps. The preferred method for timed-exposure maps is known as the 'full-frame' method and proceeds in three stages.

1. A series of M conditioning frames are examined, and the minimum value of each CCD pixel, p_o, is chosen as the zeroth order approximation, b_o, of that pixel's true bias.

2. This process will emphasize the presence of any anomalously low pixel values, which are occasionally generated by the DEA video amplifiers when recovering from being saturated by very high values of readout charge. To remove these corrupted b_o values, a median filter can be applied at this time. The filter examines each value in the approximate bias map. Any values that are lower than 7 of their 8 neighbors by more than a specified value are replaced by the median of those neighbors. This filter has been employed while computing most timed-exposure bias maps since June 2001.

3. Finally, a series of N frames is examined. A pixel that is more than a prescribed value above its corresponding approximate bias value is deemed an X-ray or background event and it and its 8 neighbor pixels are ignored. The remainder are used to improve the bias value itself. For example, if at the nth accumulation frame, the pixel $p_{n,}$ is less than a fixed threshold value above the approximate bias value b_n, and if p_n's neighbors are similarly constrained with respect to their own bias values, the new 'better' bias value is approximated by $b_{n+1} = [p_n + (n-1)b_n]/n$ $n=1,N$. If no p_n is rejected and integer roundoff is neglected, b_{n+1} will converge to the mean of the p's.

.

This algorithm is fast since only *M+N* frames need to be taken. Its results are sensitive to the medium-energy tail of the pixel distribution and the median filter is not always able to recognize and alleviate the effect of anomalously low pixel values (see the URL ftp://acis.mit.edu/pub/pixanom1.pdf for details). For acceptable bias maps, *M* is typically 5–10 and *N* is 10–30. Several other algorithms have been coded into the flight software which may, in some circumstances, perform better than the full-frame method, but they have the disadvantage of requiring N^2 exposure frames, where *N* is typically the same as in the full-frame case.

Continuous-clocking mode bias maps are more easily computed since they are 1-dimensional. A sample of 1024 rows of pixels are read into the FEP, which then examines each column. Up to the present time, the column bias values have been calculated by first determining the mean and variance of the 1024 samples, then rejecting those lying more than 5-σ from the mean, and taking the mean of the remainder to be the bias for that column. This algorithm has proven unsatisfactory at times of high background radiation, since it is sensitive to the extensive charge 'blooms' that occur within front-illuminated CCDs. A second method is available within the FEP — to choose the column bias as the *n*th quartile of the 1024 samples — and this may be employed in future (see the URL ftp://acis.mit.edu/pub/ccmode.pdf for details).

Once the FEP bias maps have been calculated, the BEP halts the CCD clocks and writes the bias maps into the downlink telemetry, compressing them using Huffman tables. The BEP then reloads and restarts the CCDs, commanding the FEPs to process the pixel streams in one of several possible modes, as outlined in the following table.

Table 3.5.1

Sub-Mode	Timed Exposure Mode		Continuous Clocking Mode	
	Data element	Max rate	Data element	Max rate
Event	**Faint**: 3x3 pixel arrays **Very Faint**: 5x5 pixel arrays **Graded**: 3x3 arrays, summed	170 /sec 68 /sec 375 /sec	**Faint 1x3**: 1x3 pixel arrays **Faint**: 3x3 pixel arrays **Graded 1x3**: 1x3 arrays, summed **Graded**: 3x3 arrays, summed	400 /sec 170 /sec 638 /sec 375 /sec
Histogram	**Raw**: raw pixel values **Event**: 3x3 arrays, summed	varies varies	not available	
Raw	1024 pixels + overclocks, compressed	varies	1024 pixels + overclocks, compressed	Varies

In an *event* mode, each FEP employs its FPGAs to examine the incoming pixel values, comparing them against the corresponding bias map values. If the pixel value is higher by more than a pre-determined *event threshold*, typically 38 ADU for front-illuminated CCDs and 20 ADU for back-illuminated, the pixel is marked for further examination by the FEP's Mongoose processor. This involves checking whether, after bias subtraction, it is a local maximum within its 3x3 neighborhood (or 3x1 in the case of those continuous-clocking sub-modes). When a FEP finds a local maximum, it reports the pixel and bias elements (as described in the *Data Element* columns of table 3.5.1) to the BEP. The BEP, in turn, collects these candidate X-ray events, determines their summed charge, filters them according to energy, grade code, and row and column location within the CCD, and writes those that survive to the telemetry system. Event energy is determined by subtracting the 3x3 (or 1x3) array of bias values from their corresponding pixel value, and summing those that exceed a pre-determined *split-threshold*. The grade code is determined by the geometric arrangement of pixels that exceed the split threshold criterion, as described the Chandra Proposer's Guide at http://cxc.harvard.edu (called event grades in that document). Acceptance or rejection by CCD row and column is determined by a set of *window blocks* that are loaded into the BEP before the start of the observing run.

As shown in table 3.5.1, *very faint mode*, which reports a 5x5 pixel island for each event, is the most demanding of telemetry — only 68 events/s (including a 10% overhead for housekeeping data) can be recorded without dropping data. The mode is most useful for faint source detection in deep exposures. The Chandra X-ray Center provides software to use the 5x5 pixel island to identify and reject background events. Normally, most of the cosmic rays are rejected on-orbit by excluding the events in which all 9 of the 3x3 pixel island contains charge above the split event threshold. The extra rejection from the 5x5 pixel analysis gives some added rejection below about 2 keV. As mentioned above in Section 3, the 5x5 pixel island can be used to determine a local bias that can be used to discriminate against visible light photons from an optically bright source.

.

For brighter sources, very faint mode may saturate the telemetry bandwidth. If this should occur, the BEP will signal the FEPs to begin ignoring exposures (or groups of 512 rows in continuous clocking mode) until buffer space becomes available. The result is that ACIS will only drop whole frames, never partial ones, making it an easy matter to compute the total exposure for each CCD. Since observing efficiency is reduced when exposures are dropped, brighter sources are best observed in *faint mode* or, for very bright sources, *graded mode*. In this latter mode, only the CCD location, the summed event energy and the grade code are reported.

Timed-exposure event modes may be further constrained by additional parameters. It is possible to specify an *alternating exposure mode* in which a very short exposure is followed by a number (up to 15) full-length exposures. This mode has the advantage of reducing pileup for a bright source in the short exposures, while collecting data on fainter sources in the field. A *sub-array readout mode* has been used extensively to reduce exposure time and pileup, but which also reduces observing efficiency, as explained above. The farther the region is from the output node the more time must be allowed to move the charge to the output node for processing, and the frame transfers must be staggered to keep the peak power within limits. The minimum frame time for sub-array readout of n rows from m CCDs is approximately $41 (m-1) + 2.85 n + 0.04 (m+q)$ milliseconds, where q is the number of rows to be skipped between the output node and the start of the sub-array. The actual integration time will be rounded up to the next 0.1 s. Finally, special CCD microcode can be employed to perform *on-chip charge summation*. In timed exposure modes, this causes pairs of rows and pairs of columns to be treated as one, *i.e.*, the 1024x1024 pixel CCD becomes a 512x512 pixel array, spatial resolution is halved, framestore readout time is reduced by a factor of four, and exposure times can be reduced without loss of efficiency. In continuous clocking modes, charge summation can be applied to pairs of columns (thereby halving the row readout time) and to groups of 2^n rows, $0 < n \leq 5$. In practice, these summed modes are not used for science observations since the opportunity has not arisen to calibrate the performance of the CCDs with the special microcode.

In addition to the event modes, there are two varieties of timed-exposure histogram mode. In *raw-histogram* mode, the FEPs pass raw pixel values to the BEP, which constructs a histogram of the values for each output node of each active CCD, and writes the histograms to telemetry after a predetermined number of exposures have been accumulated. In *event-histogram* mode, the FEP detects 3x3 event candidates, the BEP converts them to ADU values as in timed-exposure graded mode, filters them, and then constructs and reports histograms of the filtered energy values.

In *raw* mode, the FEP merely reads the raw pixel values from the DEA and sends them to the BEP, which compresses them using a Huffman table and writes them to the telemetry system. With the downlink bandwidth limited to 24 kbps, the BEP's output buffers are soon saturated and the FEPs are commanded to ignore subsequent frames until space becomes available. Typically, only one in 50 exposure frames is reported — proportionally fewer if more than one CCD is clocked simultaneously.

Finally, several specialized modes have been developed for calibration. They are not supported by the flight software but must, instead, be loaded whenever required as a series of DPA patches and customized DEA microcode. In *staggered readout mode* the integration region of the CCD is only partially transferred to the framestore area. A specified numbers of rows are transferred, then a new integration begins. This process is repeated producing a series of vertically displaced images. The advantage is that the observing efficiency is high, but the background is also increased, since each image is exposed to the sky for a time of about 3.24 s. The minimum exposure in this mode is 59.13 ms in which 18 rows of the CCD are used as the frame. In *squeegee mode*, a row of background charge is retained in the image area of each CCD and is run back and forth before every timed exposure so as to fill the charge traps that are thought to cause the increased CTI of front-illuminated CCDs, as described in section 3.1. Last, a mode known as *cuckoo clocking* permits a small number of columns of raw pixel values to be downlinked without dropping exposures.

4. Examples of scientific results

The CXO has been making observations for the past three years, interrupted only by ~8 hour periods once per orbit when the observatory comes within ~60,000 km of Earth, and during occasional high radiation events emitted by the active Sun, which has just gone through solar maximum. The detector backgrounds are actually lower during solar maximum than during solar minimum, since the magnetic fields and solar wind sweep a larger fraction of the background-generating. low-energy galactic cosmic rays out of the inner solar system during this time. The examples reported in this section were obtained from the ACIS Guaranteed Observer Time and serve to illustrate the wide range of science that can be accomplished using ACIS. It is possible in this limited space to present only a very small number of examples of the great science that is being done with *Chandra*. The great improvement in angular resolution and aspect

.

reconstruction of the CXO over previous observatories has opened entirely new fields of study in high-energy astrophysics.

4.1 The *Chandra* Deep Field North

The CXO has observed the region of the sky around the Hubble Deep Field North accumulating a total exposure of 2 Ms, 500 ks of which was from GTO time and the remainder from General Observer time, lead by W. N. Brandt. Near the aimpoint the survey reaches 0.5 – 2.0 and 2.0 – 8.0 keV limiting fluxes of ~2 x 10^{-17} and ~1 x 10^{-16} ergs cm^{-2} s^{-1}, respectively. About a dozen papers have been published on portions of these data, with a point source catalog paper containing over 500 X-ray sources [5]. The main conclusions from this large study are: as much as 90% of the "diffuse" extragalactic X-ray background discovered by the first rocket flight to detect extra solar X-rays [9] originates in discrete sources, probably massive black holes at high redshifts (1 – 3); the combined spectrum of the sources matches the previously measured "diffuse" extragalactic background to good accuracy (the normalization of the previous measurements vary by about 30% and represent the greatest uncertainty for the comparison); 25 percent of the optical counterparts must be fainter than R magnitude 27.7 [10]; the luminosity of Lyman break galaxies (2<z<4) obtained from "stacking" analysis gives an average luminosity of 3.2 x 10^{41} ergs s^{-1}, comparable to the most X-ray luminous starburst galaxies in the local Universe [11]. The stacking analysis results in an effective exposure of 22.4 Ms and is only effective because of the high angular resolution of CXO and the low ACIS background. Spiral galaxies in the range of 0.4<z<1.5 show that their ratio of X-ray luminosity to B-band luminosity remains essentially constant [12]; the number of extended X-ray emission regions is about what is expected based on earlier work, with the bulk of the luminosity from this class coming from clusters brighter than 10^{-14} ergs cm^{-2} s^{-1} [13]; six out of eight bright SCUBA sources have been detected as X-ray sources within an 8.4' x 8.4' area of the CDF-N, with four of the X-ray sources possessing very flat x-ray spectra [14]. This region of the sky continues to be a rich source of investigation. The GOODs survey using the HST will greatly increase our ability to study the morphology of the X-ray emitting galaxies.

4.2 SgrA* and surrounding region

For the first time it has been possible to resolve the X-ray source complex at the Galactic Center, where there are at least three sources in a small region of 7 x 12 arc seconds. The source associated with SgrA* is quite faint normally with an absorption corrected luminosity of only 2 x 10^{33} ergs s^{-1} in the 2.0 -10.0 keV band, an extremely low flux for a 2.6 million solar mass black hole. A large flare was detected on 27 October 2001 when the flux increased by a factor of about fifty at peak, and which lasted for about 3 hours [15]. The bright X-ray enhancement coincident with the radio source Sgr A East, a supernova remnant (SNR), was also detected. Based on the X-ray data Maeda [16] suggested that the supernova blast wave may have passed SgrA* some 300 years ago and evacuated the region around the massive black hole leading to its current low luminosity. The elemental abundance found in the SNR is about four times solar with the iron concentrated in a more compact region than the other lighter elements.

A recent series of observations of the SgrA* region have been completed and will be reported at the SgrA* Workshop to be held in November. Over 2000 X-ray sources were detected in the combined data. The entire 17 x 17 arc minute field surrounding SgrA* is covered with emission from a plasma with a temperature of nearly 100 MK.

4.3 The Orion Trapezium star cluster

The Orion Trapezium star cluster has been observed for 23 hours by the CXO. A total of 1075 X-ray sources were detected, 91% of which have optical counterparts and 7% are new sources deeply embedded in the dense molecular cloud.   Some of the discoveries include: rapid variability of the O9.5 31 solar mass star theta 2A Orionis and several B stars, which is inconsistent with the standard model of the X-rays originating from wind shocks; support for the hypothesis that intermediate-mass mid-B through A type stars do not emit significant X-ray emission; confirmation that low-mass G through M-type T Tauri stars exhibit powerful flaring but typically at luminosities considerably below the 'saturation' level; confirmation that the presence or absence of a circumstellar disk has no discernable effect on X-ray emission; evidence that T Tauri plasma temperatures are often very high with T > 100MK, even when luminosities are modest and flaring is not evident; and detection of the largest sample of pre-main sequence very low mass objects showing high flaring levels and a decline in magnetic activity as they evolve into L- and T-type brown dwarfs [17] [18] [19].

4.4 Supernova Remnants and Pulsars

The high angular resolution of the CXO has allowed unprecedented studies of supernova remnants (SNRs) and pulsars which are often found in or near the centers of SNRs. The monitoring of SN1987A has revealed a gradual increase in its

.

angular diameter and its intensity over the past three years. As noted in the introduction, only the CXO has sufficient angular resolution to measure the angular diameter of this young remnant [20]. The intensity increase appears to be faster than linear.

Observations of the nebula surrounding the Vela pulsar have discovered rapid changes in the morphology of the remnant over periods of months. Very interesting features appear, move and fade. The pulsar must be the driving force behind these dynamic features, but the details of the mechanism are not understood [21]. The Crab nebula also shows dramatic variable features that move with apparent velocities of 0.4c [22].

The central compact object (CCO) of RCW103 has been observed on nine different occasions by the CXO. Sometime between 26 September 1999 and 8 February 2000 the CCO underwent an outburst, increasing in luminosity by at least a factor of fifty. Over the next twenty-six months the CCO has remained at a level ten times the initial flux. The flux exhibits a complex, variable modulation pattern with a period of about 6.4 hours [23] [24], showing clear evidence that the CCO is in a binary system – one of the very few such objects in a SNR. The very good positional accuracy of the CXO placed the source at RA(J2000)= 16^h 17^m 36.2^s Decl=-51^o 02' 25.0" with an uncertainty of 0.7", which corresponds to a faint extremely-red counterpart with I>25, J~22.3, H~19.6, and Ks~18.5. The source spectrum is well fit by an absorbed black body with a temperature of ~0.55 keV for all of the observations except for the maximum flux point which had a temperature of ~0.43 keV (errors of 0.02 keV).

The oxygen-rich SNR G292.0+1.8 was observed by the CXO for 40 ks on 11 March 2000. The image is quite striking for its complex structure [25]. A 135 ms radio pulsar has been found at the position of a bright point X-ray source near the center of the remnant which is surrounded by a probable synchrotron nebula with a power law X-ray spectrum [26] [27] An X-ray spectral analysis using the ACIS S3 CCD reveals complex variations in the elemental composition of the clouds of material in this remnant.

4.5 X-ray Sources in Galaxies

Perhaps one of the most dramatic new fields of study opened up by the CXO is the study of the X-ray source populations in external galaxies. The high angular resolution of the X-ray mirrors combined with the small pixel size of ACIS and the moderate energy resolution of the CCDs creates a powerful tool to reveal the distribution and types of X-ray sources found in external galaxies. NGC 720 provides a good example where 42 point sources are detected in a 38.8 ks exposure. At a distance of 35 Mpc the source population traces out a Log N – Log L distribution that drops off sharply above 1.5×10^{39} erg s^{-1}. The faintest sources detected are more luminous than an Eddington limited luminosity for a 1.4 solar mass neutron star.[28]. This result is similar to that found for the Circinus Galaxy and M82, but because these galaxies are closer, the LogN-LogL curves can be extended below 10^{38} erg s^{-1} where another break in the spectrum occurs probably associated with the neutron star binaries [29].

4.6 Clusters of Galaxies

 Although clusters of galaxies have been studied extensively by previous X-ray observatories, the high angular resolution with moderate spectral resolution of ACIS/CXO permits the examination of the cores of the clusters in sufficient detail to place limits on the interaction cross section of dark matter that is thought to constitute 90% of the cluster mass [30]. The high positional accuracy together with the high spatial resolution and sensitivity of ACIS/CXO permits the study of QSOs, AGN and submillimeter sources lensed by clusters and galaxies. The magnification factor of the lens enables the study of much less luminous objects of the above classes [31][32].

5   Conclusions

        In this brief overview of the ACIS instrument on board the CXO, we have tried to present the essential features of the instrument and a few of its capabilities. To find out more about the instrument operation and calibration, please refer to the following website; http:// www.astro.psu.edu/xray/acis/technical/tech_documents.html

6   Acknowledgments

.

7 References

1. Weisskopf, M. C., Brinkman, B., Canizares, C., Garmire, G., Murray, S. and Van Speybroeck, L. P.; "An Overview of the Performance and Scientific Return from the *Chandra X-ray Observatory*"; 2002 ASP 114, 1-24

2. Burke, B.E., Gregory, J., Bautz, M., Prigozhin, , Kissel, S., Kosicki, B.,Loomis, A., Young, D.; "Soft X-ray CCD Imagers for AXAF"; IBEE Transactions on Electron Devices, vol. 44, no. 10, pp. 1633-1642 (October 1977)

3. Burrows, D.N., Michael, E., Hwang, U., McCray, R., Chevalier, R. A., Petre, R., Garmire, G.P., Holt, S.S., Nousek, J.A.;"The X-ray Remnant of SN1987A", 2000 ApJ 543, L149

4. Markavitch, M., Bautz, M.W., Biller, B., Butt, Y., Edgar, R., Gaetz, T., Garmire, G., Grant, C., Green, P., Juda, M., Plucinski, P., Schwartz, D., Smith, R., Vikhlinin, A., Virani, S., Wargelin, B.J., Wolk, S.; "*Chandra* Spectra of the soft X-ray diffuse background", 2002 ApJ submitted.

5. Alexander, D.M., the CDF-N Team,"Overview of the *Chandra* Deep Field North Survey"; 2003 X-ray Surveys in the light of new Observatories, Santander Spain

6. Bautz, M. W., Pivovaroff, M., Baganoff, F. Isobe, T., Jones, S., Kissel, S., LaMarr, B., Mannaing, H., Prigozhin, G., Ricker, G., Nousek, J. Grant, C., Nishikida, K., Scholze, F., Thornagel, R., and Ulm, G.;" X-ray CCD Calibration for the AXAF CCD Imaging Spectrometer"; 1998 SPIE 3444, p210

7. Townsley, L. K., Broos, P.S., Nousek, J.A., & Garmire, G.P.;" Modeling Charge Transfer Inefficiency in the *Chandra* Advanced CCD Imaging Spectrometer"; 2002 NIMPA,486**,** p751

8. Townsley, L.K., Broos, P.A., Chartas, G., Moskalenko, E., Nousek, J.A., and Pavlov, G.G.; "Simulating CCDs for the *Chandra* Advanced Imaging Spectrometer"; 2002 NIMPA 486, p716

9. Giacconi, R., Gursky, H., Paolini, F.R., and Rossi, B.B.; 1962, Phys. Rev. Lett., 9, 439

10. Barger, A.J., Cowie, L. L., Brandt, W. N., Capak, P., Garmire, G., Hornschemeier, A. E., Steffen, A. T. & Wehner, E. H. 2002 astro.ph 0206370

11. Brandt, W. M.; Hornschemeier, A. E., Schneider, D.P., Alexander, D. M., Bauer, F. E., Garmire, G. P., & Vignali, C. 2001 ApJ 558, L5

12. Hornschemeier, A. E., Brandt, W.N., Alexander, D. M., Bauer, F. E., Garmire, G. P., Schneider, D. P., Bautz, M.W. & Cartas, G. 2002 ApJ 568, 82

13. Bauer, F. E., Alexander, D. M., Brandt, W. N., Hornschemeier, A. E., Miyaji, T., Garmire, G. P. , Schneider, D. P., Bautz, M. W., Chartas, G., Griffiths, R. E., & Sargent, W. L. W. 2002 AJ 123, 1161

14. Alexander, D.M.; Bauer, F.E.; Brandt, W.N. ; Hornschemeier, A.E. ; Vignali, C. ; Garmire, G.P. ;& Schneider, D.P. ; "The Chandra Deep Field North Survey XIII on the Nature of Bright Scuba Sources"; 2002 ApJ (submitted)

15. Baganoff, F.K., Bautz, M.W., Brandt, W.N., Chartas, G., Feigelson, E.D., Garmire, G.P., Maeda, Y, Morris, M., Ricker, G.R., Townsley, L.K., Walter, F.," Rapid X-ray flaring from the direction of the supermassive black hole at the Galactic Centre"; 2001 Nature, No. 6851

16. Maeda, Y., Baganoff, F.K., Feigelson, E.D., Bautz, M.W., Brandt, W.N., Burrows, D.N., Doty, J.P., Garmire, G.P., Morris, M., Pravdo, S.H., Ricker, G.R., Townsley, L.K.;,"A Chandra Study of SGR A East: A Supernova Remnant Regulating the Activity of our Galactic Center?"; 2002, ApJ 570, 671

17. Feigelson, E.D. , Garmire, G.P. , "Magnetic Flaring in the Pre-Main-Sequence Sun and Implications for the Early Solar System"; 2002 ApJ, 572:335-349

.

18. Garmire, G.P., Feigelson, E.D.; Broos, P. ; Hillenbrand, L.A.; Pravdo, S.H.; Townsley, L; Tsuboi, Y; "Chandra X-ray Observatory Study of the Orion Nebula Cluster and BN/KL Region"; 2000 AJ

19. Feigelson, E. D., Broos, P. Gaffney III, J. A., Garmire, G. P., Hillenbrandt, L. A., Pravdo, S. H., Townsley, L. K., & Tsuboi, Y. 2002 astro.ph/0203316 ApJ (accepted)

20. Park, S., Burrows, D. N., Garmire, G. P. & Nousek, J. N. 2002 ApJ 567, 314

21. Pavlov, G. G., Kargaltsev, O. Y., Sanwal, D., & Garmire, G.P. 2001 ApJ 554, L189

22. Mori, K., Hester, J.J., Burrows, D. N., Pavlov, G. G., & Tsunemi, H. 2002, ASP 9999
        P. O. Slane and B. M. Gaensler, eds

23. Garmire, G.P., Pavlov, G. G., & Garmire, A. B. 2000 IAUC 7350

24. Sanwal, D., Garmire, G. P., Garmire, A., Pavlov, G. G., & Mignani, R. 2002 AAS 200.720

25. Park, S., Roming, P. W. A., Hughes, J. P., Slane, P. O., Burrows, D. N., Garmire, G. P., & Nousek, J. A 2002 ApJ 546, L39

26. Hughes, J. P., Slane, P. O., Burrows, D. N., Garmire, G. P., Nousek, J. A., Olbert, C. M., & Keohane, J. W. 2001 ApJ 559, L153

27. Camilo, F., Manchester, R. N., Gaensler, B. M., Lorimer, D. R. & Sarkissian, J. 2002 ApJ 567, L71

28. Jeltema, T., Canizares, C. R., Buote, D. A., & Garmire, G. P. 2002 ApJ (submitted)

29. Bauer, F. E., Brandt, W. N., Sambruna, R. M. , Chartas, G., Garmire, G. P., Kaspi, S., & Netzer, H. 2001 AJ 122, 182.

30. Arabadjis, J. S., Bautz, M. W., & Garmire, G. P., 2002 ApJ 572, 66

31 Bautz, M. W., Malm, M. R., Baganoff, F.K., Ricker, G. G., Canizares, C. R., Brandt, W. N., Hornschemeier, A. E., & Garmire, G. P., 2000 ApJ 543, L119

32. Chartas, G., Brandt, W. N., Gallagher, S. C., & Garmire, G. P. 2002 astro.ph/7196C

*contact garmire@astro.psu.edu, phone: 1 814 865-1117; fax : 1 814 865-2977; http://www.astro.psu.edu; Pennsylvania State University, 514A Davey Lab, University Park, PA 16802; ** mwb@space.mit.edu; phone: 1 617 253-0023; fax: 1 617 253-0861; Room 37-521 Center for Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139

.