The *Chandra* Proposers' Observatory Guide

Prepared by:
Chandra X-ray Center
Chandra Project Science, MSFC
Chandra IPI Teams

Version 28.0
December 2025

Contents

List of Figures

List of Tables

Chapter 1

Mission Overview

The *Chandra* X-Ray Observatory (*CXO*) combines an efficient high-resolution ($\leq$ 1/2 arcsec) X-ray telescope with a suite of advanced imaging and spectroscopic instruments. The Observatory was successfully launched by the National Aeronautics and Space Administration's (*NASA*'s) Space Shuttle *Columbia* on 1999-Jul-23, with Col. Eileen Collins commanding. Subsequently an Inertial Upper Stage and *Chandra*'s Internal Propulsion System placed the Observatory in a high elliptical orbit. *Chandra* is the X-Ray component of *NASA*'s four Great Observatories. The other components are the *Hubble* Space Telescope, the late *Compton* Gamma-Ray Observatory, and the decommissioned *Spitzer* Space Telescope.

1.1 Program Organization

The *Chandra* Project is managed by *NASA*'s Marshall Space Flight Center (MSFC). The Project Scientist is Dr. Steven Ehlert. Day-to-day responsibility for *Chandra* science operations lies with the *Chandra* X-ray Center (*CXC*), Dr. Pat Slane, Director. The *CXC* is located at the Cambridge, Massachusetts, facilities of the Smithsonian Astrophysical Observatory (*SAO*) and the Massachusetts Institute of Technology (*MIT*). The *Chandra* Operations Control Center (*OCC*) is located in Burlington, Massachusetts. The *CXC* uses the *OCC* to operate the Observatory for *NASA*.

1.2 Unique Capabilities

Chandra was designed to provide order-of-magnitude advances over previous X-ray astronomy missions with regards to spatial and spectral resolution. The High Resolution Mirror Assembly (HRMA) produces images with a half-power diameter (HPD) of the point spread function (PSF) of <0.5 arcsec. Both grating systems—the Low Energy Transmission Grating (LETG) and the High Energy Transmission Grating (HETG)—offer resolving powers

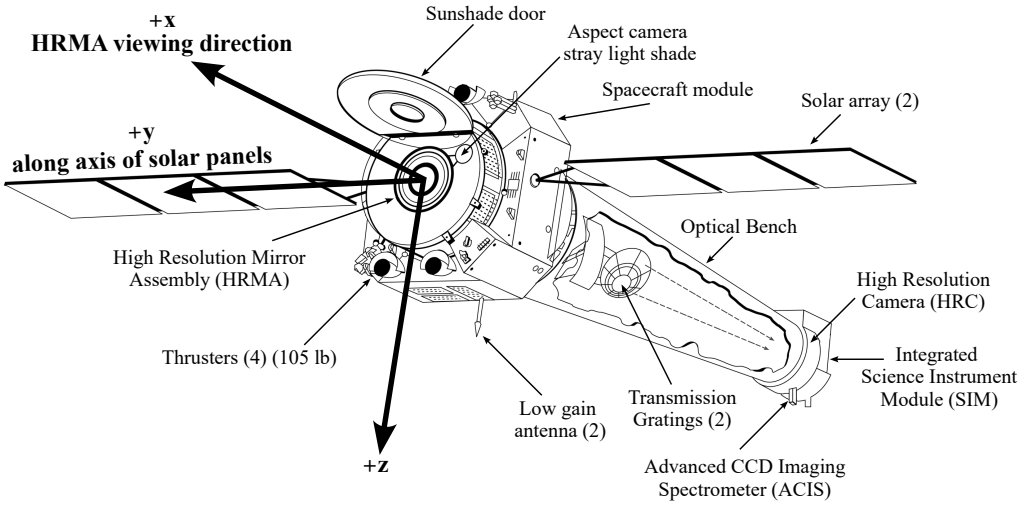

Figure 1.1: The *Chandra* Observatory with certain subsystems labeled.

well in excess of 500 over much of their bandwidth that, together, cover the range from $\leq$0.1 keV to 10 keV.

1.3 Observatory Overview

An outline drawing of the *Chandra* X-ray Observatory is shown in Figure 1.1. *Chandra* consists of a spacecraft and a telescope/science-instrument payload. The spacecraft provides power, communications, command, data management, and pointing control and aspect determination. The principal elements of the observatory that will be discussed in this document are:

- The High Resolution Mirror Assembly (HRMA; Chapter 4)
- The Aspect System (Chapter 5)
- The Focal-plane Science Instruments (SIs):
  - The Advanced CCD Imaging Spectrometer (ACIS; Chapter 6)
  - The High Resolution Camera (HRC; Chapter 7)
- The Objective Transmission Gratings:
  - High Energy Transmission Grating (HETG; Chapter 8)
  - Low Energy Transmission Grating (LETG; Chapter 9)

These and related elements of the *Chandra* Project are introduced briefly in the remainder of this chapter.

1.4 Pointing Control and Aspect Determination (PCAD)

The PCAD system controls the pointing and dithering of the observatory and provides the data from which both the relative and absolute aspect are determined. Dithering is imposed to spread the instantaneous image over many different pixels of the focal-plane detector to smooth out pixel-to-pixel variations. The dither pattern is a Lissajous figure (and can be seen quite clearly in the un-aspect corrected data from bright point sources). The amplitude, phase, and velocity depend on which instrument (ACIS or HRC) is in the focal plane.

Key elements of the PCAD system are the set of redundant gyroscopes, momentum wheels, and an aspect camera assembly (ACA) consisting of a four inch optical telescope with (redundant) CCD detector. The aspect camera simultaneously images a fiducial light pattern produced by light emitting diodes placed around the focal-plane instruments along with the flux from up to five bright stars that may be in the aspect camera's field-of-view. An interesting consequence is that the user may request that one of the targets of the aspect camera be at the location of the X-ray target. For bright optical counterparts, this option allows real-time optical monitoring, albeit at the price of reducing the accuracy of the aspect solution—see Chapter 5 for further details. This option will be implemented only pursuant to a feasibility analysis during the planning and scheduling process.

1.5 HRMA

The HRMA consists of a nested set of four paraboloid-hyperboloid (Wolter-1) grazing-incidence X-ray mirror pairs, with the largest having a diameter of 1.2 m (twice that of the *Einstein* Observatory). The focal length is 10 m.

The mirror glass was obtained from Schott Glaswerke; grinding and polishing was performed at Hughes Danbury Optical Systems; coating at Optical Coating Laboratory; and the mirror alignment and mounting at Eastman Kodak Co. The mirrors weigh about 1000 kg. Details of the HRMA and its performance are presented in Chapter 4.

The *Chandra* Telescope Scientist was the late Dr. Leon Van Speybroeck of the Smithsonian Astrophysical Observatory.

1.6 Science Instrument Module (SIM)

The Science Instrument Module consists of the special hardware that provides mechanical and thermal interfaces to the focal-plane scientific instruments (SIs). The most critical functions from an observer's viewpoint are the capability to adjust the telescope focal length and the ability to move the instruments along an axis orthogonal to the optical axis.

The SIM houses the two focal instruments, the ACIS and the HRC. Each of these have two principal components: ACIS-I and -S and HRC-I and -S, respectively. The focal-plane instrument layout is shown in Figure 1.2. The SIM moves in both the X-axis (focus) and

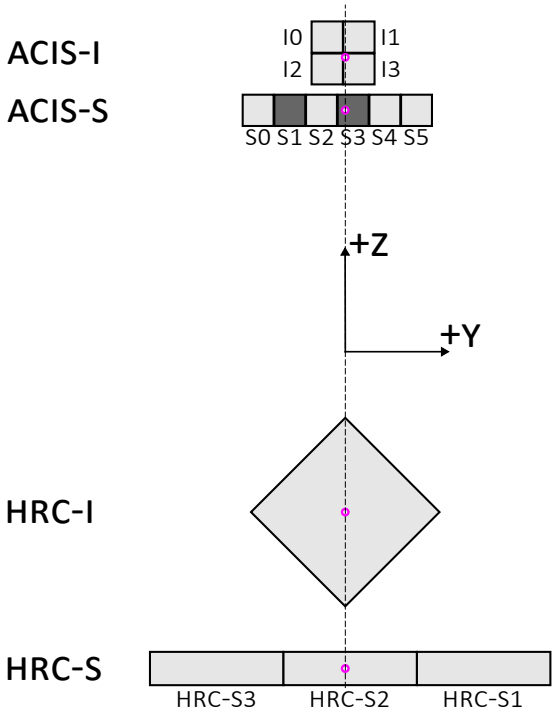

Figure 1.2: Arrangement of the ACIS and the HRC in the focal plane. The view is along the axis of the telescope from the direction of the mirrors. For reference, the two back-illuminated ACIS-S chips are shaded. Numbers indicate positions of chips I0–I3 and S0–S5. SIM motion can be used to place the aimpoint at any point on the vertical dashed line denoting the spacecraft Z axis direction. Nominal detector aimpoints are indicated with small circles.

the Z-axis (instrument and aimpoint (1.6.1) selection). Note that the Y-Axis parallels the dispersion direction of the gratings.

1.6.1 Aimpoints

Aimpoints are the nominal positions on the detector where the flux from a point source is placed. Note there is a slight (less than 20") distinction between the aimpoint and the on-axis position, which for most practical purposes can be ignored. The aimpoints are discussed in detail in the chapters about each instrument and in Chapters 4 and 5.

1.7 Ground System

The ground system consists of the *CXC* in Cambridge, Massachusetts, the *OCC* in Burlington, Massachusetts, the Engineering Support Center (ESC) at *MSFC*, and various *NASA* communications systems including the Deep Space Network operated for *NASA* by the Jet Propulsion Laboratory (JPL). See Section 2.6.2 for details.

1.8 Orbit

The *Chandra* orbit is highly elliptical and varies with time. In 2023-July the perigee altitude, after decreasing since late 2017, reached its mission minimum of 1045 km; the apogee height, after rising since 2017, reached a maximum of ∼148,000 km. The orbital eccentricity accordingly reached its maximum of ∼0.91 in 2023-July. Since then, the eccentricity has decreased, reaching ∼0.84 in mid-2025, at which time the apogee altitude was ∼142,000 km and the perigee altitude ∼6680 km. The trend for decreasing eccentricity and increasing perigee altitude will continue until ∼2029. The orbit allows for high observing efficiency as the satellite spends most of the time well above the radiation belts (∼70%) and long observations (currently ∼180 ksec) are made possible in principle by the orbital period of 63.5h (but see Section 3.3 for limitations due to spacecraft thermal considerations).

1.9 Particle Detector

There is a particle detector mounted near the telescope, called the Electron, Proton, Helium INstrument (EPHIN; see Section 2.5). This detector was used to monitor the local charged particle environment as part of the scheme to protect the focal-plane instruments from particle radiation damage; owing to performance degradation and erratic behavior, EPHIN is no longer used in this protective function as of 2013-Nov and was depowered in 2018-Sep.

The Co-Principal Investigators of the EPHIN instrument are Drs. Reinhold Muller-Mellin and Hoarst Kunow of the University of Kiel, Germany.

1.10 ACIS

The ACIS is composed of two CCD arrays: a 4-chip array, ACIS-I, and a 6-chip array, ACIS-S. The CCDs are flat, but the chips in each array are positioned (tilted) to approximate the relevant focal surface: that of the HRMA for ACIS-I and that of the HETG Rowland circle for ACIS-S. ACIS-I was designed for CCD imaging and spectrometry; ACIS-S can be used both for CCD imaging spectrometry and also for high-resolution spectroscopy in conjunction with the HETG grating.

There are two types of CCD chips. ACIS-I is composed of front-illuminated (FI) CCDs. ACIS-S is composed of 4 FI and 2 back-illuminated (BI) CCDs. The BI S3 chip is

at the best focus position and is normally used for ACIS-S imaging observations. ACIS-I is better when wider field (16 arcmin × 16 arcmin) and/or higher energy response is needed; ACIS-S imaging is better when low energy response is preferred and a smaller (8 arcmin × 8 arcmin) field of view is sufficient.

The efficiency of the ACIS instrument has been discovered to be slowly changing with time, most likely as a result of molecular contamination build-up on the optical blocking filter. The BI CCDs response extends to lower energies than the FI CCDs and the energy resolution is mostly independent of position. The low-energy response of the BI CCDs is partially compromised by the contaminant build-up. The FI CCD response is more efficient at higher energies, but the energy resolution varies with position due to radiation damage caused by protons reflecting through the telescope during radiation-zone passages in the early part of the mission. Details on the ACIS are given in Chapter 6.

The Principal Investigator is Prof. Gordon Garmire of the Huntingdon Institute for X-ray Astronomy, LLC.

1.11 HRC

The HRC is composed of two microchannel plate (MCP) imaging detectors: the HRC-I, designed for wide-field imaging, and the HRC-S, designed to serve as a read-out for the LETG. The HRC-I is placed at right angles to the optical axis, tangent to the focal surface. The HRC-S is made of three flat elements, the outer two of which are tilted to approximate the LETG Rowland circle. The HRC detectors have the highest spatial resolution on *Chandra*, matching the HRMA point spread function most closely. Under certain circumstances, the HRC-S detector also offers the fastest time resolution (16μs). Details concerning the HRC are in Chapter 7.

The current Instrument Principal Investigator is Dr. Ralph Kraft of the Smithsonian Astrophysical Observatory, who was appointed to this position following the untimely passing in 2015-Aug of the original HRC Principal Investigator, Dr. Stephen Murray of SAO.

1.12 HETG

The HETG, when operated with the ACIS-S, forms the High-Energy Transmission Grating Spectrometer (HETGS) for high resolution spectroscopy. The HETGS achieves resolving power ($E/\Delta E$) up to 1000 in the band between 0.4 keV and 10.0 keV. The HETG is composed of two grating assemblies—the High Energy Grating (HEG) and the Medium Energy Grating (MEG)—on a single structure that can, by command, be placed in the optical path just behind the HRMA. The HEG intercepts X-rays from only the two inner mirror shells and the MEG intercepts X-rays from only the two outer mirror shells. The HEG and MEG dispersion directions are offset by 10 deg so the two patterns can be easily distinguished. Details are presented in Chapter 8.

The original Instrument Principal Investigator was Prof. Claude Canizares (retired), and the current Principal Investigator is Dr. Herman Marshall of the *MIT* Kavli Institute for Astrophysics and Space Research.

1.13 LETG

The LETG, when operated with the HRC-S, forms the Low Energy Transmission Grating Spectrometer (LETGS). The LETGS provides the highest spectral resolution on *Chandra* at low (0.08–0.2 keV) energies. The LETG is composed of a single grating assembly that, on command, can be placed in the optical path behind the HRMA. The LETG grating facets intercept and disperse the flux from all of the HRMA mirror shells. Details are given in Chapter 9.

The LETG was developed at the Laboratory for Space Research in Utrecht, the Netherlands, in collaboration with the Max-Planck-Institut für Extraterrestrische Physik in Garching, Germany. The original Instrument Principal Investigator was Dr. Albert Brinkman (retired), who was succeeded as PI by Dr. Jelle Kaastra (retired) and then by Dr. Liyi Gu, all of the Laboratory for Space Research.

1.14 Effective Area Comparisons

The effective areas of the imaging instruments are shown in Figure 1.3. The ACIS curves allow for the expected degradation of the ACIS efficiency caused by molecular contamination predicted for the middle of Cycle 28. A comparison of the effective areas of the grating spectrometers are shown in Figure 1.4. Note that the data from the HEG and MEG are obtained simultaneously. The comparisons shown here are based on the most recent calibration at the time of issuance of this document and are subject to revision. The proposer is urged to read the detailed material in the appropriate chapters and examine the *CXC* web site (see Section 1.16) for updates.

1.15 Allocation of Observing Time

Observing time is awarded through the *NASA* proposal and peer review process. The prospective user must submit a proposal in which the observation is described and justified in terms of the expected results. The proposer must also show that the observation is well suited to *Chandra* and that it is technically feasible. Refer to the Call for Proposals (*CfP*, https://cxc.harvard.edu/proposer/CfP/) for more information.

1.16 How to Get Information and Help

The *CXC* web page (https://cxc.harvard.edu) provides access to documents, proposal preparation tools, and proposal submission software.

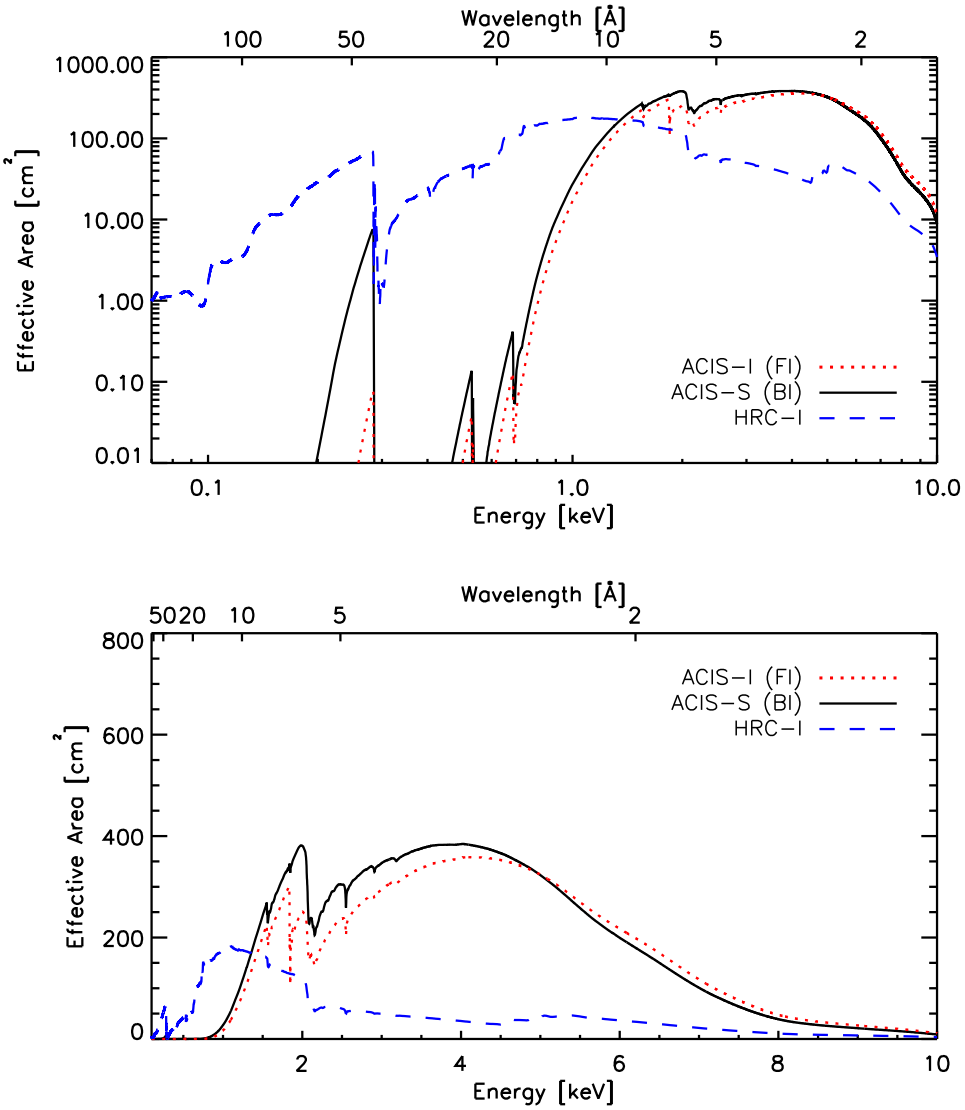

Figure 1.3: Comparison of the on-axis effective areas for observing a point source (integrated over the PSF) of the HRMA/HRC-I, the HRMA/ACIS(FI), and the HRMA/ACIS(BI) combinations. The ACIS curves show the predicted values for the middle of Cycle 28.

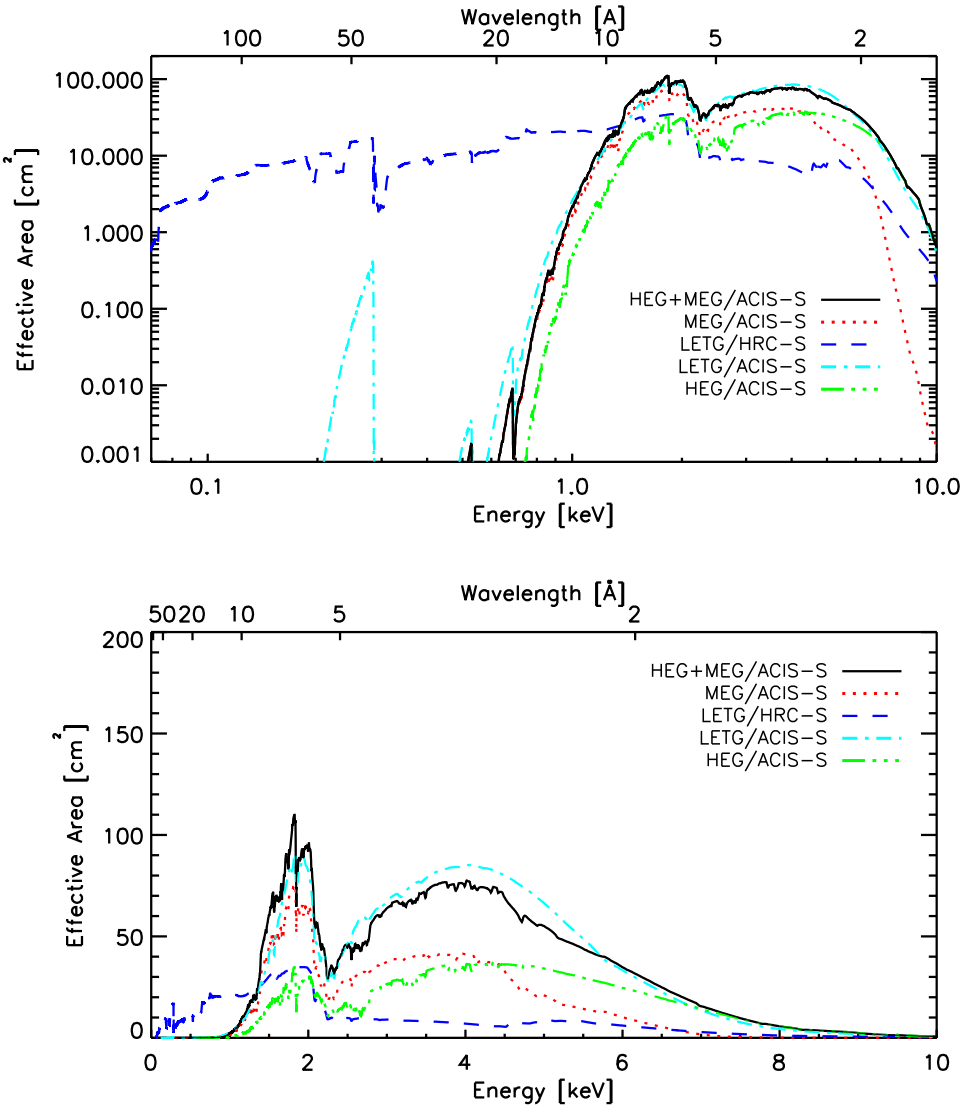

Figure 1.4: Comparison of the total first-order (positive and negative orders combined) effective areas of the LETG and HETG (HEG and MEG are shown separately and summed) spectrometers. HEG and MEG spectra are obtained simultaneously and can sometimes be usefully combined. For a given energy in the range of overlap, the resolving power of the HEG is approximately twice that of the MEG, which in turn is approximately twice that of the LETG. The LETG extends to much lower energies than reached by the HETG+ACIS-S combination, especially when used with the HRC-S detector. For full details on spectrometer performance and observation planning see Chapters 8 (HETG) and 9 (LETG).

Chapter 2

Spacecraft, Telescope, Operations, & Mission Planning

2.1 Introduction

This chapter provides a brief overview of the spacecraft, the telescope system—including the Science Instrument Module (SIM)—operations, and mission planning.

A number of observatory parameters are given in Table 2.1.

2.2 Spacecraft

An outline drawing of the Observatory was shown in Figure 1.1. The spacecraft equipment panels are mounted to, and supported by, a central cylindrical structure. The rear of the spacecraft attaches to the telescope system.

The spacecraft includes six subsystems:

1. **Structures and Mechanical** Subsystem. This subsystem includes all spacecraft structures, mechanisms (both mechanical and electro-mechanical), and structural interfaces with the Space Shuttle. Mechanisms, such as those required for the sunshade door, are also part of this subsystem.

2. **Thermal Control** Subsystem. Thermal control is primarily passive, using thermal coatings and multi-layer insulation blankets. On-board-computer-controlled electrical heaters augment these passive elements to maintain sensitive items such as the HRMA at nearly constant temperature.

3. **Electrical and Power** Subsystem. This subsystem includes all hardware necessary to generate, condition, and store electrical energy. Power is generated by solar cells mounted on two solar array wings (three panels each), sized to provide a 15% end-of-life power margin. Electrical power is stored in three NiH_2 batteries. These batteries

Table 2.1: Spacecraft Parameters

Chandra "dry" mass (incl. reserve)	4790 kg
Loaded Propellant	40 kg
Electrical Power	3 NiH_2 batteries
	Two 3-panel solar arrays
Nominal Operating Power	800–1100 W
Optical bench length	$\sim$ 10 meters
SIM focus adjustment range	$\pm$0.4 inches
SIM focus adjustment accuracy	$\pm$0.0005 inches
SIM Z-position adjustment repeatability	$\pm$0.005 inches
Solid-state recorder capacity	1.8 Gb $\times$ 2
On-board command storage	5400 command words
Nominal command storage period	72 hours
Observatory telemetry data-rate	32 kbps
Telemetry playback downlink rates	1024, 512, and 256 kbps
Nominal ground contact periods	45 to 75 minutes per 8 hours
SI telemetry rate	24 kbps
Telemetry format 1 major frame	= 32.8 sec
	= 128 minor frames
Clock error	$< 100\mu$s
Clock stability	1:10^9 per day
Clock frequency	1.024 MHz

provide spacecraft power during times when either Earth or the Moon partially or completely blocks the Sun. Even so, the battery capacity requires that certain non-critical items, including science instruments, be powered down during eclipses. These eclipses occur infrequently due to the particular nature of the *Chandra* orbit.

4. **Communication, Command, and Data Management** (CCDM) Subsystem. This subsystem includes all the equipment necessary to provide ranging, modulation, and demodulation of radio frequency transmission of commands and data to and from the Deep Space Network (DSN) *NASA* Communication System (NASCOM). The CCDM includes two low gain antennas, providing omni-directional communications, an on-board computer (OBC), a serial digital data bus for communication with other spacecraft components, the spacecraft clock, and a telemetry formatter that provides several different formats.

5. **Pointing Control and Aspect Determination** (PCAD) Subsystem. This subsystem includes the hardware and control algorithms for attitude determination and for attitude and solar array control. The solar arrays can be rotated about one axis. The PCAD subsystem also includes hardware for safing the observatory. Specific details of the PCAD subsystem especially relevant to scientific performance are discussed in Chapter 5.

6. **Propulsion** Subsystem. This subsystem consists of the Integral Propulsion System (IPS) and the Momentum Unloading Propulsion Subsystem (MUPS). The IPS contains the thrusters and fuel for control of the orbit and spacecraft orientation during orbit transfer. This system was disabled once the final orbit was achieved for observatory safety reasons. The MUPS provides momentum unloading during normal on-orbit operations. Given current usage rates there would be sufficient MUPS fuel to support ∼50 further years of operation.

2.3 Telescope System

The principal element of the telescope system is the High Resolution Mirror Assembly (HRMA, Chapter 4). The HRMA, composed of four concentric grazing incidence X-ray telescopes, focuses X-rays on the selected detector located in the Science Instrument Module (SIM, Section 2.4). The grating assemblies are also attached to the HRMA module. The telescope system also includes:

1. Optical Bench Assembly
2. Spacecraft Support Structure Assembly
3. Fiducial Transfer Optical Components
4. Spacecraft to Telescope Support Struts

5. Forward and Aft HRMA Contamination Covers

6. Magnetic Baffle Assembly

7. Stovepipe Baffle

The Optical Bench Assembly is primarily the long composite structure separating the HRMA from the SIM. The Spacecraft Support Structure Assembly includes the ring to which the spacecraft is mounted. The Fiducial Transfer Assembly Optical Components are discussed in Chapter 5. The forward and aft contamination covers were opened on-orbit and cannot be closed. The forward contamination cover also serves as the sunshade.

The Magnetic Baffle Assembly was designed to prevent low energy (up to about ~ 100 keV) electrons (reflecting through the X-ray optics) from reaching the focal plane. More details about these baffles may be found at https://wwwastro.msfc.nasa.gov/xray/spectops.

The stovepipe baffle, located inside the optical bench and at the entrance to the SIM, includes tantalum coated plates to prevent X-rays, other than those passing through the telescope, from reaching the focal plane. There are several such baffles inside the optical bench. Details of the baffles may be found at the WWW address above.

2.4 Science Instrument Module (SIM)

The SIM, shown schematically in Figure 2.1, is a movable bench on which the focal-plane X-ray detectors are mounted. Kinematic mounts (flexures) and thermal isolation are provided between the SIM and the telescope optical bench. A graphite epoxy support structure houses the translation stage.

2.4.1 SIM Motions

The focal-plane instruments are positioned by the SIM Z-axis translation stage with a repeatability to ± 0.005 inches over a translation range of 20 inches. The SIM X-axis motion sets the focus to an accuracy of ± 0.0005 inches over a range of 0.8 inches. The fine-focus adjustment step is 0.00005 inches.

2.5 Electron Proton Helium Instrument (EPHIN)

The local particle radiation environment was monitored by the EPHIN detector until 2018, when it was depowered due to erratic behavior related to increasing on-board temperatures. EPHIN consists of an array of 5 silicon detectors with anti-coincidence shielding. The instrument is sensitive to electrons in the energy range 150 keV–5 MeV, and protons/helium isotopes in the energy range 5–49 MeV/nucleon. The field of view is 83 deg and the instrument is mounted on the Sun side of the spacecraft near the HRMA. Prior to 2013 the EPHIN data rates were monitored by the OBC to activate commands to safe the

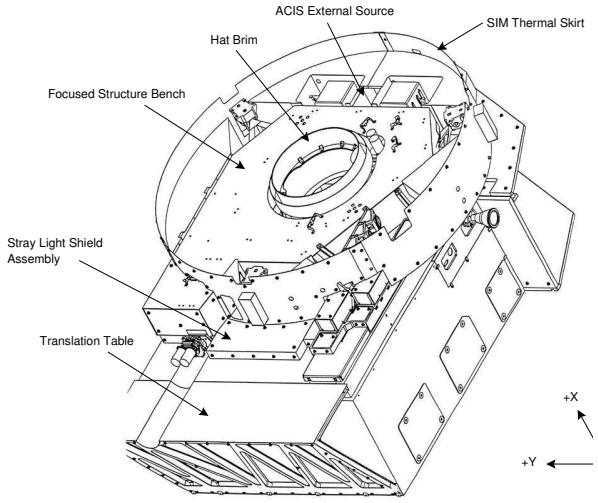

Figure 2.1: A schematic of the Science Instrument Module.

ACIS and the HRC during periods of high radiation such as a solar flare. Due to spacecraft heating, the EPHIN no longer produces reliable data and thus on-board radiation monitoring data are now provided by the ACIS instrument.

The forerunner of the *Chandra*-EPHIN was flown on the Solar and Heliospheric Observatory (SOHO) satellite. Information is available at http://www2.physik.uni-kiel.de/et/ag-heber/costep/ephin.php. The EPHIN instrument was built by the Institut für Experimentelle und Angewandte Physik at the University of Kiel, Germany. Drs. Reinhold Muller-Mellin and Hoarst Kunow are the Co-Principal Investigators.

2.6 Operations

2.6.1 Launch and On-orbit Verification

Chandra was launched on board the Space Shuttle *Columbia* from the Kennedy Space Center in Florida on 1999-Jul-23 at 12:31:00:04 a.m. Eastern Daylight Time (EDT). The Observatory was deployed from the Space Shuttle a few hours later at 8:45 a.m. EDT. Two burns of the IUS (Inertial Upper Stage) took place an hour after *Chandra* was released. A series of five burns of the Integral Propulsion System (IPS) over the period 1999-Jul-24 to 1999-Aug-7 took *Chandra* to its final orbit.

Once in final orbit, the Orbital Activation and Checkout (OAC) phase started. During this time, all systems were brought online and numerous calibrations were performed. After the contamination covers on the HRMA were opened, and after a few passages through the radiation belts under this condition, the front-illuminated (FI) ACIS CCDs showed signs of a spatially-dependent decrease of energy resolution together with increased charge transfer inefficiency (CTI), consistent with radiation damage. Steps were successfully taken to prevent further damage (see Chapter 6). Due to this situation, and because of uncertainties of the long term stability of the FI chips at that time, additional ACIS calibrations were performed and emphasis was placed on observations requiring the use of the FI CCDs. Note that the back-illuminated (BI) CCDs were unaffected, and the situation is now stable in that further degradation has been slowed to match pre-launch expectations. See Chapter 6 for further details. Normal operations started in 1999-Nov.

2.6.2 The Ground System

The *Chandra* "Ground System" is composed of facilities required to operate the spacecraft, receive and analyze the spacecraft telemetry, and provide scientific support to the user community. The ground system includes the following elements:

Deep Space Network (DSN). The DSN is used for communicating commands to the spacecraft and receiving telemetry.

***NASA* Communications** (NASCOM). NASCOM provides communications links between the DSN and the Operations and Control Center (*OCC*) and between the *OCC* and other ground facilities.

Operations and Control Center (*OCC*). The *OCC* is responsible for operating the observatory. This includes activities such as preparing command loads, processing telemetry, attitude determination, and monitoring health and safety. *OCC* personnel utilize two major software environments, the On-line System (ONLS) and the Off-line System (OFLS). The ONLS deals primarily with real-time operations such as receiving telemetry and sending commands through the DSN. The OFLS deals with functions such as mission planning and supporting engineering analysis. The Software Maintenance Facility (SMF), which maintains the flight software, is operated by Northrop Grumman Aerospace Systems (NGAS) and is located at the *OCC*.

***Chandra* X-ray Center** (*CXC*). The *CXC* is the focal point for service to the scientific community. The *CXC* is contracted to issue the Call for Proposals (*CfP*; https://cxc.harvard.edu/proposer/CfP/) and organize peer reviews. The *CXC* assists prospective observers in developing proposals, generates an observing plan from the proposals that are selected, and supplies data products to observers. The *CXC* performs on-orbit calibration and maintains the calibration database, produces response functions, etc. The *CXC* is responsible for providing software and limited assistance to observers for analyzing data. The *CXC* is also responsible for archiving *Chandra* data.

2.6.3 Commanding

All normal *Chandra* operations are preplanned. The OFLS divides the mission schedule into approximately one day segments and generates spacecraft and instrument commands to be executed that day. Once a day, this command load is uplinked to the spacecraft and stored. Three consecutive daily segments are loaded to assure autonomous operation for 72 hours. Stored command loads can be interrupted if necessary, and updated either because of an emergency or to accommodate Targets of Opportunity (TOOs). The interruption process may require up to 24 hours to complete depending on numerous factors including the availability of ground contact. In a true emergency, ground contact can almost always be scheduled.

2.6.4 Telemetry

The telemetry is formatted into major frames and minor frames, where a major frame lasts 32.8 sec and includes 128 minor frames. Each minor frame contains 1019 bytes of science and engineering data plus a 6 byte header, totalling 1025—not 1024—bytes. During normal science operations, telemetry data are generated on the Observatory at a rate of 32 kbps, of which 24 kbps are devoted to the "science stream" data from one of the focal-plane instruments and the remainder allocated to other systems, including 0.5 kbs to the "next-in-line" instrument. The data are recorded on one of two solid state

recorders for subsequent transmission. Each solid state recorder has a capacity of 1.8 Gbits, equivalent to 16 hours of operation.

The recorded data are transmitted through one of the low gain antennas to the ground at 1024 kbps (or 512 or 256 kbps) during scheduled Deep Space Network contacts every eight hours (nominally). Contacts last 45–75 minutes. The ground stations, in turn, transmit the data to JPL, which then transmits the data to the *OCC*.

2.6.5 SI Science Data

There are individual telemetry formats for HRC and ACIS data. The 24 kbps data is collected by the CCDM subsystem from each instrument as a sequence of 8-bit serial-digital words through a Remote Command and Telemetry Unit (RCTU). An additional small amount of housekeeping telemetry is always collected from each instrument independent of the selected format.

2.6.6 Event Timing and the Spacecraft Clock

The CCDM subsystem provides prime and redundant 1.024 MHz clocks, and the ($1/1.024\mu$s) pulses are utilized by the two focal-plane instruments for timing. Each instrument has electronics that counts the elapsed time since the beginning of the current telemetry major frame. The time of events recorded on *Chandra* are given in Terrestrial Time which differs from Coordinated Universal Time (UTC) by about a minute. (See https://aa.usno.navy.mil/faq/TT for a discussion.) The accuracy of the time relationship is 100 microsec. The spacecraft clock is stable to better than one part in 10^9 per day.

2.7 Mission Planning

2.7.1 The Long-Term Schedule

The *Chandra* scheduling process seeks to maximize the fraction of time on-target while minimizing risk to the spacecraft. Once the list of approved target observations for a new cycle has been finalized and targets have been reviewed in detail by the Observer/Principal Investigator (PI) via User Interface staff (https://cxc.harvard.edu/cdo/observation_scheduling.html#usp), they are scheduled by the Science Mission Planners into a Long Term Schedule (LTS). LTS observations, scheduled into bins typically one week in length, generally do not fully occupy the time available for science scheduling; a reserve of unconstrained observations are kept in a pool and used to fill in short-term schedules (STSs).

Once a new LTS is populated at the start of a Cycle, Mission Planners begin the process of scheduling into short-term bins (usually one week in duration). As the Cycle goes on, the remaining LTS is amended and posted online at https://cxc.harvard.edu/target_lists/longsched.html. Observers should note that the predictive fidelity of the LTS generally decreases farther into the future. The placement of

the unconstrained pool targets can change at any time. As the LTS is revised, non-pool targets may also be reassigned for a variety of reasons including multi-telescope coordination. Observations may be bumped or not completed because of high radiation or TOOs. Although most observations occur during the calendar year corresponding to the cycle, it can take up to two years to fully complete an observing cycle, with a roughly one year overlap between cycles.

Both the LTS and the STS web pages show sequence numbers for every observation that are hyper-linked to descriptive target pages. The STS is available online at https://cxc.harvard.edu/target_lists/stscheds/index.html. Each target page further contains a link to a plot that displays the roll, pitch, and visibility for the target for the duration of the Cycle. The target page also contains links to images of the appropriate *Chandra* instrument superposed at the correct roll on 2 deg images of the sky available from *NASA* SkyView. Any time an observation is reassigned to a new STS bin or scheduled precisely within a STS, a revised set of images is posted.

The LTS takes into account the intrinsic target visibility (based primarily on minimum Sun, Earth and Moon angles; see Section 3.3.2), additional target constraints approved by the Peer Review, and thermal and momentum limitations of the spacecraft. These additional constraints are described in Chapter 3. While user-imposed constraints can significantly enhance the science return of an observation, proposers should be aware that constraints increase the Resource Cost (RC) of the proposed observation, which will be taken into consideration during Peer Review as there is a limit on the total RC budget for each Cycle (see the *CfP*). Additionally, all constraints effectively translate into time constraints that may affect the number of STS bins available for scheduling the observation. Schedules may be interrupted unpredictably by the space radiation environment or by TOO observations. This inevitably means that the next opportunity to meet all the observing constraints can be significantly delayed if those constraints are stringent.

2.7.2 Selecting Candidates for Short-Term Scheduling

For each STS, the Mission Planning and Flight Operations Teams construct an Observation Request (OR) list. The list is composed of a combination of LTS and pool targets chosen to meet both the science requirements of the observations and the constraints of the observatory. The OR is a "short list" of targets that *can* be scheduled: not all of them will necessarily be scheduled. Well before construction of the OR list, all observing parameters must be finalized. An overview of the process follows.

- the observer is notified before the cycle begins that their proposal has been accepted, at which time it is essential that the observer verifies that the observation parameters most critical to mission planning (such as coordinates and constraints) are correct.
- the target is placed in the LTS or in the pool list.
- the observer verifies correctness of all observing parameters after a second contact from the *CXC*.

- the target is made available for scheduling.
- the target appears in an OR list as a candidate for short-term scheduling.
- the target is either scheduled for a specific time in that STS, returned to be placed in a later STS during the revision of the LTS or returned to the list of pool targets.
- the target is observed at the scheduled time or bumped to a later schedule.

Some targets may be assigned to several OR lists before they are finally scheduled.

2.7.3 The Short-Term Scheduling Process

Mission Planning assigns priorities in the OR list to emphasize constrained observations; otherwise they would rarely be scheduled for observation since they tend to have a negative impact on the observing efficiency. Whenever possible the ORs span a range of angles about the satellite-Sun line to prevent excess accumulation of momentum. In consultation with the Science Mission Planning Team, the Flight Operations Team (FOT) constructs detailed STSs and command loads for the spacecraft that combine science observations with engineering activities. Along with observing efficiency, thermal, power, momentum, and pointing constraints are all factored in, as well as minimization of maneuver error and optimal guide star acquisition. Several iterations of optimization and safety checks are not uncommon for each STS before its successful review by all teams concerned (Flight and Science Mission Planning, ACIS, HRC, ACA, PCAD, Mechanisms, Communications, Thermal, Command Management) and final approval by the Flight Director.

The *CXC* currently starts to prepare STSs 3 weeks before they begin execution on *Chandra*. Thus at any given time there may be as many as 3 STSs in various stages of preparation. Changes in any of these require a rebuild which is very labor intensive. Fast-response TOOs are currently the only allowed changes. Even small changes to a schedule typically require 24–48 hours to implement. During nominal Mission Planning, the final STS is approved and ready for upload by the Wednesday or Thursday before the STS commands begin executing Sunday night or Monday morning (UTC). Hence, given the nominal planning cycle, fast ($<$ 1 week) turn-around TOOs can most efficiently be incorporated into the STS if they are submitted to the *CXC* by mid-week. Such submission/notification will reduce the amount of disruption, allow time to meet constraints for other targets, and optimize the chances that all the observing requirements for the TOO can be met.

2.7.4 Coordinated Observations

The *Chandra* proposal process recognizes the scientific value of joining data from *Chandra* with data from other observatories. Joint observations may, but need not, be coordinated. Coordinated observations are those that must be done by all participating observatories within an observer-specified time interval. Coordinated observations must be specified as constraints in the *Chandra* Proposal Software (CPS,

https://cxc.harvard.edu/proposer/CPS.html). Detailed instructions, requirements, and recommendations for proposing joint observations are described further in the *CfP* (https://cxc.harvard.edu/proposer/CfP/).

Once coordinated observations are approved by peer review, the *Chandra* Science Mission Planning team will initiate contact with other observatories to facilitate completion of the coordination. Observers are encouraged to provide complete information to their contacts at the coordinating observatories, after which the observatories' planners will work together to achieve an effective schedule. The joint work of the planners will establish the time intervals that meet the proposal constraints and are feasible for the participating observatories. Observation times are finally fixed when coordinated observations are incorporated into short-term schedules (Section 2.7.3).

2.7.5 The Chandra Cool Targets Program

The scheduling of *Chandra* observations depends on balancing the solar pitch angle-dependent heating and cooling of multiple satellite subsystems (see Section 3.3 for additional details). Accordingly, there may be a need throughout the year to find thermally-useful science targets (generally to cool particular subsystems) in limited (and time-dependent) areas of the sky. As there may be no peer-review-approved targets in a specific area of the sky, and to avoid potential loss of observing time from pointings at blank sky, the *CXC* perceived a need for an ample source of targets broadly distributed throughout the sky at low ecliptic latitudes (within 40 degrees of the ecliptic). A call for white papers was issued in 2018-Sep with the goal of collecting suitable catalogs of targets. In addition to avoiding the always-hot regions near the ecliptic poles, the submitted catalogs/lists of targets need to be scientifically useful under the following conditions: a limited (and initially unknown) number from a particular catalog/list are observed, limited exposure times (10 to 35 ks) per target, no constraints or preferences, and with limited observing modes.

Targets from this program are scheduled only if there are no other approved targets in the desired region of the sky, and the data are made public immediately. It is the purpose of the program to make available several thousand targets distributed as uniformly as practicable over the sky at low ecliptic latitudes to meet any future operational needs. A call for white papers was concluded on 2018-Oct-22; the results are distributed on the *Chandra* Cool Targets (CCTs) website: https://cxc.harvard.edu/proposer/CCTs.html.

Chapter 3

Offset Pointing, Visibility, and other Constraints

3.1 Introduction

This chapter gathers together several topics pertaining to observation planning, irrespective of focal-plane instrument and grating configuration, to serve as additional guidelines for the preparation of proposals. Many of these topics are automatically addressed by the target visibility interface webtool (*ProVis*; https://cxc.harvard.edu/cgi-bin/provis/provis_load.cgi) or the observation visualizer software (*ObsVis*; https://cxc.harvard.edu/obsvis/) available as part of the *Chandra* Interactive Analysis of Observations (CIAO) software. The intention here is to familiarize the user with these various observation-planning topics.

3.2 Offset Pointing

The offset pointing convention for *Chandra* is that a negative offset of a coordinate moves the image to more positive values of the coordinate and vice-versa. Examples of offset pointings of the ACIS instrument are shown in Figure 3.1. Examples using the HRC are shown in Figures 3.2 and 3.3.

3.3 Visibility

There are a number of factors that limit when observations can be performed. These are discussed in the following subsections.

3.3.1 Radiation Belt Passages

High particle-radiation levels are encountered as the Observatory approaches perigee. Data acquisition ceases whenever certain particle-radiation thresholds are exceeded. A work-

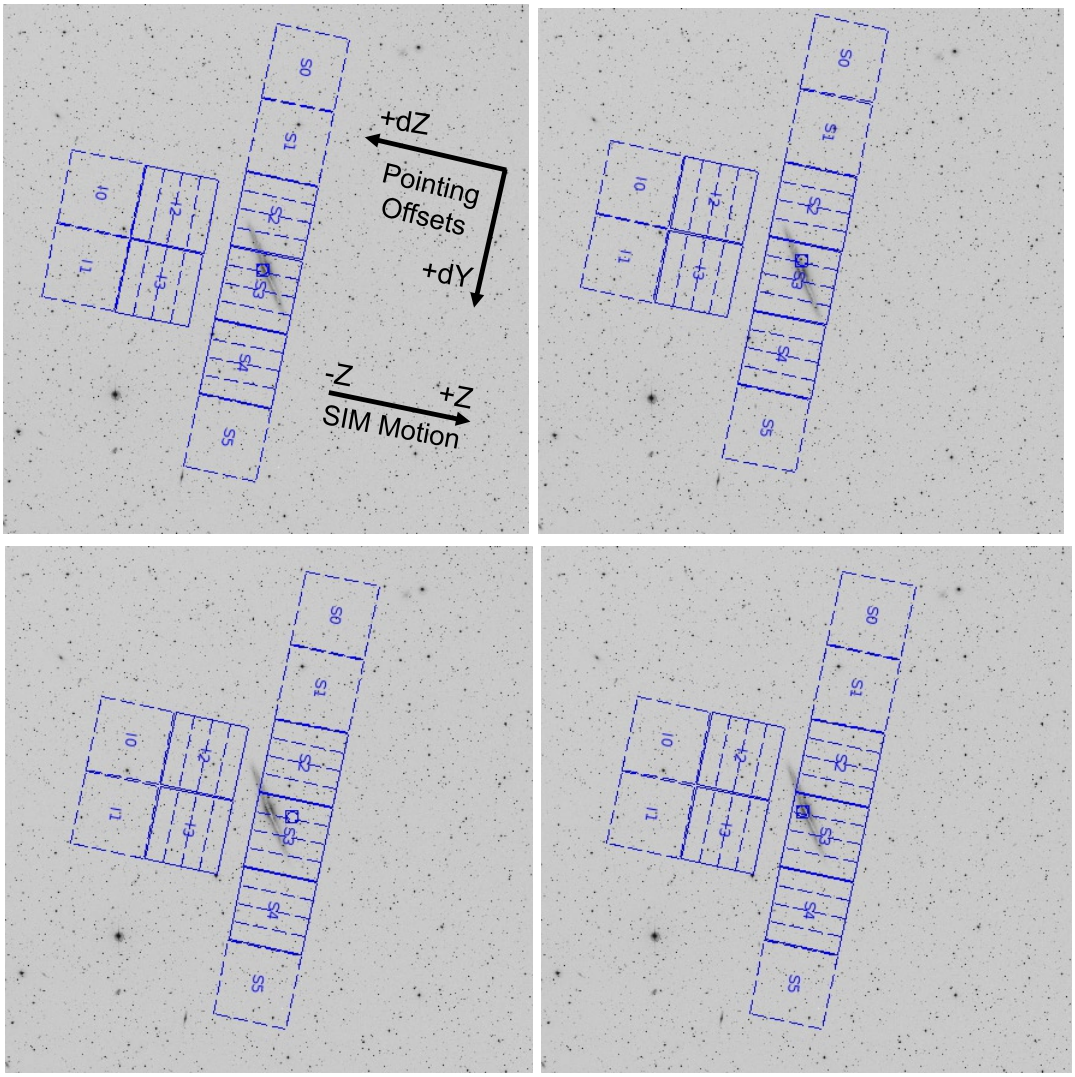

Figure 3.1: Image ($1^\circ \times 1^\circ$) created with ObsVis shows the ACIS field-of-view overlaid on the optical image of the galaxy NGC 891. North is up and east is to the left. The roll angle—the angle between celestial north and the spacecraft −Z axis projected on the sky, as measured positive from N through W—shown is 100°. Five chips are turned on (solid outlines, with dashed node boundaries shown) and five off (dashed chip outlines). The circle-in-a-square symbol shows the aimpoint. ***Upper left panel:*** Target is centered at the nominal ACIS-S aimpoint. The (dY,dZ) coordinate system used to define pointing offsets is also shown (refer to Figure 1.1 for a general view of Chandra spacecraft coordinates), as well as the SIM Z motion direction used to define the location of the aimpoint on the detector (see Figure 1.2). In the case shown, (dY,dZ)=(0,0). ***Upper right panel:*** The target has been offset by 1.5 arcmin in the negative Y direction, i.e., a (dY,dZ) target offset of (−1.5 arcmin, 0). ***Lower left panel:*** The target has been offset by 3 arcmin in the negative Z direction, i.e., (dY,dZ)=(0, −3 arcmin). ***Lower right panel:*** No target offset has been applied, i.e., (dY,dZ)=(0,0), but a SIM Z motion of +3 arcmin has been applied. Note that in the latter two cases (lower left and lower right) the target is in the same place on the detector, but in the lower right case the aimpoint moves with the target, keeping the target near the optical axis and the aimpoint for optimal imaging quality.

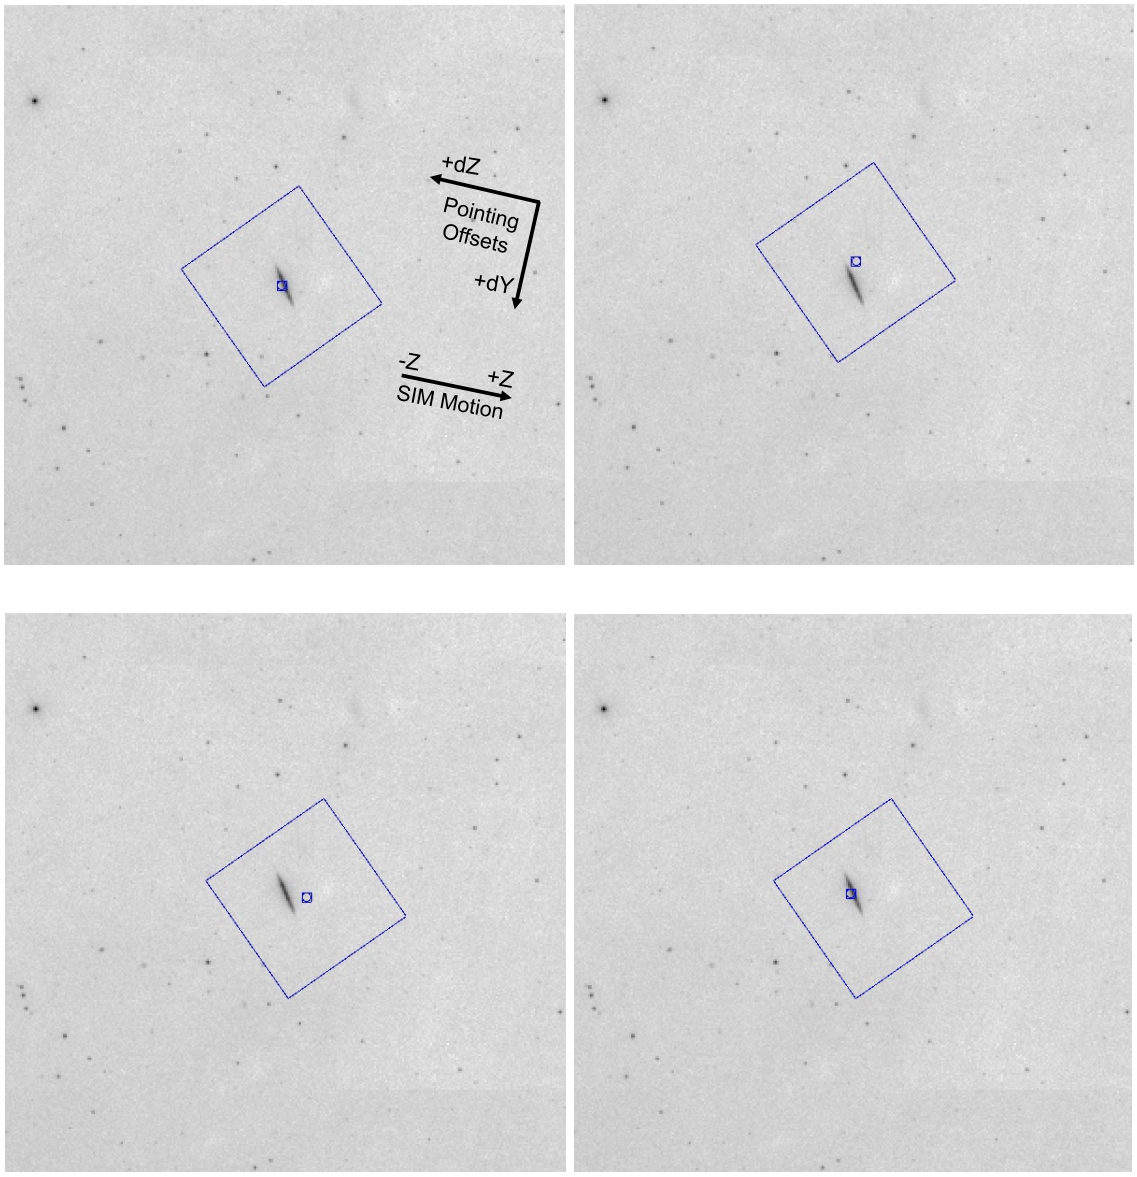

Figure 3.2: Offset pointing with HRC-I on a $2^\circ \times 2^\circ$ image of NGC 891; orientation and roll angle are the same as in Figure 3.1. ***Upper left panel:*** Target is centered at the nominal HRC-I aimpoint. The (dY,dZ) coordinate system used to define pointing offsets is also shown, as well as the SIM Z motion direction used to define the location of the aimpoint on the detector. ***Upper right panel:*** The target has been offset by 5 arcmin in the negative Y direction, i.e., a (dY,dZ) target offset of (-5 arcmin, 0). ***Lower left panel:*** The target has been offset by 5 arcmin in the negative Z direction, i.e., (dY,dZ)=(0, -5 arcmin). ***Lower right panel:*** No target offset has been applied, i.e., (dY,dZ)=(0,0), but a SIM Z motion of +5 arcmin has been applied. Note that in the latter two cases (lower left and lower right) the target is in the same place on the detector, but in the latter case *the aimpoint moves with the target.*

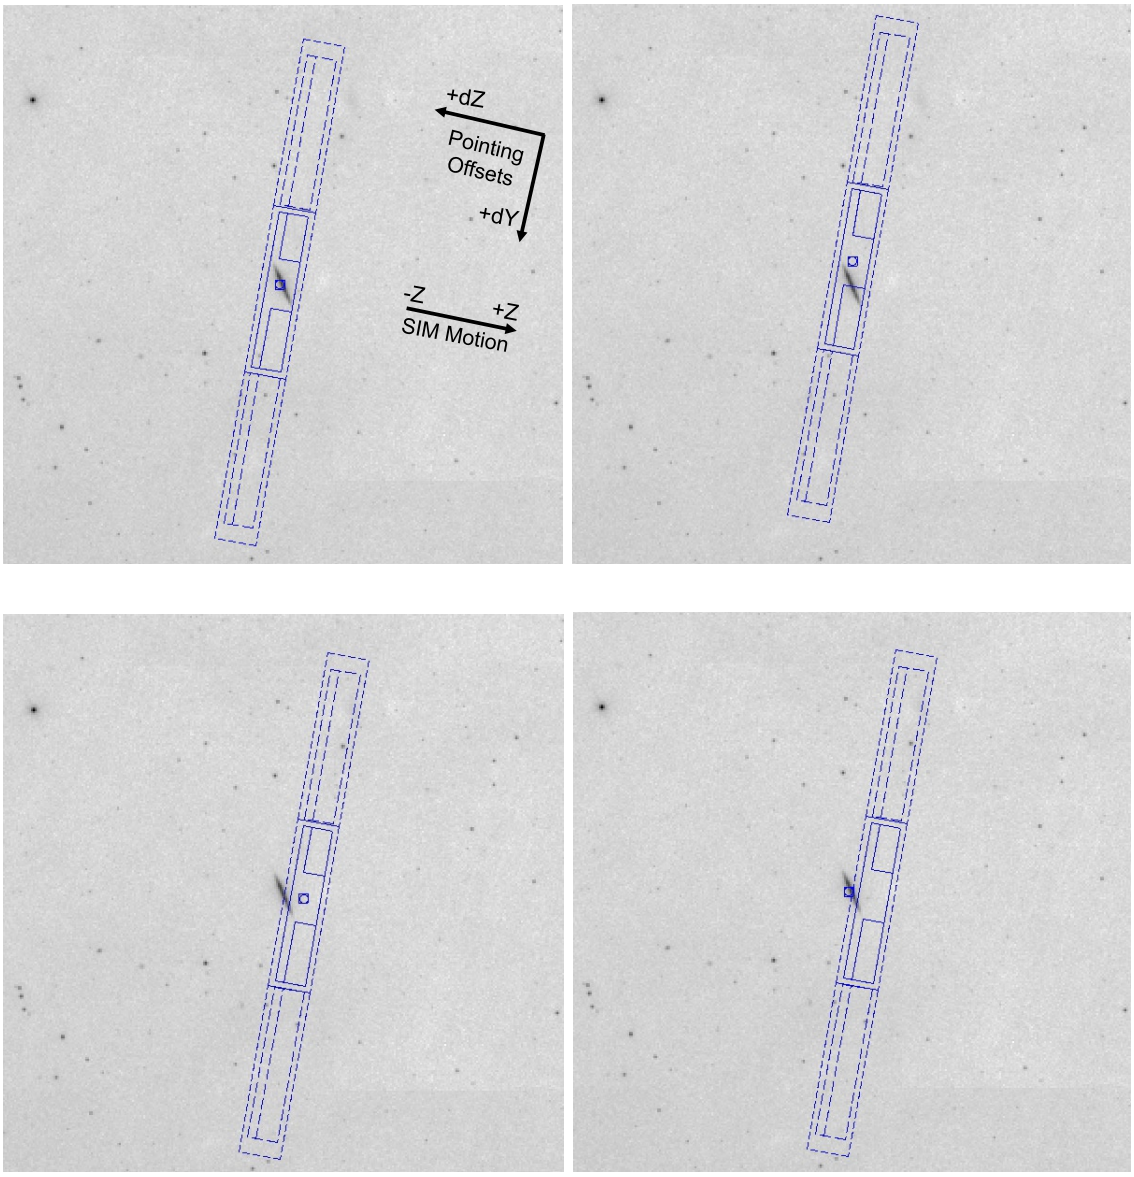

Figure 3.3: Offset pointing with HRC-S on a $2^\circ \times 2^\circ$ image of NGC 891; orientation and roll angle are the same as in Figure 3.1. ***Upper left panel:*** Target is centered at the nominal HRC-S aimpoint. The (dY,dZ) coordinate system used to define pointing offsets is also shown, as well as the SIM Z motion direction used to define the location of the aimpoint on the detector. ***Upper right panel:*** The target has been offset by 5 arcmin in the negative Y direction, i.e., a (dY,dZ) target offset of (-5 arcmin, 0). ***Lower left panel:*** The target has been offset by 5 arcmin in the negative Z direction, i.e., (dY,dZ)=(0, -5 arcmin). ***Lower right panel:*** No target offset has been applied, i.e., (dY,dZ)=(0,0), but a SIM Z motion of +5 arcmin has been applied. Note that in the latter two cases (lower left and lower right) the target is in the same place on the detector, but in the latter case *the aimpoint moves with the target.* Note also that, as shown in the lower two panels, it is quite possible to use offsets to move the target position entirely off a detector.

ing number for the altitude at which this takes place is about 60,000 km. Cessation of observations and protection of the instruments in regions of high radiation results in approximately 25% of the 63.5 hour *Chandra* orbit being unusable.

3.3.2 Avoidances

The following constraints are necessary to ensure the health and safety of the spacecraft and science instruments. Proposals which violate these constraints may be rejected.

1. Sun avoidance: viewing is restricted to angles larger than 46.2 deg from the center of the Sun. This restriction makes about 15% of the sky inaccessible on any given date, but no part of the sky is ever inaccessible for more than 3 months. This constraint cannot be overridden.

2. Moon avoidance: viewing is restricted to angles larger than 6 deg from the limb of the Moon. This restriction makes less than 1% of the sky inaccessible at any time. This avoidance can be waived, but at the price of a reduced-accuracy aspect solution (see Chapter 5).

3. Bright Earth avoidance: viewing is restricted to angles larger than 10 deg from the limb of the bright Earth. This restriction makes less than 5% of the sky inaccessible at any time, but there are certain regions which can only be viewed, continuously, for up to about 30 ks. The avoidance can be waived, but at the price of a reduced-accuracy aspect solution (see Chapter 5). Figure 3.4 illustrates the point that the Earth avoidance region is nearly stationary; this is a consequence of the combination of high elliptical orbit and the cessation of science data acquisition during radiation belt passages. This partially blocked region moves several degrees per year, reflecting the evolution of the orbital elements.

   The greatest amount of observing time is available in the vicinity of apogee, when the satellite moves most slowly and Earth and its avoidance zone occupy an approximately stationary location on the sky, visible in Figure 3.4 as the extension to the south of the Sun avoidance band.

3.3.3 Roll Angle Constraints

The spacecraft and instruments were designed to take advantage of the Observatory having a hot and a cold side. Thus, the spacecraft is preferentially oriented with the Sun on the $-Z$ side of the X-Y plane, where $+X$ is in the viewing direction, the Y-axis is parallel to the solar panel axes, and $+Z$ is in the direction of the ACIS radiator (see Figure 1.1). In this orientation there is only one "roll angle" (rotation about the X-axis, defined positive west of north) for which the solar panels can be rotated so that they are directly viewing the Sun—the nominal roll angle. Small deviations ($\sim$deg) from the nominal roll angle may be allowed, depending on the viewing geometry. The roll-angle constraint imposes further visibility restrictions. These restrictions can also be evaluated with the *ProVis* tool.

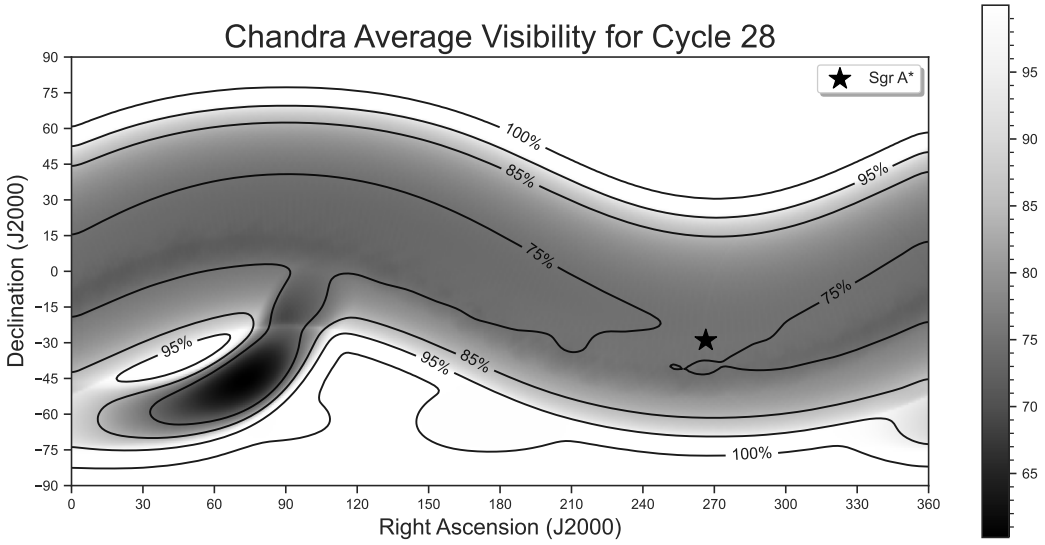

Figure 3.4: The *Chandra* visibility showing contours of fractional visibility averaged over the 12-month interval of Cycle 28. The darker the shade of gray, the lower the visibility.

3.3.4 Thermal Constraints

As the spacecraft has aged, its thermal properties have gradually changed. Degradation of its reflective multilayer insulation means that the Sun-facing side of the satellite now heats more rapidly, and the spacecraft as a whole operates at higher temperatures than earlier in the mission. Steps have been taken to reduce internal sources of heating. However, the thermal limits of the various spacecraft subsystems are an increasingly important factor in planning operations.

The solar pitch angle (i.e., the angle between the viewing direction and the direction to the Sun) affects which subsystems heat or cool during any given observation. Figure 3.5 shows the solar pitch sensitivity of various spacecraft components. Restrictions on the operating temperatures of these components limit how long the spacecraft can maintain a given pitch angle, and it is often necessary to plan operations such that components heated during one observation can cool during (or in some cases before) the next. The thermal restrictions on the various subsystems evolve with time and are too complex to describe here in detail, but some of their impacts on observations are of direct interest to observers:

1. Observations are possible across the full range of allowed solar pitch angles (46.2 to 178 degrees). However, balancing the thermal load across the whole spacecraft usually requires the breaking up of long observations into shorter exposures. The maximum continuous exposure depends on the solar pitch angle of the target, the instruments used, and the thermal history of the spacecraft. As of late 2022, observa-

tions are usually broken up into segments of duration ≤30 ks for planning purposes, and these segments may be scheduled with significant periods of time separating them. It is expected that segments will be shorter as the spacecraft continues to age.

2. The ACIS focal plane and electronics subsystems are particularly sensitive to heating at high solar pitch angles. The duration of ACIS exposures is limited to ensure that the instrument remains within its required operating temperature range. Observations which use fewer CCDs cause less heating. In Cycle 28, ACIS observers will be able to specify a maximum of 4 required chips; additional chips beyond those marked as required may be specified as optional, along with the order in which these may be turned off; this information will be used as needed during the mission planning process to control ACIS temperatures (see Section 6.22.1).

3. The duration of HRC exposures is currently limited to 14.5 ks, in order to meet operational constraints on the temperature of the instrument electronics (see Section 7.12 for further details). This maximum duration may be reduced depending on external factors such as target pointing direction and may be subject to further review. Longer HRC observations will be broken into segments up to 14.5 ks in duration and scheduled with a minimum of 30 ks separating the segments. The required cooling time between segments means that including more than 2 HRC exposures in an orbit can be difficult, and no more than 4 will be scheduled.

4. The increased operating temperature of the ACA has affected its ability to acquire faint guide stars. This can restrict observations of targets with difficult star fields to particular times of year since, if sufficiently bright guide stars occur only at the edges of the ACA field of view, they may only be visible at a limited range of roll angles. In future, some target fields may be unobservable owing to a lack of sufficiently bright guide stars.

5. Simultaneous longer-duration observations with telescopes that have a limited range of accessible solar pitch angles, such as the *X-ray Multi-Mirror Mission* (*XMM-Newton*; approximately 70–110 deg), may be difficult, or even impossible, to schedule.

Targets whose solar pitch angle varies through the year will generally be scheduled at times where these sensitivities are least restrictive. Targets for which the available solar pitch angles are always unfavorable may be broken up into smaller segments so as not to exceed the maximum dwell times for those pitch ranges. The specific impact of the various solar pitch angle sensitivities (Figure 3.5) on an observation's visibility and, therefore, scheduling difficulty is encapsulated by the Resource Cost of a given observation (see Section 4.3 of the Call for Proposals). Proposers concerned about Resource Cost and/or scheduling difficulty should contact the *CXC* HelpDesk (https://cxc.harvard.edu/helpdesk/) prior to submission.

3.4 Other Instrument and Observation Considerations

Science instrument restrictions and limitations are discussed in the chapters devoted specifically to the instruments. User-imposed constraints are discussed in the instructions for the *Chandra* Proposal Software (CPS). We summarize these here.

3.4.1 Instrument Considerations

For details on the following considerations that may be important in the planning of observations, please refer to Chapters 6 (for ACIS) and 7 (for HRC). Note that we refer to *restrictions* when degradation of instrument capabilities may result, and to *limitations* when the science objectives of an observation may be jeopardized.

- The HRC has a brightness restriction which limits the flux per microchannel plate pore. Please see Section 7.14.2.
- The HRC has a telemetry limitation. Exceeding this limit, among other consequences, reduces observing efficiency. Please see Section 7.14.2.
- The HRC has a linearity limitation. Exceeding this limitation voids the effective area calibrations. Please see Section 7.14.2.
- The ACIS has a telemetry limitation. Exceeding this, among other consequences, reduces observing efficiency. Please see Section 6.15.2.
- The ACIS is subject to the effects of pile-up. For high flux sources, when multiple photons arrive within a single CCD frametime, they may be counted as a single photon of higher total energy. Dealing with this effect requires careful planning of the observation. Please see Section 6.16.
- The ACIS has a restriction on the total amount of allowed flux in a pixel during an observation. This flux restriction only impacts a small number of potential observations, primarily those of very bright sources that request the dither to be turned off. Please see Section 6.19.

3.4.2 User-Imposed Constraints

Chandra users may need to specify a number of observing constraints particular to their observations. In general, the specification of a user-imposed constraint decreases the efficiency of the observatory and therefore should be well justified in the proposal. Note that only a limited number of constrained observations can be accommodated (see the *CfP* for details). User imposed constraints are summarized here.

Time Constraints:

Time Windows – specific time intervals in which an observation must be scheduled. Such constraints are primarily for use in coordinated observing campaigns or for arranging an observation to coincide with some time-critical aspect of the target.

Monitoring Intervals – for observing a target repeatedly, with intervals and durations specified.

Phase Interval – specific phase intervals for observing sources with long, regular periods.

Unique Phase (split observations) – specifying that unique parts of the phase need to be covered if the observation is split.

Coordinated Observations – targets specified to be observed by *Chandra* and another observatory within a given time period.

Continuity of observation – specifying that an observation be performed in a single (or the fewest possible) segment(s).

Completion Time (split observations) – specifying relative time frame within which all the segments of a split observation must be completed.

Group Observation – a target that needs to be observed within a particular time range with other targets in the program.

Roll Constraints – specifying a particular roll angle and tolerance.

Roll-dependent Pointing – an observation for which the target coordinates and/or offsets require adjustment based on roll angle.

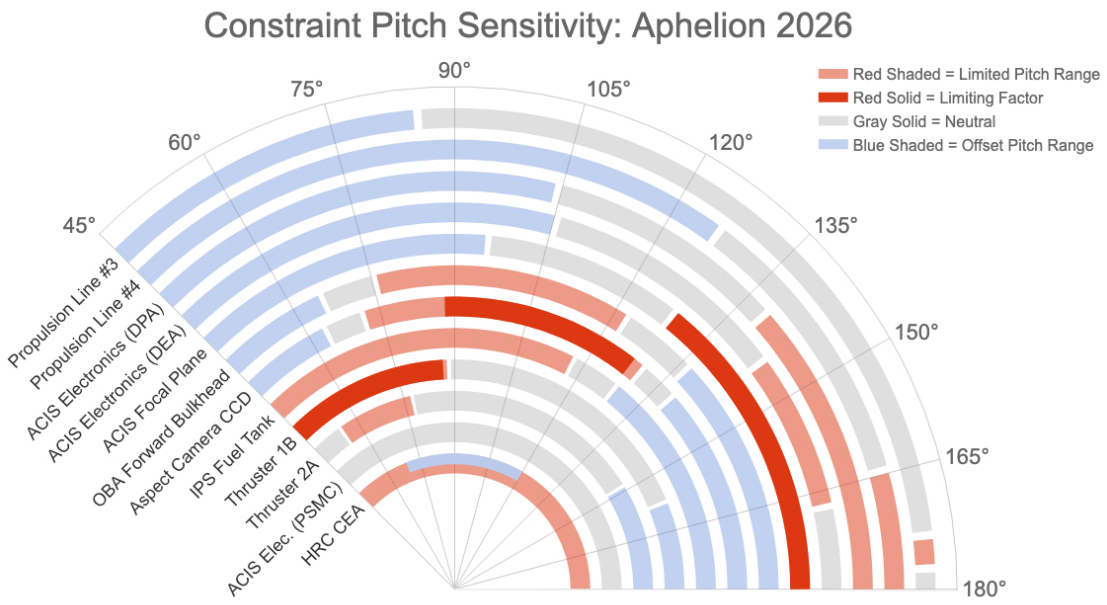

Figure 3.5: Illustration of the thermal impact of solar pitch angle on various spacecraft components with the ACIS at the focus. Pale red arcs indicate pitch ranges in which it is more difficult to maintain the component within its thermal restrictions, gray arcs show where components are thermally neutral, and pale blue arcs indicate pitch ranges in which components can recover from time spent in thermally limited pitch regions. Dark red arcs indicate where a component represents the most limiting dwell time restriction at a given pitch angle; the MUPS at low pitch angles, ACA at mid-pitch, and the ACIS at high pitch angles. The HRC heats at all pitches while switched on, and only cools efficiently in a limited range of pitches while switched off. Note that the diagram represents the predicted status during 2026 with the ACIS at the focus; thermal constraints will change over time and so this figure should not be used as a guide for planning specific observations. For detailed inquiries regarding pitch sensitivity, please contact the *CXC* HelpDesk (https://cxc.harvard.edu/helpdesk/).

Chapter 4

High Resolution Mirror Assembly (HRMA)

4.1 Introduction

The *Chandra* X-ray telescope consists of 4 pairs of concentric thin-walled, grazing-incidence Wolter Type-I mirrors (X-ray optics are reviewed by Aschenbach 1985) called the High Resolution Mirror Assembly (HRMA). The front mirror of each pair is a paraboloid (P_n) and the back a hyperboloid (H_n). The eight mirrors were fabricated from Zerodur glass, polished, and coated with iridium on a binding layer of chromium.

4.1.1 Description and Physical Configuration

The HRMA, shown schematically in Figure 4.1, contains the nested mirrors, center, forward and aft aperture plates, baffles, inner and outer cylinders, mounts, pre- and post-collimators, fiducial light transfer components, mirror support sleeves, forward and aft contamination covers, flux contamination monitors, and thermal control hardware. The outer mirror pair is number 1, and, progressing inwards, 3, 4, and 6. The original design had six mirror pairs; numbers 2 and 5 were eliminated. The pair diameters range from about 0.65 to 1.23 meters. The distance from the center of the Central Aperture Plate (CAP) separating the paraboloid and hyperboloid mirrors to the HRMA focus is 10.0548 meters, with each mirror pair varying slightly about this value. Note that this distance is close to, but not exactly, the focal length. An annular on-axis beam enters each mirror pair, is reflected from paraboloids and hyperboloids, and exits to converge to a focus. The angle θ between the direction of the reflected ray and the optical axis lies between two cone angles θ_c and θ_d. These and other important HRMA characteristics are listed in Table 4.1.

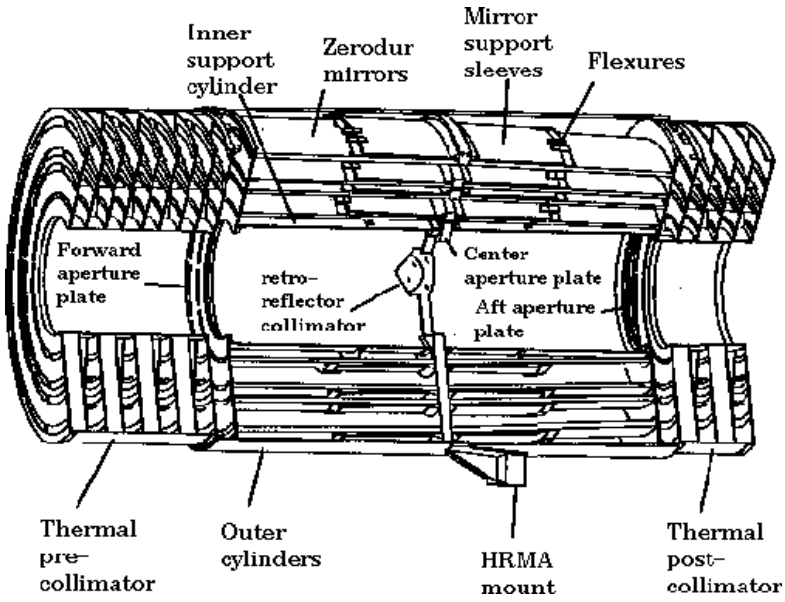

Figure 4.1: The four nested HRMA mirror pairs and associated structures.

Table 4.1: *Chandra* HRMA Characteristics

Optics	Wolter Type-I
Mirror coating	Iridium (330 Å, nominal)
Mirror outer diameters (1, 3, 4, 6)	1.23, 0.99, 0.87, 0.65 m
Mirror lengths (P_n or H_n)	84 cm
Total length (pre- to post-collimator)	276 cm
Unobscured clear aperture	1145 cm^2
Mass	1484 kg
Focal length	10.070 ± 0.003 m
Plate scale	$48.82 \pm 0.02\,\mu$m arcsec^{-1}
Exit cone angles from each hyperboloid:	
θ_c (1, 3, 4, 6)	3.42°, 2.75°, 2.42°, 1.80°
θ_d (1, 3, 4, 6)	3.50°, 2.82°, 2.49°, 1.90°
f-ratios (1, 3, 4, 6)	8.4, 10.4, 11.8, 15.7
PSF FWHM (with detector)	< 0.5 arcsec
Effective area:	@ 0.25 keV 800 cm^2
	@ 5.0 keV 400 cm^2
	@ 8.0 keV 100 cm^2
Ghost-free field of view	30 arcmin diameter

4.1.2 Sub-Assembly Calibration

Extensive measurements of the mirror shapes and of the surface characteristics were made at Hughes Danbury Optical Systems during fabrication of the mirror segments and during assembly at Eastman Kodak Co. HRMA throughput depends critically on the coating of the individual mirror elements carried out at Optical Coating Laboratory in Santa Rosa, California. Mirror flats were present in the coating chamber and coated with iridium at the same time as the HRMA mirror elements. Reflectivity of X-rays from these witness flats was measured with the X-ray beam from the synchrotron at the Brookhaven National Laboratory (Graessle et al. 1998, 2004).

4.1.3 Operating Environment

Insulation and heaters maintain the HRMA temperature at 70 °F (21 °C) on-orbit to minimize changes from the assembly and alignment environments and to minimize molecular contamination.

4.1.4 Heritage

The *Chandra* mirrors represent a logical progression from those of the *Einstein* (Giacconi et al. 1979) and Röntgensatellit (*Rosat*; Trümper 1983, Aschenbach 1991) missions. Both of these previous X-ray observatories utilized nested Wolter Type-I optics with about 4 arcsec angular resolution. The *Einstein* mirror assembly had considerably less geometric area than *Chandra*, while *Rosat* had comparable area ($1100\,\mathrm{cm}^2$) at low energies ($< 1\,\mathrm{keV}$).

To verify the technology required for the spatial resolution of *Chandra*, a Validation Engineering Test Article-I (VETA-I) was constructed and tested in 1991. VETA-I contained the P_1H_1 proto-flight mirror shells constructed to final tolerances, but uncoated and with ends uncut. The VETA-I tests included the image full width at half-maximum (FWHM), encircled energy, effective area, and ring focus properties (for azimuthal and low spatial-frequency figure). Many of the results of these tests appear in Society of Photo-Optical Instrumentation Engineers (SPIE) Proceedings 1742. A good overview of the VETA tests is given by Zhao et al. (1994).

4.2 Calibration and Performance

4.2.1 Calibration and Model

Before launch, the HRMA underwent extensive ground calibration tests at the X-Ray Calibration Facility (XRCF) at Marshall Space Flight Center (MSFC), Huntsville, Alabama, from 1996-Sep through 1997-May. The full HRMA XRCF Calibration Report is accessible at `https://cxc.harvard.edu/cal/Hrma/XRCFReport.html`. During these tests, the mirror assembly was mounted horizontally in a vacuum chamber and irradiated with X-rays from various electron-impact sources located at a distance of 524.7 meters. The data taken

at the XRCF include the effective area and image distributions as a function of incident energy and angle. The mirror performance during these tests differs from that expected in space because of gravity distortions and the finite source size and distance; consequently, the calibration data cannot be directly compared to flight observations. The approach taken was to develop a model based upon surface and assembly measurements taken before the X-ray calibration activity. The X-ray calibration data then were used to validate this model and to make minor adjustments in model parameters to achieve satisfactory agreement with the observations. Further minor modifications were made as a result of flight experience. A series of papers in SPIE Proceedings 3113 report the results of the HRMA ground calibration.

The HRMA characteristics illustrated in this chapter were generated by a raytrace program using this model. Note that this chapter typically gives characteristics of the HRMA only; unless otherwise indicated, blurring caused by the detector and the aspect solution is *not* included. These effects are *very* important for on-axis sources, and are included in the instrument chapters (Chapters 6 and 7). See also Section 4.4.

4.2.2 HRMA Effective Area

The unobscured geometric aperture of the HRMA is 1145 cm^2. The obstruction of the HRMA aperture by supporting struts is less than 10%. Since reflectivity of the mirror optics depends on photon energy as well as grazing angle, the HRMA throughput varies with X-ray energy.

The HRMA effective area is derived from the predictions of the raytrace code discussed above along with empirical corrections based on the XRCF ground-based calibration data. The initial XRCF correction—the correction used to calculate the HRMA effective area prior to Calibration Database (CALDB) 4.1.1—was derived assuming that there was no molecular contamination on the mirrors. Subsequent in-flight gratings observations of blazars showed a discontinuity in the spectrum near the Ir-M edges when reduced with the contaminant free model; this suggests that there may be molecular contamination on the mirrors on-orbit (see Figure 4.2). Using these data and raytrace simulations, it was estimated that a 22 Å layer of hydrocarbon could be present on the mirror optics. An updated HRMA effective area was released in CALDB 3.2.1 on 2005-Dec-15 based on the predictions of the raytrace code with a uniform 22 Å layer of hydrocarbon molecular contamination on all 8 pieces of optics on-orbit.

Subsequently, inconsistencies in the measurements of cluster temperatures derived by two different methods (line and continuum measures) led to a re-analysis of the data taken at the XRCF. This re-analysis provided evidence that molecular contamination was already present on the mirrors at XRCF. Thus, since the initial, ad-hoc, empirical correction of the HRMA model had already corrected for most of the effects of the molecular contamination, the addition of another 22 Å of contamination in the raytrace model post-launch (CALDB 3.2.1) was an over-correction.

During XRCF testing, a system of shutters was placed behind the HRMA so that the effective area of the 4 shells could be measured independently. These tests were essential

since the gratings intercept X-rays from different shells. Two focal-plane instruments were used during XRCF testing: 1) a flow proportional counter (FPC) and 2) a solid-state detector (SSD). Both of these instruments were non-imaging detectors and were used in conjunction with a number of different pin hole apertures. Using the data obtained from each mirror pair, the thicknesses of the contaminant on shells 1, 3, 4 and 6 were determined to be 28, 18, 20 and 27 Å respectively. Thus a new version of the HRMA effective area (CALDB 4.1.1) was released in 2009-Jan, based on the predictions of the raytrace code with the as-measured contamination depths on each shell. In addition, an energy independent correction is applied to the predictions of the raytrace code for each shell to determine the absolute effective area. The correction factor for each shell is calculated by averaging the averaged FPC line data to raytrace ratio and the averaged SSD continuum data to raytrace ratio. Note that, while a gray (i.e. energy-independent) correction is applied to each shell, the overall empirical correction for the full HRMA absolute effective area is not energy-independent, since different shells contribute a different fraction of the total effective area at different energies (see Figure 4.3). This figure reflects an even more refined analysis of the XRCF data than what was used in calculating the HRMA effective area in the CALDB 4.1.1 release. The new analysis results in a slightly lower effective area but is still consistent with the CALDB 4.1.1 data within errors.

Several HRMA effective area models have been generated with different methods for calculating the gray corrections for each shell (e.g., no gray correction or gray corrections with unequal weighting between the FPC and SSD data) to determine the systematic effect of the gray corrections on the gas temperatures derived from ACIS observations of clusters of galaxies. For cool clusters (kT $<$ 4 keV), the derived gas temperatures are essentially independent of the method used to calculate the gray correction. For hotter clusters, the derived temperatures vary by $\pm$ 2% depending on the algorithm used.

The combined HRMA/ACIS and HRMA/HRC effective areas released in CALDB 4.1.1 are shown in Figure 4.4 and the effect of off-axis vignetting on the HRMA effective area is shown in Figure 4.5 at several different photon energies. Note that this change in the effective area also serves to bring *Chandra* and *XMM-Newton* continuum measurements of cluster temperatures into closer agreement.

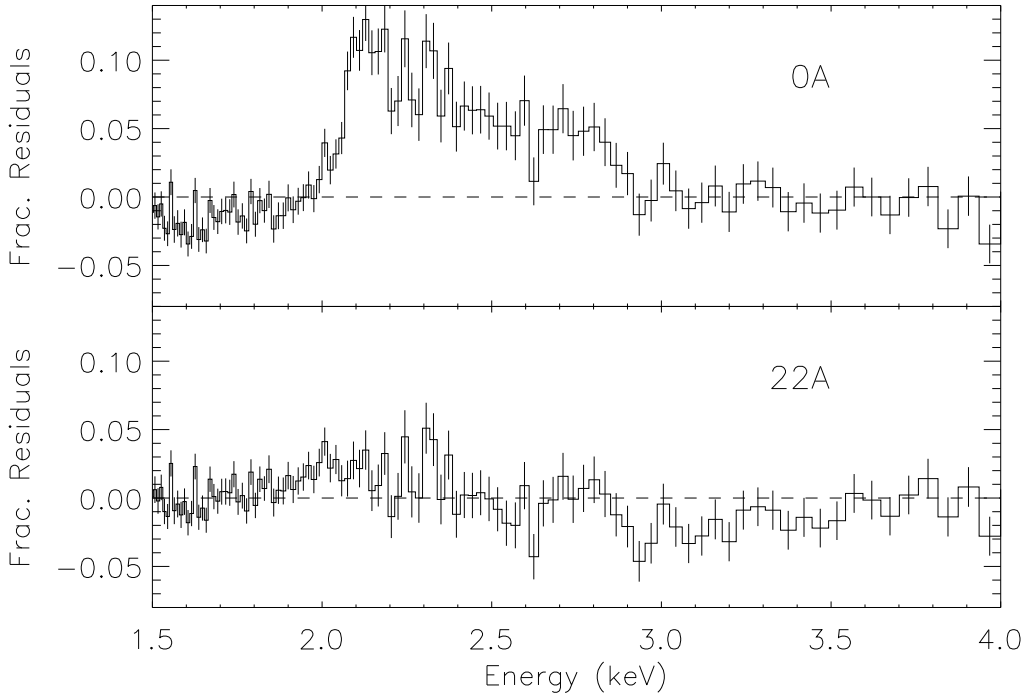

Figure 4.2: Combined residuals from power-law fits to 18 blazar observations (Marshall 2005). The residuals in the upper panel are based upon models of contaminant free optics, while those in the lower are based upon a model of the mirror surfaces which includes a thin (22 Å) hydrocarbon layer.

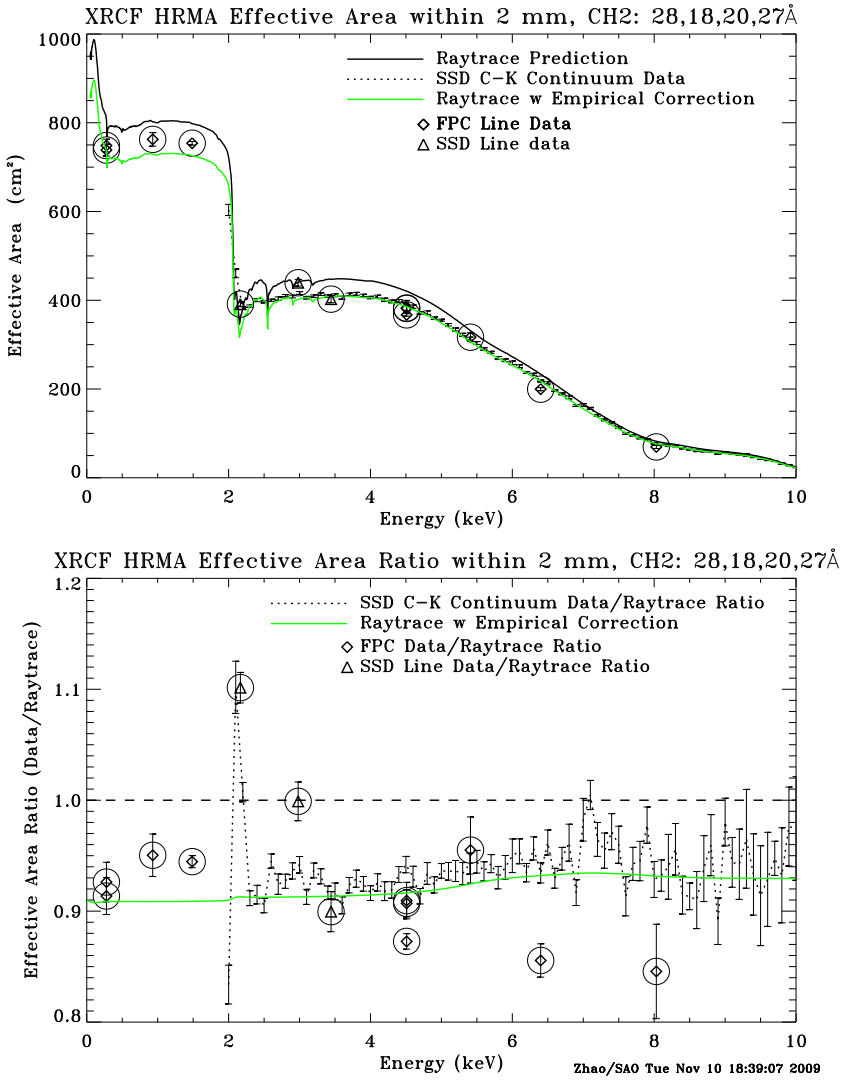

Figure 4.3: The HRMA Effective area: The top panel shows the raytrace prediction and the XRCF measurement data of the HRMA effective area as a function of energy. The bottom panel shows the ratio of the XRCF data to the raytrace. The raytrace includes the effects of molecular (CH_2) contamination with variable thickness on the mirrors. The dotted line with error bars are the C-K continuum data, taken simultaneously with a solid-state detector (SSD). Data obtained from spectral line sources with a flow proportional counter (FPC) are shown as diamonds; those obtained with an SSD are shown as triangles. All the line measurements are circled for clarity. The solid (green) line in the bottom panel shows the XRCF empirical correction to the raytrace based on the XRCF data. The solid (green) line in the upper panel shows the absolute HRMA effective area with the XRCF empirical correction.

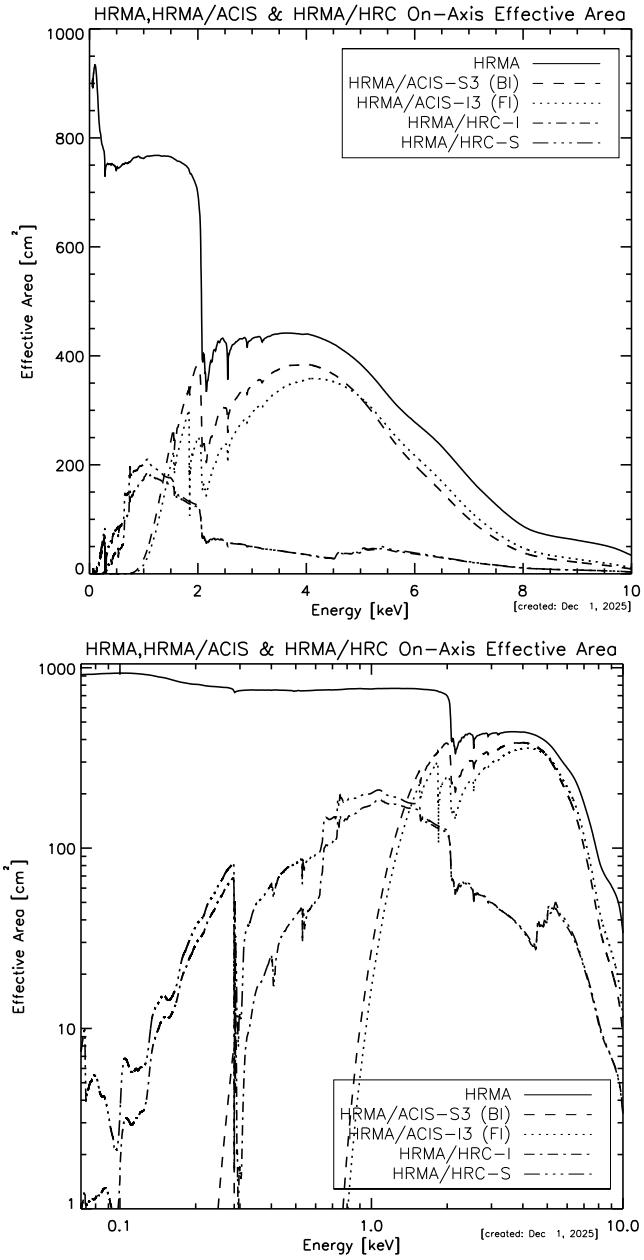

Figure 4.4: The HRMA/ACIS and HRMA/HRC effective areas versus X-ray energy in linear-linear (top) and log-log (bottom) scales. The structure near 2 keV is due to the iridium M-edge. The HRMA effective area is calculated by the raytrace simulation based on the HRMA model and scaled by the XRCF calibration data. The HRMA/ACIS effective area is the product of the HRMA effective area and the Quantum Efficiency (QE) of ACIS-I3 (front illuminated) or ACIS-S3 (back illuminated). The HRMA/HRC effective area is the product of HRMA effective area and the QE of HRC-I or HRC-S at their aimpoints, including the effect of UV/Ion Shields (UVIS). These figures are based upon the expected QEs at launch, and do not take into account any subsequent degradation or contamination while in orbit.

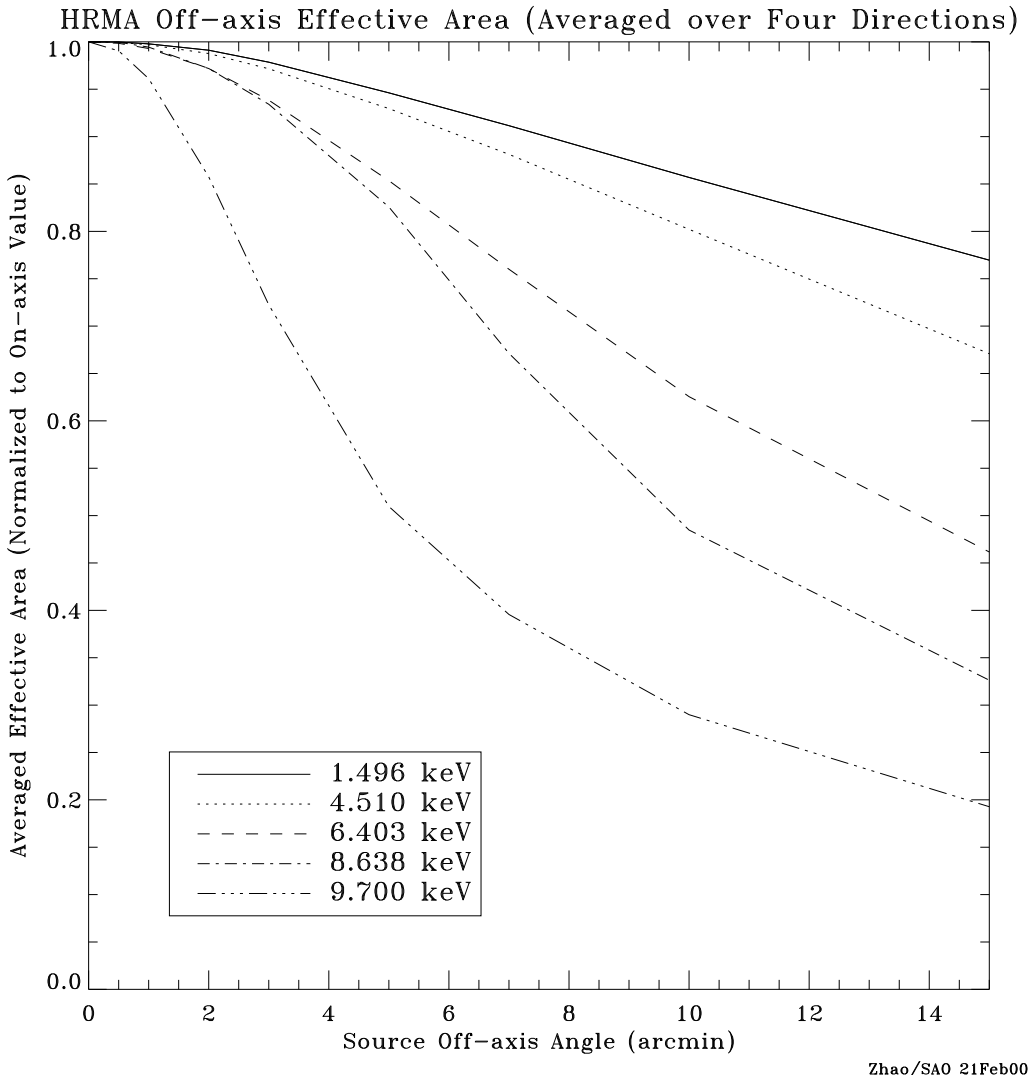

Figure 4.5: The HRMA effective area versus off-axis angle, averaged over four azimuthal directions, for selected energies, normalized to the on-axis area for that energy.

Table 4.2: HRMA Encircled Energy Performance

X-ray:		Encircled Energy Fraction	
E	λ	Diameter	
keV	Å	1 arcsec	10 arcsec
0.1085	114.2712	0.7954	0.9979
0.1833	67.6401	0.7937	0.9955
0.2770	44.7597	0.7906	0.9929
0.5230	23.7064	0.7817	0.9871
0.9297	13.3359	0.7650	0.9780
1.4967	8.2838	0.7436	0.9739
2.0424	6.0706	0.7261	0.9674
2.9843	4.1545	0.6960	0.9560
3.4440	3.6000	0.6808	0.9479
4.5108	2.7486	0.6510	0.9319
5.4147	2.2898	0.6426	0.9300
6.4038	1.9361	0.6365	0.9344
8.0478	1.5406	0.5457	0.9185
8.6389	1.4352	0.5256	0.9151
10.0000	1.2398	0.4971	0.8954

4.2.3 Point-Spread-Function and Encircled Energy Fraction

The *Chandra* HRMA point-spread function (PSF) has been simulated with numerical raytrace calculations based upon the mirror model previously discussed. A most useful parameter is the encircled energy fraction (the two-dimensional integral of the PSF) as a function of radius from the image center. The PSF and the encircled energy fraction for a given radius depend upon off-axis angle and energy. The HRMA optical axis is defined for practical purposes as—and calibrated in flight to be—the direction of the sharpest PSF. The PSF broadens and the encircled energy fraction decreases as (1) the off-axis angle increases, because of mirror aberrations; and (2) the X-ray energy increases, because of increased X-ray scattering.

On-axis PSF

Figure 4.6 shows the encircled energy fraction as a function of image radius for an on-axis point source and for different energies. The resulting increase in image size with energy is apparent. Figure 4.7 shows the radii of selected encircled energy fractions as functions of energy for an on-axis point source. Table 4.2 lists the encircled energy fraction contained within one and ten arcsec diameters for an on-axis point source at different energies.

Pre-flight measurements and images taken at the XRCF show that there is a slight (≈ 500 μm) offset between the optical axes of the paraboloids and hyperboloids, and that

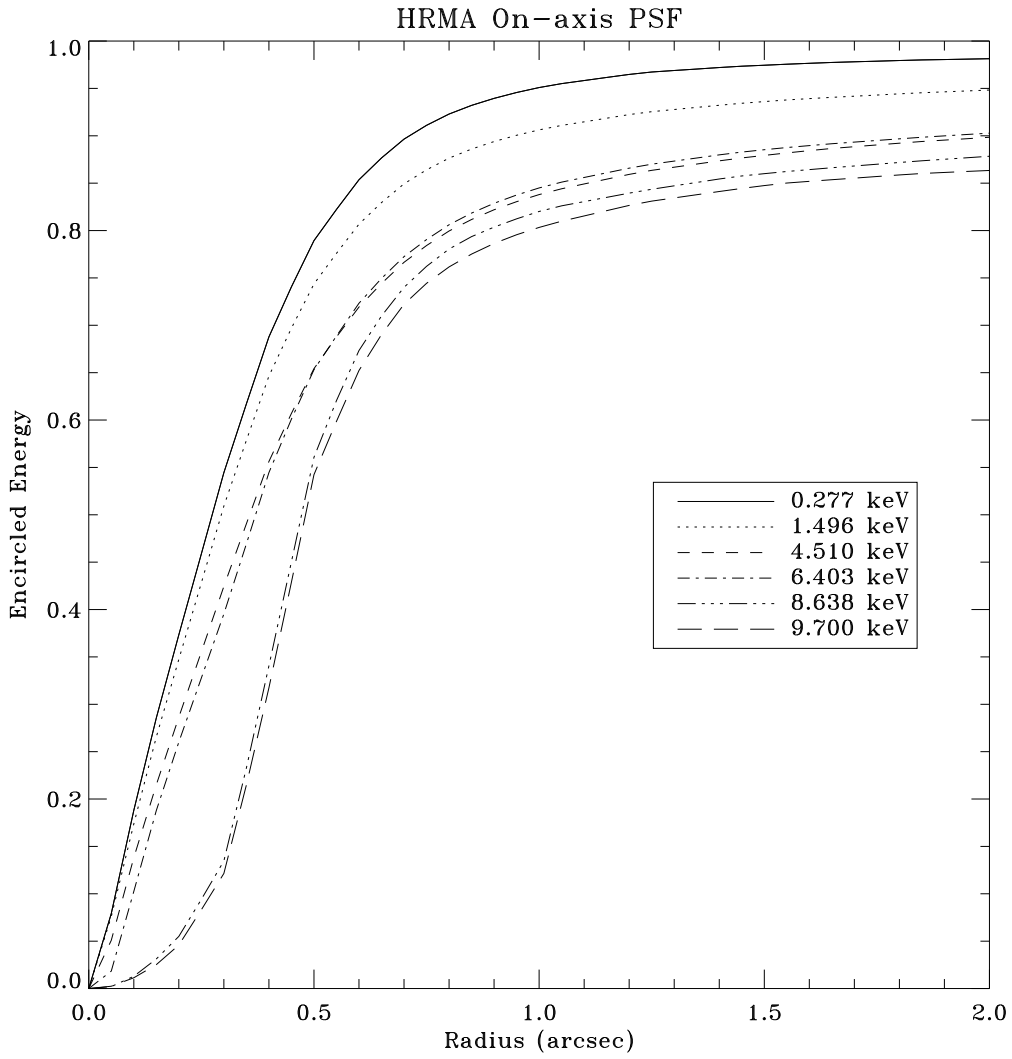

Figure 4.6: Fractional encircled energy as a function of angular radius, calculated for an on-axis point source, at selected X-ray energies. The curves are the combined response and centered at the common focus of the full HRMA, i.e. four nested mirror pairs. For higher energies (8.638 keV and 9.700 keV), the curves are broadened at small radii. This is because the focus of higher energies does not coincide with the HRMA common focus, but is offset by about 0.2 arcsec, due to a slight tilt of the HRMA mirror pair 6.

pair 6 is slightly tilted with respect to the other three. Consequently, the image from mirror pair 6 is not as symmetrical as the images produced by the other shells. The effect of this asymmetry on images depends on energy because of the different relative contribution of mirror pair 6.

Figure 4.8 shows simulated HRMA/HRC-I images at several energies. The effect of the mirror pair 6 alignment errors can be seen in the higher energy images as then mirror

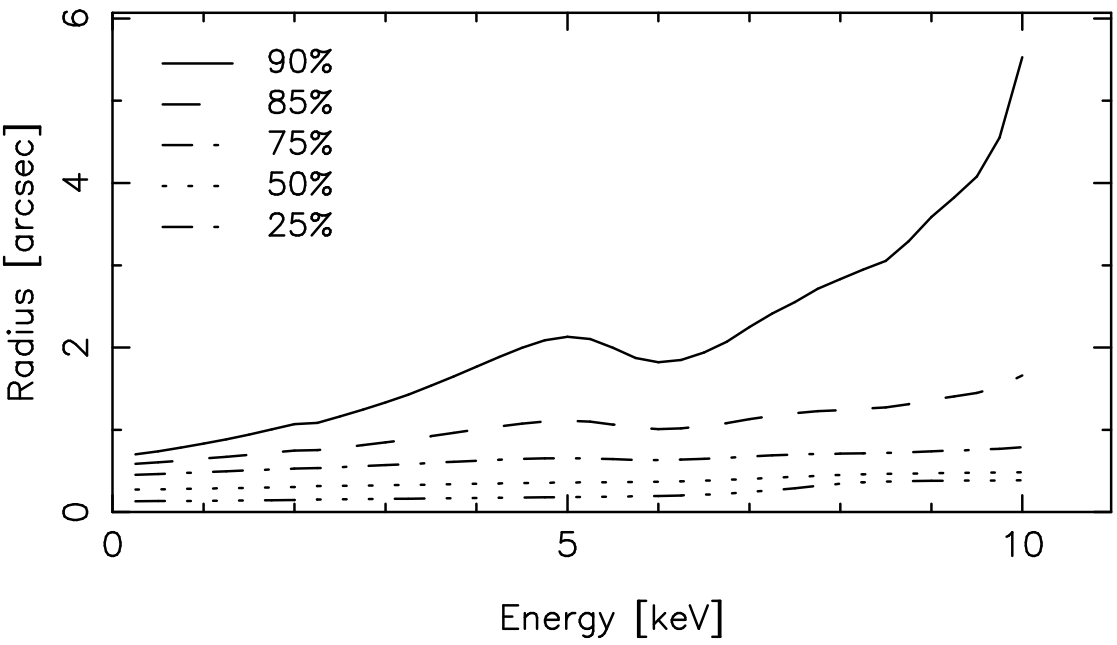

Figure 4.7: The radii of selected encircled energy fractions as functions of X-ray energy for an on-axis point source, calculated from the mirror model derived from ground-based calibration data.

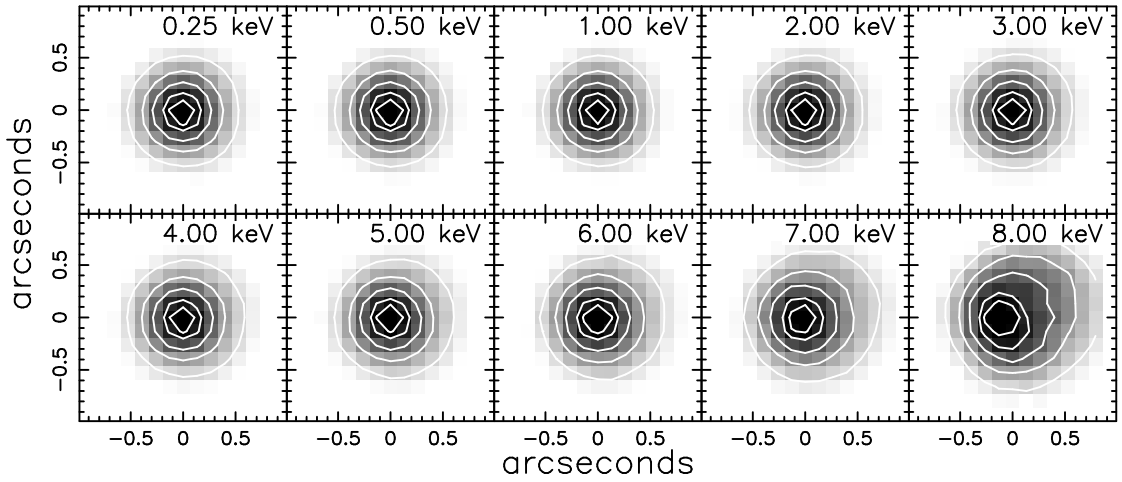

Figure 4.8: Simulated on-axis HRMA/HRC-I images of on-axis mono-energetic point sources with aspect blurring. The grayscale is a linear stretch; surface brightness contours are at 90%, 80%, 60%, 40%, and 20% of the peak brightness. The 8 keV image core is asymmetric and off-center due to the shell 6 misalignment.

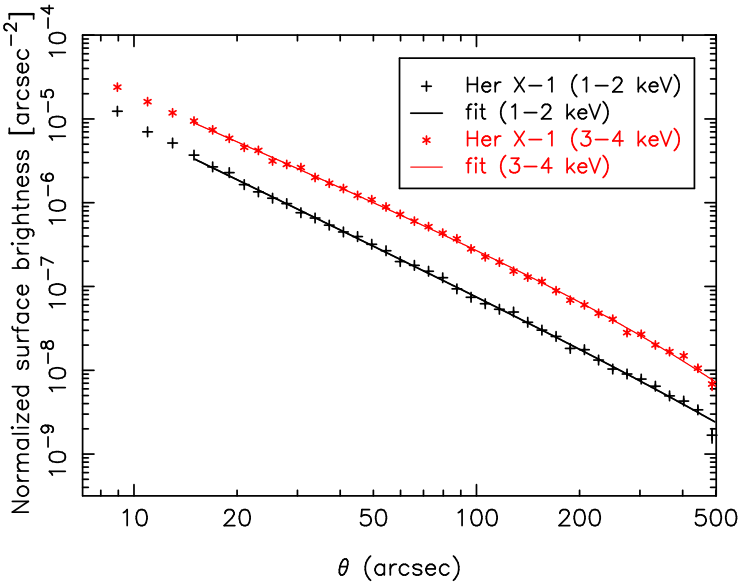

Figure 4.9: Normalized radial profiles (units arcsec^{-2}) of the Her X-1 scattering wings. The shapes are well determined beyond 15 arcsec, but the overall normalization may be uncertain by a factor of 2. The lower curve shows the 1.0–2.0 keV profile and the upper curve shows the 3.0–4.0 keV profile. In each case, the heavy solid line is a power-law plus exponential cutoff fit to the Her X-1 data.

pair 6 becomes the dominant contributor to the total effective area. Note the movement of position of the core as well as the asymmetric flaring. The ~ 0.2 arcsec core motion is comparable to other factors of image degradation encountered in flight, such as uncertainties in the aspect solution.

The HRMA PSF has a faint halo extending to large angles, resulting from X-rays scattering from micro-roughness on the mirror surfaces. This scattering is energy dependent; the spectrum of the scattered X-rays hardens significantly with increasing angle from the source. An empirical model was generated based on the ground calibration measurements; a number of systematic effects remain to be accounted for, and the uncertainties in the flux in the wings are probably at least 30–50%. This model is described more fully in `https://cxc.harvard.edu/cal/Hrma/XRCF-Wings.html`. A deep calibration observation of Her X-1 (ObsID 3662) was obtained in order to improve the understanding of the PSF wings. The SIM was shifted to move the optical axis to ~ 1 arcmin from the edge of the S3 detector furthest from the frame store; a Y-offset moved the image ~ 1 arcmin into node 0 of the detector. The resulting pointing is ~ 45 arcsec off-axis, effectively on-axis with regard to the mirror scattering properties. The analysis is discussed in more detail in `https://cxc.harvard.edu/cal/Hrma/rsrc/Publish/Optics/PSFWings/wing_analysis_rev1b.pdf`; see also `https://cxc.harvard.edu/cal/Hrma/UsersGuide.html`.

Radial profiles of the Her X-1 scattering wings are plotted for energy bands 1.0–2.0 keV and 3.0–4.0 keV in Fig. 4.9. These are surface brightness profiles normalized by the source count rate estimated from the transfer streak spectrum. The units of the normalized

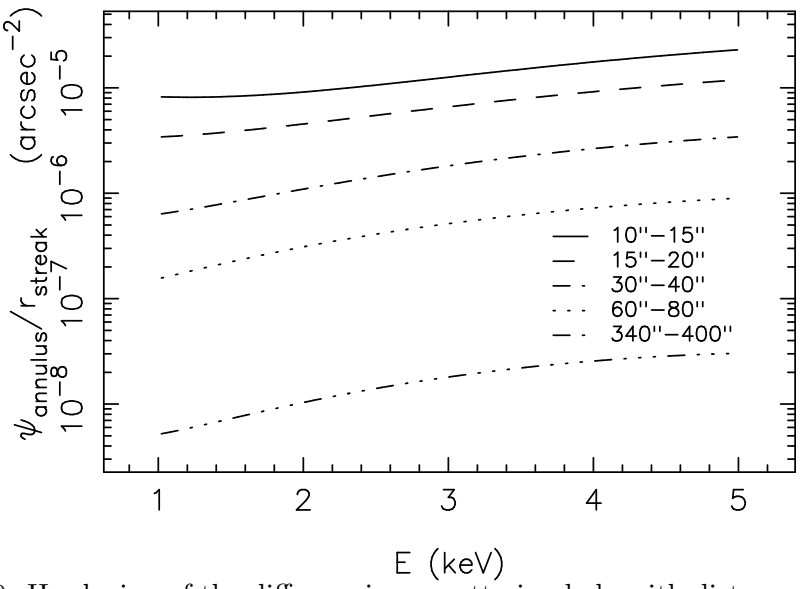

Figure 4.10: Hardening of the diffuse mirror scattering halo with distance from the direct image. The spectra for a number of annuli were normalized by the area of the extraction regions (taking into account chip edges) giving $\psi_{annulus}$ cts s^{-1} keV^{-1} arcsec^{-2}. These were each divided by a spectrum evaluated using the transfer streak events for the direct image, r_{streak} cts s^{-1} keV^{-1}. The legend indicates annular radii in arcsec.

profiles are in arcsec^{-2}. A fit (power-law plus exponential cutoff) is over-plotted, and the fit applies for $\theta > 15$ arcsec. Inside 15 arcsec, the profiles are increasingly depressed because of the effects of pile-up in this very bright source. The shape of the wing profile is well represented (beyond 15 arcsec), but the overall normalization may be off up to a factor of two.

Because the mirror scattering is in part diffractive, the diffuse mirror scattering halo is energy dependent. Spectra extracted from the diffuse mirror scattering wings of the PSF are significantly modified from the spectrum of the incident source X-rays. Generally, the scattering halo spectrum becomes harder with increasing angle from the source. Fig. 4.10 shows the ratio of diffuse spectra extracted from annuli centered on the specular image of Her X-1 (normalized by extraction region area) to the corresponding spectrum extracted from the ACIS transfer streak for the source; the transfer streak spectrum is thought to be ∼4% piled up.

Off-axis PSF

The PSF broadens for off-axis sources, and there is considerable distortion in the image even if the HRMA were perfect. This distortion is due to the aberrations of Wolter Type-I optics and to the different focal surfaces (Figure 4.11) for the four mirror pairs. The increase in image size with off-axis angle is greatest for the inner shell, and hence is larger for higher X-ray energies.

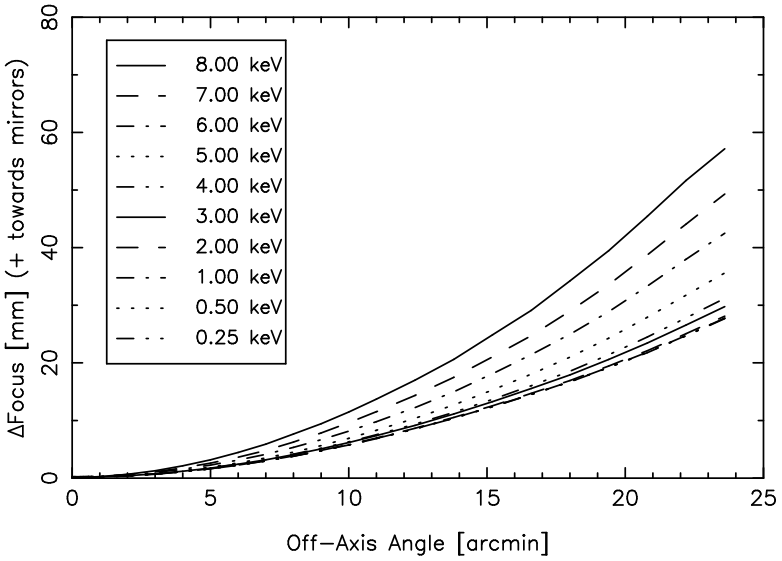

Figure 4.11: The HRMA focal surface (from simulations), indicating its dependence on energy and off-axis source position.

Figure 4.12 shows the dependence of encircled energy radii on off-axis angle on the HRC-I with the HRMA focus at the HRC-I aimpoint. Because the HRC-I is axially symmetric with respect to the HRMA optical axis, the off-axis encircled energy radii are almost azimuthally symmetric, except some small asymmetry due to the imperfect HRMA as mentioned above. The figure gives the averaged radii for 1.49 keV and 6.40 keV at 50% and 90% encircled energy. The blurs due to the HRC-I spatial resolution and the aspect solution, estimated to be FWHM = 0.22 arcsec, are included.

The ACIS-I surface is not axially symmetric with respect to the HRMA optical axis, because the HRMA aimpoint is located near the inner corner of one of the four ACIS-I chips: I3. Thus the off-axis encircled energy radii are not azimuthally symmetric. Figure 4.13 shows the dependence of encircled energy radii on off-axis angle on the four ACIS-I chips. The figure gives the encircled energy radii for 1.49 keV and 6.40 keV at 50% and 90% encircled energy in four azimuthal directions—from the aimpoint to the outer corners of the four ACIS-I chips. The blurs due to the ACIS-I spatial resolution and the *Chandra* aspect error are included.

Figures 4.14 and 4.15 illustrate the effect of aberrations on images of off-axis point sources at 1.49 keV and 6.4 keV. The images are simulations of the HRMA alone, projected to the HRC-I detector plane. The degradation in image quality is primarily due to the separation between the detector plane and the effective focal plane, which is a strong function of both energy and off-axis angle (see Figure 4.11). Cusps in the HRMA images are due to a slight misalignment of the parabolic and hyperbolic mirrors. The signal in these figures is much higher than what might be expected in an actual observation. Figure 4.16 shows how the morphology of an off-axis image varies with the number of counts in the image. It is very easy to mistakenly conclude that an off-axis source is extended or has several components, even with a large number of counts.

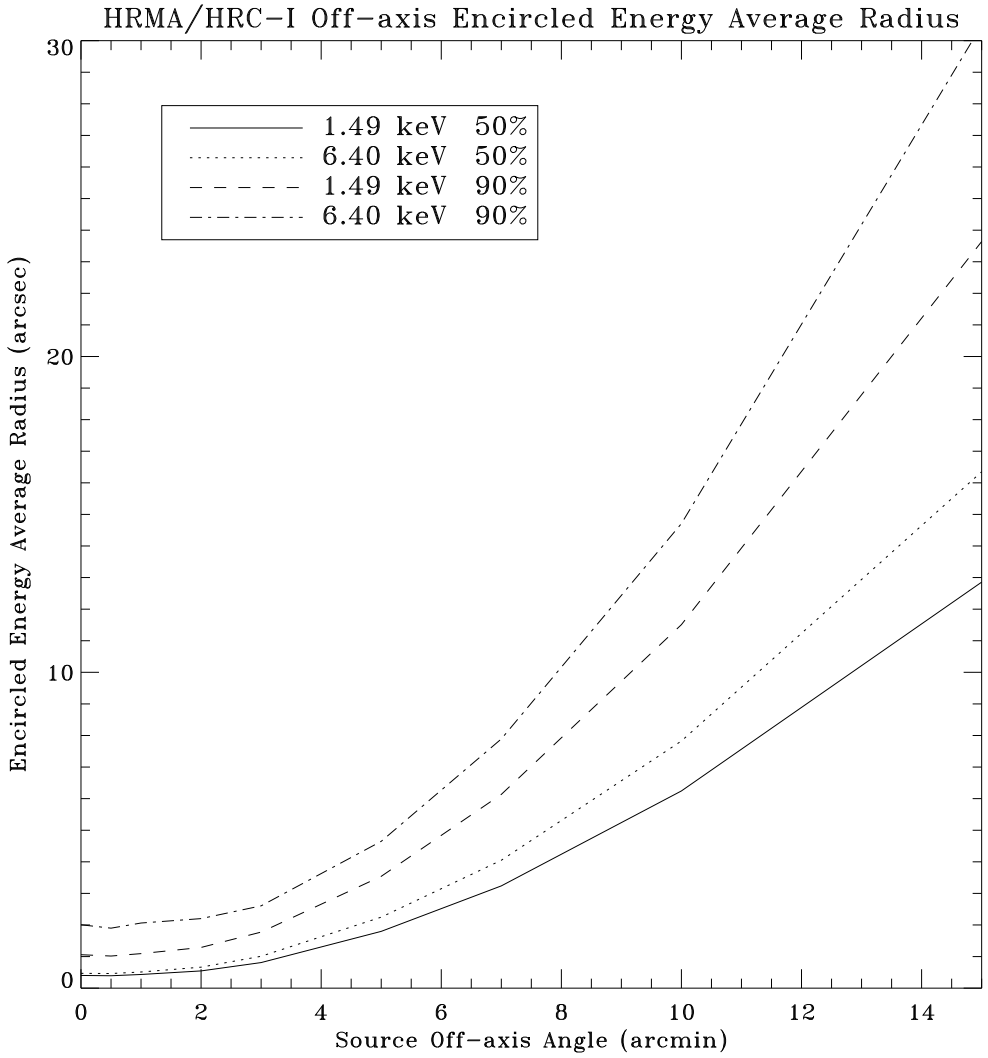

Figure 4.12: The HRMA/HRC-I encircled energy average radii for circles enclosing 50% and 90% of the power at 1.49 and 6.40 keV as a function of off-axis angle. The HRC-I surface is a flat plane perpendicular to the optical axis, which does not follow the curved *Chandra* focal plane. These curves include the blurs due to the HRC-I spatial resolution and the *Chandra* aspect error.

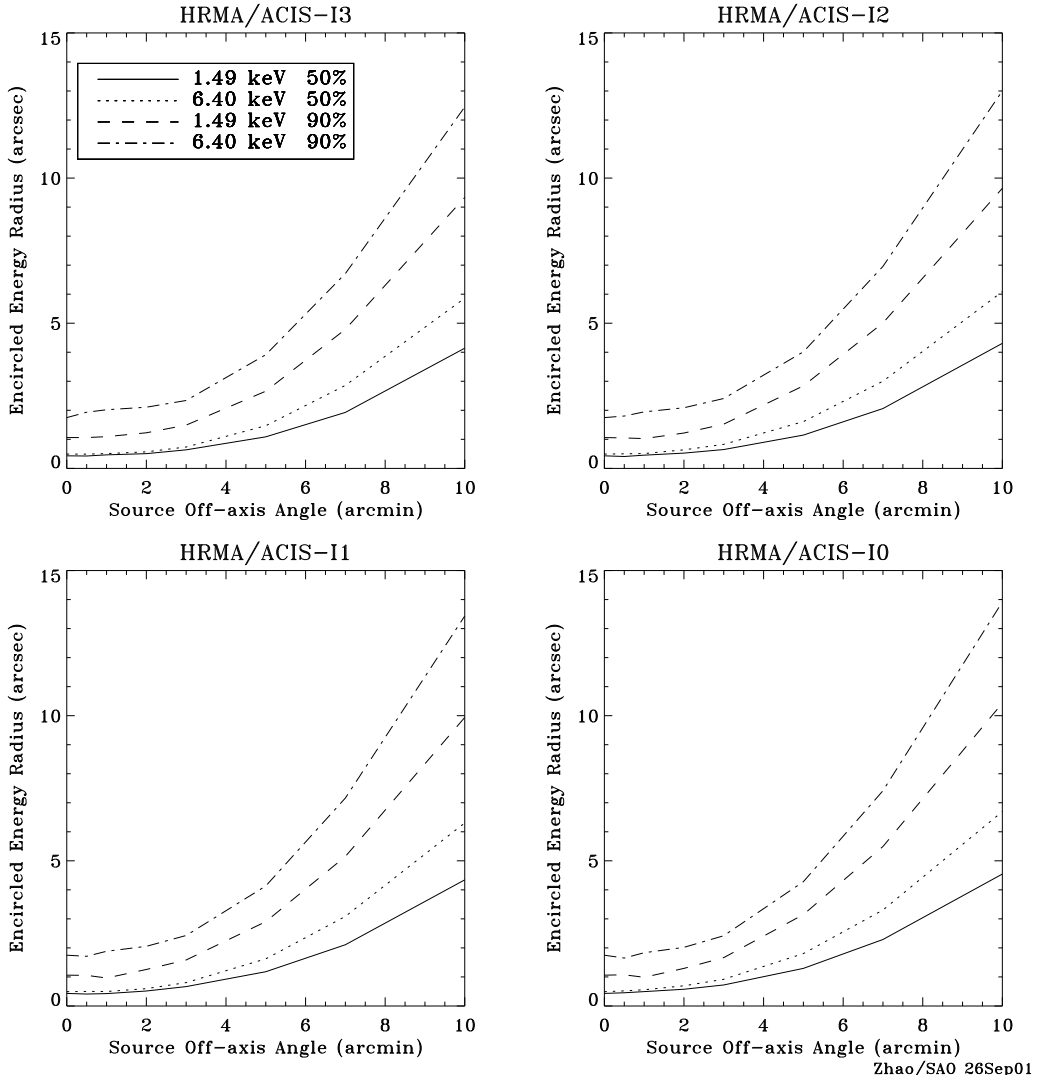

Figure 4.13: The HRMA/ACIS-I encircled energy radii for circles enclosing 50% and 90% of the power at 1.49 and 6.40 keV as a function of off-axis angle. The ACIS-I surface is composed by four tilted flat chips which approximate the curved *Chandra* focal plane. The HRMA optical axis passes near the aimpoint which is located near the inner corner of chip I3. Thus the off-axis encircled energy radii are not azimuthally symmetric. The four panels show these radii's radial dependence in four azimuthal directions—from the aimpoint to the outer corners of the four ACIS-I chips. These curves include the blurs due to the ACIS-I spatial resolution and the *Chandra* aspect error.

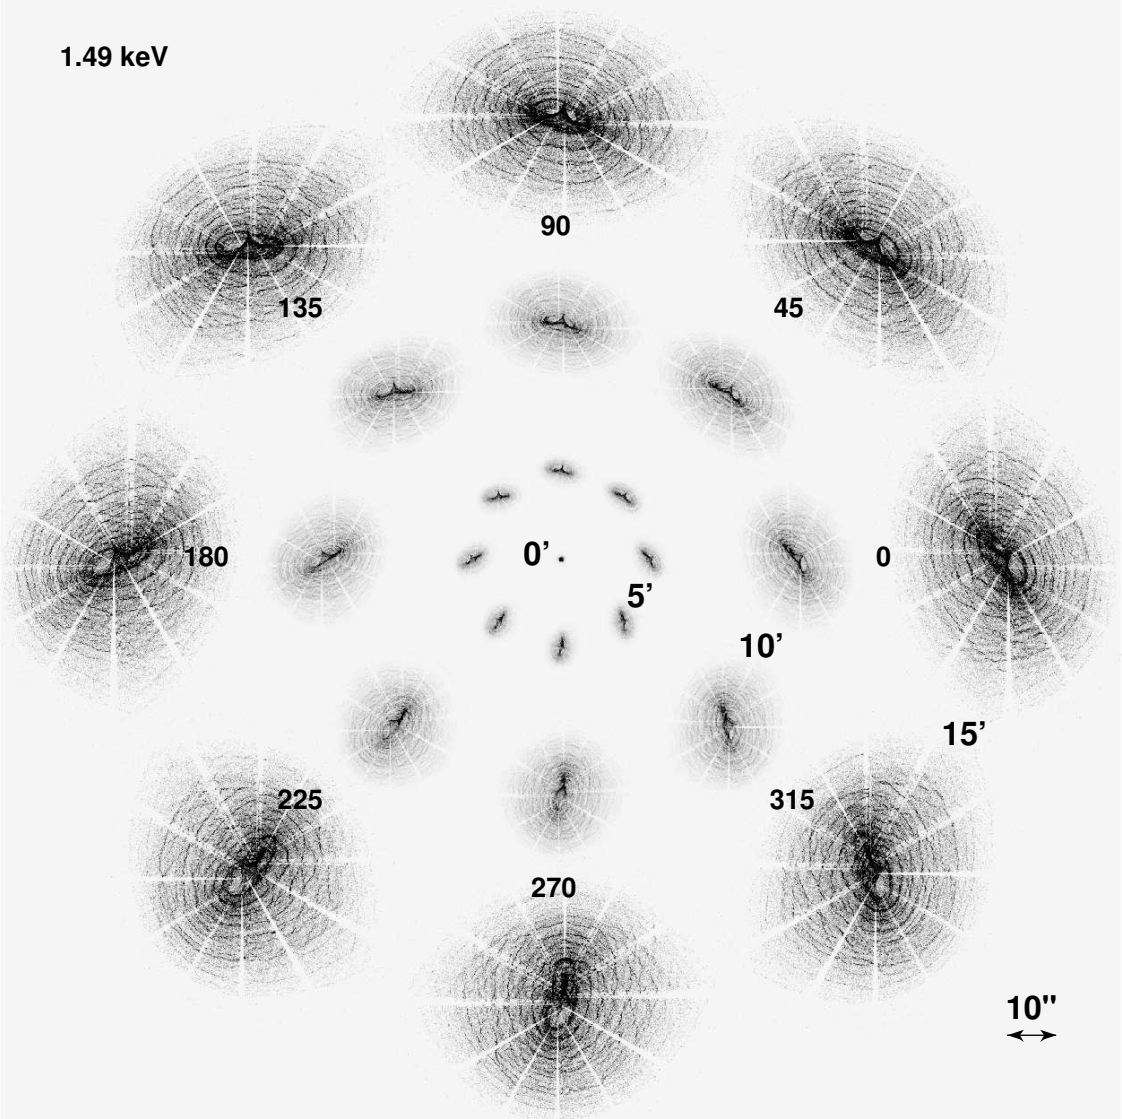

Figure 4.14: Simulated 1.49 keV images, for the HRMA only. Images are shown with a linear stretch, as they would appear on the sky, at three off-axis angles (5 arcmin, 10 arcmin, and 15 arcmin) and various azimuths. The images are all to the same scale, illustrated by the scale bar. The spacing between images is arbitrary. The surface brightness of the images at 10 arcmin and 15 arcmin has been enhanced to show structure. Spokes in the images are due to shadowing by mirror support struts. Cusps are due to a slight misalignment of the parabolic and hyperbolic mirrors. These simulations are at an effective roll of zero—observations should be "de-rolled" before comparison to these images.

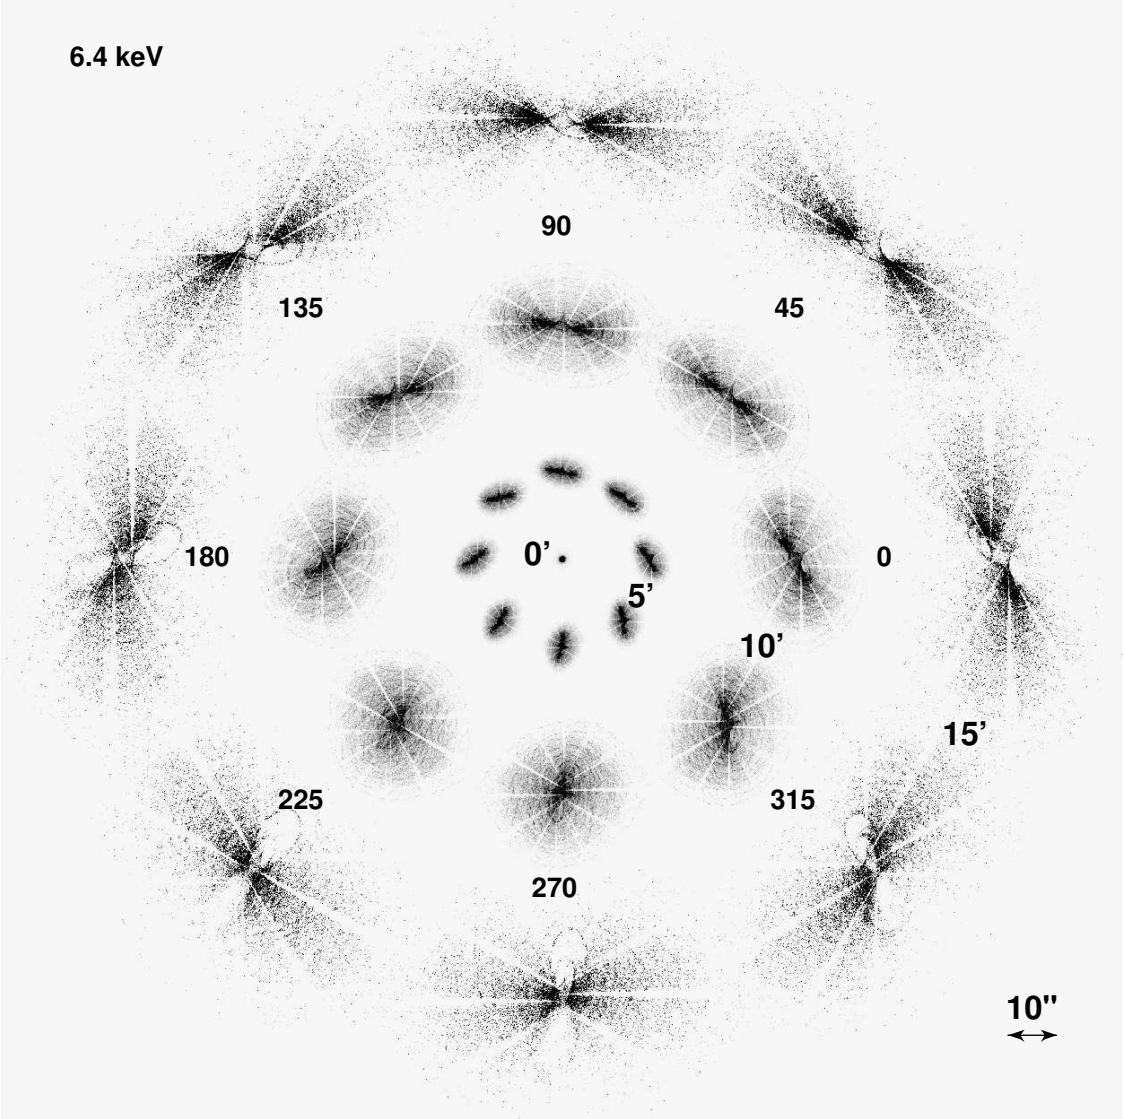

Figure 4.15: Simulated 6.4 keV images, for the HRMA only. Images are shown with a linear stretch, as they would appear on the sky, at three off-axis angles (5 arcmin, 10 arcmin, and 15 arcmin) and various azimuths. The images are all to the same scale, illustrated by the scale bar. The spacing between images is arbitrary. The surface brightness of the images at 10 arcmin and 15 arcmin has been enhanced to show structure. Spokes in the images are due to shadowing by mirror support struts. Cusps are due to a slight misalignment of the parabolic and hyperbolic mirrors. These simulations are at an effective roll of zero—observations should be "de-rolled" before comparison to these images.

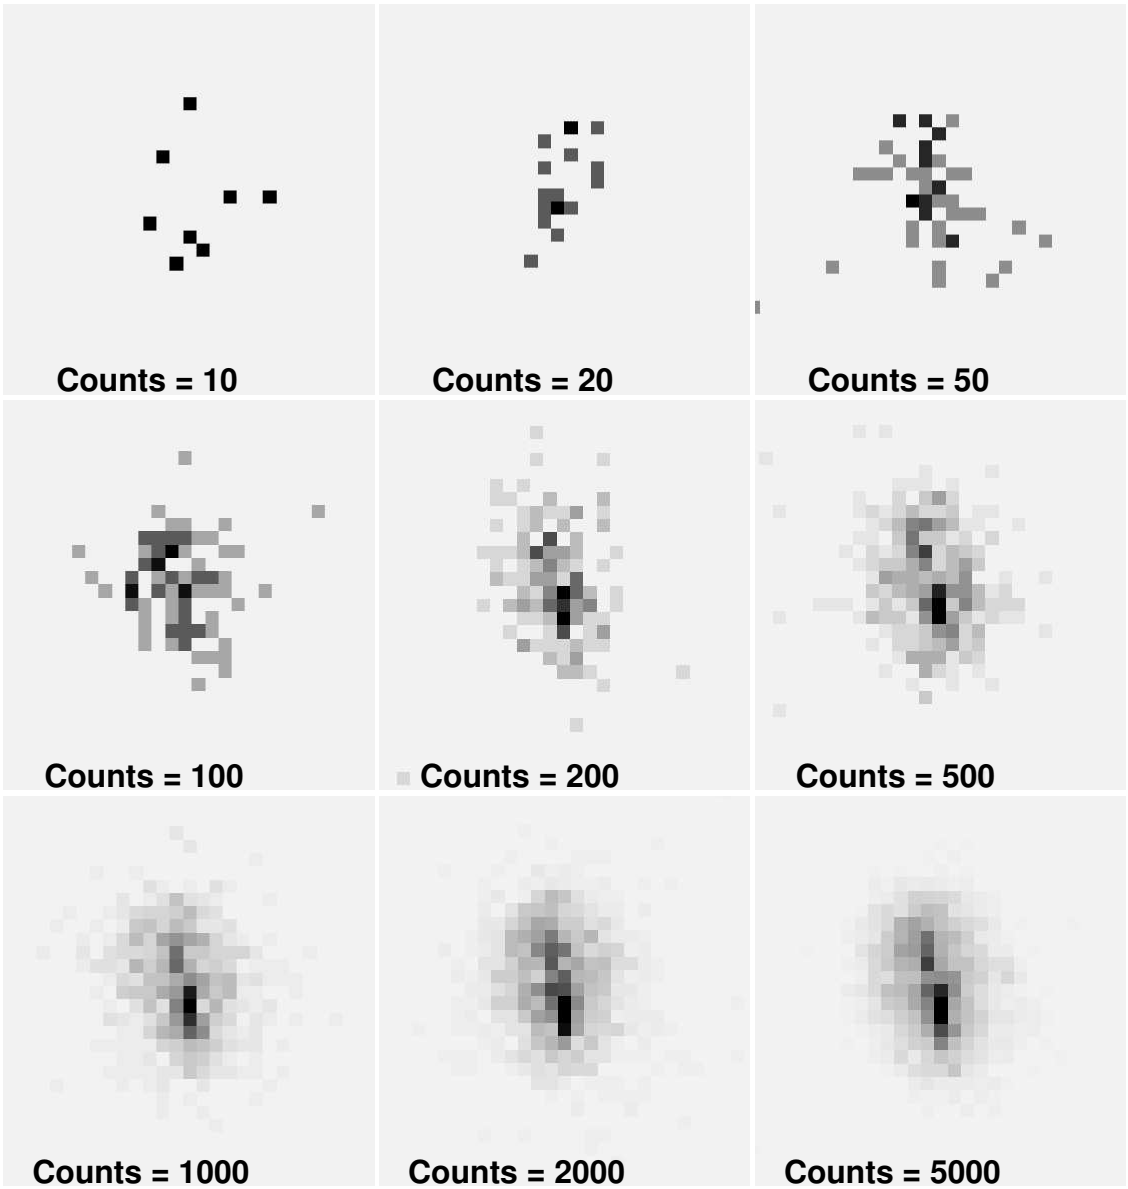

Figure 4.16: A simulated 1.49 keV point source at an off-axis angle of 5 arcmin, binned to ACIS pixels. The panels show what the source would look like with a varying number of counts. Note how the apparent morphology is a strong function of the number of counts, and how even with a large number of counts one might mistake it for an extended source or even for multiple sources.

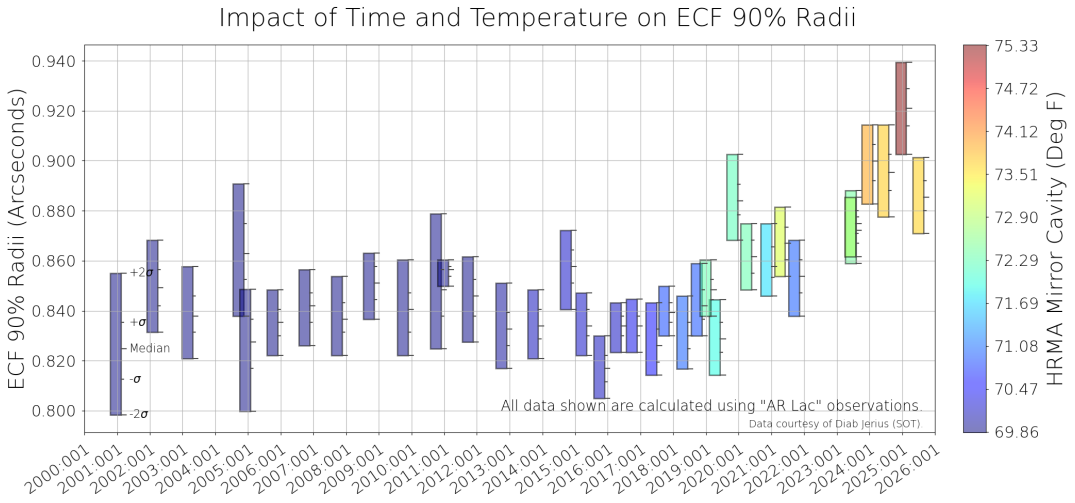

Figure 4.17: Evolution of the HRMA 90% ECF over time and its relationship to the HRMA cavity temperature. Measurements are useful only for relative longitudinal comparisons.

PSF Stability

The statistic used to track the Chandra PSF's stability is the radius enclosing 90% of the counts from a point source (the Enclosed Count Fraction, or ECF), corrected for known spacecraft and detector systematics. Changes in this statistic should be representative of intrinsic changes in the optics' performance. Determinations of the 90% ECF at different mission epochs are made utilizing annual or semi-annual on-axis observations of AR Lac on the HRC-I.

Figure 4.17 presents the 90% ECF radius annotated with the HRMA cavity temperature. The radius has broadened very slightly since 2020 by $\approx 0.05''$. While the HRMA temperature has not been immune to the degradation of thermal control of the telescope, changes in it do not appear to be correlated with changes in the PSF.

Changes in the overall PSF are small and should not measurably impact science.

PSF Anomalies

An anomaly was discovered in the observed PSF during deconvolution of on-axis HRC-I observations of AR Lac and Capella (Juda & Karovska 2010; see Section 7.6). The images showed hook-shaped residuals that developed c.2003. The anomalous features are located $\approx$ 0.8 arcsec from the centroid, towards mirror spherical coordinate, $\phi = 285 \pm 25$ (see the CIAO caveats page `https://cxc.harvard.edu/ciao/caveats/psf_artifact.html`), slightly offset from the spacecraft +Z direction (see Figure 1.2). The existence of an asymmetry in the PSF profile, consistent with the HRC results, was subsequently verified with ACIS data of low count-rate, high-exposure, on-axis point sources (Kashyap 2010;

see also Section 6.6.1). This anomaly, apparent as an asymmetry in the profile of the PSF (see Figures 4.18,4.19,4.20), is as yet unexplained and cannot be modeled with raytraces. Modeling of jets or other structure in the source aligned along this direction will be affected by this anomaly.

4.3 Ghost Images

Baffles prevent non-reflected or singly reflected photons from impinging on the focal plane within the central 30 arcmin diameter region of the field of view. Outside of this region, however, singly reflected photons from strong off-axis sources may appear. The spray of singly reflected photons is faint relative to the direct image, but can be quite complex. Each individual paraboloidal or hyperboloidal mirror can generate its own single-reflection ghosts. These form loops sweeping in toward the center of the focal plane as the source off-axis angle increases. The ghost loops from the smallest mirrors are the first to approach the central regions as source off-axis angle increases. With increasing source off-axis angle, the large mirrors come into play. As a loop approaches the central 30 arcmin diameter region of the field of view, the inner parts of the loop fade and break up.

These single-reflection ghosts can impinge on the detector even if the source itself does not fall within the detector field of view. These ghosts mainly affect the outermost portions of those detectors which extend to large off-axis angles: HRC-I, and the spectroscopy arrays, HRC-S and ACIS-S. Figure 4.21 shows simulated ghost images on the ACIS-S array. Point sources were simulated at a range of off-axis angle θ and at a fixed off-axis azimuth ($\phi = 5^\circ$). The effects discussed above (e.g. fading of the loops as they approach the central field) can be seen in comparing the ghosts in the 30 arcmin–32.5 arcmin–35 arcmin sequence, or in the 50 arcmin–52.5 arcmin sequence.

Imaging observations with HRC-I or spectroscopy observations with HRC-S or ACIS-S which are near very bright sources can be checked using *Chandra* Ray Tracer (ChaRT) or Model of AXAF Response to X-rays (*MARX*) raytraces to determine whether single-reflection ghost images are likely to be a problem.

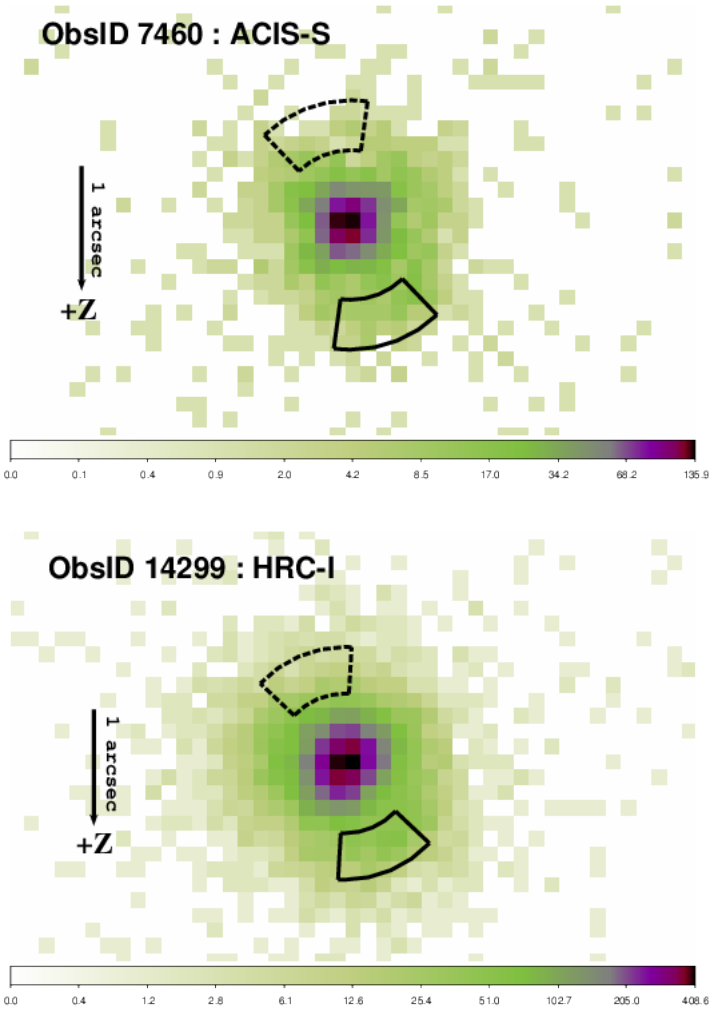

Figure 4.18: The PSF anomaly, illustrated with ACIS-S (top) and HRC-I (bottom) images. The ACIS-S image is of a Low Mass X-ray Binary in NGC 6397, and the HRC-I image is of AR Lac. The ACIS image is binned at $\frac{1}{4}$ sky-pixel resolution, similar to the HRC resolution. The datasets have been derolled such that the spacecraft +Z axis points vertically down (indicated by a vertical arrow of length 1 arcsec). The part of the PSF that is affected by the anomaly is shown by the pie-shaped region (solid boundary) computed with the CIAO tool `make_psf_asymmetry_region`. The same region, rotated diametrically across from image centroid is also shown (dashed boundary) for comparison. The asymmetry in the counts in the two regions is apparent; there are more counts within the region located over the anomaly than in the region away from the anomaly.

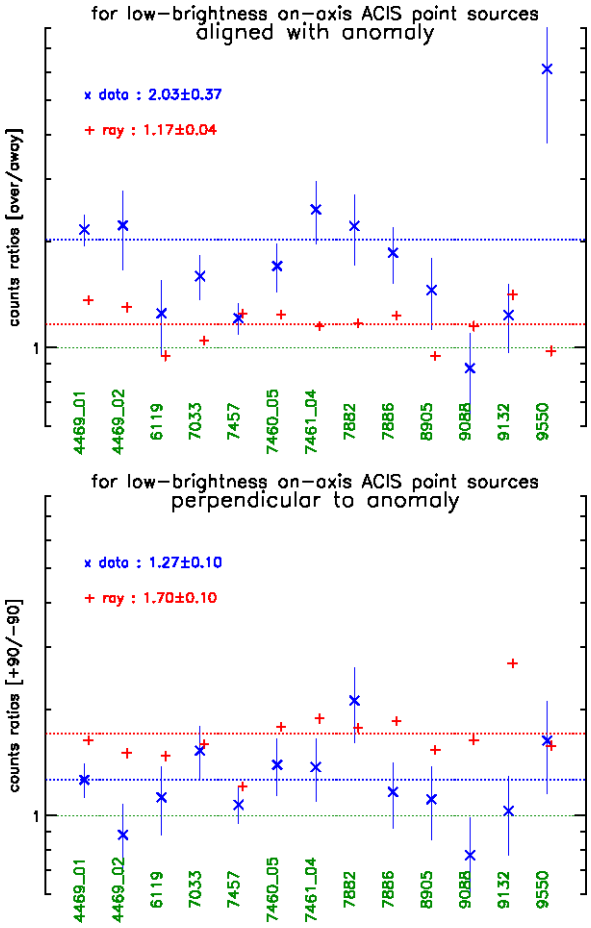

Figure 4.19: The ratio of counts collected from different segments around the PSF for a number of low-count-rate on-axis point sources (see Kashyap 2010). These are sources for which pile-up effects are minimal, and which have no previously known intrinsic features. The counts are collected in pie-shaped regions *over* and *away* from the anomaly (top; see Figure 4.18) and from similar regions rotated $\pm 90^\circ$ from them (bottom). The ratios for counts are shown for both the data ('x'; blue points) and raytrace ('+'; red points). The error bars are shown as thin vertical lines. The average of the ratios computed for all the sources are also shown, along with the standard errors on the mean. There is a clear asymmetry in the PSF profile of the data along the nominal direction of the anomaly that is neither present in the perpendicular direction, nor accountable by raytraces. Comparison with the raytrace model PSFs is necessary to establish the reality of the anomaly, since the *Chandra* PSF is not azimuthally symmetrical and displays intrinsic asymmetries. See, for instance, the ratios along the perpendicular direction, in the lower figure, where they are higher for the raytrace model than for the data, signifying that there are as yet unexplained differences between the data and the raytrace model. The mechanism by which these discrepancies arise is under investigation.

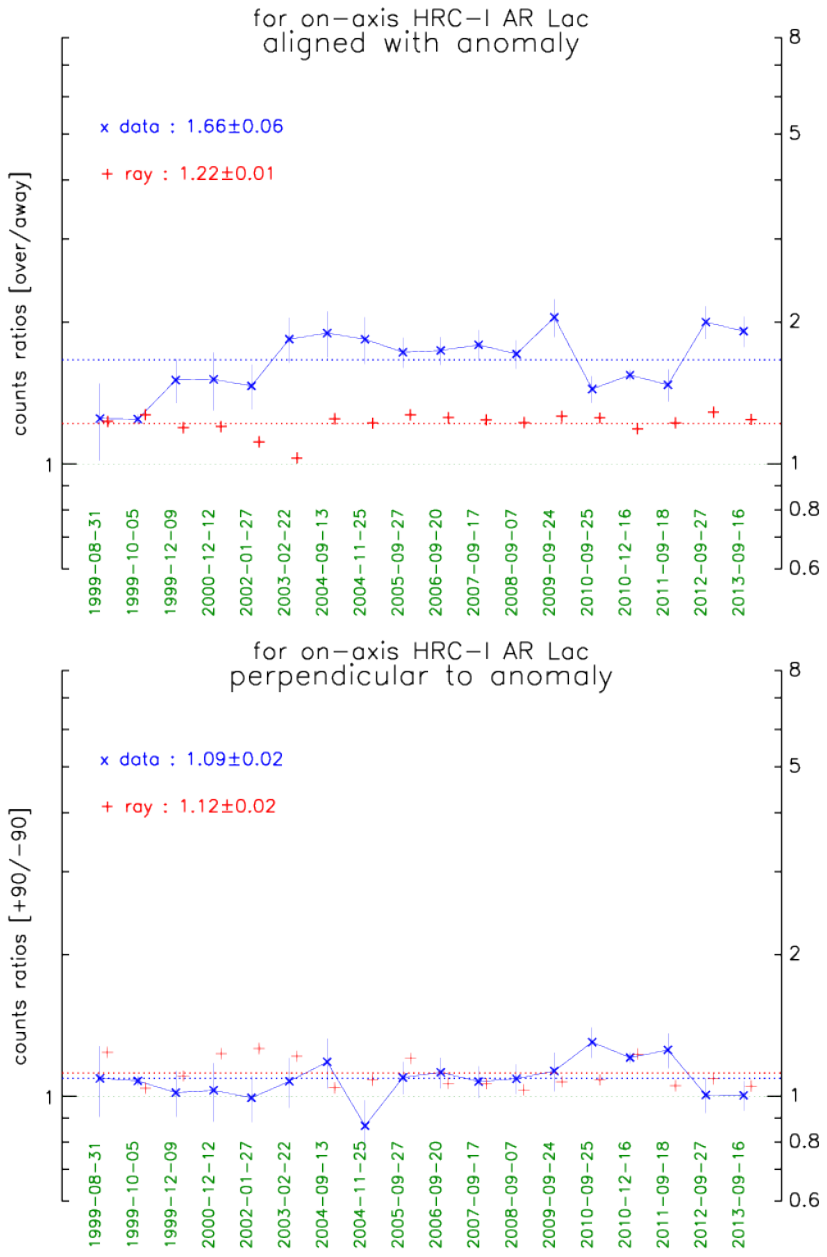

Figure 4.20: As Figure 4.19, for on-axis HRC-I observations of AR Lac. Notice that the magnitude of the anomaly was low early in the mission and appears to be fully formed by 2003-Feb. There are also indications that the asymmetry was reduced between late-2010 and late-2011. The raytrace PSF model for the HRC includes smoothing due to detector blur.

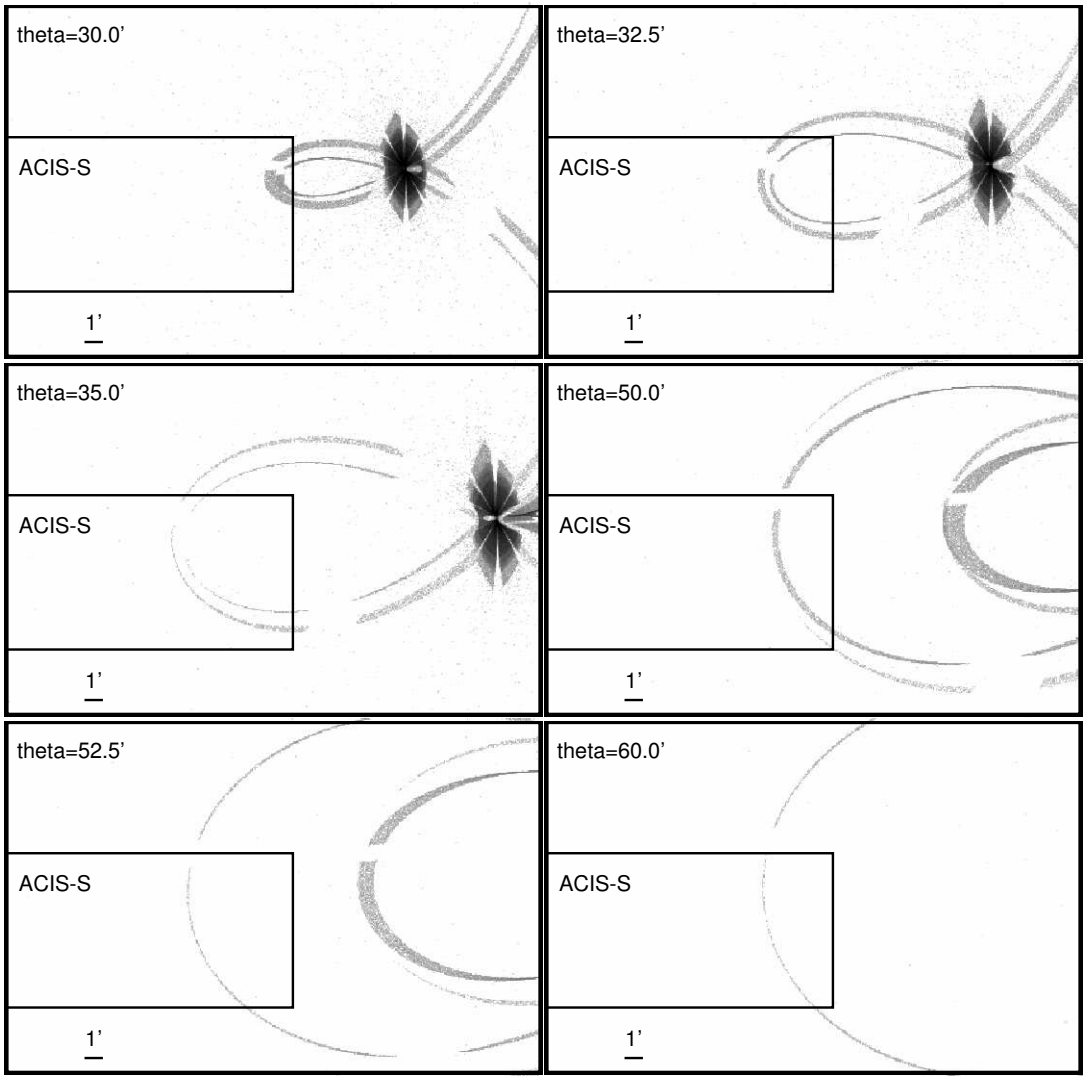

Figure 4.21: Simulated images of off-axis sources. The off-axis angle, theta (θ, in arcmin), is indicated, and all simulations were performed for the same value of ϕ (5°). The rectangle indicates the footprint for one end of the ACIS-S detector. These simulations illustrate how singly reflected photons can hit the detector even when the specular image is well outside the field of view. The surface brightness of these ghosts is low relative to the brightness of the X-ray sources, but could be relevant in planning observations near extremely bright X-ray sources.

4.4 Effects of Aspect and Instrument Uncertainties

The HRMA performance discussed in the previous sections will be slightly degraded by uncertainties in the aspect solution and the details of the imaging detector spatial response function. The ground software system also deliberately adds a small random position error to reduce image artifacts which result from instrument and data system integer location values. The randomization may be turned off in data processing if desired. These effects are illustrated for the HRC-I and HRC-S instruments in Figures 4.22 and 4.23 respectively. These figures also show the fractional encircled energy as a function of radius actually observed in flight compared to model calculations at 0.277, 1.496, and 6.403 keV. An aspect error of 0.22 arcsec (FWHM) was included in the model calculations. The observed encircled energy curve most resembles that calculated for a monoenergetic source at 1.5 keV because this is typically the energy of the average detected photon.

Similar calculations have been performed for the ACIS-S (S3) over a wider range of energies; the results are shown in Figure 4.24. The simulation accounted for the typical spacecraft jitter, so the location of the instrument pixel boundaries has little effect. There is, however, a small effect of the location of the source compared to the data system pixel boundaries. These particular calculations were performed for a point source centered on the boundary between two data system pixels. The ACIS-I instrument response is similar.

Figures 4.22, 4.23, and 4.24 may be compared with Figure 4.6 to estimate the image performance degradation due to non-HRMA effects.

ACIS PSF Broadening at Low Energies

There is some evidence that the soft ACIS PSF (up to $\sim 800\,\mathrm{eV}$) has been broadening since 2015. This broadening is discussed in more detail in Section 6.6.2.

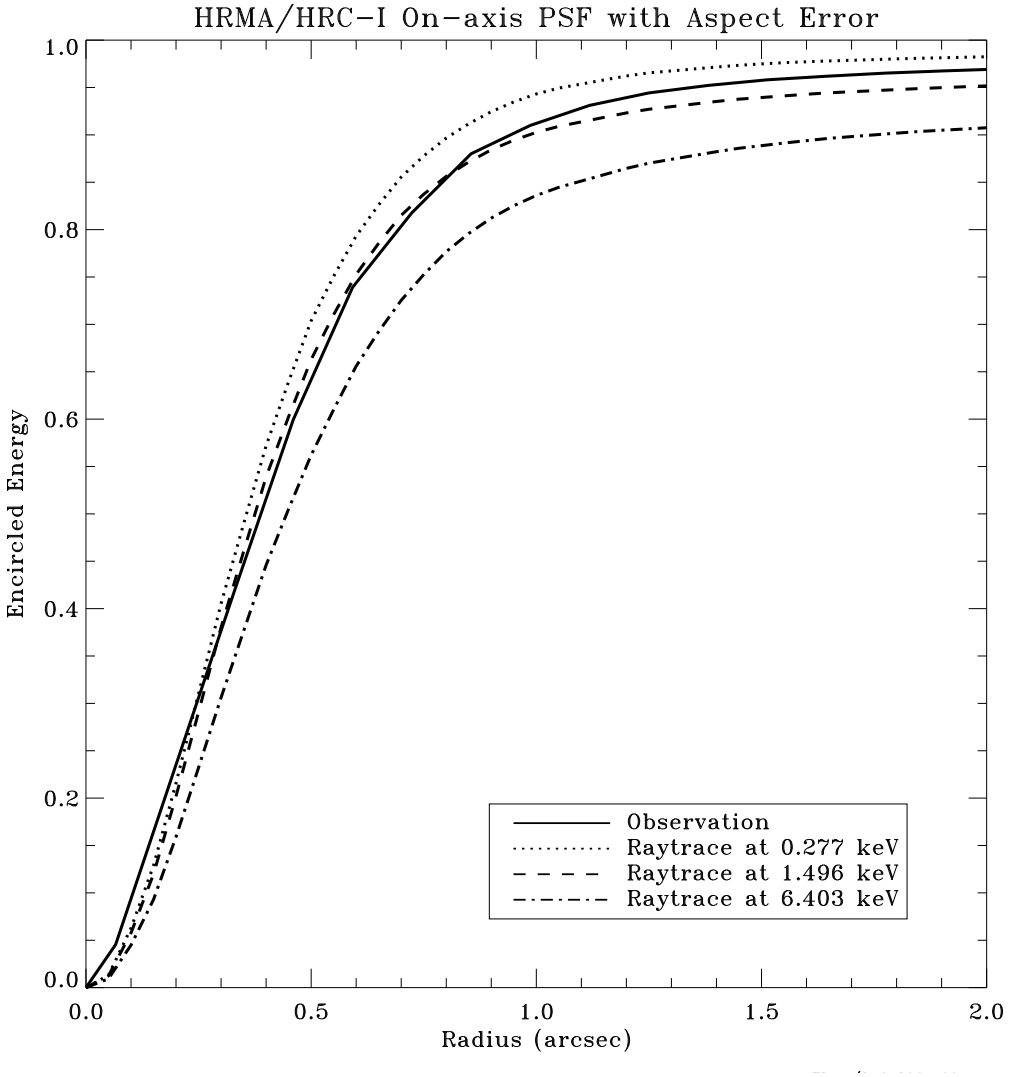

Figure 4.22: The HRMA/HRC-I on-axis fractional encircled energy as a function of angular radius from a point source (Ar Lac) observed in flight compared to raytrace simulations for an on-axis point-source at selected X-ray energies, including the aspect uncertainties and the HRC-I pixelization effects.

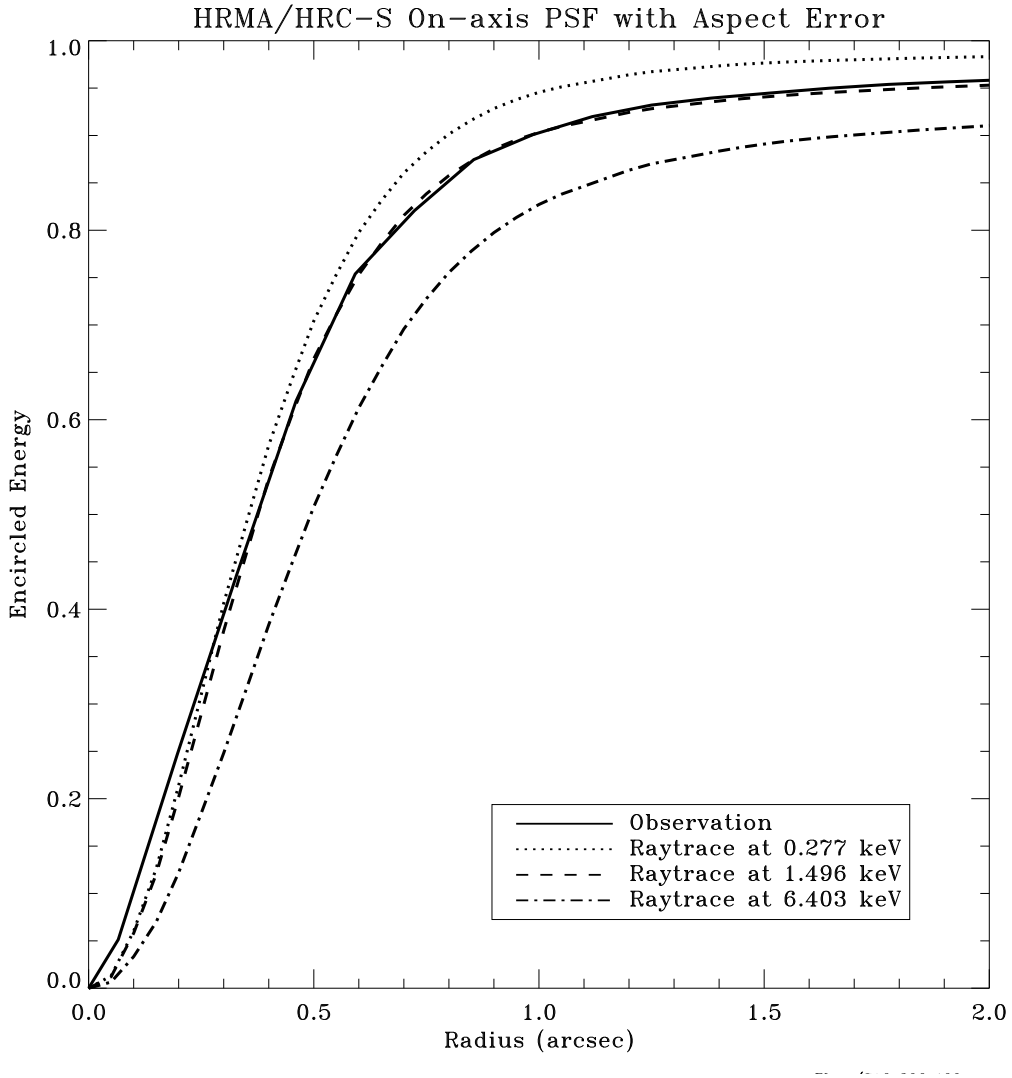

Figure 4.23: The HRMA/HRC-S on-axis fractional encircled energy as a function of angular radius from a point source (LMC X-1) observed in flight compared to raytrace simulations for an on-axis point-source at selected X-ray energies, including the aspect uncertainties and the HRC-S pixelization effects.

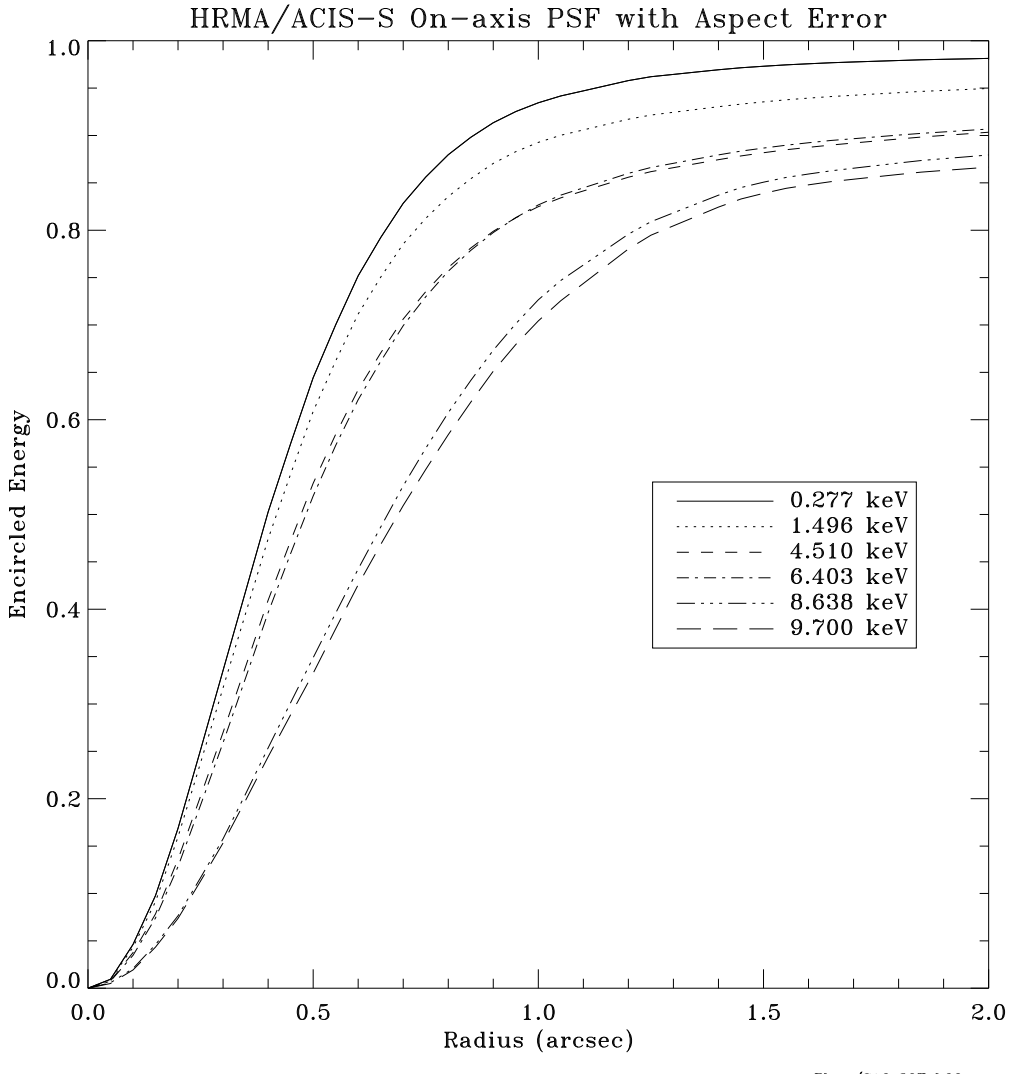

Figure 4.24: The fractional encircled energy as a function of angular radius expected for in flight ACIS-S (S3) measurements for an on-axis point-source at selected X-ray energies. The curves are the combined response of the four nested mirror pairs, typical aspect uncertainties, and the ACIS response function.

4.5 Optical Axis and Aimpoint

Tracking the positions of the *Chandra* optical axis and aimpoint with respect to the detectors is the most relevant measurement for understanding the telescope stability. This section explains the properties of the telescope optical axis, aimpoint, and their drifts during the mission along with the measurements taken to mitigate the potential misalignments caused by these factors to ensure optimal imaging resolution and a successful mission.

4.5.1 Definitions

- **HRMA Optical Axis**: For ideal Wolter Type-I optics, the optical axis is defined by the axis of symmetry of the paraboloid and hyperboloid surfaces. For the actual HRMA, which is not made of perfect surfaces, the optical axis is a straight line passing through the geometric center of the four nested HRMA parabolic and hyperbolic surface mirror pairs and the HRMA focal point.
  - **HRMA Focal Point** (AKA HRMA Focus): Point on the focal plane where the sharpest PSF is located.
  - **HRMA Focal Plane**: An imaginary plane perpendicular to the HRMA optical axis and intersecting it at the focal point. The focal plane is not defined by the detector surface, although placing the detector surface on or near the focal plane obtains the optimal on-axis imaging spatial resolution.
- **ACA Reference Frame**: Frame based on the optical telescope Aspect Camera Assembly (ACA) pointing. The ACA reference frame is used to point the spacecraft. Any drift in the ACA to HRMA alignment will directly affect the actual location on the detector where an on-axis source lands.
- **HRMA Aimpoint**: Point on the focal plane where the image of an on-axis target is located (without detector offset).
  - **Effective Aimpoint**: Location on each detector where the image of an on-axis target actually landed. The effective aimpoint drifts during the mission.
  - **Default Aimpoint**: Fixed chosen locations on detectors to avoid the chip gaps and node boundaries while still maintaining the optimal PSF. Before cycle 18, default aimpoints on ACIS-I and ACIS-S were managed by setting certain pointing offsets annually to compensate for the slow drift of the effective aimpoint. Starting with Cycle 18 observations, a permanent default aimpoint (PDA) is chosen for each detector without using the offset (see Section 4.5.4).

For ideal mirrors, the optical axis intersects the focal plane at the best focus which is both the focal point and the aimpoint. For the actual HRMA, the focal point and the aimpoint are not the same, although they differ only slightly. With respect to a coordinate

system associated with each detector, the aimpoint has been drifting slowly since launch, due to the alignment between the HRMA and ACA.

4.5.2 Aimpoint on Chandra Detectors

For each *Chandra* observation, one of four *Chandra* detectors is aligned with the HRMA optical axis by moving the entire Science Instrument Module (SIM) to that detector's nominal SIM-Z position. Figure 4.25 shows the layout of *Chandra* detectors in the SIM plane with the aimpoints indicated at their nominal positions.

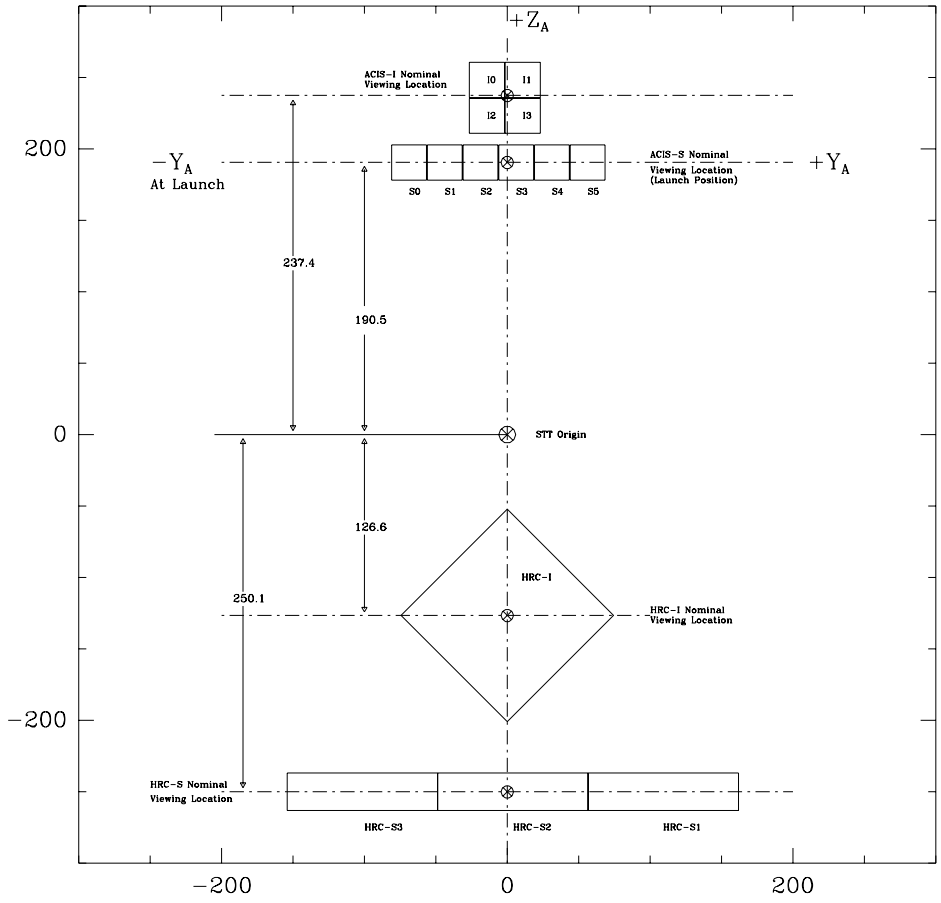

Figure 4.25: Layout of *Chandra* Detectors: The schematic of the SIM plane shows the *Chandra* focal plane instrument to scale (mm). SIM +Y is along the x-axis; SIM +Z is along the y-axis. ⊗ on each detector marks the aimpoint position with nominal SIM-Z position and zero pointing offset.

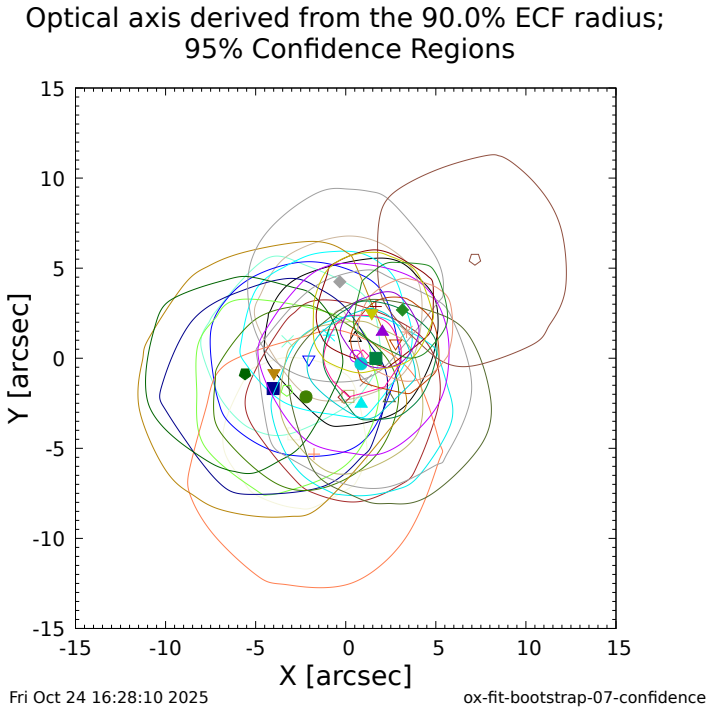

Figure 4.26: Relative locations of the optical axis since launch, and their 95% confidence regions.

4.5.3 Optical Axis

Were the *Chandra* optical system perfect, its components would be symmetrically distributed around a single axis, denoted the *optical axis*. The tightest PSF would be located at its intersection with the focal plane, designated as the on-axis location.

Unfortunately *Chandra*'s optics, while exquisite, are not perfect, and the optical axis is not well defined. It is the on-axis location in the focal plane which is of actual practical import, as it anchors the coordinate systems used by a number of calibration products, such as the vignetting map. For historical reasons, the term "optical axis" will be continued to be used, but what actually is being measured is the location of smallest PSF size.

The design of the optics leads to the size of the PSF varying approximately quadratically with off-axis angle. An estimate of the optical axis can be obtained by sampling the PSF size—$r(x, y)$, where x and y are detector coordinates—across the focal plane and fitting a 2-dimensional quadratic function to it. The optical axis is defined to be at the origin of the fit quadratic.

Figure 4.26 shows the determined optical axis locations and their 95% confidence regions, as derived from HRC-I observations of Ar Lac (an eclipsing RS CVn binary star) at several epochs of the mission through the year 2024. The optical axis has remained

essentially unchanged since the start of the mission.

The optical axis was determined as follows. For each observation, proxies for the PSF size are calculated from the areas—$A(x, y, \mathrm{ECF})$—of convex hulls enclosing given fractions of the object's total background subtracted counts (Enclosed Count Fraction, or ECF). Areas are converted to effective radii through the equation

$$r(x, y, \mathrm{E}CF) = \sqrt{A(x, y, \mathrm{E}CF)/\pi}. \tag{4.1}$$

The quadratic function is then fit to the set of collected radii for a given ECF, $R_{\mathrm{ECF}}(x, y)$. Confidence ranges for the fits are determined using the bootstrap method. Multiple realizations of the observations were constructed by sampling their events with substitution and then analyzed as above. Confidence regions were determined from convex hulls calculated from the resulting distribution of positions.

The results in Figure 4.26 are based on the 90% ECF radii.

4.5.4 Permanent Default Aimpoints (PDA)

The target acquisition is achieved by first pointing *Chandra*'s optical telescope (ACA) to a coordinate based on the telescope roll angle and with a known offset from the target position. Since this offset is fixed, the X-ray telescope (HRMA) should then be pointing to the target. However, because of the thermal changes in the alignment between the HRMA and the ACA, the *Chandra* effective aimpoint has been drifting slowly around the SIM[-Y,-Z] direction by a few arcsec each year since launch. Before Cycle 18, this drift was mitigated by setting default aimpoints on ACIS-I and ACIS-S using yearly updated point offsets. Starting with Cycle 18 observations, a permanent default aimpoint (PDA) is chosen for each detector. These PDAs are close enough to the optical axis so the PSFs are as sharp as on the optical axis while also far from the chip edges or node boundaries.

A predictive thermal model of the alignment between the ACA reference frame and HRMA is used in the planning process to compensate the pointing. However, due to residual uncertainties in this model, error boxes are given around each PDA, indicating the current pointing uncertainty. The dynamical aimpoint adjustment started on 2016-Aug-29. This process has improved the effective aimpoint stability.

The error box sizes are based on the observations taken on all the detectors during the 6 month period prior to 2021-Oct. The peak to peak swings of the short term fluctuations are 16 arcsec in the SIM +Y direction and 12 arcsec in the SIM +Z direction.

Observers also have the option to request the Y-offset, Z-offset and SIM-Z offset to put their targets at different locations on the detector. These offsets will be calculated from the PDA. Observers can use *ObsVis* to visualize the target location on the detectors with their chosen offsets. However, the error boxes will be the same size around their chosen location.

Table 4.3: Optical Axis & PDA Locations

Detector	Chip	SIM-Z [mm]	Optical Axis [pixel] ChipX	Optical Axis [pixel] ChipY	PDA [pixel] ChipX	PDA [pixel] ChipY
ACIS-I	I3	−233.587	961.1 ± 3.1	977.2 ± 3.1	953.74	958.74
ACIS-S	S3	−190.143	234.6 ± 3.1	509.1 ± 3.1	193.74	520.00
HRC-I		126.983	7610.7 ± 15.8	7761.5 ± 3.4	7590	7745
HRC-S	S2	250.466	2163.6 ± 11.4	8911.1 ± 11.4	2195	8915

4.5.5 Optical Axis and PDA positions

The current optical axis and PDA locations are listed in Table 4.3. These numbers are used in *CIAO* and *ObsVis*.

On-axis targets will be imaged near the PDA and inside the error box on each detector. Observers should use this table to check their target location and may request different pointing offset based on their sources to maximize the scientific return. If the observer does not request a specific pointing offset, their target will be put near the PDA.

Telescope dither is used to reduce pileup and protect detectors from damage from bright sources. The dither pattern used for the ACIS detectors has been enlarged in Cycle 24 to improve background analysis for the PCAD system (the HRC dither pattern extent is already sufficient). When planning an observation, it is important to take into account the pointing uncertainty in addition to the dither extent to ensure that targets are located appropriately with respect to chip edges and ACIS node boundaries. Table 4.4 provides dimensions for the dither boxes (which include aimpoint error). Note that the dither patterns (and thus the dither boxes) differ for sub-arrays ≤ 128 pixels in order to ensure that dither in the short direction of the sub-array does not move targets out of the active region of the detector.

Figures 4.27 – 4.32 show the optical axis and PDA with their error boxes on all four detectors.

More details of this topic can be found at `https://cxc.harvard.edu/cal/Hrma/OpticalAxisAndAimpoint.html` and the CXC memorandum 2021: *"Chandra Optical Axis and Aimpoint"* by Ping Zhao.

Table 4.4: Dimensions of Pointing Error & Total Dither Boxes

Detector	Coords	Pointing Error Box		Frame Size	Total Dither Box†	
		[″]	[pixel]		[″]	[pixel]
ACIS-I	Chip X,Y	$12'' \times 16''$	24.4×32.5	Full	$44'' \times 48''$	89.44×97.56
				Sub > 128		
				Sub ≤ 128	$44'' \times 32''$	89.44×65
ACIS-S	Chip X,Y	$16'' \times 12''$	32.5×24.4	Full	$48'' \times 44''$	97.56×89.44
				Sub > 128		
				Sub ≤ 128	$48'' \times 28''$	97.56×56.9
HRC-I	SIM Y,Z	$16'' \times 12''$	121.4×91.1	Full	$56'' \times 52''$	425.0×394.7
HRC-S	Chip X,Y	$12'' \times 16''$	91.1×121.4	Full	$52'' \times 56''$	394.7×425.0

† Includes pointing error

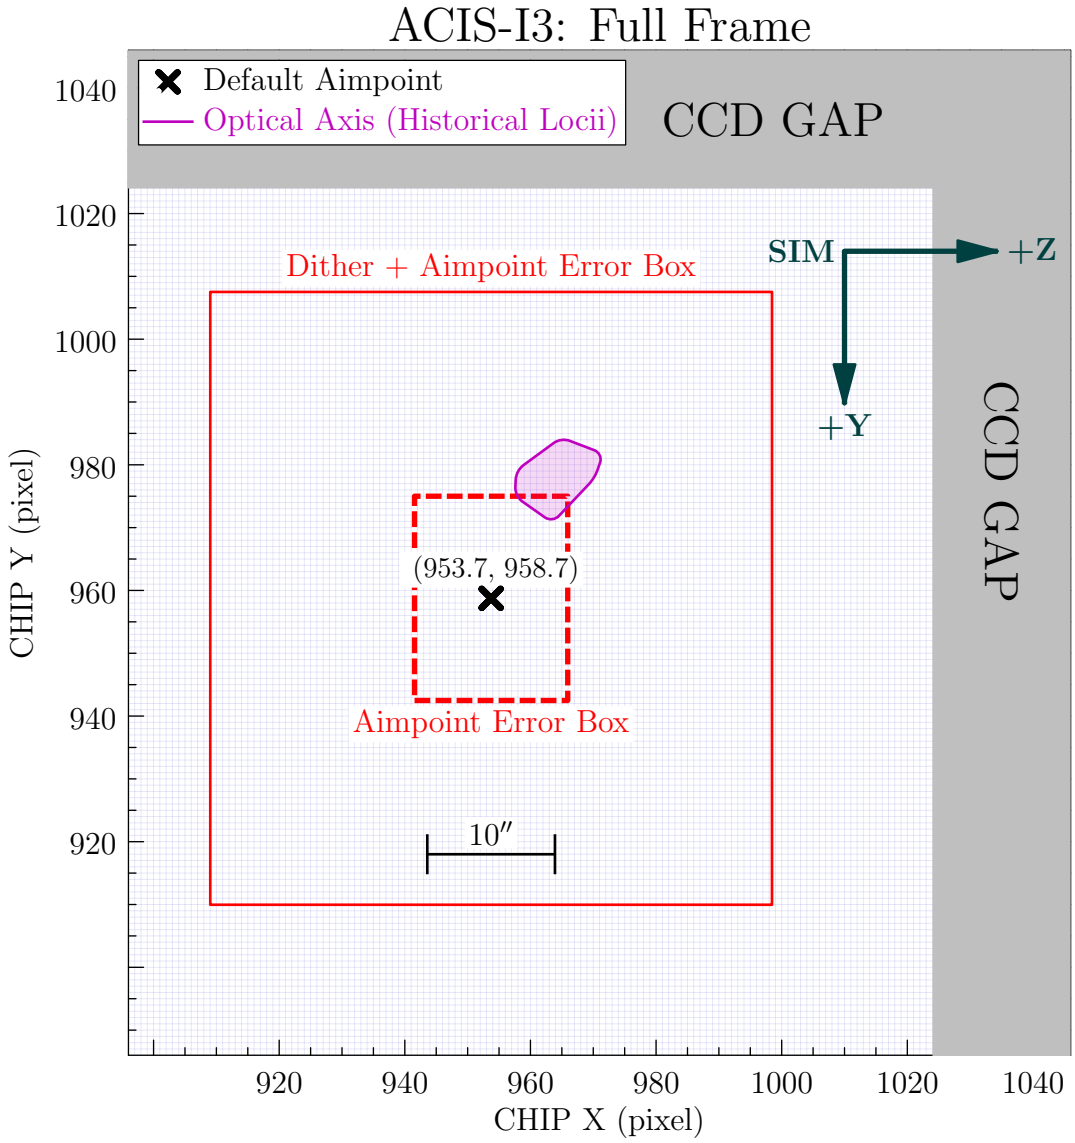

Figure 4.27: ACIS-I3 Full-Frame PDA & Dither Box. The *Chandra* historical optical axis locii are represented by the purple blob; the Permanent Default Aimpoint (PDA) by the black cross. The smaller red dashed-line box centered at the PDA is the error box of the effective aimpoint center—an on-axis target can be anywhere inside this box. The larger red solid-line box contains the dither pattern convolved with the error box.

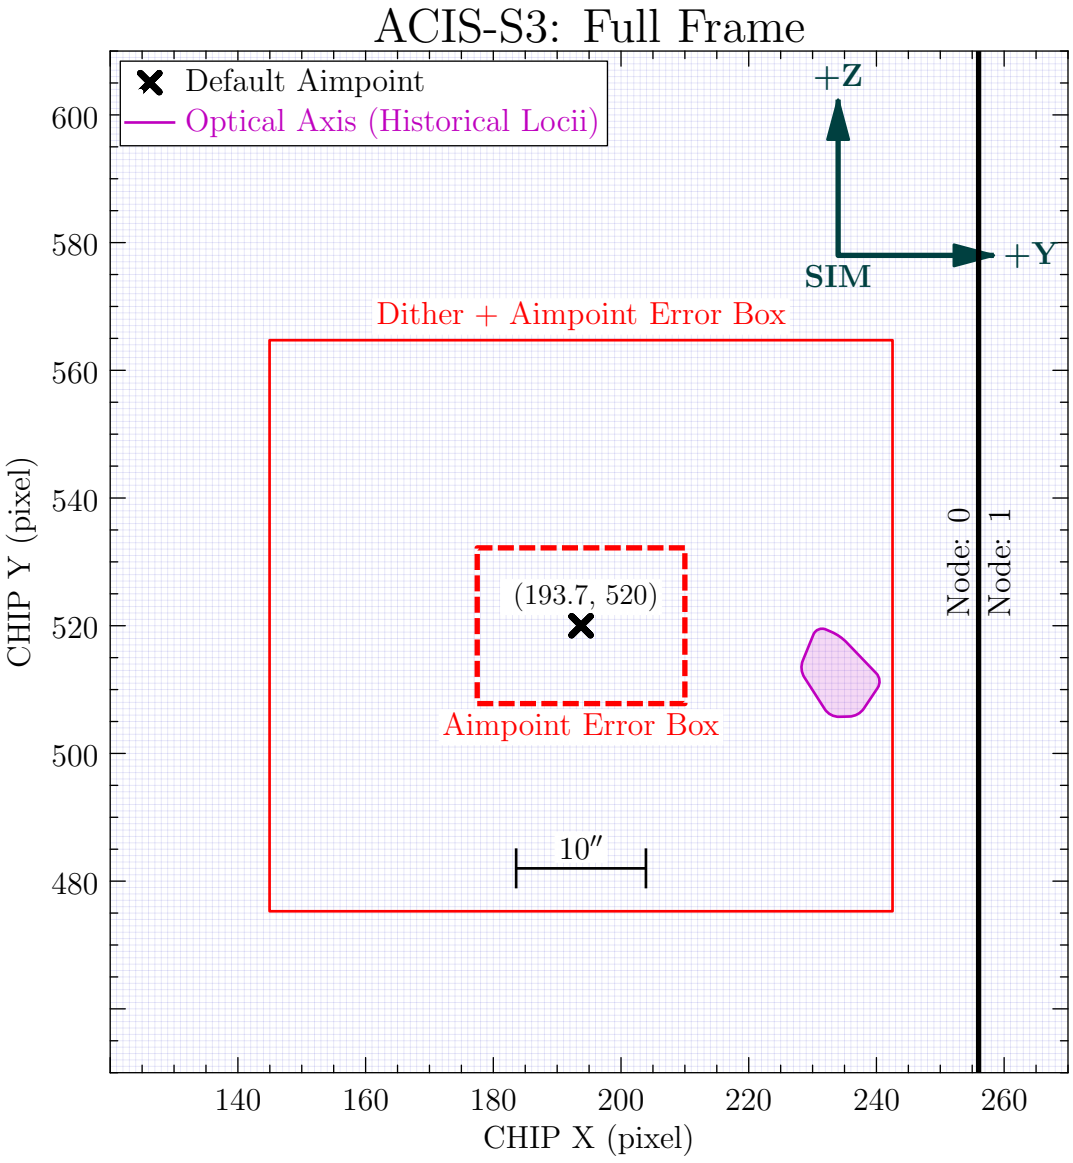

Figure 4.28: ACIS-S3 Full-Frame PDA & Dither Box. The *Chandra* historical optical axis locii are represented by the purple blob; the Permanent Default Aimpoint (PDA) by the black cross. The smaller red dashed-line box centered at the PDA is the error box of the effective aimpoint center—an on-axis target can be anywhere inside this box. The larger red solid-line box contains the dither pattern convolved with the error box. The PDA is shifted ∼12 arcsec from the optical axis to avoid the node boundary 0|1.

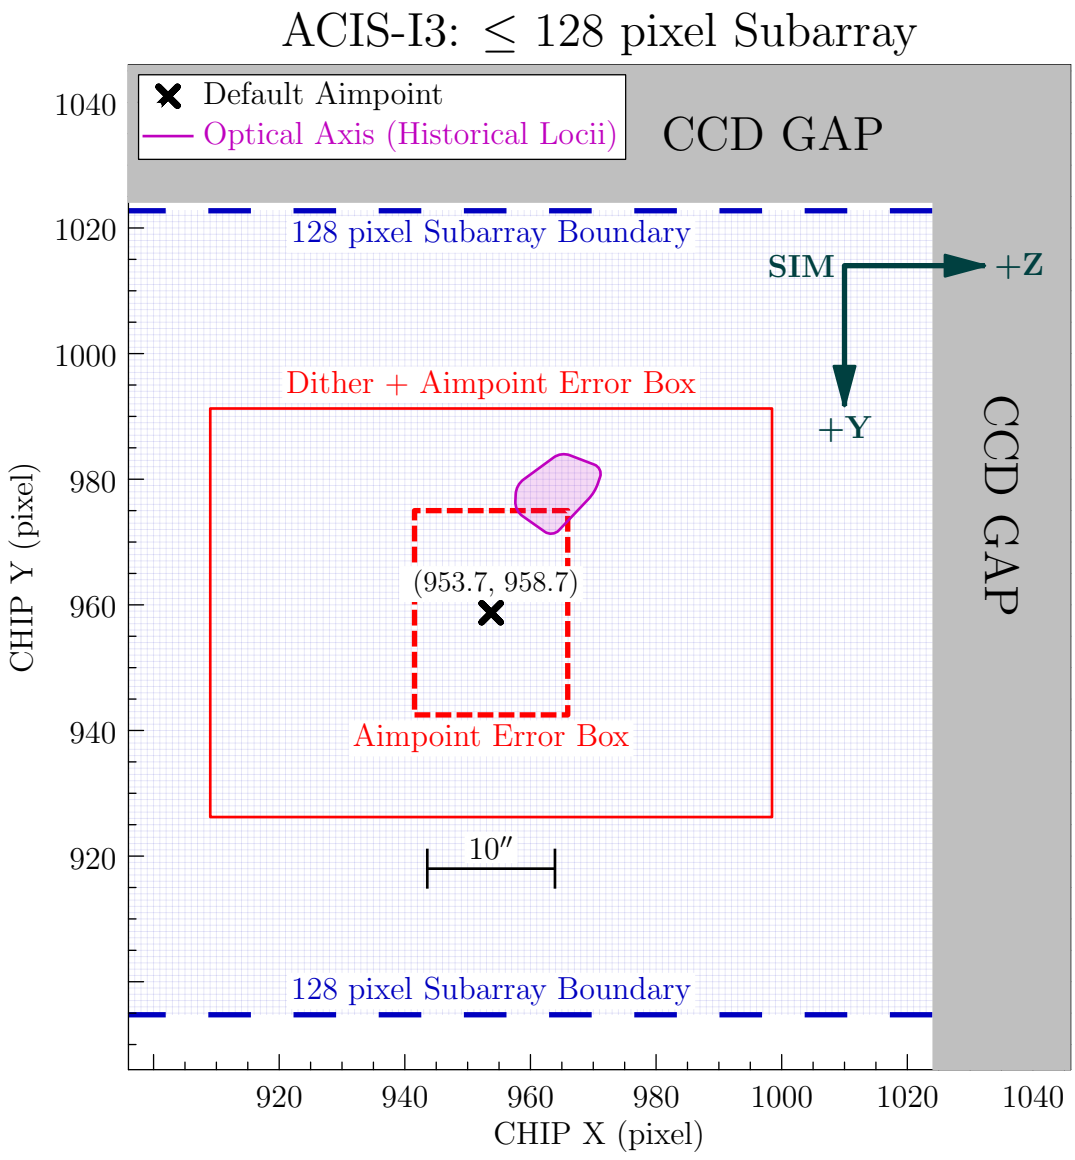

Figure 4.29: ACIS-I3 Dither Box on a subarray ≤ 128 pixels. The *Chandra* historical optical axis locii are represented by the purple blob; the Permanent Default Aimpoint (PDA) by the black cross. The smaller red dashed-line box centered at the PDA is the error box of the effective aimpoint center—an on-axis target can be anywhere inside this box. The larger red solid-line box contains the dither pattern convolved with the error box.

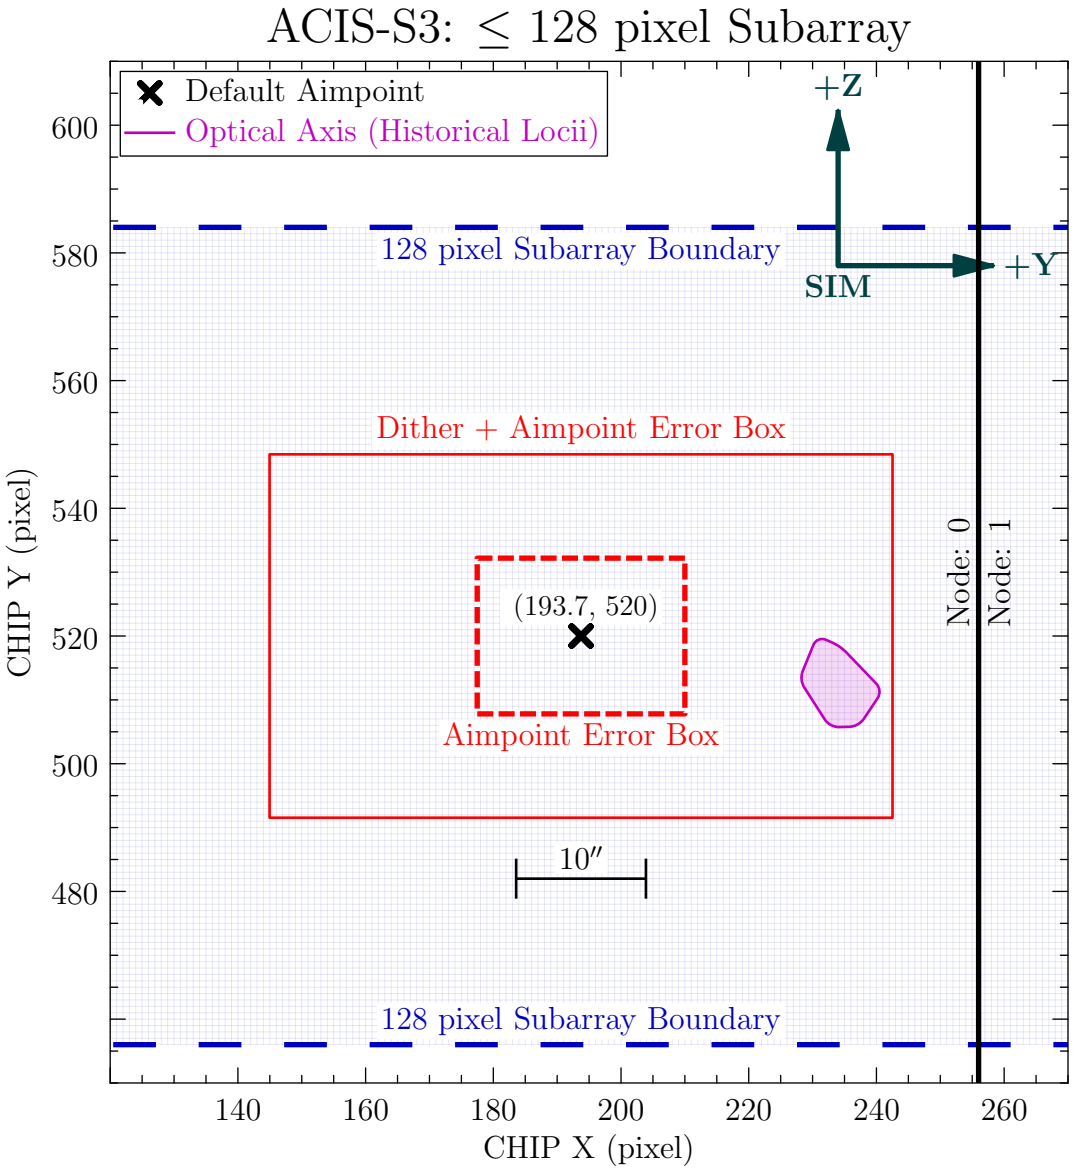

Figure 4.30: ACIS-S3 Dither Box on a subarray ≤ 128 pixels. The *Chandra* historical optical axis locii are represented by the purple blob; the Permanent Default Aimpoint (PDA) by the black cross. The smaller red dashed-line box centered at the PDA is the error box of the effective aimpoint center—an on-axis target can be anywhere inside this box. The larger red solid-line box contains the dither pattern convolved with the error box.

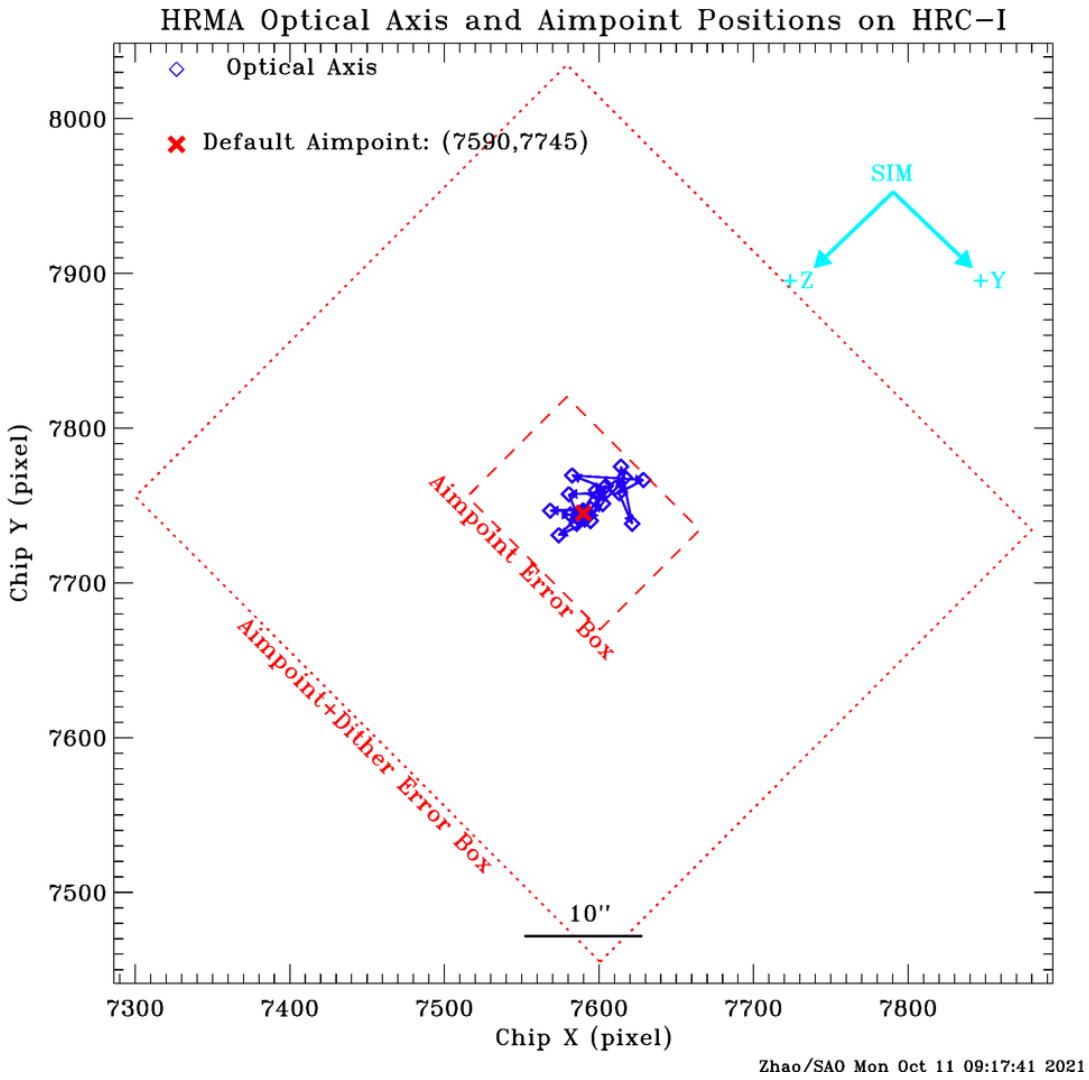

Figure 4.31: *Chandra* optical axis (blue diamonds) positions for each year and permanent default aimpoint (red cross) on HRC-I at [7590,7745]. The small box (red dashed line, 16 arcsec × 12 arcsec, or 121.4 × 91.1 pixels in SIM [Y, Z]) centered at the permanent default aimpoint is the error box of the effective aimpoint center, i.e. an on-axis target can be anywhere inside this box. The large box (red dotted line, 56 arcsec × 52 arcsec, or 425.0 × 394.7 pixels in SIM [Y, Z]) is the error box of effective aimpoint with dither.

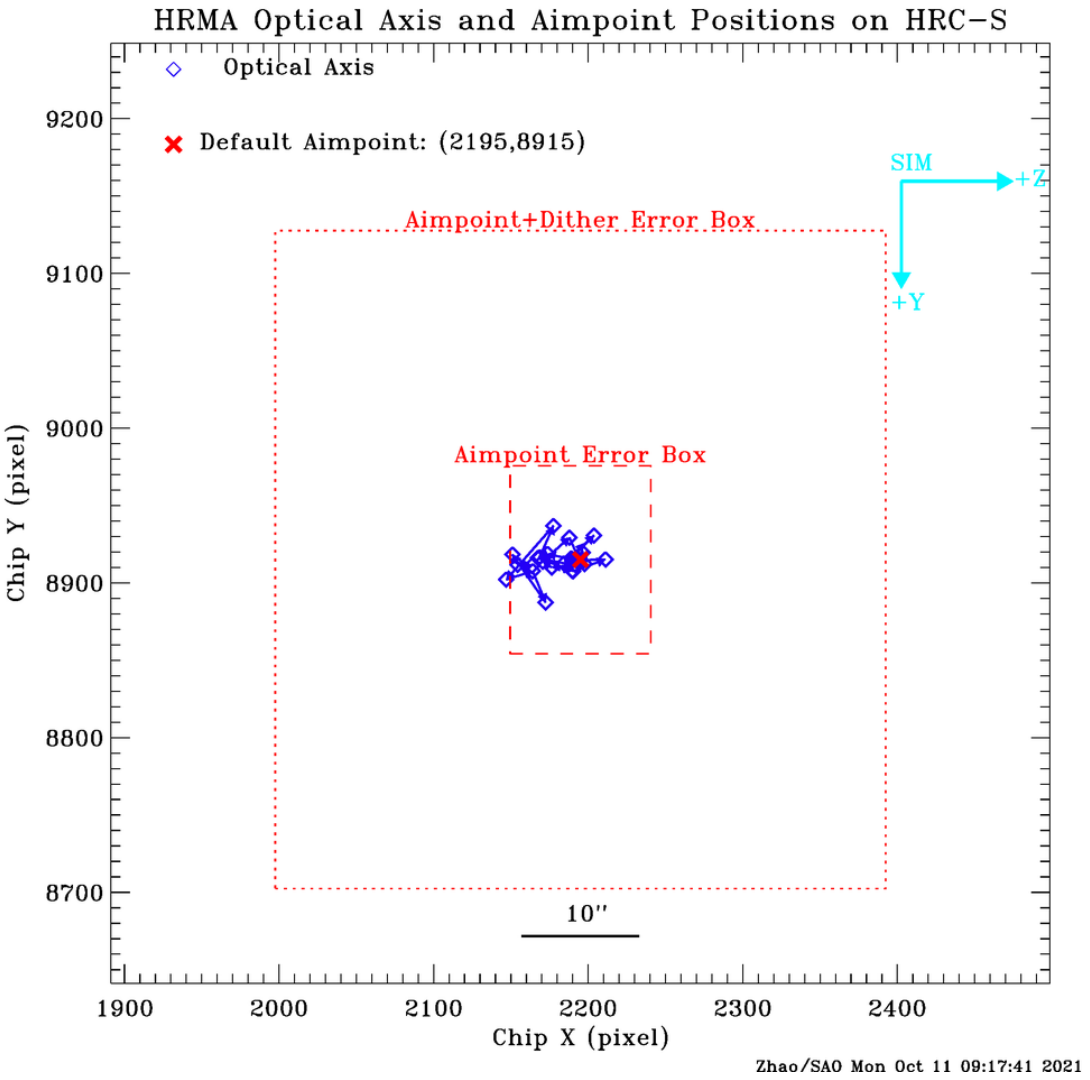

Figure 4.32: *Chandra* optical axis (blue diamonds) positions for each year and permanent default aimpoint (red cross) on HRC-S at [2195,8915]. The small box (red dashed line, 12 arcsec × 16 arcsec, or 91.1 × 121.4 pixels) centered at the permanent default aimpoint is the error box of the effective aimpoint center, i.e. an on-axis target can be anywhere inside this box. The large box (red dotted line, 52 arcsec × 56 arcsec, or 394.7 × 425.0 pixels) is the error box of effective aimpoint with dither.

4.6 References

Aschenbach, B. 1985, Rep. Prog. Phys. 48, 579
Aschenbach, B. 1991, Rev. Mod. Astron. 4, 173
Boese, F.G. 2000, Astron. Ap. Suppl. 507, 141
Edgar, R.J., et al. 1997, SPIE Proceedings, 3113, 124
Elsner, R.F., et al. 1998, SPIE Proceedings, 3444, 177
Gaetz, T.J., et al. 1997, SPIE Proceedings, 3113, 77
Gaetz, T.J., et al. 2000, SPIE Proceedings, 4012, 41
Gaetz, T.J., 2004, https://cxc.cfa.harvard.edu/cal/Hrma/rsrc/Publish/Optics/PSFWings/wing_analysis_rev1b.pdf
Gaetz, T.J. and Jerius, D. 2005,https://cxc.harvard.edu/cal/Hrma/UsersGuide.html
Gaetz, T.J., et al. 2004, SPIE Proceedings, 5165, 411
Giacconi, R., et al. 1979, ApJ, 230, 540
Graessle, D.E., et al. 1998, SPIE Proceedings, 3444, 140
Graessle, D.E., et al. 2004, SPIE Proceedings, 5165, 469
Henke, B.L., et al. 1993, Atomic Data and Nuclear Data Tables, 54, 181
Hughes, J.P., et al. 1993, SPIE Proceedings, 1742, 152
Jerius, D., et al. 2000, SPIE Proceedings, 4012, 17
Jerius, D., et al. 2004, SPIE Proceedings, 5165, 402
Jerius, D., et al. 2004, SPIE Proceedings, 5165, 433
Juda, M., & Karovska, M., 2010 AAS/HEAD Poster "Chandra's Ultimate Angular Resolution: Studies of the HRC-I Point Spread Function", https://hea-www.harvard.edu/~juda/memos/HEAD2010/HEAD2010_poster.html
Kashyap, V., 2010 CXC memorandum, "Analysis of Chandra PSF feature using ACIS data", https://cxc.harvard.edu/cal/Hrc/PSF/acis_psf_2010oct.html
Kolodziejczak, J.J., et al. 1997, SPIE Proceedings, 3113, 65
Marshall, H.L., 2005 "Improving the Relative Accuracy of the HETGS Effective Area", https://space.mit.edu/ASC/calib/heg_meg/meg_heg_report.pdf
Marshall, H.L., 2005a, private communications
O'Dell, S.L. and Weisskopf, M.C. 1998, SPIE Proceedings, 3444, 2
Olds, C.R. and Reese, R.P. 1998, SPIE Proceedings, 3356, 910
Schwartz, D.A., et al. 2000, SPIE Proceedings, 4012, 28
Trümper, J., 1983, Adv. Space Res. 2(4), 241.
Van Speybroeck, L.P., et al. 1997, SPIE Proceedings, 3113, 89
Weisskopf, M.C. and O'Dell, S.L. 1997, SPIE Proceedings, 3113, 2
Weisskopf, M.C., et al. 2000, SPIE Proceedings, 4012, 2
Zhao, P., et al. 1993, SPIE Proceedings, 1742, 26
Zhao, P., et al. 1993, SPIE Proceedings, 1742, 75
Zhao, P., et al. 1994, SPIE Proceedings, 2011, 59
Zhao, P. and Van Speybroeck, L.P. 1995, SPIE Proceedings, 2515, 391
Zhao, P., et al. 1997, SPIE Proceedings, 3113, 106

Zhao, P., et al. 1998, SPIE Proceedings, 3444, 234
Zhao, P. and Van Speybroeck, L.P. 2003, SPIE Proceedings, 4851, 124
Zhao, P., et al. 2004, SPIE Proceedings, 5165, 482
Zhao, P., 2004, Chandra Calibration Workshop, Cambridge, MA *"Chandra Telescope Optical Axis"*
https://cxc.cfa.harvard.edu/ccw/proceedings/04_proc/presentations/zhao
Zhao, P., 2005, Chandra Calibration Workshop, Cambridge, MA *"Chandra Telescope Optical Axis and Aimpoint"*
https://cxc.cfa.harvard.edu/ccw/proceedings/05_proc/presentations/zhao
Zhao, P., 2006, CXC Memorandum *"Chandra Telescope Optical Axis and Aimpoint"*
https://cxc.harvard.edu/cal/Hrma/rsrc/Publish/Optics/OpticalAxisAndAimpoint/opt_axis_memo.pdf
Zhao, P., 2007, Chandra Calibration Workshop, Huntsville AL *"Chandra Telescope Optical Axis and Aimpoint"*
https://cxc.cfa.harvard.edu/ccr/proceedings/07_proc/presentations/zhao
Zhao, P., 2009, Chandra Calibration Review, Boston, MA *"The Quality and Stability of Chandra Telescope Pointing and Spatial Resolution"*
https://cxc.harvard.edu/ccr/proceedings/2009/presentations/zhao
Zhao, P., 2011, CXC memorandum *"Chandra Aimpoint Drift and Default Offsets"*
https://cxc.harvard.edu/cal/Hrma/rsrc/Publish/Optics/OpticalAxisAndAimpoint/aimpoint_memo_2011.pdf
Zhao, P., 2012, CXC memorandum *"Chandra Optical Axis, Aimpoint and Their Drifts"*
https://cxc.harvard.edu/cal/Hrma/rsrc/Publish/Optics/OpticalAxisAndAimpoint/oxap_memo_2012.pdf
Zhao, P., 2013, CXC memorandum *"Chandra Optical Axis, Aimpoint and Their Drifts"*
https://cxc.harvard.edu/cal/Hrma/rsrc/Publish/Optics/OpticalAxisAndAimpoint/oxap_memo_2013.pdf
Zhao, P., 2014, CXC memorandum *"Chandra Optical Axis, Aimpoint and Their Drifts"*
https://cxc.harvard.edu/cal/Hrma/rsrc/Publish/Optics/OpticalAxisAndAimpoint/oxap_memo_2014.pdf
Zhao, P., 2015, CXC memorandum *"Chandra Optical Axis and Aimpoint"*
https://cxc.harvard.edu/cal/Hrma/rsrc/Publish/Optics/OpticalAxisAndAimpoint/oxap_memo_2015.pdf
Zhao, P., 2016, CXC memorandum *"Chandra Optical Axis and Aimpoint"*
https://cxc.harvard.edu/cal/Hrma/rsrc/Publish/Optics/OpticalAxisAndAimpoint/oxap_memo_2016.pdf
Zhao, P., 2017, CXC memorandum *"Chandra Optical Axis and Aimpoint"*
https://cxc.harvard.edu/cal/Hrma/rsrc/Publish/Optics/OpticalAxisAndAimpoint/oxap_memo_2017.pdf
Zhao, P., 2018, CXC memorandum *"Chandra Optical Axis and Aimpoint"*
https://cxc.harvard.edu/cal/Hrma/rsrc/Publish/Optics/OpticalAxisAndAimpoint/oxap_memo_2018.pdf

Zhao, P., 2019, CXC memorandum *"Chandra Optical Axis and Aimpoint"* https://cxc.harvard.edu/cal/Hrma/rsrc/Publish/Optics/OpticalAxisAndAimpoint/oxap_memo_2019.pdf
Zhao, P., 2020, CXC memorandum *"Chandra Optical Axis and Aimpoint"* https://cxc.harvard.edu/cal/Hrma/rsrc/Publish/Optics/OpticalAxisAndAimpoint/oxap_memo_2020.pdf
Zhao, P., 2021, CXC memorandum *"Chandra Optical Axis and Aimpoint"* https://cxc.harvard.edu/cal/Hrma/rsrc/Publish/Optics/OpticalAxisAndAimpoint/oxap_memo_2021.pdf

Portable Document Format (PDF) copies of various aspects of the HRMA calibration can be obtained from the *CXC* Optics Calibration Group at:
https://cxc.harvard.edu/cal/Hrma/XRCFReport.html

A detailed guide to using and understanding the HRMA is available at:
https://cxc.harvard.edu/cal/Hrma/UsersGuide.html

Further information can be obtained from the MSFC *Chandra* calibration page at:
https://wwwastro.msfc.nasa.gov/xray/xraycal/

The current positions of Chandra Optical Axis and Aimpoint, and all the references can be found at: https://cxc.harvard.edu/cal/Hrma/OpticalAxisAndAimpoint.html.

Chapter 5

Pointing Control and Aspect Determination System

5.1 Introduction

The system of sensors and control hardware that is used to point the observatory, maintain the stability, and provide data for determining where the observatory has been pointing is called the Pointing Control and Aspect Determination (PCAD) system. As *Chandra* detectors are essentially single-photon counters, an accurate post-facto history of the spacecraft pointing direction is sufficient to reconstruct an X-ray image.

This chapter briefly discusses the hardware that comprises the PCAD system, how it is used, and its flight performance. Further information can be found on the Aspect Information web page within the main *CXC* Science web site (https://cxc.harvard.edu/cal/ASPECT/).

5.2 Physical Configuration

The main components of the PCAD system are:

Aspect Camera Assembly (ACA) – 11.2 cm optical telescope, stray light shade, two CCD detectors (primary and redundant), and two sets of electronics

Inertial Reference Units (IRU) – Two IRUs, each containing two 2-axis gyroscopes

Fiducial Light Assemblies (FLA) – LEDs mounted near each science instrument (SI) detector which are imaged in the ACA via the FTS

Fiducial Transfer System (FTS) – The FTS directs light from the fiducial lights to the ACA via the retroreflector collimator (RRC)—mounted at the HRMA center—and a periscope

Coarse Sun Sensor (CSS) – Sun position sensor, all-sky coverage

Fine Sun Sensor (FSS) – Sun position sensor, 50 deg field of view (FOV) and 0.02 deg accuracy

Earth Sensor Assembly (ESA) – Conical scanning sensor, used during the orbital insertion phase of the mission

Reaction Wheel Assembly (RWA) – 6 momentum wheels that change spacecraft attitude

Momentum Unloading Propulsion System (MUPS) – Liquid fuel thrusters that allow RWA momentum unloading

Reaction Control System (RCS) – Thrusters that change spacecraft attitude

Since data from the CSS, FSS, and ESA are not normally used in the processing of science observations, these are not discussed. However, in the unlikely event of a complete failure of the ACA, use of the CSS and FSS data would be attempted.

5.2.1 ACA

The aspect camera assembly (Figure 5.1) includes a sunshade (∼2.5 m long, ∼40 cm in diameter), an 11.2 cm, F/9 Ritchey-Chretien optical telescope, and light-sensitive CCD detectors. This assembly and its related components are mounted on the side of the HRMA. The camera's field of view is $1.4^\circ \times 1.4^\circ$, and the sunshade is designed to protect the instrument from the light from the Sun, Earth, and the Moon, with protection angles of 20°, 6°, and 6°, respectively.

The aspect camera focal-plane detector is a 1024×1024 Tektronix CCD chip operating between -2.1 °C and -14.0 °C, with 24×24 micron (5 arcsec × 5 arcsec) pixels, covering the spectral band between 4000 and 9000 Å. The optics of the camera are defocused (point source FWHM = 9 arcsec) to spread the star images over several CCD pixels, both to increase accuracy of the centering algorithm and to reduce variation in the point response function over the field of view. There is a spare identical CCD chip, which can be illuminated by inserting a rotatable mirror.

The ACA electronics track a small pixel region (either 4 pixels × 4 pixels, 6 pixels × 6 pixels, or 8 pixels × 8 pixels) around each fiducial light and star image. There are a total of eight such image slots available for tracking. Typically five guide stars and three fiducial lights (Section 5.2.2) are tracked. With the application of the dynamic background patch, a software upgrade, in May 2023, the ACA accumulates background information from the edge pixels of the pixel regions as the spacecraft dithers. This per-pixel background is then subtracted on-board, and image centroids are calculated by a weighted-mean algorithm. The image centroids and magnitudes are used on-board by the PCAD and are also telemetered to the ground along with the raw pixel data.

The spectral response of the CCD detector (Figure 5.2) is such that faint cool stars (e.g. type N0) with visual magnitudes much fainter than selected guide stars (i.e., 10.5

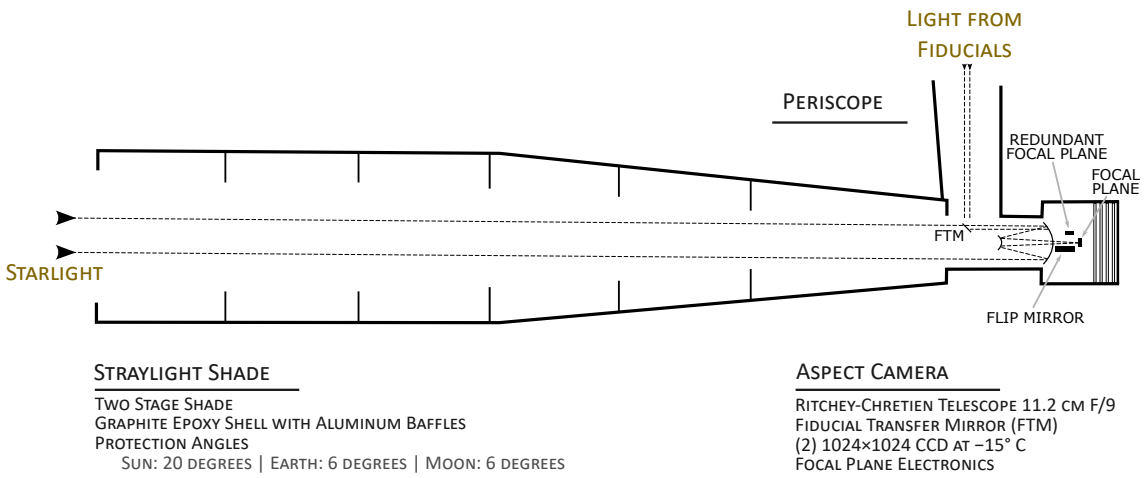

Figure 5.1: Aspect camera assembly

mag) can produce large numbers of counts. These so-called "spoiler stars" are effectively avoided in the mission planning stage.

5.2.2 Fiducial Lights and Fiducial Transfer System

Surrounding each of the SI detectors is a set of light emitting diodes, or "fiducial lights," which serve to register the SI focal plane laterally with respect to the ACA boresight. Each fiducial light produces a collimated beam at 635 nm that is imaged onto the ACA CCD via the retroreflector collimator, the periscope, and the fiducial transfer mirror (Figure 5.3).

5.2.3 IRU

Two Inertial Reference Units (IRU) are located in the front of the observatory on the side of the HRMA. Each IRU contains two gyros, each of which measures an angular rate about 2 gyro axes. This gives a total of eight gyro channels. Data from four of the eight channels can be read out at one time. The gyros are arranged within the IRUs and the IRUs are oriented such that all 8 axes are in different directions and no three axes lie in the same plane. The gyros' output pulses represent incremental rotation angles. In high-rate mode, each pulse nominally represents 0.75 arcsec, while in low-rate mode (used during all normal spacecraft operations) each pulse represents nominally 0.02 arcsec.

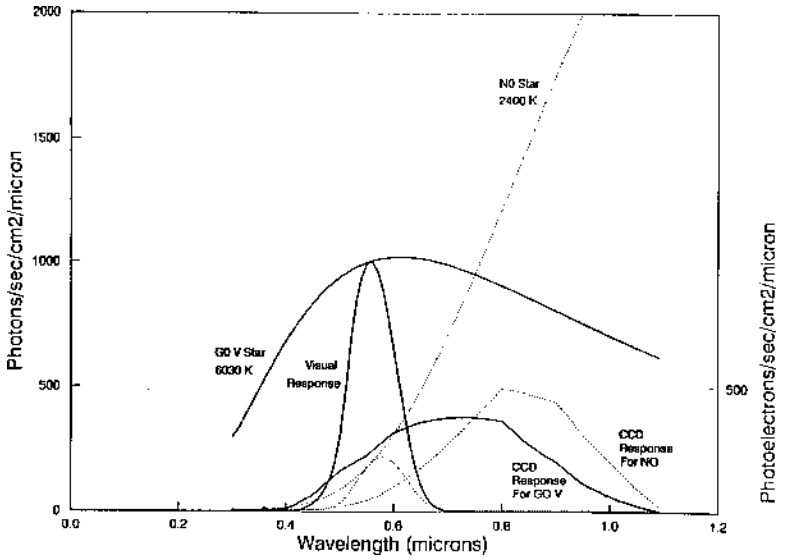

Figure 5.2: Spectral response of the ACA CCD. The same signal-to-noise is achieved for a V=11.7 magnitude N0 star as for a V=10 magnitude G0V star. Also shown are the spectra and the standard visual response for the two stars.

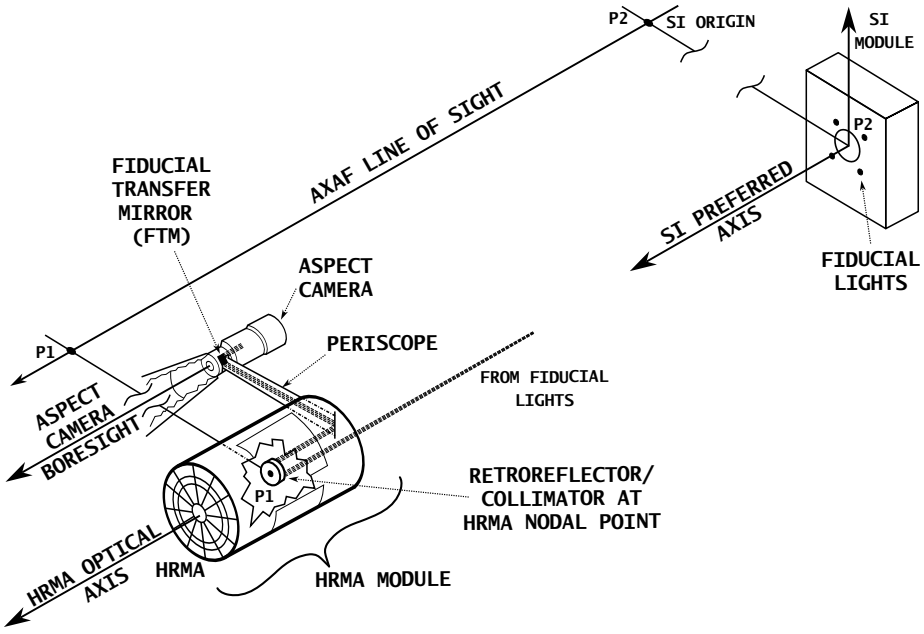

Figure 5.3: Fiducial Transfer System. Fiducial lights are imaged onto the ACA CCD via the retroreflector collimator, the periscope, and the fiducial transfer mirror.

5.2.4 Momentum Control – RWA and MUPS

Control of the spacecraft momentum is required both for maneuvers and to maintain stable attitude during science observations. Momentum control is primarily accomplished using six Teldix RDR-68 reaction wheel units mounted in a pyramidal configuration. During observing—with the spacecraft attitude constant apart from dither—external torques on the spacecraft (including gravity gradient and magnetic torques) will cause a buildup of momentum in the RWA. Momentum is then unloaded by firing the MUPS and simultaneously spinning down the reaction wheels.

5.3 Operating Principles

The *Chandra* aspect system serves two primary purposes: first, for on-board spacecraft pointing control and aspect determination and second, for post-facto ground aspect determination, used in X-ray image reconstruction and celestial location.

The PCAD system has nine operational modes (six normal and three safe modes) which use different combinations of sensor inputs and control mechanisms to control the spacecraft and ensure its safety. These modes are described in Section 5.7.1. In the normal science pointing mode, the PCAD system uses sensor data from the ACA and IRUs and control torques from the RWAs to keep the target attitude within ∼30 arcsec of the telescope boresight. This is done using a Kalman filter which optimally combines ACA star centroids (typically five) and angular displacement data from two 2-axis gyroscopes. On short time scales (∼sec) the spacecraft motion solution is dominated by the gyroscope data, while, on longer timescales, it is the star centroids that determine the solution.

Post-facto aspect determination is done on the ground and uses more sophisticated processing and better calibration data to produce a more accurate aspect solution. The suite of *CXC* tools to perform this processing is called the aspect pipeline. The key improvements over PCAD aspect come from better image centroiding and using Kalman smoothing (which uses all available data over the observation period, as opposed to historical data). In addition, the aspect pipeline folds in the position of the focal-plane instrument as determined by the fiducial light data.

5.4 Performance

This section provides information about the aspect system performance, with an emphasis on parameters that impact science and/or observation planning. Table 5.1 provides a summary of results for four key aspect requirements (Sections 5.4.1 - 5.4.4) that originate in the Observatory Level-2 Project Requirements Document. In each case, the actual performance exceeds the requirements.

Table 5.1: Aspect System Requirements and Performance

Description	Requirement	Actual
Celestial location	1.0 arcsec (RMS radius)	0.63 arcsec
Image reconstruction	0.5 arcsec (RMS diameter)	0.32 arcsec
Absolute celestial pointing	30.0 arcsec (99.0%, radial)	8 arcsec
PCAD 10 sec pointing stability	0.12 arcsec (95% RMS)	0.09 arcsec (pitch) 0.06 arcsec (yaw)

5.4.1 Celestial Location Accuracy

Celestial location accuracy measures the absolute accuracy of *Chandra* X-ray source locations. Based on observations of 315 point sources detected within 3 arcmin of the boresight and having accurately known coordinates, the 90% source location error circle has a radius of less than 0.87 arcsec overall, and less than 0.95 arcsec for each SI (Figure 5.4). Approximately 6% of sources are outside a 1 arcsec radius. The difference in astrometric accuracy for different SIs is a function of two factors: number of available data points for boresight calibration and accuracy of the fiducial light SIM-Z placement. These values apply for sources within 3 arcmin of the aimpoint and with the SIM-Z at the nominal detector value. Observations on ACIS or HRC-S at large off-nominal SIM-Z can suffer additional residual aspect offsets of up to 0.5 arcsec and, for HRC-I, this can be up to 3 arcsec. The plotted level of accuracy applies to observations after 2020-Oct-2, which were processed for this analysis with the most recent calibration data. Achieving this level of accuracy for archive analysis may require reprocessing with CALDB version 4.10.1 or later.

For further technical details about celestial location accuracy and performance, see https://cxc.cfa.harvard.edu/cal/ASPECT/celmon/.

5.4.2 Image Reconstruction

Image reconstruction performance measures the effective blurring of the X-ray PSF due to aspect reconstruction. A direct measure of this parameter can be made by determining the time-dependent jitter in the centroid coordinates of a fixed celestial source. Any error in the aspect solution will be manifested as an apparent wobble in the source location. Unfortunately this method has limitations. ACIS data are count-rate limited and only produce an upper limit: aspect reconstruction effectively convolves the HRMA PSF with a Gaussian having FWHM of less than 0.25 arcsec. HRC observations can produce acceptably high count-rates, but the HRC photon positions (at the chip level) have systematic errors due to uncertainties in the HRC de-gap calibration (See "Position modeling, de-gap corrections, and event screening" in Section 7.16). These errors exactly mimic the expected dither-dependent signature of aspect reconstruction errors, so no such analysis with HRC data has been done. An indirect method of estimating aspect reconstruction blurring is to use the aspect solution to de-dither the ACA star images and measure the residual jitter.

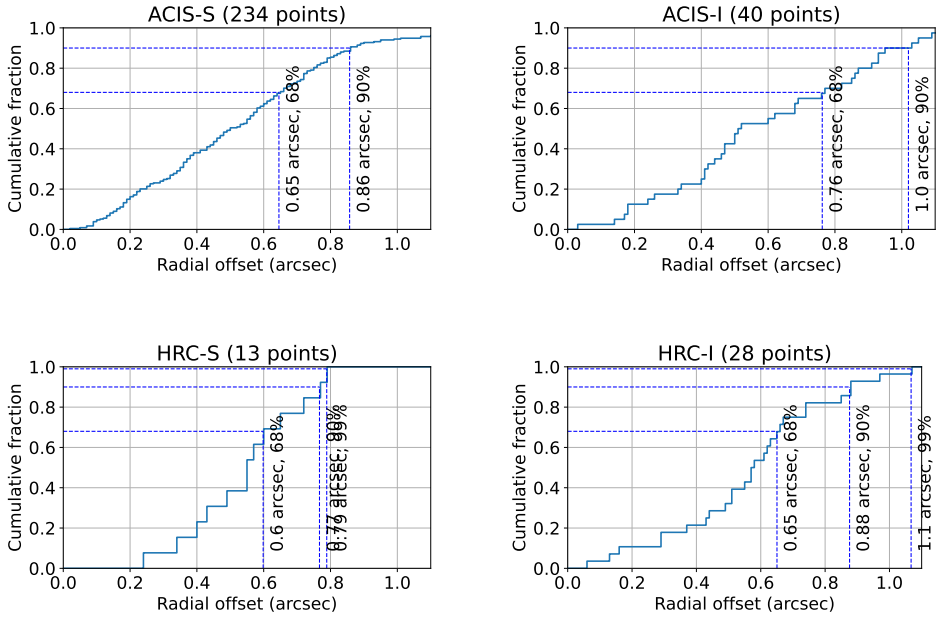

Figure 5.4: Cumulative histogram of celestial accuracy for *Chandra* X-ray source locations for each SI using observations since 2020-Oct. Radial offset is the distance in arcsec between the optical coordinate, typically from the Tycho-2 catalog, and the *Chandra* position.

Performing this for 350 observations shows that 99% of the time the effective blurring is less than 0.20 arcsec (FWHM).

Another component of the image reconstruction error is error in the registration of the SI and the ACA, specifically registration error due to drift in the apparent fiducial light positions used to establish this registration. Drift is caused by changes in the alignment of the periscope (see Figure 5.3). Registration error from this drift is reduced in ground aspect processing by a periscope drift correction applied in the aspect pipeline, but residual error of up to 0.5 arcsec may persist. Users with long (>50ks) observations of a bright, on-axis source, and a science goal that would benefit from sub-arcsec image reconstruction should see https://cxc.harvard.edu/ciao/ahelp/correct_periscope_drift.html for a correction tool.

5.4.3 Absolute Celestial Pointing and Aimpoint Stability

Absolute celestial pointing refers to the accuracy with which an X-ray source can be positioned at a specified location on the detector. Consideration of this accuracy is required in observation planning and is especially relevant for planning of observations sensitive to position on the detector, such as ACIS subarray or windowed observations. It is worth noting that the absolute celestial pointing accuracy discussed here does not have an impact on the accuracy of image reconstruction and celestial location in the processed X-ray data.

The absolute celestial pointing accuracy is determined by the precision to which the pointing system works as well as by the aimpoint stability and drift as discussed in the context of observation planning in Section 4.5. That section provides detailed information on the updated approach with default detector aimpoints and Table 4.4 and Figures 4.27–4.32 provide corresponding uncertainty boxes showing the precision to which a source can be placed on a given detector. These figures include detector-specific padding for dither and address concerns about placement in CCD pixel space.

The absolute celestial pointing performance is dependent on the stability of the aspect camera optical axis because the spacecraft pointing and ultimately the aspect reconstruction is referenced to the ACA frame. In this performance context, the absolute celestial pointing accuracy is defined as the 99% radial offset of actual aimpoints from the planned aimpoint location. This quantifies the ability to position a target near the intended aimpoint.

As of 2025-Oct the absolute celestial pointing accuracy is approximately ±8 arcsec. This is comprised of two observable components, mean absolute pointing error (±5 arcsec, per-axis) and intra-observation drift (±3 arcsec, per-axis). These error component estimates are based on review of recent and historical data. Plots of the two components in recent data may be seen in Figures 5.5 and 5.6 respectively. Both the mean absolute pointing error and the intra-observation drift are correlated with thermal conditions on the spacecraft as the aspect camera optical axis alignment (relative to the HRMA alignment) is sensitive to the ACA housing temperature.

For further technical details about aimpoint monitoring and performance see https://cxc.cfa.harvard.edu/mta/ASPECT/aimpoint_mon/.

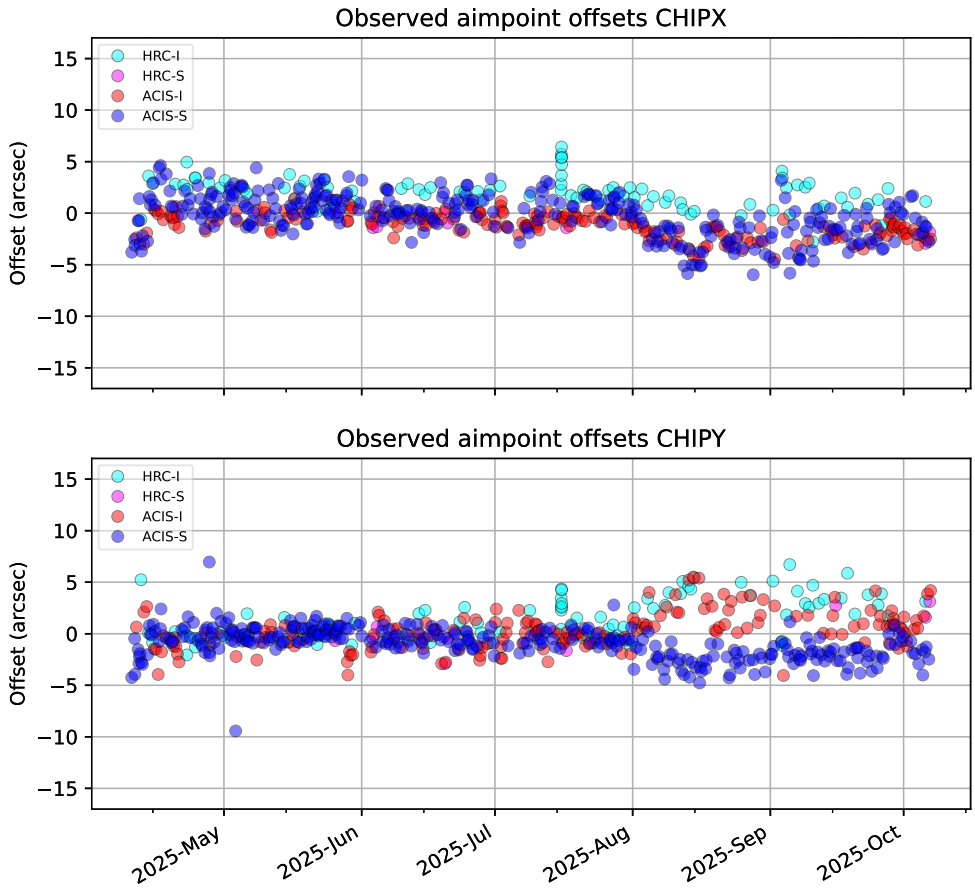

Figure 5.5: Difference between observed and planned aimpoint in CHIPX and CHIPY directions for recent observations with planned aimpoints within 100 arcsec of the nominal aimpoint.

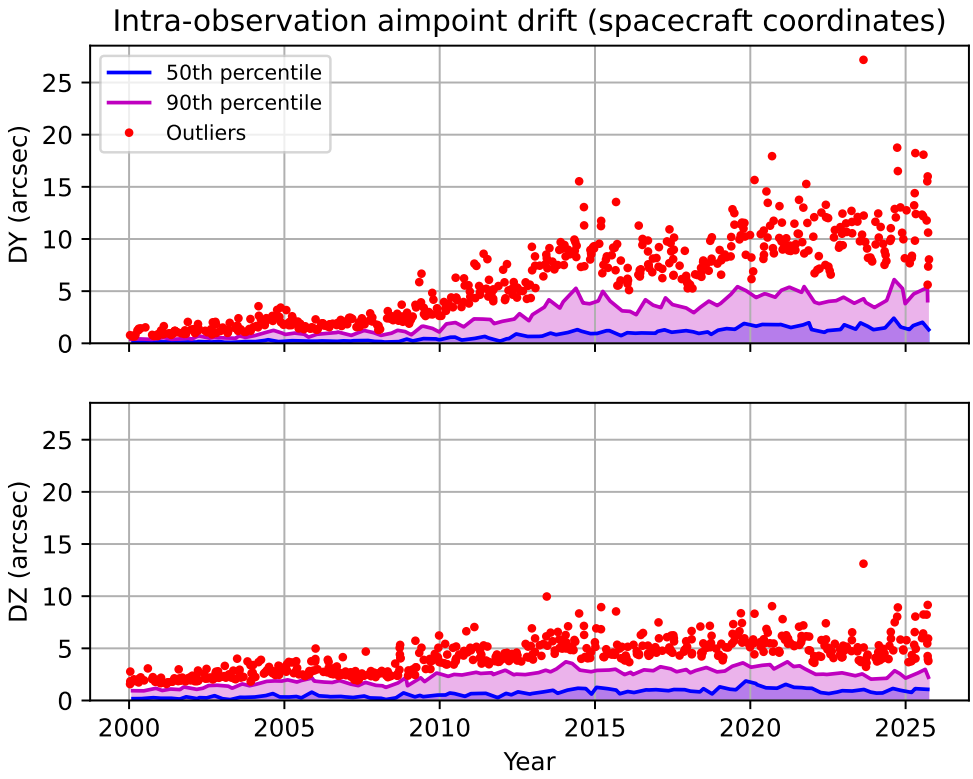

Figure 5.6: Peak to peak aimpoint drift within individual science observations. This figure shows that total intra-observation drift (DY+DZ) up to 6 arcsec during observations is possible and, in rare cases, drift over 15 arcsec can occur.

5.4.4 PCAD 10-Second Pointing Stability

The PCAD 10-second pointing stability performance is measured by calculating the root mean squared (RMS) attitude control error (1-axis) within successive 10 second intervals. The attitude control error is simply the difference between the ideal (commanded) dither pattern and the actual measured attitude. Flight data show that 95% of the RMS error measurements are less than 0.09 arcsec (pitch) and 0.06 arcsec (yaw). Systematic offsets are not included in this term.

5.4.5 Relative Astrometric Accuracy

Relative astrometric accuracy refers to the residual astrometric offsets assuming that the X-ray coordinates have been registered using well-characterized counterparts of several X-ray sources in the field. The most comprehensive dataset for measuring relative astrometry is based on the 900 ksec ACIS-I observation of the Orion Nebula. The members of the COUP (*Chandra* Orion Ultradeep Project) provided data for over 1300 X-ray sources listing the offset from a 2 Micron All-Sky Survey (2MASS) counterpart and the off-axis angle. (Full details are available in Getman et al. 2005. The 1152 sources with more than 50 counts were used for the analysis.) The left plot of Figure 5.7 shows a scatter plot of offset (arcsec) versus off-axis angle (arcmin). The right side of the figure shows cumulative histograms of the fraction of sources with relative offset below the specified value. This is broken into bins of off-axis angle as listed in the plot. In the "on-axis" (0 – 2 arcmin) bin, 90% of sources have offsets less than 0.22 arcsec. After accounting for the ∼0.08 arcsec RMS uncertainty in 2MASS coordinates, this implies the intrinsic 90% limit is 0.15 arcsec. See https://www.ipac.caltech.edu/2mass/releases/allsky/doc/.

5.4.6 On-Board Acquisition and Tracking

As described in Section 5.8, in normal operations the ACA is used to acquire and track stars and fiducial lights. Occasional failures in acquisition and difficulties in tracking are expected due to uncertainties in star position and magnitude, the presence of spoiler stars, CCD dark current noise (see Section 5.6.3), and other factors.

Table 5.2 summarizes success statistics for star acquisition and tracking during the last year. An acquisition star is "successfully" acquired if it is tracked within the acquisition interval and satisfies position and magnitude requirements. A guide star is "successfully" tracked if it spends less than 5% of the observation in the loss-of-track state. In general, fainter stars are less likely to be acquired or successfully tracked than brighter stars due to CCD dark current noise. For a nominal observation, separate sets of stars are used in the acquisition and guide catalogs.

In May 2023, the ACA on-board software was updated to improve background subtraction and mitigate the impact of increasing CCD dark current noise. This update has resulted in improved guide star tracking performance. The improved background subtrac-

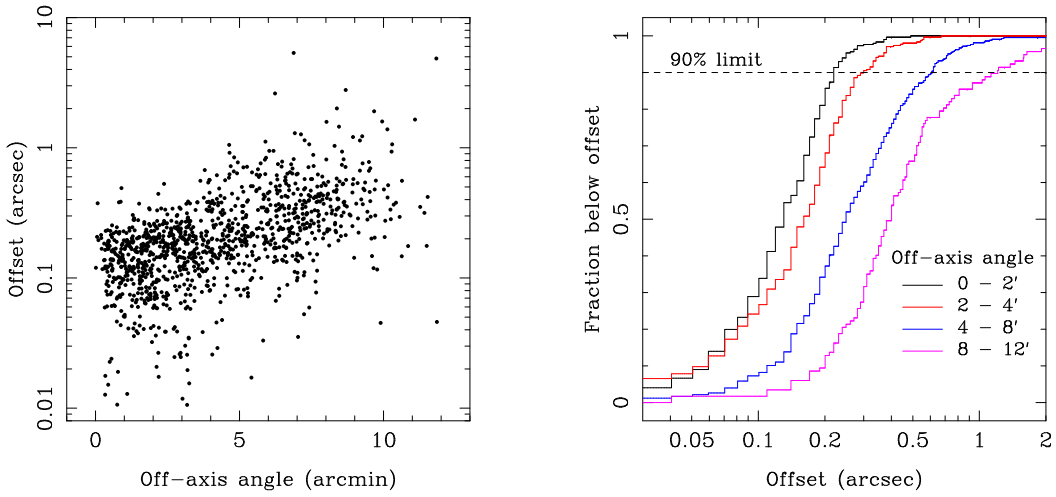

Figure 5.7: Left: scatter plot of offset (arcsec) versus off-axis angle (arcmin) for sources in the ultra-deep ACIS-I Orion observation. Right: cumulative histograms of the fraction of sources with relative offset below the specified value.

Table 5.2: Star Acquisition and Tracking Success from 2024-Oct to 2025-Oct

Catalog Star Magnitude	**Acquisition Success**	**Tracking Success**
All stars	92%	99%
10.0 – 10.3 mag stars	61%	98%
10.3 – 10.9 mag stars	55%	—*

*Insufficient data

tion also allows stars fainter than 10.3 mag to be used as guide stars, but very few have been used.

5.5 Heritage

The *Chandra* aspect camera design is based on the Ball CT-601 star tracker, which was also used for the *Rossi* X-ray Timing Explorer. The *Chandra* IRUs are nearly identical to the space-qualified Kearfott IRU (SKIRU) V, some 70 of which have been built by the manufacturer, Kearfott. These IRUs are similar to those used on the *Compton* Gamma-Ray Observatory (*CGRO*).

5.6 Calibration

5.6.1 Pre-launch Calibration

IRU component testing at Kearfott provided calibration data necessary for accurate maneuvers and for deriving the aspect solution. The key parameters are the scale factor (arcsec/gyro pulse) and the drift rate stability parameters. The stability parameters specify how quickly the gyro readout random-walks away from the true angular displacement. These terms limit the aspect solution accuracy during gyro hold observations (described in further detail in Section 5.8.2).

ACA component testing at Ball provided calibration data necessary for on-orbit pointing control and for post-facto ground processing. On-orbit, the ACA uses CCD gain factors, the plate scale factor, and temperature dependent field distortion coefficients to provide the control system with star positions and brightnesses. In ground processing, the *CXC* aspect pipeline makes use of those calibration data as well as CCD read noise, flat-field maps, dark current maps, and the camera PSF to accurately determine star positions.

5.6.2 Orbital Activation and Checkout Calibration

Orbital activation and checkout (OAC) of the PCAD occurred during the first approximately 30 days of the *Chandra* mission. During the first phase of OAC, before the HRMA sunshade door was opened, it was possible to use the ACA to observe the fiducial lights; this is referred to as period 1. After the sunshade door was opened, in period 2, it was possible to fully check the aspect camera using star light.

Chandra activation produced the following aspect system calibration data:

- Bias, alignment, and scale factor of the CSS and FSS (period 1)
- Coarse gyro bias (period 1)
- ACA CCD dark current map (period 1)
- Fiducial light intensity, image, and centroid at nominal voltage (periods 1 and 2)
- IRU bias, alignment, and scale factor (period 2)
- ACA alignment and field distortion coefficients (period 2)

5.6.3 On-orbit Calibrations

During the *Chandra* science mission, aspect system components require on-orbit calibration to compensate for alignment or scale factor drifts and to evaluate ACA CCD degradation due to cosmic radiation.

The following ACA calibrations are performed, as-needed, based on the trending analyses of aspect solution data.

IRU Calibration

The IRU calibration coefficients are evaluated quarterly and updated, as needed, approximately once every two years. These are impacted by IRU configuration changes, including the IRU gyro swap on 2020-Jul-31.

Dark Current

Cosmic radiation damage will produce an increase in both the mean CCD dark current and the non-Gaussian tail of "warm" (damaged) pixels in the ACA CCD. This is illustrated in Figure 5.8, which shows the distribution of dark current shortly after launch and in 2025-Sep. The background non-uniformity caused by warm pixels (dark current $> 100\ e^-/sec$) is the main contributor to star centroiding error, although the effect is substantially reduced by code within the aspect pipeline that detects and removes most warm pixels.

The fraction of pixels that appear to be warm during an observation is dependent on the underlying damage to the CCD and the CCD temperature during the observation. As the ACA thermoelectric cooler no longer has sufficient power margin to maintain the CCD at -19 °C, temperatures fluctuate between -14.0 °C and -2.1 °C. The effective warm pixel fraction over this temperature range is expected to be between 27 and 67% during the next year.

Dark current calibration data is gathered from as many full-frame CCD readouts as possible during engineering observations during the weekly schedule. After at least 10 full-frame readouts are collected, a dark current product based on the combined data from those readouts is created. A new calibration product from this process is evaluated and released for operational use approximately every month.

Flickering Pixels

The dark current of radiation damaged pixels is observed to fluctuate by factors of up to 25% on timescales of 1 to 50 ksec. This behavior was studied using a series of ACA monitor windows commanded during perigee passes in 2002. In 2008 and 2009, similar monitor window data were acquired and analyzed, and the flickering pixel behavior was seen to be qualitatively unchanged from that seen in 2002. An important consequence of the flickering pixel phenomenon is that the dark current pixel values obtained during the dark current calibration may not be directly subtracted from observation pixel data in post-facto processing. Instead, users of monitor window data should use a warm pixel detection algorithm such as the one implemented in ground processing (Cresitello-Dittmar et al., 2001).

Charge Transfer Inefficiency (CTI)

Radiation damage degrades the efficiency with which charge is transferred in the CCD by introducing dislocations in the semiconductor, trapping electrons and preventing their transfer. The most important consequence is a "streaking" or "trailing" of star images

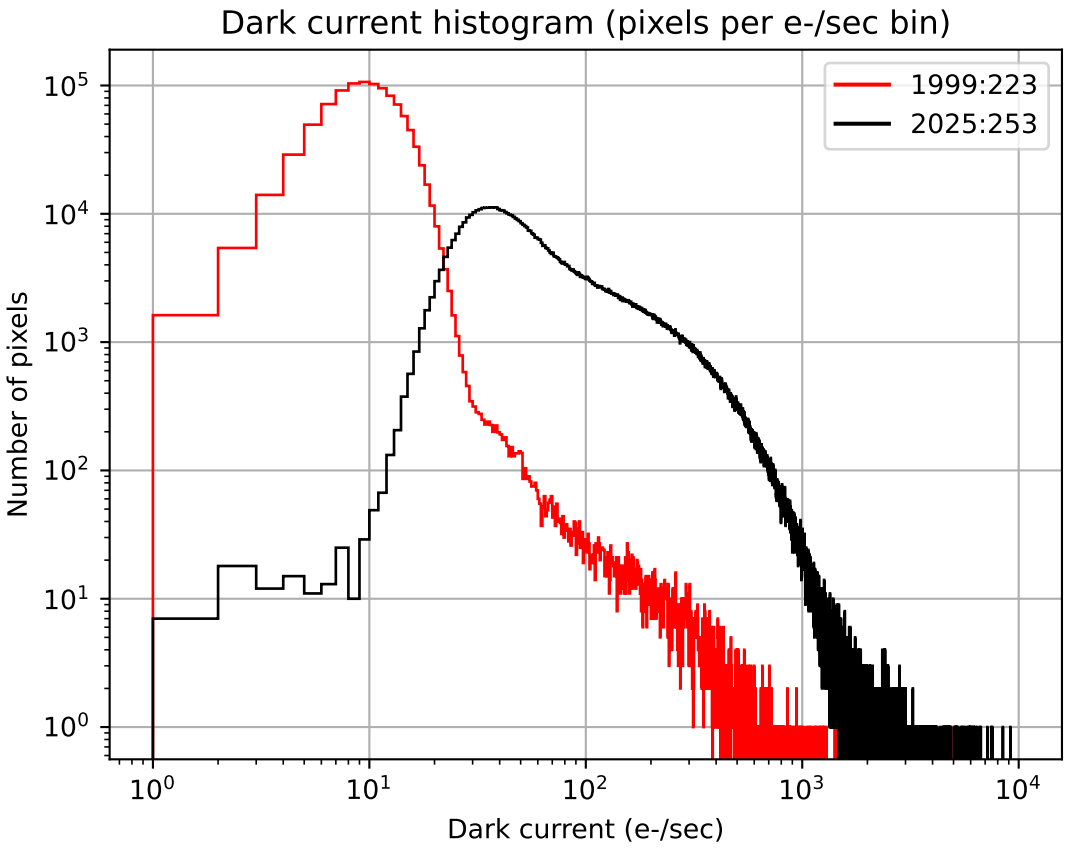

Figure 5.8: Differential histogram of dark current distribution for the ACA CCD on 1999-Aug-11 and 2025-Sep-10.

along the readout column(s), which can introduce systematic centroid shifts. These shifts depend primarily on CCD transfer distance to the readout and star magnitude.

The procedure for calibrating the mean CTI is to dither a faint star across the CCD quadrant boundary and observe the discontinuity in centroid (the CCD is divided electrically into four quadrants). In 2004, a total of 20 calibration observations were performed during perigee, each with a guide star dithering over a quadrant boundary. Despite significant concerns prior to launch, as well as notable CTI degradation in the ACIS front-illuminated chips, there is no evidence of increased CTI in the ACA CCD.

Field Distortion

The precise mapping from ACA CCD pixel position to angle relative to the ACA boresight is done with the "ACA field distortion polynomial." This includes plate scale factors up to third order as well as temperature-dependent terms. To verify that no mechanical shift in the ACA had occurred during launch, a field distortion calibration was performed during the orbital activation and checkout phase. The on-orbit calibration revealed no mechanical shift. Such a shift would have caused degraded celestial location accuracy.

The calibration was done by observing a dense field of stars with the spacecraft in normal pointing mode. Two reference stars were observed continuously, while sets of four stars each were observed for 100 sec. The calibration was completed after observing 64 stars over the ACA field of view, taking roughly 60 minutes. There are currently no plans to repeat this on-orbit calibration. Instead, the field distortion coefficients are monitored by long-term trending of observed star positions relative to their expected positions.

Responsivity

Contamination buildup on the CCD surface was predicted in pre-launch estimates to result in a mean throughput loss of 9% after five years on-orbit, although the calculation of this number has significant uncertainties. The buildup of contaminants is tracked by a trending analysis of magnitudes for stars which have been observed repeatedly throughout the mission (e.g. in the AR Lac field). To date, these trending analyses show no indication of contamination build-up. In the unlikely event that future contamination occurs and causes significant operational impact, "baking-out" the CCD on-orbit will be considered. In this procedure, the current to the CCD thermo-electric cooler is reversed so as to *heat* the device to approximately 30 °C for a period of several hours. After bake-out the CCD would be returned to its nominal operating temperature of −19 °C.

5.7 Operations

5.7.1 PCAD Modes

The PCAD system has nine operational modes (six normal and three safe) which use different combinations of sensor inputs and control mechanisms to control the spacecraft

Table 5.3: PCAD modes

Mode	Sensors	Control	Description
Standby	—	—	OBC commands to RWA, RCS, and Solar Array Drive Assembly (SADA) disabled, for initial deployment, subsystem checkout, etc.
Normal Pointing	IRU, ACA	RWA	Point at science target, with optional dither
Normal Maneuver	IRU	RWA	Slew between targets at peak rate of 2° per minute
Normal Sun	IRU, CSS, FSS	RWA	Acquire the Sun and hold spacecraft −Z axis and solar arrays to the Sun
Safe Sun	IRU, CSS, FSS	RWA	Safe mode: acquire the Sun and hold spacecraft −Z axis and solar arrays to the Sun
Derived Rate Safe Sun	IRU, CSS, FSS	RWA	Similar to Safe Sun Mode, but using only one gyro (two axes) plus Sun sensor data
Transfer orbit only—now disabled			
Powered Flight	IRU	RCS	Control *Chandra* during Liquid Apogee Engine burns
RCS Maneuver	IRU	RCS	Control *Chandra* using the RCS
RCS Safe Sun	IRU, CSS, FSS	RCS	Same as Safe Sun Mode, but using RCS instead of RWA for control

and ensure its safety. These modes are listed in Table 5.3. Normal science observations are carried out in Normal Pointing Mode (NPM), while slews between targets are done in Normal Maneuver Mode (NMM).

5.7.2 Operational Constraints

The ACA will meet performance requirements when the ACA line-of-sight is separated from: the Sun by 45 deg or more; the limb of the bright Earth by 10 deg or more; and the dark Earth or the Moon by 6 deg or more. If these restrictions are violated, the star images may be swamped by scattered background light, with the result that added noise on the star position will exceed the 0.360 arcsec requirement (1σ, 1-axis).

5.7.3 Output Data

The important output data from the ACA are the scaled raw pixel intensities in regions (4 pixels $\times$ 4 pixels, 6 pixels $\times$ 6 pixels, or 8 pixels $\times$ 8 pixels) centered on each of the star and fiducial light images. These data are placed in the engineering portion of the

telemetry stream, which is normally allocated 8 kbit/sec. During an ACA dark current calibration (Section 5.6.3), *Chandra* utilizes a 512 kbit/sec telemetry mode in real-time contact to enable read-out of the entire CCD (1024 pixels $\times$ 1024 pixels). The key data words in telemetry from the IRU are the 4 accumulated gyro counts (32 bits every 0.256 sec).

5.8 Performing an Observation

5.8.1 Star Acquisition

After maneuvering at a rate of up to 2°/minute to a new celestial location using gyroscope data and the reaction wheels, *Chandra* begins the star acquisition sequence, a process that typically takes from 1 to 4 minutes. First, the OBC commands the ACA to search for up to 8 acquisition stars, which are selected to be as isolated from nearby stars as possible. The search region size is based on the expected uncertainty in attitude, which is a function of the angular size of the slew. If two or more acquisition stars are found, an attitude update is performed using the best (brightest) pair of stars. This provides pointing knowledge to 3 arcsec (3σ per axis). Next, the guide star search begins. Depending on the particular star field configuration, the star selection algorithm may choose guide stars which are the same as the acquisition stars. In this case, the guide star acquisition time is somewhat reduced. When at least two guide stars have been acquired and pointing control errors converge, the on-board Kalman filtering is activated and the transition to Normal Point Mode is made. At this point sensing of the fiducial lights begins.

5.8.2 Science Pointing Scenarios

The on-board PCAD system is flexible and allows several different *Chandra* science pointing scenarios, described in the following sections.

Normal Pointing Mode Dither

The large majority of observations are performed using Normal Point Mode, with dither selected. In this case, the *Chandra* line-of-sight will be commanded through a Lissajous pattern. Dithering distributes photons over many detector elements (microchannel pores or CCD pixels) and serves several purposes: dithering reduces uncertainty due to pixel to pixel variation in quantum efficiency (QE); dithering allows sub-sampling of the image; and, in the case of the HRC, dithering distributes the total exposure over many microchannel pores—which is useful since the QE of a pore degrades slowly with exposure to photons. The dither pattern parameters are amplitude, phase, and period, each set independently for two axes. Each of the six parameters is separately commandable and differ for the two different instruments (See Chapters 6 and 7). The default values for these parameters are given in Table 5.4. Dither can be disabled for ACIS observations, while the minimum dither rate required to maintain the health of the HRC is 0.02 arcsec

sec^{-1}. The maximum dither rate, determined by PCAD stability requirements, is 0.22 arcsec sec^{-1}.

Table 5.4: Default dither parameters

Parameter	HRC	ACIS
Phase (pitch)	0.0 rad	0.0 rad
Phase (yaw)	0.0 rad	0.0 rad
Amplitude (pitch)	20.0 arcsec	16.0 arcsec
Amplitude (yaw)	20.0 arcsec	16.0 arcsec
Period (pitch)	768.6 sec	2000.0 sec
Period (yaw)	1087.0 sec	1414.0 sec

Normal Pointing Mode Steady

This mode is identical to Normal Pointing Mode Dither, but without the dither.

Pointing at Solar System Objects

Observations of moving solar system objects are done using a sequence of pointed observations, with the object moving through the field of view during each dwell period. Except in special circumstances, each pointing is selected so that the object remains within 5 arcmin of the *Chandra* line-of-sight. Most solar system objects move slowly enough that a single pointed observation will suffice.

Raster Scan

Survey scans of regions larger than the instrument field of view are specified simply with a grid, i.e., a list of target coordinates giving the field centers. The fields can optionally overlap, depending on the science requirements of the survey.

Offset and Gyro Hold

In special circumstances it will be necessary to perform observations without tracking guide stars. It may occur that a field has no suitable acquisition and guide stars, although this situation has not been encountered to date. A more likely situation is that a very bright object, such as Earth or the Moon, saturates the ACA CCD and precludes tracking stars. In this case, *Chandra* will first be maneuvered to a nearby pointing which has guide stars to establish fine attitude and a gyro bias estimate. A dwell time of approximately 20 minutes is needed to calibrate the bias estimate, which is the dominant term in the drift equation below. *Chandra* will then be maneuvered to the target. The default automatic transition to NPM will be disabled, and the spacecraft will hold on the target attitude in NMM.

After a maximum of 3.6 ksec, *Chandra* will be maneuvered back to the nearby field with guide stars to re-establish fine attitude and update the gyro drift rate. While holding on gyros only, the spacecraft attitude will drift due to noise in the gyros, which results in an aspect solution error. The variance of the angle drift for each gyro axis, in time t, is given by

$$\sigma^2 = \sigma_b^2 t^2 + \sigma_v^2 t + \sigma_u^2 t^3/3. \tag{5.1}$$

Ground test data for gyro noise parameters indicate worst case values of $\sigma_u = 1.5 \times 10^{-5}$ arcsec sec$^{-3/2}$ and $\sigma_v = 0.026$ arcsec sec$^{-1/2}$. Analysis of the residual Kalman filter bias estimate gives $\sigma_b = 0.002$ arcsec sec^{-1}. This results in 1σ angle drift errors of: 0.3 arcsec for 0.1 ksec; 2.2 arcsec for 1 ksec; 11 arcsec for 5 ksec; and 22 arcsec for 10 ksec.

5.8.3 PCAD Capabilities (Advanced)

Monitor Star Photometry

The ACA has the capability to devote one of the eight image slots to "monitor" a particular sky location. This allows simultaneous optical photometry of a target in the ACA field of view. However, since there are a fixed number of image slots, devoting a slot to photometry instead of tracking a guide star results in a degradation of the image reconstruction and celestial location accuracy (Section 5.4). Using a monitor slot represents a 15 - 25% increase in the aspect image reconstruction RMS diameter, depending on the particular guide star configuration. The photometric accuracy which can be achieved depends primarily on the star magnitude, integration time, CCD dark current, CCD read noise, sky background, and CCD dark current uncertainty.

The ACA capability to collect data for optical photometry does not have utility for optical sources brighter than $m_{ACA} = 5.2$ mag or sources fainter than $m_{ACA} = 10.3$ mag. The conversion from V and B magnitude to ACA instrument magnitude, based on flight data, is given approximately by

$$m_{ACA} = V + 0.426 - 1.06(B-V) + 0.617(B-V)^2 - 0.307(B-V)^3. \tag{5.2}$$

Dark current uncertainty ultimately limits the photometric accuracy at the faint end; this uncertainty results from uncalibrated pixel-to-pixel changes in dark current due to radiation damage. This includes both changing background pixels as *Chandra* dithers as well as intrinsic flickering in the radiation-damaged CCD pixels. This flickering, which occurs on time scales from less than 1 ksec to more than 10 ksec, poses fundamental problems for accurate photometry since the background dark current is a strong random function of time. With straightforward data processing, the noise introduced by the dark current variations (both spatial and temporal) at an operating temperature of -2.1 °C is approximately 1900 e$^-$ sec^{-1}. A star with an ACA magnitude of 9.9 mag produces about 7600 e$^-$ sec^{-1}, giving a signal-to-noise ratio (S/N) of 4.0. This represents the practical faint limit for ACA monitor star photometry. Somewhat improved S/N could

be obtained with a more sophisticated analysis that tracks the time-dependent dark current of each pixel. Users interested in processing ACA monitor window data should see https://cxc.harvard.edu/ciao/threads/monitor_photom/.

5.9 Ground Processing

For each science observation, the aspect system data described in Section 5.7.3 are telemetered to the ground to allow post-facto aspect determination by the *CXC* aspect pipeline as part of the standard *CXC* data processing pipeline. The important components of the pipeline are:

Gyro process: Filter gyro data, gap-fill, and calculate raw spacecraft angular rate

ACA process: Filter bad pixels, make CCD-level corrections (e.g. dark current), find spoiler stars, centroid, make camera-level corrections, convert to angle

Kalman filter and smooth: Optimally combine ACA and gyro data to determine ACA celestial location and image motion

Combine ACA and fids: Derive fiducial light solution and combine with ACA solution to generate image motion and celestial location at the focal-plane science instrument

5.9.1 Data Products

The data products which are produced by the aspect solution pipeline are listed in Table 5.5. Key data elements include: IRU accumulated counts; raw pixel data for 8 images; observed magnitudes and pixel positions of the aspect stars and fiducial lights versus time; and aspect solution versus time. The star data are used to determine the right ascension, declination, and roll (and corresponding uncertainties) of the HRMA axis as a function of time. The fiducial light images are used to track any drift of the SIM away from the nominal position. One cause of such drift is thermal warping of the optical bench assembly. The Kalman filtering routines also calculate an optimal estimate of the gyro bias rate as a function of time.

5.9.2 Star Catalog

The Aspect system uses the Advanced X-Ray Astrophysics Facility (AXAF) Guide and Aspect Star Catalog (*AGASC*) version 1.8. Further information about the *AGASC*, as well as access to catalog data, can be found on the *CXC AGASC* web page (https://cxc.harvard.edu/agasc). The *AGASC* was prepared by the *CXC* Mission Planning and Operations & Science Support groups, and is a compilation of catalogs including the *Hubble* Guide Star Catalog 1.1, the Positions and Proper Motion Catalog, the Tycho Output Catalog, the ACT Reference Catalog, the Tycho-2 Catalog, the GSC-ACT Catalog, the 2MASS Galaxy Catalog, the AAVSO Photometric All Sky Survey (DR9), and the Gaia Catalog (DR3).

Table 5.5: Aspect pipeline data products

Product	Description
ASPSOL	Final aspect solution with errors
ASPQUAL	Aspect solution quality indicators
AIPROPS	Aspect Intervals
ACACAL	ACA calibration data from the optical database (ODB) and CALDB
GSPROPS	Guide star properties, both from the AXAF Guide and Acquisition Star Catalog and as actually observed with the ACA
FIDPROPS	Fiducial light properties, as commanded and as observed
ACADATA	Aspect camera telemetry (including ACA housekeeping) and images after CCD-level correction
ACACENT	Image centroids and associated fit statistics
GYROCAL	Gyro calibration data from ODB and CALDB
GYRODATA	Gyro raw data and gap-filled, filtered data
KALMAN	Intermediate and final data in Kalman filter and smoother

5.10 References

Cresitello-Dittmar, M., Aldcroft, T. L., & Morris, D. 2001, ADASS X, 238, 439

Getman, K. V., Flaccomio, E., Broos, P. S., et al. 2005, ApJS, 160, 319

Chapter 6

ACIS: Advanced CCD Imaging Spectrometer

6.1 Introduction & Layout

The Advanced CCD Imaging Spectrometer (ACIS) offers the capability to simultaneously acquire high-resolution images and moderate resolution spectra. The instrument can also be used in conjunction with the High Energy Transmission Grating (HETG) or Low Energy Transmission Grating (LETG) to obtain higher resolution spectra (see Chapters 8 and 9). ACIS contains 10 planar, 1024×1024 pixel CCDs (Figure 6.1); four arranged in a 2×2 array (ACIS-I) used for imaging, and six arranged in a 1×6 array (ACIS-S) used either for imaging or for a grating spectrum read-out. Two CCDs are back-illuminated (BI) and eight are front-illuminated (FI). The response of the BI devices extends to energies below that accessible to the FI chips. The chip-average energy resolution of the BI devices is better than that of the FI devices.

In principle, any combination of up to 6 CCDs can be operated simultaneously. However, because of changes in the thermal environment, proposed observations can only specify up to 4 required CCDs; a 5th or 6th CCD can only be proposed as optional. Optional CCDs may or may not be available depending on prevailing thermal conditions at the time of observation (see also Section 6.20).

The original Instrument Principal Investigator for ACIS is Prof. Gordon Garmire (Pennsylvania State University, now at Huntingdon Institute for X-ray Astrophysics and Space Research). ACIS was developed by a collaboration between Penn State, the *MIT* Kavli Institute for Astrophysics and Space Research, and the Jet Propulsion Laboratory, and was built by Lockheed Martin and *MIT*. The *MIT* effort was led by Dr. George Ricker. The CCDs were developed by *MIT*'s Lincoln Laboratory.

ACIS is a complex instrument having many different characteristics and operating modes. Radiation damage suffered by the FI chips has had a negative impact on their energy resolution—the BI devices were not affected—thus affecting the basic considerations as to how to make best use of the instrument. Trade-offs are discussed in this chapter.

ACIS FLIGHT FOCAL PLANE

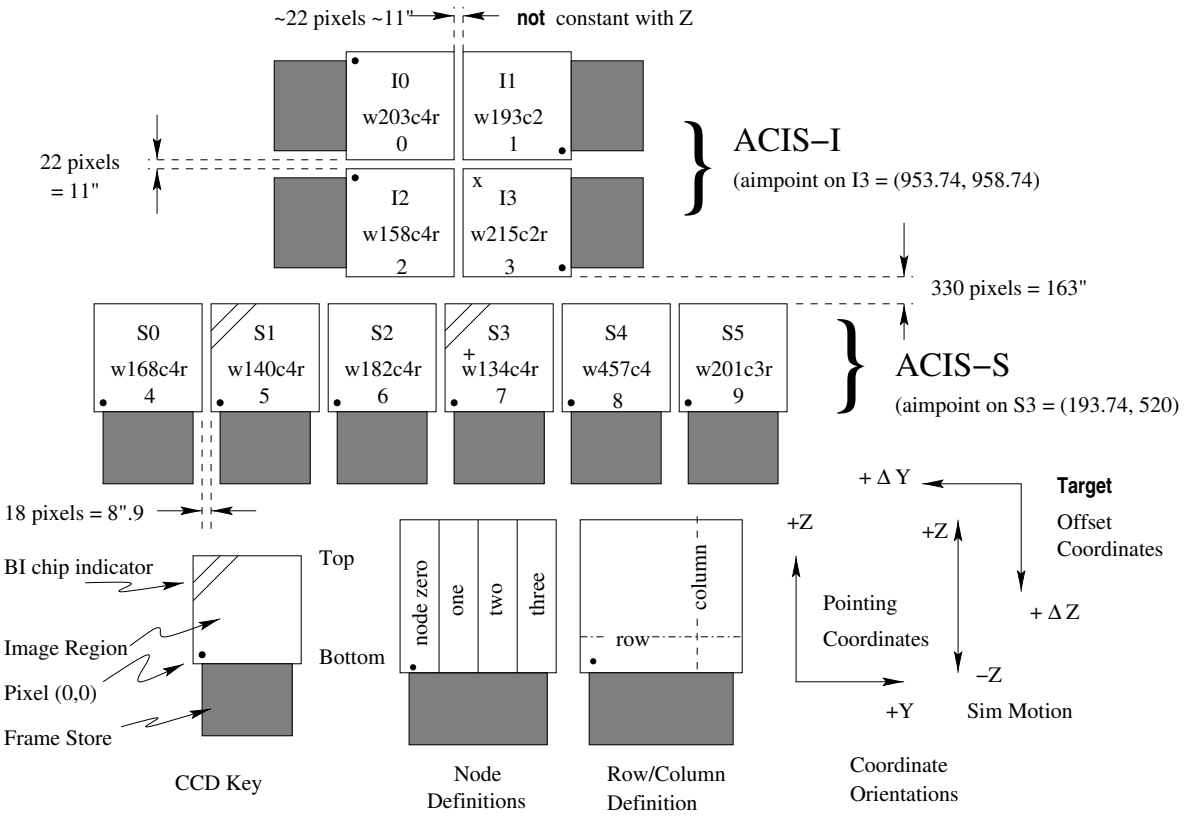

Figure 6.1: A schematic drawing of the ACIS focal plane; insight to the terminology is given in the lower left. Note the aimpoints: on S3 (the '+') and on I3 (the 'x'). Note the differences in the orientation of the I and S chips, important when using subarrays (Section 6.13.1). Note also the (Y, Z) coordinate system and the target offset convention (see Chapter 3) as well as the SIM motion (+/−Z). This view is along the optical axis, from the sky toward the detectors, (−X). The numerous ways to refer to a particular CCD are indicated: chip letter+number, chip serial number, and ACIS chip number (CCD_ID); see also Table 6.1. As indicated, S3 and S1 are back-illuminated (BI) CCDs, and the rest are front-illuminated (FI) CCDs. The node numbering scheme and the row/column directions are illustrated lower center. The row and column pixel indices run from 0 to 1023. The row direction corresponds to `CHIPX` (1 to 1024 pixels). The column direction corresponds to `CHIPY` (1 to 1024 pixels); the row at `CHIPY` = 1 is closest to the frame store.

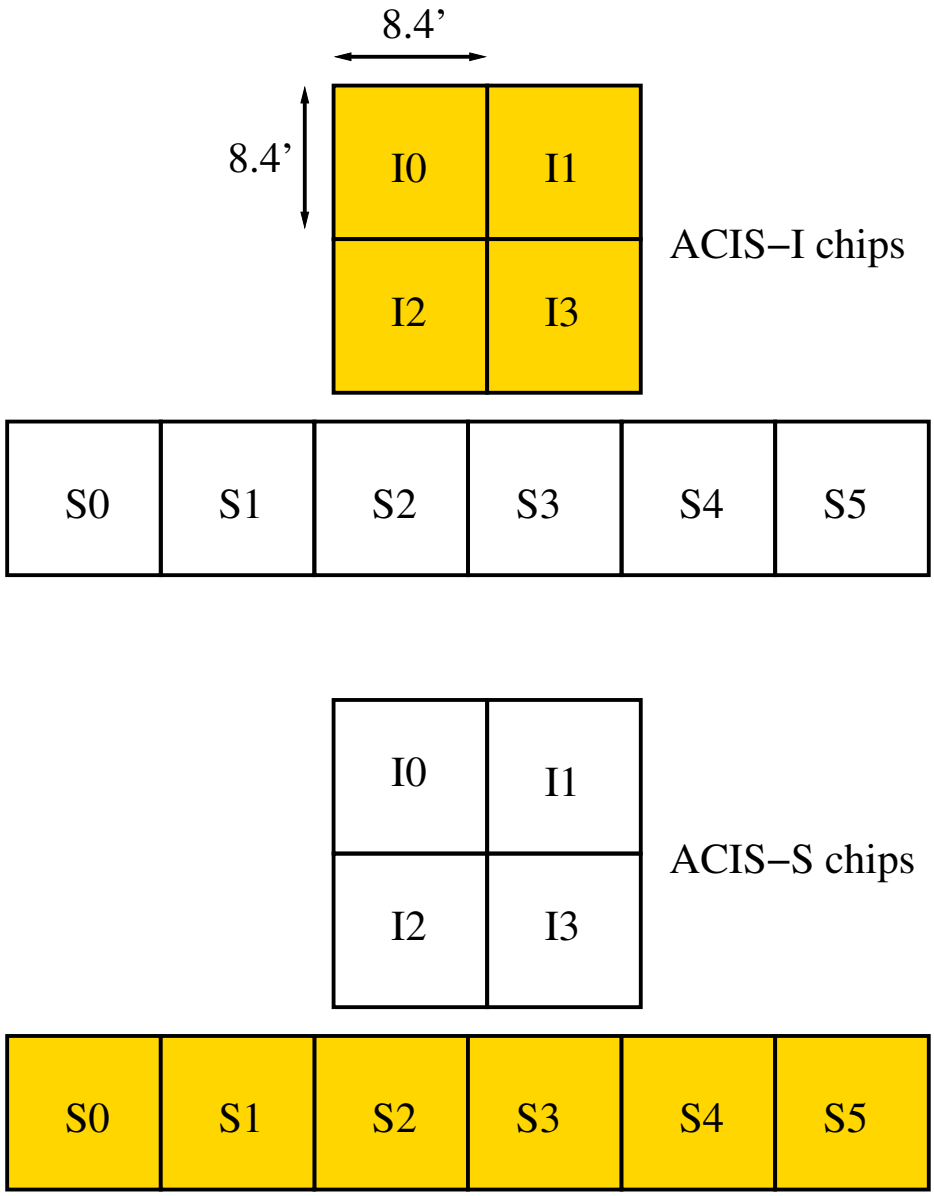

Figure 6.2: A schematic drawing of the ACIS focal plane, not to scale. The ACIS-I array consists of chips I0-I3 (shaded in the upper figure). The ACIS-S array consists of chips S0-S5 (shaded in the lower figure). The view is along the optical axis in each case, from the sky toward the detectors (−X); see also Figure 6.1 for more details. See the discussion in Section 6.20 for information on selecting CCDs.

CCD_ID	CCD Name	CCD Type*	Serial Number
0	I0, ACIS-I0	FI	w203c4r
1	I1, ACIS-I1	FI	w193c2
2	I2, ACIS-I2	FI	w158c4r
3†	I3, ACIS-I3	FI	w215c2r
4	S0, ACIS-S0	FI	w168c4r
5	S1, ACIS-S1	**BI**	w140c4r
6	S2, ACIS-S2	FI	w182c4r
7‡	S3, ACIS-S3	**BI**	w134c4r
8	S4, ACIS-S4	FI	w457c4
9	S5, ACIS-S5	FI	w201c3r

†contains I-array aimpoint; best imaging
‡contains S-array aimpoint; best imaging
*FI = front-illuminated, **BI** = back-illuminated

Table 6.1: Naming conventions for the ACIS CCDs. See Figure 6.1 for the positions and orientations in the flight focal plane.

Software methods for improving the energy resolution of the FI CCDs are discussed in Section 6.7.1. The low energy response of ACIS has also been affected by the buildup of a contaminant on the optical blocking filters and this is discussed in Section 6.5.1.

Many of the characteristics of the ACIS instrument are summarized in Table 6.2.

6.2 Basic Principles

A CCD is a solid-state electronic device composed primarily of silicon. A "gate" structure on one surface defines the pixel boundaries by alternating voltages on three electrodes spanning a pixel. The silicon in the active (depletion) region (the region below the gates wherein most of the absorption takes place) has an applied electric field so that charge moves quickly to the gate surface. The gates allow confined charge to be passed down a "bucket brigade" (the buried channel) of pixels in parallel to a serial read-out at one edge by appropriately varying ("clocking") the voltages in the gates.

The ACIS front-illuminated CCDs have the gate structures facing the incident X-ray beam. Two of the chips on the ACIS-S array (S1 and S3) have had the back sides of the chips treated, removing insensitive, undepleted bulk silicon material and leaving the photo-sensitive depletion region exposed. These are the BI chips and are deployed with the back side facing the HRMA.

Photoelectric absorption of an X-ray photon in silicon results in the liberation of a proportional number of electrons: an average of one electron-hole pair for each 3.7 eV of energy absorbed. Immediately after the photoelectric interaction, the charge is confined

Focal-plane arrays:	
I-array	4 CCDs placed tangent to the focal surface
S-array	6 CCDs in a linear array tangent to the grating Rowland circle
CCD format	1024 by 1024 pixels
Pixel size	23.985 microns (0.4920±0.0001 arcsec)
Array size	16.9 by 16.9 arcmin ACIS-I 8.4 by 51.1 arcmin ACIS-S
On-axis effective Area (Cycle 1)	110 cm^2 @ 0.5 keV (FI)
(integrated over the PSF	600 cm^2 @ 1.5 keV (FI)
to >99% encircled energy)	40 cm^2 @ 8.0 keV (FI)
Quantum efficiency†	
Front-illuminated CCD	$>$ 80% between 3.0 and 6.5 keV $>$ 30% between 0.7 and 11.0 keV
Back-illuminated CCD	$>$ 80% between 0.8 and 5.5 keV $>$ 30% between 0.4 and 10.0 keV
Charge transfer inefficiency (Cycle 1):	
(parallel)	FI: $\sim 2\times10^{-4}$; BI: $\sim 2\times10^{-5}$
(serial)	S3(BI): $\sim 7\times10^{-5}$; S1(BI): $\sim 1.5\times10^{-4}$; FI: $< 2 \times 10^{-5}$
System noise	$\lesssim$ 2 electrons (RMS) per pixel
Maximum read-out-rate per channel	$\sim$ 100 kpix/sec
Number of parallel signal channels	4 nodes per CCD
Pulse-height encoding	12 bits/pixel
Event threshold	FI: 38 ADU (∼150–350 eV) BI: 20 ADU (∼150–220 eV)
Split threshold	13 ADU
Maximum internal data-rate	6.4 Mbs (100 kbs $\times 4 \times 16$)
Output data-rate	24 kb per sec
Minimum row read-out time	2.8 ms
Nominal frame time	3.2 sec (full frame)
Allowable frame times	0.2 to 10.0 s
Frame transfer time	40 μsec (per row)
Point-source sensitivity (0.4-6.0 keV)	$4 \times 10^{-15}\,\mathrm{ergs\,cm^{-2}\,s^{-1}}$ in 10^4 s
Detector operating temperature	-90 °C to -120 °C

†Does not include contaminant layer absorption.

Table 6.2: ACIS Characteristics

by electric fields to a small volume near the interaction site. Charge in a FI device can also be liberated below the depletion region, in an inactive substrate, from where it diffuses into the depletion region. This charge may easily appear in two or more pixels.

Good spectral resolution depends upon an accurate determination of the total charge deposited by a single photon. This in turn depends upon the fraction of charge collected, the fraction of charge lost in transfer from pixel to pixel during read-out, and the ability of the read-out amplifiers to measure the charge. Spectral resolution also depends on read noise and the off-chip analog processing electronics. The ACIS CCDs have read-out noise less than 2 electrons RMS. Total system noise for the 40 ACIS signal chains (4 nodes/CCD) ranges from 2 to 3 electrons (RMS) and is dominated by the off-chip analog processing electronics.

The CCDs have an "active" or imaging section (see Figure 6.1) which is exposed to the incident radiation and a shielded "frame store" region. A typical mode of the ACIS CCD operation is: (1) the active region is exposed for a fixed amount of time (full frame $\sim$ 3.2 s); (2) at the end of the exposure, the charge in the active region is quickly ($\sim$ 41 ms) transferred in parallel into the frame store; (3) the next exposure begins; (4) simultaneously, the data in the frame store region are transferred serially to a local processor which, after removing bias (see Section 6.14), identifies the position and amplitude of any "events" according to a number of criteria depending on the precise operating mode. These criteria always require a local maximum in the charge distribution above the event threshold (see Table 6.2). The position and the amount of charge collected, together with similar data for a limited region containing and surrounding the pixel are classified ("graded") and then passed into the telemetry stream.

6.3 Optical Blocking Filter & Optical Contamination

Since the CCDs are sensitive to optical as well as X-ray photons, optical blocking filters (OBFs) are placed just over the CCDs between the chips and the HRMA. The filters are composed of polyimide (a polycarbonate plastic) sandwiched between two thin layers of aluminum. The nominal thicknesses of these components for the two arrays are given in Table 6.3. Details of the calibration of these filters may be found in Chapter 5 the ACIS calibration report at https://cxc.harvard.edu/cal/Acis/Papers/cal_report.pdf. These calibrations do not include the more recent effects of molecular contamination. Molecular contamination is discussed in Section 6.5.1.

Table 6.3: Nominal Optical Blocking Filter Composition and Thicknesses

ACIS-I	Al/Polyimide/Al	1200Å 2000Å 400Å
ACIS-S	Al/Polyimide/Al	1000Å 2000Å 300Å

The brightness threshold for optical contamination (here defined as a 1 analog-to-digital unit (ADU) shift per pixel in the signal) is based on on-orbit calibration observa-

tions of two stars: Vega and Betelgeuse. The ratio of the signals leads us to believe that there is a light leak in the I or J band in the near infrared part of the spectrum. The light leak has decreased with time; the most recent observations of Vega and Betelgeuse show a substantial further decrease. The threshold I and J magnitudes are ~ -1.0. for S3 and ~ -2.0 for I3. Brighter stars can be observed, but require special techniques in either data analysis, observation planning, or both. Details can be found on this web page: https://cxc.harvard.edu/cal/Hrma/UvIrPSF.html. The *CXC* HelpDesk (https://cxc.harvard.edu/helpdesk/) can help with detailed questions.

6.4 Calibration

Measurements of ACIS include laboratory calibrations, a system-level ground calibration of the HRMA and ACIS at the X-Ray Calibration Facility (XRCF) at MSFC, and on-orbit calibration using celestial and on-board radioactive X-ray sources. The ACIS external calibration source (ECS) consists of an ^{55}Fe source (half-life 2.7 years) and a target made of aluminum and titanium. The source emits five strong lines (Al Kα at 1.49 keV, Ti Kα and Kβ, at 4.51 and 4.93 keV, and Mn Kα and Kβ at 5.90 and 6.49 keV). A number of weaker lines are also present.

The on-orbit calibration of ACIS is a continuing activity. All calibration data are, or will be, described in detail, at https://cxc.harvard.edu/cal/. Consult the WWW site and its links for the latest information.

6.5 Quantum Efficiency and Effective Area

The quantum efficiencies near the read-out for the ACIS CCDs for the standard grade set, including optical blocking filters and molecular contamination, are shown in Figure 6.3. Note that the quantum efficiency for the FI chips varies somewhat with row number (not shown), and decreases by 5–15% farthest from the read-out at energies above about 4 keV. This is due to the migration of good grades to bad grades produced by charge transfer inefficiency (CTI), which varies with row number. The quantum efficiency (QE) variation with position for the BI chips is much smaller.

Cosmic rays tend to cause large blooms on FI chips, and much smaller ones on BI chips. This results in a 2–4% decrement of the QE for FI chips and $\sim 0.5\%$ for BI chips.

The combined HRMA/ACIS on-axis effective areas are shown in Figure 6.4 (log energy scale) and 6.5 (linear energy scale). The effective areas are for an on-axis point source and a 20 arcsec-diameter extraction region. The ACIS effective areas include a correction that accounts for the buildup of molecular contamination on the ACIS filters (see the discussion in Section 6.5.1). Figures 6.4 and 6.5 show the predicted ACIS effective areas for the middle of Cycle 28 based on the current time-dependent ACIS contamination models.

Figure 6.6 shows the vignetting (defined as the ratio of off-axis to on-axis effective area) as a function of energy at several off-axis angles. These data are from a calibration

observation of G21.5-0.9, a bright supernova remnant/pulsar wind nebula. The detector was appropriately offset for each off-axis angle so that the data were obtained at the same focal position, minimizing the effects of any spatially-dependent variations in the CCD response.

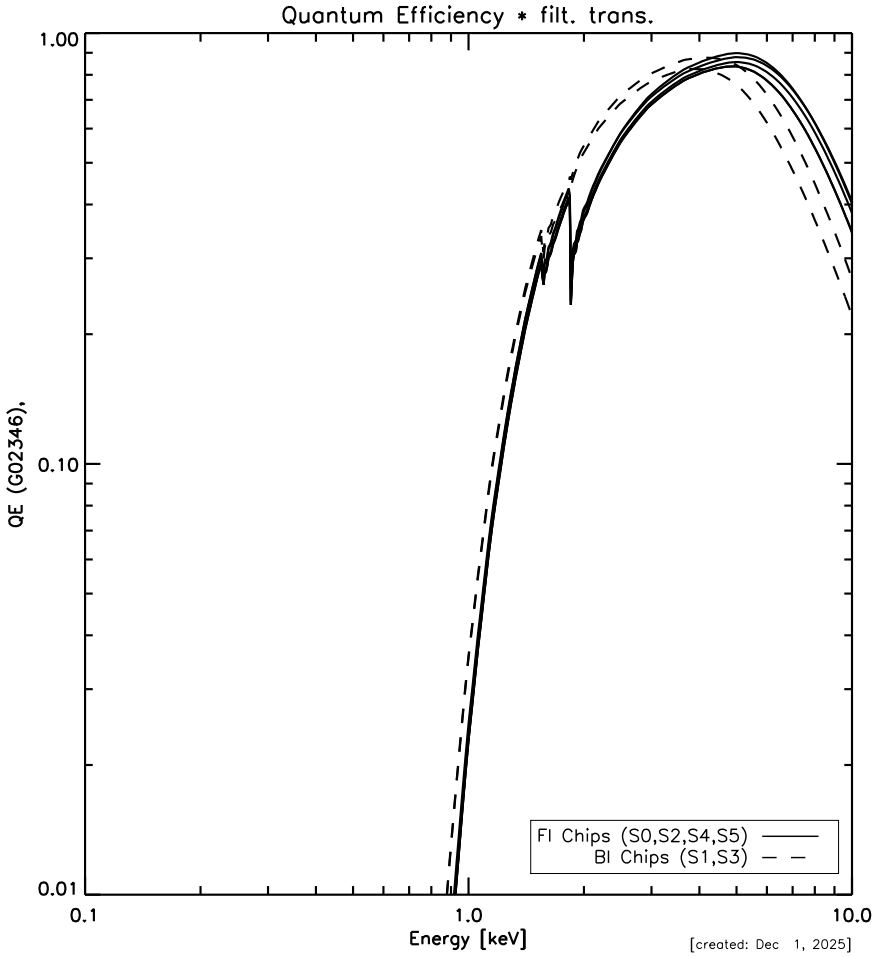

Figure 6.3: The quantum efficiency (convolved with the transmission of the appropriate optical blocking filter) of the FI CCDs (from a row nearest the read-out) and the two BI CCDs as a function of energy. S3 is somewhat thicker, hence more efficient, than S1. These curves include the effects of molecular contamination, as discussed in the text.

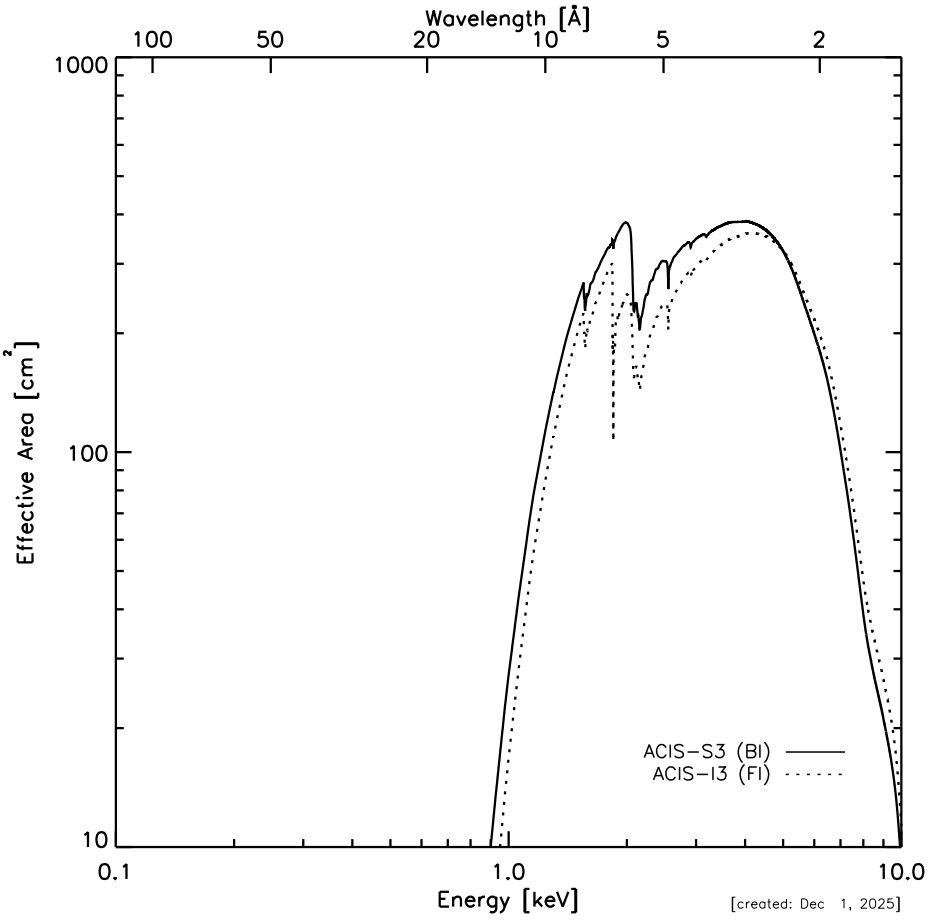

Figure 6.4: The HRMA/ACIS predicted effective area versus the energy on a log scale. The dashed line is for the FI CCD I3, and the solid line is for the BI CCD S3. These curves include the effects of molecular contamination, as discussed in the text.

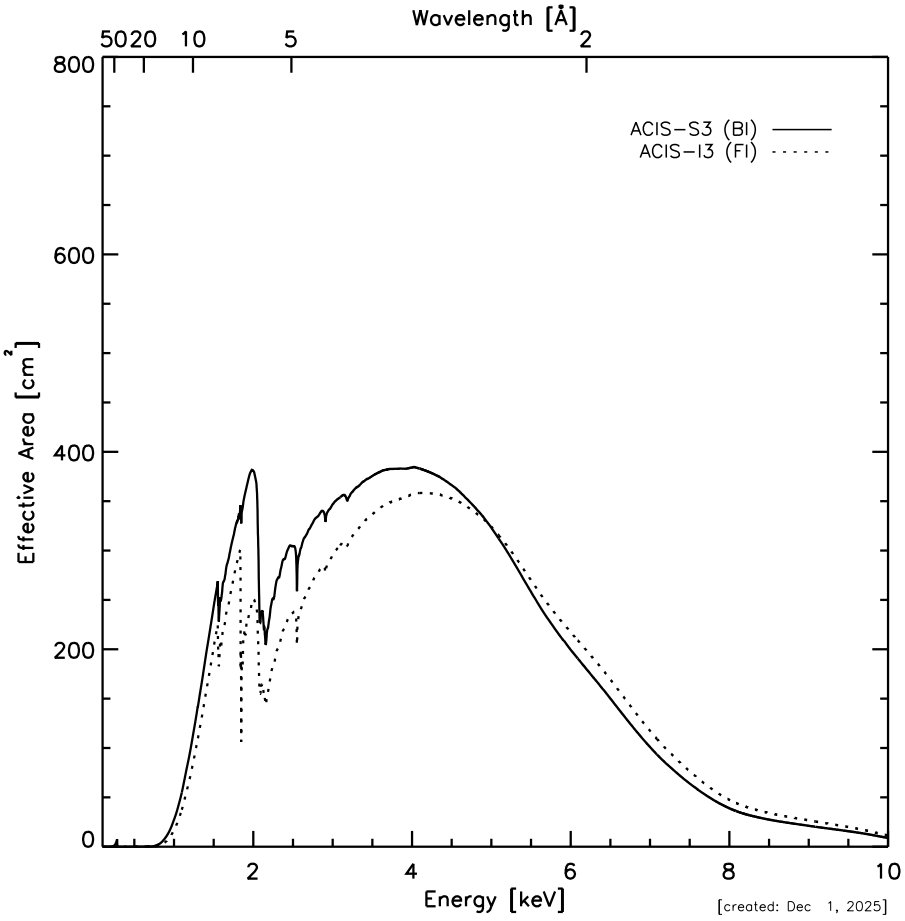

Figure 6.5: The HRMA/ACIS predicted effective area versus the energy on a linear scale. The dashed line is for the FI CCD I3, and the solid line is for the BI CCD S3. These curves include the effects of molecular contamination, as discussed in the text.

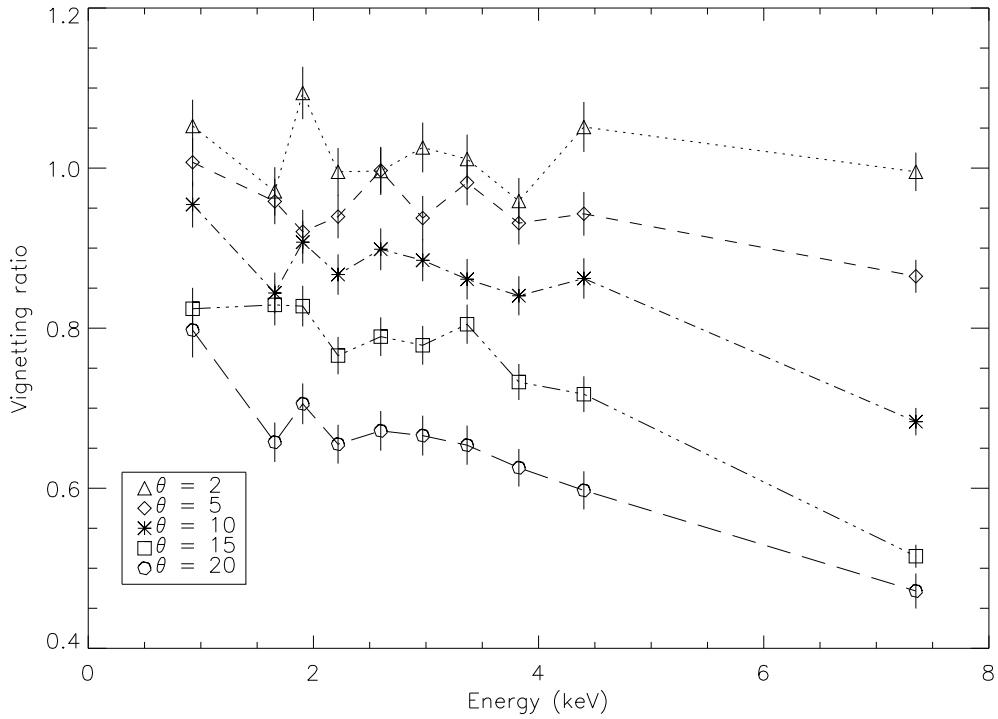

Figure 6.6: Vignetting (the ratio of off-axis to on-axis effective area) as a function of energy for several off-axis angles in arcmin.

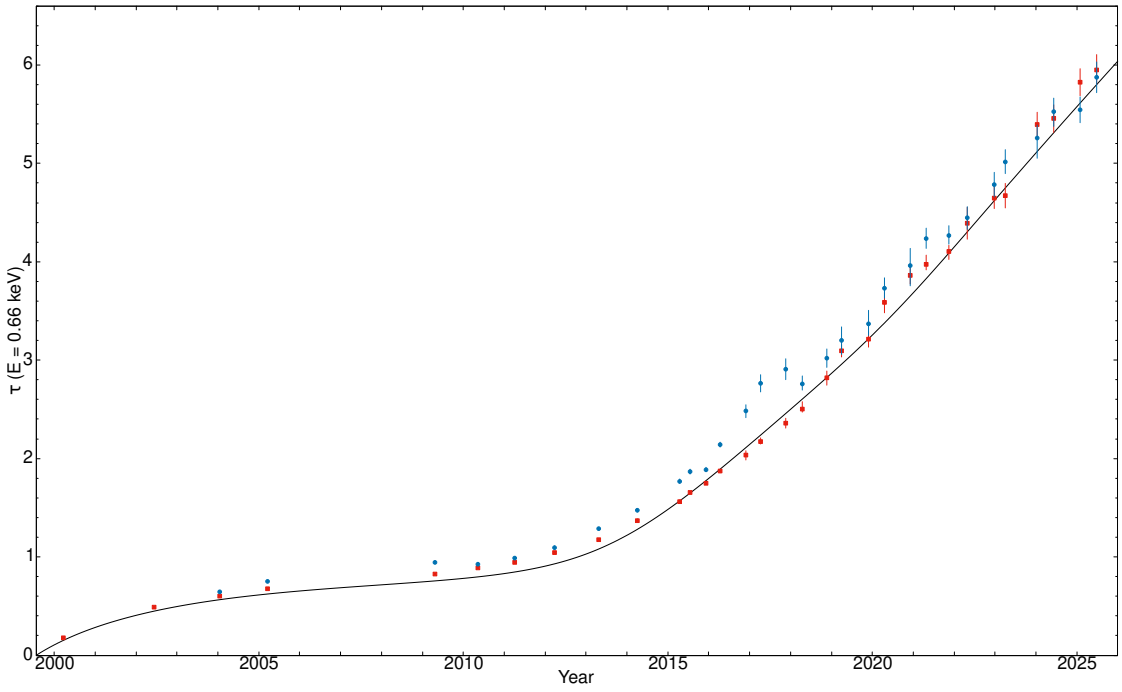

Figure 6.7: Optical depth of the contaminant measured at 0.66 keV from observations of the galaxy cluster Abell 1795. Red squares denote observations with the target on ACIS-S and blue disks denote observations with the target on ACIS-I. The solid curve shows the optical depth of the contaminant as modeled in the most recent contamination model.

6.5.1 Molecular Contamination of the OBFs

Astronomical observations and data acquired from the on-board ACIS external calibration source (ECS) show that the ACIS effective area below 2 keV has continuously declined since launch due to the build-up of out-gassed material on the cold ACIS optical blocking filters. The HRC operates at a warmer temperature and shows no sign of contamination build-up. The build-up of contaminant on the ACIS filters has been monitored with ECS observations, LETG/ACIS-S observations of the blazars Mkn421 and PKS2155-304 and the isolated neutron star RXJ1856-3754, and ACIS imaging observations of the rich cluster of galaxies Abell 1795 and the oxygen-rich supernova remnant E0102-72.3.

Typically the calibration team releases updates to the contamination model near the end of each calendar year. The update accounts for any changes in the characteristics of the contamination as measured by observations taken over the same calendar year. Most updates do not affect the analysis of data taken prior to that calendar year. The contamination model is distributed as part of the CALDB, and is described at: https://cxc.cfa.harvard.edu/ciao/why/acisqecontam.html. There are three components to the ACIS contamination model: 1) the build-up rate of the contaminant on the

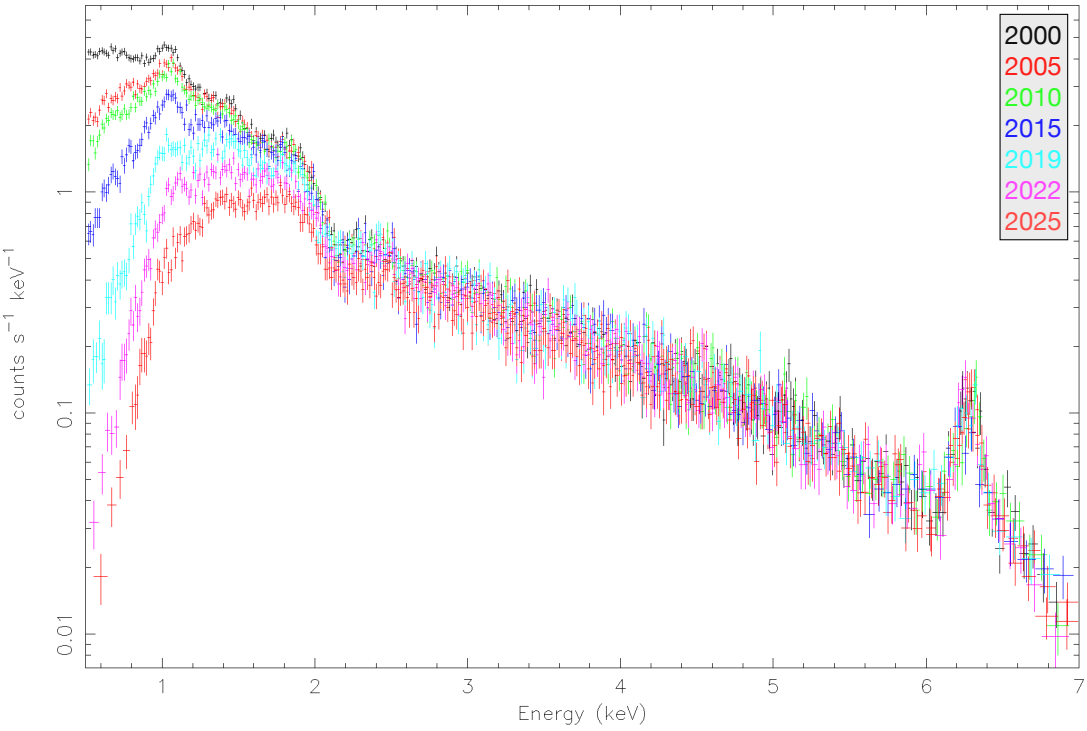

Figure 6.8: ACIS-S spectrum of the galaxy cluster Abell 1795 at seven different epochs ranging from 2000 to 2025. The decreasing count rates in the spectrum of a stable calibration source demonstrate the loss of effective area due to the build-up of the contaminant on the ACIS OBF.

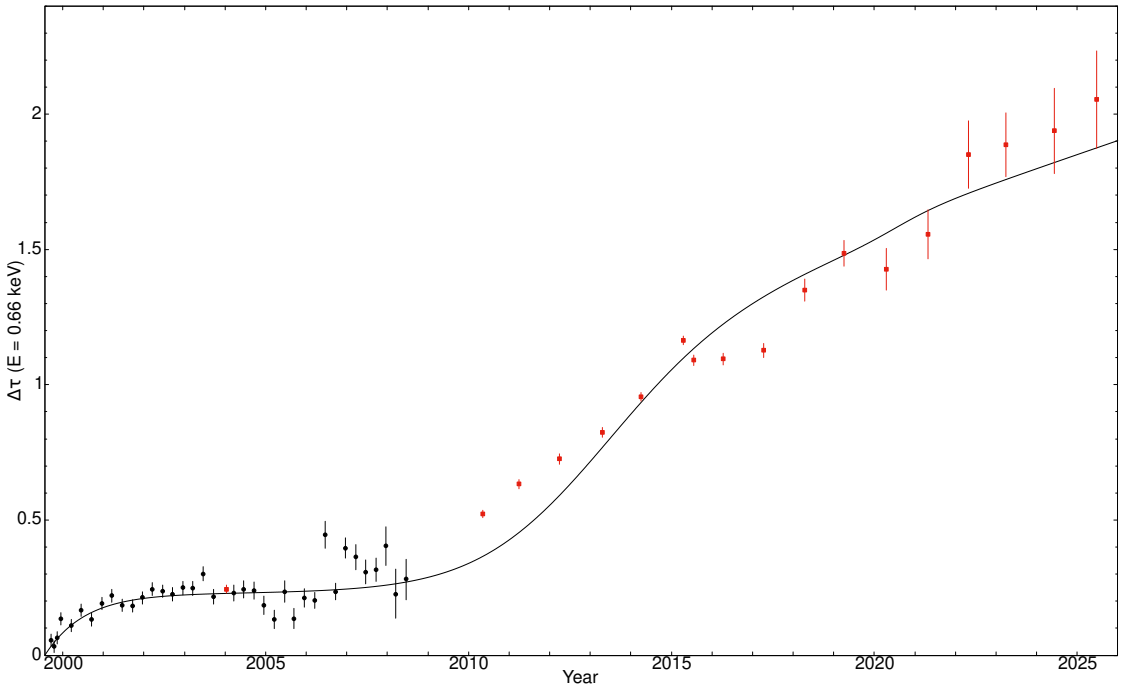

Figure 6.9: The difference between the optical depth at 0.66 keV measured near the center of the ACIS-S chip S3 and that measured near the readout at the edge of the same chip is shown against time. Black circles are measurements made using the ECS data and red squares are made from observations of the galaxy cluster Abell 1795. The solid curve shows the optical depth of the contaminant as modeled in the most recent contamination model.

ACIS filters, 2) the spatial distribution of the contaminant on the ACIS filters, and 3) the chemical composition of the contaminant and hence its energy-dependent effects. The energy-dependent effect is largest near the Carbon edge at 0.277 keV and is negligible beyond $\sim$ 1–2 keV. Figures 6.7–6.11 show that all three of these properties concerning the contaminant have changed significantly during the *Chandra* mission.

Figure 6.7 shows that the optical depth of the contaminant appeared to be leveling off prior to about 2009, at which time the condensation rate of the contaminant onto the ACIS filter began increasing. Since 2014 the contaminant accumulates on the optical blocking filters at a faster rate. The effect of the contaminant build-up can be also illustrated using the X-ray energy spectrum from the galaxy cluster Abell 1795. Fig. 6.8 shows the ACIS-S spectrum of the central regions of Abell 1795 observed at different epochs from early in the *Chandra* mission to the present day. The loss of the effective area due to the contaminant is demonstrated by the decreasing count rates at low energies.

In addition, the contaminant has always been thicker near the edges of the filters. Figure 6.9 shows the difference in the optical depth on the filter measured near the center

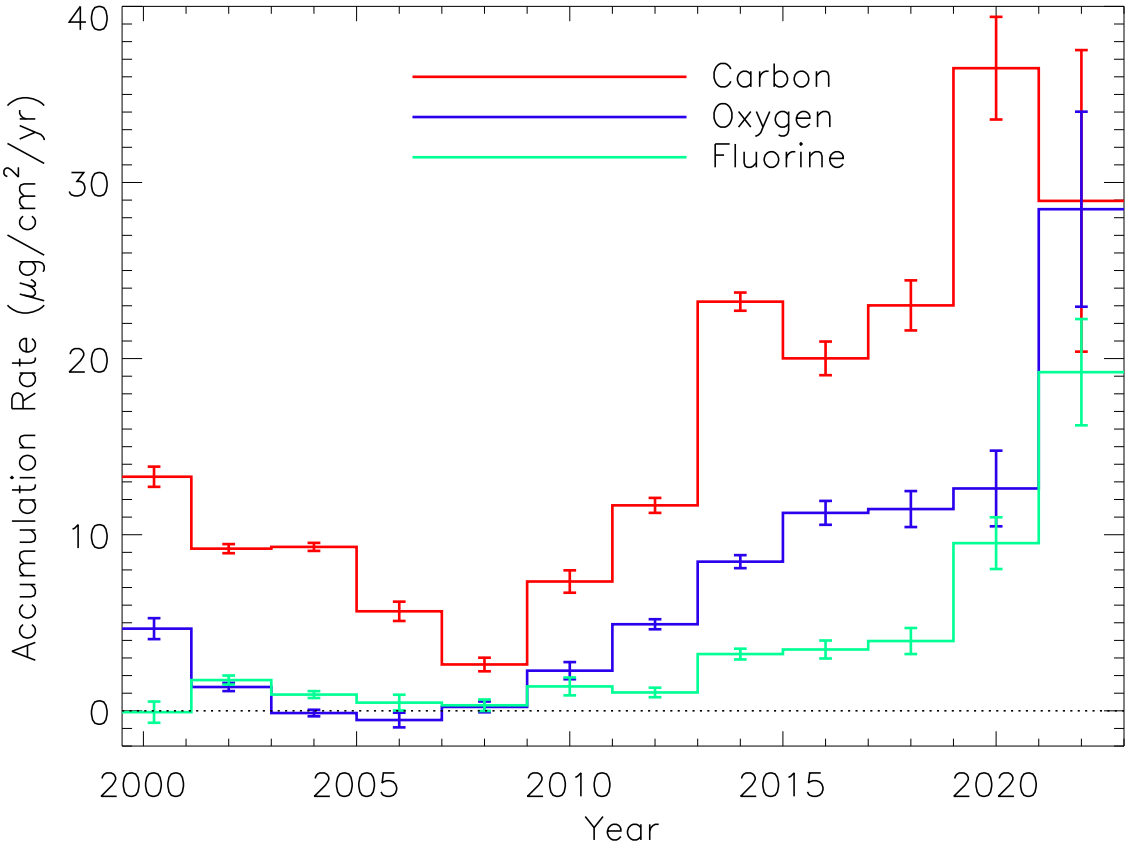

Figure 6.10: The accumulation rate for C, O, and F over the *Chandra* mission using LETG/ACIS-S observations of blazars.

of the chip and near the edge. The difference in optical depth appeared to be leveling off prior to 2009, but has steadily increased since 2009. The data for the edge of the chip are taken at CHIPY=64 (near the readout) for ACIS-S. On ACIS-I the spatial variations are similar to those on ACIS-S with more contaminant on the outer edges of the array.

LETG/ACIS-S observations show that the contaminant is composed of carbon, oxygen and fluorine. Figure 6.10 shows the accumulation rate for C, O, and F in μg cm^{-2} year^{-1}. Figure 6.11 shows that the oxygen-to-carbon ratio of the material condensing onto the ACIS filters has varied significantly during the mission indicating the presence of several different contaminants which have dominated the condensation rate onto the ACIS filters at different times.

To assist with the proposal process, Portable Interactive Multi-Mission Software (*PIMMS*) tables are generated for the ACIS effective area based on an extrapolation of the current contamination model to the middle of the next Cycle, which is an extrapolation of approximately 18 months. Based on recent changes in the behavior of the out-gassed material condensing onto the ACIS filters, proposers are advised to include an uncertainty

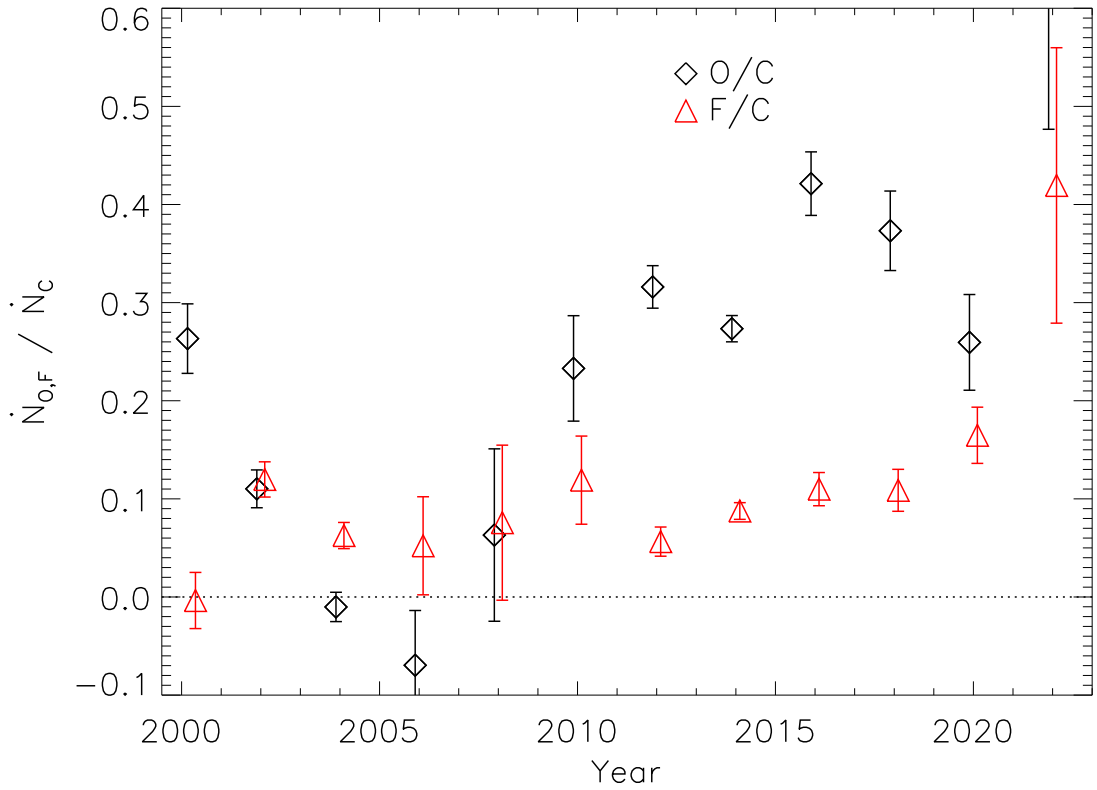

Figure 6.11: The ratio of the deposition rate in oxygen atoms and fluorine atoms per year to that for carbon over the *Chandra* mission as measured by LETG/ACIS-S observations of blazars.

of about 10% in the predicted ACIS on-axis count rates below about 1 keV, and about 20% near the edges of the ACIS-S and ACIS-I arrays.

6.6 Spatial Resolution, PSF, & Encircled Energy

The spatial resolution for on-axis imaging with HRMA/ACIS is limited by the physical size of the CCD pixels (23.985 μm square ∼0.492 arcsec) and not the HRMA. This limitation applies regardless of whether the aimpoint is selected to be the default aimpoint on I3 or S3 (Figure 6.1). Approximately 90% of the encircled energy lies within 4 pixels (2 arcsec) of the center pixel at 1.49 keV and within 5 pixels (2.5 arcsec) at 6.4 keV. Figure 6.12 shows an in-flight calibration. There is no evidence for any differences in data taken with either S3 or I3 at the nominal focus. The ACIS encircled energy as a function of off-axis angle is discussed in Chapter 4 (see Section 4.2.2 and Figure 4.13).

Off-axis, the departure of the CCD layout from the ideal focal surface and the increase of the HRMA PSF with off-axis angle become dominating factors. Since the ideal focal

surface depends on energy, observers for whom such considerations may be important are urged to make use of the *MARX* simulator to study the impact on their observation.

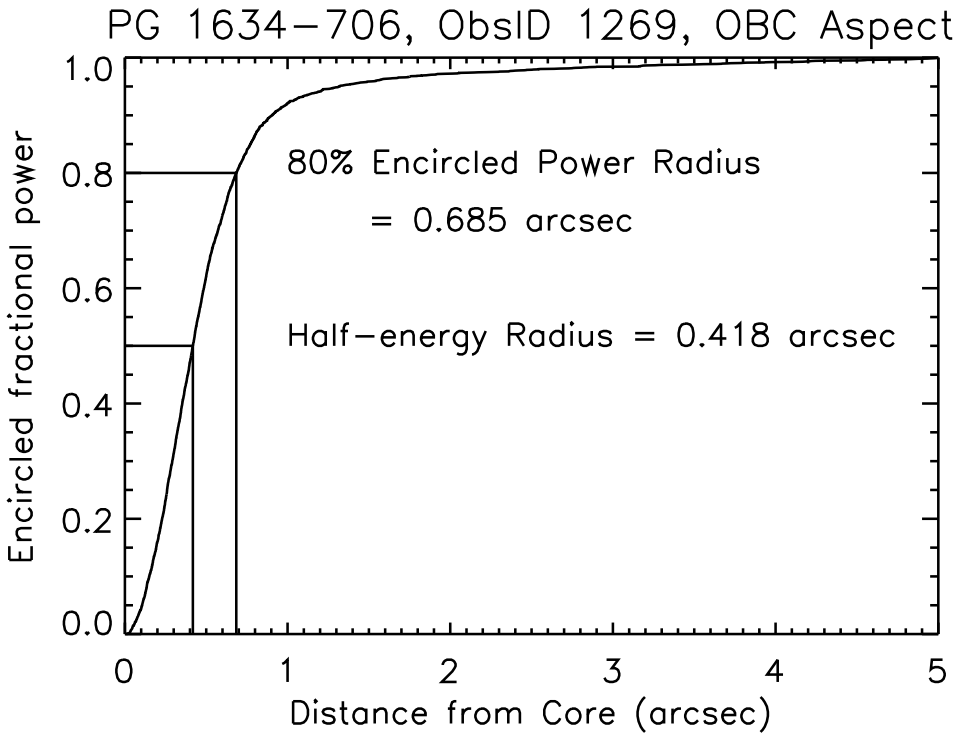

Figure 6.12: The on-orbit encircled broad-band energy versus radius for an ACIS observation of point source PG1634-706 (ObsID 1269). The curve is normalized to unity at a radius of 5 arcsec. The statistical uncertainty is estimated to be 3%, and the systematic uncertainty due to power beyond 5 arcsec is 2%. The effective energy is 1 keV.

6.6.1 PSF Anomaly

The *Chandra* PSF near the aimpoint displays an unexplained enhancement in the profile $\approx$ 0.8 arcsec from the source centroid (see Section 4.2.3). This anomaly is in excess of that expected from ray trace simulations, and is preferentially oriented towards the mirror spherical coordinate (MSC) angle of $\phi = 285°$ (see the *CIAO* caveats page https://cxc.harvard.edu/ciao/caveats/psf_artifact.html). This is approximately oriented towards the spacecraft +Z axis (see Figure 1.2). The asymmetry is illustrated for an ACIS observation of a low-mass X-ray binary (LMXB) in NGC 6397 in Figure 4.18, and the magnitude of the asymmetry is illustrated for a number of low-count-rate on-axis point sources in Figure 4.19 (see Section 4.2.3). Figure 6.13 depicts the effect of the anomaly on a number of point sources in NGC 6397.

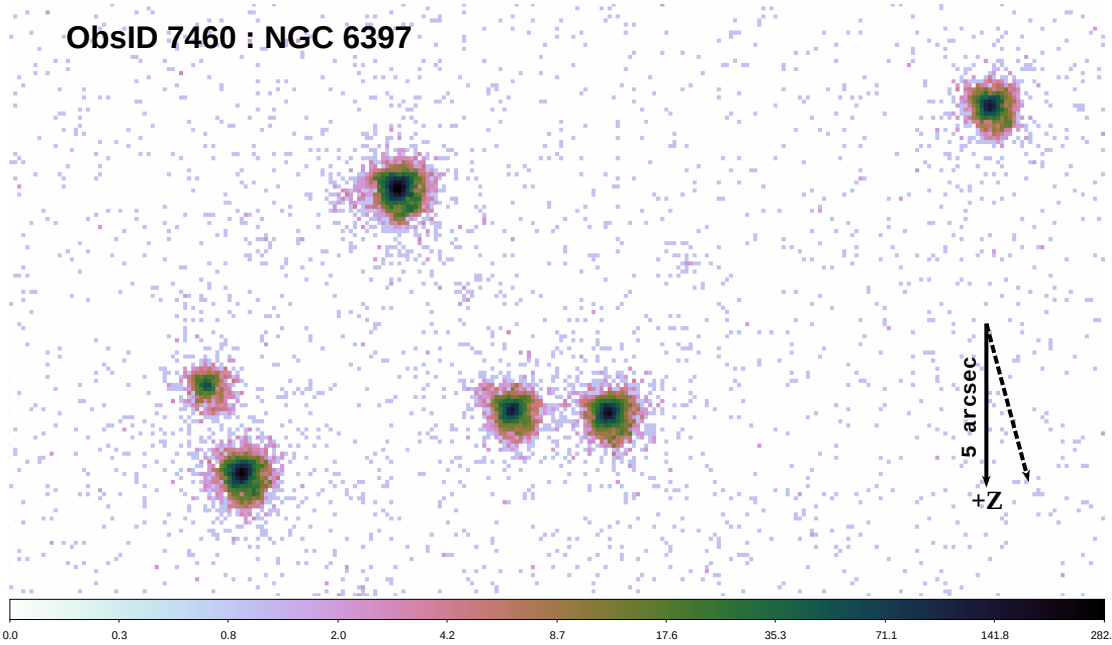

Figure 6.13: The PSF anomaly illustrated with an ACIS-S observation of NGC 6397 for a number of sources near the aimpoint. The image has been derolled such that the spacecraft +Z is pointed downwards (see solid vertical downward arrow of length 5 arcsec). The nominal direction of the anomaly is indicated by the dashed arrow (also of length 5 arcsec). The data have been binned to $1/4 \times 1/4$ ACIS sky pixels, and a logarithmic intensity stretch has been applied. Notice that all the sources in the field have discernible asymmetries in the indicated direction.

6.6.2 ACIS PSF Broadening at Low Energies

There is some evidence that the soft ACIS PSF (up to $\sim$ 800 eV) has been broadening since 2015. This is based on an analysis of the zeroth-order LETG/ACIS-S images of the isolated neutron star RXJ 1856. The calibration team has been observing RXJ 1856 annually since 2011 to monitor the build-up of contamination on the ACIS filters. The 50% encircled energy of the zeroth order image of the RXJ 1856 LETG/ACIS-S data at low energies has increased to about 4″ over the past 6 or 7 years. This effect is most likely due to the increase in the depth of the contaminant on the ACIS filters and is the most significant near and below the C K-edge (0.284 keV). Users are encouraged to consider using the HRC-I if one of their primary scientific objectives is to measure extended emission around a soft point source at scales of a few arcseconds. The HRC-I PSF has been stable throughout the mission and the HRC-I QE is greater than the ACIS QE below 1 keV.

6.7 Energy Resolution

The ACIS FI CCDs originally approached the theoretical limit for the energy resolution at almost all energies, while the BI CCDs exhibited poorer resolution. The pre-launch energy resolution as a function of energy is shown in Figure 6.14. Subsequent to launch and orbital activation, *the energy resolution of the FI CCDs has become a function of the row number*, being near pre-launch values close to the frame store region and substantially degraded in the farthest row. An illustration of the dependence on row is shown in Figure 6.15.

The loss of energy resolution is due to increased charge transfer inefficiency (CTI) caused by low energy protons encountered during radiation belt passages and Rutherford scattering through the X-ray telescope onto the focal plane. Subsequent to the discovery of the degradation, operational procedures were changed: ACIS is moved to a sheltered position during radiation belt passages. Since this procedure was initiated, no further degradation in performance has been encountered beyond that predicted from pre-launch models. The BI CCDs were not impacted, which is consistent with the proton-damage scenario—it is far more difficult for low-energy protons from the direction of the HRMA to deposit their energy in the buried channels of the BI devices, since the channels are near the gates and the gates face in the direction opposite to the HRMA. Thus the energy resolution for the two BI devices remains nearly at pre-launch values. The position-dependent energy resolution of the FI CCDs depends significantly on the ACIS focal plane (FP) temperature. Since orbital activation, the ACIS FP temperature has been lowered in steps and is now set at ~ -120 °C. As discussed in Section 6.20, the FP temperature now commonly warms above -120 °C owing to *Chandra*'s evolving thermal environment. If the optimal spectral performance is important for the scientific goals, the observer should consider operating as few CCDs as possible to minimize the FP temperature (see Section 6.22 for details on specifying the configuration).

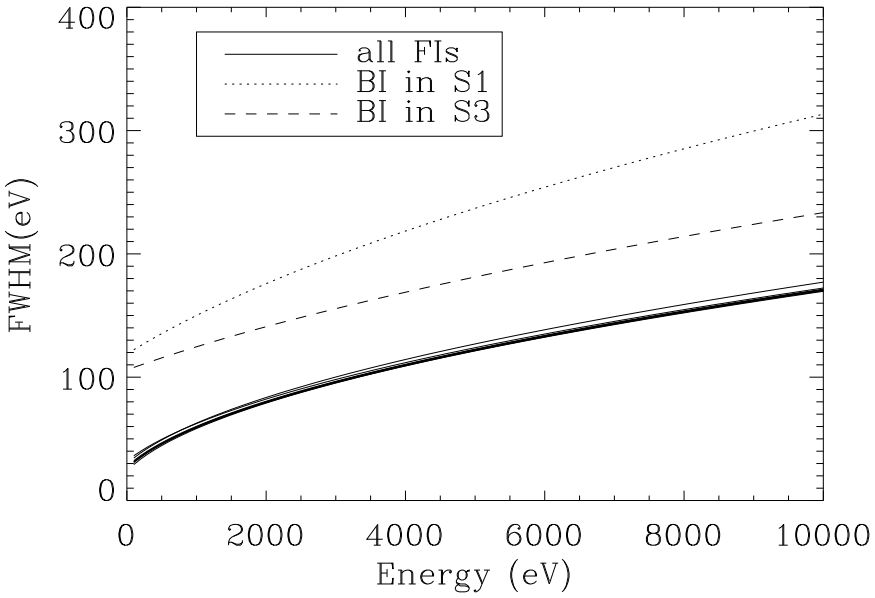

Figure 6.14: The ACIS pre-launch energy resolution as a function of energy.

6.7.1 Correcting the Energy Resolution of the CCDs

The ACIS instrument team has developed a correction algorithm that recovers much of the lost energy resolution the FI CCDs. The correction recovers a significant fraction of the CTI-induced loss of spectral resolution of the FI CCDs at all energies. The algorithm has been incorporated in the *CIAO* tool *acis_process_events* since *CIAO* 2.3. Figure 6.15 illustrates the improvement that the tool provides. As of 2006-Dec, data for the two BI chips can also be corrected in the same way, including a correction for serial CTI for the BI chips, though the effects are more subtle. The resulting response is very nearly uniform across the BI chips once this correction is made. As of 2010-Dec-15 (*CIAO* 4.3), temperature dependence was added to the CTI correction in *acis_process_events* for both FI and BI CCDs.

The ACIS energy resolution (here taken to mean the full width at half maximum [FWHM] of a narrow spectral line) varies roughly as the square root of the energy, and increases with distance from the read-out. On FI chips (the I array and all the S array chips except S1 and S3), the increase with CHIPY (row number) is dramatic, as can be seen in Figure 6.15. The spatial dependence on BI chips (S1 and S3) is much weaker, but depends on both coordinates. This can be illustrated by measuring the resolution at the aimpoints and default offsets in effect 2009 (see Table 6.4), using the Al-Kα (1.49 keV) and Mn-Kα (5.9 keV) lines in the external calibration source. Data were summed over three months in early 2012, and over an area of 128×128 pixels. Note that these values

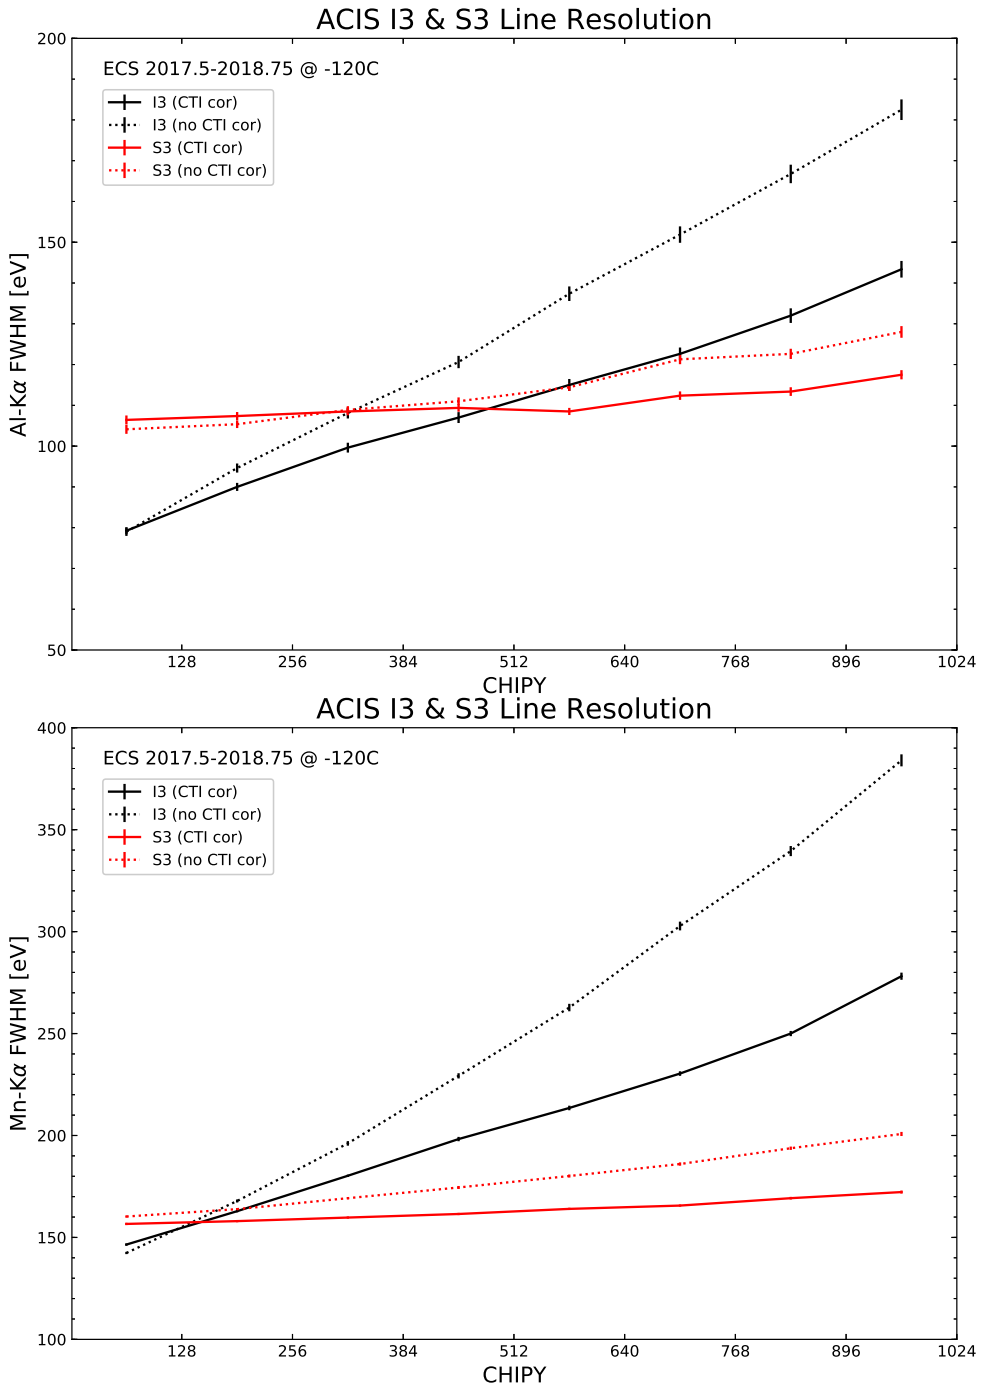

Figure 6.15: Energy resolution (FWHM in eV) versus row number (CHIPY) for several cases is plotted. The top plot shows Aluminum Kα (1.49 keV), while the bottom plot is for Manganese Kα (5.9 keV). The resolution of the I3 chip—with and without CTI correction—is shown by the black curves, and the S3 resolution is shown by the red curves.

are intended to be illustrative.

Table 6.4: Aimpoint Spectral Resolution: 2012-Apr to 2012-Jun

CCD aimpoint	CHIPX	CHIPY	Al-Kα (eV)	Mn-Kα (eV)
I3	0897 : 1024	0897 : 1024	145	267
S3	0129 : 0256	0513 : 0640	101	159

6.7.2 Effect of the Focal Plane Temperature on CCD Energy Resolution

As described below in Section 6.20, the ACIS focal plane temperature is sensitive to the solar pitch angle, as well as Earthshine into the ACIS radiator field of view (more important near perigee) and heating from ACIS electronics. Increased focal plane temperatures result in higher CTI on the ACIS CCDs, and hence poorer energy resolution. Figure 6.16 shows the FWHM of the Mn-Kα and Al-Kα lines as a function of row number (CHIPY) for several different focal plane temperatures. As in Figure 6.15, the energy resolution degrades with increasing CHIPY, but overall the energy resolution is worse at all CHIPY with increasing focal plane temperature. Updated temperature-dependent CTI correction products were released for all four ACIS-I chips and the S2 and S3 chips in CALDB 4.11 (December 2023). The updated CALDB products were calibrated using data taken with focal plane temperatures up to -105 °C. Applying the updated temperature-dependent CTI correction products required some software changes, which were released in CIAO 4.16 (December 2023).

At present there is a focal plane temperature limit which depends on the observation's aimpoint (ACIS-I array or S3). Mission Planning will not schedule an observation if the predicted temperatures exceed the temperature limit during the observation. With the release of the updated temperature-dependent CTI correction products, the maximum allowed focal plane temperature for a given observation could be relaxed. The temperature constraints were set to ensure that the statistical uncertainties in an observation would always be greater than the systematic calibration uncertainties. High S/N ACIS observations are now allowed up to a focal plane temperature of -111 °C. Very low S/N observations are now allowed up to a temperature of -105 °C.

6.8 Detector Gain

For CCD detectors, the gain is the relation between the energy of the incident photon and the magnitude of the collected charge (digitized to ADU). CTI both broadens the detector response (because charge can be trapped and released) and reduces the charge recorded (because trapped charge can be lost) for a recorded event. CTI depends on the CCD type (FI or BI). It also depends on row number, thus inducing a spatially-dependent gain effect. To account for the spatial variation of the response, the response width (`p2_resp`) and

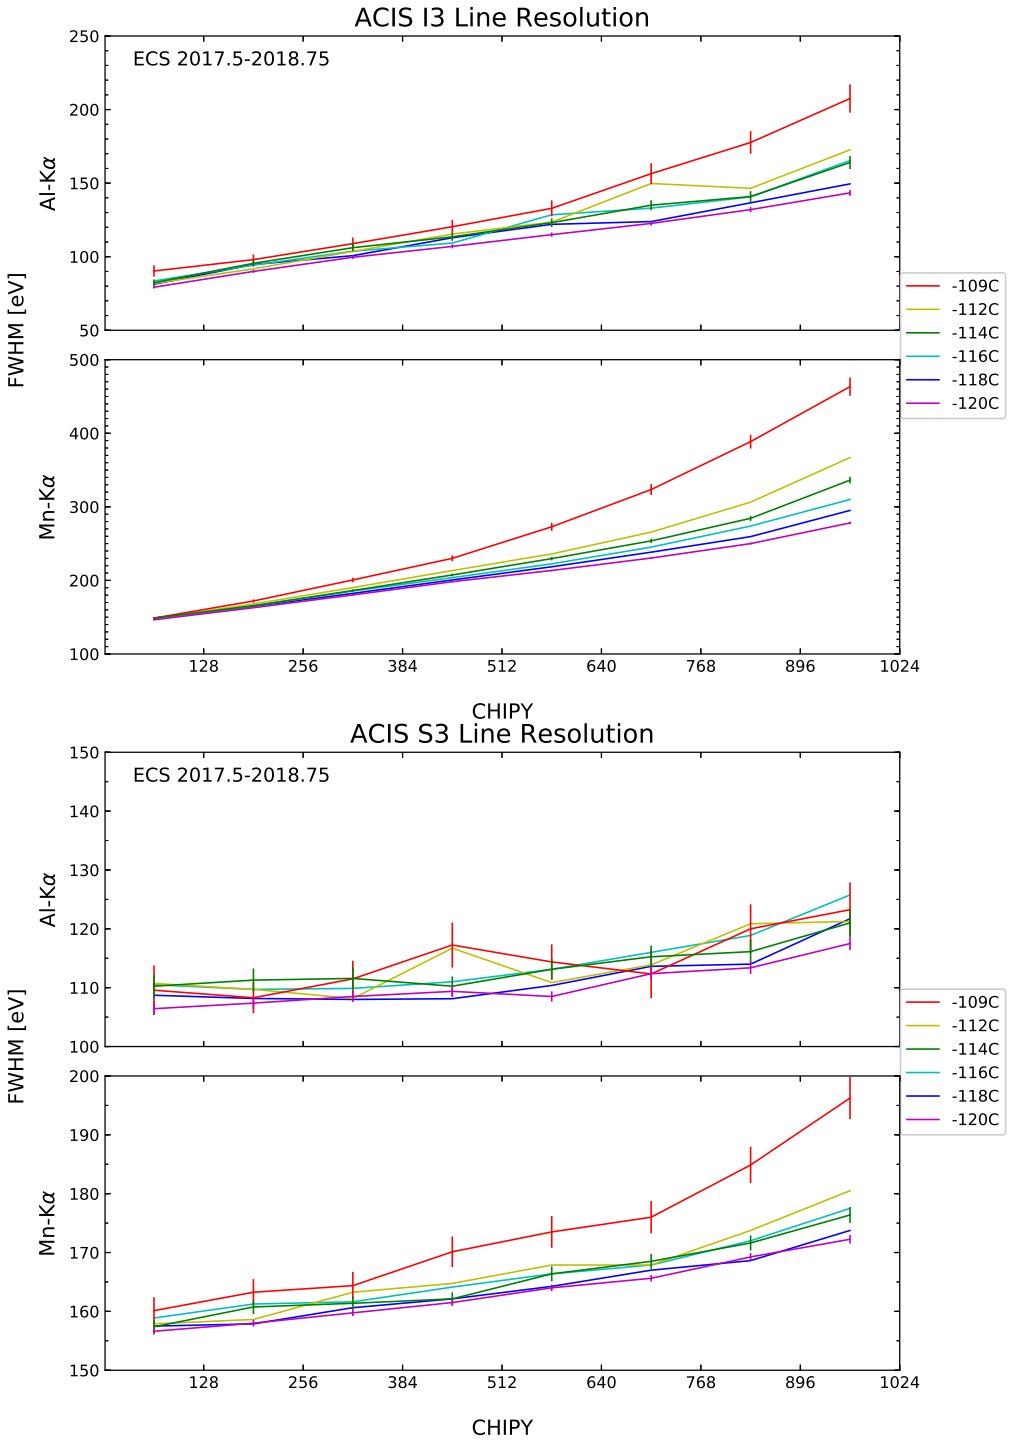

Figure 6.16: Energy resolution (FWHM in eV) versus row number (CHIPY) for Al-Kα and Mn-Kβ lines for both FI and BI CCDs for several different focal plane temperatures. The error bars are large at the warmest temperatures because of the relative lack of calibration data at these temperatures compared to lower temperatures.

gain (`detgain`) were initially calibrated for 256 × 32 [`CHIPX`,`CHIPY`] "tiles" for FI chips, and 32 × 32 tiles for BI chips.

The `detgain` "mid-chip gain droop" in the FI detectors (in which the conversion from `PHA` to energy produced energies which were too small in the 32×32 pixel (`CHIPX`, `CHIPY`) regions immediately adjacent to the node boundary at `CHIPX` 512, 513) has been corrected. Revised `detgain` and `p2_resp` calibration files were released as part of the CALDB release update 4.9.1. As of that release the `detgain` detector gains for FI chips were recalibrated on 32 × 32 tiles and the response widths were converted to 32 × 32 tiles. At higher focal plane temperatures, there is now a mid-chip gain excess, which has not yet been corrected.

The CTI correction (see Section 6.7.1) and gain (detgain) calibrations were developed based on an initial "epoch" of external calibration source (ECS) data taken between 2000-Feb and 2000-Apr. The CTI evolution is monitored with regular observations of the ECS and astrophysical sources. The CTI correction accounts for much of the gain shift, however residual changes remain that are highly dependent on the X-ray particle background. These residual changes are calibrated by collecting ECS and astrophysical source data in intervals and measuring the shift in the fitted lines relative to the baseline detgain to obtain the time-dependent gain correction ("`tgain`"). Fading of the ECS due to radioactive decay has led to longer collecting intervals (3 months, then 6 months, and currently 12 months), and the need to combine more pixels per correction region (first 32 × 32 pixel sized "tiles", now 32 × 128) due to the decreased number of spectral counts available. Time-dependent gain corrections are interpolated between `tgain` files to provide more accurate gain corrections based on the date of an observation.

Select astrophysical sources are now included to further refine the time-dependent gain corrections. The ACIS gain calibration program currently comprises:

1. The 5.9 keV Mn emission line in the ECS and the 6.6 keV Fe line from the Perseus cluster of galaxies to calibrate the high-energy gain.

2. The 1.85 keV Si emission line in the supernova remnant Cas A to calibrate the lower-energy gain.

3. Short observations of Cas A and Perseus in the middle of each chip are carried out on the six imaging chips (ACIS-I0, I1, I2, I3, S2, S3) each year.

4. A Principal Component Analysis (PCA) of all the time-dependent gain correction files (`tgain`) generated over the course of the mission is used to define the spatial structure in the `tgain` files.

5. A velocity map of Cas A is used to correct Si emission line shifts due to bulk motions. The *Hitomi* data for Perseus show that all of the gas in Perseus is at the redshift of the cluster.

6. The ECS is monitored in 3-month intervals, using coarse 256 × 256 pixel tiles, to inform and refine `tgain` calibration release dates. Periods of steady, or linear, gain

changes that can be accounted for by the interpolation algorithm can extend `tgain` update release intervals.

7. One `tgain` file is now created on an annual, or shorter, interval with a spatial binning of 32×128 [chipx, chipy] pixels using the ECS, Cas A, and Perseus observations along with the PCA results to map the gain corrections across each of the six imaging chips.

This gain calibration program results in a measured uncertainty of approximately 0.5–1.0% since 2023, and approximately $< 0.5\%$ prior. The outer ACIS-S CCDs—S0 and S5—are less rigorously corrected for time-dependent gain changes as they are primarily used with gratings, which rely on the dispersion relation to obtain photon energies. Additionally, two of the CCDs (I0 and I2) show instrumental noise affecting the gain ("anomalous gain") on shorter intervals than are able to be accurately captured by the periodic `tgain` updates, so the gain may be somewhat less accurate for these chips than for the other CCDs.

6.9 Hot Pixels and Columns

Hot pixels and columns are defined to be those which produce a spuriously high or saturated pulse-height for a large number of consecutive frames of data. These depend on operating conditions such as temperature. One should always refer to the *CXC* web site for the most recent list.

6.10 Cosmic Ray Afterglows

Cosmic ray hits sometimes deposit so much charge that they appear in the same pixels for a number of successive (or nearly successive) frames. These cosmic ray afterglows are essentially temporary hot pixels, and are removed by the hot pixel logic if they contain more than approximately 8 events.

6.11 Aimpoints

Aimpoints are the nominal positions on ACIS where the flux from a point source with no target offset or SIM-Z translation is placed. There are two aimpoints, indicated in Figure 6.1: one on the corner of I3 on the ACIS-I array (the ACIS-I aimpoint) and one near the boundary between nodes 0 and 1 on S3 of the ACIS-S array (the ACIS-S aimpoint). Because of variations in the thermal conditions in the aspect subsystem (see Section 5.4.3), the actual position on the detector of a source with no offsets or SIM-Z translations can vary somewhat from the default location. Note also spacecraft dither will move the position of the source on the detector in a Lissajous pattern centered on the source position on the detector (see Section 6.12).

As of Cycle 18, permanent default aimpoints were specified for each detector. Because of the changes to the default dither for ACIS in 2022, the default aimpoints were again shifted slightly. Note that the aimpoint is not the same as the optical axis (which is defined as the position of the narrowest PSF). The chip coordinates of the aimpoints and the current estimate for the optical axis position are listed in Table 4.3. The aimpoints are close to the optical axis, and have been adjusted to reduce the risk of dither moving a source at the default aimpoint across a chip edge (ACIS-I3) or across a node boundary (ACIS-S3) with the expected aiming uncertainties. Figures 4.27 and 4.28 show the default aimpoints, aimpoint error box, and aimpoint error box plus dither. Further information on the absolute pointing and aimpoint stability and uncertainties can be found in Sections 5.4.3 and 4.5.

Approximate contours of constant encircled energy for ACIS-I and ACIS-S observations at the aimpoints are shown in Figures 6.17 and 6.18.

When ACIS is used with one of the transmission gratings, HETG, or LETG, a SIM translation may be applied, and offsets may be applied to the default aimpoints used for non-grating observations. For further details, see Chapter 8 (HETG), or Chapter 9 (LETG).

6.12 Dither

Unless specially requested, the spacecraft is dithered in a Lissajous pattern during all observations. For observations with ACIS, the default dither pattern now spans 32 arcsec peak to peak in both the Y and Z directions. Note that for subarrays 134 pixels or smaller, the dither spans 32 arcsec peak to peak in the long direction, but only 16 arcsec peak to peak in the short direction. The dither serves two purposes: (1) to provide some exposure in the gaps between the CCDs, and (2) to smooth out pixel-to-pixel variations in the response. The effect of dither is removed during high-level ground processing of the data. The exposure time in the gaps between chips (and at the outside edges) will be less than that for the remainder of the field. Default dither parameters are listed in Table 5.4.

6.12.1 Gaps Between the CCDs, and Chip Tilts

The approximate sizes of the various gaps between chips are shown in Figure 6.1. Note that due to the way the ACIS-I array CCDs are tilted to follow approximately the HRMA focal plane curvature, the Y-gaps vary slightly with Z and the Z-gaps vary slightly with Y. The ACIS-S chips are tilted and raised slightly in order to follow approximately the grating Rowland Circle. The array is still nearly flat, with the edge of S0 raised 0.74 mm and the edge of S5 raised 0.53 mm; the length of the S array is $\sim$ 150 mm. The S-array chip gaps are nearly constant with Z.

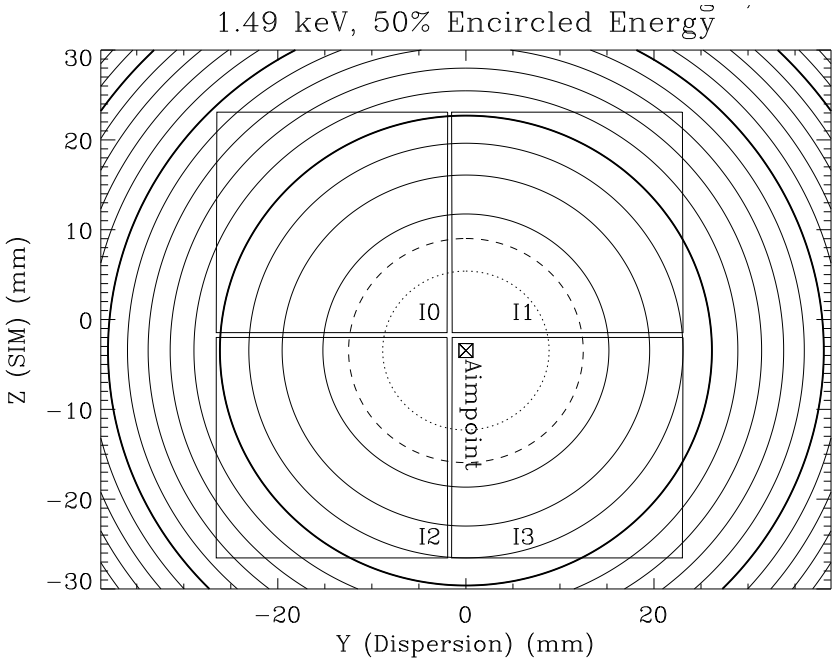

Figure 6.17: Approximate contours of constant 50% encircled energy at 1.49 keV when the ACIS-I aimpoint is selected. The dotted circle is 1 arcsec. That is, the 50% encircled energy at 1.49 keV at any point on the dotted circle is 1 arcsec. Similarly, the dashed circle is 1.5 arcsec. The remainder are at 1 arcsec intervals. The thicker solid circles highlight the 5, 10, and 15 arcsec contours.

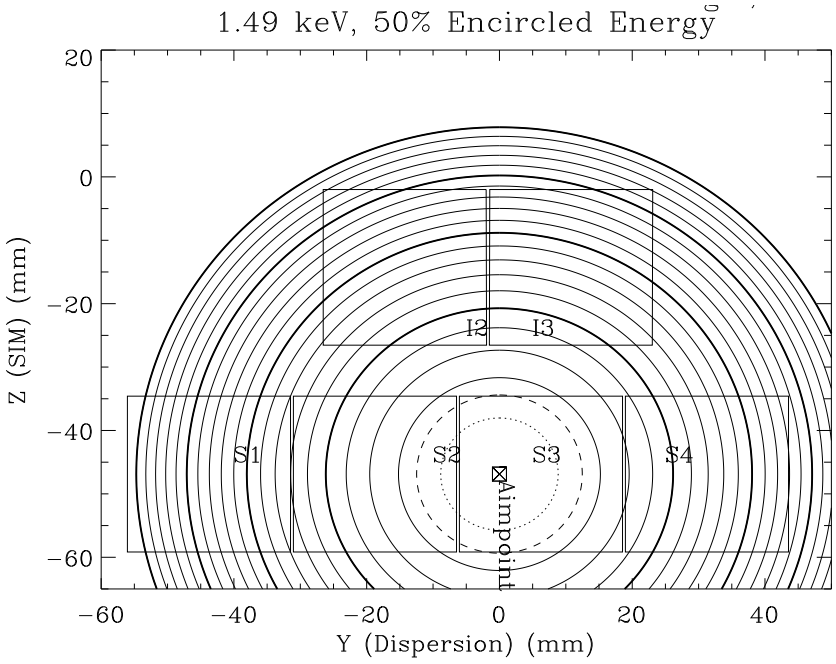

Figure 6.18: Approximate contours of constant 50% encircled energy at 1.49 keV when the ACIS-S aimpoint is selected. The dotted circle is 1 arcsec. That is, the 50% encircled energy at 1.49 keV at any point on the dotted circle is 1 arcsec. Similarly, the dashed circle is 1.5 arcsec. The remainder are at 1 arcsec intervals. The thicker solid circles highlight the 5, 10, 15, and 20 arcsec contours.

6.13 Operating Modes

ACIS has two operating modes: the Timed Exposure (TE) Mode, and the Continuous Clocking (CC) Mode. One must select one or the other for an observation as it is not possible to simultaneously operate individual CCDs in different modes during a single observation.

6.13.1 Timed Exposure (TE) Mode

Timed exposure refers to the mode of operation wherein a CCD collects data (integrates) for a preselected amount of time—the Frame Time. Once this time interval has passed, the charge from the 1024×1024 pixel imaging region is quickly ($\sim$ 41 msec) transferred to the frame store region and subsequently read out through four output amplifiers.

Frame Times—Full Frames Frame times are selectable within a range of values spanning the time interval from 0.2 to 10.0 sec. If the data from the entire CCD are utilized (full frame) then *the nominal (and optimal!)* frame time, T_{opt}, is 3.2 s. Selecting a frame time *shorter* than the nominal value (e.g. to decrease the probability of pile-up – Section 6.16) has the consequence that there will be a time during which *no* data are taken, "deadtime," as 3.2 s are required for the frame store read-out process regardless of the frame time. The fraction of time during which data are taken is simply the ratio of the selected frame time to the sum of the nominal frame time and a flush frame—*e.g.* for a new frame time of $t < T_{opt}$, the fraction of time during which data are taken is $\sim t/(T_{opt} + t)$. Note that this formula is approximate: it ignores the 41 msec image-to-framestore transfer time (which affects both t and T_{opt}) and the 18.84 msec needed to flush the initial rows to the output amplifiers and to flush the last row from the output amplifiers. Also, strictly speaking, the full-frame time depends on how many CCDs are on—see the equation below—but the differences are very small. Finally, note that selecting a frame time *longer* than the most efficient value does not reduce efficiency (there is no additional flush frame) but increases the probability of pile-up occurring and is not recommended. In addition, with the standard ACIS dither, a frame time longer than 3.2 s will add a blur to the PSF.

Frame Times & Subarrays It is also possible for one to select a *subarray*—a restricted region of the CCD in which data will be taken. A subarray is fully determined by specifying the number of rows separating the subarray from the frame store region (q) and the number of rows in the subarray (n). Examples of subarrays are shown in Figure 6.19. The nominal frame time for a subarray depends on q, n, and the total number of CCDs that are activated (m)–see Table 6.5. The nominal frame time is given by:

$$T(\text{msec}) = 41.12 \times \text{m} + 0.040 \times (\text{m} \times \text{q}) + 2.85 \times \text{n} - 32.99. \tag{6.1}$$

As with full frames, selecting a frame time less than the most efficient value results in loss of observing efficiency. Frame times are rounded up to the nearest 0.1 sec, and can range from 0.2 to 10.0 sec.

Table 6.5: CCD Frame Time (sec) for Standard Subarrays

Subarray	ACIS-I (no. of chips)		ACIS-S (no. of chips)		
	1	6	1[1]	2	6
1	3.0	3.2	3.0	3.0	3.2
1/2	1.5	1.8	1.5	1.6	1.8
1/4	0.8	1.2	0.8	0.9	1.1
1/8	0.5	0.8	0.4	0.5	0.7

[1]The minimum configurable selection for an ACIS-S3 observation is 2 chips, owing to the requirement to include at least one FI chip (Section 6.20)

When operating with only one chip, subarrays as small as 100 rows are allowed (this permits 0.3 sec frame times which pay no penalty in dead time if $q \leq 200$. Use the above equation to determine what the most efficient frame time is for the desired ACIS configuration. For multichip observations, the smallest allowed number of rows is 128. For small subarrays, the aiming uncertainties plus dither should be taken into account; see Section 6.11 and references therein. Furthermore, S3-only observations are now required to include an FI chip, as detailed in Section 6.20.

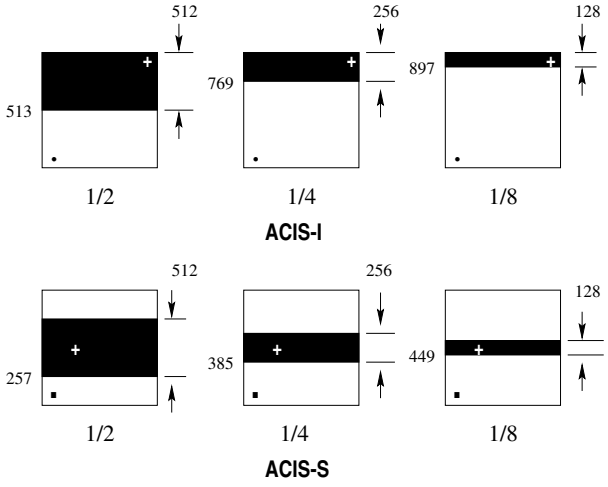

Figure 6.19: Examples of various subarrays. The heavy dot in the lower left indicates the origin.

Trailed Images It takes 40 μsec to transfer the charge from one row to another during the process of moving the charge from the active region to the frame store region. This

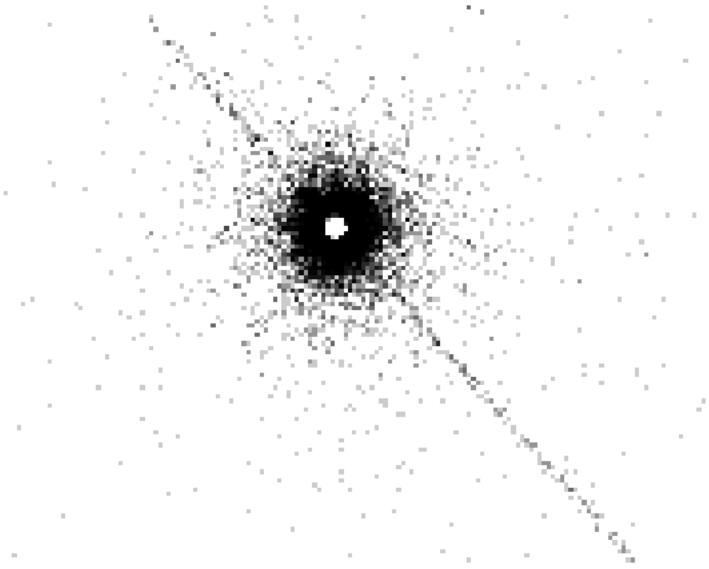

Figure 6.20: Trailed image of a strong X-ray source. The core of the image is faint due to pile-up. Most events here are rejected because of bad grades. The read-out direction is parallel to the trail.

has the interesting consequence that each CCD pixel is exposed not only to the region of the sky at which the observatory was pointing during the long (frame time) integration, but also, for 40 μsec each, to every other region in the sky along the column in which the pixel in question resides. Figure 6.20 is an example where there are bright features present so intense that the core of the PSF is suppressed because of pile-up (see Section 6.16), allowing the tiny contribution of the flux due to trailing to be stronger than the (piled-up) core of the direct exposure—hence the trailed image is clearly visible. Trailed images are also referred to as "read-out artifacts," "transfer smear," or "out-of-time images." The user needs to be aware of this phenomenon as it has implications for the data analysis, including estimates of the background. In some cases, the trailed image can be used to obtain an unpiled spectrum and can also be used to perform 40 microsec timing analysis of (extremely bright) sources.

6.13.2 Alternating Exposures

In some instances, it is desirable to have both long and short frame times. If the exposure time is made very short, pile-up may be reduced, but the efficiency of the observation is greatly reduced by the need to wait for the full 3.2 sec (if six chips are clocked) for the frame-store array processing.

In alternating exposure mode, all CCDs are clocked in unison, but have two exposure times. One (typically short) primary exposure of length $0.2 < t_p < T_{opt}$ sec is followed by

k secondary exposures of length t_s (typically the standard optimal time T_{opt}, for instance, 3.2 sec if six chips are clocked with full frames). Permissible values of k range from 1 to 15. The short exposures are used to reduce photon pile-up, and the long exposures are useful for studying the fainter objects in the field of view. For example, a typical choice of long and short exposure times might be 3.2 and 0.3 sec. If $k = 3$, ACIS would perform a flush frame followed by one 0.3 sec exposure, and then three exposures of 3.2 sec, repeating until the total observing time expires.

If the duty cycle of long exposures is $1 : k$ (primary:secondary), the observing efficiency η is then

$$\eta = \frac{t_p + kt_s}{t_p + (k+1)t_s} \tag{6.2}$$

6.13.3 Continuous Clocking (CC) Mode

The continuous clocking mode is provided to allow 3 msec timing at the expense of one dimension of spatial resolution. In this mode, one obtains 1 pixel $\times$ 1024 pixel images, each with an integration time of 2.85 msec. Details as to the spatial distribution in the columns are lost—other than that the event originated in the sky along the line determined by the length of the column.

In the continuous clocking mode, data are continuously clocked through the CCD and frame store. The instrument software accumulates data into a buffer until a virtual detector of size 1024 columns by 512 rows is filled. The event finding algorithm is applied to the data in this virtual detector and 3×3 event islands are located and recorded to telemetry in the usual manner (Section 6.15.1). This procedure has the advantage that the event islands are functionally equivalent to data accumulated in TE mode, hence differences in the calibration are minimal. The row coordinate, called CHIPY in the Flexible Image Transport System (FITS) file, maps into time in that a new row is read from the frame store to the buffer every 2.85 msec. This does have some minor impacts on the data. For example, since the event-finding algorithm is looking for a local maximum centered in a 3×3 pixel area, it cannot find events on the edges of the virtual detector. Hence, CHIPX cannot be 1 or 1024, and CHIPY cannot be 1 or 512. In other words, events cannot occur at certain times separated by 512×2.85 msec or 1.4592 sec. Likewise, it is impossible for two events to occur in the same column in adjacent time bins.

Event files for continuous-clocking mode data contain two time columns, TIME, and TIME_RO. The TIME_RO records the read-out times (effectively the time the virtual frame is processed to identify events). The TIME column is an estimate of the photon time of arrival at the detector, based on the read-out time, the aspect solution, and the specified sky location of the source. See the `acis_process_events` documentation for information on the generation of these columns. In order to ensure that the TIMEs are as accurate as possible, it is best to specify the source location to better than 0.5 arcsec. The read-out time can differ from the arrival time by about 2.9 to 5.8 s, depending on the nominal location of the source on the CCD and the dither of the spacecraft.

6.14 Bias Maps

In general, the CCD bias—the amplitude of the charge in each pixel in the absence of external radiation—is determined at every change of instrumental parameters or setup when ACIS is in place at the focus of the telescope. These bias maps have proven to be remarkably stable and are automatically applied in routine data processing.

The bias maps for continuous-clocking mode observations can be corrupted by cosmic rays. If a cosmic ray deposits a lot of charge in most of the pixels in one or more adjacent columns, the bias values assigned to these columns will be too large. As a result, some low-energy events that would have been telemetered will not be telemetered because they do not satisfy the minimum pulse height criterion and the spectrum of a source in the affected columns will be skewed to lower energies. The BI CCDs are relatively insensitive to the problem. A bias algorithm was implemented in Cycle 6 to mitigate the problem.

Occasionally a cosmic ray produces an artifact in a bias map. The pipelines search for these artifacts, and, when found, replace the bias map with another from the same epoch. Work is in progress to use long-term average bias maps, either when there are artifacts in the observation-specific bias map, or for all observations.

6.15 Event Grades and Telemetry Formats

6.15.1 Event Grades

In the first step in detecting X-ray events, the on-board processing examines every pixel in the full bias-subtracted CCD image (even in the continuous clocking mode (Section 6.13.3)) and selects as events those with values that both exceed the event threshold and are greater than all of their touching or neighboring pixels (i.e., a local maximum in a 3×3 pixel detection island). The pixel pattern, and thus the event grade, is determined by which of the outer pixels in a 3×3 grid centered on the initial pixel have values above the split-event threshold. Depending on the grade, the data are then included in the telemetry. On-board suppression of certain grades is used to limit the telemetry bandwidth devoted to background events (see Section 6.17.1).

The grade of an event is thus a code that identifies which pixels within the 3×3 pixel island, centered on the local charge maximum, are above certain amplitude thresholds. The thresholds are listed in Table 6.2. Note that the local maximum threshold differs for the FI and the BI CCDs. A "Rosetta Stone" to help one understand the ACIS grade assignments is shown in Figure 6.21, and the relationship to the Advanced Satellite for Cosmology and Astrophysics (*ASCA*) grading scheme is given in Table 6.6.

It is important to understand that most, if not all, calibrations of ACIS are based on a specific subset of ACIS grades, called g02346. This *standard* set comprises *ASCA* grades 0,2,3,4, and 6. In the absence of pile-up, this particular grade selection appears to optimize the signal-to-background ratio, but this conclusion depends on the detailed spectral properties of the source. Further, most of the scientifically important characteristics of

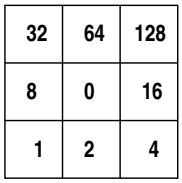

Figure 6.21: Schematic for determining the grade of an event. The grade is determined by summing the numbers for those pixels that are above their thresholds. For example, an event that caused all pixels to exceed their threshold is grade 255. A single pixel event is grade 0.

Table 6.6: ACIS flight grades and ASCA Grades

	ASCA grade	ASCA grade	
ACIS grades	CC mode graded	Other modes	Description
0	0	0	Single pixel
64 65 68 69	2	2	Vertical split up
2 34 130 162	2	2	Vertical split down
8 12 136 140	3	3	Horizontal split left
16 17 48 49	4	4	Horizontal split right
72 76 104 108	6	6	“L” & square, upper left
10 11 138 139	6	6	“L” & square, lower left
18 22 50 54	6	6	“L” & square, lower right
80 81 208 209	6	6	“L” & square, upper right
1 4 32 128	1	1	2-pixel diagonal split
5 33 132 160	1	1	3-pixel diagonal split
36 129	1	1	3-pixel diagonal split
37 133 161 164	1	1	4-pixel diagonal split
165	1	1	5-pixel diagonal split
3 6 9 20	5	5	3-pixel “L” with corner
40 96 144 192	5	5	3-pixel “L” with corner
13 21 35 38	5	5	4-pixel “L” with corners
44 52 97 100	5	5	4-pixel “L” with corners
131 134 137 145	5	5	4-pixel “L” with corners
168 176 193 196	5	5	4-pixel “L” with corners
53 101 141 163	5	5	5-pixel “L” with corners
166 172 177 197	5	5	5-pixel “L” with corners
24	7	7	3-pixel horizontal split
66	2	7	3-pixel vertical split
255	7	7	All 9 pixels
All other grades	7	7	

Table 6.7: Telemetry Saturation Limits

Mode	Format	Bits/event	Events/sec*	Number of Events in full buffer
CC	Graded	58	375.0	128,000
CC	Faint	128	170.2	58,099
TE	Graded	58	375.0	128,000
TE	Faint	128	170.2	58,099
TE	Very Faint	320	68.8	23,273

*(includes a 10% overhead for housekeeping data)

ACIS (effective area, sensitivity, point spread function, energy resolution, etc.) are grade- and energy-dependent.

6.15.2 Telemetry Formats

There are a number of telemetry formats available. Specifying a format determines the type of information that is included in the telemetry stream. The number of bits per event depends on which mode and which format is selected. The number of bits per event, in turn, determines the event rate at which the telemetry will saturate and data will be lost until the on-board buffer empties. The formats available depend on which mode (Timed Exposure or Continuous Clocking) is used. The modes, associated formats, and approximate event rates at which the telemetry saturates and one begins to limit the return of data, are listed in Table 6.7. The formats are described in the following paragraphs. Event "arrival time" is given relative to the beginning of the exposure.

Faint Faint format provides the event position in detector coordinates, an arrival time, an event amplitude, and the amplitude of the signal in each pixel in the 3×3 event island that determines the event grade. The bias map is telemetered separately. Note that certain grades may be not be included in the data stream (Section 6.17.1).

Graded Graded format provides event position in detector coordinates, an event amplitude, the arrival time, and the event grade. Note that certain grades may be not be included in the data stream (Section 6.17.1).

Very Faint Very Faint (VF) format provides the event position in detector coordinates, the event amplitude, an arrival time, and the pixel values in a 5×5 island. As noted in Table 6.7, this format is only available with the Timed Exposure mode. Events are still graded by the contents of the central 3×3 island. Note that certain grades may be not be included in the data stream (Section 6.17.1). This format offers the advantage of reduced

background after ground processing (see Section 6.17.2) for applicable observations with count rates sufficiently low as to avoid both telemetry saturation and event pile-up.

Studies (see https://cxc.harvard.edu/cal/Acis/Cal_prods/vfbkgrnd) of the ACIS background have shown that for weak or extended sources, a significant reduction of background at low and high energies may be made by using the information from 5×5 pixel islands, i.e. very faint mode, instead of the faint mode 3×3 island. This screening results in a 1–2% loss of good events after ground processing. *CIAO* 2.2 and later provides a tool to utilize the VF mode for screening background events. Please note that the Redistribution Matrix File (RMF) generation is the same for very faint mode as it is for faint mode. See the "Why Topic" https://cxc.harvard.edu/ciao/why/aciscleanvf.html. The very faint mode screening also reduces the nonuniformity of the non-X-ray background (see Section 6.17.1).

It is important to realize that VF mode uses more telemetry; the limit is ∼68.8 events/s (all telemetered grades), which includes the target flux and the full background from all chips. Check the calibration web page (https://cxc.harvard.edu/cal/Acis/Cal_prods/bkgrnd/current/background.html) for a discussion of background flares and the telemetry limit. In particular, review Section 1.3 of the memo "General discussion of the quiescent and flare components of the ACIS background" by M. Markevitch.

Starting with Cycle 11, the default upper energy cutoff has been decreased from 15 keV to 13 keV. To reduce the total background rate and the likelihood of telemetry saturation, VF mode observers should consider using no more than 4 CCDs and an energy filter with a 12 keV upper cutoff. If the target is brighter than 5-10 cts/s, either turning off more chips or using Faint mode may be necessary, depending on the total telemetered background included for the operating chips (see Table 6.10). See Section 6.22.1 for more information on selection of required and optional CCDs.

6.16 Pile-Up

Pile-Up results when two or more photons are detected as a single event. The fundamental impacts of pile-up are: (1) a distortion of the energy spectrum—the apparent energy is approximately the sum of two (or more) energies; and (2) an underestimate as to the correct counting rate—two or more events are counted as one. A simple illustration of the effects of pile-up is given in Figure 6.22. There are other, somewhat more subtle, impacts discussed below (Section 6.16.1).

The degree to which a source will be piled can be roughly estimated using *PIMMS*. Somewhat more quantitative estimates can be obtained using the pile-up models in *XSPEC*, *Sherpa*, and the Interactive Spectral Interpretation System (*ISIS*). If the resulting degree of pile-up appears to be unacceptable given the objectives, then the proposer should employ some form of pile-up mitigation (Section 6.16.3) as part of the observing strategy. In general, pile-up should not be a problem in the observation of extended objects, the Crab Nebula being a notable exception, unless the source has bright knots or filaments.

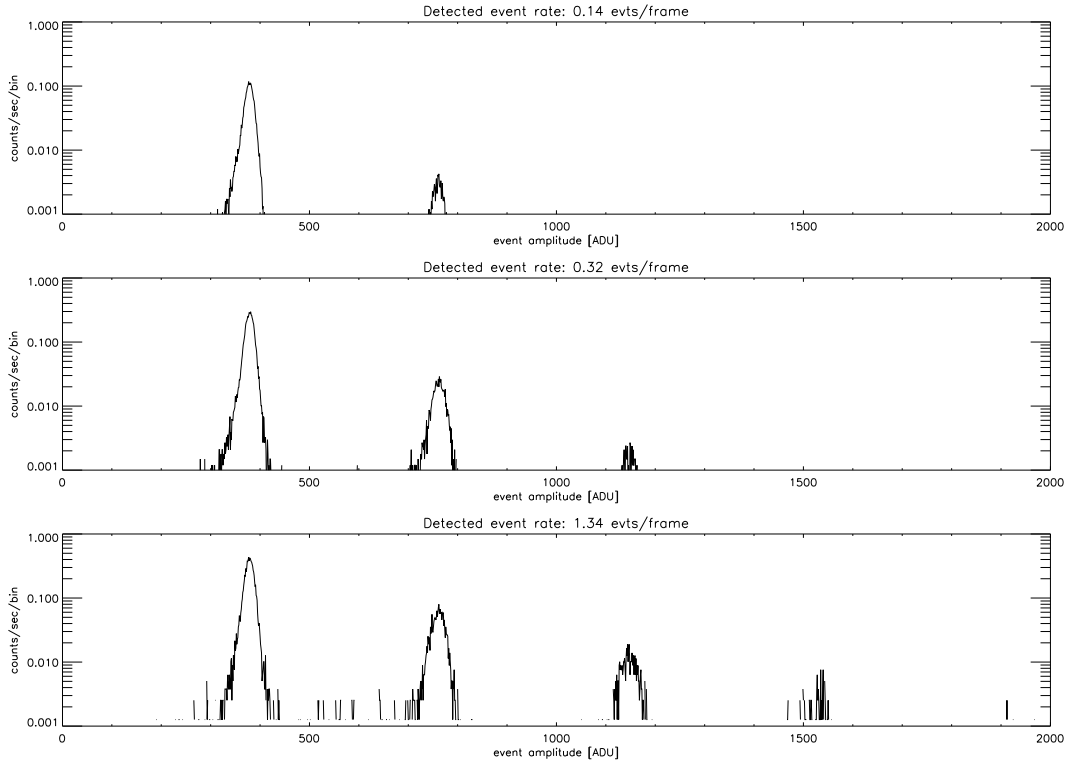

Figure 6.22: The effects of pile-up at 1.49 keV (Al Kα) as a function of source intensity. Data were taken during HRMA-ACIS system level calibration at the XRCF. Single-photon events are concentrated near the pulse height corresponding to the Al Kα line ($\sim$ 380 ADU), and events with 2 or more photons appear at integral multiples of the line energy.

6.16.1 Other Consequences of Pile-Up

There are other consequences of pile-up in addition to the two principal features of spurious spectral hardening and underestimating the true counting rate by under counting multiple events. These additional effects are grade migration and pulse saturation, both of which can cause distortion of the apparent PSF.

Grade migration Possibly the most troubling effect of pile-up is that the grade distribution for X-ray events may change. The change of grade introduced by pile-up has come to be referred to as "grade migration." Table 6.8 shows an example of grade migration due to pile-up as the incident flux is increased. In this simple test, which involved only mono-energetic photons, the largest effect is the depletion of G0 events and the increase of G7 events. In general, as the incident flux rate increases, the fraction of the total number of events occupying a particular event grade changes as photon-induced charge clouds merge and the resulting detected events "migrate" to other grades which are not at all

Table 6.8: ASCA Grade Distributions for different incident fluxes at 1.49 keV (Al Kα, based on data taken at the XRCF during ground calibration using chip I3)

Incident Flux*	G0	G1	G2	G3	G4	G5	G6	G7
9	0.710	0.022	0.122	0.053	0.026	0.009	0.024	0.035
30	0.581	0.057	0.132	0.045	0.043	0.039	0.029	0.073
98	0.416	0.097	0.127	0.052	0.050	0.085	0.064	0.108
184	0.333	0.091	0.105	0.040	0.032	0.099	0.077	0.224

*arbitrary units

necessarily included in the standard (G02346) set. If one applies the standard calibration to such data, the true flux will be under-estimated.

Pulse Saturation One consequence of severe instances of pile-up is the creation of a region with no events! In this case, the pile-up is severe enough that the total amplitude of the event is larger than the on-board threshold (typically 13 keV) and is rejected. Holes in the image can also be created by grade migration of events into ACIS grades (e.g. 255) that are filtered on-board.

PSF distortion Obviously the effects of pile-up are most severe when the flux is highly concentrated on the detector. Thus, the core of the PSF suffers more from pile-up-induced effects than the wings. Figure 6.23 illustrates this point. Because the core is suppressed, the PSF profile appears less peaked and (apparently) broader than would be the case if pile-up were negligible.

6.16.2 Pile-Up Estimation

It is clearly important in preparing a *Chandra* observing proposal to determine if the observation will be impacted by pile-up, and if so, to decide what to do about it (or convince the peer review why the specific objective can be accomplished without doing anything). There are two approaches to estimating the impact of pile-up on the investigation. The most sophisticated uses the pile-up models in *XSPEC*, *Sherpa*, and *ISIS* to create a simulated data set which can be analyzed in the same way as real data. A less sophisticated, but very useful, approach is to use the web version of *PIMMS* to estimate pile-up, or to use Figures 6.23 and 6.24.

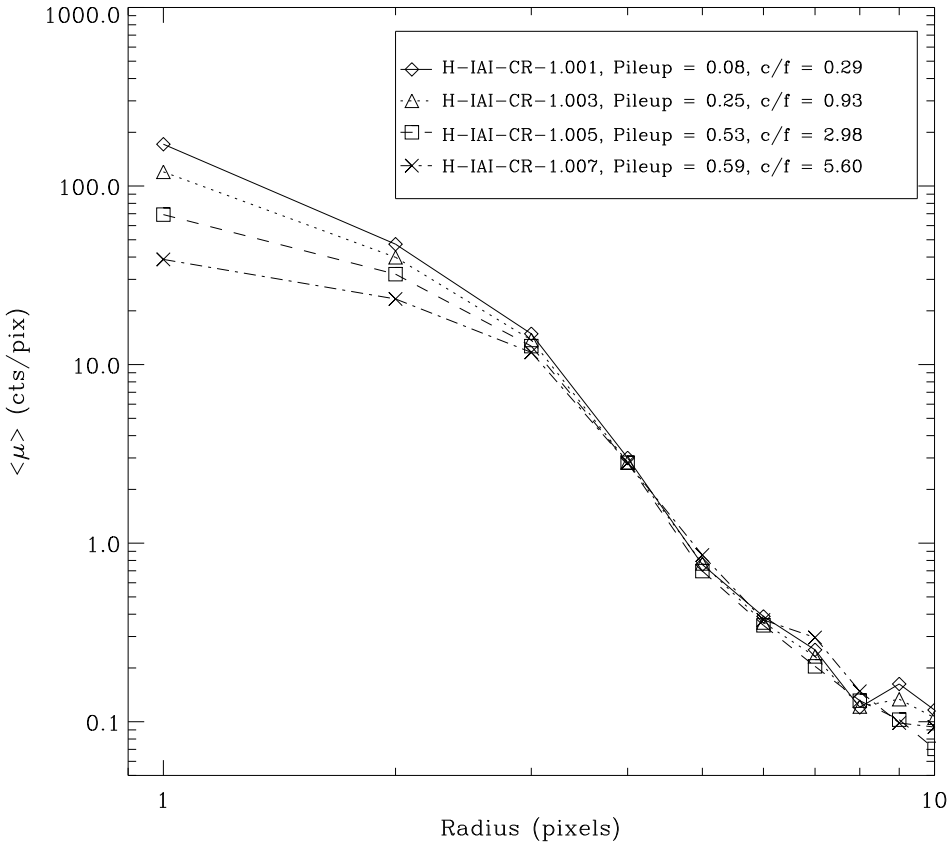

Figure 6.23: The effects of pile-up on the radial distribution of the PSF are illustrated. These data were taken during ground calibration at the XRCF. The specific "OBSIDs," the counting rate per CCD frame ("c/f"), and the "pile-up fraction" as defined in Section 6.16.2 are given in the inset.

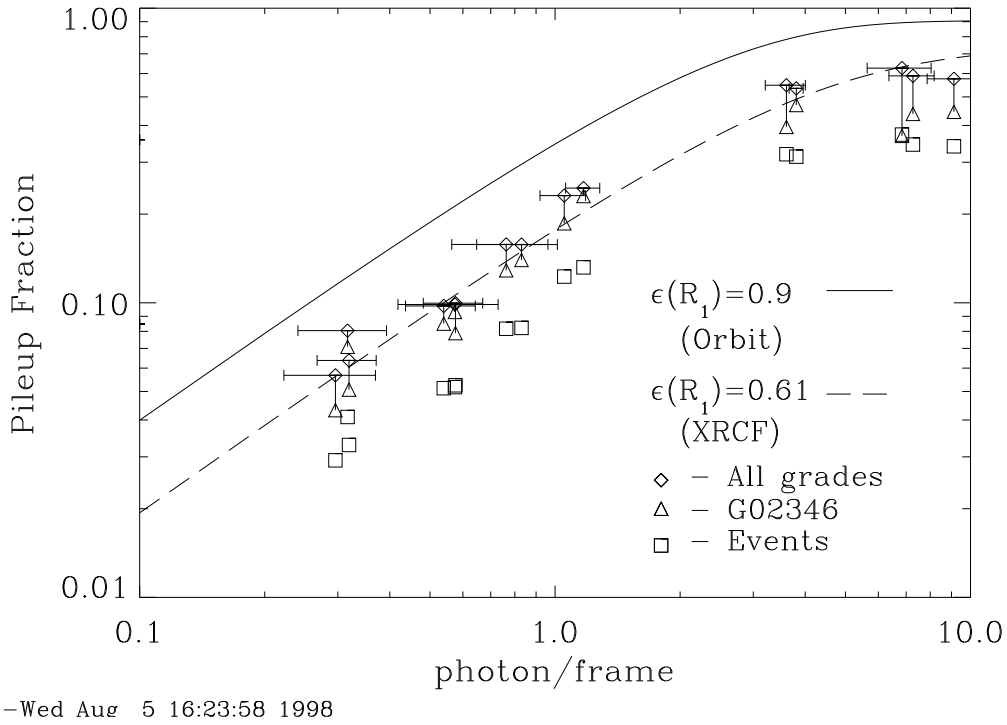

Figure 6.24: The pile-up fraction as a function of the the counting rate (in the absence of pile-up in units of photons/frame). The solid line is for on-orbit, the dashed line and the data points are for, and from, ground-based data respectively. The difference between the ground and flight functions are a consequence of the improved PSF on-orbit, where gravitational effects are negligible—see Chapter 4. Note that when pile-up occurs there are two or more photons for each event, so the fraction of events with pile-up is always less than the fraction of photons with pile-up.

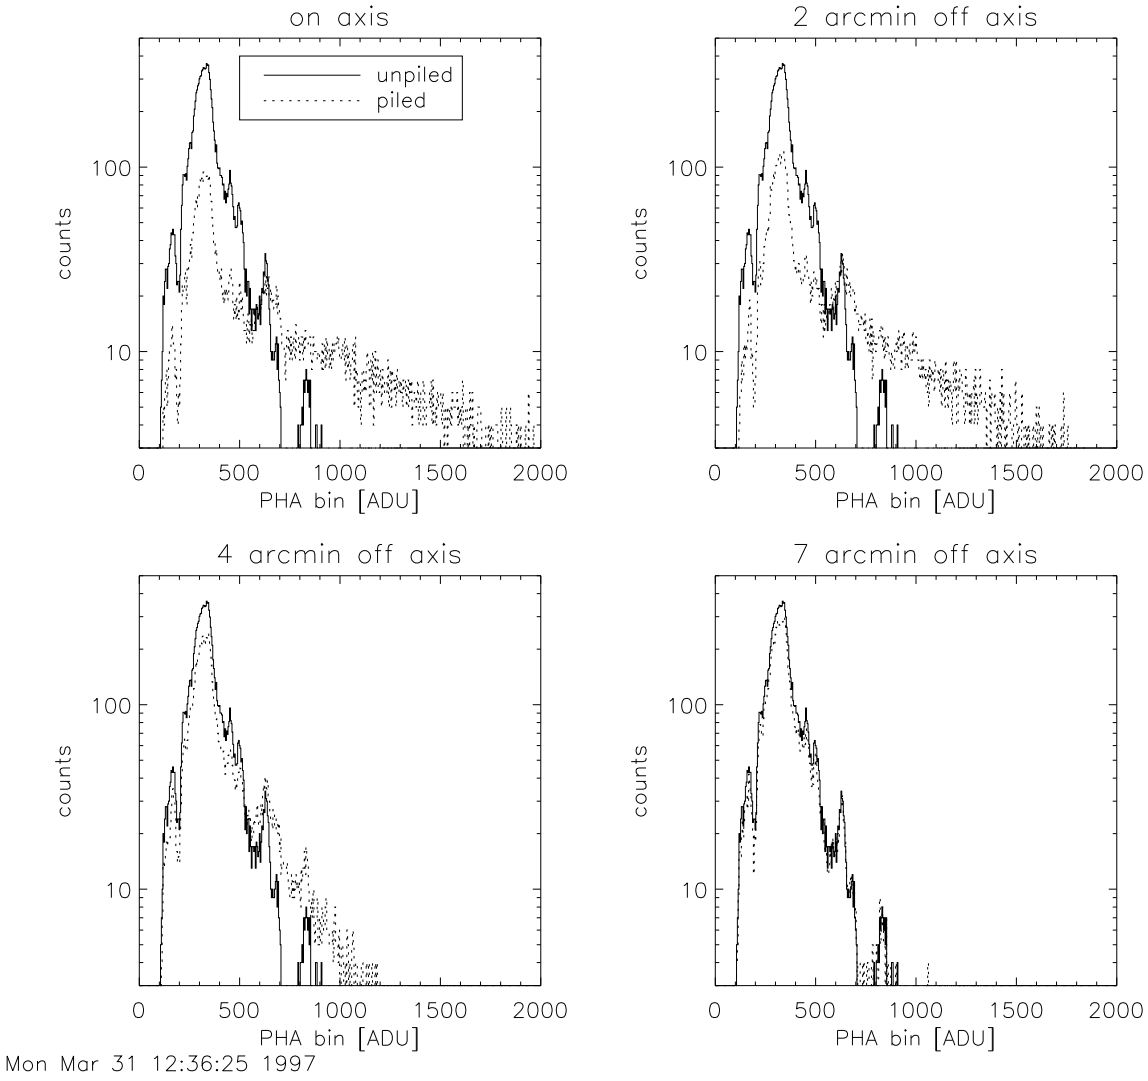

Figure 6.25: *MARX* simulations of the effect of pile-up on the shape of the spectrum. The true (solid line) and the detected (dotted line) spectra are shown for four different viewing angles. The corresponding "pile-up fractions"—see Section 6.16.2—are 46%, 40%, 15%, and 2% as the image is moved progressively further off-axis.

Simple Pile-Up Estimates The pile-up fraction is the ratio of the number of detected events that consist of more than one photon to the total number of detected events. An estimate of the pile-up fraction can be determined from Figure 6.24. The algorithm parametrizes the HRMA+ACIS PSF in terms of the fraction of encircled energy that falls within the central 3×3 pixel event detection cell, and assumes that the remaining energy is uniformly distributed among the 8 surrounding 3×3 pixel detection cells. The probabilities of single- and multiple-photon events are calculated separately for the central and surrounding detection cells and subsequently averaged (with appropriate weighting) to obtain the pile-up fraction as a function of the true count rate—the *solid* line in Figure 6.24. The model was tested against data taken on the ground under controlled conditions—also shown in Figure 6.24.

As a general guideline, if the estimated pile-up fraction is $> 10\%$, the proposed observation is very likely to be impacted. The first panel (upper left) in Figure 6.25 qualitatively illustrated the effect on a simulated astrophysical X-ray spectrum. *However, the degree of pile-up that is acceptable will depend on the particular scientific goals of the measurement, and there is no clear-cut tolerance level.* If one's scientific objective demands precise flux calibration, then the pile-up fraction should probably be kept well below 10%.

The *PIMMS* tool provides the pile-up fraction using the algorithm described here, both for direct observation with ACIS and also for the zeroth-order image when a grating is inserted.

Simulating Pile-Up John Davis developed an algorithm for modeling the effects of pile-up on ACIS spectral data. The algorithm has been implemented as of *XSPEC* V11.1 and *Sherpa* V2.2. The algorithm can be used to attempt to recover the underlying spectrum from a source, or to simulate the effects of pile-up for proposal purposes.

The algorithm has been tested by comparing CCD spectra with grating spectra of the same sources. Care should be taken in applying the algorithm—for example, using the appropriate regions for extracting source photons and avoiding line-dominated sources. A description of the algorithm can be found in Davis 2001 (Davis, J.E. 2001, ApJ, 562, 575). Details on using the algorithm in *Sherpa* are given in a *Sherpa* "thread" as of *CIAO* V2.2 on the *CXC CIAO* web page: https://cxc.harvard.edu/ciao/.

6.16.3 Reducing Pile-Up

Various methods that can be used to reduce pile-up are summarized in this section.

- **Shorten exposure time:** By cutting back on CCD exposure time, the probability of pile-up decreases. The user is advised to select the best combination of a subarray and frame time to avoid losing data as discussed in Section 6.13.1.

- **Use the Alternating Exposure option:** This option simply alternates between exposures that are subject to pile-up and those that are not. The capability was originally developed for use with certain grating observations to allow one to spend

some time obtaining useful data from a zeroth order image, which would otherwise be piled up. See Section 6.13.2.

- **Use CC mode** If two-dimensional imaging is not required, consider using CC mode (Section 6.13.3).
- **Insert a transmission grating:** Inserting either the HETG (Chapter 8) or the LETG (Chapter 9) will significantly decrease the counting rate as the efficiency is lower. The counting rate in the zero order image may then be low enough to avoid pile-up.
- **Offset point:** Performing the observation with the source off-axis spreads out the flux and thus decreases the probability of pile-up at the price of a degraded image. Figure 6.25 illustrated the impact.
- **Defocus:** The option is only listed for completeness, the option is *not* recommended or encouraged.

6.17 On-Orbit Background

There are three components to the on-orbit background. The first is that due to the cosmic X-ray background (a significant fraction of which resolves into discrete sources during an observation with *Chandra*). The second component is commonly referred to as the charged particle background. This arises both from charged particles, photons, and other neutral particle interactions that ultimately deposit energy in the instrument. The third component is the "read-out artifact" which is a consequence of the "trailing" of the target image during the CCD read-out; it is discussed in Section 6.13.1.

The background rates differ between the BI and the FI chips, in part because of differences in the efficiency for identifying charged particle interactions. Figure 6.26 illustrates why.

6.17.1 The Non-Celestial X-ray Particle Background

Beginning in 2002-Sep and continuing until 2012-Jun, observations have been carried out with the ACIS in the stowed position, shielded from the sky by the SIM structure, and collecting data in normal imaging TE VF mode at -120 °C. Chips I0, I2, I3, S1, S2, S3 were exposed. The SIM position was chosen so that the on-board calibration source did not illuminate the ACIS chips. This allowed characterization of the non-celestial contribution to X-ray background (i.e., of the effects of charged particles). The resulting spectra from different chips are shown in Figure 6.27. Chip S2 is similar to the ACIS-I chips (denoted I023 in the figure) and not shown for clarity.

In addition, in July-September 2001, *Chandra* performed several short observations of the dark Moon, which blocks the cosmic X-ray background. The dark Moon and stowed background spectra were indistinguishable (except for short periods of flares and variable Oxygen line emission in the Moon observations). No background flares have been observed

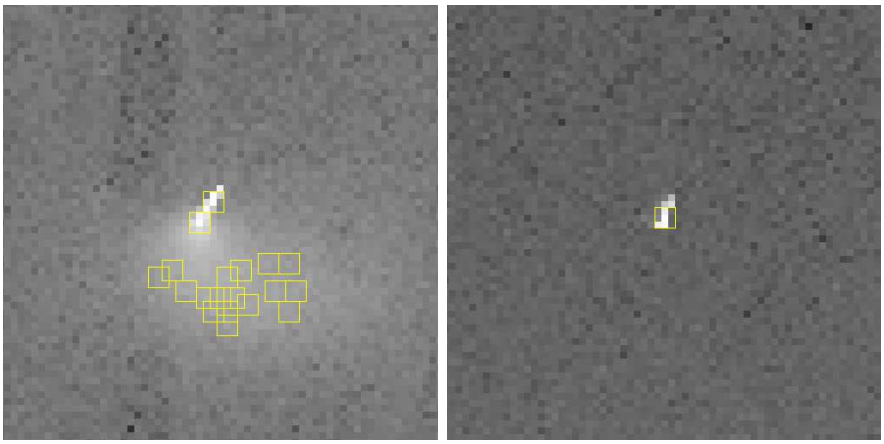

Figure 6.26: Enlarged view of an area of a FI chip I3 (left) and a BI chip (right) after being struck by a charged particle. There is far more "blooming" in the FI image since the chip is thicker. The overlaid 3×3 detection cells indicate that the particle impact on the FI chip produced a number of events, most of which end up as *ASCA* Grade 7, and are thus rejected with high efficiency. The equivalent event in the BI chip, is much more difficult to distinguish from an ordinary X-ray interaction, and hence the rejection efficiency is lower.

in the stowed position. Thus, the ACIS-stowed background is a good representation of the quiescent non-X-ray background in the normal focal position and can be used for science observations.

The flight-grade distributions in early measurements of the non-X-ray background for the two types of CCDs are shown in Figure 6.28. Although subsequent to these early measurements the CCD temperature has been lowered and the FI devices suffered the effects of the radiation damage, the background is still dominated by the same grades. Based on these data, events from flight grades 24, 66, 107, 214, and 255 are routinely discarded on-board. The total rate of the discarded events is available in the data stream. The remaining non-X-ray events telemetered to the ground are still dominated (70-95%) by other bad grades. They are not discarded on-board because some of them may turn out to be valid X-ray events after ground processing.

For data taken using the VF telemetry format (Section 6.15.2), the non-X-ray background can be reduced in data processing by screening out events with significant flux in border pixels of the 5×5 event islands. This screening leaves the data from faint sources essentially the same while reducing the FI background at different energies: a factor of 1.4 ($E > 6\,\mathrm{keV}$); 1.1 ($1 - 5\,\mathrm{keV}$); and 2 (near $\sim 0.5\,\mathrm{keV}$). For the BI chips the reductions are: 1.25 ($E > 6\,\mathrm{keV}$); 1.1 ($1 - 5\,\mathrm{keV}$); and 3 (near $0.3\,\mathrm{keV}$). In addition, the spatial nonuniformity of the non-X-ray background may be reduced by VF screening; see the background uniformity memo at https://cxc.harvard.edu/cal/Acis/detailed_info.html#background. The screening algo-

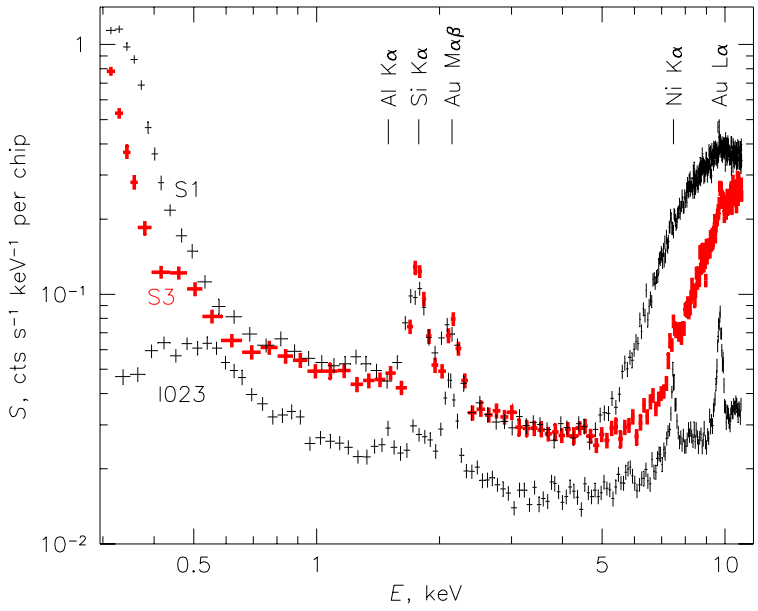

Figure 6.27: Energy spectra of the charged particle ACIS background with ACIS in the stowed position (a 50 ks exposure taken in 2002-Sep; standard grade filtering, no VF filtering). Line features are due to particle-induced fluorescence of material in and surrounding the focal plane. The "I023" points are from combined I0, I2, and I3 data.

rithm has been incorporated into the *CIAO* tool *acis_process_events*. Further discussion may be found at https://cxc.harvard.edu/cal/Acis/Cal_prods/vfbkgrnd/index.html.

Proposers should be aware that telemetry saturation occurs at lower count rates for observations using the VF format, so they may need to take steps to limit the total ACIS count rate (see Section 6.17.2). Proposers should also be aware that if there are bright point sources in the field of view, the screening criterion discussed above is more likely to remove source events due to pile-up of the 5×5 pixel event islands. Point sources should have count rates significantly less than 0.1 cts/sec to be unaffected. However, there is no intrinsic increase of pile-up in VF data compared to Faint mode, and the screening software can be selectively applied to regions, excluding bright point-like sources. Thus the use of VF mode is encouraged whenever possible.

6.17.2 The Total Background

In sky observations, two more components to the background come into play. The first is the cosmic X-ray background which, for moderately long ($\sim$ 100 ks) observations, will be mostly resolved into discrete sources (except for the diffuse component below 1 keV) but, nevertheless, contributes to the overall counting rate. The second is a time-variable "flare" component caused by any charged particles that may forward-scatter from the HRMA mirrors and have sufficient momentum so as not to be diverted from the focal plane by the magnets included in the observatory for that purpose, or from secondary particles (Section 6.17.3). Figure 6.29 compares flare-free ACIS-S3 spectra of the non-X-ray (dark Moon) background and a relatively deep pointing to a typical region of the sky away from bright Galactic features.

Estimates for the mid-2019 quiescent detector+sky background counting rates in various energy bands and for the standard good grades are given in Table 6.9. Insertion of the gratings makes little measurable difference in the background rates, but it does block the background flares. The lower-energy rates are very approximate as they vary across the sky. The rates are slowly changing on the timescale of months, so Table 6.9 can only be used for rough sensitivity estimates. Table 6.10 gives total background count rates for each type of chip, including all grades that are telemetered (see Section 6.15.1 and 6.17.1), and can aid in estimating the probability of telemetry saturation.

The trend in background rates follows approximately the inverse of the solar cycle with a lag of about a year; see Figures 6.30 and 9.21. The recent peak appears to be mid-2020, although it is quite broad. The rates declined much more rapidly starting in 2022 (compared to the decline from 2010 to 2013); the decline is more like that at the beginning of the mission (late-1999 to mid-2000). The 2025 rates are below the level of the $\sim$2014 minimum and even the 2000–2004 minimum, but the rates seem to have plateaued mid-2024. For this cycle, the quiescent background rates may begin increasing again.

For aid in data analysis and planning background-critical observations, the *CXC* has combined a number of deep, source-free, flare-free exposures (including all components of the background) into background event files for different time periods. These blank-sky

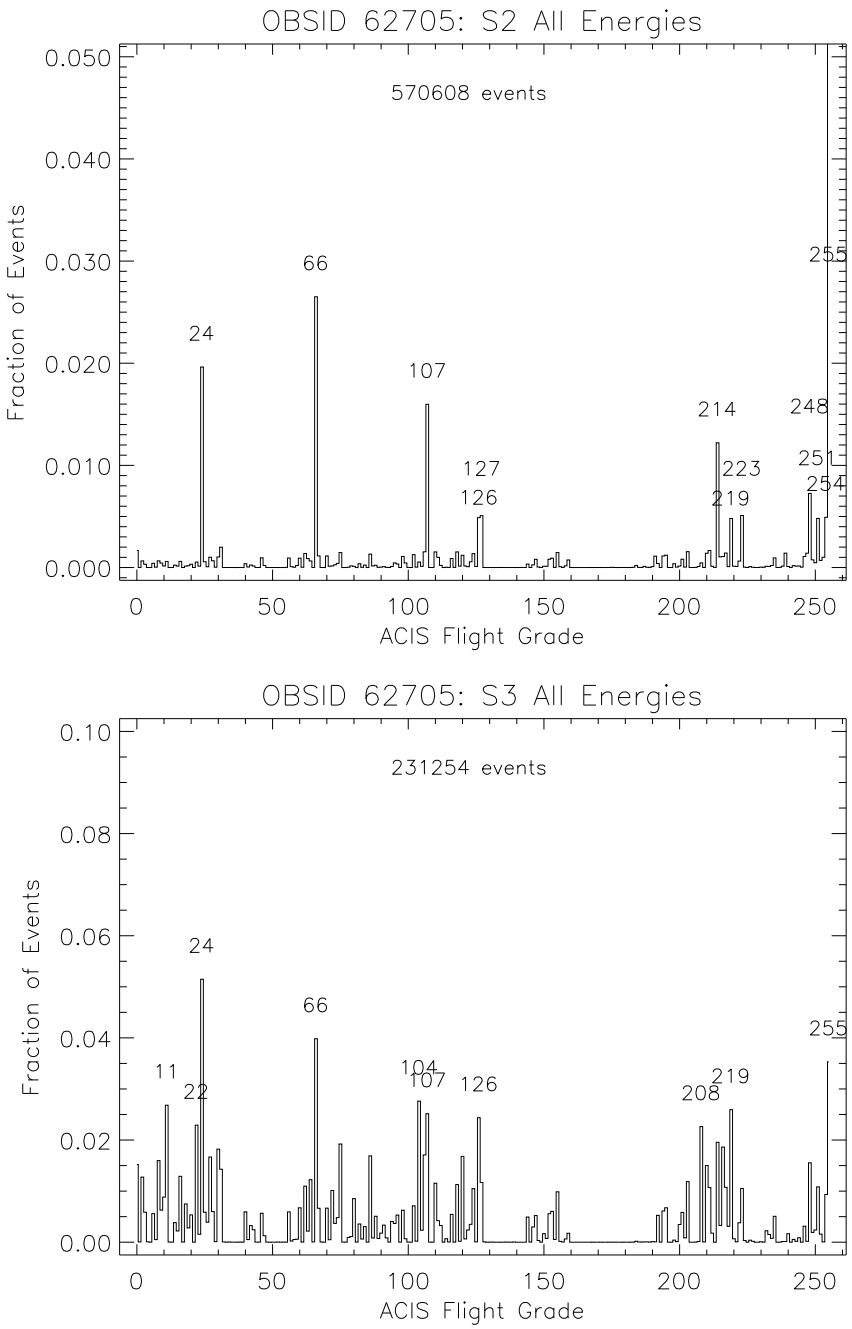

Figure 6.28: Fraction of ACIS background events as a function of grade from early in-flight data for an FI chip (S2; top) and a BI chip (S3; bottom).

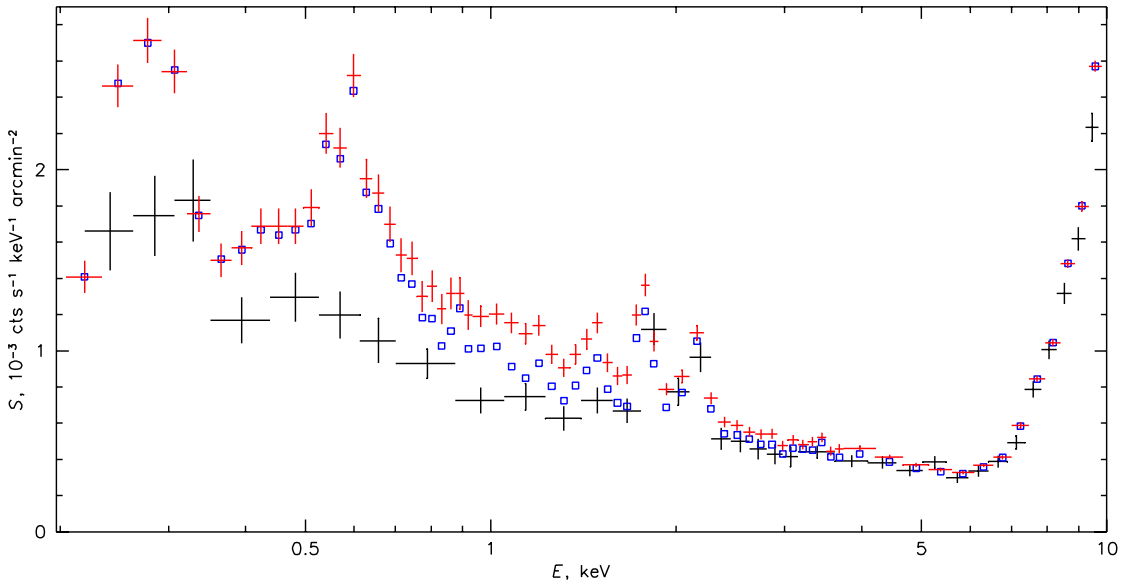

Figure 6.29: ACIS-S3 spectrum of the non-X-ray background (large crosses) overlaid on the quiescent blank sky spectrum. Small crosses show the total sky spectrum, while squares show the diffuse component left after the exclusion of all point sources detectable in this 90 ks exposure. Standard grade filtering and VF filtering are applied. The background and blank-sky spectra are normalized to the same flux in the 10–12 keV band.

	ACIS Background rates (cts/s/chip)							
Energy Band (keV)	I0	I1	I2	I3	S1	S2	S3	S4
0.3–10.0	0.31	0.32	0.31	0.34	2.35	0.45	1.21	0.44
0.5–2.0	0.05	0.05	0.05	0.06	0.21	0.08	0.18	0.09
0.5–7.0	0.18	0.19	0.19	0.20	0.64	0.26	0.47	0.25
5.0–10.0	0.17	0.17	0.16	0.18	1.64	0.25	0.73	0.23
10.0–12.0	0.07	0.07	0.07	0.08	1.34	0.15	0.93	0.14

Table 6.9: Approximate on-orbit standard grade background counting rates (2019, averaged April to September). The rates are cts/s/chip, using only *ASCA* grades 02346, no VF filtering, excluding background flares and bad pixels/columns and celestial sources identifiable by eye. These values can be used for sensitivity calculations. (See https://cxc.harvard.edu/cal/Acis/detailed_info.html#background for the most recent version of this table.)

Period	1999-Aug	2000-2003	2009	2025			
Upper E cutoff	15 keV	15 keV	15 keV	15[†] keV	13 keV	12 keV	10 keV
Chip S2 (FI)	10	6.3	10.7	3.9	2.9	2.8	2.5
Chip S3 (BI)	11	7.7	14.7	6.3	4.3	3.9	2.8

[†]Beginning in Cycle 11 (late 2009) the default high energy limit was changed from 15 keV to 13 keV. 13 kev/15 keV conversion factors of 0.82 for S2 and and 0.68 for S3. These factors were determined from observations with fairly constant background rates between 2009-Apr and 2009-Jul when the onboard high energy limit was set either to 15 keV or to 13 keV.

Table 6.10: Typical total quiescent background rates (cts/s/chip), averaged 2025-Jan through mid-2025-Oct), including *all* grades that are telemetered (and not just standard *ASCA* "good" grades), by chip type and upper energy cutoff. (See https://cxc.harvard.edu/cal/Acis/detailed_info.html#background for the most recent version of this table.)

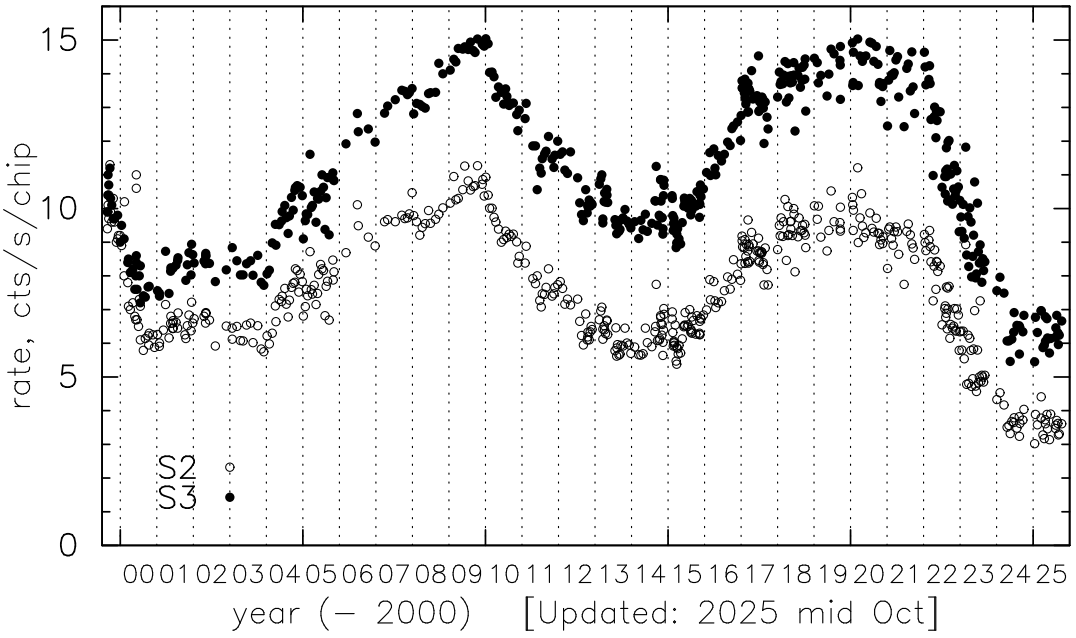

Figure 6.30: Total telemetered background rates (including all grades and the upper event cutoff at 15 keV) for chips S2 (FI) and S3 (BI) as a function of time. Vertical dashed lines are year boundaries. Data are plotted through mid-2025-Oct. Beginning in Cycle 11 (late 2009) the default high energy limit was changed from 15 keV to 13 keV. Subsequent 15 keV rate estimates were scaled from 13 keV rates using 13 kev/15 keV conversion factors of 0.82 for S2 and and 0.68 for S3. These factors were determined from observations with fairly constant background rates between 2009-Apr and 2009-Jul when the onboard high energy limit was set either to 15 keV or to 13 keV. (See https://cxc.harvard.edu/cal/Acis/detailed_info.html#background for the most recent version of this figure.)

datasets, along with the detector-only (ACIS-stowed) background files (Section 6.17.1), can be found in the CALDB.

H. Suzuki et al. (2021, Astronomy & Astrophys. 665, 116) examined the individual ACIS-stowed observations, and developed models describing the spatial and time variation of the background. A tool for generating background spectral models is available at https://github.com/hiromasasuzuk/mkacispback.

6.17.3 Background Variability

In general, the background counting rates are stable during an observation. Furthermore, the spectral shape of the non-X-ray background has been remarkably constant during 2000-2025 for FI chips and, to a lower accuracy, for BI chips, even though the overall background rate showed secular changes by a factor of 1.5. (For chip S3, the shape has been constant during 2000-2005, but a small change has been observed since late 2005.) When the *quiescent* background spectra from different observations are normalized to the same rate in the 10-12 keV interval, they match each other to within ±3% across the whole *Chandra* energy band. The previous discussion assumes that the upper threshold is set to 13 keV.

Occasionally, however, there are large variations (flares), as illustrated in Figure 6.31. Figure 6.32 shows the frequency of such variations when compared to the quiescent background. An average fraction of the exposure affected by flares above the filtering threshold used for the blank-sky background datasets (a factor of 1.2 above the nominal rate) was about 6% for FI chips and up to 1/3 for BI chips during the first few years of the mission. The average fraction of exposure affected by flaring has declined with time, and was practically zero for a long stretch. Recently, the flare frequency has been increasing, but at present it has not reached the frequency seen early in the mission. Thus, given that the quiescent background in FI chips is also lower than that in S3, background-critical observations may best be done with ACIS-I.

Several types of flares have been identified, including flares that are seen only in the BI chips, and flares that are seen in both the FI and BI chips. Figure 6.33 shows the spectra of two of the most common flare species. Both flares have spectra significantly different from the quiescent background. In addition, the BI flares have spatial distribution very different from that of the quiescent background. The BI flares produce the same spectra in S1 and S3.

Users should note that the total counting rate can significantly increase during a flare (although flare events are almost exclusively good-grade so the total rate does not increase by as large a factor as the good-grade rate; details can be found at https://cxc.harvard.edu/cal/Acis/Cal_prods/bkgrnd/current). If the probability of telemetry saturation is significant, users of ACIS-I might consider turning off the S3 chip. However, if ACIS-S is used in imaging mode, and flaring is a concern, the *CXC* recommends for extended sources appreciably filling S3 that both BI chips be turned on if thermal constraints permit (see Sections 6.20 and 6.22). The advantage is that for

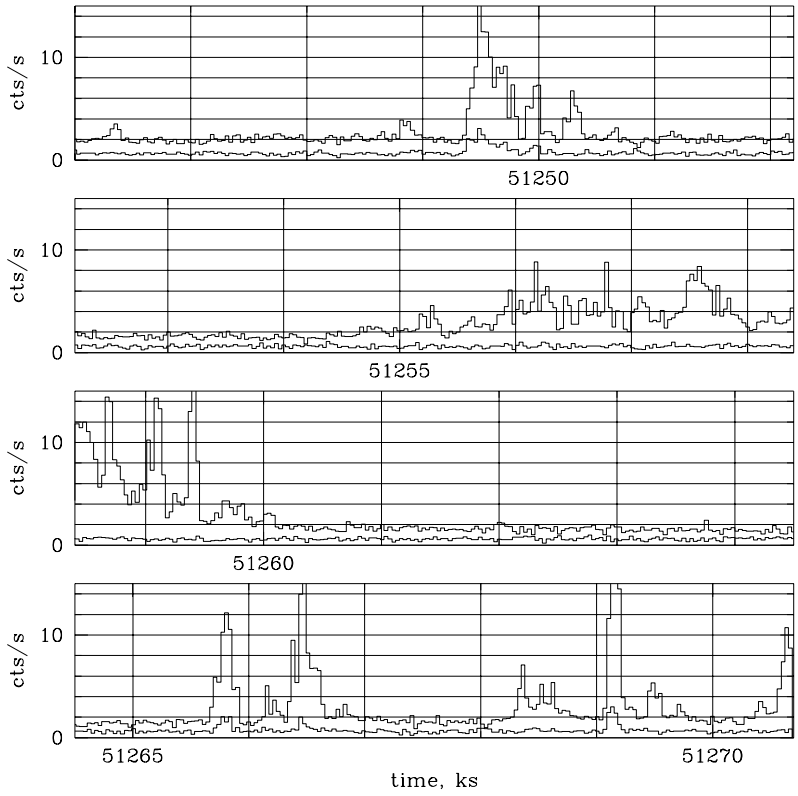

Figure 6.31: An example of the ACIS background counting rate versus time for a BI chip (S3; top curve) and an FI chip (I2; bottom curve). These are for the standard grades and the band from 0.3–10 keV.

most types of flares S1 can be used to determine the flare spectrum which can then be subtracted from the spectrum obtained with S3.

6.17.4 Background in Continuous Clocking Mode

Apart from compressing the data into one dimension (Section 6.13.3), there is essentially no difference in the total background in CC mode and that encountered in the timed exposure mode. The background per-sky-pixel, however, will be 1024 times larger, since the sky-pixel is now 1×1024 ACIS pixels.

6.18 Sensitivity

The sensitivity for detecting objects is best estimated using the various proposal tools such as *PIMMS*, *MARX*, etc. The "*Chandra* Proposal Threads" web page gives detailed

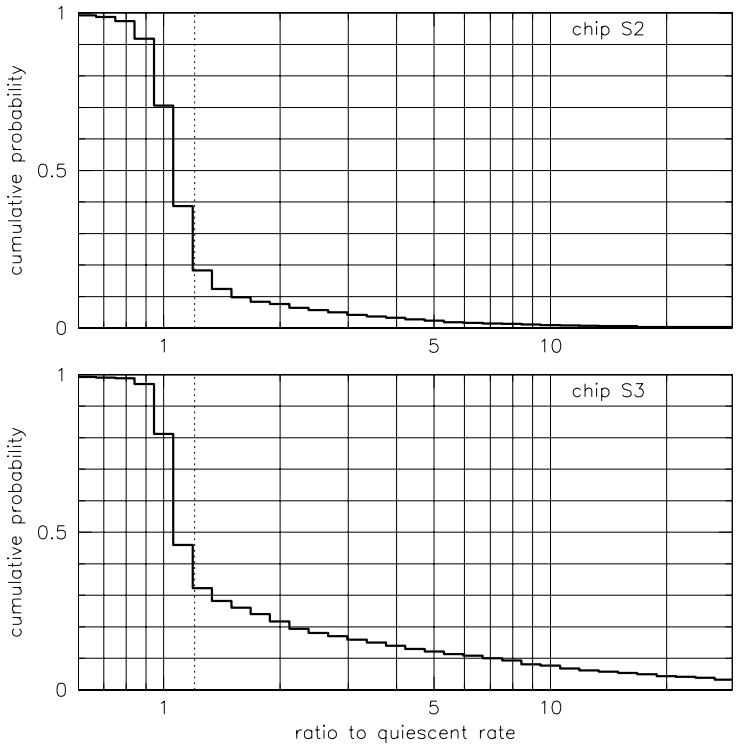

Figure 6.32: An estimate of the cumulative probability that the ratio of the background counting rate to the quiescent background counting rate is larger than a given value. The upper plot is for a representative FI chip (S2) and the lower plot is for a representative BI chip (S3). The vertical dotted line is a limiting factor 1.2 used in creating the background data sets. These probabilities are relevant for the archival data from the first 2–3 years of the mission. After declining with time to almost zero for a number of years, the amount of flaring has increased with solar maxima, but there has been less flaring than was seen in the first few years of the mission.

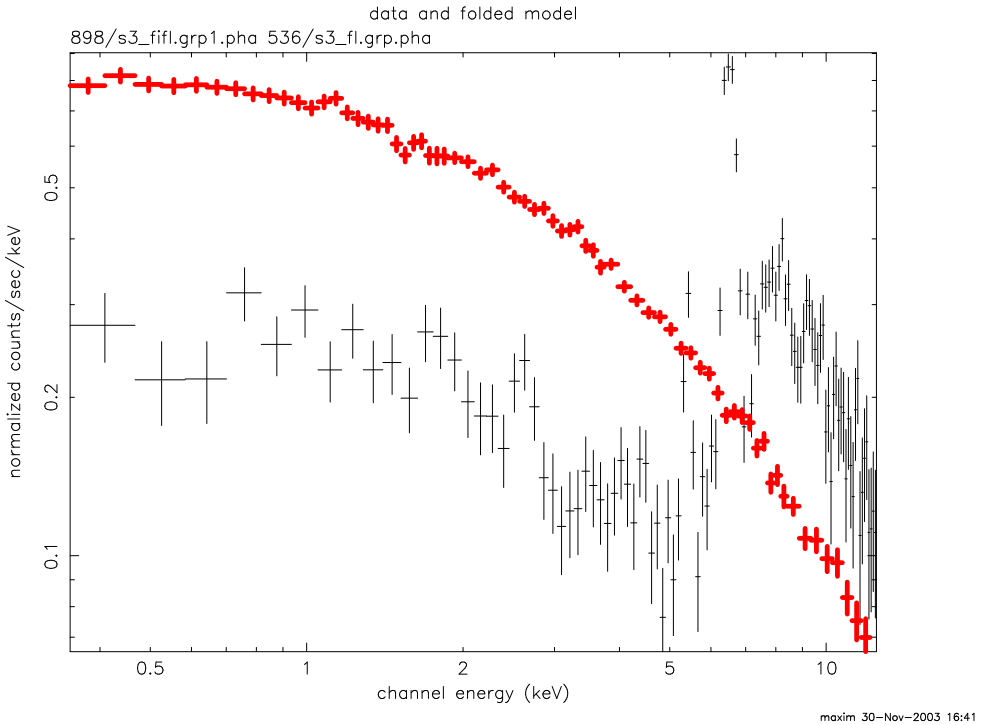

Figure 6.33: Spectra of different background flares in chip S3. Thick crosses show a common flare species that affects only the BI chips. Thin crosses show one of the several less common flare species that affect both the BI and FI chips. Note how both these spectra are different from the quiescent spectrum (see Figure 6.29).

examples of how to use these tools (https://cxc.harvard.edu/proposer/threads).

6.19 Bright Source X-ray Photon Dose Limitations

Pre-Flight radiation tests have shown that ∼200 krads of X-ray photon dose will damage the CCDs. The mechanism for the damage is the trapped ionization in the dielectric silicon oxide and nitride separating the gates from the depletion region. Since the charge is trapped, the damage is cumulative. Because the structure of the BI CCDs differs significantly from that of the FI CCDs, the two types of chips have different photon dose limitations. Specifically, the BI CCDs are more than 25 times as tolerant to a dose of X-ray photons as compared to the FI CCDs since the former have $\sim 40\mu$m of bulk Si protecting the gate layer.

Simulations of astrophysical sources have yielded a very conservative, spectrally-averaged, correspondence of 100 cts/pix = 1 rad. In this context 'counts' means all photons that impinged on the detector, whether or not they were piled-up or discarded on-board.

In consultation with the instrument principal investigator (IPI) team, the *CXC* has adopted the following mission allowances, per pixel of the two types of chips:

FI chips: 25 krads 2,500,000 cts/pix
BI chips: 625 krads 62,500,000 cts/pix

If an observation calls for observing a bright point-like source close to on-axis, it is suggested that the *MARX* simulator (with the parameter DetIdeal=yes & dither, typically, on) be used to calculate whether the observation may reach 1% of the above mission limits in any one pixel. If so, please contact the *CXC* HelpDesk (https://cxc.harvard.edu/helpdesk/) to custom design an observational strategy which may accommodate the science aims, while maintaining the health & safety of the instrument.

6.20 Limitations on the Number of Required and Optional CCDs

As of 2022-Jan, the ACIS threshold-crossings ("txings") algorithm is now the primary radiation monitor aboard *Chandra*. The txings rates on FI chips are much higher than on BI chips, and the FI chips are accordingly more sensitive to radiation. Thus, *at least* one FI chip is required for all ACIS observations. This change has the greatest impact for observers interested in achieving short frame times using a subarray on S3. The increase in frame time is largest for an observation which would have been performed on S3-only with a 128-row subarray, but which now requires an additional chip. In this case, the minimum frame time increases from 0.4 s to 0.5 s. **As of 2025-Nov, S3-only observations are no longer permitted.**

Many of the spacecraft components have been reaching higher temperatures over the course of the mission because of changes in the insulating layers on the exterior surfaces of the *Chandra* spacecraft. The ACIS electronics and Focal Plane (FP) temperatures can reach their operational limits depending on the orientation of the spacecraft and the number of operating CCDs. The number of operating CCDs and/or the duration of observations is limited due to more restrictive thermal constraints. Starting in Cycle 21, the proposal submission software allows the proposer to specify a maximum of 4 required CCDs. If a 5th and/or 6th CCD are desired by the proposer, they must be specified as optional at the time of proposal submission. This restriction is made necessary due to the difficulty in meeting the thermal constraints.

Please note that the solar pitch restrictions specified here, while they correspond to the pitch range boundaries given in other *CXC* documents, are practical representations of more complex physical behavior and should be understood as approximate.

Several components on the *Chandra* spacecraft have reached elevated temperatures at a variety of pitch angles. Figure 6.34 displays approximate pitch ranges and the components sensitive in those ranges. Within ACIS, the Power Supply and Mechanism Controller (PSMC) heats at pitch angles less than 90° and the Focal Plane (FP), Detector Electronics Assembly (DEA), and Digital Processing Assembly (DPA) heat at angles larger than 100°. Observations at pitch angles larger than about 120°, longer than about 30 ks, and with 5 or 6 CCDs operating are likely to approach or exceed the DPA and DEA thermal limits (Figures 6.36 and 6.37). Finally, the ACIS FP temperature is now most often several degrees above the desired operating temperature of −119.7 °C, particularly for observations with pitch angles larger than 100°. Most of these operating temperatures are reduced by powering fewer CCDs. Therefore, the Operations team will turn off optional CCDs as required to optimize ACIS performance and balance thermal constraints (see section 6.22.1). For this reason, observations must employ 4 or fewer (required on) CCDs whenever possible. Figure 6.35 illustrates the increase in PSMC temperature for observations at low pitch angles. Observations using 6 chips are plotted as green triangles, 5 chips as red circles, and 4 chips as blue crosses. Variations in the maximum temperature at a particular pitch angle and chipset correspond primarily to variations in the starting temperatures for the observations. The yellow and red lines indicate the yellow and red limits for the PSMC temperature. It is evident from Figure 6.35 that using one fewer CCD can reduce the temperature by a few degrees. In tail-Sun orientations (pitch angles larger than 100°), the ACIS FP temperature, the DEA temperature, and the DPA temperature can warm outside of the desired range. These temperatures can be reduced by using fewer chips. Roughly speaking, reducing the number of chips used by one will reduce the asymptotic temperature by a few to several degrees. Figure 6.36 displays the DPA temperature as a function of pitch angle for 4, 5, and 6 CCD configurations, and Figure 6.37 shows the corresponding data for the DEA. The FP temperature as a function of pitch angle for 4, 5, and 6 CCD configurations is shown in Figure 6.38.

The ACIS PSMC, DEA, DPA and FP temperatures are kept within operational limits through the use of thermal models which predict the temperature for a given week,

given the mix and timing of observations. If the predicted temperatures exceed the planning limits, adjustments are made such as turning off more optional CCDs, splitting an observation, or rescheduling an observation at a more favorable pitch angle.

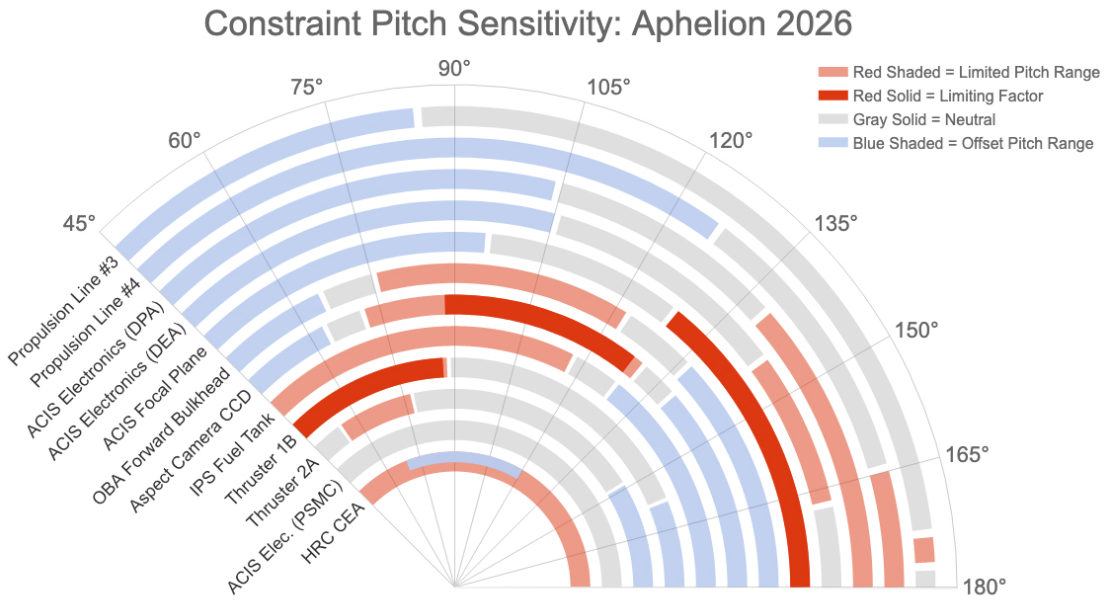

Figure 6.34: Pitch sensitivity of spacecraft components. Please see Chapter 3 for additional information about the various spacecraft component pitch sensitivities.

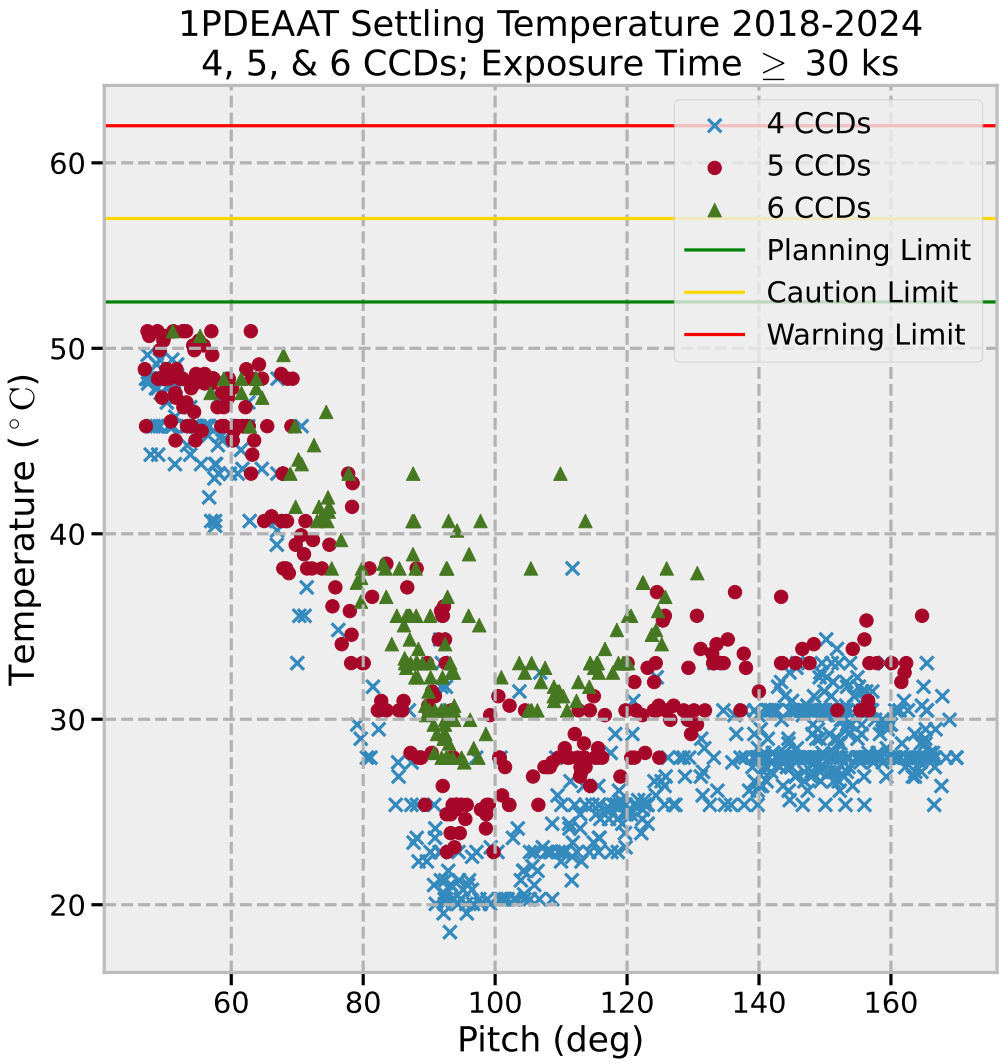

Figure 6.35: ACIS PSMC temperature as a function of spacecraft pitch angle. 6 CCD observations are indicated by green-filled triangles, 5 CCD observations by red-filled circles and 4 CCD observations are indicated by blue ×s.

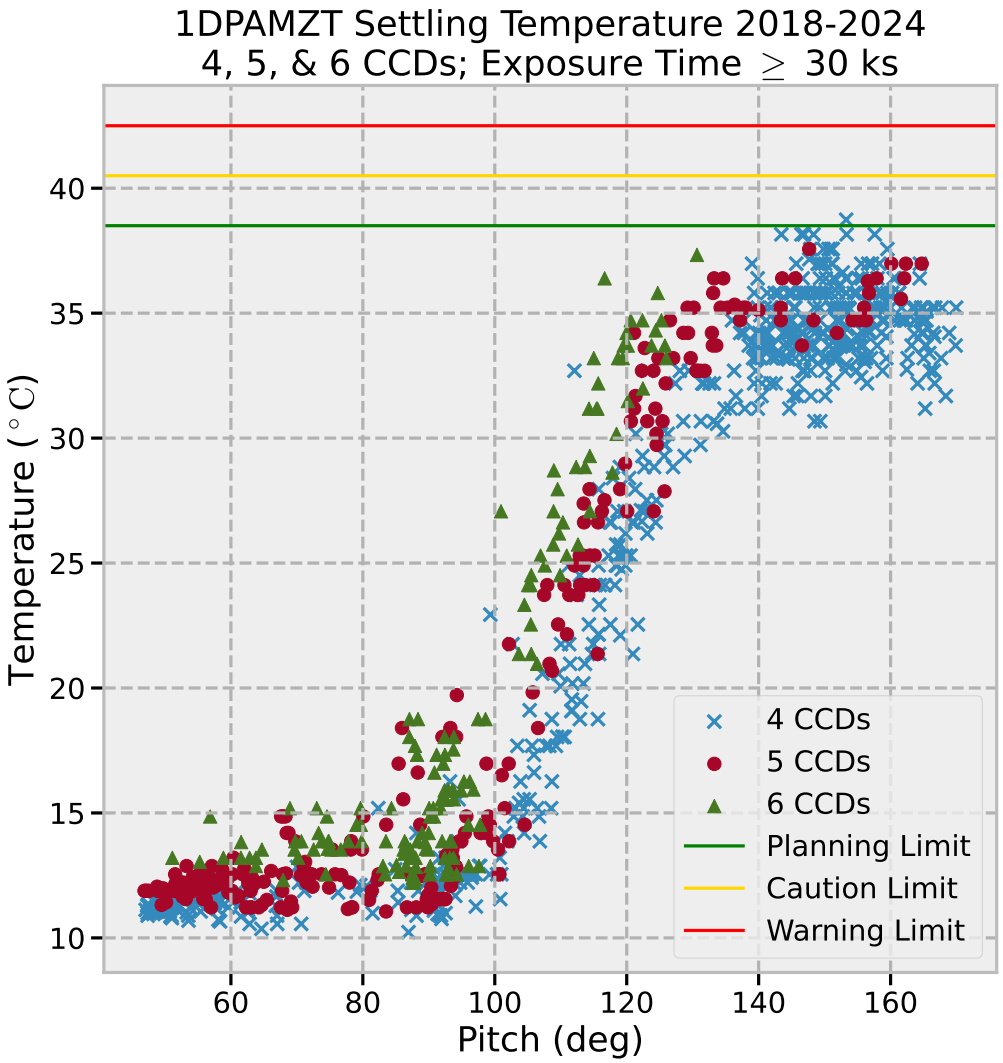

Figure 6.36: DPA temperature as a function of spacecraft pitch angle. 6 CCD observations are indicated by green-filled triangles, 5 CCD observations by red-filled circles and 4 CCD observations are indicated by blue ×s.

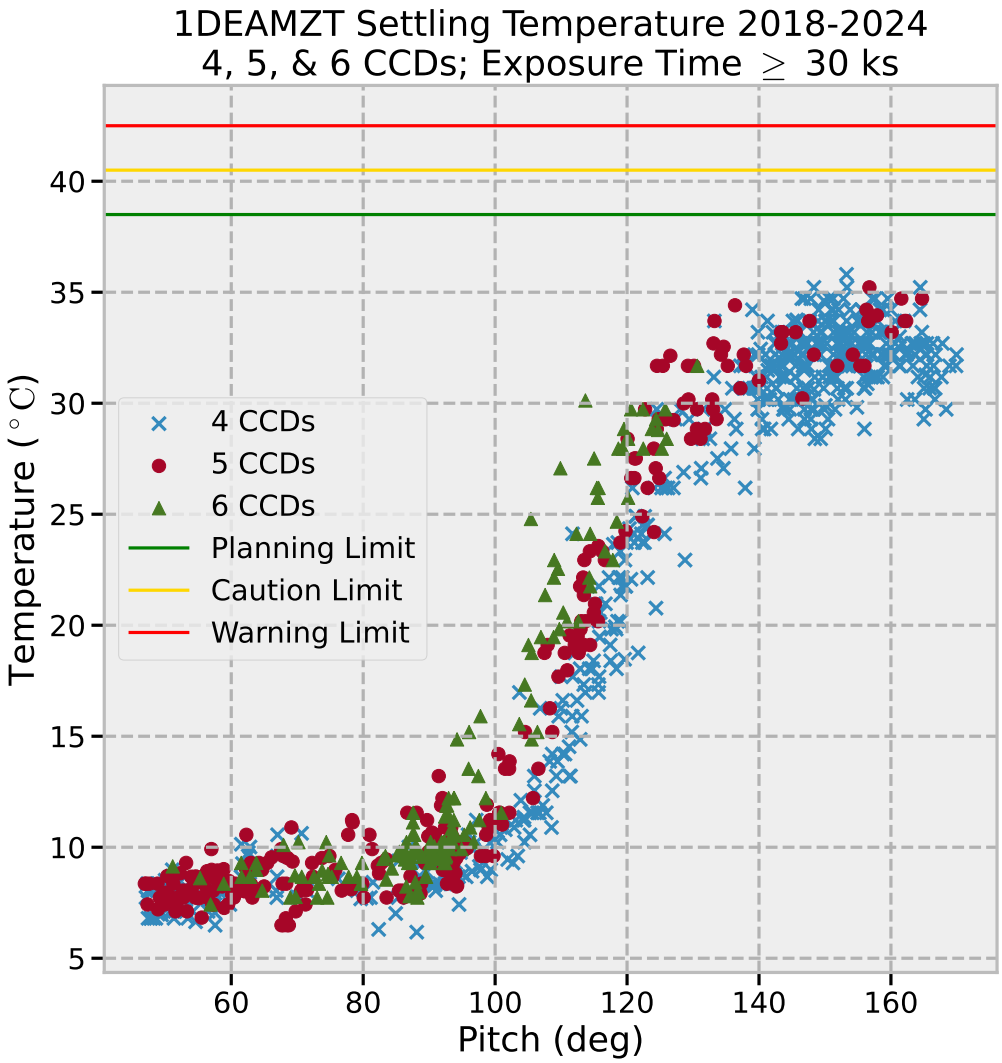

Figure 6.37: DEA temperature as a function of spacecraft pitch angle. 6 CCD observations are indicated by green-filled triangles, 5 CCD observations by red-filled circles and 4 CCD observations are indicated by blue ×s.

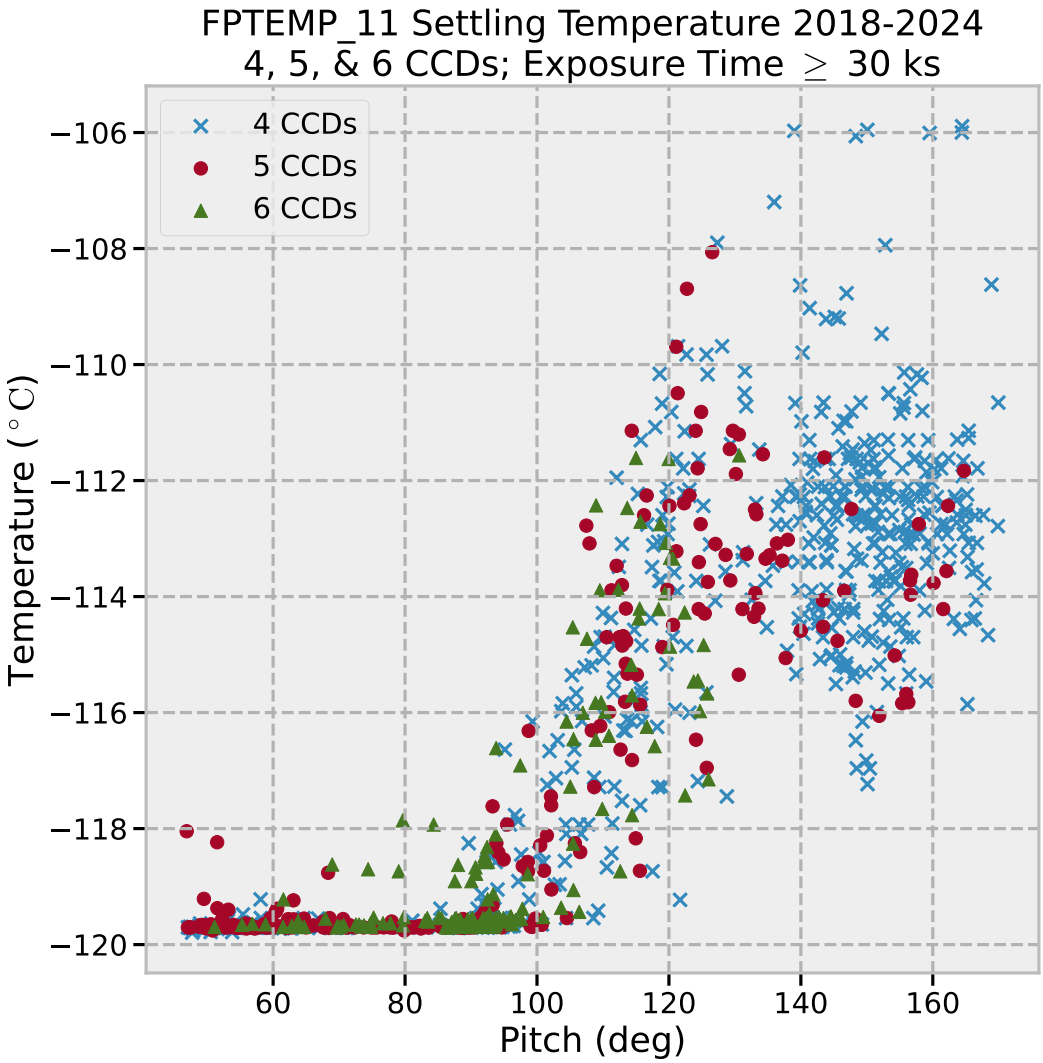

Figure 6.38: ACIS FP temperature as a function of spacecraft pitch angle. 6 CCD observations are indicated by green-filled triangles, 5 CCD observations by red-filled circles and 4 CCD observations are indicated by blue ×s.

6.21 Observing Planetary and Solar System Objects with ACIS

Chandra has successfully observed several solar system objects, including Venus, the Moon, Mars, Jupiter, and several comets. Observations of planets and other solar system objects are complicated because these objects move across the celestial sphere during an observation and the optical light from the source can produce a significant amount of charge on the detectors (this is primarily an issue for ACIS-S observations). Some information regarding observation planning and data processing is given here. Users are encouraged to contact the *CXC* for more detailed help.

6.21.1 The Sun, Earth, and the Moon

Chandra cannot observe the Sun for obvious reasons. *Chandra* has conducted observations of the Moon earlier in the mission, but observations of the Moon with ACIS are currently not allowed. The concern is that the bright flux of optical and ultraviolet (UV) photons could potentially polymerize the contaminant on the ACIS filters. Observations of the dark portion of the Moon are not allowed since there is a risk that the Sun-illuminated portion of the Moon might encroach upon the FOV during the observation. For similar reasons, ACIS observations of Earth (including the dark portion) are not allowed. See Section 3.3.2 in Chapter 3 and Chapter 5 for further discussion on avoidances and constraints.

6.21.2 Observations with ACIS-I

Any solar system object other than the Sun, Earth, the Moon, and Mercury can be observed with ACIS-I, subject to the avoidances discussed in Section 3.3.2 and Chapter 5. Previous solar system observations with ACIS-I have not revealed significant contamination from optical light. However, proposers are encouraged to work with the *CXC* when planning the specifics of a given observation. Since the source moves across the celestial sphere in time, an image of the event data will exhibit a "streak" associated with the source. The *CIAO* tool *sso_freeze* can be used to produce an event data file with the motion of the source removed.

6.21.3 Observations with ACIS-S

Any solar system object other than the Sun, Earth, the Moon, and Mercury can be observed with ACIS-S, subject to the avoidances discussed in Section 3.3.2 and Chapter 5. The ACIS-S array can be used with or without a grating. The BI CCDs are more sensitive to soft X-rays than the ACIS-I array CCDs, but the entire ACIS-S array suffers from the disadvantage that its OBF is thinner than for ACIS-I and may transmit a non-negligible flux of visible light onto the CCDs. It is thus necessary to estimate the amount of charge produced in the CCDs due to the optical light. More detailed information can be found

at https://cxc.harvard.edu/cal/Hrma/UvIrPSF.html and from the *CXC* via HelpDesk (https://cxc.harvard.edu/helpdesk/).

If the optical light leak is small enough, it can be mitigated by simply shortening the frame time. This leads to a linear drop in the number of ADU due to optical light. If possible, VF mode should be used, since in this mode the outer 16 pixels of the 5×5 region allows a "local" bias to be subtracted from the event to correct for any possible light leakage. However, see the warnings in Section 6.15.2.

The optical light also invalidates the bias taken at the beginning of the observation if a bright planet is in the field. It is therefore desirable to take a bias frame with the source out of the field of view. This bias map is useful even when processing 5×5 pixels in VF mode since it can be employed as a correction to the local average "bias" computed from the 16 outer pixels, thereby correcting for hot pixels, cosmetic defects etc.

A more sophisticated approach to dealing with excess charge due to optical light is to make an adjustment to the event and split thresholds. Event grades are described in more detail in Section 6.15.1. Excess charge (in ADU) due to optical light will be added to the event and split counters on-board. Without an adjustment to the thresholds (or a large enough threshold), many of the X-ray events may have all nine pixels of a 3×3 pixel event detection cell above the split threshold, in which case the event will not be telemetered to the ground. If the adjustment is too large, X-ray events may not be detected because they may not exceed the event threshold.

Users should be aware that if the detection thresholds are adjusted, standard *CXC* processing of planetary data will give inaccurate estimates of event pulse heights and grades. To analyze such data, a thorough understanding of the energy calibration process and manual massaging of the data will be required.

6.22 Observing with ACIS—the Input Parameters

This section describes the various inputs that either must be, or can be, specified to perform observations with ACIS. The subsections are organized to match the *Chandra* Proposal Software (CPS) form at https://cxc.harvard.edu/proposer/CPS.html and includes implications of the possible choices. As emphasized at the beginning of the Chapter, ACIS is moderately complex and the specific characteristics of the CCDs and their configuration in the instrument lead to a number of alternatives for accomplishing a specific objective—*detailed trade-offs are the responsibility of the observer*. For example, it might seem obvious that observations of a faint point source may be best accomplished by selecting the ACIS-S array with the aim point on S3, the BI device that can be placed at the best focus of the telescope, and the CCD with the best average energy resolution. On the other hand, perhaps the science is better served by offset pointing (by a few arcmin) the target onto S2, very near to the frame store, where the FI energy resolution is better than that of S3. Or, if the object is very faint—so that the total number of photons expected is just a handful (not enough to perform any significant spectroscopy)—the advantage of S3 may not be so obvious considering the smaller field of view and its higher background rate, and perhaps the ACIS-I array, which would optimize the angular resolution over a larger field, may be more attractive.

6.22.1 Required Parameters

Some ACIS input parameters must be specified: the number and identity of the CCDs to be used, the Exposure Mode, and the Event Telemetry Format. For imaging observations using ACIS-I or ACIS-S with no grating, the maximum number of counts expected to be used in a spectral analysis of the source(s) must be supplied. If pile-up and telemetry saturation are not expected to be a problem for the observation, then these are the only parameters that need to be specified.

Number and Choice of CCDs The CPS requires the observer to specify the desired aimpoint and to identify the CCDs they want to use. Prior to Cycle 13, the use of 6 CCDs was encouraged to facilitate serendipitous detections. Starting in Cycle 20, the CPS limits the number of CCDs that may be specified as required to a maximum of 4 for thermal reasons. The proposer may specify additional optional CCDs to bring the total of required plus optional CCDs to as many as 6 CCDs, but the proposer should be aware that optional CCDs may be turned off. If the science objectives of the proposal require 5 or 6 CCDs, the proposer must work with their Chandra Uplink Support Scientist after the proposal is selected in order to specify a 5th or 6th CCD as required. But proposers should be aware that 5 or 6 CCD observations are more difficult to schedule and are more likely to be split into multiple short observations if the exposure is long.

Using fewer CCDs is beneficial in keeping the ACIS electronics and the ACIS focal plane (FP) temperatures within the required operating ranges. See Section 6.20 for further

information on thermal limitations and selection requirements for ACIS and the number of operating CCDs.

The selection of CCDs on the CPS form is discussed in the next section.

Choosing Optional CCDs & Optional CCD Policy The observer may specify that a given CCD must be on for an observation by entering "Y" for that CCD at the appropriate place in the CPS form. If there are CCDs that the observer does not require for their science, they should enter "N." If there are CCDs that the observer would prefer to have turned on should thermal conditions allow it, the rank-ordered designations "OFF1," "OFF2," up to "OFF5" should be used. The CCD designated as "OFF1" would be the first one to be turned off, and the CCD designated as "OFF5" would be the last that would be turned off.

Even if the science requires 5 or 6 CCDs, the observer must set the designation for the four most useful as "Y" and the least useful CCD(s) to optional status "OFF1" ("OFF2"). In these instances, the proposer should include a comment in the CPS form that 5 or 6 CCDs are required for the science. If the proposal is accepted, the observer may work with their Chandra Uplink Support Scientist to change "OFF1" ("OFF2") to "Y." The *CXC* will make its best effort to try to schedule the observation under the appropriate thermal conditions. The observer should discuss the configuration with Uplink Support and, if there are difficulties in assessing which CCDs should be optional, please contact the *CXC* HelpDesk (https://cxc.harvard.edu/helpdesk/). Should it be possible to accommodate the observer's request for 5 or 6 CCDs, the observation will most likely take place at solar pitch angles less than 130°.

Recommended Chip Sets Observers should specify the chip set that is best for their primary science. The following suggestions have proven to be popular and would facilitate a more useful and homogeneous archive. Given the current thermal performance of the spacecraft, it is possible that any number of optional CCDs would be turned off.

Some Recommended Chip Sets: ACIS-I Aimpoint (Imaging) The rationale for the first ACIS-I imaging configuration (Figure 6.39) is that, in the unlikely event of major background flares, telemetry might saturate more rapidly if S3 were on. In addition, since S2 is further from the ACIS-I aimpoint on I3, data on S2 may provide nearby local (FI) background if the target contains diffuse emission which does not extend as far as S2.

For the second ACIS-I imaging configuration (Figure 6.40), the rationale is that S3 is generally more sensitive and closer to the ACIS-I aimpoint, and so more sensitive to serendipitous source detection.

The rationale for the third ACIS-I imaging configuration (Figure 6.41) is that it is desired to have both S2 and S3 on, but it is not required. The CCD specified "OFF1" would be the first to be turned off.

Some Recommended Chip Sets: ACIS-S Aimpoint (Imaging) For deciding on chipsets for ACIS-S imaging, several factors come into play. In general, chips closest to the S3 aimpoint would be selected as required, while those farthest from the aimpoint (where the PSF is degraded) would be selected as optional.

The first ACIS-S imaging configuration (Figure 6.42) is the basic default ACIS-S configuration. The aimpoint is on S3, a BI chip. If a BI chip is used, it is required that an FI chip also be on in order to provide more reliable monitoring of the radiation environment. S2 has several advantages as an FI chip to be used with the S3 aimpoint, notably that is closest to the S3 aimpoint, providing relatively good imaging performance. With only two required CCDs powered on, the focal plane heating is reduced.

If a larger field of view is needed, up to six total chips can be specified as potentially on, of which only a maximum of four can be set as required. That is, a setup with four required chips could have up to two optional chips (tagged as "OFF1" and "OFF2"). As noted above, thermal constraints may require that optional chips be turned off.

The rationale for the second ACIS-S imaging configuration (Figure 6.43) is that the optional S1 and S4 allow for the possibility of a larger imaging area, and S4 is preferred over S1 if one chip must be dropped: if necessary, S1 will be turned off first. Thermal constraints may require one or both of the optional chips be turned off.

The rationale for the third ACIS-S imaging configuration (Figure 6.44) is that the optional S1 and S4 allow for the possibility of larger imaging area. As having S1 on allows the S3 background spectrum to be modeled if diffuse emission covers all or part of S3, S1 is preferred over S4 in the situation that thermal constraints require a single chip must be dropped. If it important to the science that S1 be available for this purpose, the user should instead specify a configuration with three required chips (S3+S2+S1).

The rationale for the fourth ACIS-S imaging configuration (Figure 6.45) is that the optional I2 allows for the possibility of larger imaging area. S1 will have a higher count rate than an FI CCD in the event of a background flare and thus it might be desirable to turn S1 off.

The rationale for the fifth ACIS-S imaging configuration (Figure 6.46) is that the observer may want to use S1 to model the background on S3, and it is desirable to have S4 turned off. S4 also has significant noise streaks with resulting decreased sensitivity, so some users may prefer to specify I2 as optional instead of S4.

The rationale for the sixth ACIS-S imaging configuration (Figure 6.47) is that for this observation it would be desirable to have 6 CCDs on. It is probable that one or two optional CCDs would be turned off.

Recommended Chip Set: ACIS-S Aimpoint (Spectroscopy) The optimum ACIS-S spectroscopy chip set depends strongly on the expected spectrum of the target. Typically the maximum signal is desired, so the HETG and LETG observer is most likely to request 6 ACIS-S chips (with at most 4 specified "ON"). If the science does depend strongly on the flux received on the S0 and S5 CCDs, the observer would need to work with their Chandra Uplink Support Scientist to specify all 6 ACIS-S CCDs as required

on (see Figure 6.48). However, the observer should be aware that the amount of useful flux on S0 and S5 with the LETG is typically quite low given they are both FI CCDs (see Figure 9.6) and also because of the loss due to the optical blocking filter contamination layer.

ACIS Counts Information Proposers not using gratings are asked for the maximum number of counts expected from the brightest source to be analyzed in their observation. This answer is used by mission planners to schedule observations during the proper thermal conditions for the perceived science goals. If the proposer is interested in the optimal spectral performance of the ACIS CCDs, they should seriously consider using the minimum number of CCDs necessary for their science goals. See also Section 6.22.5 for information on "warm" ACIS operations.

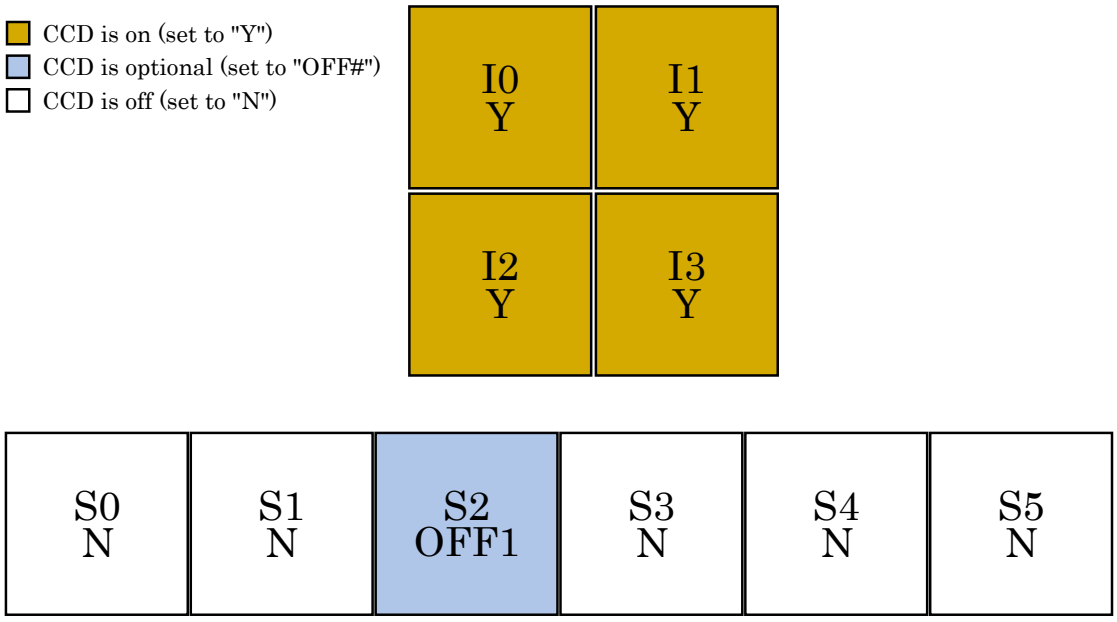

Figure 6.39: ACIS-I imaging, nominal I3 aimpoint (4 Required CCDs, 1 Optional CCD)

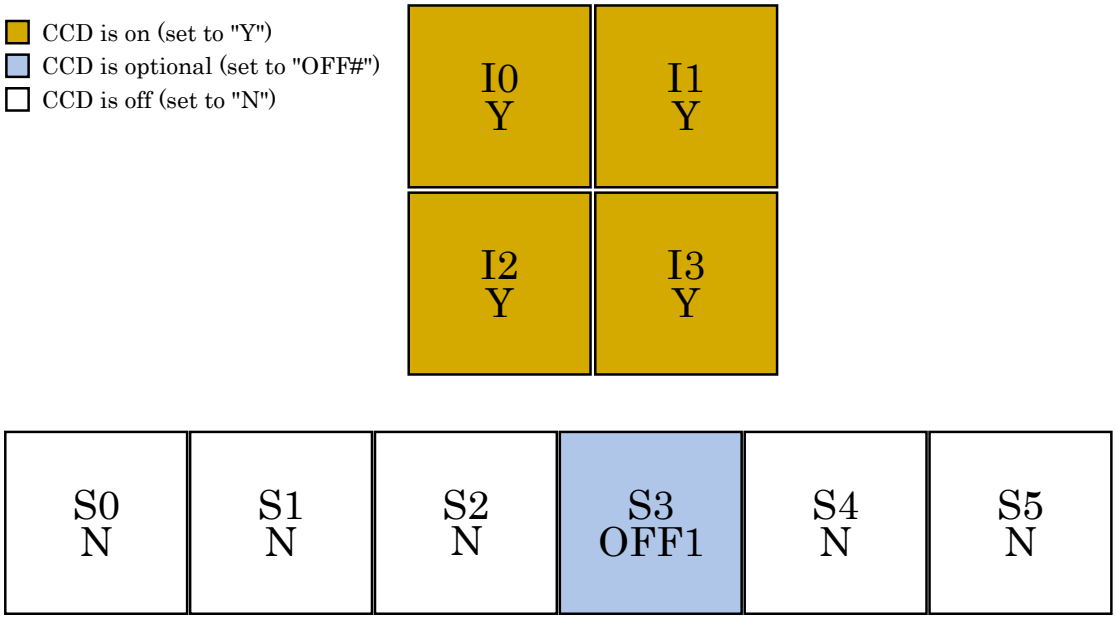

Figure 6.40: ACIS-I imaging, nominal I3 aimpoint (4 Required CCDs, 1 Optional CCD)

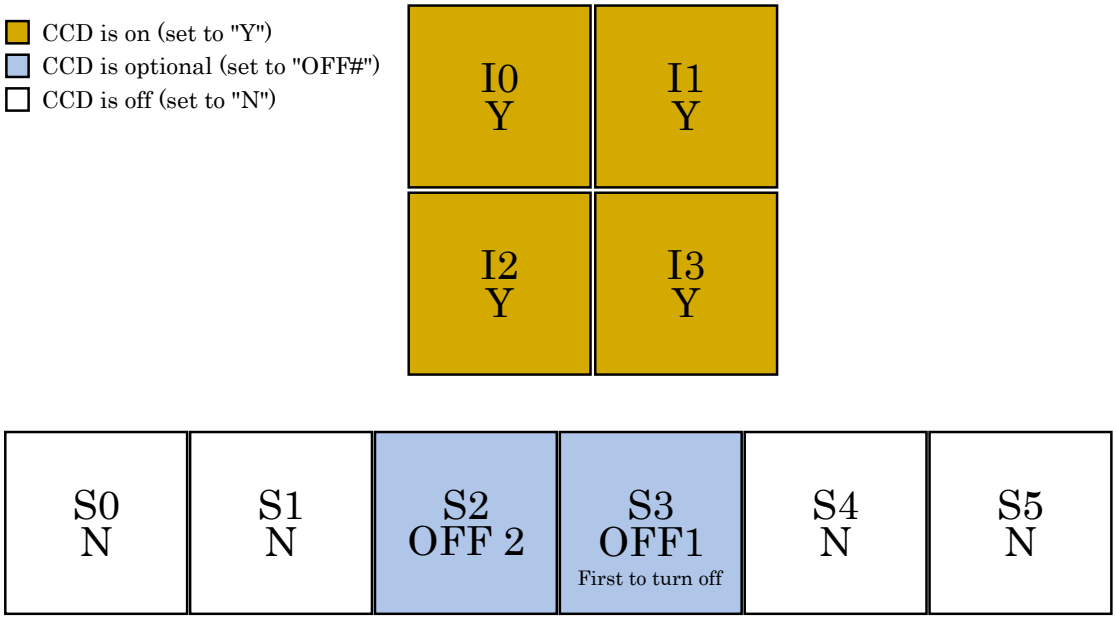

Figure 6.41: ACIS-I imaging, nominal I3 aimpoint (4 Required CCDs, 2 Optional CCDs)

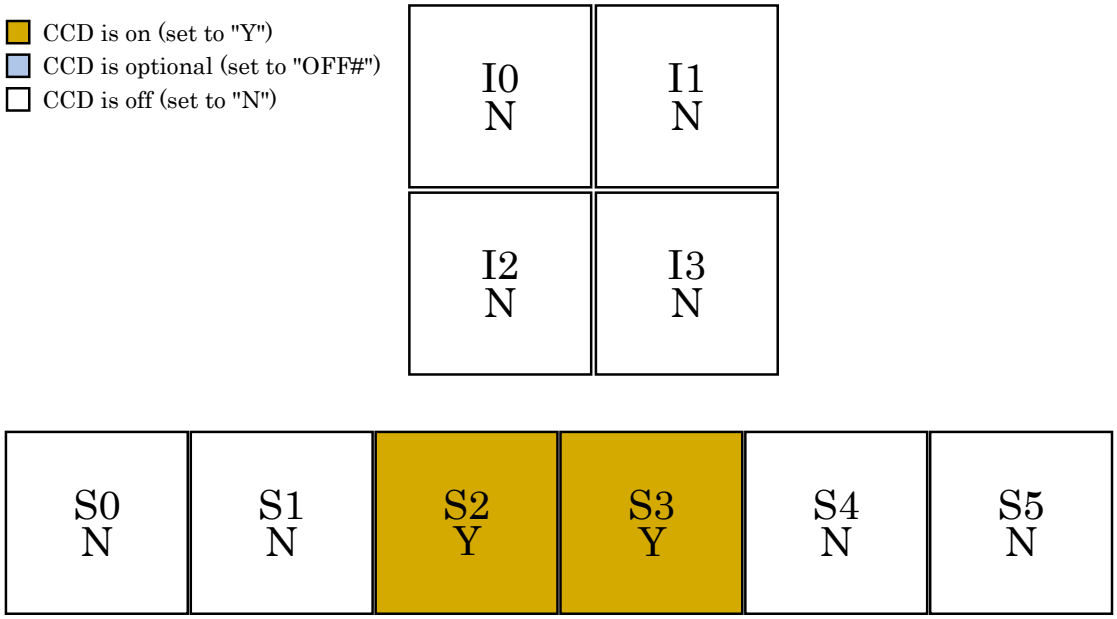

Figure 6.42: ACIS-S imaging, nominal S3 aimpoint (2 Required CCDs)

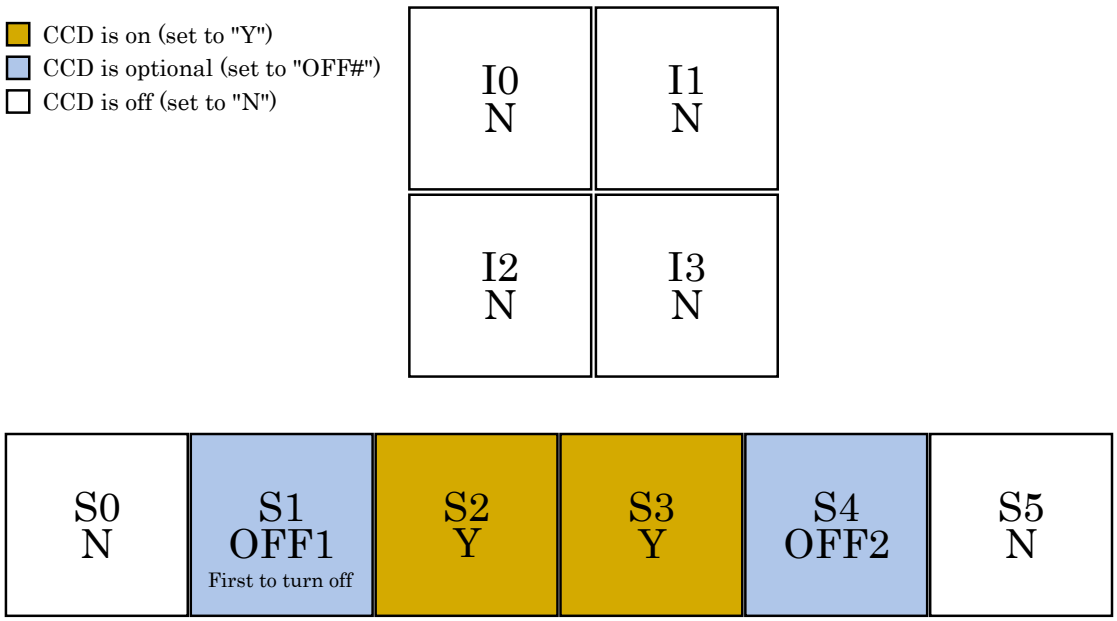

Figure 6.43: ACIS-S imaging, nominal S3 aimpoint (2 Required CCDs, 2 optional CCDs)

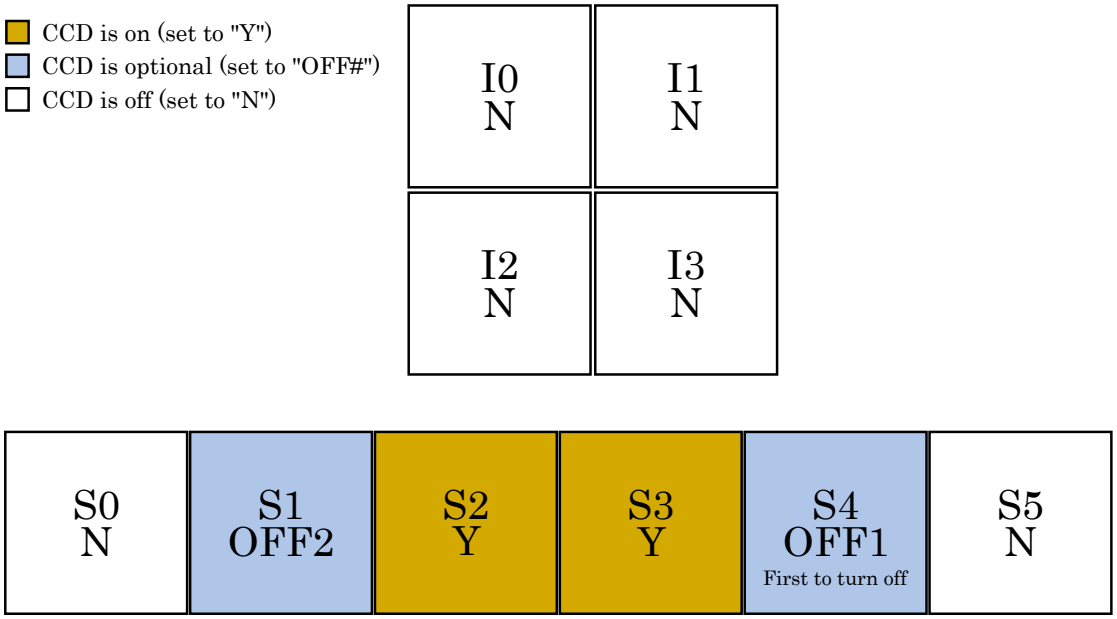

Figure 6.44: ACIS-S imaging, nominal S3 aimpoint (2 Required CCDs, 2 optional CCDs)

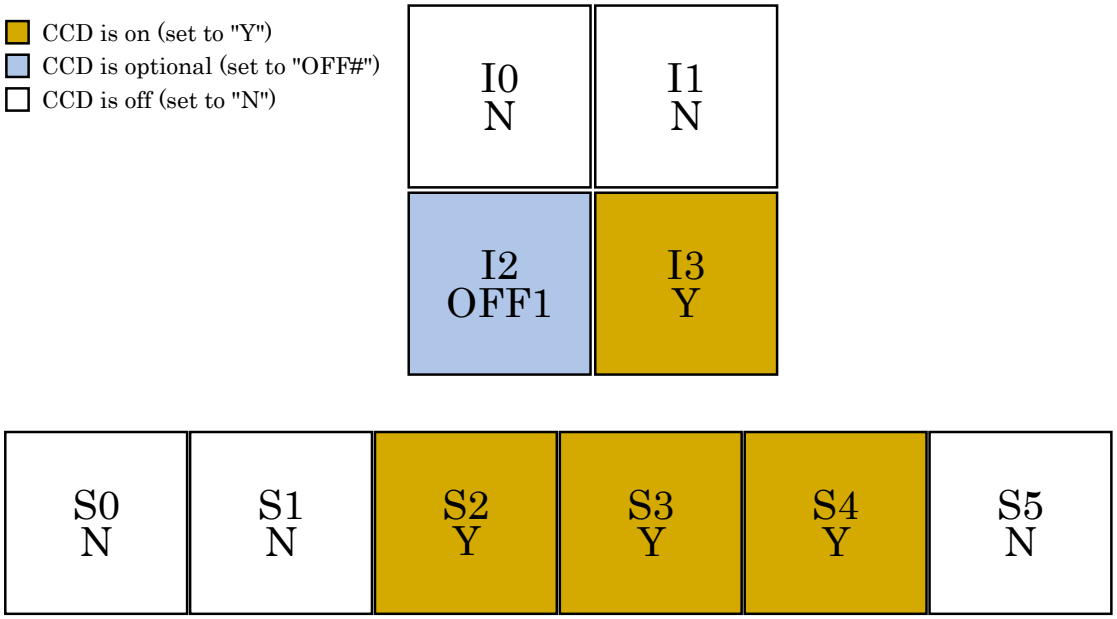

Figure 6.45: ACIS-S imaging, nominal S3 aimpoint (4 Required CCDs, 1 Optional CCD)

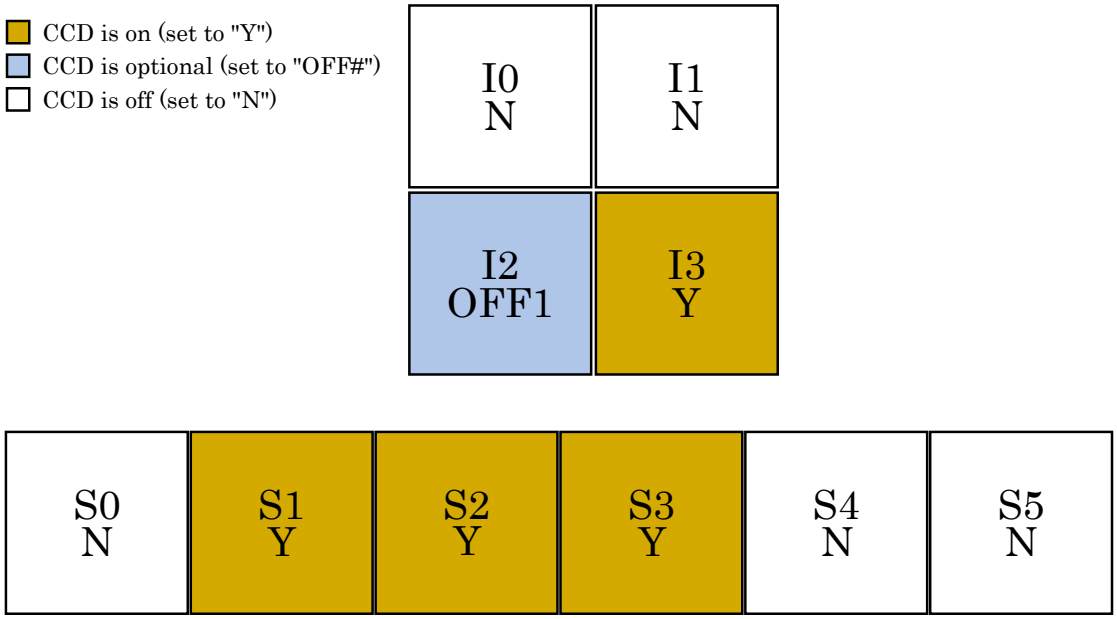

Figure 6.46: ACIS-S imaging, nominal aimpoint (4 Required CCDs, 1 Optional CCD)

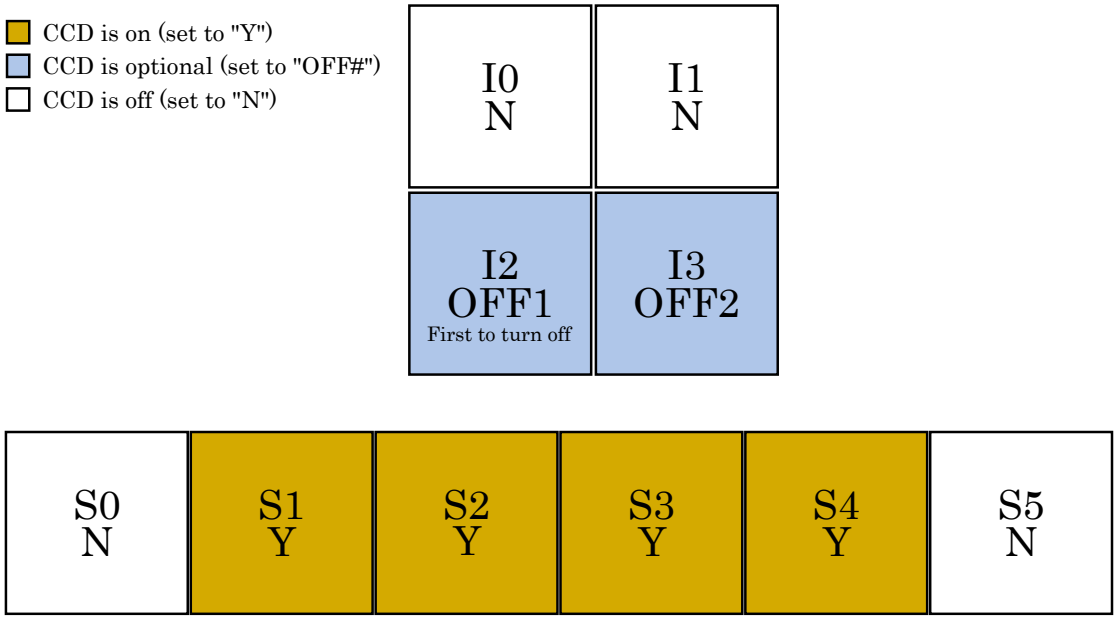

Figure 6.47: ACIS-S imaging, nominal aimpoint (4 Required CCDs, 2 Optional CCDs)

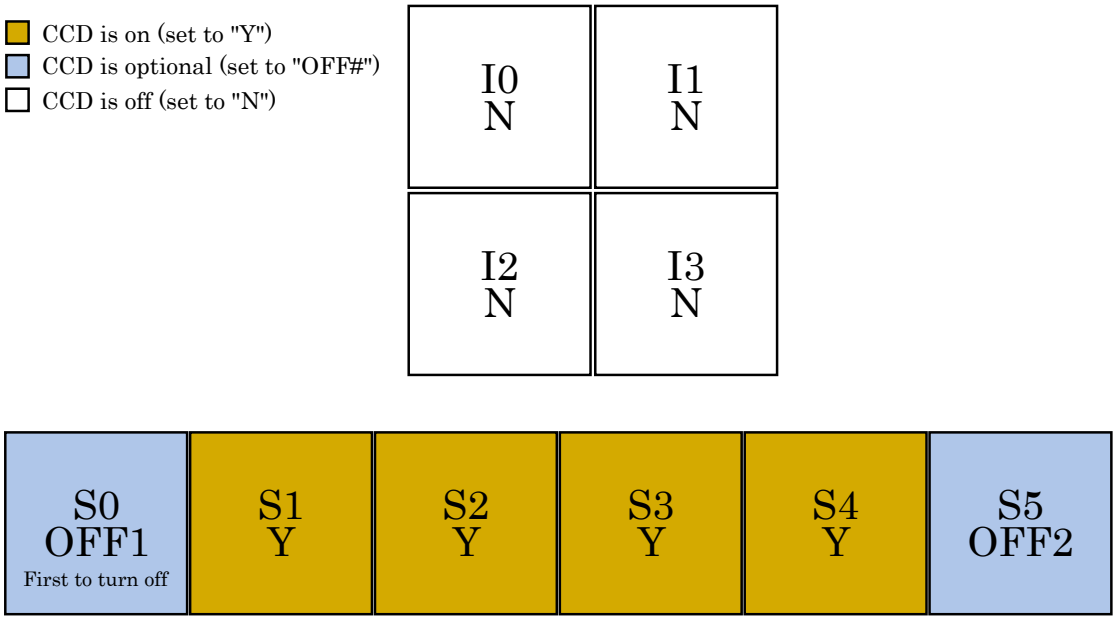

Figure 6.48: ACIS-S spectroscopy, nominal aimpoint (4 Required CCDs, 2 Optional CCDs)

Optional Parameters that affect Pile-Up See Section 6.13.1 for information on how the selection of the number of operating CCDs and the size and position of the subarray affects the minimum frame time. In general, using fewer CCDs will permit a faster read-out. For example, using the smallest subarray of 128 rows in the middle of I-array CCDs gives a minimum frame time of 0.5 s with 1 CCD operating and a minimum frame time of 0.7 s with 6 CCDs operating. In almost all cases, the observer should leave the parameter on the CPS target form that asks if the most efficient frame time should be used at the default value of "Y." This parameter should be changed only if the observer fully understands the impact on efficiency of their observation. Please contact the *CXC* HelpDesk (https://cxc.harvard.edu/helpdesk/) if you are considering a non-default setting and are unsure of the CCD selection and CCD Frame exposure time to use.

Exposure Mode There are only two choices: Timed Exposure (TE) mode (see Section 6.13.1) or Continuous Clocking (CC) mode (see Section 6.13.3).

Timed Exposure Mode The Timed Exposure (TE) mode with the default nominal (and optimal) frame time of 3.2s is the typical mode for ACIS observations. Note that the option of selecting frame times shorter than nominal reduces observing efficiency, and hence the number of photons collected for a given observation time. (Note that the value of the nominal frame time can differ a bit from 3.2s depending on how many CCDs are used; see Section 6.13.1.)

Continuous Clocking Mode The Continuous Clocking (CC) mode is useful when timing data are so critical and/or pile-up is such a problem that the sacrifice of one dimension of spatial data is warranted. The use of continuous clocking may also lead one to consider specifying a particular satellite roll orientation (see Chapter 3) to avoid having two different sources produce events in the same CCD column. (See also Section 6.22.4 below.)

6.22.2 Optional Parameters

Alternating Exposures This option applies *only* to Timed Exposure (TE) mode. The parameters specifying an alternating exposure are:

- the number of secondary exposures per primary exposure (1–15)
- the primary exposure frame time

Frame times and efficiencies in TE mode are discussed in Section 6.13.1, and the Alternating Exposure option is discussed in Section 6.13.2.

Energy Filtering It is possible to remove events from the telemetry stream, and thus avoid telemetry saturation, by specifying an energy acceptance filter within which detected events will be telemetered. The default discards events above 3250 ADU (nominally 13 keV). The total per-chip background rates for different upper energy cut-offs are in Table 6.10.

Starting 2006-Sep, a new PHA to energy conversion is used for observations with energy filters. Two sets of conversions are used, depending on the aimpoint of the observation. Observations with ACIS-S at the aimpoint use a conversion tailored for the BI CCDs and those with ACIS-I at the aimpoint use FI CCD-specific conversions. The BI and FI specific conversions are more accurate for each type of CCD than the conversion used in previous cycles. The assumption is that it is desirable to have the most accurate gain conversion for the CCD on which the HRMA default aimpoint falls. Note that the conversion only impacts the on-orbit energy filtering. *Ground data processing will always apply the appropriate PHA to energy conversion.*

The observer should be aware that for observations which mix CCD types (*i.e.* both BI and FI CCDs on), the selected conversion (based on aimpoint as above), will nevertheless apply to all selected CCDs. This will not affect the observation if the low energy threshold for the energy filter (the "Event filter: Lower" parameter) is 0.5 keV or less, as the use of either conversion at these energies results in essentially no difference in the number of accepted events. However, for selection of a low energy threshold above 0.5 keV, the conversions are significantly different. Proposers who need an energy filter lower limit above 0.5 keV are encouraged to contact the *CXC* HelpDesk (https://cxc.harvard.edu/helpdesk/) to discuss their plans with an instrument scientist.

Spatial Windows A more sophisticated approach to removing data from the telemetry stream, and thus avoiding telemetry saturation, is by the use of a Spatial Window. This option offers a good deal of flexibility. One may define up to 6 Spatial Windows per CCD. Each window can be placed anywhere on the chip. Note there is a significant difference between a Spatial Window and a Subarray (Section 6.13.1): Subarrays affect the transmission of CCD data to the on-board ACIS processors; Spatial Windows select events detected by the processors and only telemeters a filtered subset of events. The user may also specify the window energy threshold and energy range.

Spatial windows can specify the sample rate for events inside them. A sample rate of 0 excludes all events; the default rate of 1 includes all events; a rate of $n > 1$ telemeters one out of every n events in the window. As one example of a use case, a spatial window could be used to eliminate a bright, off-axis source that would otherwise overwhelm the telemetry stream. The order in which the spatial windows are specified is important if they overlap. The first specified window which includes a given pixel will be applied to events at that pixel.

6.22.3 Non-ACIS Parameters Relevant to an Observation with ACIS

There are a small number of additional parameters that need to be considered in specifying an observation with ACIS: (1) the off-axis pointing (if required), which reduces the flux, and spreads out the image; (2) the roll angle (Chapter 4); (3) time constraints (if any); and (4) time monitoring intervals (if any).

6.22.4 Choosing CC Mode for Bright Source Observation

The Continuous Clocking (CC) mode (see Section 6.13.3) can be used, with or without gratings, to mitigate pile-up of very bright sources (see Section 6.16.2). CC mode has two standard telemetry formats: "faint" and "graded." Faint mode retains the event-island information needed for the application of the CTI correction but requires a larger data volume per event compared to graded mode. Thus, graded mode is less likely to saturate telemetry but spectral features needing CTI correction cannot be recovered. Note, however, that gratings can be used to further reduce throughput while also preserving discrete spectral structures such as emission and absorption lines and edges in the dispersed spectrum regardless of telemetry format.

The calibrations of TE and CC modes are very similar and only small differences are present with respect to ACIS gain, response, and integrated grade distribution, with CC faint mode versus CC graded mode choice resulting in less than 3% differences at the most. Starting in 2009-Nov, the CC faint and graded modes have been altered to include some of the flight grades that were previously rejected on-board: ACIS now telemeters all flight grades except 24, 107, 127, 214, 223, 248, 251, 254, and 255. It is recommended to use graded mode only in extreme cases.

See Section 8.5.2 in Chapter 8 for for recommended ACIS modes, CCD choices, sub-arrays, *etc*; that section also discusses some aspects of analyzing the data.

6.22.5 Warm ACIS Observations

Proposers should consider the following options to ensure the lowest possible focal plane temperature during their observation:

1. minimizing the number of CCDs so as to reduce the potential for thermally-induced gain drifts that might broaden lines over long observations

2. locating targets near the frame store regions (low CHIPY) using offset pointing (see Chapter 3) to maximize energy resolution (see Figures 6.14 and 6.15) and/or

3. using the BI chip S3, which has a more spatially-uniform energy resolution across the CCD and better energy resolution than the FI chips at high CHIPY values.

Warm FP temperatures do not affect ACIS imaging capabilities. Thus, proposed science goals that do not require precision spectroscopy need not be concerned with FP temperature. Since 2020, Mission Planners have used the user-supplied source count information to help identify proposed observations that do not have high precision spectroscopy as a science goal and thus may be observed at warmer focal plane temperatures, which can reduce the need to split an observation.

For observations using HETG (or LETG) and ACIS the impact of higher FP temperatures is insignificant since the energy calibration is based on the position of the event in the dispersed spectrum. Accordingly, HETG observations may also be conducted at warm FP temperatures. (Note, however, imaging spectroscopy for the zero-order will be affected by warmer temperatures.)

Chapter 7

HRC: High Resolution Camera

7.1 Status of the HRC Detectors

The voltages of both HRC-I and HRC-S were increased by one voltage step in 2024-Sep. The HRC-S experienced an anomaly in late 2025-Oct that has resulted in a temporary hiatus in observations. The HRC-I is operating nominally. However, due to ongoing thermal operational constraints, HRC observations are currently limited to durations of up to 14.5 ks in length. As with other spacecraft operational constraints, this maximum duration may depend upon external factors such as target pointing direction, and may be subject to further review.

7.2 Science Use Case

The High Resolution Camera (HRC) is a microchannel plate (MCP) instrument composed of two detectors, one optimized for imaging (HRC-I), and one (HRC-S) which serves as a read-out for the Low Energy Transmission Grating (LETG) discussed in Chapter 9. The HRC-I provides the largest field-of-view ($\sim$ 30 arcmin $\times$ 30 arcmin) of any detector aboard *Chandra*, and its response extends to energies below the sensitivity of ACIS (Chapter 6), albeit without comparable spectral resolution in imaging mode—although it can be used with both LETG and HETG to obtain high-resolution spectral data at $\lesssim$ 1.5 keV. The time resolution of the HRC detectors (16 μsec) is the best on the observatory, but can only be utilized under certain conditions as discussed in Section 7.12.

In imaging mode, the HRC detectors have often been used to obtain high timing resolution data from pulsars and for locating strong sources without the hindrance of pileup effects. The locations of events on the detector are not pixelated, leading to a higher fidelity description of the Point Spread Function at sub-arcsec resolution; this property is advantageous for high-spatial resolution and deconvolution studies, such as in detecting motions of jet knots or the proper motions of compact objects in supernova remnants. The long mission timeline of *Chandra* is also advantageous in this regard, as small proper motions become more easily detectable with the longer available baselines.

The low energy sensitivity allows studies of charge exchange emission from solar system objects and of coronal emission from stars, which is consequential for studies of XUV fluence on the habitability of exoplanets.

The spectral resolution of the bare detectors is adequate to distinguish between hard and soft sources, and the spatial structure of the UVIS filter on the HRC-S has been used to obtain rudimentary information on the low-energy spectra from stars by placing a source under the thick and thin filter areas of the detector (see below) and comparing their count rates. With the LETG, the HRC-S allows observations of low-energy spectra up to ≈180 Å, which allows study of ultra-soft extragalactic sources, white dwarfs, and neutron stars. The Fe-L shell emission at 15–17 Å, the O `VIII` resonance line, the density sensitive O `VII` triplet at 21.6-22.1 Å, and absorption lines from several species are also accessible at high spectral resolution, also with HETGused with HRC-I.

The particle background in the HRC is inversely correlated with the strength of the solar activity cycle; with the cycle approaching its peak now, the background will be at its lowest levels since c.2015, and so targeting of fainter sources will be beneficial.

7.3 Instrument Details

A schematic of the HRC layout is shown in Figure 7.1, and a summary of the characteristics is given in Table 7.1. A cross-section of the HRC-S layout and the relationships to the optical axis and the LETG Rowland circle are shown in Figure 7.2.

The HRC is a direct descendant of the Einstein (Giacconi *et al.* 1979) and ROSAT High Resolution Imagers (HRIs) (David *et al.* 1996). The ROSAT HRI had the same coating (CsI) as the HRC.

The Instrument Principal Investigator is Dr. Ralph Kraft of the Smithsonian Astrophysical Observatory.

7.3.1 Basic Principles

Figure 7.3 illustrates the features of the HRC MCPs. X-rays enter through a UV/Ion shield, necessary to reduce/avoid signals from UV light, ions, and low energy electrons. Most of these X-rays are then absorbed in the CsI-coated walls of the first (input) of two consecutive MCPs. The axes of the millions of tubes that comprise the input and output MCPs are not parallel to the optical axis but are canted ("biased") at an angle of 6° in opposite directions as shown in Figure 7.3. This bias improves the probability of an interaction. The CsI coating enhances the photoemission over that from a bare MCP. The resulting photoelectrons are then accelerated by an applied electric field. The next interaction with the walls releases several secondary electrons and so on, until a cascade of electrons is produced.

One purpose of the second (output) MCP is to provide additional gain. In addition, reversing the direction of the second MCP's bias angle with respect to the first removes a clear path for positive ions, and hence reduces the possibility of (positive) ion feed-

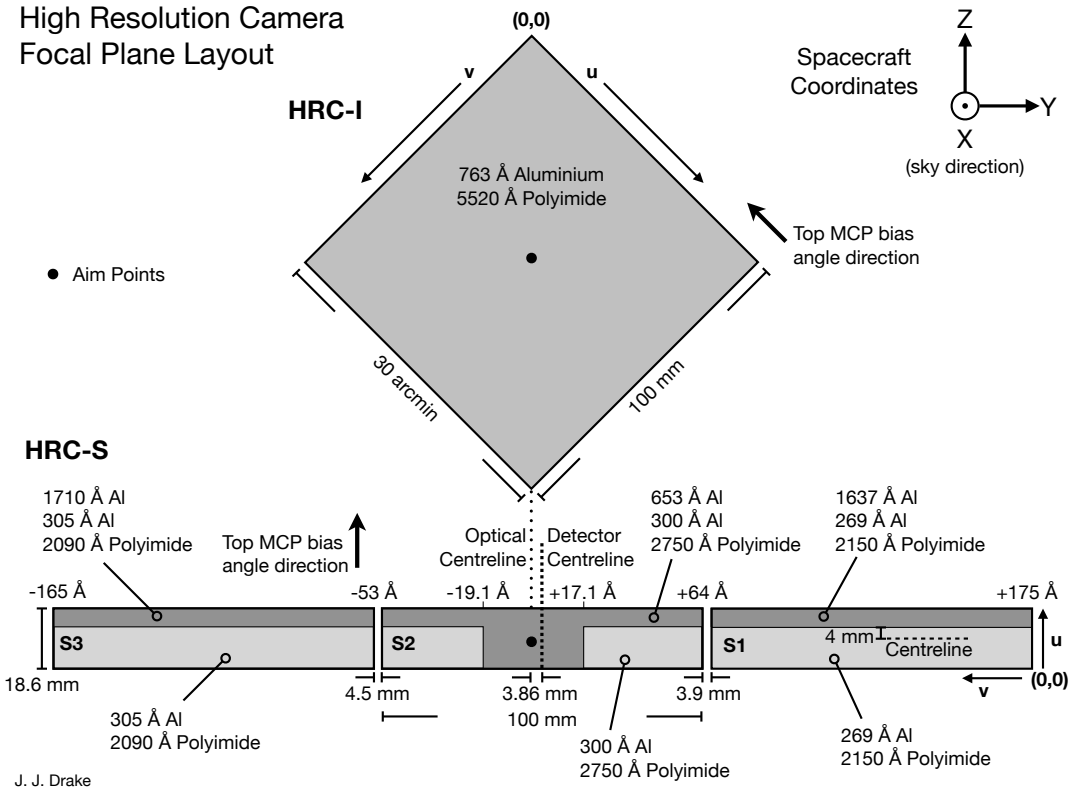

Figure 7.1: A schematic of the HRC focal-plane geometry as viewed along the optical axis from the telescope towards the focal plane. See https://hea-www.harvard.edu/HRC/calib/hrccalib_a_180998.ps

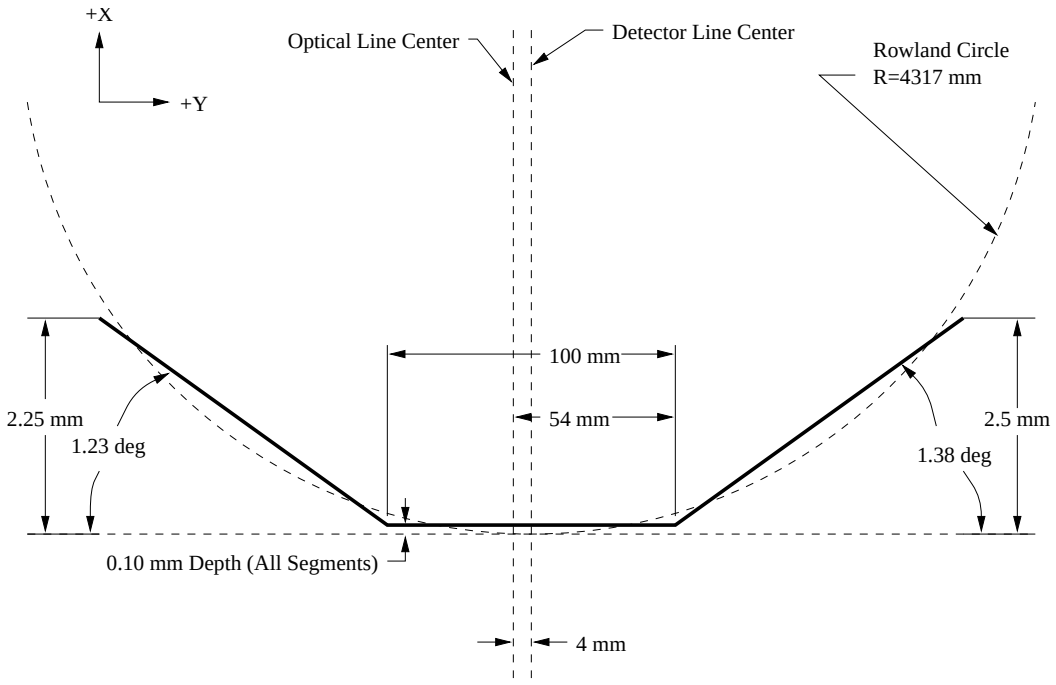

Figure 7.2: A schematic cross-section of the HRC-S MCP (not to scale). The HRC-S is shifted 0.1 mm forward of the tangent plane so the Rowland circle intersects each segment at two points.

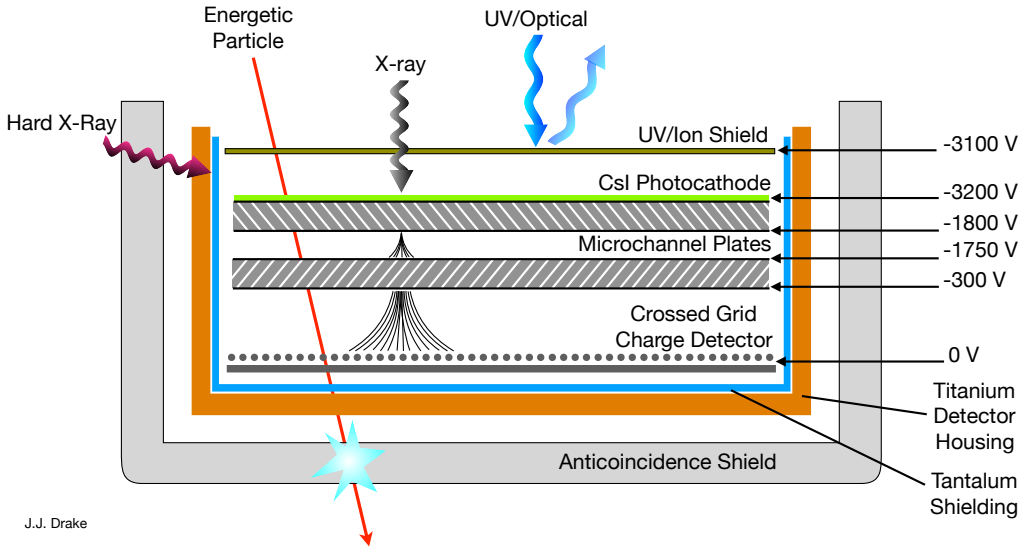

Figure 7.3: A schematic of the HRC microchannel plate detector.

Table 7.1: HRC Parameters

Focal-Plane Arrays		
HRC-I:	CsI-coated MCP pair	90×90 mm coated (93×93 mm open)
HRC-S:	CsI-coated MCPpairs	3-100×20 mm
Field of view	HRC-I:	$\sim 30 \times 30$ arcmin
	HRC-S:	6×99 arcmin
MCP Bias angle:		6°
UV/Ion Shields:		
	HRC-I:	5520 Å Polyimide, 763 Å Al
	HRC-S:	
	Inner segment (S2)	2750 Å Polyimide, 300 Å Al
	Inner segment "T" (S2)	2750 Å Polyimide, 953 Å Al
	Outer segment (S1; +1)	2150 Å Polyimide, 269 Å Al
	Outer segment (LESF) (S1; +1)	2150 Å Polyimide, 1906 Å Al
	Outer segment (S3; −1)	2090 Å Polyimide, 305 Å Al
	Outer segment (LESF) (S3; −1)	2090 Å Polyimide, 2015 Å Al
Spatial resolution	FWHM	$\sim 20\mu$m, ~ 0.4 arcsec
	HRC-I: pore size	10μm
	HRC-S: pore size	12.5μm
	HRC-I: pore spacing	12.5μm
	HRC-S: pore spacing	15μm
	pixel size (electronic read-out)	6.42938μm [0.13175 arcsec pixel^{-1}]
	pixel size (default binning size)	0.1318 arcsec pixel^{-1}
Energy range:		0.08–10.0 keV
Spectral resolution	$\Delta E/E$	~ 1 @1keV
MCP Quantum efficiency		30% @ 1.0 keV 10% @ 8.0 keV
On-Axis Effective Area:	HRC-I, @ 0.277 keV	100 cm^2
	HRC-I, @ 1 keV	193 cm^2
	HRC-S, @ 0.277 keV	85 cm^2
	HRC-S, @ 1 keV	194 cm^2
Time resolution		16 μsec (see Section 7.12)
Expected Quiescent background during Cycle 22 in level 2 data (see Sec 7.11)	HRC-I	3.5×10^{-5} cts s^{-1} arcsec^{-2}
	HRC-S	7.4×10^{-5} cts s^{-1} arcsec^{-2}
Intrinsic deadtime		50 μs
Constraints:	telemetry limit	184 cts s^{-1}
	maximum counts per aimpoint source	450000 cts
	linearity limit (on-axis point source)	
	HRC-I	~ 5 cts s^{-1} (2 cts s^{-1} pore^{-1})
	HRC-S	~ 25 cts s^{-1} (10 cts s^{-1} pore^{-1})

back—where an accelerated ion moving in the opposite direction as that of the electrons ends up causing the release of electrons and starts the process all over again.

The electron cloud—typically about 2×10^7 electrons per photon—that emerges from the output MCP is accelerated towards a position-sensitive charge detector. The HRC employs two types of charge detectors: the HRC-I uses a crossed grid charge detector, while the HRC-S uses a hybrid where one axis is composed of wires and the other has gold lines deposited on a ceramic substrate. Adjacent wires (or lines) are resistively connected and every eighth wire is attached to a charge-sensitive amplifier, referred to as a "tap", as illustrated in Figure 7.4.

The X-ray position is determined by calculating the centroid of the charge cloud exiting the rear MCP via the "three tap algorithm." In short, the three tap algorithm determines the charge cloud centroid using a combination of digital and analog electronics and off-line processing. Fast discriminators and logic circuits first determine a "coarse" position, which is based on the amplifier with maximum detected charge. Analog switches then select the three amplifiers centered on that coarse position and steer them to analog-to-digital converters. The coarse position and three digitized values are then telemetered to the ground and used off-line to calculate the event position. This process is performed for each axis. The reconstructed X-ray position can then be written as the sum of a coarse position and a charge centroid term centered on the coarse position:

$$pos = cp_i + \left(\frac{Q_{cp_{i+1}} - Q_{cp_{i-1}}}{Q_{cp_{i-1}} + Q_{cp_i} + Q_{cp_{i+1}}}\right) \times \Delta \tag{7.1}$$

where cp is the coarse position, $Q_{cp_{i+1}}$ is the charge measured on the cp_{i+1} tap, and Δ is the distance between taps. Since the charge cloud extends beyond the two outer taps, each of the outer amplifiers underestimates the amount of charge needed to calculate the true centroid. For an event perfectly centered on the middle tap, the amount of charge missed by the two outer taps cancel in the equation. If however, the event position is not over the center of a tap, the fractional amount of missing charge is different and produces a small systematic error in the reconstructed position. The small systematic positional error combined with the coarse position logic produce "gaps" in the HRC images. These gaps are perfectly aligned with the detector axes and correspond to positions exactly half-way between amplifier taps. The gaps are systematic and are removed in data processing.

The three-tap position algorithm described above can be improved upon by making use of the predictability of the shape of the charge cloud exiting the rear MCP. The spatial distribution of the charge cloud leaving the rear of the second MCP has a very specific shape for X-ray induced events. This shape has often been modeled as the combination of a Gaussian and a Lorentzian distribution. Due to this specific shape, it has been observed and simulated via Monte Carlo techniques that the fine position term:

$$\left(\frac{Q_{cp_{i+1}} - Q_{cp_{i-1}}}{Q_{cp_{i-1}} + Q_{cp_i} + Q_{cp_{i+1}}}\right) \tag{7.2}$$

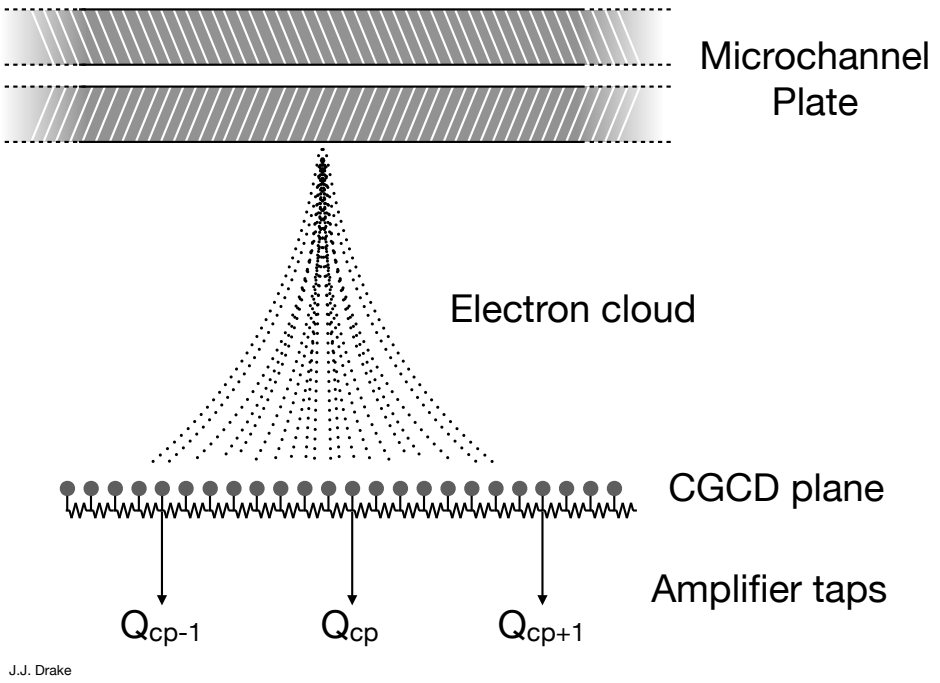

Figure 7.4: Schematic representation of event position determination for one axis of the crossed grid charge detector (CGCD). The electron cloud is divided between several amplifiers. The position of the event relative to the central coarse position, f_p, is calculated from the difference between the signals on either side of th e coarse position p (Q_{cp-1}, Q_{cp+1}), divided by the sum of the three signals, $f_p = \frac{Q_{cp-1}-Q_{cp+1}}{Q_{cp-1}+Q_{cp}+Q_{cp+1}}$.

and the complementary term:

$$(\frac{Q_{cp_i}}{Q_{cp_{i-1}} + Q_{cp_i} + Q_{cp_{i+1}}}) \tag{7.3}$$

are highly correlated. In fact, a scatter plot of these two quantities for X-ray induced events closely describes a hyperbola. Non X-ray events, primarily those due to the passage of charged particles, produce charge distributions that are often larger and more spatially extended and complex. As such, it is possible to remove many non-X-ray background events by filtering out those events that do not fit the hyperbola. Furthermore, since the charge distribution is centrally peaked, the complement Q_{cp_i} term is larger and less susceptible to noise-induced errors than the $Q_{cp_{i+1}} - Q_{cp_{i-1}}$ difference term. It is therefore possible to use the complement term, and the best fit hyperbolic locus, to correct those events where instrumental noise has compromised the three-tap fine position. A much more detailed explanation of this technique is presented in Murray, *et al.* (2000).

For more details concerning the HRC see Murray & Chappell (1989) and Zombeck *et al.* (1995).

7.3.2 Aimpoints

The aimpoints are the positions on the instrument where the flux from a point source with no commanded offsets is placed. Note that unlike ACIS, where the aimpoint position is offset by $\approx 10 - 20$ arcsec from the optical axis[1], there are no offsets for the HRC. Offsets may be set for spectroscopic observations on the HRC-S (see Section 9.3.1). There are two nominal aimpoints (see Table 4.3) as indicated in Figure 7.1 - one at the approximate center of the HRC-I, and the other slightly off-center on HRC-S. The HRC-S aimpoint Z-offset places the LETG-dispersed image along the centerline of the thinner part of the UV/Ion Shield (the two white rectangles in the diagram; see Section 7.9). The HRC-S aimpoint Y-offset is slightly off-center, so that the boundaries between the three HRC-S segments correspond to different wavelengths of the grating-dispersed spectrum (see Chapter 9 for details).

7.3.3 Drift Correction

The Chandra aimpoint is known to drift relative to the optical axis by 10-25 arcsec (Sec 5.4.3; see also Sec 4.5). This drift has been attributed to temperature fluctuations in the Aspect Camera Assembly. As a result, the nominal aimpoint is offset from the optical axis by the amount of the drift. This offset has no effect on the accuracy of the astrometry. Nominal aimpoints are expressed in terms of a permanent default aimpoint with an associated error box, as explained in 4.5. An automatic offset correction is applied to all ACIS observations to position the nominal pointing near the permanent default aimpoint (see 6.11). However, such corrections are not applied to HRC observations, since

[1] See `https://cxc.harvard.edu/cal/Hrma/OpticalAxisAndAimpoint.html`

the HRC does not have instrumental structures like node boundaries or chip gaps that require careful positioning of the source on the detector. Displacement of HRC pointings due to uncorrected aimpoint drift is comparable to the magnitude of aspect dither, and has a negligible effect on the measured rate and the shape of the PSF.

7.4 Shutters

Attached to the HRC are two mechanical blades that serve as shutters. These shutters were used to block out portions of the incident flux to aid in focusing the HRC. The blade position settings are variable and were designed to allow one to block the zero-order image of a grating observation. Currently only one blade is functional, and we do not offer use of this shutter as an observing option.

7.5 Dither

The spacecraft is dithered during all observations in a Lissajous figure. For observations with the HRC, the dither amplitude is 40 arcsec peak-to-peak, with nominal periods of 1087 (in Y) and 768 (in Z) sec. Dithering serves to average out pixel-to-pixel variations in the detector response. It also ameliorates gaps in spectral coverage with the LETG/HRC-S combination caused by the HRC-S intersegment spaces near -50 Å and $+60$ Å (see Figure 7.1). Large amplitude dither has also been used to carry out low-resolution differential filter photometry across the thick/thin filter boundary on the HRC-S (Drake et al. 2020). Observers who need to use non-nominal dither amplitudes should consult with the instrument team. The effects of dither are corrected for during ground processing.

7.6 Spatial Resolution & Encircled Energy

Imaging with the HRC is best performed with the HRC-I because of the much lower background (Section 7.11) and larger field of view. The intrinsic PSF of the HRC is well modeled by a Gaussian with a FWHM of ~ 20 μm (~ 0.4 arcsec). The HRC pixels, determined by the electronic read-out (*not* the pore size), are 6.42938 μm (0.13175 arcsec). The HRC response is thus well matched to the intrinsic HRMA resolution (Chapter 4).

Approximately 90% of the encircled energy lies within a 14 pixel diameter region (1.8 arcsec) from the center pixel for the observation of AR Lac shown in Figure 7.5. The measured PSF is as good or better than the simulations because a very conservative pre-flight estimate of the aspect solution was used in the simulations.

Deconvolution of aimpoint AR Lac and Capella observations carried out at different parts of the detector show that an anomalous feature developed c.2003 (Juda & Karovska 2010). While initially suspected to be due to detector blur, it has since been verified to be present in ACIS data as well (Kashyap 2010). The *Chandra* PSF shows an unexplained enhancement in the profile at distances of ≈ 0.8 arcsec from the source centroid (see

Section 4.2.3). This anomaly is in excess of that expected from ray trace simulations, and is preferentially oriented towards the mirror spherical coordinate angle of $\phi = 285^\circ$ (see the CIAO caveats page https://cxc.harvard.edu/ciao/caveats/psf_artifact.html), and is approximately oriented towards the spacecraft +Z axis (see Figure 1.2). The asymmetry is illustrated for an HRC-I observation of an on-axis pointing of AR Lac in Figure 4.18, and the magnitude of the asymmetry is illustrated for a number of low–count rate on-axis point sources in Figure 4.20. Figure 7.6 depicts the anomaly for Capella observations carried out at different parts of the detector. The HRC read-out blurs event locations, and contributes an additional broadening of the HRMA PSF. Based on an analysis of transient hotspots, the intrinsic detector blur has been modeled as a combination of a Gaussian and an offset Beta profile, and is incorporated in the ray trace model.

We have constructed an empirical PSF by combining on-axis HRC-I observations of AR Lac (Kashyap 2025, Kashyap & Jerius 2016; https://cxc.harvard.edu/cal/Hrc/PSF/empPSF.html). This PSF minimizes blurring caused by detector effects by excluding events that are affected by tailgating (see Section 7.12) and excluding events with SAMP< 15, which are thought to have larger positional uncertainties.

The imaging resolution of the HRC-I degrades off-axis for two reasons: the HRMA PSF increases in size with increasing off-axis angle, and the deviation increases between the flat HRC-I detection surface and the curved HRMA focal surface. The off-axis imaging behavior of the HRC-I is shown in Figure 7.7. The nominal best-focus of the HRC-I is chosen to provide the best image quality in the center of the field-of-view.

7.7 Energy Resolution

7.7.1 Non-Dispersive Energy Resolution

The intrinsic energy resolution of the HRC is poor. Even though the pulse-height amplitude (PHA) of each event is telemetered, spectral fitting cannot be usefully carried out for sources observed with the HRC. For low-resolution spectra such as these, standard spectral analysis techniques or visualization strategies cannot be used. In particular, analysis must only be done in *channel* space, and plots made in *energy* or *wavelength* space will convey no useful information.

7.7.2 Dispersive Energy Resolution

The HRC-S is optimized for use with the Low Energy Transmission Grating to become the Low Energy Transmission Grating Spectrometer (LETGS; see Chapter 9). The HRC-I can also be used but, because the HRC-I surface is flat and does not follow the Rowland circle, sharp spectral features cannot be measured as well as on the HRC-S. Furthermore, the wavelength coverage of the HRC-I is smaller than that of the HRC-S because of its smaller lateral extent.

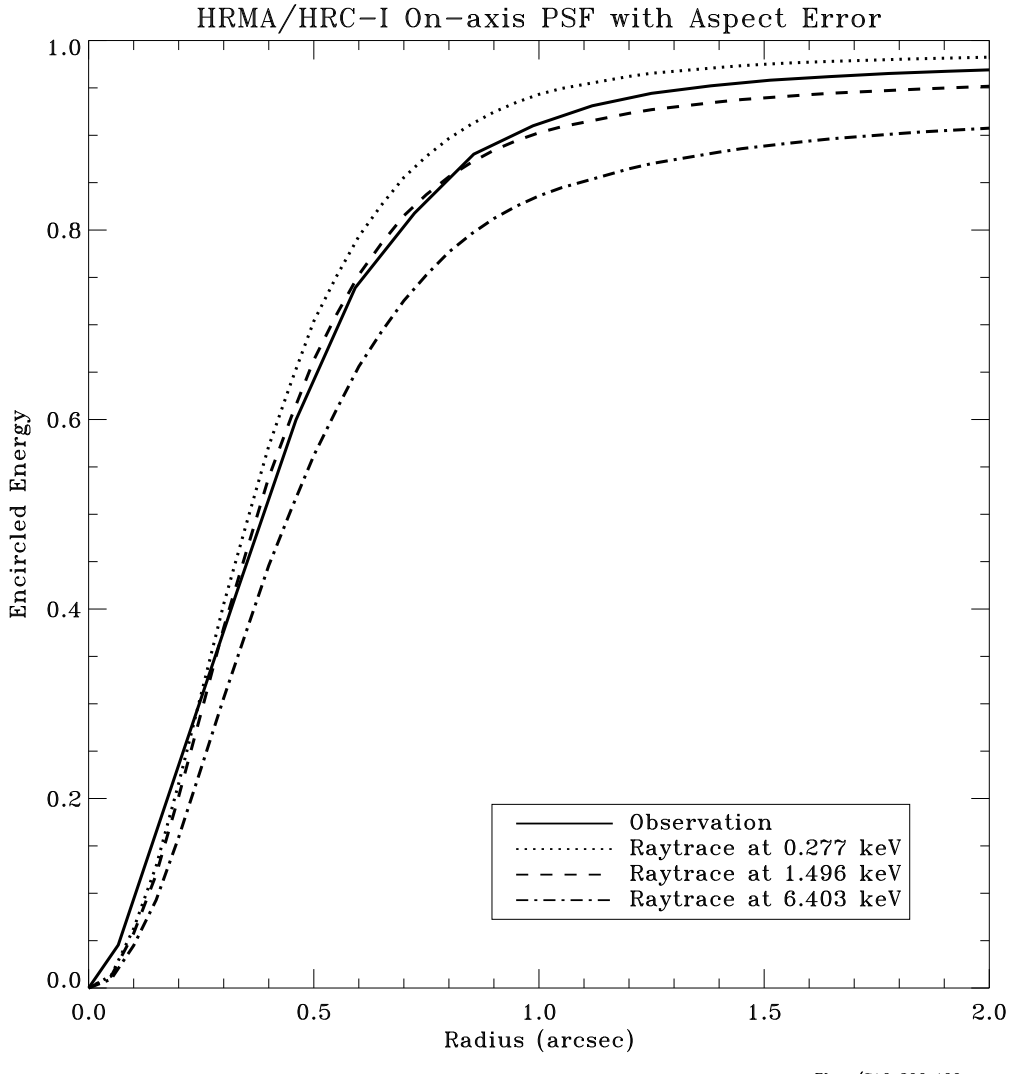

Figure 7.5: The predicted and observed fractional encircled energy as a function of radius for an on-axis source observed with the HRMA/HRC-I. The solid black curve shows the empirical PSF, based on accumulated AR Lac observations, representing the current best sharpness achievable. Raytrace models at different energies (dashed red for 0.28 keV, dash-dotted green for 1.5 keV, and dotted blue for ≈6.4 keV) are also shown for qualitative comparison. The raytraces incorporate HRC-I detector blur and aspect dither for a set of representative on-axis AR Lac observations but do not include other effects.

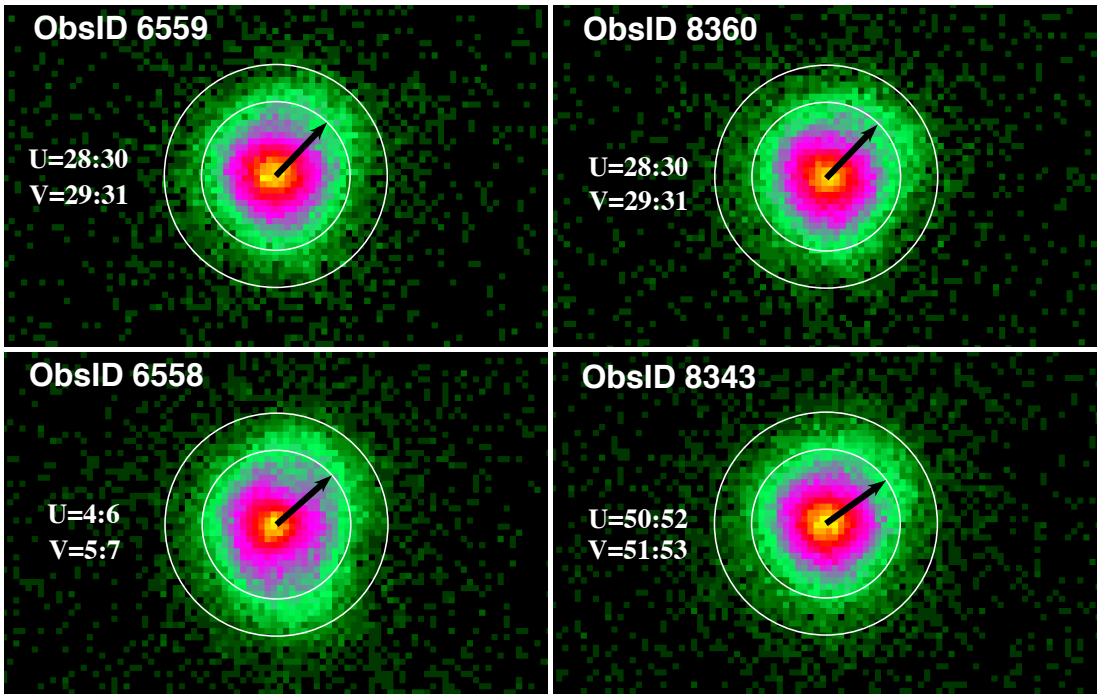

Figure 7.6: The PSF anomaly illustrated with HRC-I observations of Capella over different locations on the detector. Two observations close to the nominal aimpoint (ObsIDs 6559, 8360; top row) are shown along with two at the extreme ends of the detector (ObsIDs 6558, 8343; bottom row). Each panel is also labeled with the detector (U, V) coordinates that the sources span. The observations have slightly different roll angles, and the direction of the anomaly is indicated with an arrow of length 0.8 arcsec. An asymmetry is discernible within each of the annuli (centered on the sources and spanning 0.8–1.2 arcsec); there is an excess of counts in the direction of the anomaly. ObsID 6558 also shows an asymmetry pointing towards the bottom right, towards the detector $-U$ direction, which is due to residual ghost events (see Section 7.11).

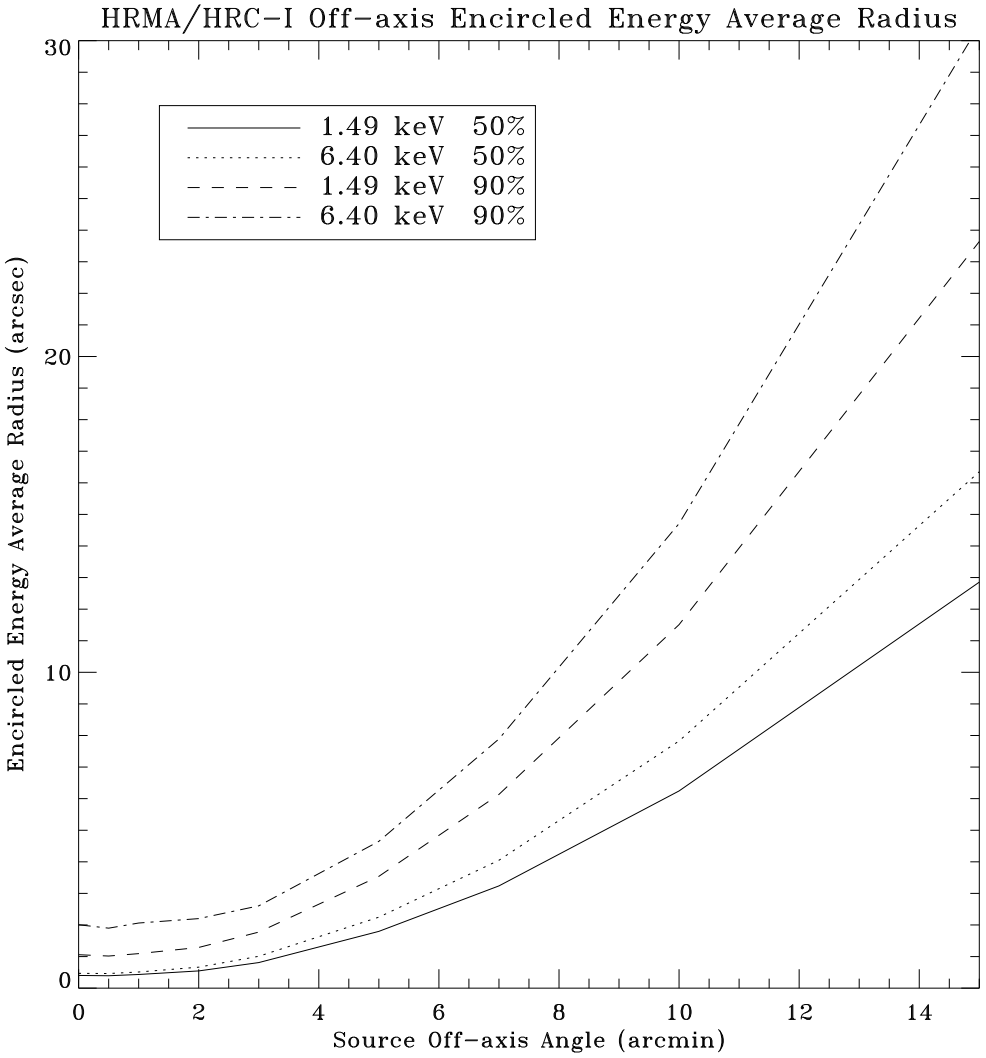

Figure 7.7: Encircled energy as a function of source off-axis angle for 50% and 90% encircled energy for 1.49 and 6.40 keV for the combined HRMA/HRC-I. A conservative contribution from the aspect solution is included (FWHM = 20 μm (0.41 arcsec)). A plot for the HRC-S would be almost identical since the PSFs of the two instruments are virtually identical.

The High Energy Transmission Grating (HETG) is unsupported in CIAO for use with HRC at this time. Nevertheless, observations have been carried out in the HETG+HRC-I configuration (see Chapter 8).

7.8 Gain Variations

There are significant gain variations present in both instruments (see Figure 7.8). Previously (prior to CIAO v4.2/CALDB v4.2), the gain was corrected such that spatial (for both HRC-I and HRC-S) and temporal variations (for HRC-I) were corrected using a multiplicative map that transformed PHA values to Pulse Invariant (PI) values that were comparable across epochs. After that, the gain maps were applied to the scaled sum of the amplifier signals (SAMP[2]) rather than PHA to generate a better behaved distribution (Wargelin 2008, Posson-Brown & Kashyap 2009). However, this is no longer a useful transformation. Indeed, since 2021-May, for HRC-S observations the gain map correction is not applied to PHA, and PI are set equal to SAMP; for HRC-I observations a gain map is applied to PHA, but this is solely to compute a nominal PI in order to maintain fidelity with older archived data and should not now be used for filtering. Because of the decline in gain, in recent years more events fall below the lower level discriminator (typically set at PHA=8, and set at PHA=21 in the HRC-S timing mode), leading to a loss of effective area of as much as 10–15%. As such, simple multiplicative corrections to the PHA or SAMP distribution do not restore the shape of the distribution to what it was early in the mission.

A more accurate correction may be computed by matching the observed SAMP>50 distribution to a template distribution for a source with a soft or hard spectrum—like HZ 43 or AR Lac—using a linear transformation of SAMP, i.e.

$$\mathrm{SAMP}_{\mathrm{Corrected}} = g_C + g_S \cdot \mathrm{SAMP}, \tag{7.4}$$

and then obtaining a temporal polynomial fit to describe g_C and g_S as functions of time (see, e.g., McEntee et al. 2022). For instance, transforming the profiles of the coronal source AR Lac in HRC-I to match that of ObsID 1385 over various epochs yields

$$\begin{aligned}
g_C^{\mathrm{AR\,Lac}} &= 20.1 + 0.374 \cdot Y - 0.0696 \cdot Y^2 + 4.255 \times 10^{-3} \cdot Y^3 \quad \text{for } 0 \le Y < 21.12 \\
&= 3588.4 - 324.36 \cdot Y + 7.377 \cdot Y^2 \quad \text{for } 21.12 \le Y < 23.5 \\
&= -122.66 + 6.894 \cdot Y \quad \text{for } 23.5 \le Y \\
g_S^{\mathrm{AR\,Lac}} &= 0.884 + 0.102 \cdot Y - 0.0117 \cdot Y^2 + 5.68 \times 10^{-4} \cdot Y^3 \quad \text{for } 0 \le Y < 21.12 \\
&= -18.3 + 1.4856 \cdot Y - 0.0242 \cdot Y^2 \quad \text{for } 21.2 \le Y < 23.5 \\
&= -5.053 + 0.355 \cdot Y \quad \text{for } 23.5 \le Y\,,
\end{aligned} \tag{7.5}$$

[2] $\mathrm{SAMP} = \frac{\mathrm{SUMAMPS}}{\text{scaling factor}}$, with $\mathrm{SUMAMPS} = \sum_{j=1,2,3} \sum_{X=U,V} \mathrm{A}X_j$, where $\mathrm{A}X_j$ are the amplifier tap signals and scaling factor = 148 and 128 for HRC-I and HRC-S, respectively.

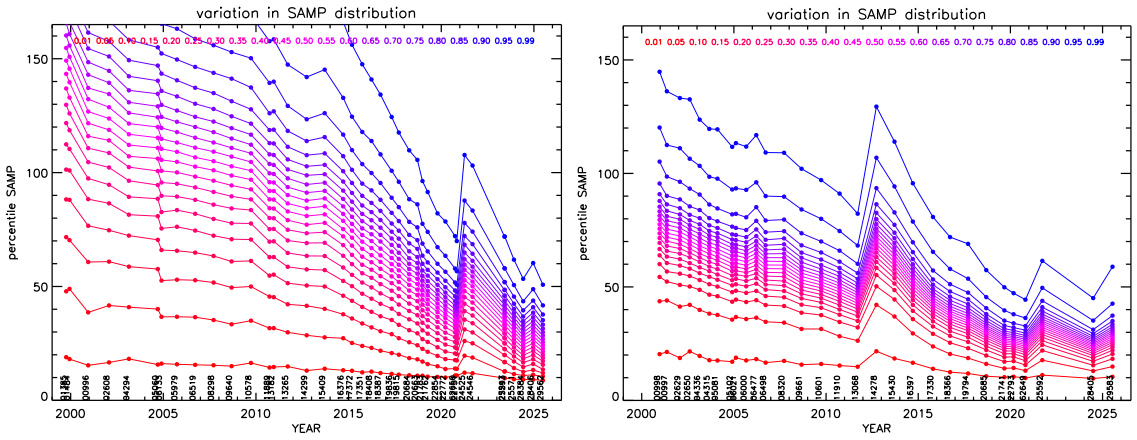

Figure 7.8: The evolution of the HRC SAMP distributions. The percentiles of the SAMP spectra of AR Lac are shown for both the HRC-I (left) and HRC-S (right) over the duration of the mission. Each line represents a percentile (color coded and listed along the top of the plots), with the corresponding ObsID marked along the bottom. The distribution inexorably narrows and shrinks to lower values of SAMP over time, except when the high voltage setting is increased.

where $Y = (\text{Decimal Year} - 2000)$. Similarly for the soft source HZ 43 in HRC-I to match ObsID 59,

$$
\begin{aligned}
g_C^{\text{HZ 43}} &= 0.032 - 0.362 \cdot Y + 0.0602 \cdot Y^2 + 2.096 \times 10^{-4} \cdot Y^3 \quad \text{for } 0 \le Y < 21.12 \\
&= 4156.8 - 376.04 \cdot Y + 8.524 \cdot Y^2 \quad \text{for } 21.12 \le Y < 23.5 \\
&= 37.1 - 0.557 \cdot Y \quad \text{for } 23.5 \le Y \\
g_S^{\text{HZ 43}} &= 0.973 + 0.0962 \cdot Y - 0.0116 \cdot Y^2 + 5.74 \times 10^{-4} \cdot Y^3 \quad \text{for } 0 \le Y < 21.12 \\
&= 109.4 - 9.853 \cdot Y + 0.227 \cdot Y^2 \quad \text{for } 21.2 \le Y < 23.5 \\
&= -8.131 + 0.489 \cdot Y \quad \text{for } 23.5 \le Y \,. \qquad (7.6)
\end{aligned}
$$

In order to restore the SAMP distributions to what they were earlier, and thus partially mitigate the QE loss that is occurring due to gain decline, the voltage on the HRC detectors has been increased. The HRC-S voltage was increased four times: 2012-Mar, 2021-May, 2023-Sep, and 2024-Jun; the voltage on the HRC-I was increased in 2021-Mar, 2023-Sep, and 2024-Jun. See also Chapter 9 for details. Since the power supply anomaly (see Section 7.12), the gain drop has continued and the corresponding QE loss is ≈10%. We therefore recommend carrying out gain profile corrections using methods as described in Equations 7.5 and 7.6 above.

7.9 UV/Ion Shields

The placement, composition, and thickness of the various UV/Ion shields (filters) are shown in Figure 7.1. Details of the UVIS transmission as a function of energy can be found at https://cxc.harvard.edu/cal/Hrc/detailed_info.html#uvis_trans.

The shields suppress out-of-band (outside the X-ray band) radiation from the ultraviolet through the visible. The detector response to out-of-band light for an object in its field-of-view is a possible source of unwanted signal. Suppressing out-of-band radiation is particularly important for observing sources which have bright EUV and UV fluxes. The HRC has strongly reduced sensitivity in this spectral region, as shown in Figure 7.9. As part of the in-flight calibration program the bright A star Vega (A0V, $U = 0.02$, $B = 0.03$, $V = 0.03$) was observed with both the HRC-I and HRC-S. The predicted count rate for HRC-I was 7×10^{-4} cts s^{-1}. From monitoring observations of Vega, an upper limit to the UV rate of 8×10^{-4} cts s^{-1} is calculated (Pease *et al.* 2005). The image of Vega was also placed on three regions of the HRC-S—the inner segment "T", the thin aluminum inner segment, and on one of the thin aluminum outer segments. The predicted count rates were 1, 400, and 2000 cts s^{-1} respectively. The corresponding observed rates were 0.2, 240, and 475 cts s^{-1}. Sirius was observed with the HRC-S/LETGS to obtain a soft X-ray spectrum of Sirius B (white dwarf) and Sirius A (A1V, $V = -1.46$, $B - V = 0.01$) was seen in zeroth order at about the expected count rate. Based on these sets of observations, the UV/Ion shields are performing as designed. Ongoing monitoring observations of Vega indicate no change in the UV response of HRC-I and HRC-S since launch. For a detailed discussion of the out-of-band response of the HRC to stars, see https://hea-www.harvard.edu/HRC/calib/palermopaper.ps, which allows one to determine the out-of-band count rate produced by a blackbody source with known T_{eff}, m_V, and N_H.

Scattered UV, far ultraviolet (FUV), and extreme ultraviolet (EUV) light from the Sun or the bright Earth may cause a background dependence on viewing geometry. The spacecraft was designed to limit the contribution from stray scattered radiation to 0.01 cts cm^{-2} s^{-1} (2.4×10^{-7} cts arcsec^{-2} s^{-1}) on the HRC. The imaged components of scattered radiation are dependent on the solar cycle, but are at most ~ 0.01 cts cm^{-2} s^{-1} for most lines of sight.

7.10 Quantum Efficiency and Effective Area

The efficiency of the HRC detector is the product of the appropriate UV/Ion shield transmission and the quantum efficiency of the CsI coated MCP. Pre-flight flat field measurements show a 10% variation in the efficiency across the HRC-I. The HRC-S also exhibits efficiency variations of the same magnitude, with the complex structure of the HRC-S UVIS contributing to the spatial variations. In-flight observations of the steady coronal source Capella show that the HRC-I variation is known to better than $\sim 2\%$ at high energies. There are time-dependent decreases in the QE for both HRC-I and HRC-S. The HRC-S

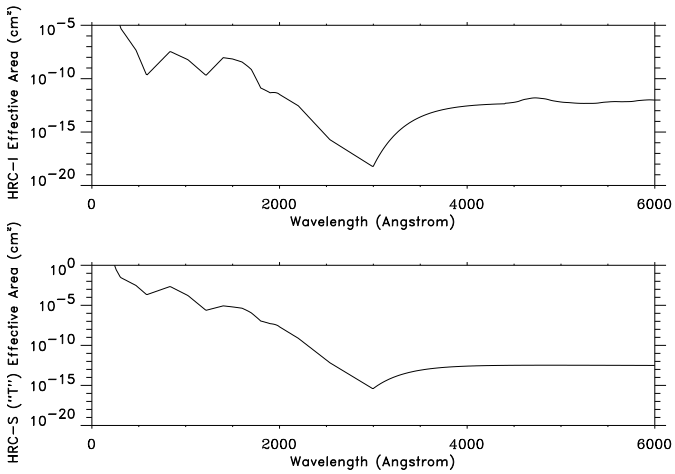

Figure 7.9: The HRC-I (top) and the center section of the HRC-S (bottom) UV/Ion shield effective area as a function of wavelength.

decline is wavelength independent and amounts to $\sim$ 15–20% over the course of the mission. These corrections are incorporated into the corresponding CALDB files. The HRC-I QE has been stable at short wavelengths, but has declined by $\sim$ 10% at long wavelengths based on measurements of the white dwarf HZ 43. The HRC-I also shows some fluctuations at low energies at large offset locations, $\sim$ 10 arcmin away from the nominal aimpoint (see https://cxc.harvard.edu/ccw/proceedings/2007/presentations/possonbrown3/).

The combined HRMA/HRC effective areas—the product of the HRMA effective area, the quantum efficiency of the HRC-I or the HRC-S and the transmission of the appropriate UV/Ion shield, integrated over the point spread function—are shown in Figure 7.10. Monitoring of the efficiency of both detectors is continuing. The charge extracted since launch has resulted in a decrease in gain in both detectors. (See https://cxc.harvard.edu/cal/Hrc/).

The HRC-S QE has been declining at an average rate of $\approx$2.4% per year. This decline has been steady, and is weakly dependent on wavelength, with higher energies affected less than lower energies. The QE decline in the HRC-I accelerated after c.2015, and reached $\approx$35% prior to the voltage increase; after the voltage increase, the QE relative to the beginning of the mission is estimated to be down $\approx$15% at soft energies and $\approx$5% at high energies. Similarly, the low-energy QE in the HRC-S is currently down 14% from launch. These effects are incorporated in the latest QE and quantum efficiency uniformity (QEU) maps in the current CALDB.

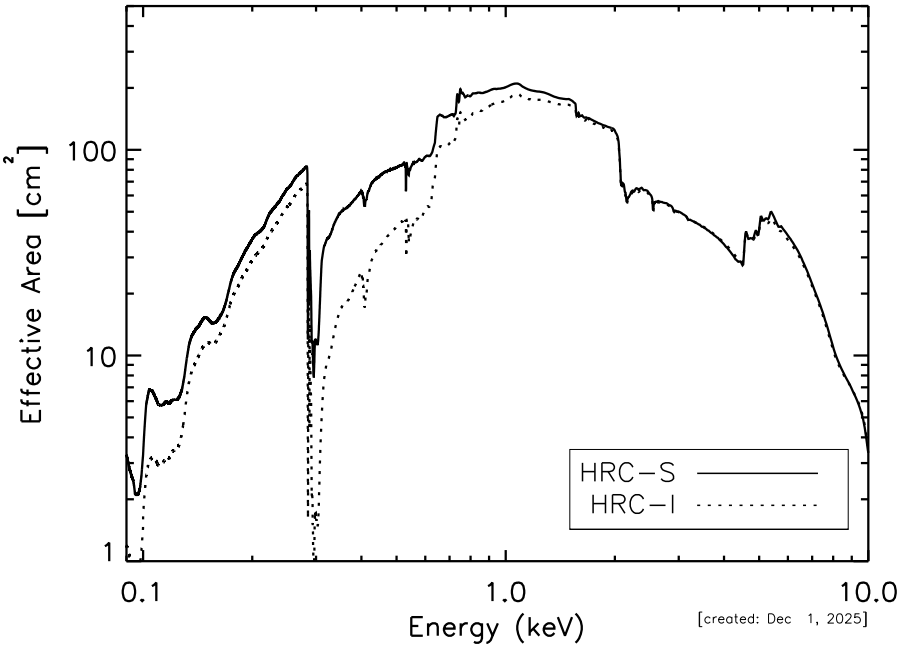

Figure 7.10: The effective area of the HRMA/HRC-I (dashed line) and the central segment of the HRMA/HRC-S in imaging mode (solid line) integrated over the full PSF. Absorption edges are due to the iridium coating on the mirrors, the CsI MCP coating, and the polyimide/Al of the UVIS.

7.11 On-Orbit Background

7.11.1 HRC-I

The HRC-I anti-coincidence shield reduces the on-orbit valid event rate by about a factor of 5 to ∼100-150 cts s^{-1} over the field; without on-board anti-coincidence vetoing the rate would greatly exceed the telemetry limit of 184 cts s^{-1}. After standard processing, the Level 2 event file background rate is $\approx 3.5 \times 10^{-5}$ cts s^{-1} $arcsec^{-2}$, though variations of 20% are possible. The background varies smoothly over the field with no more than a 10% difference between the center (lower) and edges (higher) of the detector. The background is not azimuthally symmetric (see Isobe & Juda 2009). Note, the total event rate remains unchanged, but detector events in coincidence with anti-coincidence events no longer enter the telemetry data stream. Before launch the expected rate, after vetoing the effects of cosmic rays, was 10-20 cts s^{-1} composed of mainly the internal rate of the MCPs (10-15 cts s^{-1}), and a small contribution from cosmic rays due to anti-coincidence shield inefficiency. The anti-coincidence shield inefficiency has increased in recent years, and the quiescent background has increased by ≈10%. For observations where there is a chance that the telemetry saturation limit might be hit, the HRC Instrument Team has devised a new instrument mode that excludes the events recorded in the two outermost taps from being telemetered. This mode can be selected in consultation with Chandra uplink support contact scientists when the observation is scheduled. There is additional background in the HRC-I that is not well understood. For point source detection and exposure times of 100 ks or less the background is virtually negligible. However, for extended low surface brightness objects this relatively low rate can become significant depending on the specific details of the source.

Ground-based filtering further reduces the non-X-ray background in the HRC detectors (see Murray *et al.* 2000, Juda *et al.* 2000 and Wargelin *et al.* 2001; https://cxc.harvard.edu/cal/Letg/Hrc_bg/). After filtering the non-X-ray background for HRC-I data is reduced by ∼ 40% while the corresponding reduction in X-ray events is less than a few percent. For the HRC-S, the non-X-ray background is decreased by ∼ 50% and the X-ray loss is 1–2%. Furthermore, filtering makes the spatial distribution of the detector background flatter. Filtering also removes saturated events responsible for faint secondary "ghost" images (see Section 7.12).

7.11.2 HRC-S

The anti-coincidence shield of the HRC-S does not work because of a timing error in the electronics. The error is not correctable. As a result the event rate is very high and exceeds the telemetry rate limit. To cope with this problem the HRC team has defined a "spectroscopy region" which is about 1/2 of the full width and extends along the full length of the HRC-S detector. The spectroscopy region (∼10 mm) is implemented using the edge blanking feature of the electronics. With this change, the telemetered quiescent background rate is about 120 cts s^{-1}.

The background can be further reduced in ground data processing by using pulse height filtering that preferentially selects X-rays over cosmic ray events. A reduction in background by a factor of about three is possible for dispersed spectra. Thus there are two relevant background rates for the HRC-S: a telemetry rate of 120 cts s^{-1} and a post-processing rate for calculating signal-to-noise. The latter is discussed in detail in Section 9.3.6 (see especially Figure 9.22).

7.11.3 Temporally Variable Background

Both the HRC-I and HRC-S experience occasional fluctuations in the background due to charged particles. These periods of enhanced background are typically short (a few minutes to a few tens of minutes) and are anywhere from a factor of two to ten over the quiescent rates. The increased background appears to be uniformly distributed over the detector and introduces no apparent image artifacts. On average it seems that no more than about 20% of the observing time is affected by these events, and they are easily recognized in the secondary science rate data and so can be filtered out if desired. An example of this behavior is shown in Figure 7.11. See Juda *et al.* (2002) for more information on the HRC background.

The background increases due to increased cosmic ray flux during solar cycle minima (see Figure 7.12). The difference between cycle minimum and maximum can be as much as a factor of two. On-orbit non-sky background spectra and events lists are available in CALDB for analysis and modeling. The spectra are obtained during yearly AR Lac observations (see Figure 7.13). A recipe for their use in filtering on PI to reduce background is described in the CIAO thread https://cxc.cfa.harvard.edu/ciao/threads/hrci_bg_spectra/.

Furthermore, much of the background can be alleviated by filtering out events with PI < 20 and PI > 350 on the HRC-I. The background is reduced by $\sim 20\%$ even as only $\approx$1–2% of the source counts are lost (see Figures 7.14 - 7.16). The CIAO thread https://cxc.harvard.edu/ciao/threads/hrci_bg_spectra/ describes how to perform this estimate for different source models background spectra. This thread can also be followed to compute background reduction factors for non-grating HRC-S sources, provided that they are extended or are observed off-axis, and user generated background spectra are used. This approach for improving signal-to-noise is not recommended for sources observed on-axis with the HRC-S, since gain near the aimpoint varies significantly on very small scales and is not well calibrated, potentially leading to undesired filtering effects. The approach should also not be used for data obtained with a grating in place; see Section 9.3.6 for a discussion of background reduction using PI filtering for dispersed spectra.

7.11.4 Limiting Sensitivity

The limiting sensitivity for the detection of a source is dependent on both the background count rate and the estimated counts from the source. In Tables 7.2 and 7.3, we compute the source counts for on-axis point sources with different spectral shapes and estimate the flux that would be detected at 3σ in a 50 ksec observation. The source counts are

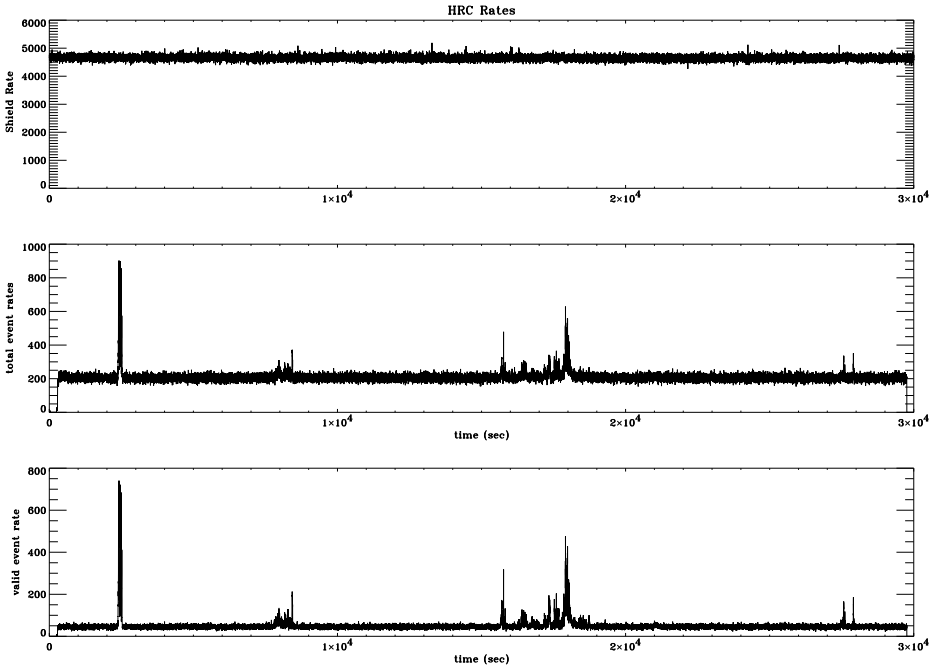

Figure 7.11: An example of the background variability during a ∼30 ks HRC-I observation of the supernova remnant (SNR) SNR G21.5-09 taken on 1999-Oct-25. The total event rate (middle) and valid event rate (bottom) show correlated bursts up to ∼800 cts s^{-1}. The bursts are uniformly distributed over the detector. The anti-coincidence shield (top) exhibits no correlated enhancements. The total and valid rates differ by ∼200 cts s^{-1} due primarily to cosmic ray events that are vetoed and don't appear as valid events in the telemetry.

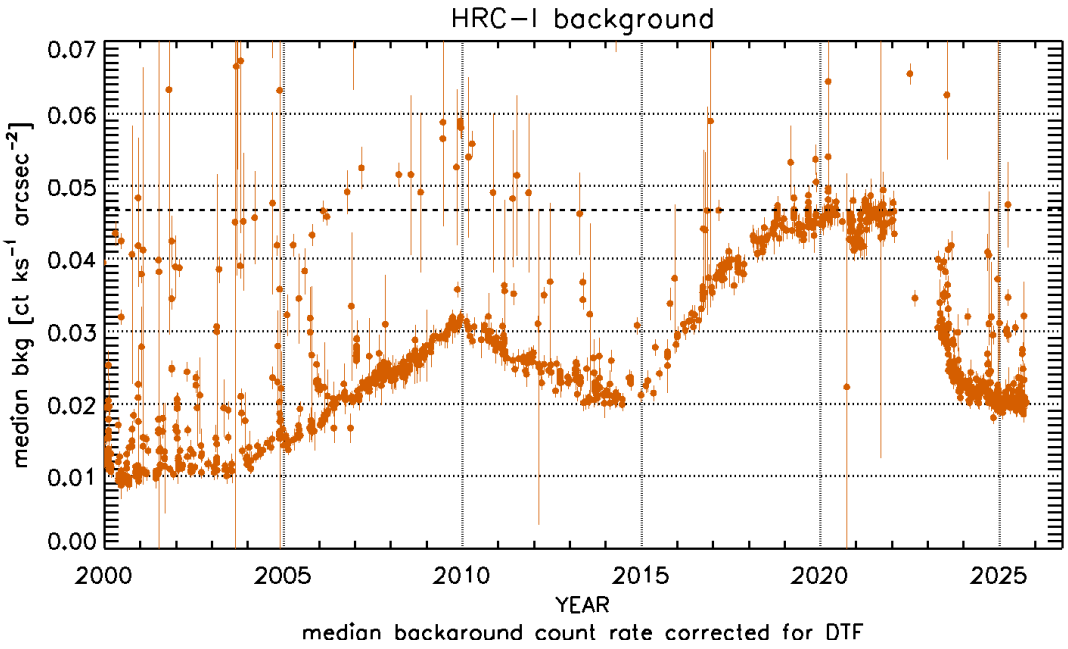

Figure 7.12: The change in HRC-I background rate with time. The deadtime corrected detector averaged counts per ksec per square arcsec is shown with error bars for each HRC-I observation. The lower envelope represents the quiescent rate, and the fluctuations are due to the occasional background flaring. The HRC-S data show a similar behavior of anti-correlation with solar activity (see Figure 9.21)

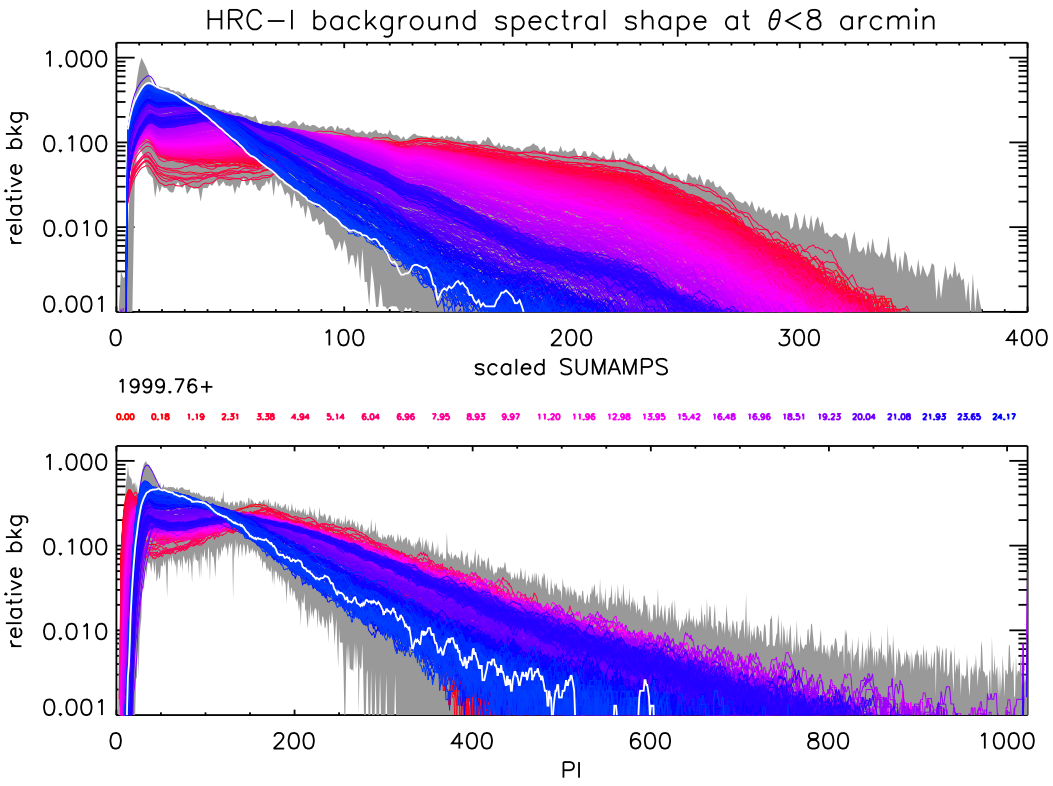

Figure 7.13: The shape of the background spectrum varies with time on both the HRC-I and HRC-S. This plot shows the HRC-I spectra of both the SAMP (top panel) and Pulse Invariant channels (PI; bottom panel), color coded by years into the mission (red is early, blue is current). The spectral shapes are obtained from the central 8 arcmin region of the HRC-I during regular AR Lac calibration observations, with the source photons and observations affected by high background excluded. The gray shaded region represents the full envelope, and the individual curves are the smoothed spectra, normalized by total counts and rescaled to a sample maximum of 1. The white curves represent the latest observation.

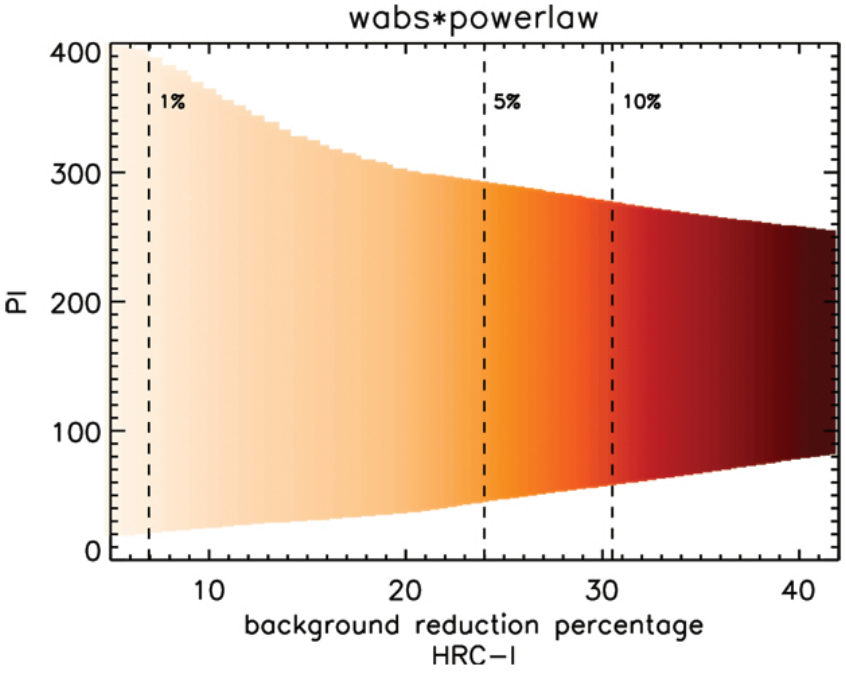

Figure 7.14: An estimate of the range of PI values that should be included to reduce the background by a given percentage is shown in the figure as a shaded band for sources with absorbed power-law spectra. (see `https://cxc.cfa.harvard.edu/ciao/threads/hrci_bg_spectra/`). The depth of the shading indicates how much of the source events are expected to be lost. The vertical dashed lines indicate the background reduction for source event losses of 1%, 5%, and 10%. As shown in Figure 7.13, the background varies with time. Here, for the sake of definiteness, we use the background from the year 2008.

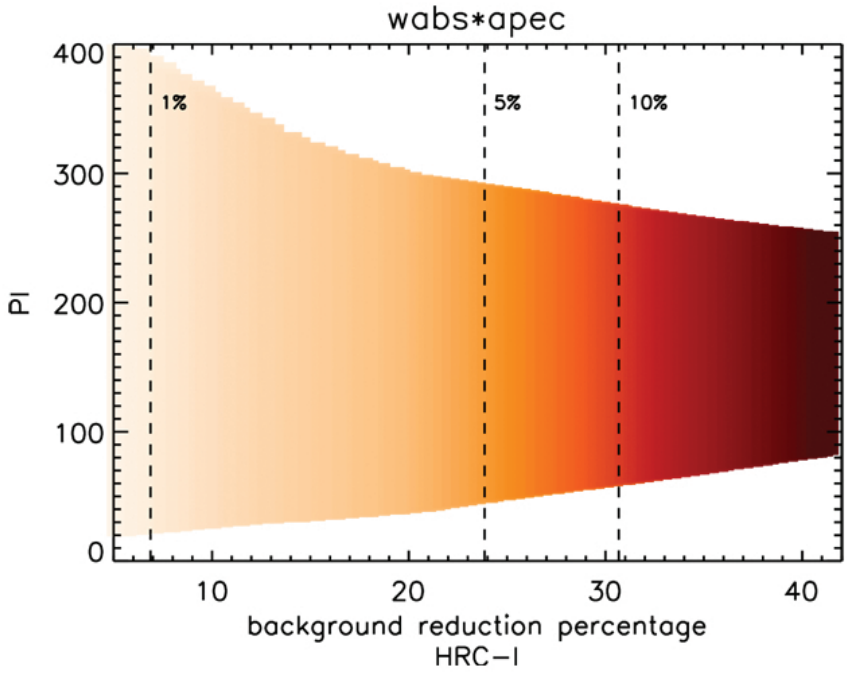

Figure 7.15: As Figure 7.14, but for sources with thermal spectra.

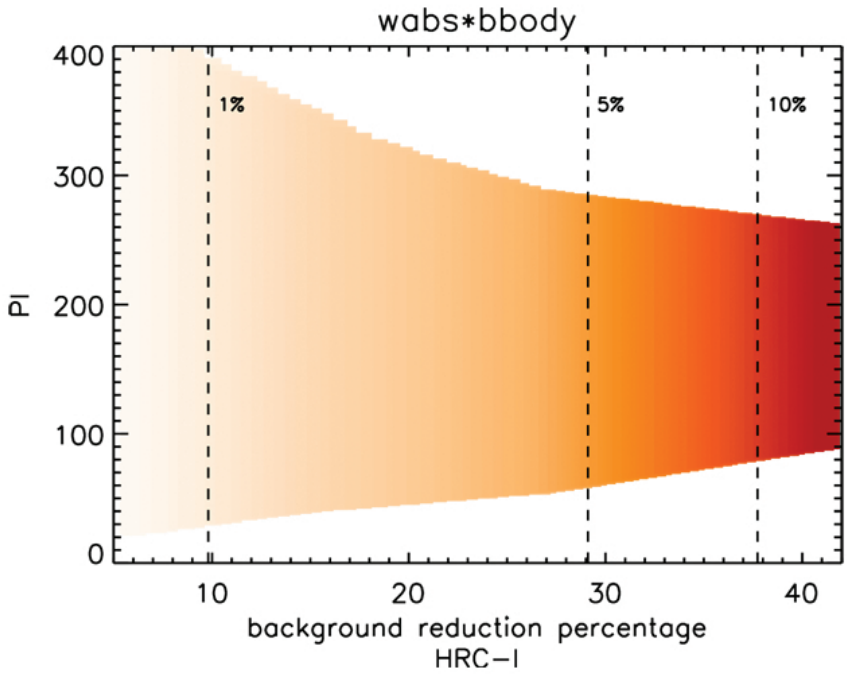

Figure 7.16: As Figure 7.14, but for sources with blackbody spectra.

assumed to be collected in a circle of radius 2.6 arcsec, and the background is assumed to be collected in an area 10× larger. The source extraction radius is not optimized for source detection and is intended to provide ballpark estimates of the limiting sensitivity. For each spectral model, the counts-to-energy conversion factor is computed using effective area curves derived from recent calibration observations (ObsID 25593 for HRC-I, and ObsID 25634 for HRC-S).

7.12 Instrument Anomalies

Power supply anomalies: On 2022-Feb-9, the HRC instrument experienced an electronics anomaly. This was similar in characteristics to the anomaly experienced on 2020-Aug-24, when there was a drop in power being supplied to the instruments. After the 2020-Aug anomaly, the HRC was switched over to a redundant set of power supplies and command and processing electronics ("Side B"), and had been operating nominally prior to the anomaly on 2022-Feb-9. Subsequent investigations showed that the anomalies differed in detail, and that switching to the original set of power supplies and electronics ("Side A") would restore operational functionality. An attempt was made on 2022-Mar-11 to restore Side B functionality, but this was unsuccessful. The switch to Side A was made on 2022-May-19, which successfully restored the instrument functionality. Several check-out observations were carried out after that, and analysis shows no departure from known trends in instrument characteristics. The degap is unaffected, and gain and QE drops (see Sections 7.8,9.3.6) are consistent with previous trends. New thermal constraints have been added that limit observation durations to a maximum of 14.5 ks at a time, with a substantial cooling duration following each observation. HRC science observations resumed on 2023-Apr-10.

PSF "ghost": Early observations with the HRC-I showed a faint secondary "ghost" image. This "ghost" image was a displaced, weaker ($\sim 3\%$) image ~ 10 arcsec on one side of every source in the HRC-I field of view, generally along the negative U axis of the instrument (Figure 7.1). The cause of this imaging anomaly is saturation of the fine position amplifiers. A change in the HRC-I operating high-voltage reduced the occurrence of saturating events and the previously mentioned event processing algorithms, which are now part of the *CXO*/HRC data pipeline, label these events and filter them out. The combination of the HV change and filtering have reduced the relative intensity of the ghost image to $< 0.1\%$, effectively eliminating it. If the location of the ghost image interferes with features of the source, the *CIAO* tool `obsvis` can be used to determine a roll angle that places the source features away from the ghost image. A similar ghost image existed in the HRC-S but at a much reduced intensity.

Timing Error: The HRC has a hardware problem that corrupts the data from the position taps under a specific set of conditions: 1) the amplifier scale factor is switched to the least sensitive scale, 2) an even number of taps on the axis have signals that are

Table 7.2: HRC-I sensitivity

Model Parameters	$N_{\rm H}$ [$\times 10^{20}$ cm^{-2}]	Limiting Flux† [$\times 10^{-15}$ erg s^{-1} cm^{-2}] at Background level [counts s^{-1} arcsec^{-2}]			
		5×10^{-6}	10^{-5}	5×10^{-5}	10^{-4}
		`xspowerlaw`			
	0.01	6.8	8.2	17.	27.
$\alpha = 1.4$	4	8.3	10.	21.	33.
	100	22.	27.	57.	89.
	0.01	5.1	6.1	13.	20.
$\alpha = 1.7$	4	6.1	7.3	15.	24.
	100	18.	21.	45.	71.
	0.01	4.0	4.8	10.	16.
$\alpha = 2.1$	4	4.4	5.2	11.	17.
	100	13.	16.	33.	52.
	0.01	3.8	4.6	9.7	15.
$\alpha = 2.4$	4	3.6	4.4	9.2	14.
	100	11.	13.	27.	42.
		`xsapec`			
	0.01	7.9	9.5	20.	32.
kT= 0.086 keV	4	3.0	3.6	7.7	12.
	100	2.0	2.4	5.0	7.8
	0.01	2.0	2.4	5.0	7.8
kT= 0.432 keV	4	1.8	2.2	4.6	7.2
	100	2.5	3.0	6.4	10.
	0.01	2.3	2.8	5.8	9.2
kT= 0.862 keV	4	2.0	2.4	5.2	8.1
	100	3.4	4.1	8.6	13.
	0.01	3.4	4.1	8.6	13.
kT= 1.72 keV	4	3.3	3.9	8.2	13.
	100	6.8	8.2	17.	27.

†: For a 3σ detection of an on-axis source in a 50 ks observation with a 2.6 arcsec-radius source circle and background collected in an area 10× larger.

Table 7.3: HRC-S sensitivity

Model Parameters	$N_{\rm H}$ [$\times 10^{20}$ cm^{-2}]	Limiting Flux† [$\times 10^{-15}$ erg s^{-1} cm^{-2}] at Background level [counts s^{-1} arcsec^{-2}]			
		5×10^{-6}	10^{-5}	5×10^{-5}	10^{-4}
		`xspowerlaw`			
	0.01	6.7	8.1	17.	27.
$\alpha = 1.4$	4	8.2	9.9	21.	33.
	100	24.	29.	61.	95.
	0.01	4.9	5.9	12.	19.
$\alpha = 1.7$	4	6.0	7.2	15.	24.
	100	19.	23.	48.	75.
	0.01	3.8	4.5	9.5	15.
$\alpha = 2.1$	4	4.2	5.0	11.	17.
	100	14.	17.	35.	56.
	0.01	3.5	4.2	8.9	14.
$\alpha = 2.4$	4	3.4	4.1	8.6	13.
	100	11.	14.	29.	45.
		`xsapec`			
	0.01	3.5	4.2	8.9	14.
kT= 0.086 keV	4	2.5	3.0	6.3	9.9
	100	2.0	2.3	5.0	7.8
	0.01	2.0	2.4	5.2	8.1
kT= 0.432 keV	4	1.9	2.3	4.8	7.5
	100	2.6	3.2	6.7	10.
	0.01	2.3	2.8	5.9	9.2
kT= 0.862 keV	4	2.1	2.5	5.3	8.3
	100	3.5	4.2	8.9	14.
	0.01	3.3	3.9	8.2	13.
kT= 1.72 keV	4	3.2	3.8	8.1	13.
	100	7.1	8.5	18.	28.

†: For a 3σ detection of an on-axis source in a 50 ks observation with a 2.6 arsec-radius source circle and background collected in an area 10× larger.

above a set threshold, and 3) the event occurs on the negative side of the tap. When these conditions are met the tap signals are sampled while the amplifiers are still ringing after switching from the initial guess for the event coarse position to the correct one. The ringing results in offsets on the telemetered tap values from their true values, with the smallest signal of the triplet for an axis being most affected. When the event position is calculated from corrupted data, positions are incorrectly determined and can be off by a few pixels. This ringing is partially corrected for in ground processing (Juda *et al.* 2000). These corrections are implemented via the *CIAO* tool `hrc_process_events`. Observers, if they are concerned that the ringing may be producing artifacts, can apply additional filtering to remove events with AMP_SF=3.

A wiring error in the HRC causes the time of an event to be associated with the *following* event, which may or may not be telemetered. The result is an error in HRC event timing that degrades accuracy from about 16 microsec to roughly the mean time between events. For example, if the trigger rate is 250 events/sec, then the average uncertainty in any time tag is less than 4 millisec.

The HRC team has developed a special operating mode that allows high precision timing to be achieved (see Section 7.15.1). This timing mode uses only the central segment of the HRC-S. Disabling the outer two segments lowers the total count rate by two-thirds, dropping it below the telemetry saturation limit for most sources. Thus, there is a high probability that all events will be telemetered. In this case, once the time tag of each event has been appropriately shifted in ground processing, the original timing accuracy (16 microsec) can be recovered. When using this approach, it is prudent to be sure that the total count rate (source plus background) is below the telemetry saturation limit to avoid telemetry saturation due to statistical fluctuations in the count rate. Note also that due to the gain drop in both the HRC-I and the HRC-S, a large number of events will fall below the lower-level discriminator set for this mode, which could reduce the source count rate by as much as 45%. Observers should consult the Instrument or Calibration Teams if their program requires the use of this special mode.

Saturation Effects: In addition to the primary science data for individual events, the rate of microchannel plate triggers (total rate) and triggers that pass on-board validity tests (valid rate) are telemetered to the ground. The valid rate is used to correct the primary rate for deadtime and telemetry saturation effects. As long as the primary rate is below saturation, the primary rate itself can be used to make the small (<1%) correction, since the event processing deadtime is known. However, when the event rate exceeds saturation, a fairly common occurrence because of background flaring from low energy protons, the valid rate is necessary to correct the event rate. Unfortunately, the total and valid event rates are overestimated by about 15% for normal operation of the HRC-S. This problem is caused by an overshoot in occasional large trigger pulses, resulting in double counting in the total and valid event on-board scalers. The primary science event is not affected, since once event processing starts with the initial trigger pulse, a gate rejects further pulses until processing is complete. The HRC-I does not have this overshoot problem. The HRC-

S valid event rate is corrected in standard processing, using the fraction of event pulse amplitudes that are above a given (segment dependent) threshold.

Tailgating: The HRC PSF also suffers from a tailgating effect, where photons within the area of the PSF that are recorded rapidly after a previous photon have less accurate positions, leading to the PSF for these events being puffier (Juda 2012, https://cxc.cfa.harvard.edu/contrib/juda/memos/hrc_pileup/index.html). Photons with arrival time differences of < 0.05 sec are affected.

Secondary Science Corruption: Sometimes the processed data may appear to have a large number of Good Time Intervals ($\gg$20) despite the full-chip background rate not showing any evidence of telemetry saturation. This is caused by a byte-shift anomaly in the secondary housekeeping files. Observers should note that the science data are not corrupted in any way, and are still available in the level 1 events list. All the events lost to time filtering may be recovered by simply not including the affected intervals in the time filtering step recommended during reprocessing. A new deadtime factor should be computed using the *CIAO* tool `hrc_dtfstats` as described in the thread *Computing Average HRC Dead Time Corrections* (`https://cxc.cfa.harvard.edu/ciao/threads/hrc_dtfstats/`).

7.13 Calibration

Calibration of the HRC included laboratory calibrations, a system-level ground calibration with the HRMA and HRC at the X-ray Calibration Facility (XRCF) at MSFC, and on-orbit calibration using celestial X-ray sources. The on-orbit calibration of the HRC is an on-going activity. See Tables 7.4,7.5 for a list of HRC calibration targets. All calibration analysis is described in detail at (https://cxc.harvard.edu/cal/Hrc).

7.14 Operational Considerations and Constraints

In addition to the general *Chandra* observatory level constraints (Chapter 3), there are a few HRC-specific considerations and constraints that must be taken into account when planning an observation.

7.14.1 Total Count Limits

Both the gain and the quantum efficiency are adversely affected by the total amount of charge extracted from the MCP at the point of extraction. To minimize such effects, the high voltage on the detector is lowered during passage through the radiation belts and at times of very high particle radiation. To limit the impact from X-ray sources themselves, a 450,000 count limit, per source, distributed over the dither pattern from an on-axis source at a given aimpoint has been imposed. Any source with flux approaching the vicinity of

Table 7.4: Current and past HRC-I calibration targets

Target	Frequency (per Cycle)	Cycle	Grating	Purpose
2REJ1032+532	11	1	LETG/None	PSF calibration
31 Com	1	>8	None	ACIS undercover; off-axis PSF & gain uniformity
3C273	1	1	None	Cross-calibration with ACIS
AR Lac	21	1-19	None	Monitor gain at aimpoint & 20 offset locations
AR Lac	42	20+	None	Monitor gain and optical axis at aimpoint & 40 offset locations
Betelgeuse	1	3–6	None	Monitor UV/Ion Shield
Capella	20	7–8	None	Improve de-gap corrections
Cen A	3	1	None	imaging capabilities
Cas A	2	1–8	None	Monitor QE; cross-calibration
Cas A	1/2	8–11	None	Monitor QE; cross-calibration
Cas A	1	22	None	Monitor degap
Coma Cluster	4	1	None	Monitor temporal variations & calibrate de-gap
Coma Cluster	1	2,3,16	None	Monitor temporal variations & calibrate de-gap
E0102-72.3	1	1,14	None	Cross-calibration with ACIS
G21.5-0.9	1	5–8,19+	None	Monitor QE; cross-calibration
G21.5-0.9	2	2–4	None	Monitor QE; cross-calibration
G21.5-0.9	3	0–1	None	Monitor QE; cross-calibration
G21.5-0.9	0.5	9–18	None	Monitor QE; cross-calibration
HR 1099	7	20	None	Monitor PSF size
HR 1099	1	21	None	Monitor PSF size
HR 1099	1	21	HETG	Monitor PSF size, calibrate 0th order and HETGS+HRC-I configuration
HR 1099	63	1	LETG/None	PSF and Wavelength calibration
HR 1099	30	20	None	PSF calibration
HZ 43	1	8–15	LETG	Monitor low energy response
HZ 43	2	4–7,16+	LETG	Monitor low energy response
HZ 43	3	1	LETG	Monitor low energy response
HZ 43	4	2–3	LETG	Monitor low energy response
LMC X-1	16	1	None	PSF calibration
M82	1	1	None	detector imaging
N132D	2	1	None	Cross-calibration with ACIS
NGC 2516	3	1	None	Boresighting & plate scale
NGC 2516	1	2	None	Boresighting & plate scale
PKS2155-304	2	2	LETG	Monitor low energy QE; cross-calibration
PKS2155-304	1	4	LETG	Monitor low energy QE; cross-calibration
Procyon	1	4–5	None	ACIS undercover; off-axis PSF& gain uniformity
Proxima Cen	2	14,17	None	ACIS undercover; off-axis PSF & gain uniformity
PSRB0540-69	5	1	None	Verification of fast timing capability
Ross 154	1	8	None	ACIS undercover; off-axis PSF & gain uniformity
RT Cru	1	14	None	PSF calibration
RXJ1856.5-3754	2	4–7	None	ACIS undercover; off-axis PSF & gain uniformity
Vega	2	1–7	None	Monitor UV/Ion Shield
Vega	4	7–11	None	Monitor UV/Ion Shield
Vega	1	>11	None	Monitor UV/Ion Shield
Vela SNR	1	3	None	low energy QE uniformity
Vela SNR	2	4	None	low energy QE uniformity

Table 7.5: Current and past HRC-S calibration targets

Target	Frequency (per Cycle)	Cycle	Grating	Purpose
3C273	1	1	LETG	Cross-calibration
3C273	1	3	LETG	Cross-calibration
AR Lac	2×21	3,5–8	None	Monitor gain at aimpoint & 20 offset locations
AR Lac	21	1–2,4,9–19	None	Monitor gain & 20 offset locations
AR Lac	1	20+	None	Monitor gain and PSF at aimpoint
AR Lac	7	20	None	Monitor PSF in cross-dispersion direction
Betelgeuse	4	1-6	None	Monitor UV/Ion Shield
Betelgeuse	2	7	None	Monitor UV/Ion Shield
Capella	10	1	LETG	Monitor gratings
Capella	1	1–7,10–13,17+	LETG	Monitor gratings
Cas A	5	1	None	Cross-calibrate HRC MCPs
Cas A	1/2	>9	None	Cross-calibrate HRC MCPs
G21.5-0.9	1	5–8,15,19	None	Monitor QE; cross-calibration
G21.5-0.9	2	2,4	None	Monitor QE; cross-calibration
G21.5-0.9	3	0,3	None	Monitor QE; cross-calibration
G21.5-0.9	7	1	None	Monitor QE; cross-calibration
G21.5-0.9	1	24	LETG	Monitor QE; cross-calibration
HR 1099	1,21	1	LETG	Wavelength calibration
HZ 43	1	8,10–12,14–15	LETG	Monitor low energy response
HZ 43	2	0–7,20+	LETG	Monitor low energy response
HZ 43	3	13,19	LETG	Monitor low energy response
HZ 43	4	18	LETG	Monitor low energy response
HZ 43	14	14	None	Monitor gain across the detector
HZ 43	16	20	None	Monitor gain across detector, locate thin/thick filter transition
LMC X-1	26	1	None	PSF calibration
Mkn 421	1	8	LETG	Monitor ACIS contamination, cross-calibration
NGC 2516	1	1	None	Boresight and plate-scale
PKS2155-304	1	1–4	LETG	Gratings calibration, monitor ACIS contamination
PKS2155-304	2	4–8	LETG	Gratings calibration, monitor ACIS contamination
Procyon	3	1	LETG	Gratings calibration, cross-calibration
PSRB0540-69	7	1	None	Verification of fast timing capability
PSRB1821-24	1	7	None	Timing calibration
Sirius B	3	1	LETG	Calibrate LETG low-energy QE
Vega	2×4	1–6	None	Monitor UV/Ion Shield
Vega	4×4	6–11	None	Monitor UV/Ion Shield
Vega	1×4	>11	None	Monitor UV/Ion Shield
var. blank sky	≈8	22+	None	Background monitoring at aimpoint during cold ECS pointings

10^{-10} erg cm^{-2} s^{-1} should be checked. *Users anticipating to exceed this count limit should so note in the comments section of the CPS form when submitting their proposal.* In this case, the *CXC* will establish new aimpoints as necessary. Offsets in the pointing may be imposed, if necessary, to limit the accumulated dose to a given region of the MCP.

7.14.2 Count Rate Limits

There are two limits on count rates, arising from telemetry and linearity.

Telemetry Limit

The maximum telemetered count rate is 184 cts s^{-1}. This is a limitation on the total count rate received over the full field-of-view rather than for one individual source within the field. It is possible to exceed this limit and to subsequently correct the total count rate by using the secondary science rates, which keep track of the actual detected rate, to determine the deadtime correction (see Section 7.12). The resulting deadtime fraction increases rapidly with valid event rates above 184 cts s^{-1}. For example, at 200 cts s^{-1} the deadtime fraction is 8%, at 250 cts s^{-1} 26%, and at 300 cts s^{-1} 39%. Listed below are some methods for dealing with situations where the telemetry limit is exceeded.

1. *Bright target:*
    - Insert either the LETG or HETG and analyze the zeroth-order image. This solution may be so dramatic as to substantially increase the required observing time.
    - Offset aimpoint. To be effective, this solution may result in substantially reduced spatial resolution.
2. *Bright nearby source*
    - Depending on the proximity, an appropriate choice of roll angle and/or offset can position the offending source(s) off the detector. Flux from bright sources could be blocked with the HRC shutters, but note that only one blade is functional and this option is unavailable in the standard setup.
    - Request a rectangular window for on-board data so that events produced by the nearby bright source(s) do not contribute to the telemetry limit.

There are of course, other combinations and situations that can lead to telemetry saturation - numerous faint sources on the field, a too-bright extended source, etc.

Linearity limit

During ground calibration, the HRC-I was verified to be linear for incident photon rates at ~ 2 cts s^{-1} pore^{-1}, which translates to ~ 5 cts s^{-1} for an on-axis point source (see Kenter

et al. 1997, Figure 7). The HRC-S was found to be linear for rates five times greater. At much higher incident fluxes, the measured rate will be lower than expected (see Pease & Donnelly 1998; https://cxc.harvard.edu/cal/Hrc/detailed_info.html#ctrt_lin). Observations of the coronal point source Capella with the HRC-I show that the data are consistent with the nominal correction for a point source of intensity ∼19–22 cts s^{-1}.

It is important to be aware that *avoiding telemetry saturation does not guarantee that linearity limits are not exceeded.* There are only three approaches to assure that the linearity limit is not exceeded:

- Offset aimpoint to smear the image out.
- Insert a transmission grating to reduce the flux and offset aimpoint (if also necessary).
- Defocus—this option is not recommended and is only mentioned for completeness.

Note that sources with high count rates will also have smaller photon arrival time differences, which will cause the PSF to be broader (see Section 7.6).

7.15 Observing with HRC - Operating Modes

For many observations, it is only necessary to specify the instrument, the exposure time, and the target coordinates. However, there are a number of optional parameters that might be invoked to optimize a particular observation. Tools such as PIMMS and *MARX* can be used to plan an observation, e.g., to account for the background when estimating sensitivity. These tools may be found at https://cxc.harvard.edu/proposer/.

7.15.1 Timing Mode

The HRC-S is normally operated in spectroscopy mode, where signals from any of the three MCP segments can be recognized as triggers. An alternate mode of operation (timing) ties the signals from the outer segments to ground so that only signals from the center MCP generate triggers. A key distinction of this mode from using an edge-blanked region (described below) to select only the center MCP segment is that the timing mode selects events without using the on-board veto logic. This preferred method of doing high-precision timing observations reduces the active detector area, minimizing the total trigger rate. Provided that this rate is below telemetry saturation, all events will then be telemetered and the event time tags can be correctly assigned in ground processing (see Section 7.12).

The HRC-S, when used in this mode, provides about a 6 arcmin × 30 arcmin field of view. Note that the lower discriminator trigger is set higher in this mode and, due to gain decline, it will also result in reduced count rates by as much as 50%.

7.15.2 Edge and Center Blanking

It is possible to define a rectangular region other than the default region on both the HRC-I and the HRC-S. Events from either inside (edge-blanking) or outside (center-blanking) the defined regions are telemetered. This could be done, for example, to prevent events from a nearby bright source from contributing to telemetry (see Section 7.14.2). If a proposer wishes to define such a rectangular region, they should state this request in the "Remarks" field of the CPS form to prompt discussions with a *CXC* Support Scientist.

7.16 References

General

David, L.P., Harnden, F.R. Jr., Kearns, K.E, and Zombeck, M.V., *The ROSAT High Resolution Imager (HRI) Calibration Report*, revised (1999).
https://hea-www.harvard.edu/rosat/hricalrep.html

Drake, J.J., Kashyap, V.L., Wargelin, B.J., and Wolk, S.J., *Pointing Chandra toward the Extreme Ultraviolet Fluxes of Very Low Mass Stars* (2020), ApJ, 893, 137.

Fraser, G., "X-ray Detectors in Astronomy", 1989, Cambridge University Press.

Giacconi, R., *et al.*, 1979, Ap. J., 230, 540.

Murray, S.S., Chappell, J.H., Elvis, M.S., Forman, W.R., Grindlay, J.E., Harnden, F.R., Jones, C.F., Maccacaro, T., Tananbaum, H.D., Vaiana, G.S., Pounds, K.A., Fraser, G.W., and Henry, J.P., "The AXAF High Resolution Camera (HRC) and its use for observations of Distant Clusters of galaxies" Astro. Lett. Comm., 26, 113-125, 1987.

Murray, S.S., *et al.*, "In-flight Performance of the Chandra High Resolution Camera", SPIE, 4012, 2000. https://cxc.cfa.harvard.edu/cda/SPIE/smurray2000.pdf

Zhao, P., 2014, CXC Memo, Nov 1 2014, *Chandra Optical Axis, Aimpoint and Their Drifts* The optical axis of the HRC-I is currently measured to be at (ChipX,ChipY)= $(7586.7\pm14.2, 7741.8\pm4.7)$, and that of the HRC-S is at (ChipX,ChipY)= $(2194.6\pm10.6, 8913.4\pm10.6)$ (Section 4.5.5).

Zombeck, M.V., Chappell, J. H , Kenter, A, Moore, R., W., Murray, S. S., Fraser, G.W., Serio, S.,"The High Resolution Camera (HRC) on the Advanced X-ray Astrophysics Facility (AXAF)", Proc. SPIE, 2518, 96, 1995.

Position modeling, de-gap corrections, and event screening

Juda, M., *et al.*, "Improving Chandra High Resolution Camera event positions via corrections to cross-grid charge detector signals", SPIE Proceedings, 4140, 2000. https://cxc.cfa.harvard.edu/contrib/juda/memos/spie2000_tap_correction.html

Juda, M., & Karovska, M., "Chandra's Ultimate Angular Resolution: Studies of the HRC-I Point Spread Function", AAS/HEAD 2010. https://hea-www.harvard.edu/ juda/memos/HEAD2010/HEAD2010_poster.html

Karovska, M., 2011, "Followup Study of the PSF Asymmetry", CXC Memo, 2011-Jun.

Kashyap, V., "Analysis of Chandra PSF feature using ACIS data", CXC Memo, 2010-Oct. https://cxc.harvard.edu/cal/Hrc/PSF/acis_psf_2010oct.html

Kashyap, V., *et al.*, 2005, "HRC-S Degap Corrections". https://cxc.harvard.edu/cal/Letg/Hrc_disp/degap.html

Kenter, A., "Degap as a Transformation of Probability Distribution Problem", 1999-Mar-01. https://cxc.harvard.edu/cal/Hrc/Documents/degap.pdf

Murray, S.S., Chappell, J.H., 1989, SPIE 1159, 460-475. "Position Modeling for the AXAF High resolution Camera (HRC)"

Murray, S.S., *et al.*, "Event Screening for the Chandra X-ray Observatory High Resolution Camera (HRC)", SPIE Proceedings, 4140, 2000. https://cxc.harvard.edu/cda/SPIE/smurray2000b.pdf

https://cxc.harvard.edu/ciao/caveats/psf_artifact.html (CIAO Caveats page on the PSF Artifact)

Count rate limitations and linearity

Juda, M and Dobrzycki, A, "HRC Deadtime and Telemetry Saturation", 1999-Jun-18. https://cxc.harvard.edu/contrib/juda/memos/tlm_sat.html

Juda, M., "Telemetered vs. Processed Events", memo, 2001-Dec-07. https://cxc.harvard.edu/contrib/juda/memos/proc2valid/index.html

Juda, M., "HRC-S Double Pulse Fraction", memo, 2002-Jun-27. https://cxc.harvard.edu/contrib/juda/memos/proc2valid/pha_fraction.html

Kenter, A.T., Chappell, J.H. Kobayashi,K.,Kraft,R.P., Meehan, G.R., Murray, S.S., Zombeck, M.V., Fraser, G.W., Pearson, J.F., Lees, J.E., Brunton, A.N. and Pearce, S.E. Barbera, M., Collura, A., Serio, S., "Performance and Calibration of the AXAF High Resolution Camera I " SPIE 3114, 1997.

Pease, D.P., & Donnelly, H., memo, 1998-May.
https://cxc.harvard.edu/cal/Hrc/detailed_info.html#ctrt_lin

Zombeck, M. V., "Secondary Science Rate Double Counting", memo, 2002-Dec-02.
https://hea-www.harvard.edu/HRC/calib/doublecount.html

Calibration

https://cxc.harvard.edu/cal (CXC calibration site)

https://hea-www.harvard.edu/HRC/calib/calib.html (HRC IPI Team calibration site)

https://cxc.harvard.edu/cal/Hrc/ (HRC CXC Cal team site)

Juda, M., 2012, CXC Memo, *"Pile-up" Effect on the HRC PSF*
https://cxc.cfa.harvard.edu/contrib/juda/memos/hrc_pileup/index.html

Kenter, A.T., Chappell, J., Kobayashi, K., Kraft, R.P., Meehan, G.R., Murray, S.S., Zombeck, M.V., "Performance and Calibration of the AXAF High Resolution Camera: I. Imaging Readout", SPIE, 3114, 26, 1997.
https://cxc.cfa.harvard.edu/cal/spie/spie97_kenter.pdf

Kenter, A., *et al.*, "In-flight Performance and Calibration of the Chandra High Resolution Camera Spectroscopic Readout (HRC-I)" SPIE, 4012, 2000.
https://hea-www.harvard.edu/HRC/calib/hrci.spie2000.ps

Kraft, R.P., Chappell, J., Kenter, A.T., Kobayashi, K., Meehan, G.R., Murray, S.S., Zombeck, M.V., "Performance and Calibration of the AXAF High Resolution Camera: II. the Spectroscopic Detector", SPIE, 3114, 53, 1997.
https://hea-www.harvard.edu/HRC/calib/spie97_kraft.ps

Kraft, R., *et al.*, "In-flight Performance and Calibration of the Chandra High Resolution Camera Spectroscopic Readout (HRC-S)" SPIE, 4012, 2000.
https://cxc.harvard.edu/cda/SPIE/kraft2000.pdf

Meehan, G.R., Murray, S.S. , Zombeck, M.V., Kraft, R.P., Kobayashi, K., Chappell, J.H., and. Kenter, A.T., "Calibration of the UV/Ion Shields for the AXAF High Resolution

Camera", SPIE, 3114, 74, 1997.
https://hea-www.harvard.edu/HRC/calib/spie97_meehan.ps

Meehan, G, "Calibration of the HRC-I UV/Ion Shield", 1999-Oct-13.
https://hea-www.harvard.edu/HRC/calib/hrci_cal_report.ps

Meehan, G.,"Calibration of the HRC-S UV/Ion Shields", 1999-Oct-13.
https://hea-www.harvard.edu/HRC/calib/hrcs_cal_report.ps

Murray, S.M., et al., "HRC Ground Calibration Revision A, SAO-HRC-CR-98-345, 1998,
https://hea-www.harvard.edu/HRC/calib/hrccalib_a_180998.ps
Murray, S. S.; Chappell, J.H.; Kenter, A. T.; Kobayashi, K.; Kraft, R. P.; Meehan, G. R.; Zombeck, M. V.; Fraser, G. W.; Pearson, J. F.; Lees, J. E.; Brunton, A. N.; Pearce, S, E.; Barbera, M.; Collura, A.; Serio, S., "AXAF High-Resolution Camera (HRC): calibration and recalibration at XRCF and beyond", SPIE, 3114, 11, 1997.

Background

Isobe, T., and Juda, M., memo, 2007-Sep-11.
https://cxc.harvard.edu/contrib/cxchrc/Stowed_study/hrc_stowed_position_study.html

Isobe, T., and Juda, M., 2007, "High Resolution Camera Stowed Background Study", Proc. of Chandra Calibration Workshop, October 2007, Huntsville, AL.
https://cxc.harvard.edu/ccw/proceedings/07_proc/presentations/isobe/

Isobe, T., and Juda, M., 2009, "How to Create a Background Map for an Observation", memo, 2009-Jan-27.
https://cxc.harvard.edu/contrib/cxchrc/Stowed_study/hrci_image_correction.html

Isobe, T., and Juda, M., 2009, "How to Create a Background Map for an Observation", Proc. of Chandra Calibration Review, September 2009, Boston, MA.
https://cxc.harvard.edu/ccr/proceedings/09_proc/presentations/isobe/

Juda, M., "Time History of the HRC Background", memo, 2001-May-22.
https://cxc.harvard.edu/contrib/juda/memos/hrc_bkg/time_history.html

Juda, M., "HRC Rates and High Solar Activity", memo, 2001-May-21.
https://cxc.harvard.edu/contrib/juda/memos/hrc_bkg/high_solar.html

Juda, M., *et al.*, "Characteristics of the On-Orbit Background of the Chandra X-ray Observatory High Resolution Camera", Proc. SPIE 4851, August 2002
https://cxc.harvard.edu/contrib/juda/memos/spie2002/spie2002.html,

Detector coordinate systems

McDowell, J., "Coordinate Systems for Analysis of On-orbit Chandra Data, Paper I: Imaging", https://cxc.cfa.harvard.edu/ciao/download/doc/coords60.pdf

Counts lifetime

Kenter, A.T., K.A. Flanagan, G. Meehan, S.S. Murray, M.V. Zombeck, G.W. Fraser, J.F. Pearson, J.E. Lees, A.N. Brunton, and S.E. Pearce, "Microchannel plate testing and evaluation for the AXAF high resolution camera (HRC)", Proc. SPIE, 2518, 356, 1995.

Gain, spectral response, out-of-band response

McEntee, S.C., Jackman, C.M., Weigt, D.M., Dunn, W.R., Kashyap, V., Kraft, R., Louis, C.K., Branduardi-Raymont, G., Gladstone, G.R., & Gallagher, P.T., 2022, "Comparing Jupiter's equatorial X-ray emissions with solar X-ray flux over 19 years of the Chandra mission", JGR Space Physics, 127, e2022JA030971

Kashyap, V. Posson-Brown, J., 2009, "The Imaging and Spectral Performance of the HRC", Chandra Calibration Review, 2009.14,
https://cxc.harvard.edu/ccr/proceedings/09_proc/presentations/kashyap/

Kashyap, V., Posson-Brown, J., 2005, "Spectral Response of the HRC-I", Chandra Calibration Workshop, 2005-Oct-31 to 2005-Nov-01,
https://cxc.harvard.edu/ccw/proceedings/05_proc/presentations/kashyap2/

Pease, D.O., Drake, J.J., & Kashyap, V.L., 2005, "The Darkest Bright Star: Chandra X-Ray Observations of Vega", ApJ, 636, 426

Pease, D., Kashyap, V., Drake, J., Juda, M., 2005, "Monitoring the HRC-S UV Rate: Observations of Vega", CXC Memo, 2005-May,
https://cxc.cfa.harvard.edu/cal/Hrc/Documents/hrcs_vega_uv05.pdf

Posson-Brown, J., Kashyap, V., 2005, "Monitoring the Optical/UV Transmission of the HRC with Betelgeuse", CXC Memo, 2005-Jun,
https://cxc.harvard.edu/cal/Hrc/Documents/betelgeuse.pdf

Posson-Brown, J., Kashyap, V., 2005, "Monitoring the Optical/UV Transmission of the HRC with Betelgeuse", Chandra Calibration Workshop, 2005-Oct-31 to 2005-Nov-01,
https://cxc.harvard.edu/ccw/proceedings/05_proc/presentations/possonbrown/

Posson-Brown, J., Kashyap, V., 2009, "SUMAMPS-based gain maps and RMF for the HRC-I", Chandra Calibration Review, 2009.16,

https://cxc.harvard.edu/ccr/proceedings/09_proc/presentations/possonbrown2/

Wargelin, B., 2012, CXC Memo, *HRC-S Voltage Change*,
https://cxc.cfa.harvard.edu/cal/Letg/newHRCShv/

Wilton, C., Posson-Brown, J., Juda, M., Kashyap, V., 2005, "The HRC-I Gain Map", Chandra Calibration Workshop, 2005-Oct-31 to 2005-Nov-01 ,
https://cxc.harvard.edu/ccw/proceedings/05_proc/presentations/wilton/

Zombeck, M.V., "HRC-I out of band response."
https://hea-www.harvard.edu/HRC/calib/hrci_cal.html

Zombeck, M.V., "HRC-S out of band response."
https://hea-www.harvard.edu/HRC/calib/hrcs_cal.html

Zombeck, M.V., *et al.*, "Vega calibration observations."
https://hea-www.harvard.edu/HRC/calib/vega/vega.html

Zombeck, M.V., *et al.*, "The Out-of-band Responses of the HRC on Chandra", X-ray 2000 Proceedings, Palermo, 2000.
https://hea-www.harvard.edu/HRC/calib/palermopaper.ps

Zombeck, M. V., "Response of the HRC to Vega", memo, 2002-Oct-28.
https://hea-www.harvard.edu/HRC/calib/vega/vega_trend.html

Chapter 8

HETG: *Chandra* High Energy Transmission Grating

8.1 Instrument Overview

HETG is the High-Energy Transmission Grating (Canizares *et al.* 2005). In operation with the High Resolution Mirror Assembly (HRMA) and a focal-plane imager, the complete instrument is referred to as the HETGS—the High-Energy Transmission Grating Spectrometer. The HETGS provides high resolution spectra (with $E/\Delta E$ up to 1000) between 0.4 keV and 10.0 keV for point and slightly extended (few arcsec) sources. Although HETGS operation differs from proportional counter and CCD spectrometers, standard processing of an HETGS observation produces familiar spectrometer data products: pulse height amplitude (*PHA*) files, auxiliary response files (ARFs), and redistribution matrix files (RMFs). These files can then be analyzed with standard forward-folding model fitting software, e.g., *Sherpa*, *XSPEC*, or *ISIS*.

The HETG itself consists of two sets of gratings, each with different period. One set, the Medium Energy Grating (MEG), intercepts rays from the outer HRMA shells and is optimized for medium energies. The second set, the High Energy Gratings (HEG), intercepts rays from the two inner shells and is optimized for high energies. Both gratings are mounted on a single support structure and therefore are used concurrently. The two sets of gratings are mounted with their rulings at different angles so that the dispersed images from the HEG and MEG will form a shallow X centered at the undispersed (zeroth order) position; one leg of the X is from the HEG, and the other from the MEG. The HETG is designed for use with the spectroscopic array of the *Chandra* Advanced CCD Imaging Spectrometer (ACIS-S) although other detectors may be used for particular applications. A summary of characteristics is given in Table 8.1.

The Instrument Principal Investigator for the HETG is Dr. Herman Marshall of the MIT Kavli Institute for Astrophysics and Space Research. See Canizares *et al.* (2005) for a thorough description of the instrument and its performance.

Table 8.1: HETG(S) Parameters

HETGS Range:	0.4–10.0 keV, 31–1.2Å
HEG Range:	0.8–10.0 keV, 15–1.2 Å
MEG Range:	0.4–5.0 keV, 31–2.5 Å
Effective Area (see Figures 8.7 8.8):	$7\,\mathrm{cm}^2$ @ 0.5 keV
(MEG+HEG first orders,	$59\,\mathrm{cm}^2$ @ 1.0 keV
with ACIS-S)	$200\,\mathrm{cm}^2$ @ 1.5 keV
	$28\,\mathrm{cm}^2$ @ 6.5 keV
Resolving Power ($E/\Delta E$, $\lambda/\Delta\lambda$)	
HEG:	1070–65 (1000 @ 1 keV, 12.4 Å)
MEG:	970–80 (660 @ 0.826 keV, 15 Å)
Resolution:	
ΔE:	0.4–77 eV FWHM
$\Delta\lambda$, HEG:	0.012 Å FWHM
$\Delta\lambda$, MEG:	0.023 Å FWHM
Absolute Wavelength Accuracy: (w.r.t. "theory")	
HEG	±0.006 Å
MEG	±0.011 Å
Relative Wavelength Accuracy: (within and between obs.)	
HEG	±0.0010 Å
MEG	±0.0020 Å
HEG angle on ACIS-S:	$-5.235^\circ \pm 0.01^\circ$
MEG angle on ACIS-S:	$4.725^\circ \pm 0.01^\circ$
HETGS Rowland spacing	8632.65 mm (flight installed)
Wavelength Scale:	
HEG	0.0055595 Å / ACIS pixel
MEG	0.0111200 Å / ACIS pixel
HETG Properties:	
Diffraction Efficiency:	2.5% @ 0.5 keV (MEG)
(single-side, first order)	19% @ 1.5 keV (MEG & HEG)
	9% @ 6.5 keV (HEG)
HETG Zeroth-order Efficiency:	4.5% @ 0.5 keV
	8% @ 1.5 keV
	60% @ 6.5 keV
Grating Facet Average Parameters	
HEG and MEG bar material:	Gold
HEG / MEG period:	2000.81 Å / 4001.95 Å
HEG / MEG Bar thickness:	5100 Å / 3600 Å
HEG / MEG Bar width:	1200 Å / 2080 Å
HEG / MEG support:	9800 Å / 5500 Å polyimide

8.1.1 Examples of Observations with the HETGS

An example of an HETGS observation is presented in Figure 8.1 using data from an observation of Capella, ObsID 1318. The top panel shows an image of detected events on the ACIS-S detector with the image color indicating the ACIS-determined X-ray energy.In this detector coordinate image (TDETX, TDETY), the features are broad due to the nominal dither motion which serves to average over detector non-uniformities. The ACIS-S chips are numbered S0 to S5 from left to right, with the aim point in S3 where the bright zeroth-order image is visible and includes a vertical frame-transfer streak (a trailed image). The HRMA optical axis passes through S3 approximately 6 mm from the S2–S3 chip gap. For further information see Figure 6.1 and related text.

HETG-diffracted photons are visible in Figure 8.1 forming a shallow "X" pattern; the full opening angle between the HEG and MEG spectra is 9.96°. The back illuminated (BI) chips are S1 and S3. The S1 location was chosen to enhance the first order MEG spectrum since back illumination provides higher efficiency below 1 keV. The location of the zeroth-order for any particular observation, however, may be adjusted by offset pointing in order to select the energies of the photons that will be placed in the gaps between the chips. Details on gaps are presented in Section 8.2.1.

The middle panel of Figure 8.1 shows an image after the data have been aspect corrected and data filters applied to include only valid zeroth and first-order events. Note that this image was created using Sky coordinates that were rotated and had their y-axis sign "flipped" to match the detector coordinates view in the top panel. The lower set of panels shows an expanded view of the MEG minus-first-order spectrum with emission lines clearly visible. Wavelengths are assigned based on the diffraction angle of the events, that is, how far the events are from the zeroth-order image. Using the grating equation in Section 8.1.3, absolute wavelengths can be assigned based on the dispersion angle. A spectrum of the source is then created by binning the events into energy or wavelength bins; the spectrum from another Capella observation is shown in Figure 8.2.

Note: The dispersion distance on the detector is essentially linear in wavelength. Thus, wavelength is the natural unit for this high-resolution X-ray spectrometer. The conversion between energy and wavelength is provided by the relation: $E \times \lambda = hc = 12.39852$ keV Å.

Each of the "arms" of the HETGS diffracted X pattern yields a first-order spectrum identified by type (HEG or MEG) and sign of the order (plus or minus). Using ARFs and RMFs these spectra can be analyzed in an *XSPEC*-like framework. Additionally, the *CXC* software package *ISIS* (https://space.mit.edu/ASC/ISIS/) can be used to identify spectral lines, e.g., as seen in Figure 8.2.

8.1.2 Scientific Objectives and Grating Heritage

The HETGS allows one to probe the physical parameters of emitting regions of all classes of X-ray sources, including stars, X-ray binaries, supernova remnants, galaxies, clusters of galaxies, quasars, and interstellar and intergalactic material. Plasma diagnostic tech-

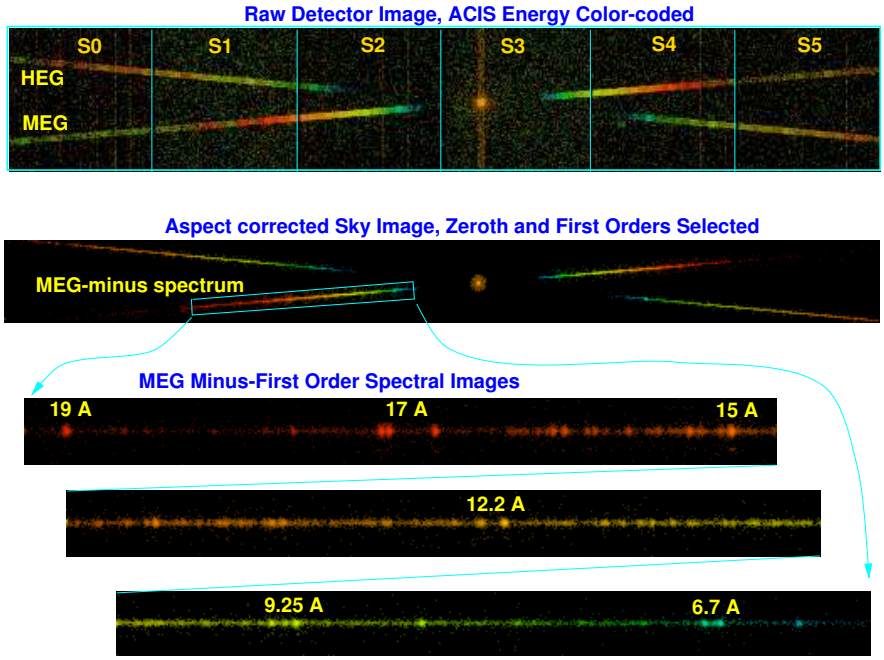

Figure 8.1: HETGS observation of Capella, ObsID 1318. The top panel shows an image of detected events on the ACIS-S detector with the image color indicating the ACIS-determined X-ray energy. The bright zeroth-order image is visible on CCD S3 and includes a trailed image (the vertical frame-transfer streak). Diffracted photons are visible forming a shallow "X" pattern; the HEG and MEG spectra are indicated. The images are broad due to dither of the spacecraft. The middle panel shows an image after the data have been aspect corrected and selections applied to include only valid zeroth and first-order events; note that the Y axis has been flipped from the normal Sky view to match the detector coordinates view in the top panel. Finally, the lower panel shows an expanded view of the MEG minus-first-order spectrum with emission lines clearly visible.

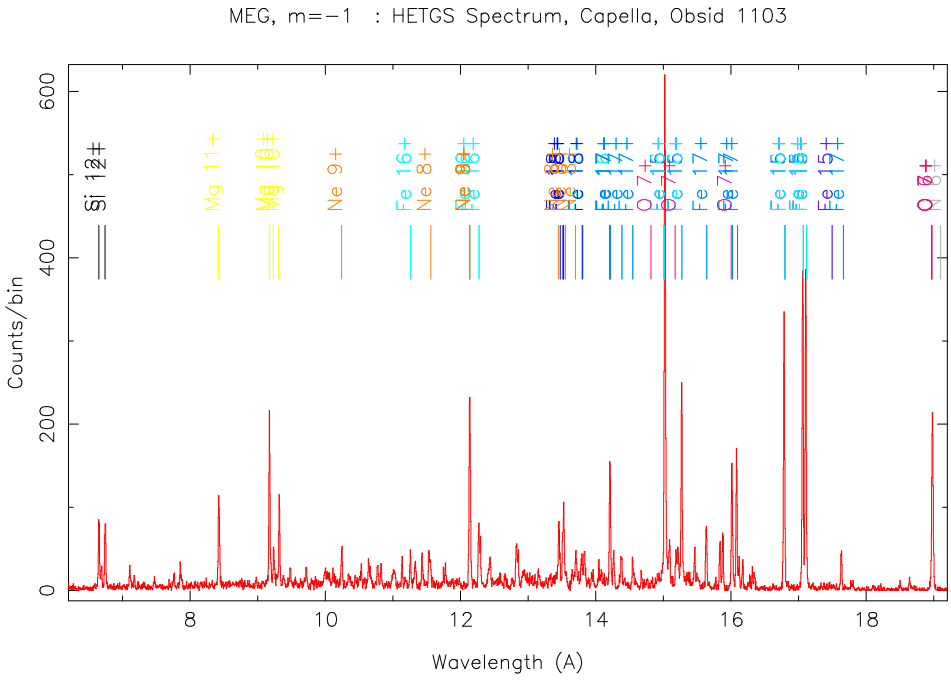

Figure 8.2: HETGS Capella spectrum, MEG $m = -1$, ObsID 1103. The first-order events identified in the MEG minus-side "arm" of the HETG X pattern are assigned wavelengths by *CXC* pipeline software according to the grating equation and known instrument parameters. These values are then binned to produce a pulse height analysis spectrum (`pha2.fits` file) which is plotted here. The `ISIS` software package available from the *CXC* has also been used to indicate the location of expected emission lines based on a simple source model.

niques applied to emission lines, absorption lines and absorption edges will convey source properties such as temperatures, ionization states, densities, velocities, elemental abundances, and thereby structure, dynamics, and evolution of various classes of sources. The energy band amenable to observation is extremely rich in lines from both coronal and photo-ionized plasmas, containing the L-shell lines from ionization stages of Fe XVII to Fe XXIV and the K-shell lines of hydrogenic and helium-like ions of oxygen through nickel. The 6 keV Fe K lines are well within the observable band. The highest resolutions available will also allow detailed study of motions through Doppler line shifts in supernova remnants, X-ray binaries, turbulent intra-cluster or intra-galactic gas, or early-type galaxies in clusters.

Although gratings have flown on *Einstein* and *EXOSAT*, the HETGS shares only the basic operating principles with these. Advanced grating technology has enabled achievement of greater efficiency and increased dispersion. The Rowland geometry of the grating plate and spectroscopic arrays maintains the telescope focal properties in the dispersion direction by minimizing dispersed image aberrations and hence contributes to improved spectral resolution.

8.1.3 HETGS Operating Principles

The HETG is mounted, and can be inserted, just aft of the HRMA as shown in the schematic of the HRMA-HETG-detector system, Figure 8.3. The HETG provides spectral separation through diffraction. X-rays from the HRMA strike the transmission gratings and are diffracted (in one dimension) by an angle β given according to the grating equation,

$$\sin\beta = m\lambda/p, \tag{8.1}$$

where m is the integer order number, λ is the photon wavelength in angstroms, p is the spatial period of the grating lines, and β is the dispersion angle. A "normal" undispersed image is formed by the zeroth-order events, $m = 0$, and dispersed images are formed by the higher orders, primarily the first-order, $|m| = 1$.

The HETGS-faceted Rowland design is shown in Figure 8.4. The "Rowland circle" is a circle whose diameter is simply the distance from the grating that would lie on the optical axis to the point in the focal plane where the zeroth order image is placed. The "Rowland torus" is formed by rotating the circle about the line in the dispersion direction going through the on-axis focus. Individual grating facets are mounted such that their centers lie on the torus. In the figure, the axis of the torus is perpendicular to the page for the side view and lies in the plane of the top view. Ideally, the detector is shaped to follow the counterpart Rowland torus in the image plane. The result is that the telescope focal properties in the dispersion direction are maintained for a large range of diffraction angles, β, thereby minimizing any grating-added optical aberrations.

An important parameter of the HETGS is the Rowland spacing, the distance from the outer intersection of the HETG axis and Rowland Circle to the HRMA focus. This

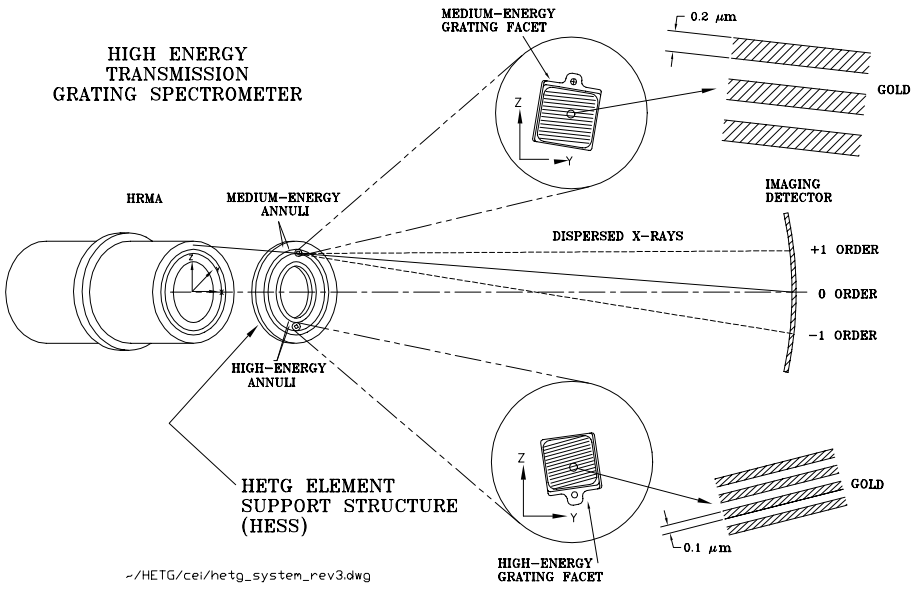

Figure 8.3: A schematic layout of the High Energy Transmission Grating Spectrometer. (Dimensions shown are approximate.)

Rowland spacing is what determines the value of β in the grating equation. The value of the Rowland spacing is listed in Table 8.1.

Order overlap and source confusion can be discriminated by the intrinsic energy resolution of the CCD detector (ACIS-S is the preferred detector for HETG spectroscopy since it has intrinsic energy resolution and so can separate orders; the HRC may be used if high time resolution is desired, but this choice is at the price of using the focal-plane detector's energy resolution to aid in order separation). The form of a spectral image on the ACIS-S array is shown in Figure 8.1. The spectroscopic array spans about 8 arcmin $\times$ 48 arcmin of the sky, though image quality and resolving power degrade rapidly for sources more than about 4 arcmin off-axis. For an on-axis source, the detector edge in the dispersion direction causes a low energy cutoff of the spectrum at about 0.4 keV for the MEG and 0.8 keV for the HEG. Order selection and chip gaps are described more fully in Section 8.2.1.

8.1.4 HETG Physical Configuration

The HETG support structure (HESS) is a circular aluminum plate (110 cm diameter by 6.35 cm thick) which can be swung into position behind the HRMA. Mounted on the HESS are 336 grating facets, each about 25 mm square. The position and orientation of the HESS mounting surfaces have been designed and machined to place each grating center on a Rowland torus of diameter 8633.69 mm. A detailed drawing of the HETG

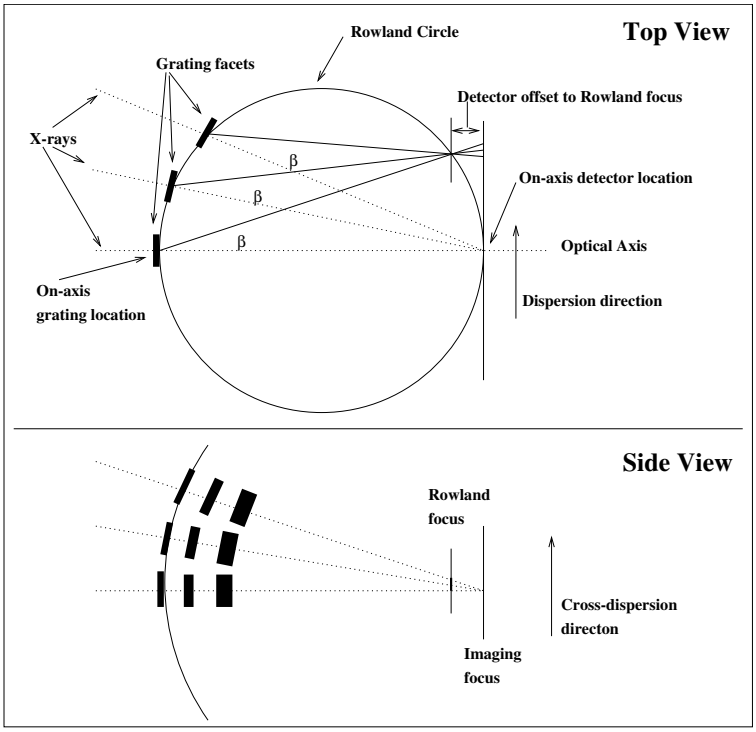

Figure 8.4: The Rowland geometry is shown schematically. The "Top" view looks across the dispersion direction. The diffraction angle is β. The geometry is such that converging rays diffracted at a specific angle by the gratings (which are located on the Rowland circle) will converge to a point that is also on the Rowland circle. The dotted lines represent zeroth-order rays and the solid lines represent the grating-diffracted first-order rays. The bottom panel ("Side" view) looks along the dispersion direction at rays from a set of gratings arranged perpendicularly to those in the "Top" view and schematically shows the astigmatic nature of the spectrally focused image: since the converging rays have not yet reached the imaging focus, they are extended across the dispersion (by less than 100 microns).

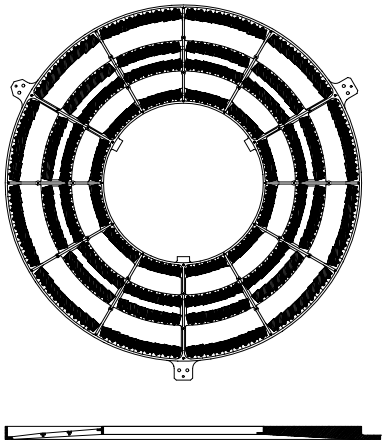

Figure 8.5: A front (upper) and side (lower) view of the HETG support structure (HESS). The grating facets are mounted to intercept the X-rays as they exit the HRMA; the front view is from the HRMA i.e., what an approaching X-ray would see. In the side view, the left cross-section shows that the four support rings are in different planes due to the Rowland curvature. The right cross-section is through a radial rib at one of the three mounting "ears."

(HESS plus facets) is shown in Figure 8.5.

The gratings cover the annular regions through which the X-rays pass. The 192 grating facets on the outer two annuli (MEG) have a period of 4001.95 Å. Tiling the inner two annuli are 144 (HEG) gratings, which have a period of 2000.81 Å (see Table 8.1). The two sets of gratings are mounted with their rulings at different angles so that the dispersed images from the HEG and MEG will form a shallow X centered at the undispersed (zeroth order) position; one leg of the X is from the HEG, and the other from the MEG. See Figure 8.1 for an example.

The HETG grating facets are composed of electro-plated gold bars supported on a polyimide substrate, as shown schematically in Figure 8.6. The grating bar design parameters, height and width, are nominally chosen to reduce zeroth-order and maximize first-order intensities. Choosing to have the bar width one-half of the grating period suppresses even orders and provides maximum 1st order efficiency for a rectangular profile; this is closely achieved for the MEG gratings. For the HEG gratings, the bar is wider and results in a higher 2nd order efficiency and reduced 3rd order efficiency. The bar height choice "tunes" the efficiency peak in energy by allowing X-rays to constructively interfere in first order in the region where the gold is partially transparent, primarily above 1.2 keV.

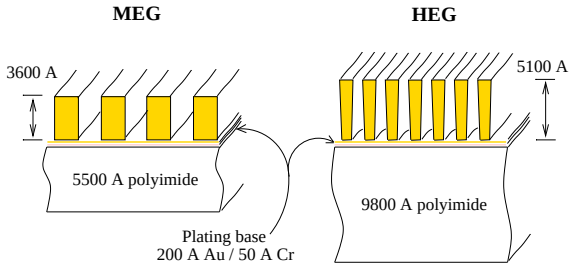

Figure 8.6: Cross-sections of the MEG and HEG membranes. The soap-bubble-thin membranes of the HETG consist of gold bars attached to a polyimide support layer. The MEG grating bars are close to rectangular, typically with a height of 3600 Å and a bar-to-period fraction of 52%. The HEG bars have a crudely trapezoidal shape, narrower on the polyimide side as shown, typically 5100 Å high with an effective bar-to-period fraction around 60%.

8.2 Instrument Characteristics

When observing a point source, the HETGS can be viewed as a black-box spectrometer characterized by its *Effective Area* and *Resolution*. More specifically an HETGS count spectrum produced by standard analysis (a *PHA* file) can be related to the source spectrum through a grating ARF and grating RMF; because of the spatially-dispersive operation of the HETGS, the RMF is also referred to as the Line Response Function, LRF. Four first-order spectra are obtained from an observation corresponding to the four whiskers of the dispersed "X" pattern: the plus and minus first-orders of the HEG and MEG gratings. Standard CIAO tools produce these *PHA* files, ARFs, and RMFs for observers' use.

In the sections that follow, information on the HETGS Effective Area (ARFs) and Resolution (RMFs/LRFs) are given with an examination of the components and effects that contribute to them. In addition two other characteristics of the HETGS are briefly presented: the *Background* event rate for an extracted spectrum and the *Absolute Wavelength* accuracy.

8.2.1 HETGS Effective Area

The HETGS effective area, as encoded in the ARF, depends on the HETG efficiency coupled with the HRMA effective area and the ACIS efficiency. Additional effects can arise from the process of selecting events, the effect of chip gaps, and the use of "ACIS ENERGY" to do order sorting. In this chapter, the term ACIS ENERGY is used to describe the energy deduced from the ACIS pulse height.

Nominal HETGS ARFs

Combining the HETG diffraction efficiencies with the HRMA effective area and the ACIS-S detection efficiency produces the system effective area as a function of energy, de-

scribed by an "ancillary response file" or ARF. Plots of HETGS ARF's are shown in Figures 8.7 and 8.9 which are plotted with log vertical axes; the same plots with linear vertical scale are shown in Figures 8.8 and 8.10. The values are plotted from ARF files created by the *CXC CIAO* tool `fullgarf`. The effective area includes the effect of molecular contamination on the ACIS filter (projected to the middle of the Cycle).

The nominal plots shown here are for qualitative reference only; because the `fullgarf` tool also accounts for a variety of other effects, *e.g.*, dither motion, bad pixels, QE non-uniformity, etc., grating ARFs are custom made for a given observation. The details of the calibration of the ARF are discussed in Section 8.3.

Since first-order photons from both the HEG and MEG gratings provide information, to compare the HETGS with other instruments, it is useful to plot the total HETGS effective area (the combined plus and minus first-order areas of both the HEG and MEG); this is shown by the solid curve in Figure 8.11. During an observation the zeroth-order photons from HEG and MEG form a single zeroth-order image; the effective area for this zeroth-order image is also plotted on this figure (dotted line).

Note the dips caused by the gaps between chips in these figures. The observatory is dithered to spread the signal across a large number of pixels. For HETGS observations, sinusoidal motions with 16 arcsec amplitude create a Lissajous pattern (see Section 5.8.2). When the combination of the chip gaps and dither are accounted for, a "pitch fork" dip occurs at each gap region in the ARFs. Although effects of this motion are removed in on-ground processing, observers are advised to *avoid placing spectral features of interest near the gaps*. More information concerning gaps are in the next section. The effective areas shown in Figure 8.7-8.11 are based on an integration over the full LSF. Most of the flux in a line will be contained within a circle of diameter 4 arcsec. The user might wish to note that in data processing, the pipeline software keeps only events that are in a spatial window that lies within 3 arcsec of the dispersion axis. This aperture guarantees that a high fraction, 97–99%, of the signal flux is retained while minimizing the contribution of the background. Further discussion of the spatial distribution of events can be found for the HRMA PSF in Chapter 4 and for the HETGS in Section 8.2.2 below.

HETG Grating Efficiency

The HETG contribution to the effective area comes in through the efficiencies of the HETG gratings; the values of these are shown in Figure 8.13. All calibration data support the modeling assumption that the positive order efficiencies are equal to the negative order efficiencies. These efficiencies are primarily based on laboratory measurements of each facet, synchrotron reference grating corrections, improved polyimide transmission models, and updated gold optical constants as described in Flanagan et al., (2000). Slight adjustments to the HETG efficiencies have been determined using in-flight data by comparing the HEG and MEG spectra of many sources. The adjustments are mostly less than 10%; see Marshall (2005) for details. In an update (Marshall 2012) included in CalDB 4.4.7 and later, the HEG and MEG agree now at the $\sim 1\%$ level.

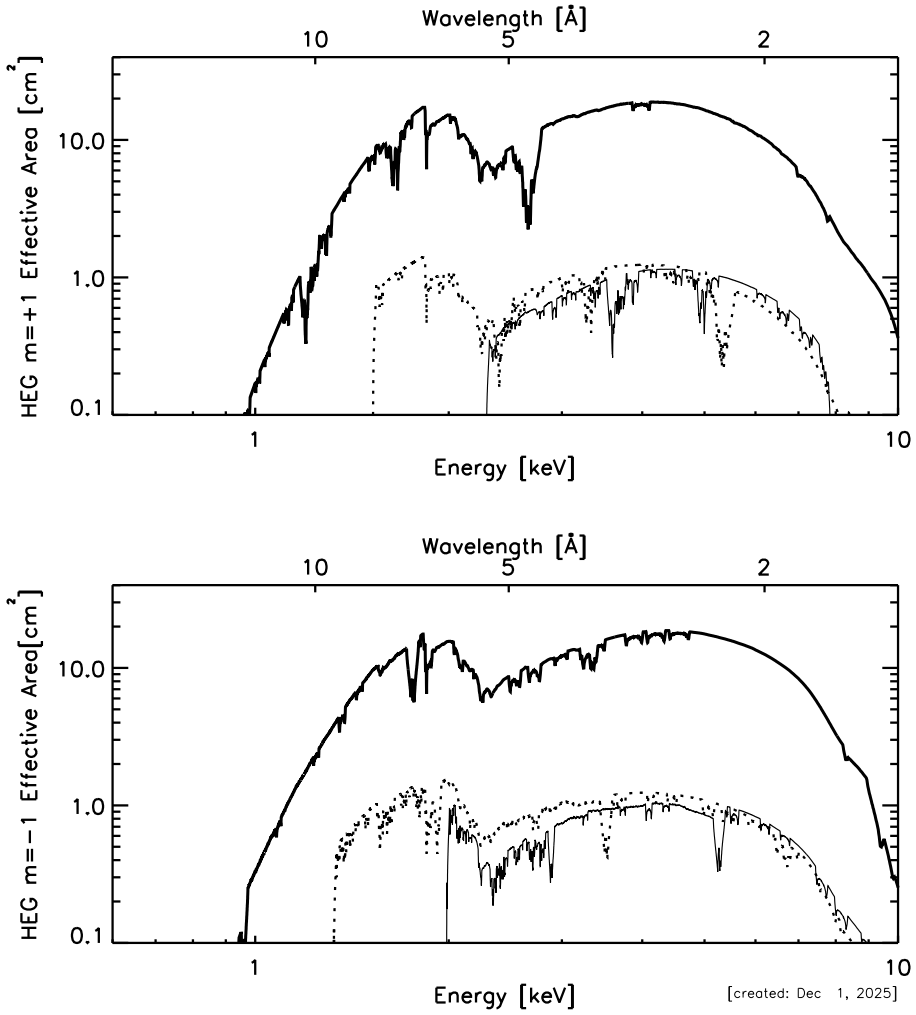

Figure 8.7: The HETGS HEG effective area, integrated over the PSF, is shown with energy and wavelength scales. The $m = +1, +2, +3$ orders (falling on ACIS chips S5, S4, S3; left to right) are displayed in the top panel and the $m = -1, -2, -3$ orders (falling on ACIS chips S0, S1, S2; left to right) are in the bottom panel. The thick solid lines are first order; the thin solid line is third order; and the dotted line is second order.

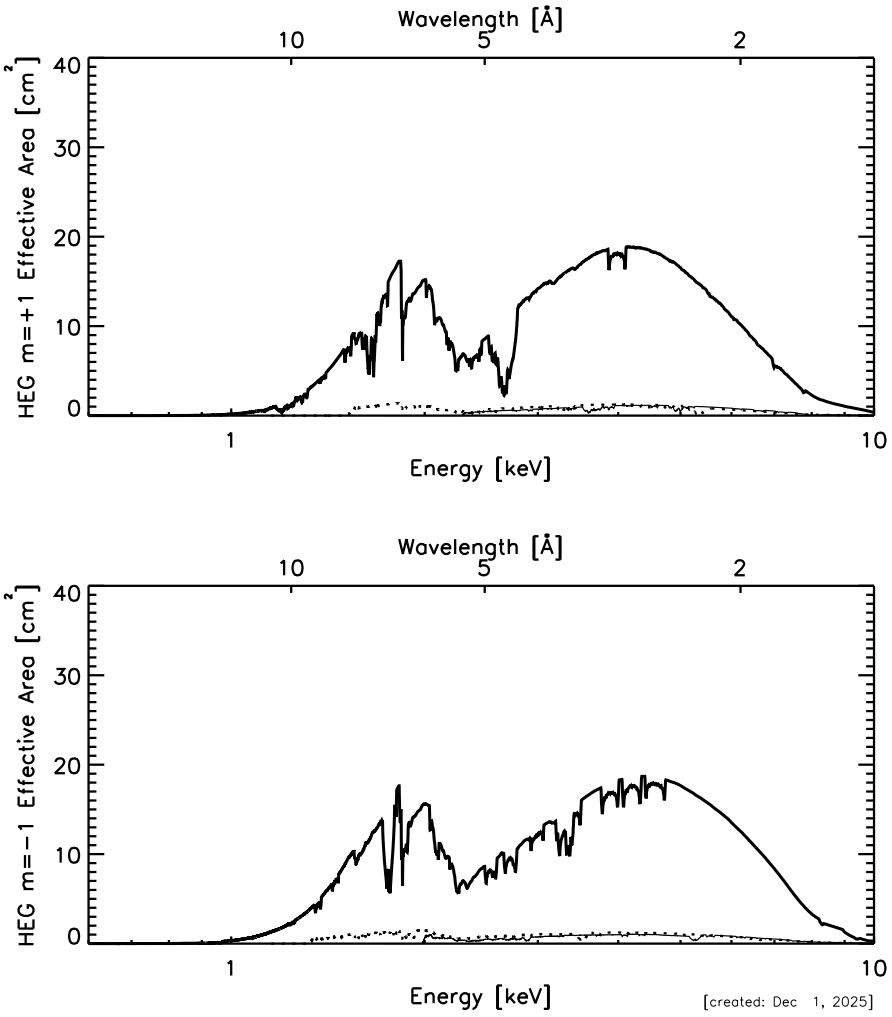

Figure 8.8: The HETGS HEG effective area: same caption as previous figure, except the vertical scale is now linear.

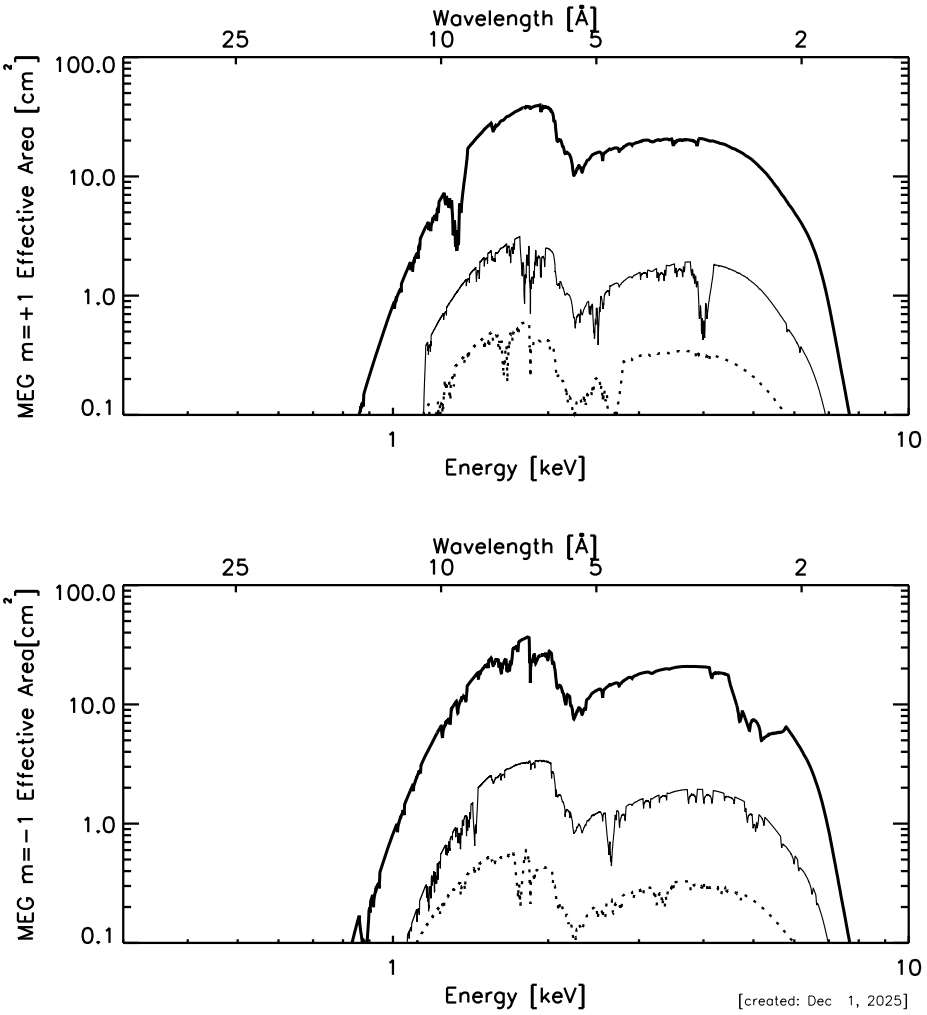

Figure 8.9: The HETGS MEG effective area, integrated over the PSF is shown with energy and wavelength scales. The $m = +1, +2, +3$ orders (falling on ACIS chips S5, S4, S3; left to right) are displayed in the top panel and the $m = -1, -2, -3$ orders (falling on ACIS chips S0, S1, S2; left to right) are in the bottom panel. The thick solid lines are first order; the thin solid line is third order; and the dotted line is second order.

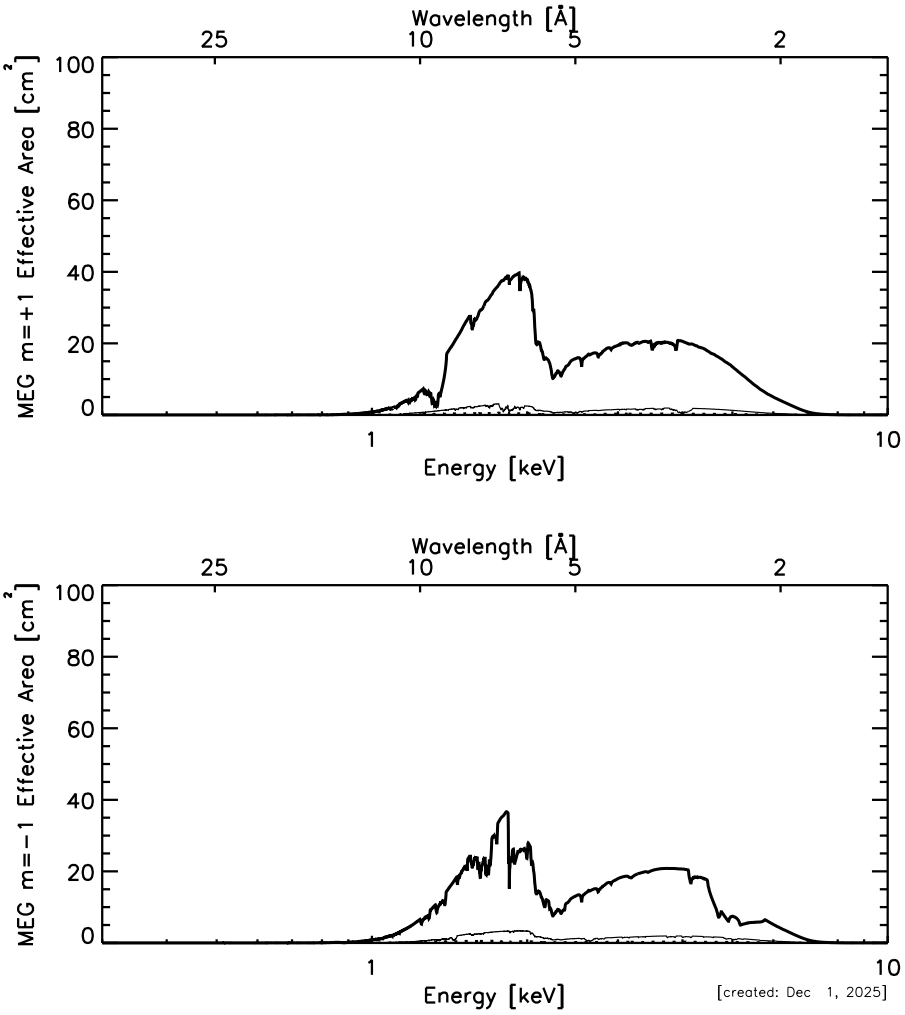

Figure 8.10: The HETGS MEG effective area: same caption as previous figure, except the vertical scale is now linear.

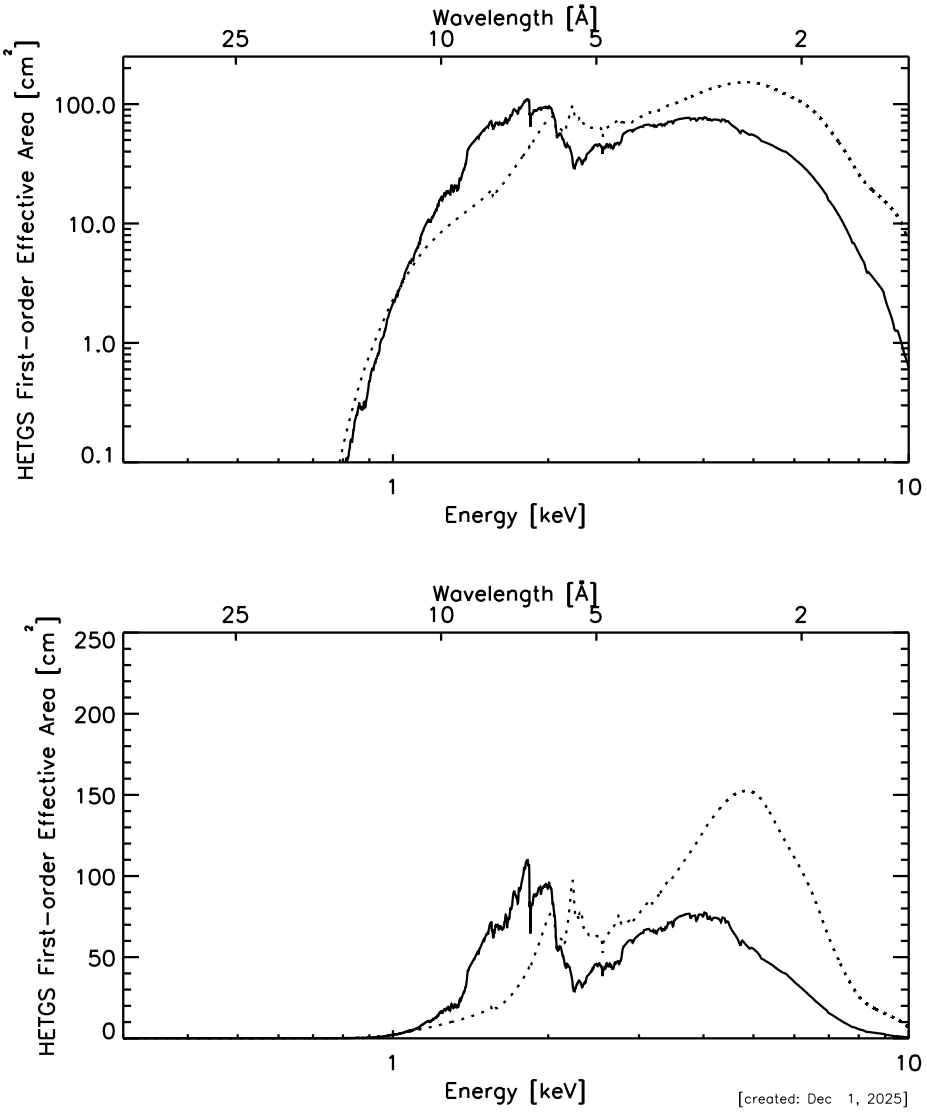

Figure 8.11: The modeled total first-order (solid curve) and zeroth-order (dotted curve) effective area, integrated over the PSF, of the HRMA-HETG-ACIS-S combination, as a function of energy. The first-order data are the same as those plotted in Figures 8.7 and 8.9. The plotted first-order values are the sums of the area at a particular energy from both orders (+/−) of both MEG and HEG spectra. Both a log-log and a log-linear version are shown.

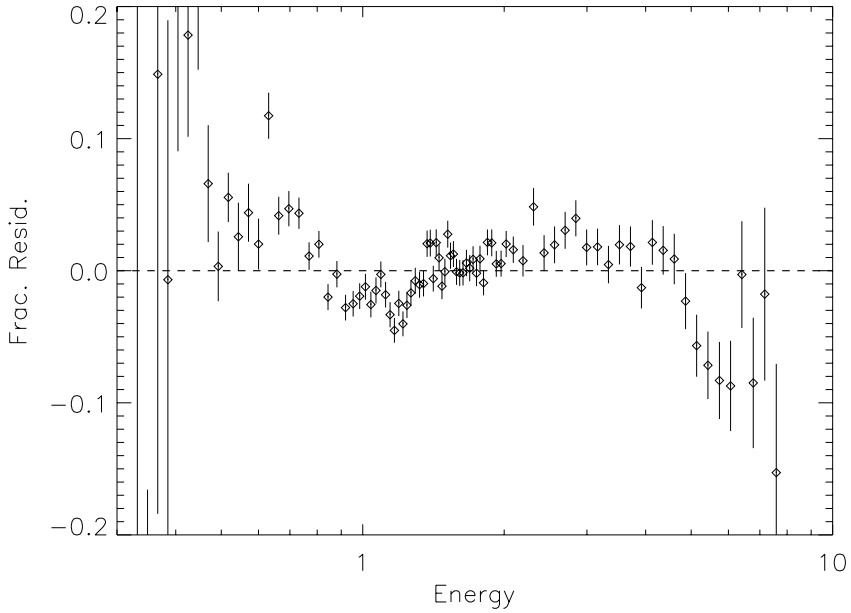

Figure 8.12: The average residuals to fits to the HETGS data for BL Lac objects using logarithmic parabola continuum models. For most of the HETGS range, the systematic deviations are not significant or are less than 3%. See Marshall (2012) for details.

A number of systematic effects at the 10% level have now been accounted for in the HETGS effective area via adjustments to the grating efficiencies. Observations of blazars (Marshall 2012) have been used to verify these corrections. Relative systematic errors from fits to logarithmic parabolas (with fixed, source-dependent column densities) are generally less than 3% (0.5 –8.0 keV) on small scales (see Figure 8.12). Preliminary results from cross-calibration with *XMM-Newton* indicate that measured fluxes agree to better than 10% (Smith & Marshall 2012). Cross-calibration with *XMM-Newton* and other X-ray telescopes (*NuSTAR*, *Swift*, and *Suzaku*) was addressed for a limited energy range by Madsen et al. (2016), who found differences of order 10% or less above 3 keV and as much as 15% against the *XMM-Newton* pn.

ACIS-S Order Sorting Effects

One of the advantages of using the ACIS-S as the HETG read-out detector is the ability of ACIS to determine the energy of detected X-rays. This crude (by HETGS standards)

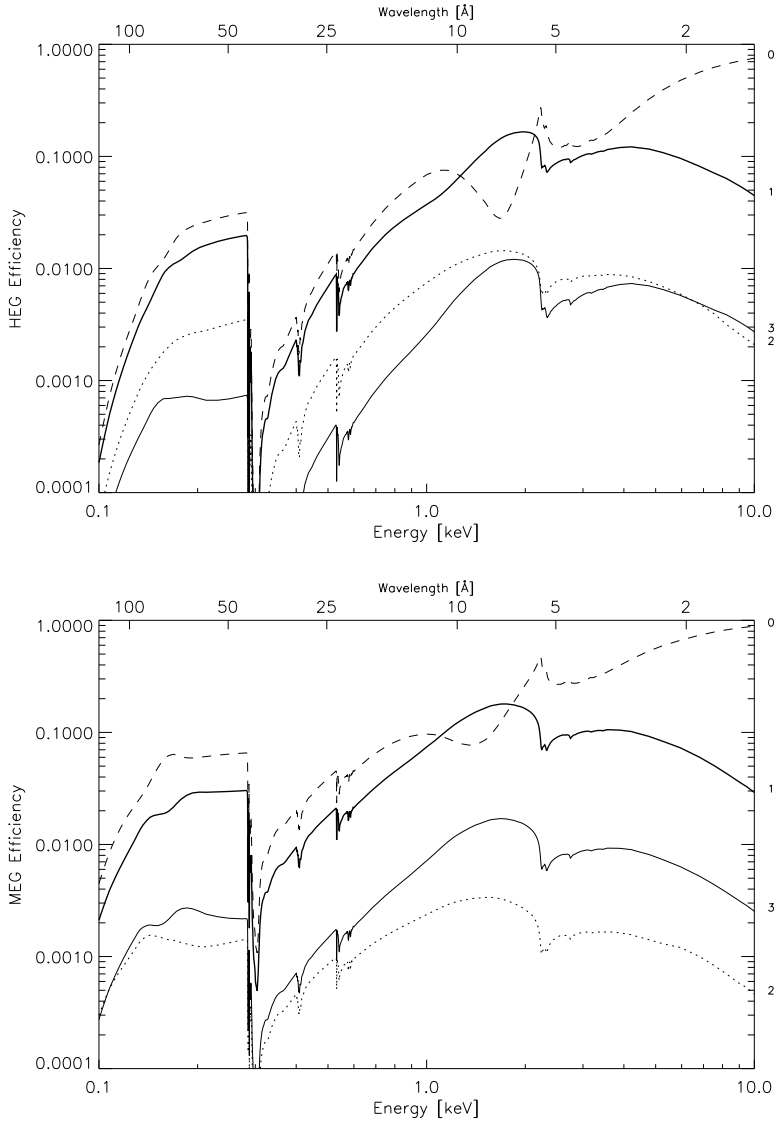

Figure 8.13: HEG (upper panel) and MEG (lower panel) efficiencies as a function of energy. The values plotted are the mirror-weighted efficiency into a single plus, minus, or zero order (labelled on the right edge). The dashed line is zeroth order; the thick solid line is first order. Note that the relative strengths of the third orders (thin solid lines) are comparable, whereas the second order strengths (dotted lines) are significantly different between the HEG and MEG.

energy measure can be used to determine the diffraction order of the photon, i.e., perform "order sorting," as shown in the "banana plot" of Figure 8.14.

During data analysis, this filtering is accomplished by utilizing two of the data columns supplied in the level 1.5 (or 2.0) FITS data file: the ACIS-determined energy, ENERGY, and the dispersion distance, $m\lambda$ = TG_MLAM. Ideally this order sorting would have perfect efficiency, that is, all first-order events would be correctly identified. In practice, a high sorting efficiency is achieved by accurately calibrating the ACIS ENERGY values and by accepting events in a large ENERGY range. The slight efficiency corrections that do arise are included in the ARF through values in an order sorting integrated probability (OSIP) file.

ACIS-S Pile-Up Effects

Figure 8.15 shows a closeup of the "banana plot" (ACIS-determined energy versus the dispersion distance in units of wavelength) for MEG minus-order events for an observation which exhibits pile-up (see Section 6.16) and thus mimics higher-order photons. One can encounter pile-up even in the dispersed spectra. The effect is most likely seen in first order spectra when observing bright continuum sources such as those found in the Galactic bulge. Pile-up, when it occurs, is most usually found in the MEG first order spectrum near the iridium edge at 2 keV where the HETGS effective area is the highest. Users analyzing data should note that not correcting for pile-up may introduce an artificial absorption edge. In these cases users may well wish to examine the spectrum in the third order to either salvage or correct a result.

ACIS-S BI / FI QE Effects

The ACIS-S array is made up of both back-illuminated (BI) and front-illuminated (FI) CCDs; chips S1 and S3 (see Figure 8.1) are BI devices and the rest are FI devices. These devices have different quantum efficiencies (QE) with the BI having greater sensitivity at lower energies; most notably the S1 BI device gives an increased effective area in the MEG minus first order between about 0.5 to 0.8 keV. The grating ARFs are created taking these QE differences into account.

ACIS-S Chip Gap Effects

The nominal ACIS-S aim point is on chip S3, about 2.0 arcmin from the gap between chips S2 and S3. Energies of gap edges in both dispersed spectra for the default aim point and for 3 offsets in both (+/−) Y directions are given in Table 8.2. For example, with zeroth order at the −0.66 arcmin Y-offset position, the gap between chips S1 and S2 spans the energy range 0.829–0.842 keV in the MEG spectrum (lower energies on S1). The observer is advised to *try to avoid placing known features of interest within three gap widths of the tabulated gap edges.* All HETGS observations are nominally dithered with an amplitude of ±8 arcsec. There will be reduced coverage in the spectral regions within

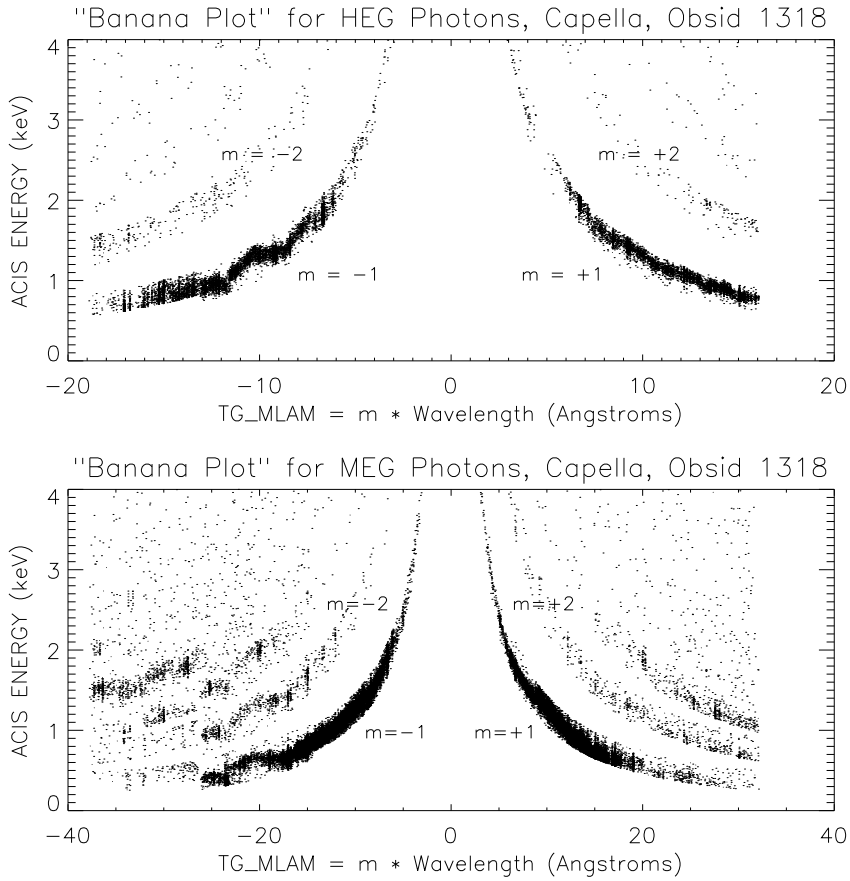

Figure 8.14: HEG (upper panel) and MEG (lower panel) "Banana Plots." A useful look at the HETGS data is obtained by plotting the ACIS-measured event ENERGY as a function of $m\lambda$ = TG_MLAM (or versus dispersion distance). These "banana plots" are shown here for HEG and MEG parts of the ObsID 1318 Capella observation. The various diffraction orders show up as hyperbolae. Events can be assigned to a diffraction order based on their location in this space. By accurately calibrating the ACIS ENERGY and by taking an appropriate acceptance region, events can be sorted by order with high confidence and efficiency. A "zig-zag" in the $m = -1$ events pattern is visible around -10 Å in the HEG plot and is due to uncorrected serial charge transfer inefficiency in the BI device S1, which produces a slow variation of gain across a node.

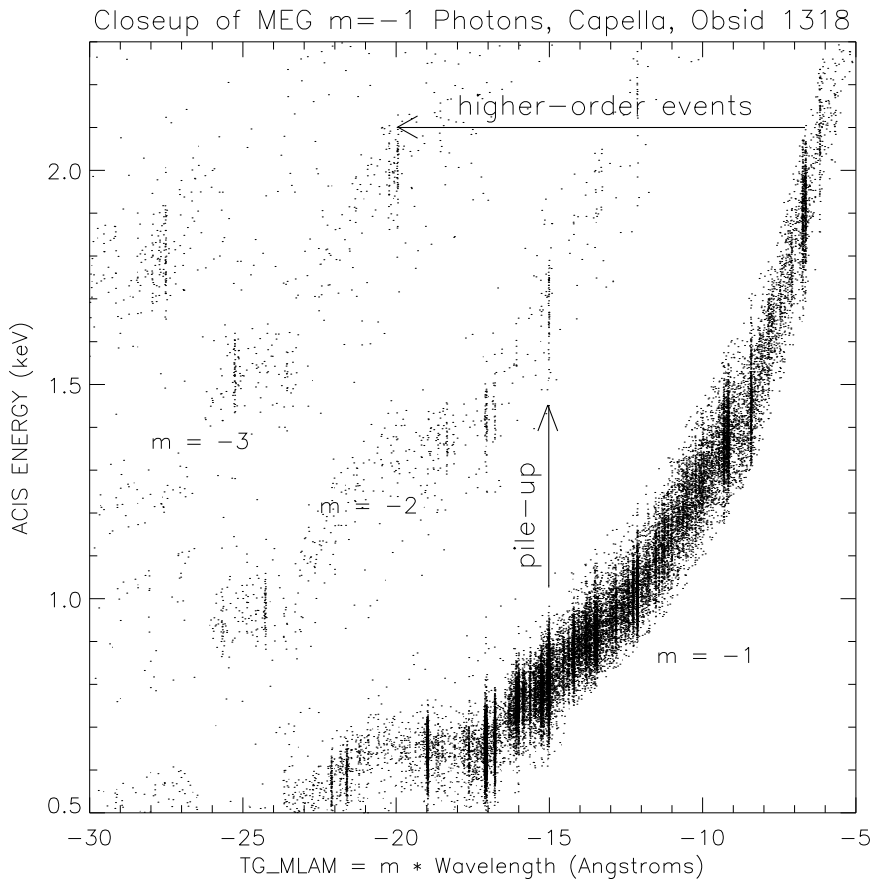

Figure 8.15: HETGS pile-up and higher-order events. Taking a close look at the MEG "banana plot" demonstrates how the ACIS ENERGY can be used to identify higher-order events and pile-up in an HETGS spectrum. The 3rd order of the ≈ 6.7 Å lines are clearly visible; the lines are only weakly present in 2nd order because the MEG 2nd order is suppressed. In comparison, the 15 Å line (and others) are so bright in 1st order that a fraction of the events (≈ 6 % here) pile-up and produce events with twice the ACIS ENERGY. Note that the 6.7 Å lines are better resolved in the high order spectrum.

one gap-width on either side of the gaps. The web-based Spectrum Visualization Tool (https://cxc.harvard.edu/cgi-gen/LETG/alp.cgi) displays where spectral features fall on the ACIS-S detector as a function of Y-offset and source redshift.

The values in the table are based on an effective gap size of 0.502 mm, corresponding to 10 arcsec on the sky. It is "effective" in the sense that the gap includes columns 1 and 1024 of the devices from which no events are reported. This value for gap size is approximate and accurate to about 2 pixels. The actual gap sizes vary slightly; more accurate values of the ACIS-S chip geometry are given in a *Chandra* X-ray Center data systems (CXCDS) CALDB file 'telD1999-07-23geomN006.fits' (and higher versions) and incorporated in *MARX* version 3.0 and higher. Relative to S3, where zeroth order is normally placed, the ACIS-S chip locations are calibrated to better than 0.2 pixels allowing accurate relative wavelengths.

8.2.2 HETGS Line Response Function

A high-resolution spectrum is created by the projection of events along the dispersion axis and binning the events into energy or wavelength bins as shown in Figure 8.2. The HETGS line response function (LRF) at a given wavelength is the underlying distribution which would result if the source were monochromatic at that wavelength. The LRF function is encoded in the grating RMF files. Examples from flight data are shown in Figures 8.16 and 8.17. To a good first approximation the core of the LRF can be modeled as a Gaussian, parameterized by a *Resolution*, ΔE or $\Delta\lambda$, given as the full-width at half-maximum of the Gaussian, 2.35σ. For the HETG the resolution is roughly constant when expressed as a wavelength. The *Resolving Power*, $E/\Delta E = \lambda/\Delta\lambda$, is a useful dimensionless measure of the spectrometer performance. Plots of the HETGS resolving power are presented in Figure 8.19.

The HETGS LRF is not simply a Gaussian and, as for other spectrometers, the response can be encoded at a higher level of fidelity through the use of response matrix files, RMF's. As explained below, the LRF (RMF) of the HETGS depends on all system components as well as the source spatial properties. Thus, LRF creation is carried out using a system model, e.g., the *MARX* ray trace software. A set of RMF's for a point source and nominal telescope properties can be created based on the latest LRF library in the *Chandra* CALDB which includes two Gaussian and two Lorentzian components to describe the LRF as derived from realistic *MARX* simulations; for examples, see the fitted LRF models in Figures 8.16 and 8.17.

The line response function can be decoupled approximately into three contributing components: the telescope PSF, the HETG effects in the dispersion direction, and HETG effects in the cross-dispersion direction. These are described below. With the exception of "HEG scatter," all effects described here are included in *MARX* version 3.0 (and higher) ray trace software.

The HRMA PSF "anomaly" has not been found to cause a problem in analysis. For zeroth order images, it is a very subtle feature, the HETGS is generally used with ACIS

Table 8.2: Table of HETGS Gap Locations

Y Offset		Grating	HETGS Gaps (keV)						
arcmin.	mm		S0	S0–S1	S1–S2	S2–S3	S3–S4	S4–S5	S5
1.00	2.93	MEG		0.508	0.969	10.724	1.183	0.561	0.367
			0.346	0.512	0.987	13.439	1.157	0.555	
		HEG		1.014	1.937	21.434	2.364	1.120	0.734
			0.692	1.024	1.973	26.859	2.313	1.109	
0.66	1.93	MEG		0.498	0.935	7.656	1.238	0.573	0.372
			0.342	0.503	0.952	8.946	1.209	0.566	
		HEG		0.995	1.869	15.302	2.474	1.144	0.744
			0.683	1.005	1.903	17.880	2.417	1.132	
0.33	0.97	MEG		0.489	0.905	5.992	1.296	0.585	0.378
			0.337	0.494	0.920	6.755	1.265	0.578	
		HEG		0.978	1.808	11.976	2.590	1.169	0.755
			0.674	0.987	1.839	13.500	2.528	1.156	
0.00	0.00	MEG		0.481	0.876	4.923	1.360	0.597	0.383
			0.333	0.485	0.891	5.426	1.326	0.591	
		HEG		0.961	1.751	9.838	2.718	1.194	0.765
			0.666	0.970	1.780	10.844	2.650	1.181	
−0.33	−0.97	MEG		0.472	0.849	4.177	1.430	0.611	0.388
			0.329	0.477	0.863	4.534	1.393	0.604	
		HEG		0.944	1.697	8.348	2.859	1.220	0.776
			0.658	0.953	1.724	9.061	2.784	1.206	
−0.66	−1.93	MEG		0.465	0.824	3.627	1.509	0.624	0.394
			0.325	0.469	0.837	3.893	1.467	0.617	
		HEG		0.928	1.646	7.250	3.015	1.248	0.787
			0.650	0.937	1.672	7.781	2.932	1.233	
−1.00	−2.93	MEG		0.457	0.799	3.194	1.599	0.639	0.400
			0.322	0.461	0.811	3.399	1.552	0.632	
		HEG		0.913	1.597	6.384	3.195	1.278	0.799
			0.643	0.920	1.621	6.793	3.102	1.263	

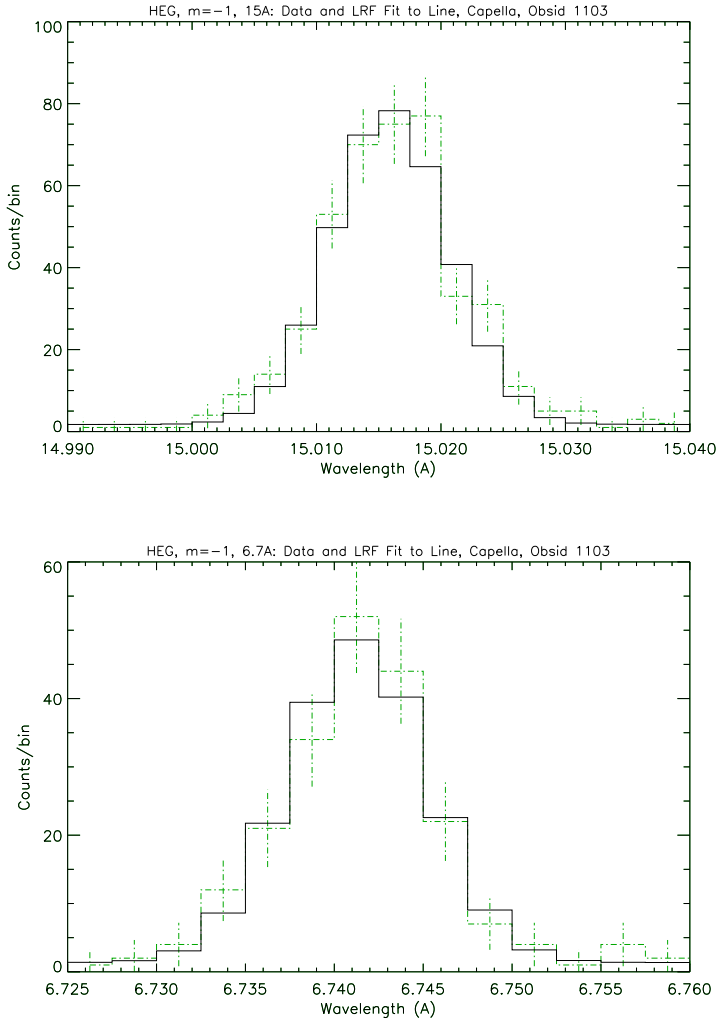

Figure 8.16: Representative Line Response Functions at two wavelengths for the HEG; 15 Å top, 6.7 Å bottom. Two of the bright lines in the HEG counterpart to the MEG Capella spectrum shown in Figure 8.2 have been fit by the instrumental LRF. The LRF is encoded in the HEG RMF created using *CXC* software and calibration parameters.

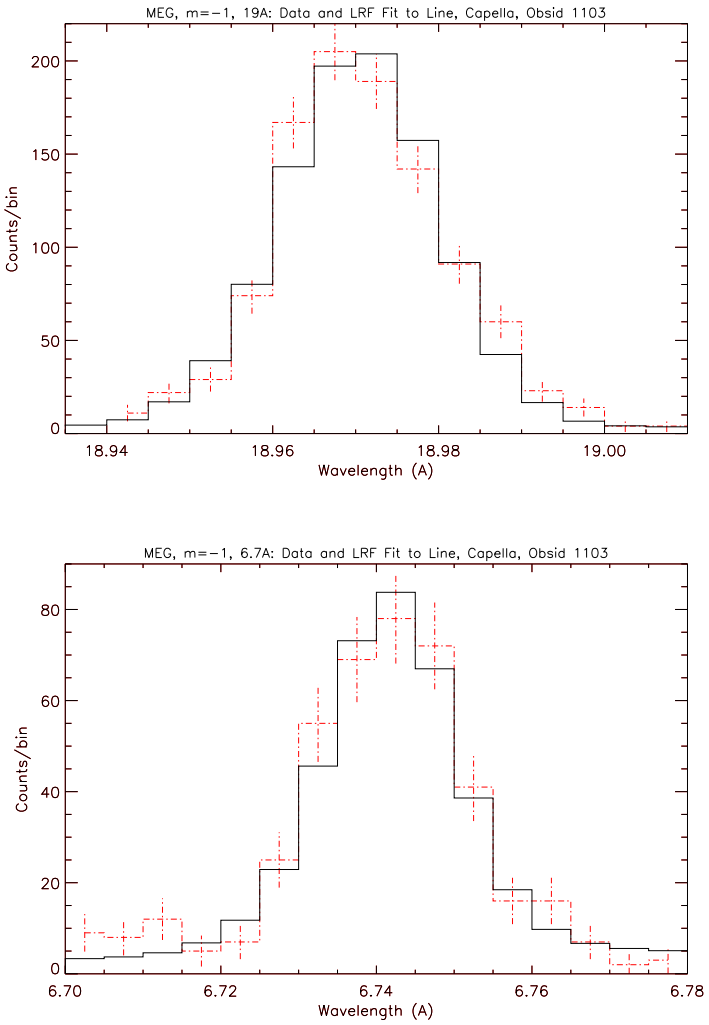

Figure 8.17: Representative Line Response Functions at two wavelengths for the MEG; 19 Å top, 6.7 Å bottom. Two of the bright lines in the MEG Capella spectrum shown in Figure 8.2 have been fit by the instrumental LRF. The LRF is encoded in the MEG RMF created using *CXC* software and calibration parameters.

and targets that are bright enough for grating spectroscopy generally show significant pile-up in the zeroth order images. See the ACIS chapter (Section 6.6.1) for more discussion of the possible impact on ACIS images. For dispersed events, the line response function has additional complication due to astigmatism, grating misalignment, and grating scatter that would make it very difficult to observe any effect specific to the PSF anomaly.

LRF: Telescope PSF and Zeroth Order

The HETG itself does not focus the X-rays emerging from the HRMA. Rather, *the Rowland design attempts to maintain the focal properties of the HRMA in the dispersion direction even as the focus is deflected by the diffraction angle* β. The 1-D projection of the telescope PSF onto the dispersion axis is thus at the heart of the HETGS LRF and can be thought of as the "zeroth-order LRF." Ground testing showed no measurable effect on the telescope PSF due to the HETG insertion; this was used to good advantage for the now famous image of the Crab Nebula and its pulsar, ObsID 168, where the jet and swirling structure are seen in the zeroth-order HETGS image. Thus, the zeroth-order image in an HETGS observation can be used to determine the telescope contribution to the LRF.

Image quality depends on many factors, and so, while a nominal LRF can be modeled, the detailed LRF will be observation dependent at some level. Factors in the telescope PSF performance include: source size and spectrum, HRMA properties, focus setting, detector effects (e.g., pixel quantization), aspect solution and reconstruction effects, and data analysis operations (e.g., pixel randomization). While all of these effects can be modeled, the "proof of the pudding" is in the as-observed zeroth order image.

As an example, results from the observation of Capella (ObsID 1318, Figure 8.1) are shown in Figure 8.18. Both the zeroth-order event distribution and its 1-D projection indicate that the zeroth-order is heavily piled up with an unpiled event rate of order 10 events per frame time (per few square ACIS pixels). The wings of the PSF are visible but the core shape and intensity have been severely distorted. However, because the ACIS-S CCDs have their columns perpendicular to the (average) dispersion axis, the "frame-transfer streak" events or "trailed image" (see Section 6.13.1 in Chapter 6) can be used to create an accurate zeroth order LRF that is not affected by pile-up, as shown. For point-sources such as Capella, measurements of the FWHM of the zeroth-order streak events for selected observations over the first two years of HETGS operation show FWHM values generally in the range 1.46 to 1.67 ACIS pixels with an average of 1.57 ACIS pixels. Thus, by appropriately examining the zeroth-order image and its LRF, one can get a good idea of the expected width of a truly monochromatic spectral line, and determine whether or not any broadening seen in a dispersed order is a spectral property of the X-ray source.

LRF: Dispersion Direction

As mentioned, the profile in the dispersion direction defines the instrument spectral resolution, ΔE or $\Delta\lambda$. The resolution function has two main terms with different dependences

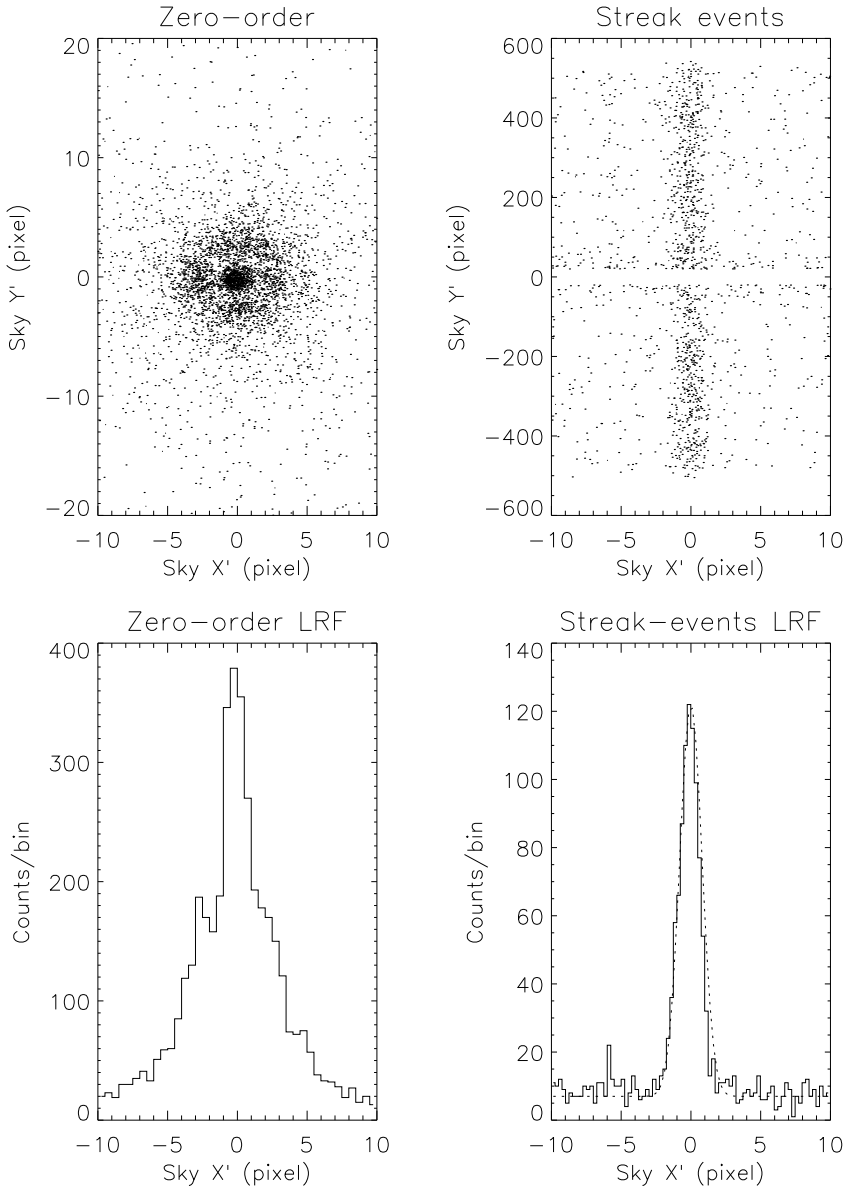

Figure 8.18: HETGS zero order and Frame transfer Streak (trailed image) for ObsID 1318 of Capella. The sky coordinates, X, Y, have been rotated so that the frame-transfer streak is along the Y' axis, hence Y' is parallel to the CCD detector Y axis (`CHIPY`) and X' is approximately along the average HEG-MEG dispersion axis. The left-side panels show the detected zero-order events and their 1-D projection; pile-up is evident by the enhanced wings relative to the suppressed PSF core. The right-side panels show the frame-transfer streak events and their 1-D projection; the dotted line is a Gaussian fit to the data.

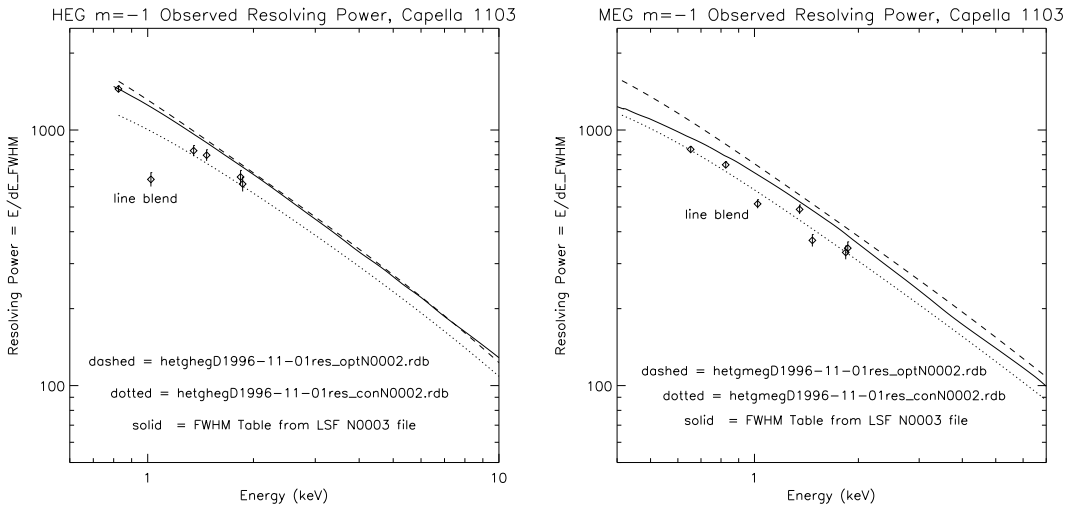

Figure 8.19: HEG and MEG resolving power ($E/\Delta E$ or $\lambda/\Delta\lambda$) as a function of energy for the nominal HETGS configuration. The resolving power at high energies is dominated by the telescope PSF; at low energies grating effects enter but do not dominate. The "optimistic" dashed curve is calculated from pre-flight models and parameter values. The "conservative" dotted curve is the same except for using plausibly degraded values of aspect, focus, and grating period uniformity. The cutoff at low-energy is determined by the length of the ACIS-S array. Measurements from the HEG and MEG $m = -1$ spectra, e.g., Figure 8.2, are typical of flight performance and are shown here by the diamond symbols. The values plotted are the as-measured values and therefore include any natural line width in the lines; for example, the "line" around 12.2 Å is a blend of Fe and Ne lines. The solid line gives the resolving power encoded in the RMFs generated using the current CALDB.

on energy: the image blur from the mirror described above and that caused by grating period variations which come in through the dispersion relation and are described here.

From the grating equation, $m\Delta\lambda = p\Delta\beta cos\beta + \Delta psin\beta \approx y\Delta p/R_s + p\Delta y/R_s$, where p is the grating period, β is the dispersion angle, y is the dispersion distance and R_s is the (fixed) Rowland spacing. The two terms of interest are on the right side of $\lambda/\Delta\lambda = (\Delta p/p + \Delta y/y)^{-1}$. The grating fabrication process produced tightly distributed grating periods ($\Delta p/p < 2.5 \times 10^{-4}$) so that the first term is important in the spectral resolution only at very high dispersion (low energy). The mirror point response function has a nearly constant size Δy and dominates the resolution over most of the HETGS band, as shown in Figure 8.19. At very low energies there is a contribution from variation in the grating periods. These variations are taken into account in the *MARX* simulator.

During ground testing, a low level of incoherent dispersion (or "scattering") was discovered in HEG spectra. This scattering effect distributes a small amount of the flux along the dispersion direction. The total power involved is only 1.0% of the total in first order

but the light is irregularly distributed between the coherently dispersed orders. Assuming that the power distribution scales with the first order dispersion distance, there is no more than 0.02% of the first order flux in any bin of width 0.01λ. There has been no scattering detected in the MEG spectra to a level of order 100 times fainter than in the HEG. See the HETG Ground Calibration Report listed at the end of this chapter for further details. The effects of scattering from the grating are likely to be negligible for most observations.

A spatially resolved source may appear to show spectrally resolved emission lines. This effect is described in some detail in a *Chandra* Newsletter article (Marshall 2017). Quantitatively, one may model a Gaussian line with several contributions to assess the effect of spatial broadening on line width. Suppose that all effects have profiles given by Gaussians, with the instrumental broadening given by σ_i in detector space, Doppler broadening given by σ_v in velocity space, and spatial broadening given by σ_θ, in imaging space. Converting each to their effect on the total broadening in physical coordinates x on the detector for wavelength λ and remembering the grating equation (see Eq. 8.1) for small dispersion angles gives

$$\sigma_x^2 = \sigma_i^2 + (\frac{Rm\lambda}{pc})^2\sigma_v^2 + (F\sigma_\theta)^2 \tag{8.2}$$

where m is the grating order, p is the grating period (see Table 8.1), R is the Rowland spacing of the HETGS (see Table 8.1), and F is the focal length of the HRMA.

Eq. 8.2 shows that imaging is more important than instrumental broadening when $F\sigma_\theta \gg \sigma_i$; i.e., when a source is resolved. Instrumental and imaging terms dominate when $m\sigma_v$ is small but as m increases, the line width is dominated by Doppler broadening. Thus, spatial extent can be ignored only if $\sigma_v \gg \frac{3400}{m\lambda}\sigma_\theta$ km/s for λ in Å and σ_θ in arcsec. For a resolvable source with $\sigma_\theta = 1$ arcsec and examining the Fe Kα line at 1.94 Å using the high energy gratings, then spatial extent is comparable to Doppler broadening when $\sigma_v = 1750$ km/s. In third order, however, this cross-over value drops to 600 km/s. It is clearly advantageous to examine the high order HETGS data where feasible.

LRF: Cross-Dispersion Direction

The profile in the cross-dispersion direction is dominated by three effects: mirror blur, grating roll variations, and astigmatism (as a by-product of the Rowland design which optimizes spectral resolution). The cross-dispersion profile that results from astigmatism is slightly edge brightened, but quasi-uniform, with a length at the Rowland focus of $2R_f y^2/R_s^2$, where y is the dispersion length and R_f is the radius of the ring of facets on the HETGS structure and dominates the size of the cross dispersion profile at large dispersion.

The spread of facet roll angles (defining the dispersion direction for each facet, and not to be confused with the spacecraft roll angle), $\Delta\phi$, contributes a cross-dispersion term of order $y\Delta\phi$. Sub-assembly measurements predicted $\Delta\phi$= 0.42 arcmin RMS. However, analysis of ground test measurements lead to a somewhat larger and more complex roll angle distribution for the gratings. In addition, six misaligned MEG facets were discovered

during ground testing. The inferred facet roll angles were misaligned from the average dispersion direction by 5–23 arcmin. On average, each facet contributes only 1/192 of the flux at any given energy, so the cross dispersion profile has small deviations in the form of peaks displaced from the main distribution.

To include explicitly the MEG misaligned gratings, *MARX* uses "sector" files which allow the specification of grating alignment and period parameters for certain regions (sectors) of each of the four shells. Using these files, the agreement between ground calibration and flight data is very good. For the MEG the misaligned gratings are explicitly included and the rest of the gratings' $\Delta\phi$ term is modeled as the sum of two Gaussian distributions centered at $+1$ and -1 arcmin with respect to the nominal axis, each with an RMS value of 1.5 arcmin. For the HEG a more pronounced bi-gaussian distribution is observed and modeled: the Gaussians are offset by -1.35 and $+1.65$ arcmin, each with a 1 arc minute RMS, and in a relative ratio of 55:45.

In each case, these effects are accurately included in *MARX* version 3.0 (and above). Flight data for the Crab pulsar (ObsID 168) are shown in Figure 8.20 for the MEG and in Figure 8.21 for the HEG. Note that these profiles are on top of a significant baseline due to the presence of the Nebula. The asymmetry in the MEG profile caused by misaligned gratings is quite clear at large dispersions.

Finally, Figure 8.22 shows how the total observed flux depends on the width of the extraction region in the cross-dispersion direction. The figure can be used to estimate the reduction in flux if analysis using a narrow extraction window, smaller than the nominal 4 arcsec full width, is planned.

Extended and Off-Axis Targets

The observation of extended sources with the HETGS adds complexity. Chiefly, the position of an event in the focal plane is not a unique function of the position within the source and the photon energy. The source extent, measured by the zeroth-order image size, can effectively increase in several ways: the telescope is out of focus, the source is off-axis, or there is a natural extent to the astrophysical source. Figure 8.23 illustrates the chief consequence for extended sources—a degradation of the apparent spectral resolution. In Figure 8.24 similar resolution curves are shown as a function of the source off-axis angle.

The discussion and plots above assumed that the source has no spatially dependent variations in the spectrum. The more general case of extended sources with spatially varying spectra is briefly discussed below in Section 8.5.4.

8.2.3 Background

Since the HETG is always used in conjunction with a focal-plane detector, spectra from the HETGS will have background events determined by the detector's intrinsic and environmental backgrounds. The cosmic background, folded through the HETGS response, will likewise contribute background events. In addition to these detector-dependent backgrounds, there are additional grating-dependent effects such as scattering from the gratings

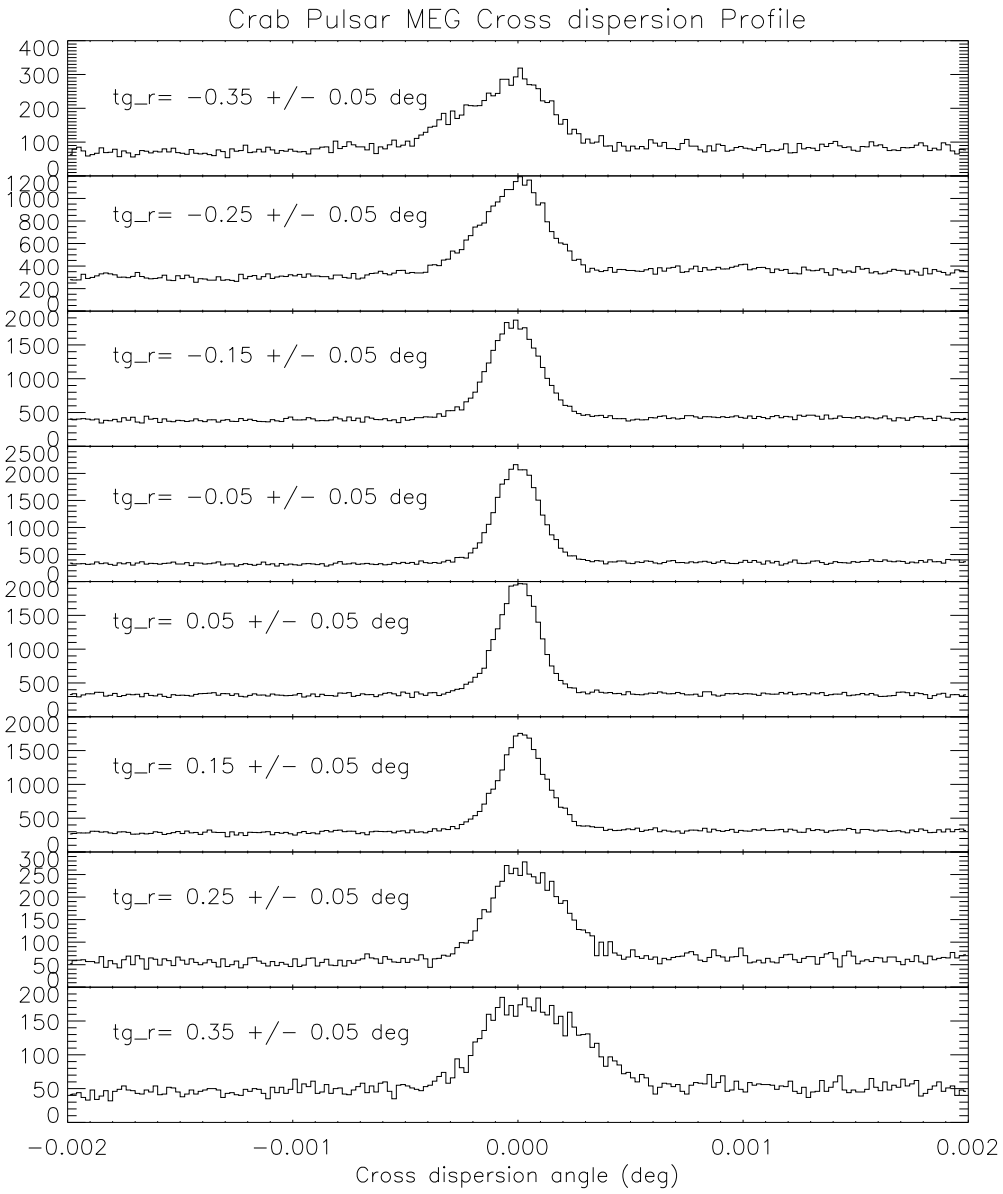

Figure 8.20: The cross dispersion profile is shown for eight slices of the dispersed MEG spectrum of the Crab pulsar. There is an asymmetry caused by misaligned gratings that becomes most evident at large dispersion.

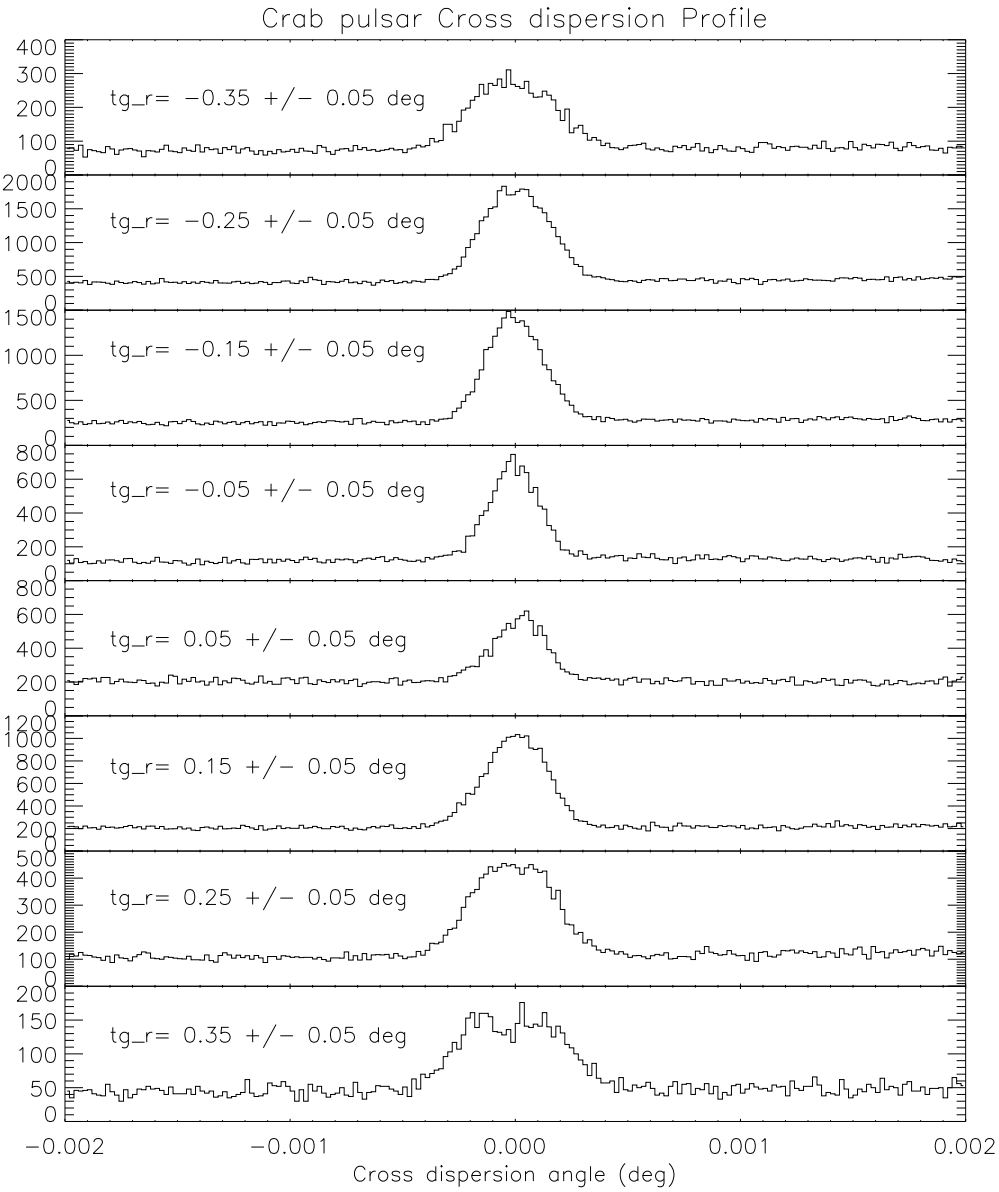

Figure 8.21: As in Figure 8.20, the cross dispersion profile is shown for the HEG spectrum of the Crab pulsar. The profile is symmetric but broadens significantly at large dispersion.

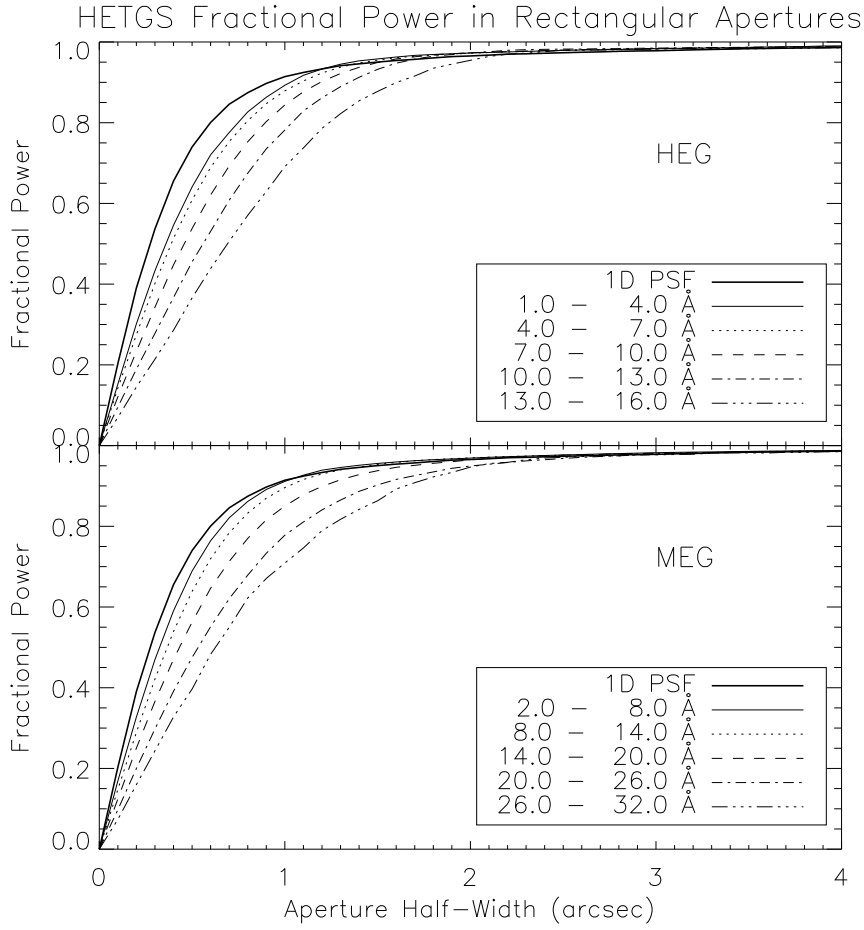

Figure 8.22: Enclosed power distributions are computed for five wavelength intervals for both the HEG (top) and the MEG (bottom). The observation of Mk 421 (observation ID 1714) was used.

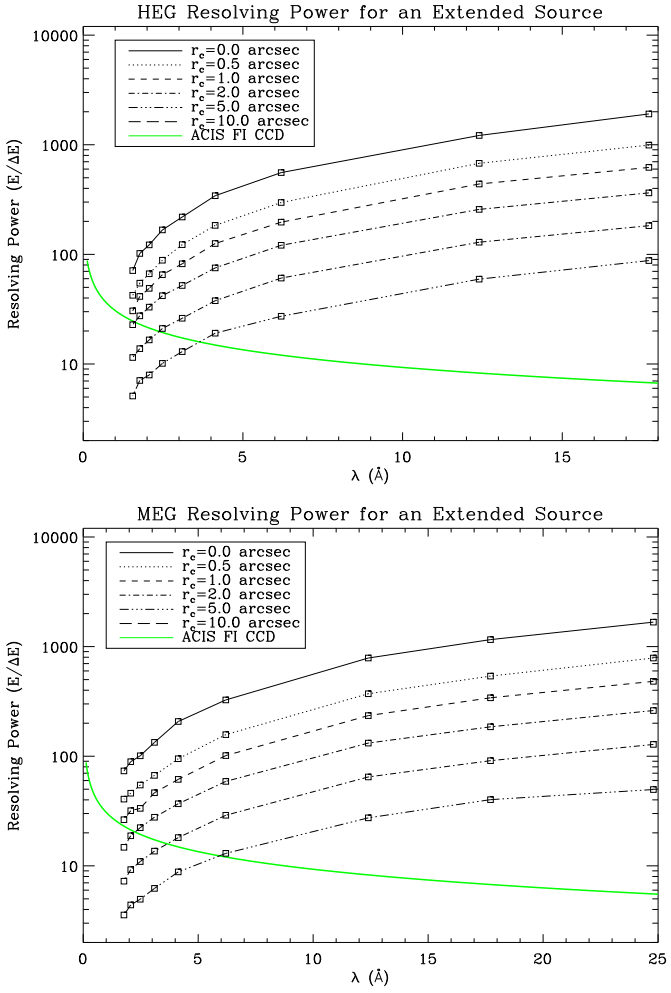

Figure 8.23: The effects of source size on the apparent HETGS spectral resolution. This *MARX* simulation uses a cluster (of galaxies) Beta model for the surface brightness profile. The Beta model is parameterized by a core radius (r_c) which represents the extension of the source. The effect on the apparent resolving power (E/ΔE) is shown as a function of photon energy for source sizes of 0, 0.5, 1, 2, 5, and 10 arcsec. The spectral resolution of an ACIS FI CCD near the framestore region is shown for comparison.

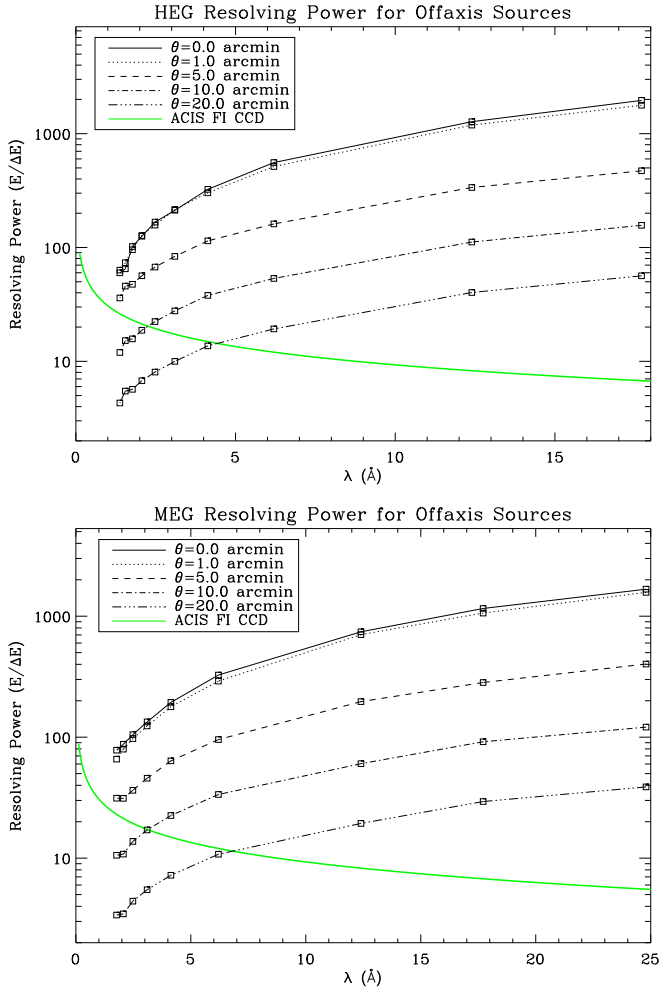

Figure 8.24: The effects of off-axis pointing on the HETG grating spectral resolution. *MARX* was used to simulate an observation of a point source at increasing off-axis positions. The effect on the resolving power (E/ΔE) is shown as a function of photon energy for off-axis angles of 0, 1, 5, 10, and 20 arcmin. The spectral resolution of an ACIS Front-Illuminated (FI) CCD at a point near the framestore region is shown for comparison.

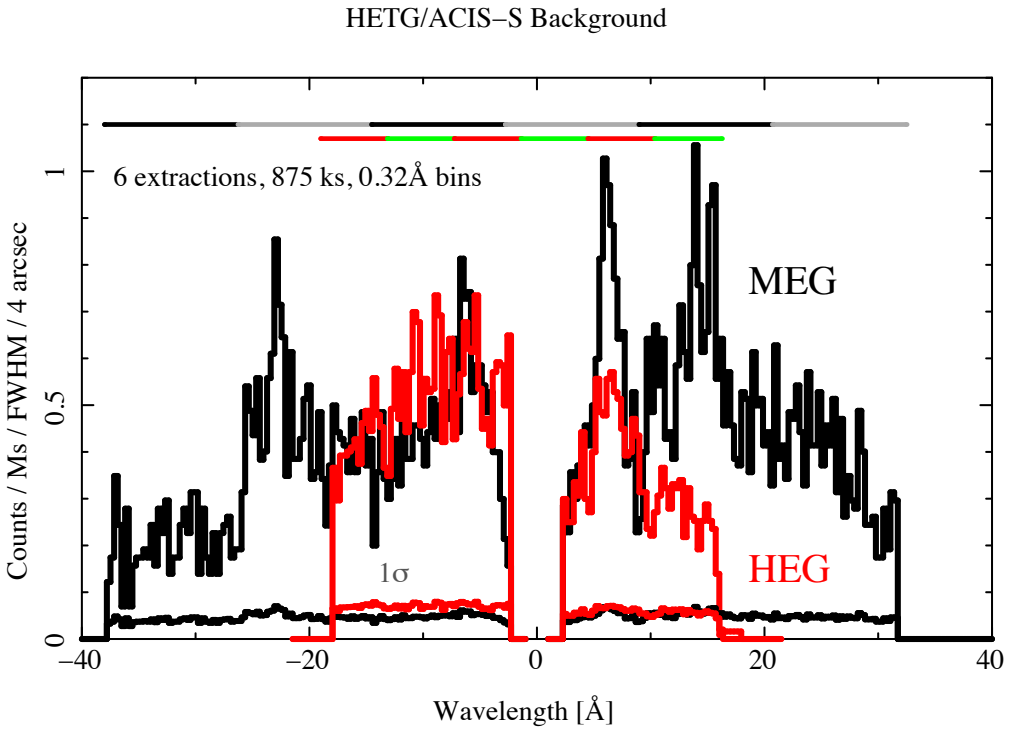

Figure 8.25: The first orders' HETGS background count rate spectra, scaled to the spectral resolution along the dispersion axis and for the default cross-dispersion extraction width. The 1σ uncertainties due to counting statistics are shown at the bottom. The approximate locations of each CCD, from S0 at the left to S5 at the right, are shown by the alternating color bars for the MEG (upper bar) and HEG (lower bar). This figure is adapted from a memo by Huenemoerder (2017).

which will produce extraneous photons in locations unexpected on the basis of the simple grating equation. One such effect is the scattering along the dispersion direction described in Section 8.2.2.

Figure 8.25 shows the HEG and MEG count spectra of the background from an analysis by Huenemoerder (2017). Data from six long observations of faint or heavily obscured sources were used in a standard extraction, applying normal ACIS energy and grade filtering and spatial selection to avoid sources. This plot can be used to estimate the background in a dispersed spectrum at a particular wavelength for proposal purposes. "Streaks," short-lived events observed in the S4 detector, have been removed; otherwise, the background in +1 order would be significantly higher and would show more structure. This background is comparable to that obtained when the detector is stowed, indicating that the background is primarily instrumental, after filtering on ACIS energy appropriate for HETGS observations.

Table 8.3: Comparison of HETGS 0th and 1st Orders

Energy Range (keV)	r
0.5 – 1.1	0.98 ± 0.05
1.1 – 2.4	1.02 ± 0.02
2.4 – 7.7	0.987 ± 0.013

8.2.4 Absolute Wavelength

The HETGS-measured wavelength depends, as the grating equation implies, on knowing the diffraction angle, the diffraction order, and the grating average period. The angle depends on knowing the HETGS geometry, specifically the Rowland spacing and the ACIS-S pixel size and configuration. Preliminary comparisons between measured and expected emission line wavelengths indicates an agreement to the accuracies listed in Table 8.1. Systematic wavelength errors are now at the 100 km/s level.

8.2.5 Comparing 0th and 1st Orders

The 0th and 1st order spectra were compared using six sources from deep observations of the bulge of M 31 (ObsIDs 15273-78 and 15561-3). These were chosen because the sources were faint enough that pileup is not a dominant effect and the total exposure in the dispersed spectra was long enough for accurate spectral fits. Table 8.3 gives the results of comparing the combined spectrum obtained from 0th order to that of 1st order (both $+1$ and -1); the ratio r is defined as the factor applied to the 1st order HETG spectral fit in order to match the 0th order ACIS-S spectral fit. The data were divided into energy ranges that isolate the background uncertainty due to extended emission in the M 31 bulge that primarily affects the lower energy range. In all cases, the results are consistent with unity within (90%) uncertainties, indicating that the 0th order effective area is consistent with that of 1st order.

8.3 Calibration Status

The calibration of the HETGS is based on extensive laboratory tests, system-level ground measurements, and flight data and analyses. Because the HETGS involves the HRMA, HETG, and ACIS-S as well as aspect system properties, calibration of all these components is important to the HETGS calibration. Details of the present state of the HETGS calibration are available at https://space.mit.edu/ASC/calib/hetgcal.html; see also Marshall *et al.* (2004), Weisskopf *et al.* (2004), and Marshall (2012).

Updates to the HETGS calibration web page over the past year include a memo on a Si-K edge residual (Schulz 2022) and a note about how the HETGS calibration is not significantly affected when the ACIS-S focal plane is "warm" (Section 8.4.8). See Section

6.22.5 for details of warm operations. In progress is an update to the HETGS high order efficiencies that should provide better agreement with 1st order spectra.

8.4 HETG Operations

8.4.1 Flight Events and Anomalies

There have been no flight anomalies with the HETG *per se*. There have been some problems with the HETG and LETG grating insertion/retraction mechanisms. To date these have been limited to failure of some of the limit switches which are used to sense the gratings' position. In 2000-May there was failure of the HETG A-side electronics retraction limit switch indicating that the HETG was not retracted, when in fact it was. Switching to the redundant B-side limit switch worked until 2000-June when it too would not indicate that the HETG was retracted. Subsequently operational procedures have been changed to determine when the gratings are properly retracted. There have been no impacts to the science program.

Because the HETG insert limit switches continue to function and because the HETG is inserted against a hard stop, **these anomalies have had no effect on the HETG wavelength scale.**

8.4.2 Operational Constraints

With the exception of operational constraints on the focal-plane detectors, there are no operational constraints for the use of the HETGS from the proposer's point of view. The HETG is placed in the stowed position during passage through the radiation belts, a time when no data can be taken. Additional functional constraints include preventing both the HETG and LETG from being simultaneously commanded into position, which could cause a mechanical interference. Finally, a "failsafe" command, once used, will permanently retract the grating. A decision to issue the failsafe command will not be taken without a thorough review including the *Chandra* Project and NASA Headquarters.

8.4.3 Output Data

There are no data from the HETG itself. The data are generated by the focal-plane imager in its format (e.g. Figure 8.1).

8.4.4 Performance Monitoring, Health and Safety

The HETG itself has only a few thermal and mechanical switch sensors associated with it. These sensors are examined routinely as part of the health and safety monitoring of the Observatory. HETGS performance is monitored by means of the calibration observations (Section 8.3).

8.4.5 Thermal Response Time

There is a negligible thermal time constant for the HETG to equilibrate from "near-wall" storage to "in-use" temperature environments. The temperature dependence of the resolution and energy scales have been minimized through use of low-expansion material ("Invar") and single-point-mounted facet-frames. Thus, the support structure may expand or contract, but the facets will not.

8.4.6 Observation Frequency/Duty Cycle

Currently, the longest continuous observations are limited to the time between passage through the radiation belts (section 3.4.2).

8.4.7 Radiation Considerations

The main radiation concern for the HETG concerns the polyimide support material. Thin membranes of this material, used for proportional counter windows operating under a pressure differential, have been tested for the effects of radiation damage on leak rate. No increased leak rate was encountered after a dosage of 9 krad. In these tests, the mechanical integrity of the material—the key issue for the HETG—was severely tested by the ability of the window to maintain the pressure differential of order one atmosphere. Loss of mechanical integrity has been reported in the literature, but only after exposures of 1000 MRads. The estimated proton dose to the HETG polyimide is of order 1 kRad per orbit when the HETG is inserted, and much lower values when HETG is in its stowed position. Current practice is to have the HETG retracted during radiation passages; however, even if it were left inserted the total exposure would be ≈ 1 MRad over 10 years, well below the 1000 MRad level.

A secondary concern would be changes to the gold grating bars (which, when in place, face the HRMA) due to sputtering by particles, particularly for the high-aspect ratio HEG gratings. Diffraction order ratios are sensitive to these changes. To date, after flight experience and laboratory radiation tests, there is no evidence that this concern is anything other than intellectual.

8.4.8 Operating with a Warm ACIS-S Array

Generally, the temperature of the ACIS-S array has little effect upon the HETGS data. There is no concern that the spectral resolution of the HETGS would depend on focal plane temperature, as the resolution depends purely on the HRMA point response function and the physical properties of the gratings themselves, both of which are thermally decoupled from the focal plane. An analysis in 2020 indicated that HETGS observations are not significantly affected as ACIS-S is operated warm. Marshall (2020) gives details on how ACIS RMF changes due to running ACIS-S warm have a negligible effect upon order selection in HETGS data. Schulz (2023) examined HETGS data for ACIS-S temperature of -105 °C, finding that spectra are not affected for SIM-Z locations below -6 mm; at

the nominal aimpoint, there appears to be a loss of flux below 3 Å in the HEG −1 order spectrum. See Section 6.22.5 for more information about warm ACIS operations.

8.5 Observation Planning

The following sections provide assorted information and topics relevant to planning an HETGS observation. See also the HETGS observation planning web page: https://space.mit.edu/ASC/calib/hetg_GO_info.html

8.5.1 General Considerations

- Orientation of multiple (or extended) sources: One may need to specify a restricted range of spacecraft roll-angles to avoid overlapping spectra from multiple targets in the field, or to arrange that the dispersed spectra from particular features of an extended target do not similarly overlap. Note that roll angle constraints usually will lead to restrictions on the dates dates of target availability. See Chapter 3.

- Offset pointing: Pointing offsets may be specified and used to include or exclude nearby sources, to keep an important spectral feature clear of the gaps between chips, to put a particular low-energy feature on the higher efficiency BI chip S1, etc. Offsets greater than one or two arcmin will, however, degrade the image quality—Chapter 4—which in turn broadens the LRF (Section 8.2.2).

8.5.2 Choice of Focal-Plane Detector

The HETG was designed for use with the ACIS-S detector (see ACIS-S Considerations), but can be used with other detectors including HRC-I (see HRC Considerations). Details concerning the detectors may be found in Chapters 6 and 7.

ACIS-S Considerations

Some considerations specific to ACIS-S are:

- Operating mode of ACIS-S: The ACIS-S array can operate in many modes, giving control over e.g., the read-area, pixel-binning, and read-frequency. The selection of the appropriate operating mode and its ramifications for the experiment is one of the most important that the user faces. A careful reading of Chapter 6 is recommended. The proposer should *pay special attention to the pros and cons of designating optional CCDs* for their HETGS observation, as opposed to simply requesting the entire ACIS-S chip set (even with the optional settings). See Section 6.22.1 for details on designating optional CCDs. See Table 8.2 to determine the energies that correspond to various chips; for example for a Y offset of 0.0, the energy range of the MEG spectrum that would be found on S0 is 0.334–0.483 keV while the HEG spectral

range is 0.668–0.965 keV. These energy ranges will not be found in the $m = -1$ spectra if S0 is dropped.

- Selecting the aimpoint: The capability of moving the SIM along the spacecraft Z-axis (the cross-dispersion direction) is useful for placing the image (dispersed spectra and zeroth order) closer to the ACIS chip read-outs. This placement minimizes the effects of the row-dependent energy resolution of the FI chips. However, due to contamination buildup, which is more substantial toward the edges of the array, it is no longer recommend that the nominal aimpoint be shifted from the center of the detector. The aimpoint must be shifted, however, if an off-center subarray is selected (for example, in order to decrease the frame time).

- ACIS subarray modes: One might wish to reduce the ACIS-S frame time e.g. to minimize the effects of pile-up. The user might consider using a subarray with the HETG, as described at https://space.mit.edu/ASC/calib/hetgsubarray.html. If the source is point-like, then a -3 mm SIM shift can be used. In this case, there are at least 6 mm of ACIS-S rows that may not be scientifically interesting (unless the user desires data from serendipitous sources). The reduced array could have $1024 - 250 = 774$ rows starting at row 1 thus reducing the frame time to 2.5 s. The size of the minimal subarray depends on the low energy cut-off, E, below which the spectra are not of interest. To understand this better, please refer to Figure 8.1. The subarray must be large enough to encompass both HEG and MEG "arms." Larger subarrays are needed at lower energies where the arms are furthest apart. The minimal subarray size is $y_{\rm sub} = 2 \times y_{\rm bg} + 32 + 389/E$, where E is in keV, 32 pixels allow for dither, and $y_{\rm bg}$ is the size of a background region on either side of the spectrum. The background region might be of order 70 pixels (about 10 times the spectrum's extraction width). The SIM should be shifted to center the spectra in the subarray by SIM_Z $= 0.024 \times (y_{\rm sub}/2 - 511 + 122\theta_z)$ mm, where 511 is the row of the ACIS-S aimpoint and θ_z is an optional telescope Z offset. For $E = 1$ keV and $\theta_z = 0.0$ arcmin, then $y_{\rm sub} = 561$ rows and the SIM shift is -5.53 mm.

- Example of a set of parameters: It is instructive to examine the observation catalog (ObsCat) entry for ObsID 9703 which shows observation parameters and values. (See https://cda.harvard.edu/chaser/startViewer.do?menuItem=details&obsid=9703.) The main target is a quasar. The AGN is mildly absorbed, so there is flux of interest down to the MEG limit, hence chip S0 (where the low energy flux will fall) is designated as "yes," while S5, which is redundant with S1 and S0, is designated as "optional." A subarray is used to reduce pile-up. The SIM has been shifted in the Z direction to center the spectrum in the subarray and bring the zeroth order closer to the optical axis. This shift reduced the impact of row-dependent CTI in the FI-chips but contaminant absorption was larger for the MEG -1 order. Finally, a Y offset of 0.0 arcmin is used to keep the Fe-Kα region of the spectrum out of the S2–S3 chip gap in the HEG.

- Use of continuous clocking (CC) mode: this mode can be applied to mitigate pile-up in very bright sources. Most high resolution features will be unaffected. However, there are some additional science issues that should be examined. Based on HETGS observations, X-ray sources with fluxes less than 30 mCrab (1 Crab = 2.4×10^{-8} erg cm^{-2} s^{-1} in the 2–10 keV band) and with absorption columns of less than 1.5×10^{21} cm^{-2} should not show any significant differences in spectra taken in either TE mode using a 512 row subarray or CC mode. Such a subarray is regularly used for moderately bright sources observed in TE mode.

  Brighter and more absorbed source spectra in CC mode, however, may show a number of artifacts stemming from secondary dispersive exposures, which are a result of the collapse of the entire HETG dispersive image in the y-pixel dimension. For further details about CC mode, see ACIS section 6.22.4 and links therein. The largest additional contribution comes from a scattering halo around bright and highly absorbed X-ray sources, which also disperses in the HETG spectrometer. Determining background in CC mode is more complicated because there are no spatial extraction windows that can separate source from background. Therefore, in CC mode, the halo's softer spectrum overlaps the dispersed source spectrum but cannot be spatially separated from it. Considerable science modeling may be required to extract useful continuum spectra. If detailed continuum modeling is anticipated it is recommended to use *MARX* to simulate the level of expected continuum deviations and consider a companion small subarray TE mode companion observation. Some specific issues related to this modeling are:

  - The Si K edge and all other edges can have a distorted shape and incorrect optical depth
  - A broad feature may appear at about 5.4 Å (2.3 keV)
  - There can be significant charge losses below 3 Å ($>$ 4 keV)
  - The spectrum may have the wrong shape above 10 Å ($<$ 1.2 keV)

  For bright sources up to about 300 mCrab, TE faint mode with four CCDs (S1 – S4) and a 350 row subarray is recommended. The data will be simpler to analyze but come with the following trade-offs: loss of the spectral range above 12 Å ($<$1 keV), some pile-up ($<$ 10%) in the 1.5–2.0 keV range, and frames may be dropped due to telemetry saturation. Telemetry saturation can be mitigated by using TE graded mode in some cases. In CC mode, the halo has less impact for harder spectra. For example, a spectrum that decreases in flux from 2 Å (6.4 keV) to 12 Å (1 keV) by more than two orders of magnitude is only mildly affected by the dispersed halo spectrum and the soft CC mode background, practically independent of source flux. Softer spectra, as observed in sources such as Cyg X-1, Cyg X-2, or GX 339-4, on the other hand, are much more gravely affected.

  For the brightest sources, other mitigating actions can be taken. It is now routine to place a 10% window over zeroth order to reduce the fraction of telemetry occupied

by zeroth order events (which are strongly affected by pile-up for bright sources anyway). In addition, one may turn off S0 and S5 or move zeroth order slightly off the array (i.e. using Z-SIM = −11 mm) so that only the HEG −1 and MEG +1 orders are observed. Using such techniques, sources up to 10 Crab can be observed in graded mode. If CC mode is also required for fast timing or to reduce pile-up in the dispersed spectra, significant modeling of the dust scattering halo may be required to interpret the resultant spectra for soft sources. Finally, flare events are not currently flagged in CC mode data, increasing background in FI CCDs. Also, if the particle background is particularly high, the flare event rate can be large enough to fill available telemetry and cause telemetry dropouts.

HRC-I Considerations

The HRC has been used rarely with the HETG so calibration and analysis of data from this configuration are unsupported. The HRC is not generally recommended for use with the HETG because the detector does not provide energy resolution for order sorting, the background can be higher than for ACIS, and its effective area is less above 1 keV. However, there are circumstances that may warrant use of the HETG with the HRC; there is greater effective area below 1 keV, pileup is not an issue, and the resolving power of the HETG with the HRC can be much better than that provided by the LETG with the HRC. Of the HRC options, the HRC-I is most likely to be useful as longer wavelengths of the HETG spectra are observable. See Figure 8.27 for an example of an HETG observation with the HRC-I.

The HRC-I was used in a calibration observation of Capella, as reported in issue 25 of the Chandra Newsletter (Günther 2018). Because the HRC-I is flat, it does not conform to the Rowland circle for optimal spectral resolution over the entire spectrum; one has to choose a small detector focus offset to obtain the best spectral resolution at a specific wavelength, as shown in Figure 8.26 (Günther 2018). The data for that observation (ObsID 19837) show that the measured spectral resolution is consistent with expectations (Figure 8.27, from Günther 2018).

8.5.3 Complications from Multiple Sources

Multiple sources in the field of view can also lead to effects which impact the observation.

Faint Background Sources

The position of a faint second source might be such that the zeroth-order image falls directly on the dispersion pattern from the prime target. In this case, the zeroth-order image of the second source appears as a line in the dispersed image of the prime target. The ACIS energy spectrum can be used to minimize the contribution to the measured dispersed spectrum of the target. Also, the lack of a feature in the other side of the dispersion pattern will indicate that the "line" is spurious.

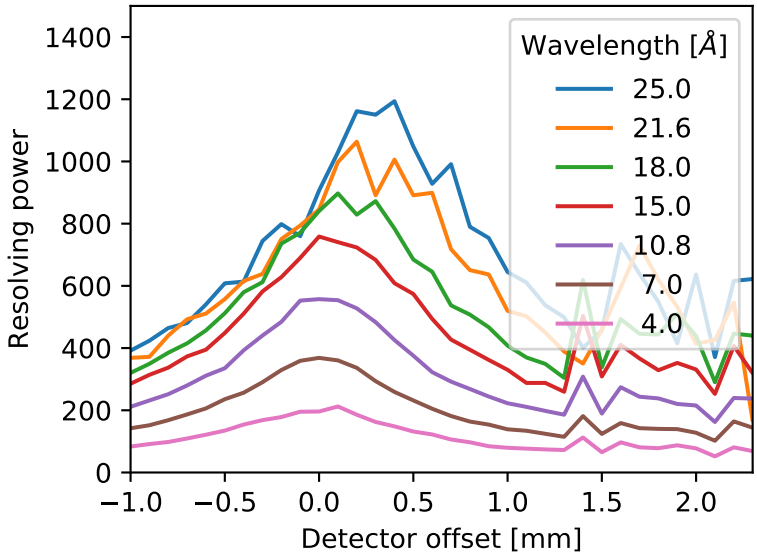

Figure 8.26: Expected spectral resolving power for HETG with the HRC-I detector for different focus offsets, based on ray-trace simulations using *MARX* (see Section 8.6). Offsets between 0.0 and 0.4 mm provide the best spectral resolution for this configuration, and depend on the wavelength of interest. This figure is from Günther (2018).

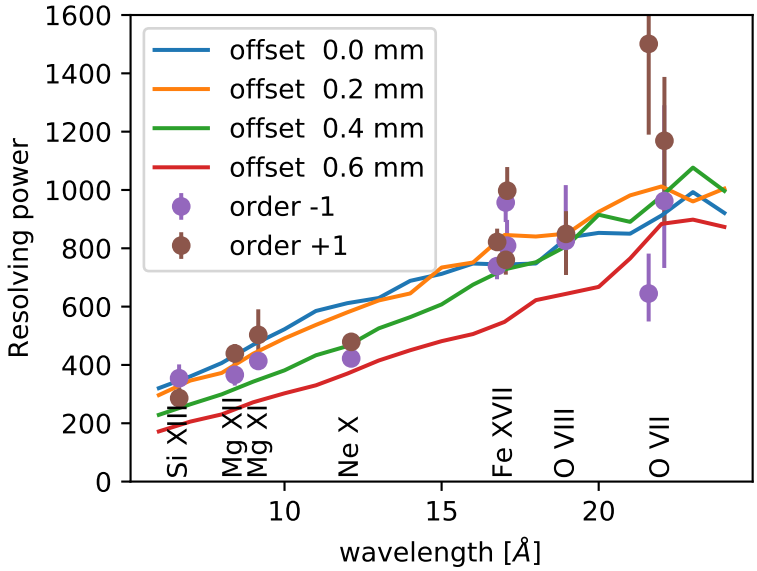

Figure 8.27: Predicted and observed spectral resolving power for HETG with the HRC-I detector for Capella (observation ID 19837), where the focus offset was +0.24 mm. The resolving power measurements are consistent with the predicted values. This figure is from Günther (2018).

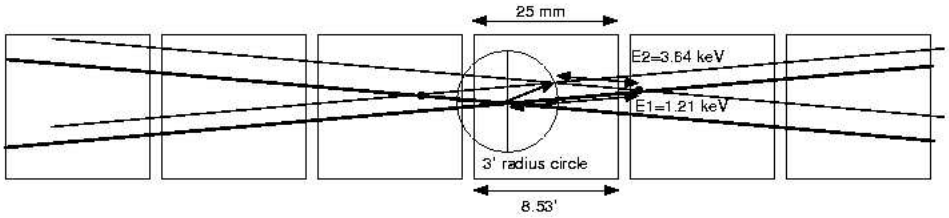

Figure 8.28: An idealized sketch of a 'collision' between two sources separated by 3 arcmin. At the 'collision' point, third-order photons from the on-axis source will have an energy 3x1.21=3.63 keV and ACIS can not distinguish these from the second source's first order photons, at 3.64 keV.

Two Point Sources of Comparable Intensity

The dispersed MEG and HEG spectra of two sources will cross if the objects are fairly close. When the two targets are less than about 3 arcmin apart, both will be nearly in focus, so the spectra appear like two flattened "X"s. Normally, the ACIS-S pulse heights of the events will be significantly different in the regions of overlap, so that one may distinguish the events from two sources in data analysis. There are specific roll angles, however, where the identification of the source is ambiguous; a rare occurrence, but one the user should be aware of.

An example is shown in Figure 8.28, where the MEG spectrum of the brighter object (source 1) overlaps the HEG spectrum of the fainter target (source 2). The first order energies at the overlap positions are a factor of 3 apart, so that $E_2 = 3 \times E_1$. An ambiguity arises from 3rd order photons from source 1 at $3 \times E_1$, which cannot be discriminated by ACIS from photons of about the same energy but from source 2. For a given angular distance between sources, it is possible to specify the observatory roll angle so that collisions like the one shown in the top of Figure 8.29 are avoided.

A Strong Source Lying Outside the Field

The proposer should also take into consideration sources, other than the target, that are within the field of view of the telescope, but out of the field of view of the detector. Parts of the image of the dispersed spectrum may still fall onto the detector. If this presents a problem, a sensible choice of a range of allowable roll angles might ameliorate the situation.

8.5.4 Extended Sources and Spatial-Spectral Effects

The case of a simply extended, spectrally homogeneous source was described in Section 8.2.2, under the heading, "Extended and off-axis targets." Here more complex cases are briefly considered; generally these must be treated on a case-by-case basis.

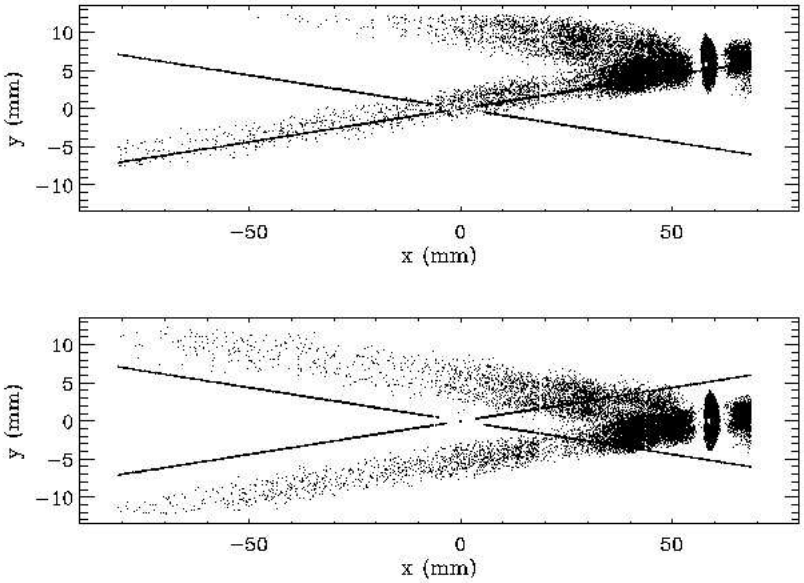

Figure 8.29: A simulation of spectral contamination caused by a second source in the field. The image of the dispersed spectrum from the second source is seen in the upper right-hand corner for particular choice of roll angle. Note that the image is highly extended as the source is 20 arcmin off-axis. For this roll angle, there is significant overlap of the two images. The lower panel shows the same situation, but for a different choice of roll angle. Here the overlap of the images is minimal and data analysis will be further aided through the use of energy discrimination provided by the ACIS-S detector.

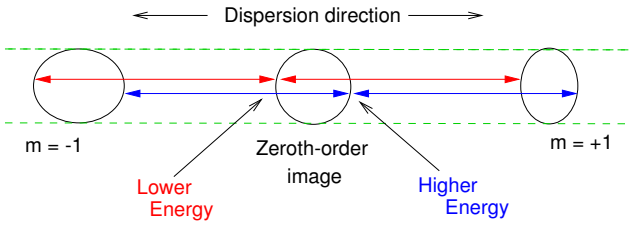

Figure 8.30: HETGS spatial-spectral effect example. In this schematic, a zeroth-order ring image emits at an energy which varies across the ring's diameter in the dispersion direction emitting lower-energy photons on the left and higher-energy photons on the right. The resulting diffracted images in ±1st orders have different appearances due to the spatial-spectral interaction. In the cross-dispersion direction, however, the images have the same extent.

For extended sources with multiple condensations, careful selection of the roll angle (see e.g. Section 8.5.3) might make the data easier to analyze and interpret. It may also be possible to model the spectrum given information from the zero order image and/or a short ACIS exposure with the grating retracted. The ACIS spectrum can then be used as an initial guess in modeling the dispersed HEG and MEG spectra.

The diffracted images of extended objects which lead to position-dependent spectra are complicated. The complexity indicates that information is present but extracting the information is more difficult than for a point or an extended source with a uniform spectrum. For example, the plus and minus order images may not have the same appearance. An example of this effect was seen in ground test data using the double crystal monochromator source; e.g., test image H-HAS-EA-8.003 which is schematically presented and described in Figure 8.30. For astrophysical sources, variations in temperature, abundances, Doppler velocities, cooling flows, etc. can all create spatial-spectral variations. For these complex objects general analysis techniques are not available and forward folding of the spatial-spectral model through *MARX* is the best way to study these effects and to plan potential observations.

8.5.5 Optimizing Detection of Isolated Emission Lines: Choice of Spectrometer

If the scientific objectives require detecting emission lines against a moderately bright source continuum, then the signal/noise ratio depends on the effective area of the instrument in combination with the spectrometer resolving power. The relative merits of each *Chandra* spectrometer are computed in this content. Three cases where this analysis will not apply are when: (1) detecting weak lines that may blend with stronger lines, (2) observing significantly extended sources, and (3) observing lines that are substantially broadened. In case 1, the highest resolving power at the energy of interest would be indicated. Case 2 will require that the reduction of the grating resolution for extended

sources, discussed in Section 8.2.2, be included.

When a line is isolated and appears against a "background" due primarily to the source continuum, then the signal/noise ratio is given by:

$$\frac{C_L}{\sigma_C} = \frac{A_E T W n_E}{[A_E T (dE)_E n_E]^{1/2}} \tag{8.3}$$

where C_L is the number of counts in the emission line, σ_C gives the uncertainty in this number, A_E is the instrumental effective area, T is the integration time, n_E is the photon flux in the continuum in units of photon cm^{-2} s^{-1} keV^{-1}, W is the equivalent width of the line in keV, and $(dE)_E$ is the spectral resolution of the spectrometer in keV. The signal-to-noise ratio per fractional equivalent width, $W_f = W/E$, is then:

$$\frac{C_L/\sigma_C}{W_f} = (T n_E E)^{1/2} [\frac{A_E\ E}{(dE)_E}]^{1/2} \tag{8.4}$$

This last instrument-specific term is a figure of merit for the spectrometers:

$$F_E \equiv [A_E\ (E/dE)_E]^{1/2} \tag{8.5}$$

which can be compared for different instruments at the desired energy. Of course, all these considerations are tempered by the additional features of each instrument setup. For example, this calculation does not take into account instrumental background effects nor the additional continuum that may result from higher energy flux detected in higher orders when the LETG is used with the HRC-S. The reduction of the line detectability then depends on the source spectrum.

8.6 Simulations with *MARX*

For sources with spatial or spectral complexity, observation planning is best carried out using the *MARX* simulator to create a simulated data set. These data can then be analyzed with the same tools as flight data to demonstrate the feasibility of extracting useful results from a proposed observation.

MARX is a suite of programs designed to simulate the on-orbit performance of *Chandra*. It is built around a core program or engine which performs a ray trace of photon paths through all elements of the *Chandra* observatory. The user specifies a file containing the spectral energy distribution of the source to be simulated and then selects a model for the spatial distribution of the source, which can be a FITS image. More complicated "user source models" allow simulation of sources with spatial-spectral variations.

Once the source has been specified, *MARX* traces the path of photons through a model of the HRMA. Models for the High Energy Transmission Grating (HETG) and Low Energy Transmission Grating (LETG) can also be included and, in the focal plane, the user has the choice of all four *Chandra* detectors. The result of the simulation is converted with `marx2fits` into a FITS event file which can then be processed with standard CIAO tools.

The latest capabilities and instructions for use of *MARX* are given in the *MARX* User Guide at the *MARX* web site, https://space.mit.edu/CXC/MARX/ and in analysis threads available at https://cxc.harvard.edu/ciao/threads/index.html.

8.7 REFERENCES

WWW resources:

https://cxc.harvard.edu/cal/- *CXC* Instruments and Calibration page
https://cxc.harvard.edu/caldb/ - *CXC* CALDB page
https://space.mit.edu/HETG/- HETG home page
https://space.mit.edu/CXC/ - *CXC* at MIT, focuses on gratings
https://wwwastro.msfc.nasa.gov/xray/xraycal/ - MSFC ground cal site

Canizares, C.R., Schattenburg, M.L. and Smith, Henry I. 1985, "The High Energy Transmission Grating Spectrometer for AXAF", SPIE, **597**, 253.

Canizares, C.R. *et al.* 2000, "Initial Results from the *Chandra* High Energy Transmission Grating Spectrometer", Atomic Data Needs for X-ray Astronomy, M.A. Bautista, T.R. Kallman, and A.K. Pradhan, eds. https://heasarc.gsfc.nasa.gov/docs/heasarc/atomic/

Canizares, C.R. *et al.* 2000, "High Resolution X-ray Spectra of Capella: Initial Results ferom the *Chandra* High Energy Transmission Grating", ApJ, **539**, L41.

Canizares, C.R. *et al.* 2005,"The *Chandra* High Energy Transmission Grating: Design, Fabrication, Ground Calibration, and Five Years in Flight", PASP, **117**, 1144.

Dewey, D., Drake, J.J., Edgar, R.J., Michaud, K., and Ratzlaff, P., 1998, "AXAF Grating Efficiency Measurements with Calibrated, Non-imaging Detectors", SPIE, **3444**, 48.

Davis, J.E., H.L. Marshall, M.L. Schattenburg, and D. Dewey, 1998, "Analysis and Modeling of Anomalous Scattering in the AXAF HETGS", SPIE, **3444**, 76.

Dewey, D., Humphries, D. N., McLean, G. Y., and Moschella, D. A. 1994, "Laboratory Calibration of X-Ray Transmission Diffraction Gratings", SPIE, **2280**, 257.

Edgar, R.J., 2003, *Chandra* X-ray Center memo dated 6/9/2003, https://cxc.harvard.edu/cal/Acis/Cal_prods/qe/ACIS_QE_O_S23.ps.

Edgar, R.J., & Vikhlinin, A.A., 2004, *Chandra* X-ray Center memo dated 8/11/2004, https://cxc.harvard.edu/cal/Acis/Cal_prods/qe/qe_memo.ps.

Flanagan, K. A., Dewey, D. and Bordzol, L. 1995, "Calibration and Characterization of HETG Grating Elements at the MIT X-ray Grating Evaluation Facility", SPIE, **2518**, 438.

Flanagan, K.A., et al., 2000, "Modeling the *Chandra* High Energy Transmission Gratings below 2 keV" SPIE, **4140**, 559.

Günther, H.M., 2018, *Chandra* Newsletter Issue 25, Spring 2018, "HETG Update", https://cxc.harvard.edu/newsletters/news_25/hetg.pdf.

Huenemoerder, D.P., 2017, *Chandra* X-ray Center memo dated 7/7/2017, https://space.mit.edu/cxc/docs/hetg_bg/hetg_background.pdf.

Madsen, K.K. et al., 2016, "IACHEC Cross-Calibration of *Chandra*, NuSTAR, Swift, *Suzaku*, and *XMM-Newton* with 3C 273 and PKS 2155-304", AJ, **153**, 2.

Markert, T.H., Canizares, C.R., Dewey, D., McGuirk, M., Pak, C., and Schattenburg, M.L. 1994, "The High Energy Transmission Grating for AXAF", SPIE, **2280**, 168.

Markert, T. H. et al., 1995, "Modeling the Diffraction Efficiencies of the AXAF High Energy Transmission Gratings", **SPIE**, **2518**, 424

Marshall, H.L. et al. 1997, "Towards the Calibration of the HETGS Line Response Function", SPIE, **3113**, 160.

Marshall, H.L., Dewey, D., Schulz, N.S., and Flanagan, K.A., 1998, "Spectral Features in the AXAF HETGS Effective Area using High-signal Continuum Tests", SPIE, Vol. **3444**, 64.

Marshall, H.L., D. Dewey, and K. Ishibashi, 2004, "In-Flight Calibration of the *Chandra* High Energy Transmission Grating Spectrometer", Proc. SPIE, Vol. **5165**, 457.

Marshall, H.L., Tennant, A., Grant, C., Hitchcock, A.P., O'Dell, S., and Plucinsky, P.P., 2004, "Composition of the ACIS Contaminant", SPIE, **5165**, 497.

Marshall, H.L., 2005, *Chandra* X-ray Center memo dated 10/14/2005, https://space.mit.edu/ASC/calib/heg_meg/meg_heg_report.pdf.

Marshall, H.L., 2012, "Updating the *Chandra* HETGS Efficiencies using In-Orbit Observations", SPIE, **8443**, 48.

Marshall, H.L., 2017, *Chandra* Newsletter Issue 24, Spring 2017, "HETGS Update", https://cxc.harvard.edu/newsletters/news_24/marshall.pdf.

Marshall, H.L., 2020, *Chandra* X-ray Center memo on effects of warm ACIS on the HETGS Effective Area dated 6/16/2020, https://space.mit.edu/CXC/calib/rmf_effects.pdf.

Schattenburg, M.L. et al. 1991, "Transmission Grating Spectroscopy and the Advanced X-ray Astrophysics Facility", Optical Engineering, **30**, 1590.

Schattenburg, M.L., Ancoin, R.J., Flemming, R.C., Plotnik, I., Porter, J., and Smith, H.I. 1994, "Fabrication of High Energy Transmission Gratings for AXAF", SPIE, **2280**, 181.

Schulz, N.S., Dewey, D., Marshall, H.L. 1998 "Absolute Effective Areas of HETG", SPIE, **3444**, 160.

Schulz, N.S., 2022, "Instrumental Features Near the Si K Edge in the Chandra HETG 1st Order", *Chandra* X-ray Center memo, https://space.mit.edu/CXC/calib/sikedge_final_doc.pdf.

Schulz, N.S., 2023, "HETG Spectra at Focal Plane Temperature −105 C", *Chandra* X-ray Center report, https://space.mit.edu/CXC/calib/fptemp105_062723.pdf.

Smith, M.J.S., & Marshall, H.L., 2012, "*XMM-Newton* – *Chandra* Blazar Flux Comparison", presented at the 7th IACHEC meeting https://iachec.org/wp-content/presentations/2012/Smith.pdf.

Weisskopf, M.C. et al. 2004, "An Overview of the Performance of the *Chandra* X-Ray Observatory", Experimental Astronomy, Vol. **16**, no. 1, 1.

"X-ray Spectroscopy in Astrophysics", 1999, Springer-Verlag, ed. van Paradijs, J., & Bleeker, J.A.M.

Chapter 9

LETG: Low Energy Transmission Grating

9.1 Instrument Description

The Low Energy Transmission Grating (LETG) was developed under the direction of Dr. A.C. Brinkman in the Laboratory for Space Research (SRON) in Utrecht, the Netherlands, in collaboration with the Max-Planck-Institut für Extraterrestrische Physik (MPE) in Garching, Germany. The grating was manufactured in collaboration with Heidenhain GmbH.

The Low Energy Transmission Grating Spectrometer (LETGS) comprises the LETG, a focal-plane imaging detector, and the High Resolution Mirror Assembly discussed in Chapter 4. The *Chandra* High Resolution Camera spectroscopic array (HRC-S) is the primary detector designed for use with the LETG. The spectroscopic array of the *Chandra* Advanced CCD Imaging Spectrometer (ACIS-S) can also be used, though with lower quantum efficiency longward of $\sim$ 10 Å and a smaller detectable wavelength range than with the HRC-S. The High Energy Transmission Grating (HETG) used in combination with ACIS-S offers superior energy resolution and quantum efficiency above 0.78 keV. The HRC is discussed in Chapter 7, the ACIS in Chapter 6, and the HETG in Chapter 8.

The LETGS provides high-resolution spectroscopy ($\lambda/\Delta\lambda > 1000$) between 80 and 175 Å (0.07–0.15 keV) and moderate resolving power at shorter wavelengths. The nominal LETGS wavelength range accessible with the HRC-S is 1.2–175 Å (0.07–10 keV). ACIS-S has some effective area out to roughly 60 Å ($\sim$0.20 keV); however LETG-HRC-S has more area at all wavelengths beyond $\sim$10 Å, and no non-calibration LETG/ACIS-S observations have been made since 2015.

A summary of LETGS characteristics is given in Table 9.1.

Table 9.1: LETGS Parameters

Wavelength range	1.2–175 Å (HRC-S) 1.2–60 Å (ACIS-S)
Energy range	70–10000 eV (HRC-S) 200–10000 eV (ACIS-S)
Resolution ($\Delta\lambda$, FWHM)	0.05 Å
Resolving Power ($\lambda/\Delta\lambda$)	≥ 1000 (50–160 Å) $\approx 20 \times \lambda$ (3-50 Å)
Dispersion	1.148 Å/mm
Plate scale	48.82 μm/arcsec
Effective area (1st order)	1–25 cm^2 (with HRC-S) 1–100 cm^2 (with ACIS-S)
Background (quiescent)	LETG/HRC-S: 26 (90) cts/0.07-Å/100-ks @ 50 (175) Å for typical L2 background rate (0.10 cts/pixel/100-ks; see Figure 9.21) LETG/ACIS-S: $\ll$ 0.01 cts/pixel/100-ks (order sorted)
Detector angular size	3.37' × 101' (HRC-S; standard 6-tap-wide spectroscopy region) 8.3' × 50.6' (ACIS-S; full 6-chip array)
Pixel size	6.43 × 6.43 μm (HRC-S) 24.0 × 24.0 μm (ACIS-S)
Temporal resolution	16 μsec (HRC-S in Imaging Mode, center segment only) $\sim$10 msec (HRC-S in default mode) 2.85 msec–3.24 sec (ACIS-S, depending on mode)
Rowland diameter	8637 mm (effective value)
Grating material	gold
Facet frame material	stainless steel
Module material	aluminum
LETG grating parameters	
Period	0.991216 $\pm$0.000087 μm
Thickness	0.474 $\pm$0.0305 μm
Width (at bar middle)	0.516 $\pm$0.0188 μm
Bar Shape	symmetric trapezoid
Bar Side Slope	83.8 $\pm$2.27 deg
Fine-support structure	
Period	25.4 μm
Thickness	2.5 μm
Obscuration	$< 10\%$
Dispersion	29.4 Å/mm
Material	gold
Coarse-support structure	
Triangular height	2000 μm
Width	68 μm
Thickness	< 30 μm
Obscuration	$< 10\%$
Dispersion	2320 Å/mm
Material	gold

9.1.1 Scientific Objectives

The LETGS provides the highest spectral resolving power (> 1000) on *Chandra* at low (0.07–0.2 keV) energies. High-resolution X-ray spectra of optically thin plasmas with temperatures between 10^5 and 10^7 K, such as stellar coronae, reveal a wealth of emission lines that provide diagnostics of temperature, density, velocity, ionization state, and elemental abundances and allow precise studies of structure, energy balance, and heating rates. Absorption features provides similar information in cases where bright compact X-ray sources are embedded in cooler, extended gas clouds.

The high resolution ($\Delta\lambda \approx 0.05$ Å) of LETGS spectra at longer wavelengths ($\gtrsim 100$ Å) also permits detailed studies of spectral line *profiles* in the X-ray region. These studies may provide non-thermal velocities of stellar coronae, flow velocities along active-region loops, orbital velocities in X-ray binaries, and upflow velocities in stellar flares. The LETGS also allows time resolved spectroscopy, 1-D spatially resolved spectra, and spectra of multiple point sources within its $\sim$4 arcmin field of best focus.

Since the ultimate spectral resolution can only be achieved for point sources, the prime candidates for study in the Galaxy mainly comprise stellar coronae, white dwarf atmospheres, X-ray binaries, and cataclysmic variables. Extragalactic sources include relatively bright active galactic nuclei (AGN) and cooling flows in clusters of galaxies.

9.1.2 Heritage

Flat transmission gratings were flown aboard *Einstein* and *EXOSAT*. The LETG grating elements are produced using a technique similar to that used for production of the *EXOSAT* gratings. However, the LETG shares only basic operating principles with earlier instruments. Advanced grating technology has enabled the achievement of greater efficiency and increased dispersion. The Rowland geometry (see Figure 8.4) of the grating plate and spectroscopic arrays reduces dispersed image aberrations and hence contributes to improved spectral resolution.

9.1.3 Operating Principles

When inserted behind the HRMA, the LETG diffracts X-rays into a dispersed spectrum according to the grating diffraction relation, $m\lambda = p \sin\theta$, where m is the integer order number, λ the photon wavelength, p the spatial period of the grating lines, and θ the dispersion angle. Parameters are summarized in Table 9.1. The grating facets are mounted on an aluminum support plate which has been machined so that the centers of individual grating facets lie on a Rowland torus. The grating facets are aligned to produce a single dispersed image. Spectral resolution is determined, among other factors, by grating line density, line density variations, HRMA point-spread function, pointing stability, alignment accuracy, pixel size of the readout detector, and detector geometry.

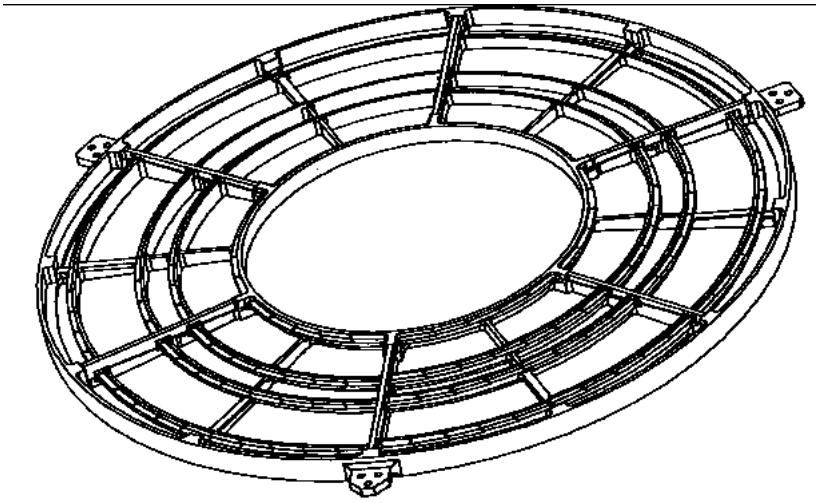

Figure 9.1: LETG Grating Element Support Structure, a machined aluminum plate approximately 110 cm in diameter that holds grating modules on a Rowland torus behind the *Chandra* mirrors.

9.1.4 Physical Configuration

When the LETG is used, the Grating Element Support Structure (GESS), an aluminum frame approximately 110 cm in diameter and 6 cm thick, is inserted ∼300 mm behind the exit aperture of the HRMA and 1.4 m behind the HRMA mid-plane. The GESS holds approximately 180 trapezoidal grating modules, which measure about 13 × 50 mm. A design drawing of the full GESS is shown in Figure 9.1; a closer view, showing some mounted modules, is seen in Figure 9.2. Figure 9.3 shows empty grating modules mounted on the GESS. Each grating module holds three circular grating facets, each of which comprises approximately 80 of the triangular grating elements seen in Figure 9.4.

In contrast to the HETG gratings, which have a thin polyimide substrate, the LETG gratings are free-standing wires held by a support mesh. Within each grating facet the grating bars are supported by perpendicular "fine support" bars and triangular "coarse support" bars. The parameters of these structures are given in Table 9.1. A schematic of the grating structure is shown in Figure 9.4. Both the fine and coarse grating supports act as long-period transmission gratings themselves. As discussed in Section 9.3.7 and illustrated in Figures 9.23, 9.24, and 9.27, the fine support produces a dispersion pattern perpendicular to the grating dispersion direction and the coarse support produces a six-pointed star pattern.

Since the gratings are produced from a single master mask, there is negligible variation in the period between facets. The thickness of the gold of the grating bars on top of the support mesh determines the "phasing," or efficiency of redistribution of photons into each spectral order in wavelengths where the gold is partially transparent. The thickness

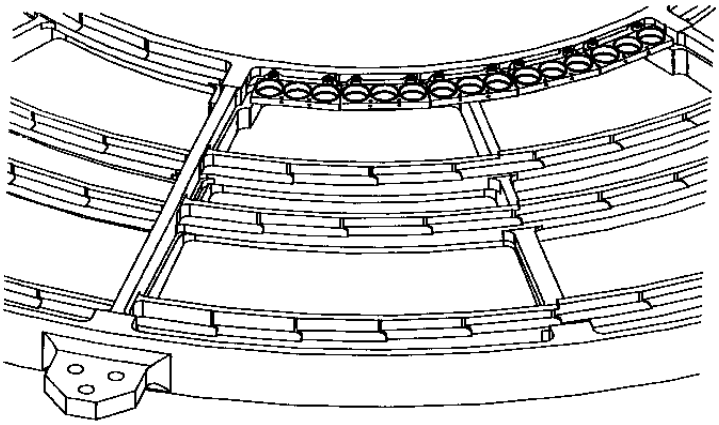

Figure 9.2: Detail of the LETG Grating Element Support Structure showing grating modules mounted on the inner annulus.

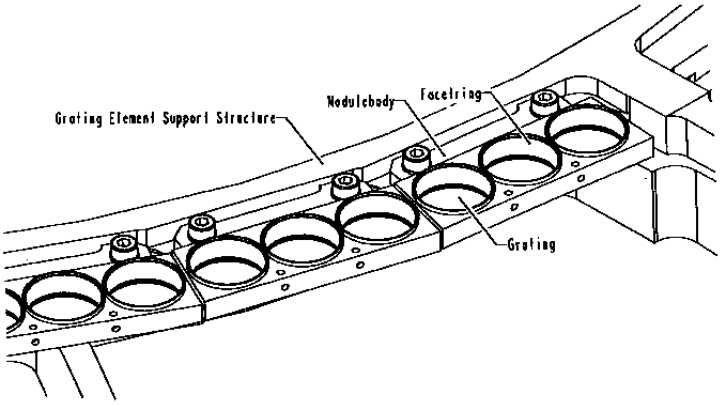

Figure 9.3: A closeup view of the LETG GESS showing nearly three complete grating modules. Each module holds three circular grating facets, and each facet contains approximately 80 triangular grating elements.

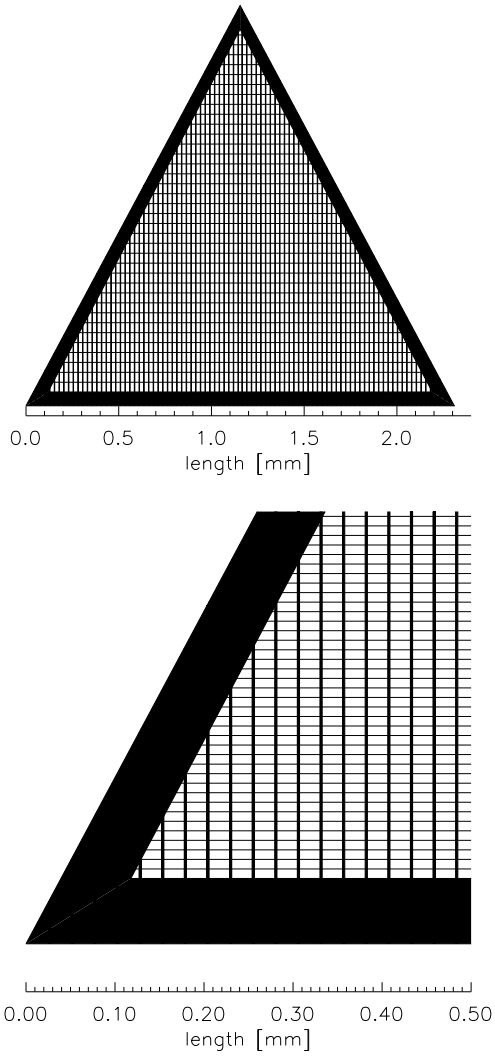

Figure 9.4: LETG facet structure schematic showing the basic shape of the individual grating elements and the relative sizes of the support structures. The upper view shows the complete grating element, which consists of the triangular coarse support, the vertical fine supporting bars, and the (horizontal) grating bars. The grating bars themselves are not shown to scale. In the upper view every 50^{th} grating bar is drawn, in the lower view every 10^{th} bar.

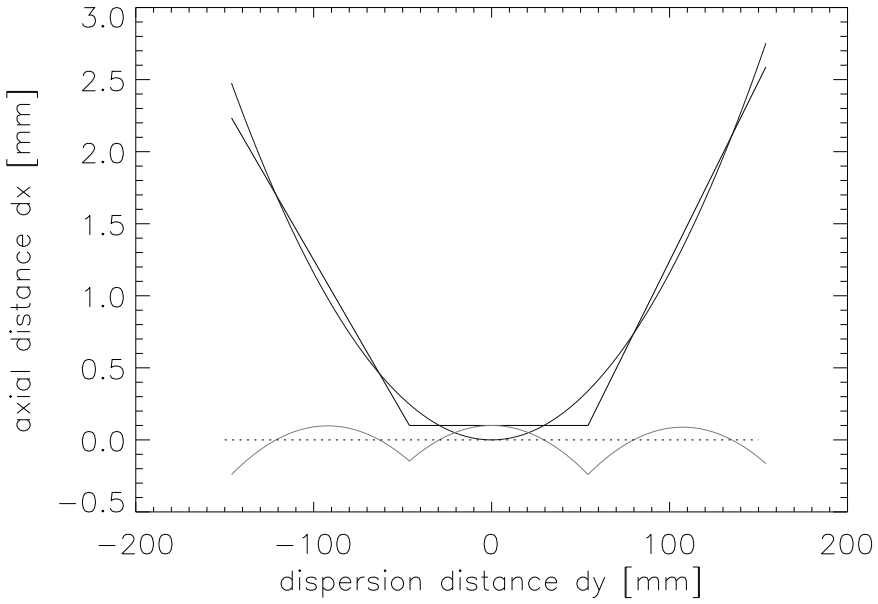

Figure 9.5: The front surfaces of the HRC-S detector segments (straight lines) and their relationship to the Rowland circle (smooth curve) are shown schematically. The scalloped line beneath them is the difference between the detector surface and the Rowland circle.

is designed to optimize the 1st order response at energies of interest.

To reduce aberrations, the GESS is shaped to follow the Rowland torus. The basics of the Rowland geometry are shown in Figure 8.4. The primary readout detector (HRC-S) is made of three tilted array segments which also follow the Rowland circle in the image plane (see Figures 7.2 and 9.5). Because the detector array elements are flat, the distance from the Rowland circle changes with position, and so the spectral resolution changes very slightly with wavelength. The secondary readout detector, ACIS-S, has 6 CCDs, each of which is only one-quarter the length of an HRC-S segment, so the ACIS-S array follows the Rowland circle even more closely.

9.2 Calibration

9.2.1 Pre-Launch Calibration

Prior to assembly, individual grating elements were tested using a visual light spectrograph at the MPE. Laboratory calibration of grating period and resolution was performed for individual grating elements at optical wavelengths and extrapolated to the X-ray range. Grating efficiencies at X-ray wavelengths were modeled using near-infrared

Table 9.2: Routine LETGS Calibration Monitoring Observations

Target	Freq. (yr^{-1})	Purpose
Capella	1	LETG/HRC-S LRF, dispersion relation, QE, EA
PKS 2155-304	1	LETG/ACIS-S EA, ACIS-S contam, cross calib with *XMM*, *Suzaku*
Mkn 421	1	LETG/ACIS-S and LETG/HRC-S EA, ACIS-S contam, HRC-S gain
RX J1856-3754	1	LETG/ACIS-S EA and ACIS-S contam
HZ 43	1	LETG/HRC-S EA and HRC-S gain

spectrophotometry and verified by X-ray measurements of a sample of facets. Grating facet and module alignment was also tested. LETGS efficiency, resolution, and line response function were tested at the X-ray Calibration Facility at MSFC in Huntsville, Alabama, for both ACIS-S and HRC-S configurations. Absolute energy scale and off-axis response were also measured. Efficiency and the Line Spread Function (LSF) of the LETG and HRMA/LETG subsystem were characterized using a detector system designed for HRMA calibration, the HRMA X-ray Detection System (HXDS). Details may be found in the XRCF Final Report at https://cxc.harvard.edu/cal/Hrma/XRCFReport.html.

9.2.2 In-Flight Calibration

In-flight calibration of the LETGS, along with its primary detector, the HRC-S, is planned and executed by the *CXC* LETG team. LETG first-light and focus observations were of the active late-type binary Capella whose coronal spectrum is rich in narrow spectral lines (see Figure 9.26). Spectra of Capella, the late-type star Procyon (F5 IV), and serendipitous Guest Observer targets have been used in calibration of the LETG dispersion relation, resolving power, line response function, and grating-detector alignment.

Calibration of the LETGS effective area (EA) and HRC-S quantum efficiency (QE) at energies above the C-K edge (0.28 keV, 44 Å) relies primarily on observations of the quasar continuum source PKS 2155-304 and, earlier in the mission, of 3C 273 and Capella. PKS 2155-304 also has been regularly observed with LETG/ACIS-S to monitor contamination buildup on ACIS (see Section 6.5.1) and for cross calibration with other X-ray missions. Mkn 421 largely replaced PKS 2155-304 as an ACIS contamination monitoring source starting in 2010. RX J1856 (neutron star continuum source) is also used for contamination monitoring. For calibration of the LETG/HRC-S EA at longer wavelengths, observations of the hot DA white dwarfs HZ 43 and Sirius B are used.

Early in the mission, most calibration targets were observed at least twice per year to monitor LETGS operation, but calibration observations are now less frequent. Regularly observed sources and monitoring frequencies are listed in Table 9.2. Other targets are occasionally observed to meet specific calibration needs that might arise, and ACIS-S and HRC-S (see Table 7.5) are routinely calibrated by themselves (i.e., without gratings).

9.3 LETGS Performance

9.3.1 Usage

Overview

The primary use of the LETG is for on-axis observations of point sources, which produce a zero-order image and a dispersed spectrum. Typical LETGS observations range from a few tens to several hundred ks. To reduce the (small) risk that the grating mechanism might fail, its frequency of use is minimized by grouping grating observations into consecutive time blocks whenever possible.

The net LETG transmission is ∼28% at energies below ∼1 keV (about 12.5% for zeroth order, the same for 1st order, and a few percent for all other orders) so count rates are usually not a concern with respect to exceeding detector limits or telemetry saturation. However, some bright sources (e.g. Sco X-1), if observed for long exposure times, could cause significant charge depletion in the HRC MCPs (see Chapter 7, especially Section 7.14), and even moderate rates may cause pile-up problems when using ACIS-S (see Sections 6.16 and 9.4.1). Some observers may find it useful to insert the LETG for imaging observations simply to reduce the detected photon count rate.

Detectors

In standard operation, the LETGS uses either the HRC-S or ACIS-S as its detector. The LETG+HRC-S covers a wavelength range of approximately −165 Å to +175 Å in 1st order for on-axis sources. This wavelength range can be shifted somewhat by offset pointing, but image quality degrades substantially beyond about 2 arcmin. The HRC-S does not have sufficient energy resolution to allow sorting of overlapping spectral orders.

In rare cases it might be useful to use the HRC-S Low Energy Suppression Filter (LESF), as discussed in Section 9.4, in order to obtain a predominantly higher-order ($m > 1$) spectrum. The LESF is a region on the HRC-S UV/Ion Shield (UVIS) where the aluminum coating is relatively thick, corresponding to the upper part of the "T" in Figure 7.1. Note that the Al coating on the LESF is thicker on the outer plates than on the central plate. See Figures 9.17 and 9.19 for the effect of the LESF on 1st and higher order effective areas.

When used with the ACIS-S detector, the effective LETGS wavelength coverage is reduced because of the smaller detector size in the dispersion direction (ACIS-S is only half as long as the HRC-S) and the fact that the three outermost chips (S0, S4, and S5) now have essentially zero QE for detecting 1st order LETG photons because of absorption by the contaminant on the ACIS filters (see Figure 9.6). Another consideration is that ACIS has worse temporal resolution than HRC, which may be important when observing periodic or rapidly varying sources. On the other hand, ACIS has a much lower background and its intrinsic energy resolution allows easy separation of diffraction orders. Note that the CTI-degraded energy resolution of the ACIS FI CCDs (Section 6.7) does not pose a problem

for LETG point-source observations, because order separation effectively eliminates energy uncertainties.

In some special cases, the HRC-I may also be used with the LETG. A detailed discussion of the various merits of LETGS detector choices from a point of view of proposal planning is given in Section 9.4.

Off-Axis and Multiple Sources

Because the LETGS is essentially an objective-grating system, it is possible to do multi-object spectroscopy, although, as noted above, the point-spread function degrades rapidly off-axis. To include or reject secondary sources, or to avoid overlapping diffraction from multiple sources, observers may specify the orientation (roll angle) of the grating dispersion direction on the sky (see Chapter 3). Observations of extended sources are also possible, but at the expense of resolving power and with the loss of the simple relation between position and energy. In angular extent, the standard HRC-S spectroscopy readout region is 3.37 arcmin × 101 arcmin, and the full ACIS-S array covers 8.3 arcmin × 50.6 arcmin. In special cases, a different HRC-S detector "window" (up to twice as wide in cross-dispersion) may be selected, as described in Section 9.4.

9.3.2 Wavelength Coverage and Dispersion Relation

The active extent of the HRC-S in the dispersion direction is 296 mm, almost exactly twice that for the ACIS-S. The nominal zeroth-order aimpoints for each detector are slightly offset from the detector center so that gaps between the three HRC-S segments (six ACIS-S segments) will occur at different wavelengths in negative and positive orders. A Y-axis offset (along the dispersion axis) of +1.1 arcmin is recommended for most LETG+ACIS-S observations in order to shift coverage of longer wavelengths onto the back-illuminated S1 and S3 chips, which have higher QE at low energies than the other ACIS-S chips.

With a dispersion of 1.148 Å/mm for the LETG, the standard wavelength range of the LETGS with HRC-S is −163 Å to +175 Å. Physical coverage with ACIS-S extends from −88 Å to +84 Å when using Y-offset=+1.1 arcmin, but the poor low-energy response of the outlying front-illuminated chips limits the effective wavelength range to −59 Å for negative 1st order and less than about +15 Å for the positive order (see Figure 9.6). Outlying chips may be useful, however, for collecting higher-order spectra.

Off-Axis Pointing and Detector Gaps

Wavelength coverage can be adjusted (increasing the wavelength range on one side and decreasing on the other) by changing the central offset (the observatory Y coordinate—see the discussion in Chapter 3) although spectral resolution degrades rather quickly beyond about 4 arcmin from the optical axis. From the information in Table 9.1, one can derive the relationship between angular offset (in the dispersion direction) and wavelength as 3.36 Å per arcmin, so an offset of 10 arcmin would stretch the positive order HRC-S coverage

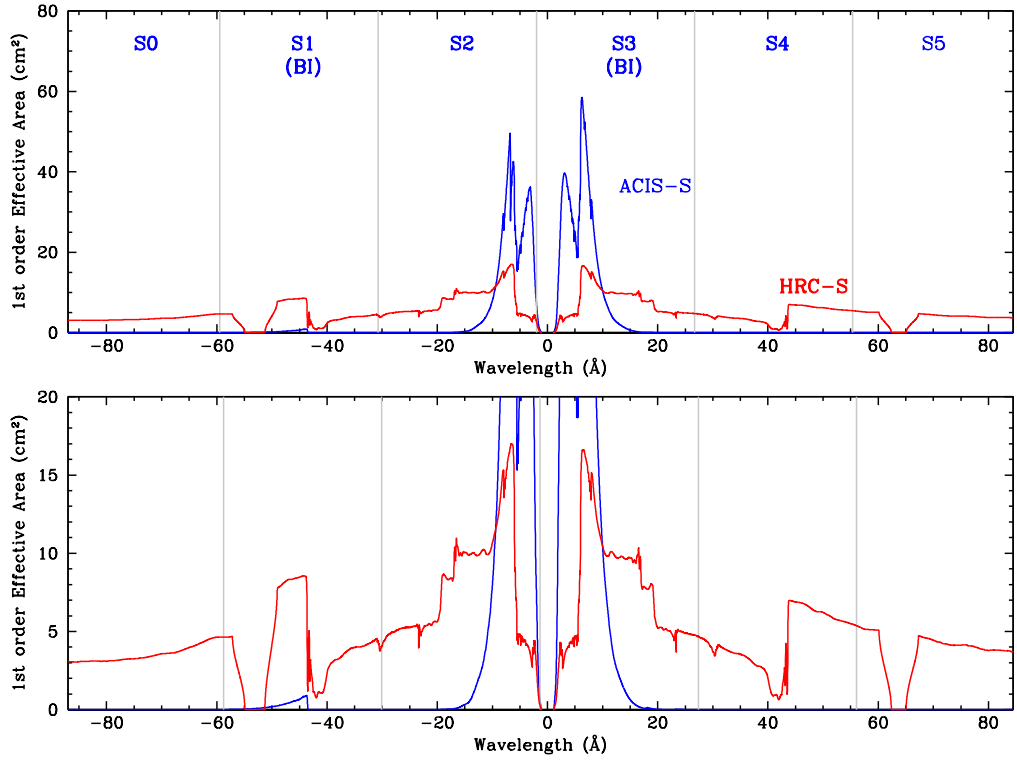

Figure 9.6: LETGS 1st order effective area (EA) with ACIS-S and HRC-S, showing plus and minus orders separately; lower panel shows low-EA regions in more detail. The effects of dither and ACIS bad columns are included. Gray vertical lines mark ACIS chip boundaries, and HRC plate gaps appear near −53 Å and +63 Å. ACIS curve is for Y-offset=+1.07 arcmin (with current dither aimpoint; see Section 6.11) and HRC curve is for Y-offset=0 arcmin. See Section 9.4.2/SIM-Z Offsets and Offset Pointing for more information regarding the choice of ACIS row and Y-offset. The LETG+HRC EA is currently decreasing a few percent per year at all wavelengths, while the LETG+ACIS EA is slowly falling because of increasing contamination on the ACIS Optical Blocking Filters. See also Figure 9.17, which plots effective areas for combined plus and minus orders. Both EA curves are taken from observations in summer 2025.

Detector	Section	Energy (eV)	Wavelength (Å)
HRC-S	UVIS Inner T (thick Al)	(650) – 730	(19) – 17
HRC-S	seg-1 (neg. $m\lambda$)	(75) – (221)	(164) – (56)
HRC-S	seg0	(245) – 203	(50) – 61
HRC-S	seg+1 (pos. $m\lambda$)	188 – 71	66 – 175
ACIS-S	S0 (neg. $m\lambda$)	(141) – (207)	(87.9) – (59.7)
ACIS-S	S1 (Back-illuminated)	(210) – (400)	(59.2) – (31.0)
ACIS-S	S2	(407) – (5300)	(30.4) – (2.3)
ACIS-S	S3 (Back-illuminated)	(7300) – 469	(1.7) – 26.4
ACIS-S	S4	460 – 225	27.0 – 55.1
ACIS-S	S5 (pos. $m\lambda$)	223 – 148	55.7 – 83.8

Table 9.3: LETG Position-Dependent Spectral Coverage (1st Order). Energies and wavelengths for the negative order are given in parentheses. Listed values are for the most commonly used pointing (on-axis for HRC-S and Y-offset = +1.2 arcmin for ACIS-S) without dither. Standard dithering affects ±1.1 Å on the edge of each HRC-S segment and ±0.90 Å on the edge of each ACIS chip. Uncertainties are of order 0.5 Å and arise mostly from thermally induced pointing errors, which have been increasing over time; shifts of ∼1 Å have been seen. Note that the QE for FI ACIS chips is much lower than for BI chips at low energies, rendering S0, S4, and S5 of limited use (see Figure 9.6). Please see https://cxc.harvard.edu/cal/Letg/ACIS_params for up-to-date and detailed information regarding offset pointing, wavelength ranges, and pointing errors.

to approximately 210 Å (60 eV). While the vast majority of LETG observations have been made with offsets of less than 2 arcmin, flight LETG calibration data have been collected at 5 arcmin off-axis (for resolution testing) and 10 arcmin off-axis (for effective area calibration).

As noted in Section 9.3.2, there are gaps between detector segments which create corresponding gaps in the wavelength coverage of each order. The gaps in + and − orders do not overlap so that the combined wavelength coverage is continuous. The locations of the gaps (neglecting the effects of dither) are listed in Table 9.3, which also lists the locations of the HRC-S UV/Ion Shield inner "T" filter edges. Dithering the spacecraft will partially smooth these gaps, but observers may wish to adjust the source pointing if a line of interest falls in a gap, or to tune the wavelength coverage of the higher-QE back-illuminated (S1 and S3) ACIS-S chips. Standard HRC dither amplitude (full width, in both directions) is 40 arcsec (1.95 mm), which covers 2.3 Å, and standard full-width ACIS dither (changed in 2023) is 32 arcsec (1.56 mm), or 1.8 Å.

Please see the "Checking your LETG/ACIS Obscat Setup" web page at https://cxc.harvard.edu/cal/Letg/ACIS_params. An interactive tool for visualizing spectral feature placement on the ACIS array as a function of Y-offset and source redshift is available there.

Dispersion Calibration

Overall wavelength calibration is accurate to a few parts in 10,000 across the full wavelength range of the instrument, and the Rowland diameter has remained stable since launch. The RMS deviation between observed and predicted wavelengths for a set of relatively unblended lines observed in Capella spectra amounts to 0.013 Å, after correction for the relative spacecraft and Capella radial velocity differences. Through analysis of accumulated calibration and general observer (GO) observations, some remaining differences between predicted and observed line wavelengths have been found to be caused by small event position errors at some locations on the HRC-S detector.

Position errors occur in both dispersion and cross-dispersion axes, though for spectroscopy the latter are usually not important because data are generally summed in the cross-dispersion direction. The magnitude of position errors in the dispersion axis for the central HRC-S plate range from 0 to ∼0.05 Å, with typical errors of ∼0.01–0.02 Å. The outer plates ($>$ 60 Å) tend to exhibit larger errors of typically ∼0.02–0.08 Å. The size of the position errors changes over spatial scales of the HRC readout taps, which are 1.646 mm apart (see Chapter 7 for details of the HRC). Spacecraft dither moves dispersed monochromatic photons of a given order over a region of the detector that is roughly 2mm square. Within any given dither region, then, monochromatic light will fall on detector regions that have different position determination errors. As a result, a narrow spectral line could suffer some distortion of its line profile and/or a small shift of its apparent wavelength. Such distortions or shifts could occur at the spacecraft dither frequency (see Chapter 5); observers should therefore exercise caution in interpretation of such periodic effects. Care should also be taken when interpreting the results of combined spectra from + and − orders, since these effects are not symmetrical about zeroth order.

The position errors appear to be stable *in detector coordinates* and to repeat in different observations for which the aimpoints are very similar. An empirical correction for the effect based on accumulated calibration and GO observations has been derived and is implemented in CALDB 3.2.0 as part of the HRC-S degap map. Details on the derivation of the wavelength corrections can be found on the "Corrections to the Dispersion Relation" page at https://cxc.harvard.edu/cal/Letg/Corrlam. After correction, the RMS deviation between observed and predicted line wavelengths observed in Capella are reduced from 0.013 to 0.010 Å across the entire wavelength range of the instrument. Corrections for the outer plates ($>$ 60 Å) are much less effective owing to a lack of reference lines with adequate signal-to-noise ratio. For the central plate alone, the RMS deviation amounts to 0.006 Å.

For the purposes of observation planning, it should be assumed that observed line wavelengths for individual lines could be in error by up to about 0.02 Å for $\lambda < 60$ Å, and 0.05 Å for $\lambda > 60$ Å. Despite the observed repetition in pattern from observation to observation, observers are reminded that the exact wavelength error for any given line depends on the exact position of the target on the detector. Small differences in actual aimpoint that occur naturally between observations as a result of uncertainties in aspect and target acquisition (see Chapter 5) mean that wavelength shifts for a specific line are

not generally repeatable from one observation to another and might also be subject to small secular trends.

9.3.3 Resolving Power

The dominant contribution to the LETGS line response function (LRF) and instrument resolving power is the HRMA point-spread function (PSF), which is ∼25 μm FWHM, depending on energy. The HRC-S has an intrinsic resolution of ∼20 μm FWHM, which is digitized into 6.43-μm-wide pixels, while ACIS has true pixels that are 24 μm wide. Uncertainties in correcting photon event positions for the observatory aspect, which occurs during ground data processing, adds a small contribution of order a few μm. Finally, the small errors in event position determination resulting from HRC-S imaging non-linearities described in Section 9.3.2 can lead to some distortion of spectral line profiles. These effects are difficult to quantify in detail but are estimated to affect the line FWHM by less than 25%.

When all these effects are combined, the LETGS line response function is generally ∼40 μm FWHM when using either HRC-S or ACIS-S as the detector. With a conversion of 1.148 Å/mm for the LETG, a good figure of merit for LETGS resolution is therefore 0.05 Å. Because the three segments of the HRC-S can not perfectly follow the Rowland circle (see Figure 9.5), however, resolution varies slightly along the detector, and is lowest near the ends of each detector segment. Resolution degradation is almost negligible when using the ACIS-S, since its six segments more closely follow the Rowland circle, although the coarser ACIS pixel size (24 μm vs. ∼40 μm FWHM LRF) means that line profiles are barely adequately sampled. A plot of LETGS resolving power for an on-axis point source, based on results from an observation of Capella, is shown in Figure 9.7.

Plots of fits to the LETG+HRC-S LRF at zeroth order and of Fe XVII and XVIII lines at ∼17 and 94 Å are given in Figures 9.8 and 9.9. The fitted form is a Moffat function:

$$I(\lambda) = \frac{I_0}{\left[1 + \left(\frac{\lambda - \lambda_0}{\Gamma}\right)^2\right]^\beta} \tag{9.1}$$

where λ_0 is the wavelength of the line center, Γ is a measure of the line width, and I_0 and β are fit parameters. The relation between Γ and the line FWHM depends on β. For a Lorentzian profile, or $\beta = 1$, the profile FWHM is $\Gamma/2$. For the value $\beta = 2.5$ recommended here, FWHM $\approx 1.13\Gamma$.

Figure 9.10 illustrates the χ^2 of fits to the zeroth order profile versus β, and shows a best-fit profile with an index of ∼2.5. (Note that $\beta = 1.0$ yields a Lorentzian profile.) The best fit to the very high signal-to-noise-ratio zeroth order profile is far from being statistically satisfactory. However, spectral lines seen in first order generally contain orders of magnitude fewer counts than in the zeroth order of the well-exposed calibration spectrum shown, and the Moffat function nearly always provides a good match. Line response functions can also be generated within CIAO in the RMF FITS format. These are based

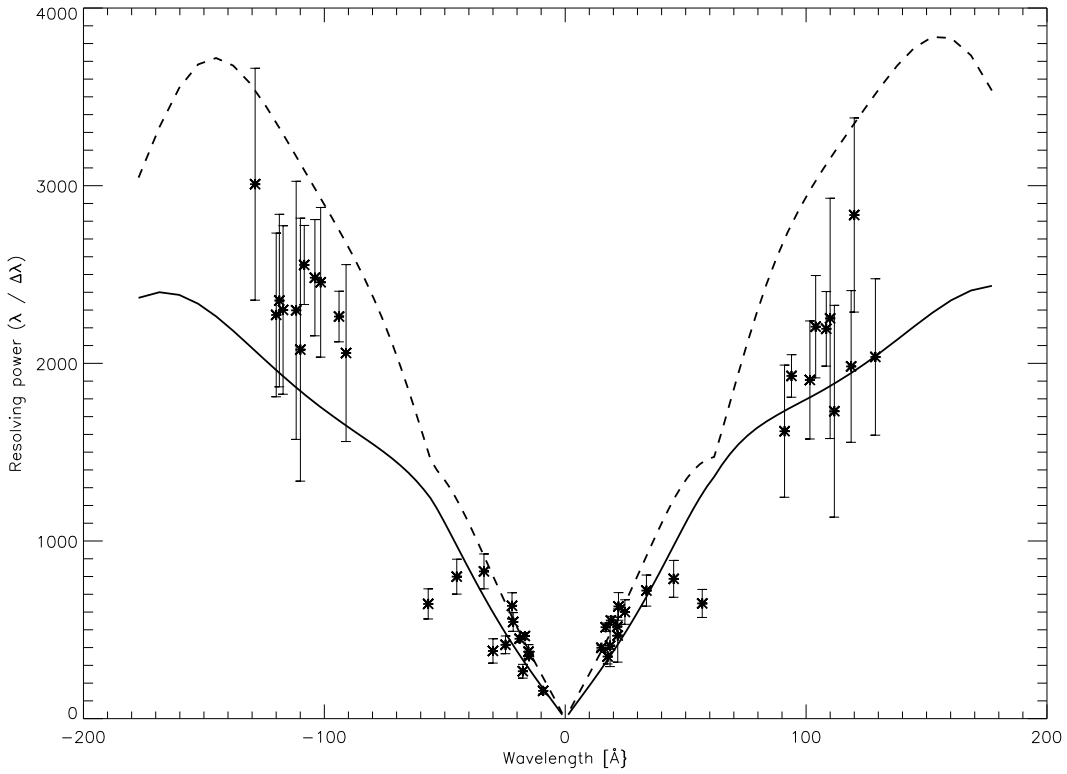

Figure 9.7: LETG spectral resolving power, as derived from observations of Capella (ObsIDs 1248, 1009, 58) and Procyon (ObsIDs 63, 1461) with the HRC-S. The analysis is based only on spectral lines thought not to be affected significantly by blending at the LETGS resolution. Measured line widths were corrected for source orbital, rotational, and thermal motions. The dashed line is an optimistic error budget prediction calculated from pre-flight models and instrument parameters. The conservative solid curve is based on in-flight values of aspect, focus, and grating period uniformity. The deviations from approximate linearity near ±60 Å and at the longest wavelengths arise from deviations of the HRC surface from the Rowland circle (see Figure 9.5). Deviations in the experimental data from a smooth curve are likely caused by hidden blends not predicted by the radiative loss model and by detector imaging nonlinearities discussed in Section 9.3.2.

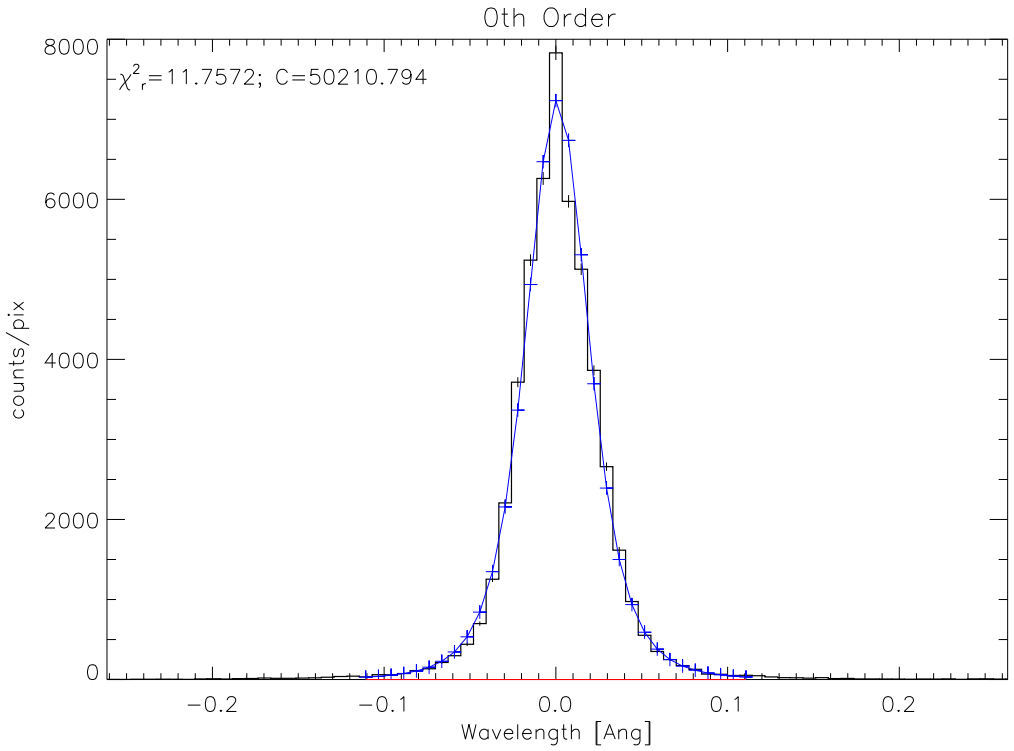

Figure 9.8: Observed LETG zeroth-order LRF from in-flight calibration observations of the active late-type binary Capella. The model profile—the continuous curve—is a Moffat function (see Equation 9.1) corresponding to the best-fit value of $\beta = 2.5$. While this function represents a statically-poor fit to this extremely high S/N zeroth-order profile, it is a good approximation to the LRF for lines in the dispersed spectrum containing many fewer counts.

on ray trace simulations using the *MARX* program and generally match observed line profiles to a level of 10% or better once intrinsic source broadening terms have been taken into account. Observers wishing to use line profile shapes as a diagnostic tool should keep in mind that, in the case of LETG+HRC-S observations, non-linearities in the HRC-S imaging can lead to significant distortions of observed line profiles (see Section 9.3.2).

Extended Sources

If a source is extended, there is no longer a unique mapping between the position of an event in the focal plane and wavelength, and this results in the apparent degradation of spectral resolving power. For very large sources, the grating resolution may be no better than the intrinsic ACIS energy resolution.

The effect of increased source size on the apparent LETG spectral resolving power has

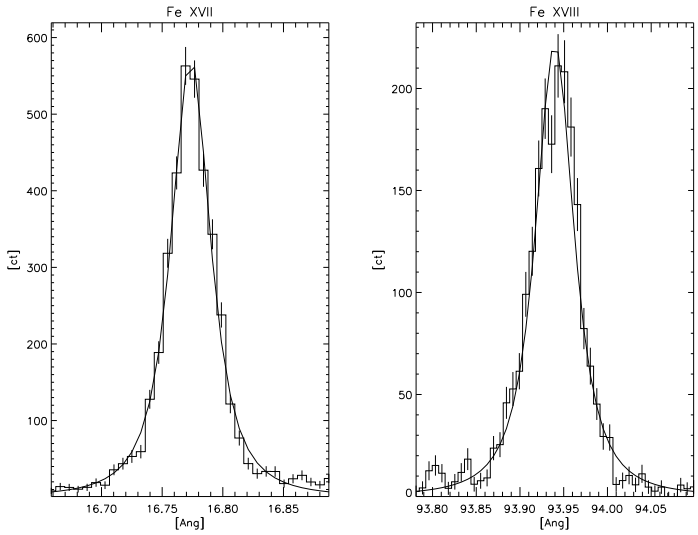

Figure 9.9: LETGS line response function as illustrated by two bright Fe lines (Fe XVII at ∼17 Å and Fe XVIII at ∼94 Å) using in-flight calibration observations of Capella. The solid curves are best-fit Moffat functions with $\beta = 2.5$.

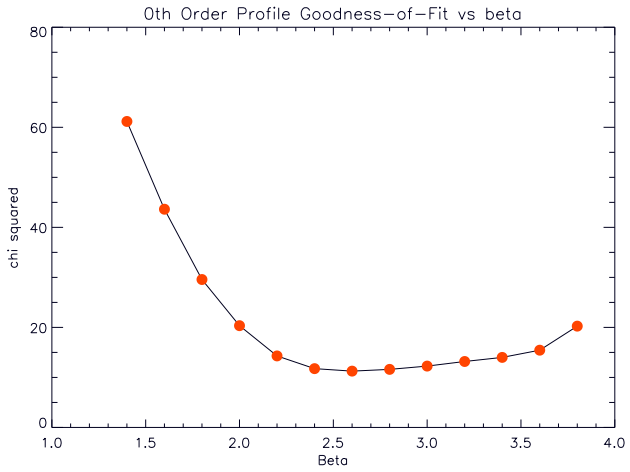

Figure 9.10: Goodness of fit for an LETGS zeroth order profile using a Moffat function (Equation 9.1); the best-fit profile is at an index of $\beta \sim 2.5$.

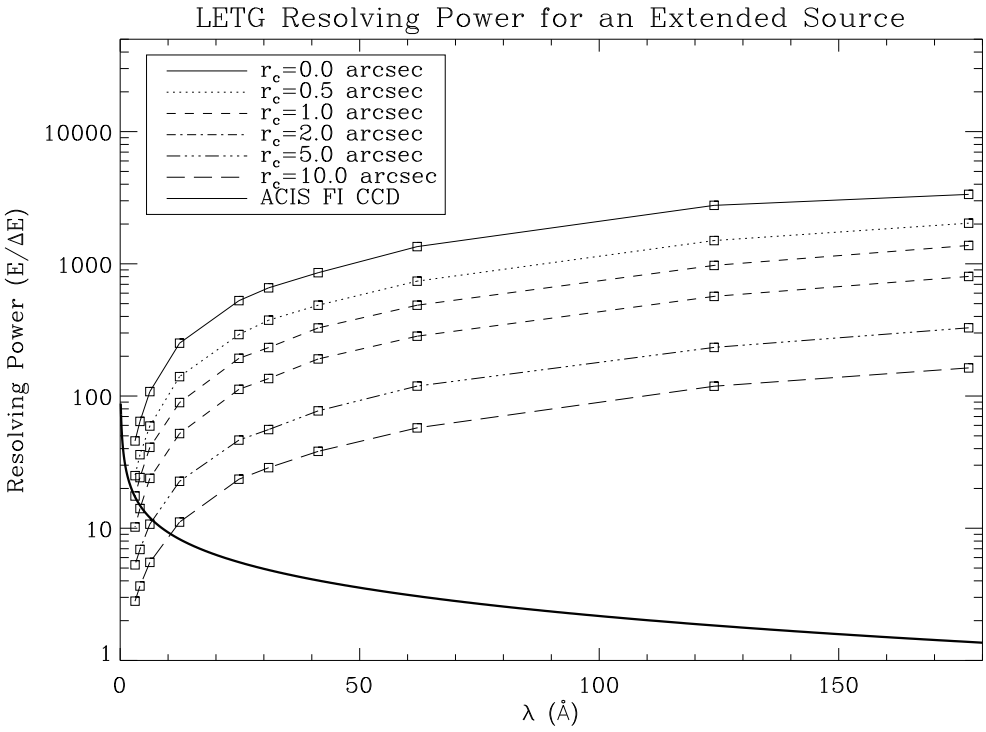

Figure 9.11: LETG spectral resolving power for extended sources on-axis. The predicted LETG resolving power ($E/\Delta E$) is shown versus wavelength for several source sizes. The *MARX* simulator has been used, and the source is represented by a β model (Equation 9.2). For comparison, the spectral resolution for ACIS front-illuminated CCD chips is shown (thick solid line). Note that the ACIS FI curve does not include the effects of CTI, which progressively degrades resolution away from the readout edge (most of this degradation can be compensated for in data analysis).

been simulated using *MARX*, and results are shown in Figure 9.11. Another illustration of the effect of source extent may be seen in Figure 9.30 (Section 9.4), which shows model spectra over a small wavelength range.

In each case, extended sources were modeled using a Beta model for the surface brightness profile. Beta models are often used to describe the distribution of emission in galaxies and clusters of galaxies, and have an identical form to the Moffat function used to describe the line profile above (Equation 9.1), except that the intensity dependence is radial and not with wavelength:

$$I(r) = \frac{I_0}{\left[1 + \left(\frac{r}{r_c}\right)^2\right]^{\beta}}. \tag{9.2}$$

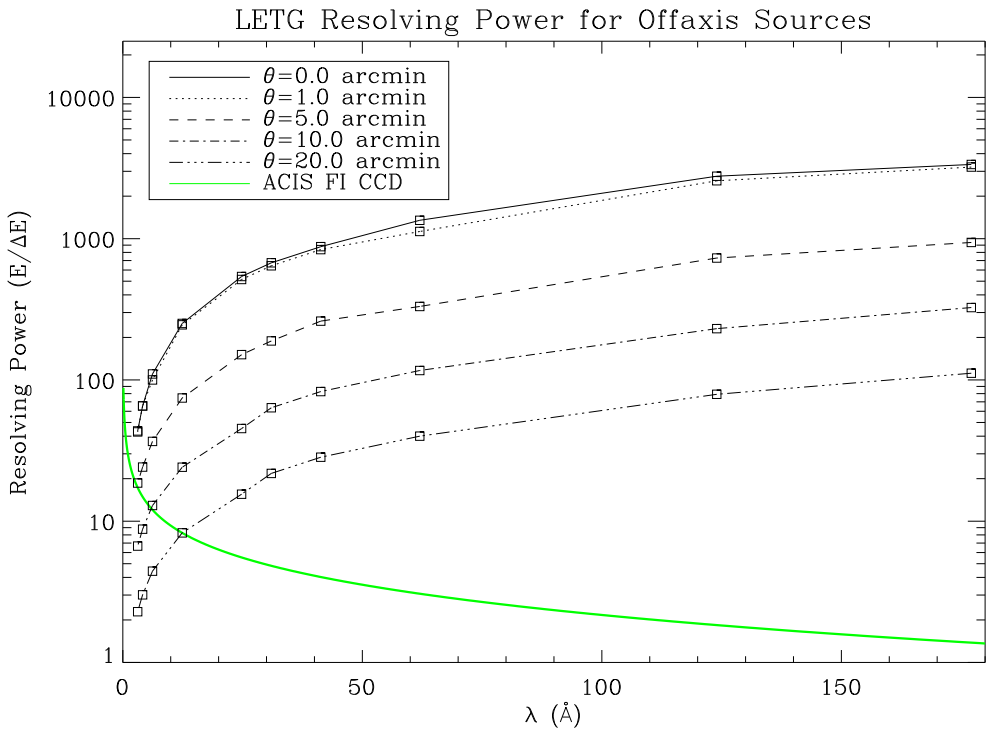

Figure 9.12: LETG spectral resolving power for off-axis sources. The predicted LETG resolving power ($E/\Delta E$) is shown versus wavelength for various off-axis distances. For comparison, the spectral resolution for ACIS front-illuminated CCD chips is shown. Note that the ACIS FI curve does not include the effects of CTI, which progressively degrades resolution away from the readout edge.

Here $I(r)$ is the surface brightness, r is the radius, and r_c characterizes the source extent. The value of β was set to a typical value of 0.75, and simulations were performed for different values of the source extent, r_c.

Off-Axis Sources

Similarly, for sources off-axis, the increased point-spread function decreases the spectral resolving power. The effect on off-axis sources has been simulated with *MARX* and is shown in Figure 9.12.

As with extended sources, an ACIS pulse-height spectrum may, in extreme cases, provide energy resolution comparable to or better than the LETG for a source far off-axis.

9.3.4 Grating Efficiency

Fine and Coarse Support Structure Diffraction

As explained in Section 9.1, the LETG has fine and coarse support structures which are periodic in nature and have their own diffraction characteristics. The fine support structure disperses photons perpendicularly to the main spectrum, with about 1/26 the dispersion of the main grating. The coarse support is a triangular grid, and creates a very small hexagonal diffraction pattern which is generally only discernible in zeroth order or for very bright lines. Examples of this secondary diffraction are visible in Figures 9.23, 9.24, and 9.27, all in Section 9.3.7.

The two support structures each diffract roughly 10% of the X-ray power, but the coarse-support diffraction pattern is so small that essentially all its photons are collected along with the primary spectrum during spectral-region extraction in data analysis. A significant fraction of the fine-support diffraction pattern, however, may lie outside the spectral extraction region, resulting in a loss of several percent of the total X-ray intensity (see, *e.g.*, Figure 9.23). The fractional retention of X-ray power in the source extraction region is referred to as the spectral extraction efficiency and is discussed in Section 9.3.5.

Total Efficiencies

The zeroth and 1st order grating efficiencies are based on a rhomboidal grating bar analytical model and verified by ground calibration. Higher orders have been calibrated relative to 1st order using flight data. Plotted values (Figure 9.13) are for the total diffraction efficiency (including photons diffracted by the coarse and fine support structures), with negative and positive orders summed. Even orders are generally weaker than odd orders up through roughly 6th order.

The wiggles near 80 Å, and the stronger features near 6 Å, arise from partial transparency of the gold grating material to X-ray photons. Note that there are no absorption-edge features from C, N, or O in the LETG efficiency as there are in the HETGbecause the LETG does not use a polyimide support film.

9.3.5 Effective Area

The LETGS effective area for any diffraction order is the product of HRMA effective area, net LETG efficiency for that order (including spectral extraction efficiency), and overall detector efficiency (which varies slightly depending on exactly where the diffracted spectrum falls on the detector). For LETG/ACIS-S an additional factor is the Order Sorting Integrated Probability (OSIP), which is determined by the width of the ACIS-S energy filter for each diffraction order. All these quantities vary with wavelength.

Of the contributors listed above, the HRMA EA (see Chapter 4) is the best calibrated within the LETGS energy band. The largest contributor to the LETG/HRC-S effective area uncertainty is the efficiency of the HRC-S, especially at longer wavelengths (> 44 Å; < 0.28 keV) where ground calibration is very difficult or impossible. In-flight calibration

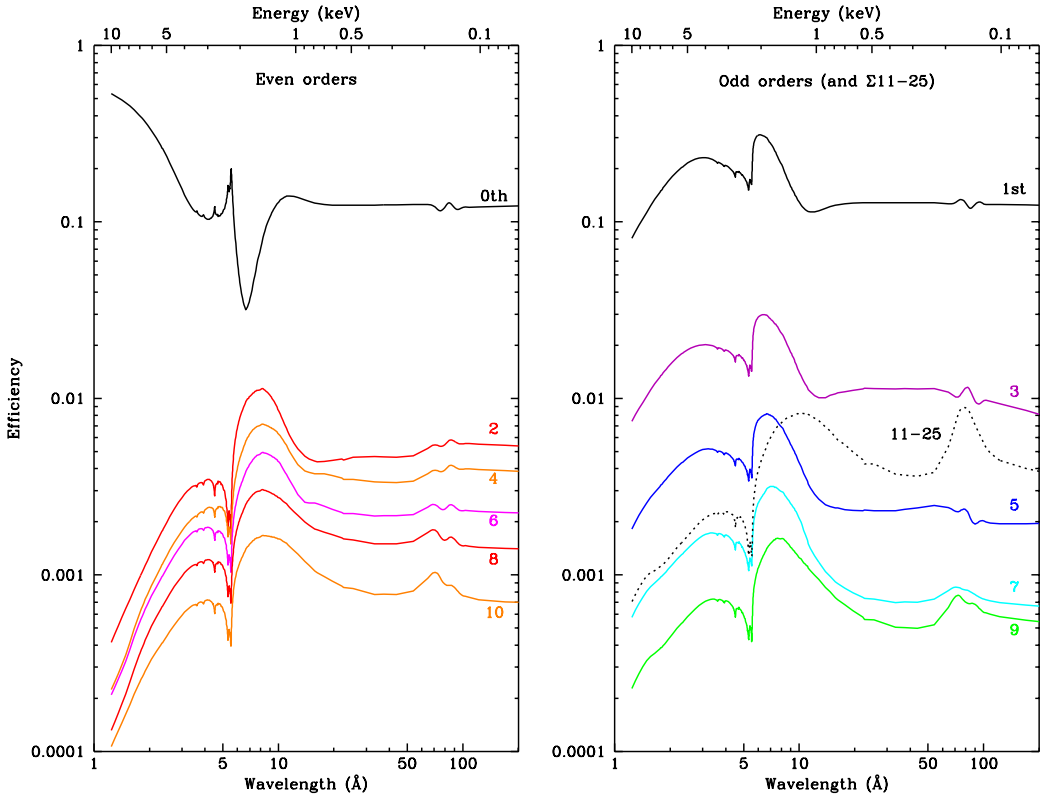

Figure 9.13: LETG grating efficiency. Combined positive and negative order efficiency is plotted versus wavelength and energy. For clarity, even and odd orders are plotted separately; the sum of orders 11–25 is also shown. Plotted values include all support structure diffraction. Net efficiency when using a spectral extraction region (see Figures 9.14 and 9.15) is 10–15% lower. Features near 6 and 80 Å are due to the partial transparency of the gold grating material at these wavelengths.

(see Section 9.2.2), particularly of the net 1st order effective area, has provided the best and most extensive data, and the effective area is now believed to be accurate to a level of approximately 10–15% or better across the entire bandpass. Effective areas for orders 2–10 have been calibrated relative to 1st order to an accuracy of 5–10% (best for 3rd order and generally worsening with increasing m) using flight data. Uncertainties for $\lambda \lesssim 6$ Å, $m\lambda \gtrsim 80$ Å, and orders beyond 10th may be larger but are usually unimportant.

The QE of both the ACIS-S and HRC-S detectors is slowly decreasing over time (see Sections 6.5.1 and 7.10), and so the LETGS effective areas are also time dependent, as reflected in the CALDB (beginning with CALDB 4.4.5 for the HRC-S). The ACIS-S QE decrease is caused by a thickening layer of contaminant on the detector's Optical Blocking Filter, and is most apparent at low energies. As of 2016, it is recommended that LETG+ACIS-S spectra be centered on row 520 (SIM-Z=0) instead of row 180 (SIM-Z=−8 mm) in order to sample a region of ACIS-S with thinner contaminant and therefore higher effective area (see Figure 9.6 and Section 9.4.1/ACIS-S). The HRC-S decrease of ∼2% per year is largely independent of energy and believed to be caused by charge extraction across the entire detector (mostly from cosmic ray background events), with a slightly higher rate of decrease around the aimpoint and at wavelengths longer than ∼ +140 Å. The HRC-S operating voltage was raised first in 2012 and again in following years to counteract decreasing QE.

Instrument Spectral Features

In addition to fixed-position detector features (primarily detector segment gaps—see Section 9.3.2), there are instrumental spectral features which occur at fixed energies because of absorption edges in the materials comprising the HRMA, LETG, and HRC-S or ACIS-S. The edges are tabulated in Table 9.4 and can be seen in the effective area curves (such as Figure 9.17) as decreases or increases in effective area depending on whether the material is part of the mirror, the filter, or the detector. Every effort has been made to adequately calibrate *Chandra* over its entire energy range, but it should be understood that effective areas near absorption edges are extremely difficult to quantify with complete accuracy and uncertainties in these regions are inevitably higher.

Spectral Extraction Efficiency

In practice, it is impossible to "put back" photons which undergo secondary diffraction (from the coarse and fine support structures) in a real observation. Instead, one defines an extraction region for the observed spectrum and adjusts the derived spectral intensities to account for the fraction of total events that are contained within the extraction region.

The default extraction region for the LETG+ACIS-S configuration is a rectangle; that for the LETG+HRC-S configuration is "bow tie" shaped, comprising a central rectangle abutted to outer regions whose widths flare linearly with increasing dispersion distance (see Figure 9.14). The shape of the bow tie has been roughly optimized to match the astigmatic cross-dispersion that is a feature of Rowland-circle geometry, with the

Table 9.4: Instrumental Absorption Edges

Instrument	Element	Edge	Energy (keV)	Wavelength (Å)
HRC	Cs	L	5.714	2.170
HRC	Cs	L	5.359	2.313
HRC	I	L	5.188	2.390
HRC	Cs	L	5.012	2.474
HRC	I	L	4.852	2.555
HRC	I	L	4.557	2.721
LETG	Au	M	3.425	3.620
LETG	Au	M	3.148	3.938
HRMA	Ir	M	2.909	4.262
LETG	Au	M	2.743	4.520
HRMA	Ir	M	2.550	4.862
LETG	Au	M	2.247	5.518
LETG	Au	M	2.230	5.560
HRMA	Ir	M	2.156	5.750
HRMA	Ir	M	2.089	5.935
ACIS	Si	K	1.839	6.742
HRC, ACIS	Al	K	1.559	7.953
HRC	Cs	M	1.211	10.24
HRC	I	M	1.072	11.56
HRC	Cs	M	1.071	11.58
HRC	Cs	M	1.003	12.36
HRC	I	M	0.931	13.32
HRC	I	M	0.875	14.17
HRC	Cs	M	0.7405	16.74
HRC	Cs	M	0.7266	17.06
ACIS	F	K	0.687	18.05
HRC	I	M	0.6308	19.65
HRC	I	M	0.6193	20.02
HRC,ACIS	O	K	0.532	23.30
HRMA	Ir	N	0.496	25.0
HRC	N	K	0.407	30.5
HRC,ACIS	C	K	0.284	43.6
HRC	Al	L	0.073	170

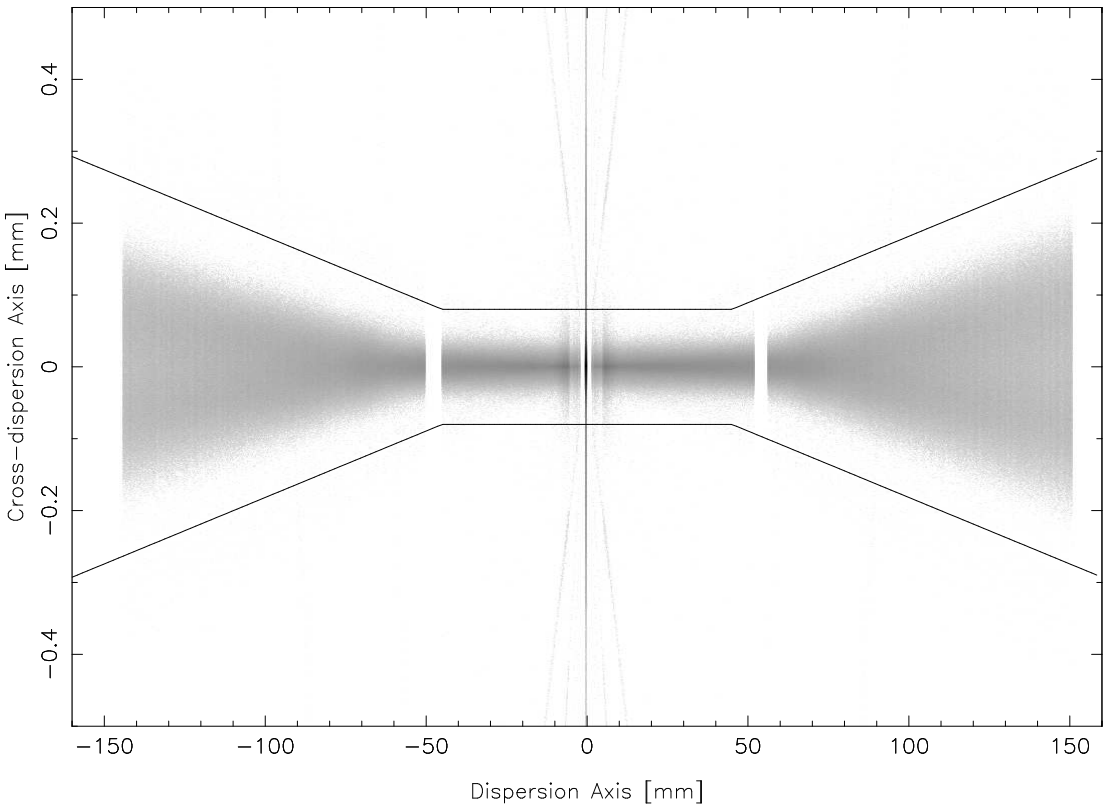

Figure 9.14: A *MARX* simulation of a flat spectrum illustrating the broadening of the LETG+HRC-S profile in the cross-dispersion direction and showing the "bow tie" spectral extraction window. Note that the vertical axis is highly stretched and that the true diffraction pattern has some minor asymmetries. The default extraction region for LETG+ACIS-S spectra (not shown) is rectangular with a half-width of 0.12 mm.

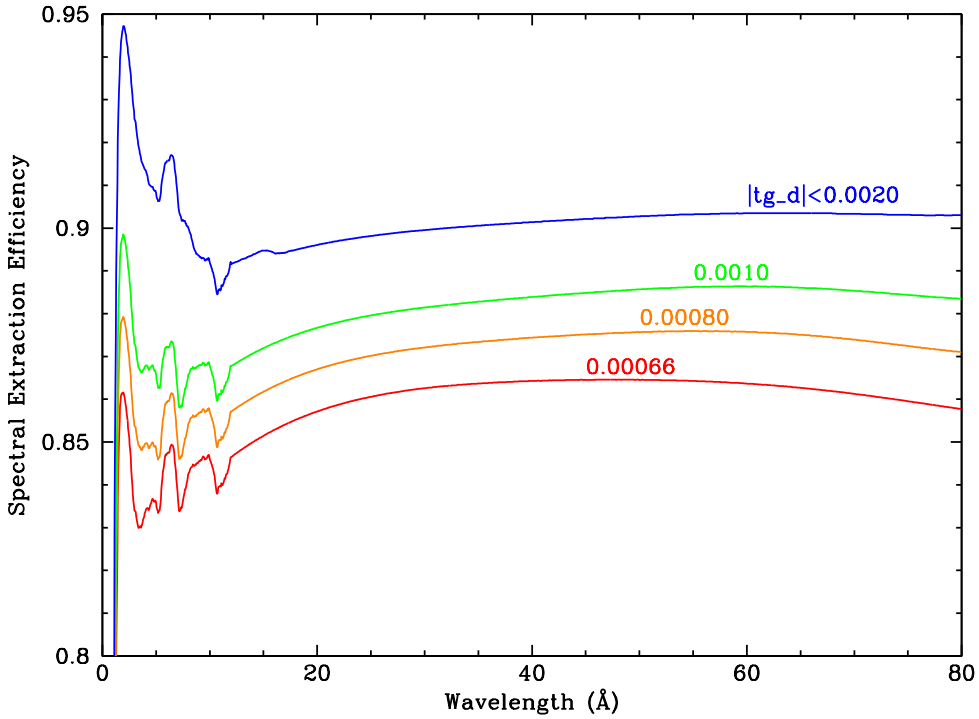

Figure 9.15: Spectral extraction efficiencies for LETG+ACIS for various extraction region half-widths (tg_d units are deg). The default spectral extraction region is $\mid tg\_d \mid < 0.0008$ deg. Approximately 10% of the total power is diffracted by the fine-support structure and most of it falls outside the extraction region; the peaks toward shorter wavelengths reflect the inclusion of progressively higher orders of cross-dispersed events within the extraction region. The otherwise general trend of declining efficiency toward short wavelengths is caused by increased scattering. The slight fall-off beyond ∼60 Å is due to astigmatism, i.e., an increasing fraction of events are lost as the dispersion pattern broadens (see Figure 9.14).

goal of including as much of the diffracted spectrum as possible while minimizing the included background. Extraction efficiencies for LETG+ACIS are illustrated in Figure 9.15. The extraction efficiency for the LETG+HRC bow tie is 85–86% at most wavelengths. For custom analysis (such as when narrower or wider extraction regions are needed), *CIAO* permits adjustment of spectral and background regions by the user. A narrower extraction region optimized for the analysis of low-S/N spectra is available at https://cxc.harvard.edu/cal/Letg/LetgHrcEEFRAC. Whatever spectral region is chosen, *CIAO* computes the extraction efficiency and includes it as a factor in the *Response Matrix File* (RMF). Users then apply this RMF together with the effective area from the *CIAO*-computed *Ancillary Response File* (ARF) during subsequent analysis.

Zeroth and First-Order Effective Areas

Although the HRC-S is the default detector for the LETG, other detector configurations are possible. Figures 9.16 and 9.17 show effective areas for the LETG zeroth and 1st orders, respectively, when using the HRC-S, HRC-S with LESF, or ACIS-S as the readout detector. Based upon these and other plots, the various trade-offs as to the use of each detector are thoroughly discussed in Section 9.4. Users interested in the low energy response using ACIS should read carefully the discussion in Section 6.5.1.

Off-Axis and Extended Sources

Differences in the LETGS effective area for off-axis and significantly extended sources compared to the on-axis point source case are primarily determined by the HRMA vignetting function (see Chapter 4).

High-Order Diffraction Effective Areas

Although the LETG (and HETG) have been designed to reduce complications from higher-order diffraction by suppressing even orders, many grating spectra will have overlapping diffraction orders. When ACIS-S is used as the detector, its intrinsic energy resolution can be used to separate orders. The situation is more complicated, however, with HRC-S, which has very little energy resolution. Detector options and various data analysis techniques are described in Section 9.4.

The relative contribution of higher-order photons with different detector configurations can be estimated by inspection of Figures 9.18, 9.19, and 9.20. As an example, an observer planning to use the LETG/HRC-S configuration may want to determine the intensity of a line at 45 Å but knows that that line may be blended with the 3rd order of a 15 Å line with 10 times the emitted intensity of the 45 Å line. Looking at Figure 9.18, the 1st- and 3rd-order curves at $m\lambda = 45$ Å show that the 3rd-order value is about one-tenth the 1st-order value. Multiplying by 10 (the ratio of the emitted intensities of the 15 and 45 Å lines), $\sim$50% of the feature at $m\lambda = 45$ Å will come from the 15 Å line. A fuller explanation, with

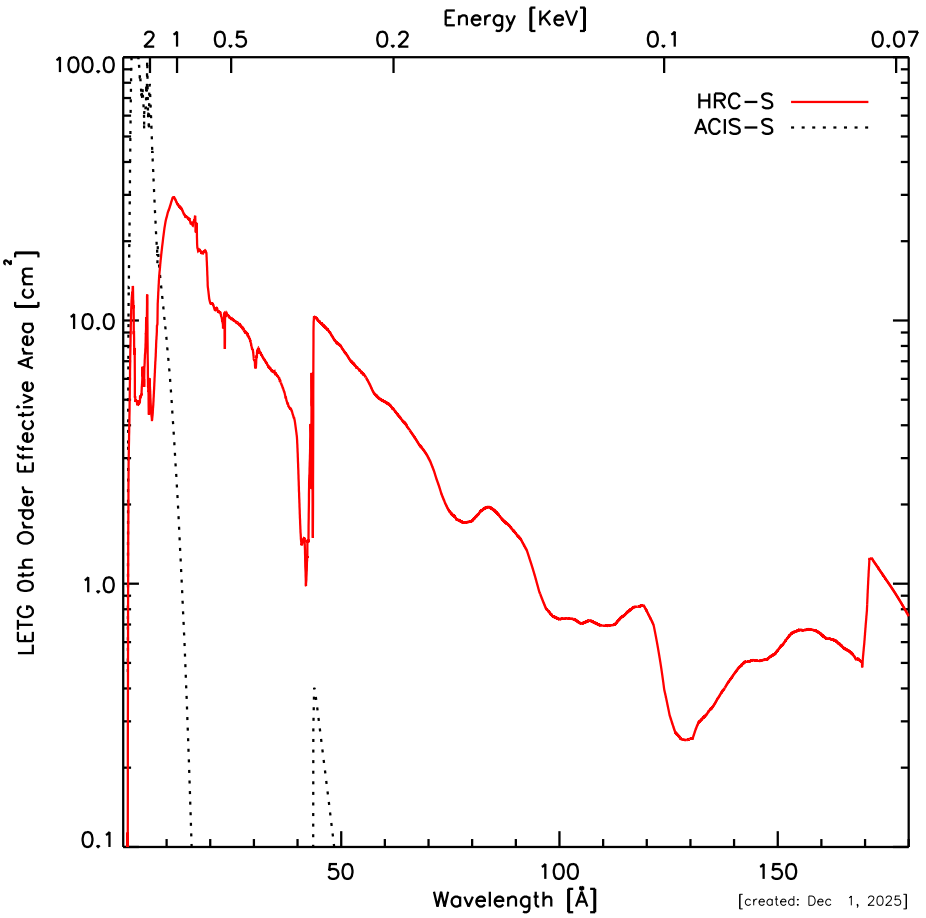

Figure 9.16: LETGS zeroth-order effective area for an on-axis point source for the LETG with HRC-S and ACIS-S detectors. The zeroth-order effective area for the HRC-S/LESF combination is the same as for the HRC-S. LETG+ACIS-S areas include the effects of time-dependent contamination.

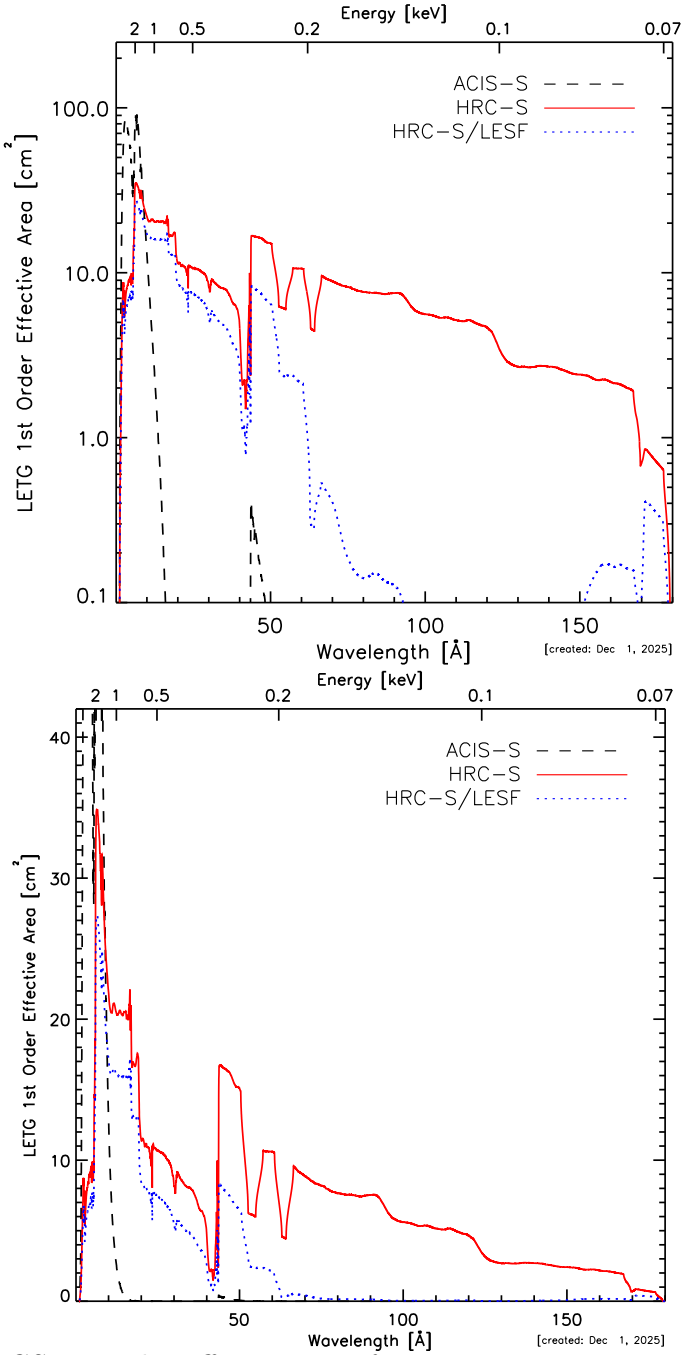

Figure 9.17: LETGS 1st-order effective area for an on-axis point source, with HRC-S, HRC-S/LESF, and ACIS-S detector configurations with log (top) and linear (bottom) scaling. Positive and negative orders are summed. LETG+ACIS-S areas include the effects of time-dependent contamination. Note that the vertical scale of the linear plot has been truncated.

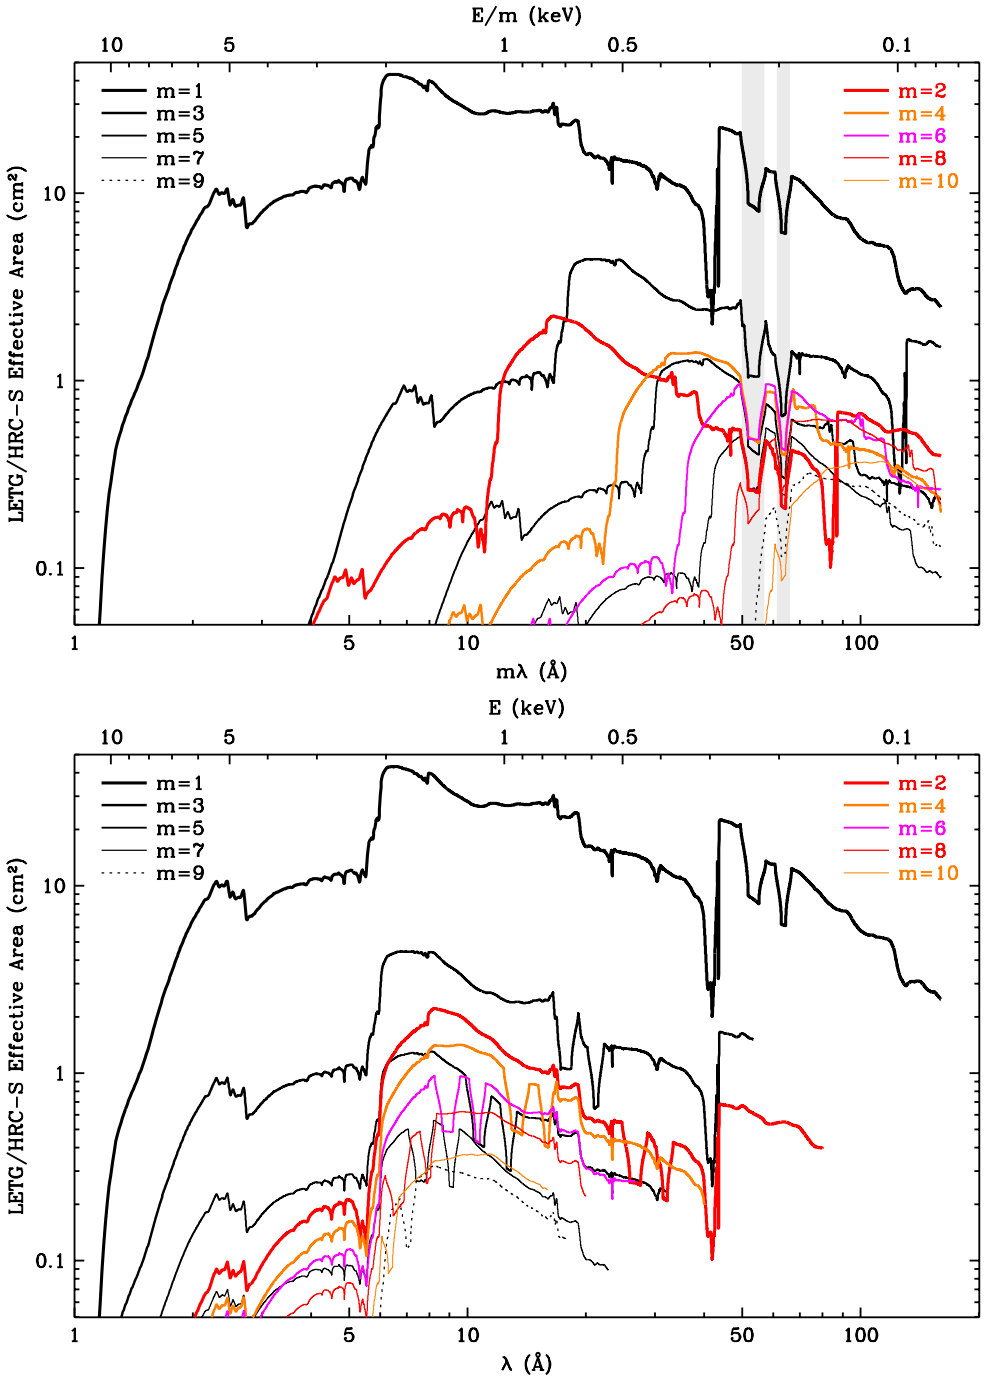

Figure 9.18: The combined LETG/HRC-S effective area, illustrating the relative strengths of 1st and higher orders, plotted versus $m\lambda$ (top) and λ (bottom). Positive and negative orders are summed. Current effective areas are about 30% lower than shown, but the ratios between orders are correct. Light shading in the top panel marks plate gaps around $m\lambda = -53$ Å and $+64$ Å. See the text for an example of how to determine the relative strength of overlapping lines from different orders and https://cxc.harvard.edu/ciao/threads/hrcsletg_orders/ for further information.

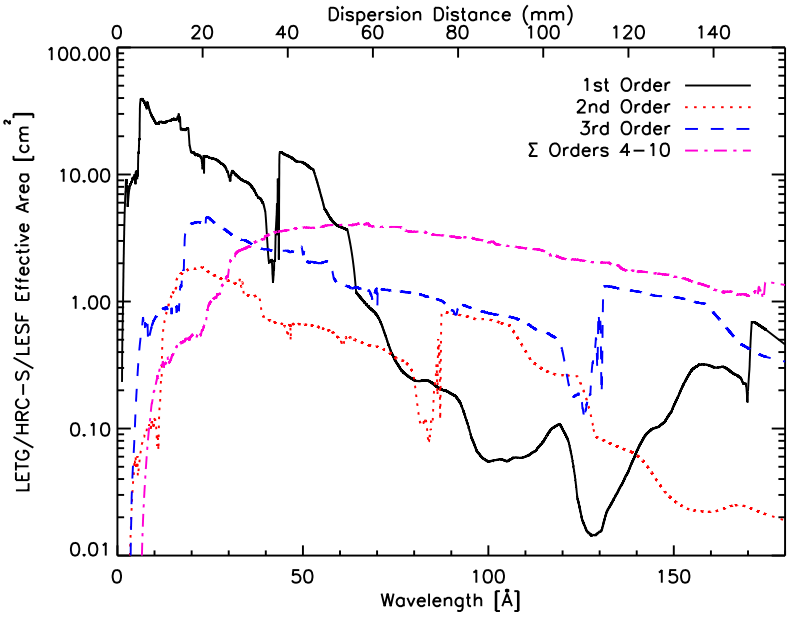

Figure 9.19: The combined HRMA/LETG/HRC-S/LESF effective areas for 1st and higher orders. Positive and negative orders are summed.

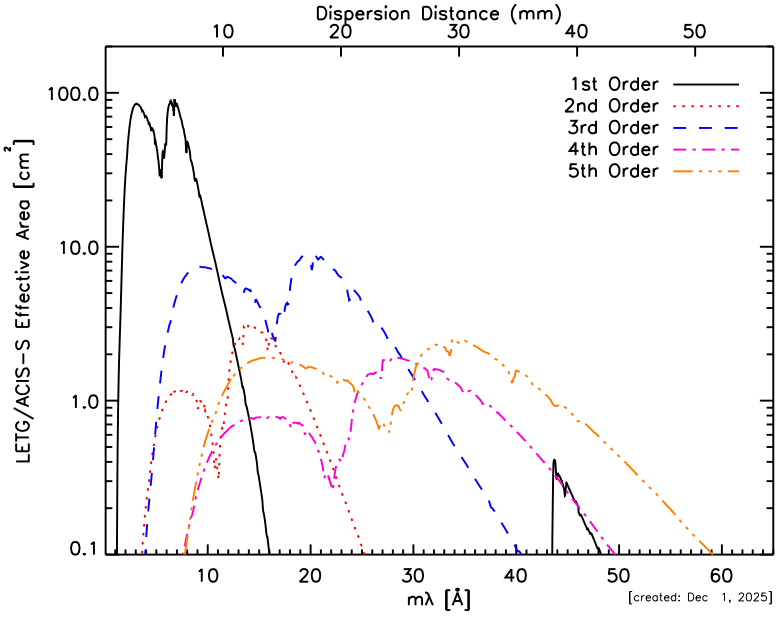

Figure 9.20: The combined HRMA/LETG/ACIS-S effective areas for 1st and higher orders. LETG+ACIS-S areas include the effects of time-dependent contamination. Positive and negative orders are summed.

color figures and more examples for the LETG/HRC-S with line and continuum sources, can be found at https://cxc.harvard.edu/ciao/threads/hrcsletg_orders/.

9.3.6 Background

The LETG is always used in conjunction with a focal-plane detector, so LETGS spectra will exhibit that detector's intrinsic, environmental, and cosmic background. The components of the background of the HRC are discussed in Section 7.11. The quiescent background rate over the full detector varies with the solar cycle (see Figure 9.21) but is always a significant fraction of the 183 cts s^{-1} telemetry limit. Imposition of the HRC-S spectroscopy window reduces the rate to between 55 and 130 cts s^{-1}, as discussed below.

HRC-S Exposure Windows, Deadtime, and Timing Resolution

To avoid constant telemetry saturation, the HRC-S is operated in a default windowed down "edge-blanking" configuration, in which data from only 6 of the 12 coarse taps in the center of the detector in the cross-dispersion direction are telemetered (see Section 7.11.2). The edge-blanking creates an active detector area slightly less than 10 mm (3.4 arcmin) in the cross-dispersion direction. This window easily accommodates the (dithered) dispersed spectra of point sources; other windows may be specified for extended sources or other special cases.

As long as the total count rate is below the 183 cts s^{-1} telemetry limit, detector deadtime is negligible (and recorded as a function of time in the secondary science `.dtf` files—net exposure time is recorded in the image FITS file header). During background "flares" arising from an increased flux of solar wind particles, however, the background rate may rise above the 183 cts s^{-1} telemetry limit. During these times detector deadtime may become significant. Current data processing algorithms correct for this deadtime with a typical accuracy of $\sim$10% or better.

Time resolution approaching 16 μsec can be achieved with the HRC if the data rate is below telemetry saturation. To leave ample margin for telemetry in case of background flares, a special Imaging Mode (see Section 7.12) is used for high-resolution timing observations. This mode utilizes only the central region of the HRC-S detector and provides a field of view of approximately 7 arcmin $\times$ 30 arcmin. However, if telemetry saturation does take place then the time resolution with HRC-S is approximately the average time between events.

HRC-S Background Reduction Using Pulse-Height Filtering

The quiescent background rate in HRC-S Level 2 data varies by more than a factor of two over a solar cycle (see Figure 9.21). A "typical" rate of 1.0×10^{-6} cts s^{-1} $pixel^{-1}$ corresponds to 5.76×10^{-5} cts s^{-1} $arcsec^{-2}$, or 0.10 cts/pix in 100 ks. The extent of a dispersed line in the LETGS spectrum is approximately 0.07 Å (9.5 pixels) and the

spectral extraction region is 25 to 85 pixels wide in the cross-dispersion direction; the typical background rate therefore yields 24 to 81 background counts beneath the line in a 100 ks exposure. However, the HRC-S pulse height distribution is sufficiently narrow that a large fraction of pulse-height space can be excluded from the data to further reduce the background, which has a relatively broad pulse-height distribution.

A new HRC-S gain map and associated LETG/HRC-S background filter were released in 2020, with an update in mid-2021 (CALDB 4.9.6) that applies to all observations prior to 2021-May-14, when the HRC-S high voltage was raised to restore gain and QE. The background filter removes more than half of the Level 2 background at wavelengths longer than 20 Å, with a loss of only 1.0% of X-ray events (see Figure 9.22). Gain calibration details may be found at https://cxc.harvard.edu/cal/Letg/Hrc_bg.

There is no calibration of the time dependence of the HRC-S gain after 2021-May-14 and therefore the LETG/HRC-S background filter cannot be applied to observations after that date.

Relevance for Higher Orders The peak of the pulse-height distribution increases weakly with photon energy, so that the background filter, which is tuned to remove 1% of 1st order events, will remove increasing fractions of events for higher orders (e.g., roughly 30% of 8th order). This extra filtering of higher orders will have negligible effect for nearly all analyses, but should be considered when studying wavelength ranges with very heavy higher order contamination. A figure showing the effects of filtering on orders 1 to 12 is shown near the bottom of https://cxc.harvard.edu/cal/Letg/Hrc_bg.

Relevance for Observation Planning There are two backgrounds relevant for the LETG/HRC-S. The first is the Level 1 data event rate over the 3-plate standard spectroscopy region, which is 55–130 cts s^{-1} during quiescence (see Figure 9.21), but can rise during background "flares" to cause telemetry saturation when the total (background plus sources) rate reaches 183 cts s^{-1}. The other is the filtered Level 2 background rate in the extracted spectrum. A "typical" rate of 1.0×10^{-6} cts s^{-1} $pixel^{-1}$ (see Figure 9.21) corresponds to ∼10–25 cts (depending on wavelength—see Figure 9.22) in a 0.07 Å spectral bin per 100,000 s integration during quiescence, which may be used for estimating signal-to-noise (see also the discussion in Section 9.5.1).

There is a third background rate which will be of interest when high time resolution (sub-msec) is required, which is the count rate before any on-board screening is applied. See Sections 7.12 and 7.15.1 for more information on the HRC-S Timing Mode.

ACIS-S Background

As with the HRC-S detector, background rates in ACIS are somewhat higher than expected, but lower than in the HRC. Pulse-height filtering applied during order separation further reduces the effective ACIS-S background to extremely low levels when used with gratings. The reader is directed to Chapter 6 for further discussion.

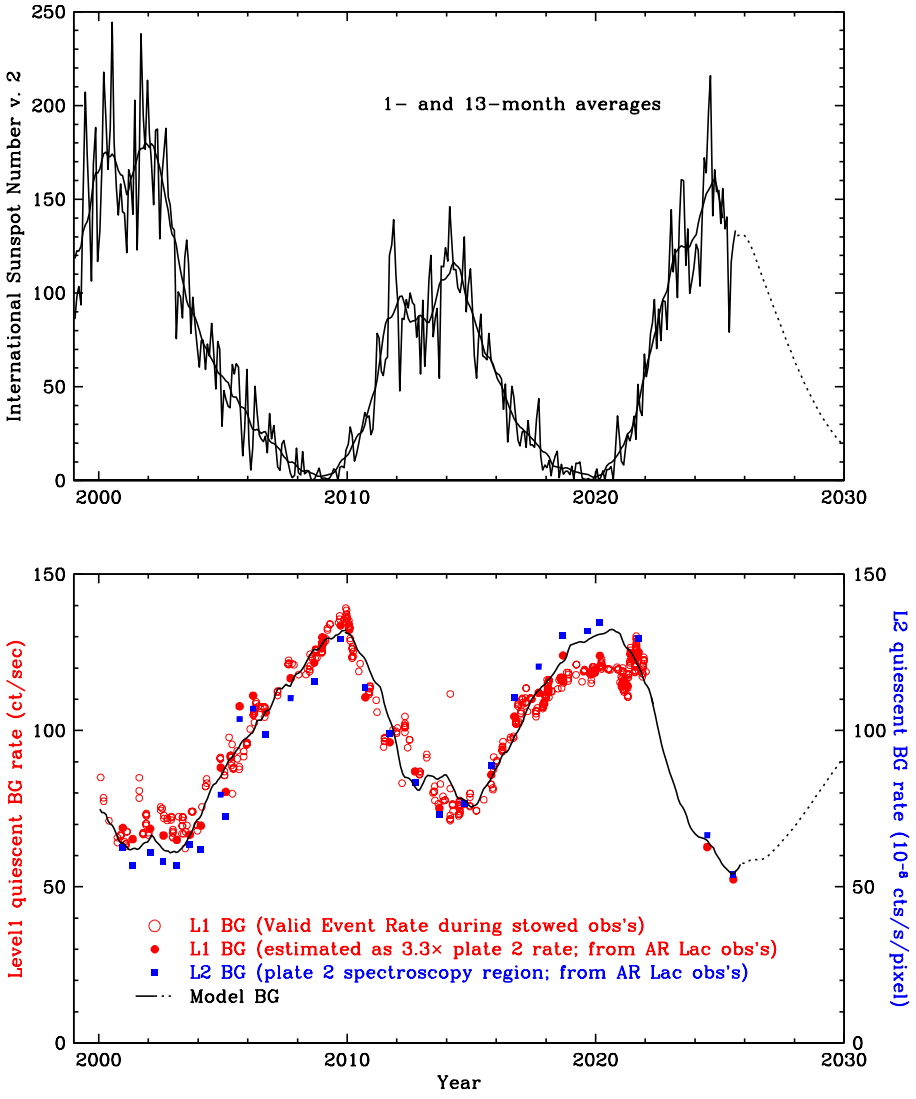

Figure 9.21: Solar cycle and HRC-S background. Top: Monthly sunspot numbers in 1- and 13-month averages; the dotted curve is the predicted 13-month curve. Bottom: Level 1 (red) and Level 2 (blue) background rates and modeled background; the Level 1 rate is for the LETGS standard spectroscopy region, derived from AR Lac monitoring data (without a grating) and from NIL-mode observations when the HRC-S is in the stowed position. The "typical" Level 2 (L2) rate is 1.0×10^{-6} cts s^{-1} pixel^{-1}, but this varies with the solar cycle. L1 and L2 background rates are not exactly correlated because of changes in the effects of L1-to-L2 filtering over time, particularly during the period of increasingly low gain prior to the May 2021 high voltage increase. The continuous black curve represents a simple model of the HRC-S background rate based on sunspot number with a 1-year lag (or 0.7 years after the 2009 solar minimum) corresponding to the approximate time for the solar wind to reach the heliopause, where its magnetic field helps deflect cosmic rays from entering the solar system. Rates in Figure 9.22 correspond to a Level 2 rate of 1.1×10^{-6} cts s^{-1} pixel^{-1}.

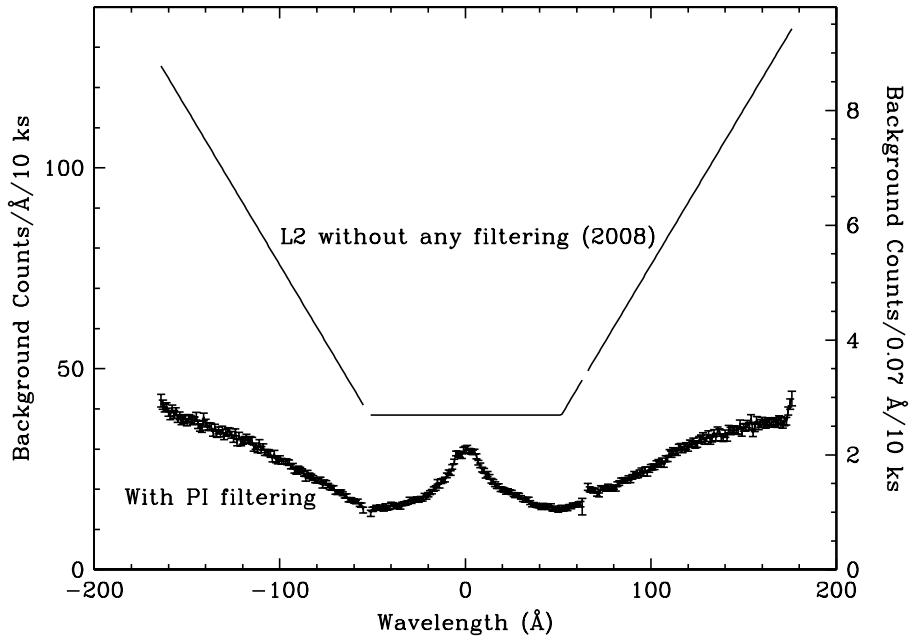

Figure 9.22: The 2008-Jan LETG+HRC-S background rate with and without pulse-height (PI) filtering. These Level 2 rates correspond to 1.1×10^{-6} cts s^{-1} $pixel^{-1}$ (see Figure 9.21). Data are shown using the standard "bow tie" spectral extraction region (see Figure 9.14). X-ray losses from the observed spectrum as a result of the filtering for 1st order are ∼1.0%; losses will be somewhat larger for higher orders. See https://cxc.harvard.edu/cal/Letg/Hrc_bg for more information.

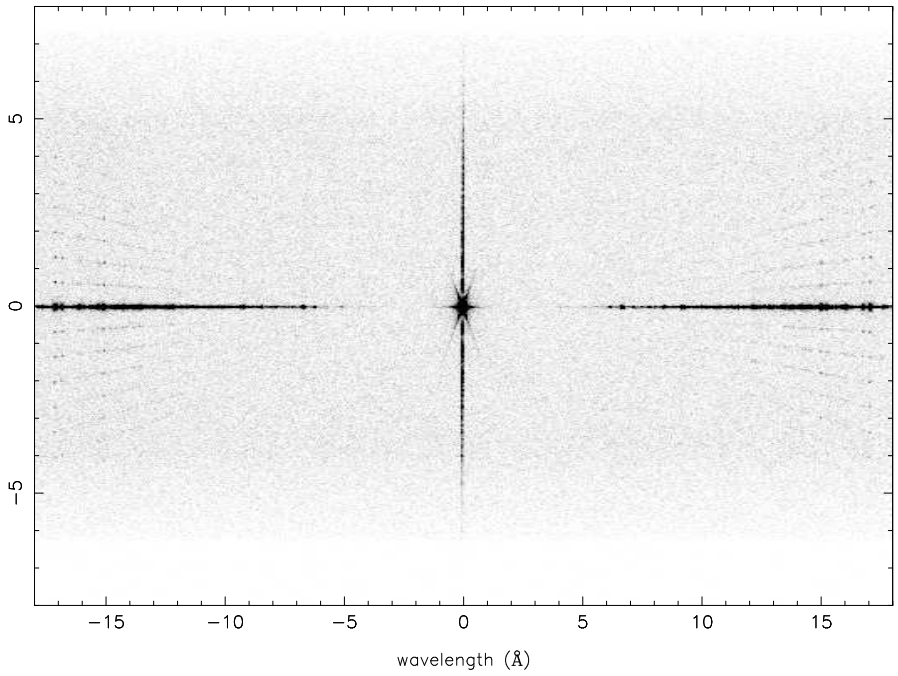

Figure 9.23: HRC-S detector image of LETGS observation of Capella. Vertical axis is stretched by a factor of 1.6 relative to horizontal; only the central 30 mm of the central plate is shown. The full extent of the telemetered six-tap cross-dispersion window is shown and measures 9.9 mm. The areas of reduced background at top and bottom are due to dither effects. Star-shaped coarse support structure diffraction is seen around zeroth order, and "cat's whiskers" fine support structure diffraction is seen above and below the primary dispersion axis, as well as in the vertical line through zeroth order.

9.3.7 Example Data and Support Structure Diffraction

Figure 9.23 is a detector image from an 85 ks LETGS observation of Capella (ObsID 1248). The central 30 mm of the dispersion axis and the full extent of the telemetered cross-dispersion window (9.9 mm) are shown. Lines radiating from zeroth order above and below the primary dispersion axis are due to fine-support structure diffraction, and the six-pointed star around zeroth order is from coarse-support diffraction (see Section 9.1.4 for a description of the LETG support structures). Figure 9.24 is a close-up of the bright Fe XVII, Fe XVIII, and O VIII lines between ∼15 and 17 Å, in which many orders of fine-support diffracted flux can be seen.

Figure 9.25 is an HRC-S image of a second Capella observation (ObsID 1420, 30 ks), showing positive order dispersion. The increasing cross-dispersion extent of lines at longer wavelengths is due to astigmatism in the HRMA/LETG system (see also Figure 9.14). The positive order HRC-S plate gap is seen at ∼63 Å. An extracted Capella spectrum (ObsID 62435, 32 ks), is shown in Figure 9.26.

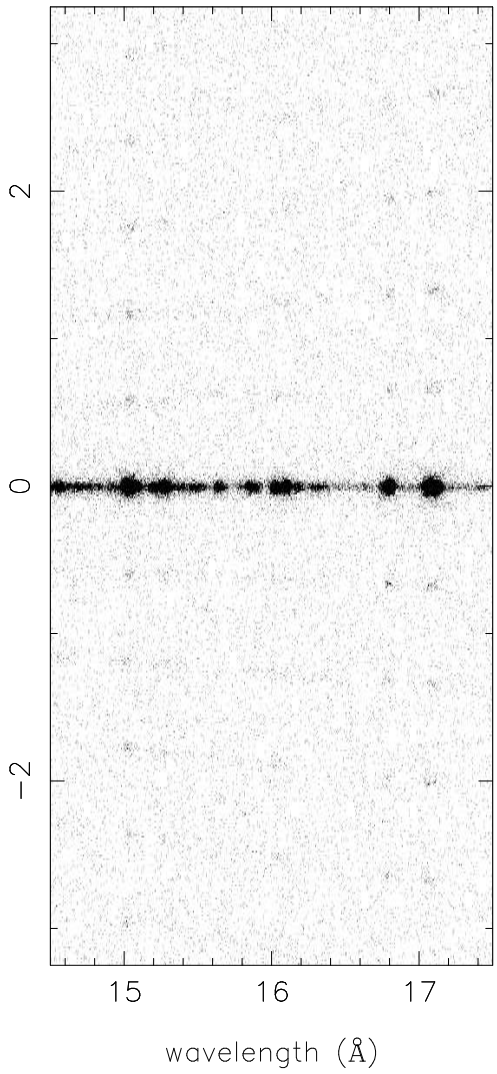

Figure 9.24: Detail of Figure 9.23, showing a zoomed-in region around the LETG/HRC-S image of bright lines in Capella. The vertical and horizontal scales in this image are equal. The Fe XVII lines at ∼15 and 17 Å are the brightest lines in the LETG Capella spectrum. Faint features above and below the primary spectrum are due to fine support structure diffraction.

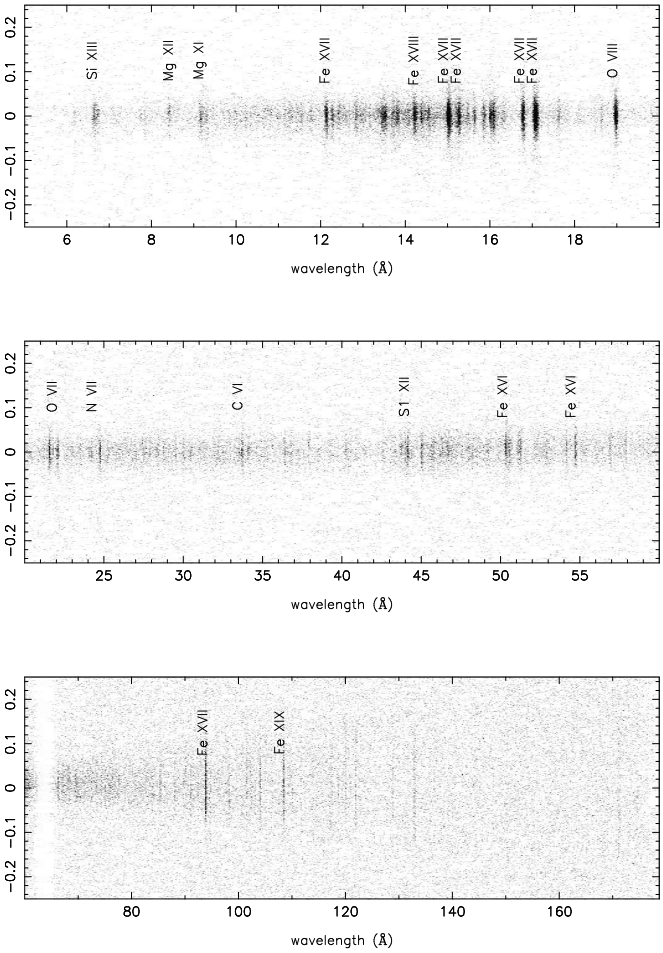

Figure 9.25: HRC-S detector image of a Capella observation, showing positive order dispersion. Images in the top, middle, and bottom panels are horizontally compressed by factors of 11, 28, and 84 respectively. The increasing cross-dispersion extent of lines at longer wavelengths is due to astigmatism in the HRMA/LETG system. The positive-order HRC-S plate gap is at ∼63 Å.

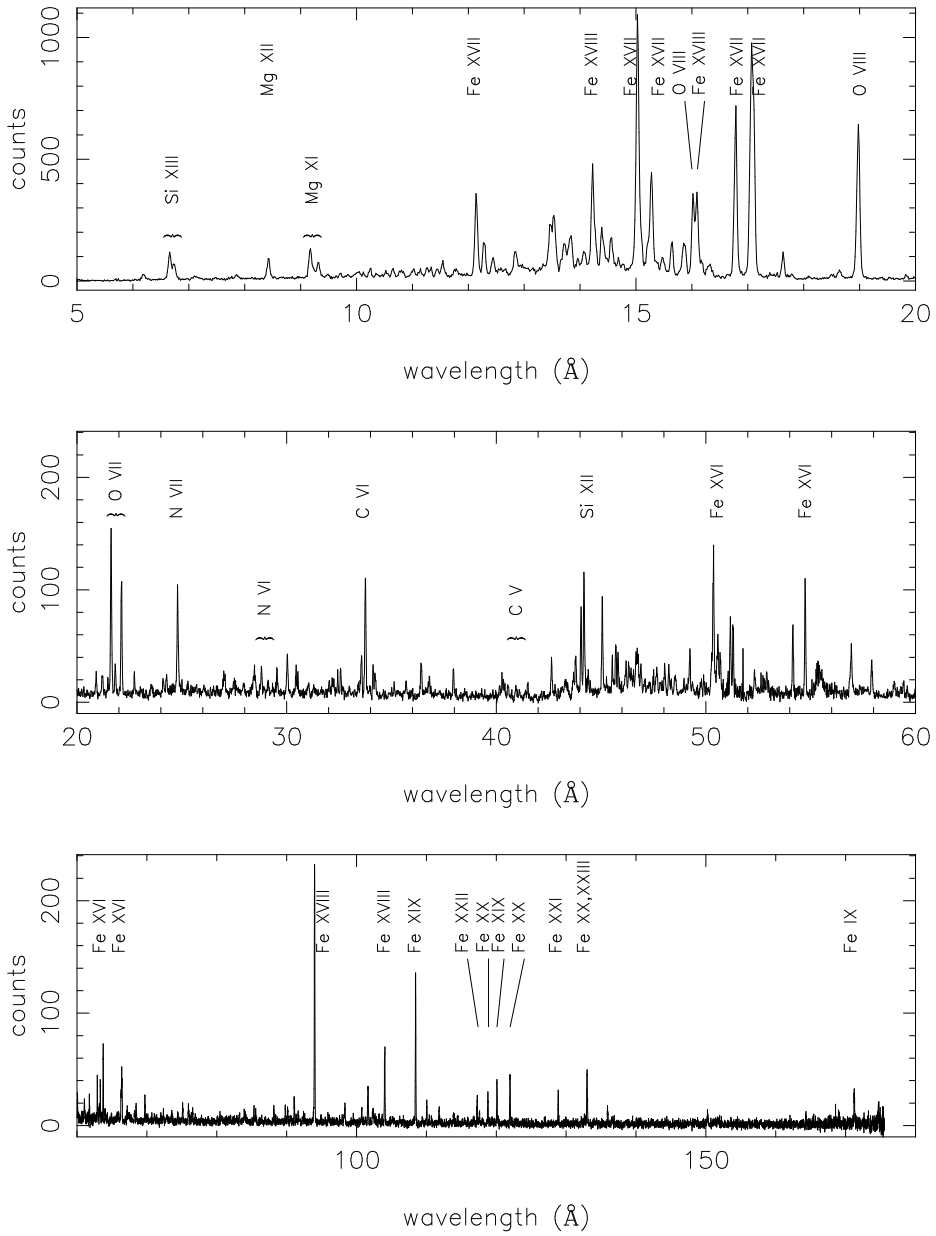

Figure 9.26: Extracted LETGS spectrum of Capella with some line identifications (from Brinkman et al. 2000, ApJ, 530, L111). Many of the lines visible between 40 and 60 Å are 3rd order dispersion of the strong features seen in 1st order in the uppermost panel. Positive and negative orders have been summed.

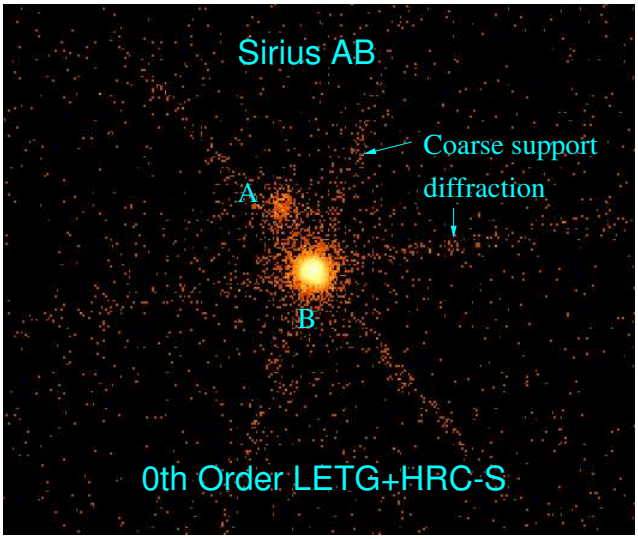

Figure 9.27: LETG/HRC-S zeroth order image of Sirius A and B. The two stars are separated by ∼4 arcsec. Flux from Sirius A is due to the small but finite UV response of the detector. The star-shaped structure is due to coarse support diffraction.

Figure 9.27 is a zeroth-order image of summed Sirius AB observations (ObsIDs 1421, 1452, 1459) with a total exposure time of 23 ks. The star-shaped pattern is due to coarse support structure diffraction. Sirius A and B are separated by ∼4 arcsec. The flux from Sirius A is due to the small, but non-zero, UV response of the detector.

9.4 Observation Planning

The purpose of this Section is to provide further information directly related to planning LETGS observations that is not explicitly presented in Sections 9.1 and 9.3, and to reiterate the most relevant issues of instrument performance that should be considered when preparing an observing proposal. Those proposing an LETG+ACIS-S observation should read the https://cxc.harvard.edu/cal/Letg/ACIS_params page. After an LETG proposal is accepted, CXC User Support scientists will work closely with observers to ensure an optimal instrument set-up.

Due to revised operational constraints enacted in 2023, HRC observations are currently limited to durations of no more than 14.5 ks, with separations of at least 30 ks. Minimum separations and maximum durations may be further constrained by external factors such as target pointing direction.

9.4.1 Detector Choices

The best choice of detector will depend on the exact application; some considerations are listed below. For further details concerning the HRC and ACIS detectors, refer to Chapters 7 and 6, respectively. Note that contamination build-up on the ACIS filters has significantly reduced the effective area of the LETG+ACIS-S combination for wavelengths > 20 Å compared to the effective area at launch. In this regard please be sure to read the discussion in Section 6.5.1.

HRC-S

- The HRC-S provides wavelength coverage from 1.2–175 Å (10–0.07 keV).
- The HRC-S QE is smaller than that of ACIS-S in the 1.2 to ∼10 Å (1.2–10 keV) range, but it is larger at longer wavelengths (see Figures 9.6 and 9.17).
- The HRC-S provides the highest time resolution at 16 μs when telemetry saturation is avoided. The probability of avoiding saturation is significantly improved if only the central plate is utilized (see Section 7.12).
- The HRC-S suffers from small position non-linearities which may slightly distort or shift spectral features (see Section 9.3.2).
- HRC-S has essentially no intrinsic energy resolution and so overlapping spectral orders cannot be separated.
- The Level 1 HRC-S background count rate is typically 100 cts s^{-1} in its windowed-down spectroscopic configuration during quiescence (see Figure 9.21 for how this varies over the solar cycle). However, this rate can rise to exceed the HRC telemetry saturation limit of 183 cts s^{-1} during background "flares." During telemetry saturation, deadtime is determined to an accuracy of 5–10%. Background flares can also be filtered out using *CIAO* or other software tools. These flares have been seen to affect 10–20% of some observations. Typical fractions are smaller than this; larger fractions are rare. If observing a very bright source where telemetry saturation is a concern, one can use a smaller-than-standard region of the detector (see Section 9.4.2).
- The LESF filter region in principle can be used to obtain a higher-order spectrum relatively uncontaminated by 1st order for wavelengths above 75 Å ($E < 0.17$ keV; see Figure 9.17). This could be useful either for observing features in a high order for high spectral resolution that cannot be easily observed with the HETG/ACIS-S combination, or for providing a direct observation of higher order contamination in conjunction with an LETG+HRC-S observation in its nominal configuration. NB: the LESF configuration has never been used in flight.

Summary HRC-S is probably the best detector choice for spectroscopic observations in which one or more of the following observational goals apply:
—signal longward of 25 Å is of significant interest (see Figure 9.17);
—the highest time resolution is required.

HRC-I

- When used with the LETG, the HRC-I provides wavelength coverage from 1.2–73 Å (10–0.17 keV).

- The raw HRC-I quiescent background event rate per unit area is lower than that of the HRC-S by about a factor of 4. After moderate filtering in both detectors, the ratio is about a factor of 2. This may be important for very weak sources.

- The HRC-I imaging capabilities are similar to those of the HRC-S, but its single flat detector plate cannot follow the Rowland geometry as well as the HRC-S. At nominal focus this results in slightly poorer spectral resolution for wavelengths > 50 Å. In principle, small focus offsets can be used to optimize the focus of the dispersed spectrum either within a specific wavelength range, or to average defocus blurring over the full wavelength range. No simple focus optimization prescription currently exists; interested readers should contact the CXC for assistance.

- The details of the LETG+HRC-I effective area have been less well-studied in general than for the LETG+HRC-S combination. The polyimide used in the HRC-I UV/Ion shield is about twice as thick as that used with the HRC-S, resulting in lower transmission at long wavelengths.

- The HRC-I offers a broad detector in the cross-dispersion direction, and this might be a consideration for observation of sources with extended components exceeding ∼2 arcmin. Note, however, that the *Chandra* spectrographs are slitless, and the effective spectral resolution is severely degraded for large sources (see Figure 9.11).

Summary HRC-I is possibly the best detector choice for sources in which signal longward of 73 Å (0.17 keV) is not of primary interest, accurate effective area knowledge for > 44 Å (< 0.28 keV) is not a strong concern, and one or more of the following observational goals apply:

—the source is very weak with interesting spectral features at wavelengths beyond the limit of the LETG+ACIS-S coverage;
—high resolution timing is *not* required;
—a larger detector area in the cross-dispersion direction than is provided by the HRC-S is required.

ACIS-S

Although the information in this section is still generally correct, in practice the LETG+ACIS-S effective area is currently too low for effective use, and no non-calibration LETG+ACIS-S observations have been made since 2015.

- Contamination build-up on the ACIS-S detector has significantly reduced the effective area of the LETG+ACIS-S combination at energies less than ∼1 keV since launch (see Section 6.5.1 for details). ACIS-S provides an effective LETG 1st order wavelength limit of about 60 Å (0.20 keV) because longward of this the ACIS-S QE is essentially zero (see Figures 9.6 and 9.17). ACIS-S is not as well-calibrated for wavelengths longward of the C edge (∼44 Å).

- The intrinsic energy resolution of ACIS-S allows for discrimination between different and otherwise overlapping spectral orders. For dispersion longward of $m\lambda \sim 60$ Å the LETG+ACIS-S response is dominated by higher order throughput (Figure 9.17) and ACIS-S can therefore be useful for observing these higher spectral orders.

- ACIS-S allows several modes of operation (see Section 6.13) including continuous clocking (CC) if high time resolution is desired or to avoid pile-up.

- If using the full frame 3.2 s exposure of ACIS-S in TE mode, photon pile-up can be a serious consideration, especially in zeroth order. Proposers should also be aware that there is a potential for pile-up in bright lines and continua and not assume that, because of dispersion, the flux is too spread out to be affected. For observations using Y-offset= +1.1 arcmin, the fraction of 1st order events lost to pile-up in a continuum spectra can be estimated as event rate (in cts/frame/dispersion-axis-pixel) times 3 for FI chips; times 4 for BI chips. As an example, if the count rate of the 1st order spectrum is expected to be 0.01 cts/frame in a wavelength interval of 0.0275 Å (one pixel wide along the dispersion axis), about 3% of those events will be lost to pile-up (in FI chips). While the example rate is observed only in very bright continuum spectra (such as Mkn 421 near 6.5 Å), pile-up can be a concern for bright features in line-dominated spectra. Some of these events can be recovered by examining higher order spectra, but most pile-up events will have "migrated" out of the standard grade set. Pile-up can affect both the shape of the PSF and the apparent spectral energy distribution of your source. Pile-up may be reduced by opting for a "subarray" that reads out a smaller area of the detector for a decrease in the frame time. See Section 9.4.2 and also Section 6.16 for details concerning pile-up, its effects, and how best to avoid it.

- ACIS-S time resolution is lower than that of HRC and depends on the control mode adopted. In timed exposure (TE) mode the full frame exposure is 3.2 s. This is reduced when using a subarray due to the shorter readout time for the smaller detector region (see Section 6.13 for details). Using fewer chips (e.g., dropping S0 and S5 because of their negligible effective area for 1st-order photons) may also

slightly reduce the frame time. The highest time resolution possible with ACIS-S (2.85 ms) is obtained in continuous clocking (CC) mode, but imaging information in the cross-dispersion direction is lost and the background will be higher due to the implicit integration over the entire cross-dispersion column of the detector.

- The energy resolution of the FI chips degrades as distance increases from the CCD readout because of CTI (see Section 6.7). The LETG dispersion axis is parallel to the ACIS-S readout and the spectrum of a point source can be placed close to the readout such that the energy resolution degradation is no longer a significant problem; a default SIM-Z offset of −8 mm is routinely applied to LETG+ACIS-S *point source* observations. If observations with extended sources are under consideration, or if for other reasons a SIM offset is undesirable, the resolution in the FI CCDs of the ACIS-S array might be a point to consider. From an LETG perspective, the effects of concern are a degradation of the CCD energy resolution that is employed for order sorting, grade migration that can make for difficult calibration of detector quantum efficiency, and, at longer wavelengths ($\gtrsim$ 50 Å), a loss of events that have pulse heights below that of the ACIS event lower-level discriminator. These effects render the effective area at wavelengths longward of the C edge ($\sim$44 Å) less well-calibrated than at shorter wavelengths.

- The ACIS-S energy resolution enables removal of the vast majority of background events in LETG spectra; the effective ACIS-S background is consequently much lower than that of HRC-S or HRC-I.

Summary ACIS-S is possibly the best detector choice for sources for which signal longward of 25 Å (0.5 keV) is of little interest and one or more of the following observational goals apply:
—Particular spectral features of interest occur where the LETG+ACIS-S effective area is higher than that of LETG+HRC-S
—High time resolution beyond the 3.2 s exposure of TE mode (less if a subarray is used), or the 2.85 ms of CC mode (if applicable), is not important
—A low resolution zeroth order spectrum from the S3 BI chip is of high scientific value, in addition to the dispersed LETG spectrum
—Order separation is important
—Pile-up can either be avoided or mitigated or is not likely to be a problem.

9.4.2 Other Focal-Plane Detector Considerations

Instrument Features and Gaps

Attention should be paid to the locations of instrument edge features and detector gaps to make sure that spectral features required to achieve science goals are not compromised by these. Calibration in the vicinities of these features is generally much more uncertain. These features and gaps are listed for both HRMA+LETG+HRC-S and

HRMA+LETG+ACIS-S combinations in Tables 9.3 and 9.4. Note that intrinsic instrumental features, such as edges, are not affected by dithering and offset pointing (see below), but chip gaps in ACIS-S and HRC-S plate gaps as well as the boundaries between "thick" and "thin" regions of Al that make the "T" shape of the HRC-S UVIS, are affected.

Dither

The standard LETG+HRC-S dither amplitude is 20 arcsec (40 arcsec peak-to-peak; 2 mm in the focal plane or 2.3 Å). As of Cycle 24, LETG+ACIS-S dither is 16 arcsec (32 arcsec peak-to-peak; 0.16 mm or 1.8 Å) for both axes (or half that in the cross-dispersed direction when using a 128-row subarray). Spectral features in dispersed LETG spectra will experience the same dither pattern, and allowance for the size of the dither must be made when considering if spectral features of interest will encounter detector gaps.

In special cases, different dither amplitudes can be specified by the observer, though it must be kept in mind that detector safety constraints, such as accumulated dose per pore in the HRC (see Section 7.14), must not be violated.

SIM-Z Offsets

The SIM permits movement of the focal-plane detectors in the spacecraft z direction (perpendicular to the LETG dispersion axis). This can be used to better position a source on the ACIS-S or HRC detectors, for example to accommodate multiple sources, or to place a source over the HRC-S LESF filter region. The default SIM-Z value is 0, which places the spectrum along the center of the ACIS-S array where contaminant on the Optical Blocking Filter is thinnest and effective area is highest. Placing the LESF (which has never been used in flight) at the aimpoint requires a SIM-Z offset of +7 mm.

Offset Pointing

Because of the changing thermal environment of the telescope, the aimpoint has drifted over time (see Section 6.11). For several years, observers wishing to make observations at the best focus had to specify small offsets in Y and/or Z pointing. As of late 2015, however, this is no longer necessary as aimpoint position and error are reviewed every two weeks and automatically incorporated into mission planning software. The default Z-offset is therefore 0.

Pointing off-axis in the observatory Y axis can be used to change the wavelengths at which detector gaps occur or to change the wavelength corresponding to the ends of the detectors. Examples of offset pointings are shown in Chapter 3. When choosing offsets, an increase of +1 arcmin in Y-offset corresponds to a shift of +3.36 Å in wavelength. As an example, by invoking a +2 arcmin offset pointing (see Chapter 3 for the convention), the long-wavelength cut-off of the HRC-S can be extended in the + order from approximately 176 Å for on-axis pointing to 183 Å. This, of course, is obtained at the expense of a commensurate shortening of coverage in the − order.

Offset pointing leads to degradation of the PSF, and consequently the spectral resolution—see Figure 9.12. For offsets of 2 arcmin or less this degradation is very small. For offsets of > 4 arcmin, spatial and spectral resolution will be considerably degraded.

In the case of the LETG+ACIS-S configuration, certain offsets might be useful, e.g., to place features of interest on (or off) back-illuminated chips for better low-energy quantum efficiency. Table 9.3 and Figure 9.6 can be used to determine what offsets are required. **There is also an extremely useful visualization tool on the 'Checking Your LETG/ACIS-S Obscat Setup' page (https://cxc.harvard.edu/cal/Letg/ACIS_params), which incorporates the latest detector aimpoint calibration (see Section 6.11).** Keeping in mind that exact values may change over time (and will be adjusted as necessary before any observation occurs), three particularly important Y-offsets for LETG/ACIS-S (subsequent to the change in default dither and aimpoint in Cycle 24) are the following:

+0.0 **arcmin:** This is the default Y-offset value and provides the highest possible resolution while avoiding dither across the node0/node1 boundary on the S3 chip.

+1.07 **arcmin:** This is the most commonly recommended offset, as it keeps O K-edge features on the S3 back-illuminated chip. Spectral resolution is ∼15% worse than with Y-offset=0.

+1.67 **arcmin:** This approximately centers zeroth order in the gap between S2 and S3 (but keep in mind that pointing errors can be up to ∼12 arcsec, with errors of 5 arcsec or less about 90% of the time); some zeroth order events will fall on each chip because of dither. This offset can be used to reduce CCD exposure and telemetry demands caused by very bright sources and also provides uninterrupted spectral coverage (apart from chip-edge dither effects around 29 Å) by the BI chips (S1 and S3) from 0 to 57 Å. Spectral resolution is degraded by ∼25% (see Figure 9.12).

ACIS-S Modes

When using the LETG+ACIS-S configuration, a mode must be selected for the ACIS detector. The ACIS detector is very flexible, but deciding the best set-up can be complicated. Prospective observers considering using ACIS-S for the focal-plane detector are urged to read Chapter 6 carefully. The most common modes used for LETG+ACIS-S observations are those using sub-arrays. The shorter frame times of subarrays can be a good way to both increase the time resolution and decrease pile-up. Care must be taken when defining subarrays to make sure that the choice of SIM-Z plus any offset pointing in the z direction places the source comfortably inside the subarray. Modes with 256 rows ($\frac{1}{4}$ subarray) or fewer are recommended for observations of point (or small) sources, provided that the source position is known to within a few arcsec. Subarrays permit the use of shorter frame times, and thus less pile-up. As an example, a $\frac{1}{8}$ subarray (using 6 chips and with the standard SIM-Z offset of -8 mm) requires a frame time of 0.7 s (vs 3.2 s for the full array).

Frame times as short as 0.6 s can be used with four ACIS-S chips, or 0.5 s with three chips. See Section 6.13.1 for further information on ACIS frame times. Users should always refer to "Checking your LETG/ACIS-S Obscat Setup" (https://cxc.harvard.edu/cal/Letg/ACIS_params) for the most current information regarding subarray setups.

Optional ACIS-S Chips and Scheduling Flexibility

As described in more detail in Section 6.20, the continuing gradual degradation of the *Chandra* thermal environment means that a steadily increasing fraction of observations are experiencing scheduling constraints with regard to the allowable exposure time per orbit. One way to increase the scheduling flexibility of LETG+ACIS-S observations is to use fewer than the full array of 6 CCDs, which reduces the temperature of the ACIS electronics and focal plane. To achieve this, users may denote some ACIS chips as "Optional." In cases where thermal limits might otherwise be breached, Optional chips will be turned off for an observation; this allows Mission Planning more flexibility when setting up observing schedules and makes the process more efficient. Problematic observations with Optional chips are thus less likely to be split into multiple pieces than they would be if all the chips are required.

Note that chips S0, S4, and S5 have negligible QE for LETG 1st order dispersed photons (see Figure 9.6); those outer chips are only useful for collecting higher-order spectra, or perhaps in other special circumstances. The web page "Checking your LETG/ACIS-S Obscat Setup" (https://cxc.harvard.edu/cal/Letg/ACIS_params) and the Spectrum Visualizer tool linked there are helpful in determining which CCDs may be unnecessary. **The S0 and S5 chips will be turned off in all LETG+ACIS-S observations** unless the observer explicitly requests and justifies their use. It is also recommended that S4 be turned off.

HRC-S Windowing

As described in Section 9.3.6, the HRC-S has a default spectroscopic "window" defined that limits the detector area from which events are telemetered to the ground. The window is a rectangle based on coarse position tap boundaries; the default rectangle comprises the central 6 taps in the cross-dispersion direction (corresponding to ∼9.9 mm) and the whole detector length in the dispersion direction. This window can be defined to suit special observational goals, such as if the source is extended and the width of the readout region must be increased (with increased risk of telemetry saturation during background flares).

Likewise, a smaller detector region can be used if, for example, the source is very bright and telemetry saturation is a concern. In such a case, the number of cross-dispersion taps could be reduced, since the source only dithers across 2 or 3 taps. There would be less area for the background extraction region, but with a bright source this would not be a problem.

Observers can also specify that a subset of taps along the dispersion direction be used, particularly if the source spectrum is expected to be cut off at long wavelengths and higher-order spectra are not needed. The most commonly used (though still rare) non-default configuration is simply to use only the central plate with the standard 6 cross-dispersion taps, usually referred to as the Imaging Mode.

When considering defining a special HRC-S window, it is reasonable to assume that the detector background is spatially uniform for the purposes of computing the total source plus background count rate. The telemetry capacity of 183 cts s^{-1} should be kept in mind to avoid telemetry saturation by using a window that is too big.

9.4.3 General Considerations

PSF Artifact

Recently, a minor PSF anomaly has been identified (see Section 4.2.3). It contains only a few percent of the source flux at a separation of 0.8 arcsec from the main peak and is roughly aligned with the cross-dispersion axis. This anomaly has negligible impact on LETG observations.

Complications from Other Sources

Field sources coincident with the target source dispersed spectrum should be avoided. While some flexibility in roll angle of order several arcmin can often be accommodated in observation execution, this avoidance is most rigorously accomplished by imposing a roll angle constraint (see below). Note that it is also desirable to retain a pristine region on either side of the dispersed spectrum to enable an accurate estimation of the background within the spectral extraction window.

In some circumstances, photons from bright sources outside of the direct field of view of the HRC or ACIS might be dispersed by the LETG onto the detector.

Particular attention should be paid to optically-bright and UV-bright sources, even if these are some distance off-axis. The ACIS-S and HRC-S filters are much more transparent to optical and UV light than are those of HRC-I and ACIS-I (the HRC-S central "T" segment is closer in performance to that of HRC-I, but has completely different thicknesses of polyimide and Al layers). As an example, an observation of the bright A0 V star Vega ($V = 0.03$) in one of the outer HRC-S UVIS segments gave a count rate of about 475 cts s^{-1}.

The energy resolution of the ACIS-S detector enables removal by filtering of all photons except those in a fairly narrow wavelength or energy range corresponding to the wavelength or energy of photons in a spectrum dispersed by the LETG. This means that contamination of the dispersed spectrum by, for example, the zeroth order or dispersed spectrum of other sources *might* not be a significant problem.

However, a much better solution to problems of source contamination is, if it is possible within other observation constraints, to choose a roll angle (Chapter 3) that avoids the source contamination issue.

Roll Angle Considerations

Roll angle constraints can be specified to avoid contamination by off-axis sources, as described above, or to help separate the dispersed spectra of multiple sources in the cross-dispersion direction. The maximum separation between dispersed LETG spectra of two sources is obviously one that places the sources in a line perpendicular to the dispersion axis.

Owing to spacecraft thermal degradation and increasing difficulty in scheduling constrained observations, proposers are reminded that only a very limited number of constrained observations can be accommodated. Technical justification for requested constraints must be provided. It is also important to remember that roll angle constraints will also impose restrictions on the dates of target availability as discussed in Chapter 3. Exact restrictions depend on celestial position. Their impact can be examined using the observation visualizer tool, downloadable from the *CIAO* home page at https://cxc.harvard.edu/obsvis.

High Order Throughput

It is expected that the majority of observations with the LETG will make use of the HRC-S as the readout detector because of its wavelength coverage and high quantum efficiency at long wavelengths. Since the HRC-S has very little energy resolution, the overlapping of spectral orders could be a significant issue and prospective observers should assess the degree to which their observation might be affected. The following list summarizes some considerations:

Scientific Utility: Higher spectral orders provide higher spectral resolving power than the 1st order spectrum by the approximate factor of the order number m. For observations in which features are expected to be seen in higher orders, this capability for higher spectral resolution could be scientifically useful.

LESF: The LESF (the region of thicker Al coating on the HRC-S UVIS) is untested in flight, but it could be useful for obtaining a spectrum containing mostly higher order flux.

Source Spectrum: For some sources higher orders will contain very little flux and thus will not be an issue. Typical examples are hot white dwarfs or relatively cool stellar coronae with $T \sim 10^6$ K. Sources whose spectra are fairly weak in the region where the effective area of the LETG+HRC-S is highest (∼8–20 Å; 1.5–0.6 keV) but gain in strength toward longer wavelengths will also be less affected by higher order throughput. Typical

examples are blackbody-type spectra with temperatures $T \sim 10^6$ K or less, such as might characterize novae or isolated neutron stars.

Estimates: Figure 9.18 can be used to estimate high order contamination. *PIMMS* can be used for gross estimates of higher order count rates; the *PIMMS* higher order calculation uses an effective area curve for orders $m > 1$ combined.

Instrumental Capabilities: Order separation is straightforward with ACIS-S. With HRC-S orders cannot be separated.

Spectral Modeling and Higher Orders: Unlike ACIS-S, the HRC-S does not have enough energy resolution to allow order sorting of LETG spectra. However, by folding a spectral model through an LETG+HRC-S instrument response that includes all significant higher orders (generally ≤ 10), the whole spectrum can be modeled at once. The capability to generate and simultaneously utilize response matrices for multiple orders (which can be created for plus and minus orders 1–25) is available within *Sherpa*/*CIAO*, and these response matrices can also be used with other spectral analysis software such as *XSPEC*. Note that the response matrices do not include the small-scale wavelength distortions discussed in Section 9.3.2, and hence care must be taken when analyzing some details of line-dominated spectra, particularly line profiles.

Interstellar Medium (ISM) Absorption

The long wavelength cutoff of the LETGS in tandem with the HRC-S detector is ∼175 Å (although this can be extended with offset pointing as described in Section 9.4.2), which reaches well into the extreme ultraviolet (EUV). In this wavelength regime, the spectra of even very nearby sources with relatively low ISM absorbing columns can be appreciably attenuated by H and He bound-free photoionizing transitions. Therefore observers should be aware that the effective long wavelength cutoff for anything but the nearest sources (∼100 parsec or less) will probably be determined by ISM absorption. It is also important to remember that neutral and once-ionized He can dominate the ISM absorption cross-section in the 44–200 Å (0.28–0.062 keV) range, and consideration of the neutral H absorption alone is generally not sufficient. Shortward of the C edge near 44 Å (0.28 keV), metals become the dominant absorbers. For illustration, the ISM transmittance for a "typical" mixture of neutral and ionized H and He with H:He:He^+ ratios of 1:0.1:0.01 is shown in Figure 9.28 for the 5–200 Å range with different values of the neutral hydrogen column.

The *CXC* web page has a tool *Colden* (https://cxc.harvard.edu/toolkit/colden.jsp) that provides the total *Galactic* neutral hydrogen column for a given line-of-sight. The Ahelp page for *Colden* is located at https://cxc.harvard.edu/ciao/ahelp/colden.html. An IDL routine from the PINTofALE data analysis package for computing the ISM optical depth is available from https://hea-www.harvard.edu/PINTofALE/pro/ismtau.pro.

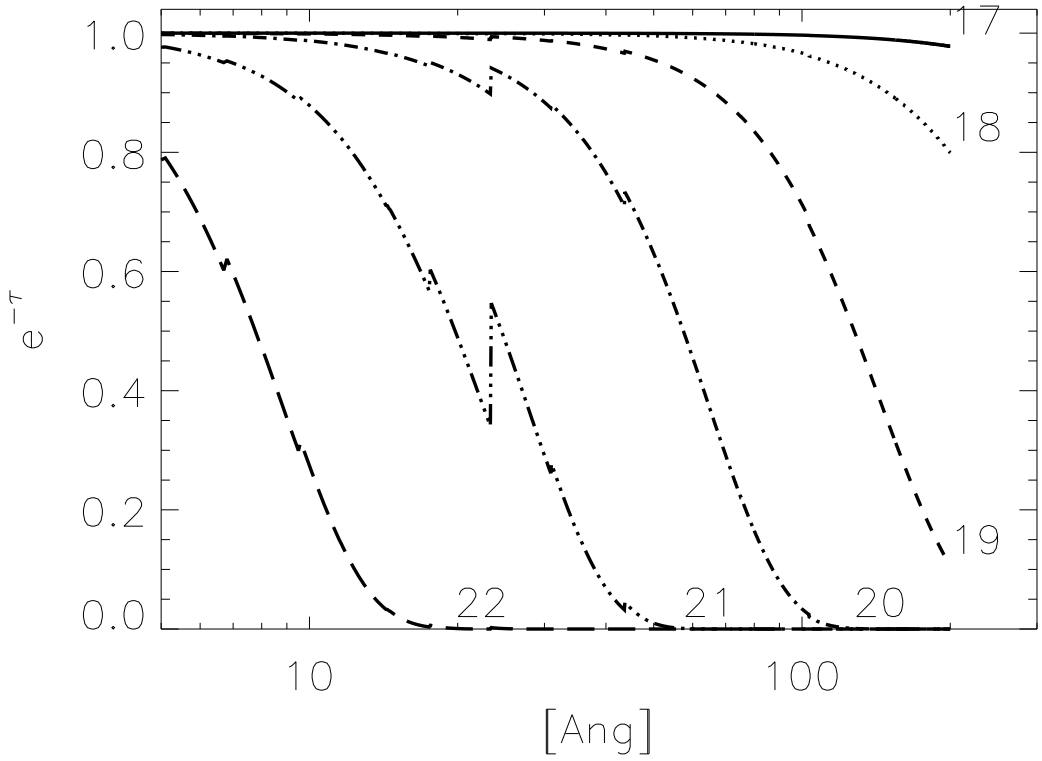

Figure 9.28: The ISM transmittance within the LETGS bandpass for different values of neutral hydrogen column density 10^{17}–10^{22} cm^{-2}.

9.5 Technical Feasibility

Proposers should always be aware of possible limitations in the physical models and methods they are using for observation planning purposes. For example, older *XSPEC* versions might not include ISM absorption edges or spectral models at the high resolutions appropriate for *Chandra* grating observations. Available optically thin, collision-dominated plasma radiative loss models have also generally only been tested in any detail for strong lines of abundant elements. Some prominent transitions in Fe ions with n=2 ("L-shell") and n=3 ("M-shell") ground states are not yet well-represented by some models. The spectral region 25–75 Å (0.5–0.17 keV) still remains largely unexplored, and both total radiative loss and predicted strengths of lines in this region are more uncertain than at shorter wavelengths.

Some additional technical limitations in *MARX* modeling of LETG+HRC-S spectra are detailed below.

9.5.1 Simple Calculation of Exposure Times and Signal-to-Noise Ratio for Line and Continuum Sources

The detection of an isolated emission or absorption line was previously discussed in Section 8.5.5 in the context of the HETG, but it is also relevant to the LETG+HRC-S combination. This discussion is based on line equivalent width, which is appropriate for broadened lines and continuum features but which is more difficult to apply to simple modeled estimates of expected line fluxes. Additional formulae that are simple to apply are presented below. The units are Å rather than keV, as Å is a much more natural unit of choice for dispersed spectra, particularly for the LETG range.

Emission Line Sources The source signal S in a bin is the difference between the total counts and the background counts B. The estimated standard deviation of the source counts S in a spectral bin is given by Poisson statistics as:

$$\sigma_S = \sqrt{S + B} \tag{9.3}$$

This estimate assumes that there is effectively no additional uncertainty in the estimation of the background B. Such an assumption may only be valid if, for example, the detector region used to estimate the background within the spectrum extraction window is much larger than the window itself.

Spectrometer count rates for emission features are given by

$$s_l = A_{eff}(\lambda)\mathcal{F}_l(\lambda) \tag{9.4}$$

where s_l is the source count rate in the resolution element centered at λ, in cts s^{-1}, A_{eff} is the effective area in cm^2, and $\mathcal{F}_l$ is the source flux at the telescope aperture, in photons cm^{-2} s^{-1}. For A_{eff}, it is reasonable to use the total area obtained from the sum of + and

– orders as illustrated in Figure 9.17. Raw source counts are estimated by multiplying this instrument count rate by an integration time.

Using Equations 9.3 and 9.4, the signal-to-noise ratio for an integration time t is then

$$\frac{S}{\sigma_S} = \frac{s_l\sqrt{t}}{\sqrt{s_l + b}} \tag{9.5}$$

where b is the background count rate *within the spectrum extraction window* (i.e. "underneath" the spectrum) in the same resolution element centered at λ, in cts s^{-1}. Equation 9.5 provides the expected relation that is valid in the limit where the background count rate b is small compared with the source count rate s_l and that the signal-to-noise ratio scales with the square root of the exposure time.

The exposure time required to achieve a given signal-to-noise ratio is then provided by inversion of equation 9.5,

$$t = \left(\frac{S}{\sigma_S}\right)^2 \frac{s_l + b}{s_l^2}. \tag{9.6}$$

The exposure time estimate thus depends on the background count rate, b. Since the spectrometer does not have infinite resolution, the flux from an otherwise narrow spectral line is spread over a typical line width, w_l. The background rate b is therefore given by the quantity $b = w_l b'$, where b' is the background rate in units of cts $Å^{-1}$ s^{-1}. For LETG+HRC-S spectra, a good estimate for w_l is 0.07 Å; this value is somewhat larger than the FWHM value of 0.05 Å listed in Table 9.1, but it is more appropriate for calculations of signal-to-noise because it includes more of the line flux. For lines that are additionally broadened, w_l should be increased to cover the region under the feature of interest. Background spectra for LETG+HRC-S are illustrated in Figure 9.22, which can be used to estimate b or b'. Two scales are shown in that Figure, one corresponding to b' and one corresponding to b where a width $w_l = 0.07$ Å was assumed; in both cases, displayed rates are per 10 ks.

Using the signal count rate s_l, provided by the product of source flux (at the telescope aperture) and effective area as stated in equation 9.4, results in two equations: for the signal-to-noise ratio S/σ_S resulting from an exposure time t,

$$\frac{S}{\sigma_S} = \frac{\mathcal{F}_l A_{eff} t}{\sqrt{\mathcal{F}_l A_{eff} t + b' w_l t}} \tag{9.7}$$

and for the exposure time t required for a signal-to-noise ratio S/σ_S

$$t = \left(\frac{S}{\sigma_S}\right)^2 \frac{\mathcal{F}_l A_{eff} + b' w_l}{(\mathcal{F}_l A_{eff})^2}. \tag{9.8}$$

These simple equations, which include the effects of instrumental background, can also be easily applied to observations of lines on top of continua, as well as to situations in which features of interest lie on top of higher (or lower) spectral orders (HRC). In these cases, the continuum or higher order flux acts as an additional background term—the count rate in cts $Å^{-1}$ due to these additional terms is simply added to b'.

Continuum Sources Model fluxes for continuum sources can be expressed as flux densities in units of photons cm^{-2} s^{-1} Å^{-1} . To compute instrument count rates s_c from a continuum source spectrum, the A_{eff} function and spectrum must be partitioned with some bin size w, large enough to give adequate count rates. The product of the source spectrum with the A_{eff} function is then summed over some wavelength region of interest. Equation 9.4 becomes the sum

$$s_c(\lambda) = \sum_{j=1}^{N} \mathcal{F}_c(\lambda_j)\, A_{eff}(\lambda_j)\, w \tag{9.9}$$

where $\mathcal{F}_c$ is the model source flux in photons cm^{-2} s^{-1} Å^{-1}, $A_{eff}(\lambda_j)$ is the effective area of the jth bin in cm^2. The region of interest spans bins 1 through N, and w is the bin width in Å. In using this formula for planning purposes, proposers must choose a spectral bin width that will demonstrate the viability of the program proposed. For fairly narrow spectral ranges in which A_{eff} is nearly constant, the sum over 1–N reduces to

$$s_c(\lambda) = \mathcal{F}_c(\lambda_j)\, A_{eff}(\lambda_j)\, Nw \tag{9.10}$$

In this case one can of course simply chose a new bin size $w' = Nw$.

The difference between the continuum and the emission line case above lies in the units of $\mathcal{F}_c$, which is a flux density. The equations corresponding to the line source Equations 9.7 and 9.8 are, for the signal-to-noise resulting from an exposure time t

$$\frac{S}{\sigma_S} = \frac{\mathcal{F}_c A_{eff} wt}{\sqrt{\mathcal{F}_c A_{eff} wt + b'wt}} \tag{9.11}$$

and for the exposure time t required for a signal-to-noise ratio S/σ_S

$$t = \left(\frac{S}{\sigma_S}\right)^2 \frac{\mathcal{F}_c A_{eff} w + b'w}{(\mathcal{F}_c A_{eff} w)^2}. \tag{9.12}$$

Note also in the above equations that the background b' is in units of cts Å^{-1} s^{-1}.

PIMMS for Rough Planning Purposes

PIMMS is best suited to performing rough estimates of total or zeroth order count rates, or estimating the fraction of zeroth order events that would be piled up. Some degree of caution should accompany *PIMMS* calculations of detailed quantities such as count rates within narrow spectral bands using the Raymond-Smith model. For example, line positions and intensities in this model were only designed to represent total radiative loss and do not stand up to high resolution scrutiny. Calculations using power-law and featureless continua are not prone to such difficulties, but are susceptible to other *PIMMS* limitations. One particular limitation concerns the background model for HRC-S, which in *PIMMS* is assumed to be a single average number per spectral resolution element of 20 cts/100 ks. This approximation overestimates the background at medium energies, and underestimates the background at lower and higher energies; see Figure 9.22.

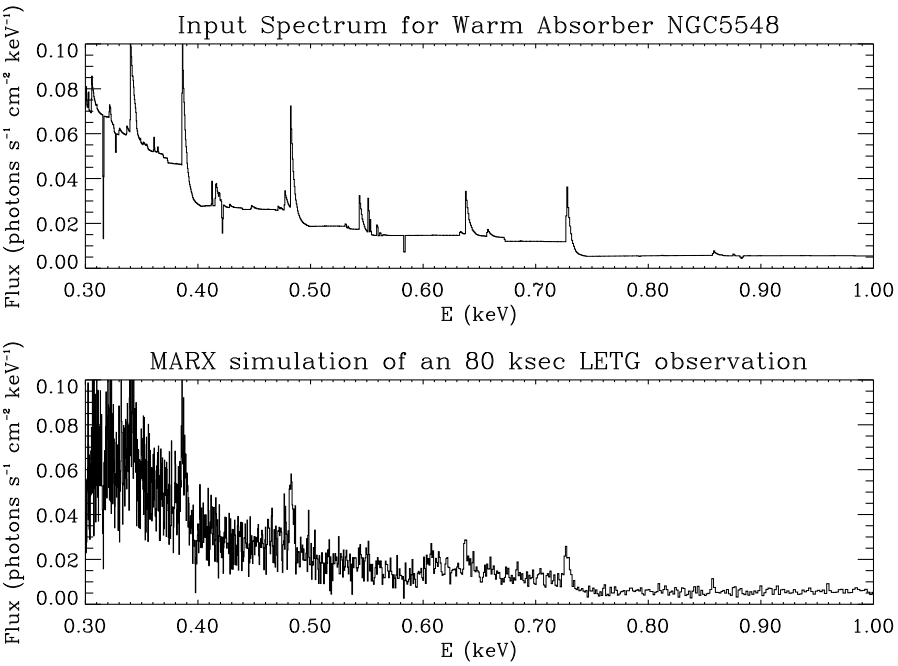

Figure 9.29: This figure shows the extracted 1st order spectrum for an 80 ks observation of the AGN NGC5548. The input spectrum consists of a power-law plus a "warm" absorber (shown in the top panel). The simulated spectrum (bottom panel) has been corrected for the instrument response to give the flux from the source.

MARX Simulations

The best tool for simulating LETG observations, including cases where there are multiple targets in the field of view, is the *MARX* ray trace simulator, which permits study of any of the available *Chandra* instrument combinations. For cases requiring higher fidelity modeling of the PSF, simulations can be performed by first simulating the response of the HRMA using ChaRT (https://cxc.harvard.edu/chart) and then feeding those ray trace results into *MARX* to simulate the LETG and detector responses.

MARX has an important limitation for LETG+HRC-S observations in that instrument, sky, and particle backgrounds are almost always significant (see Section 9.3.6 and Figure 9.22) but are *not* directly included and need to be simulated or otherwise accounted for by the user. Background can be simulated by approximating it as a flat field and adding this simulation to that of the source (see https://space.mit.edu/CXC/MARX/). Another way of simulating background is to scale the background in Figure 9.22 after adjusting for solar cycle variations in the background rate.

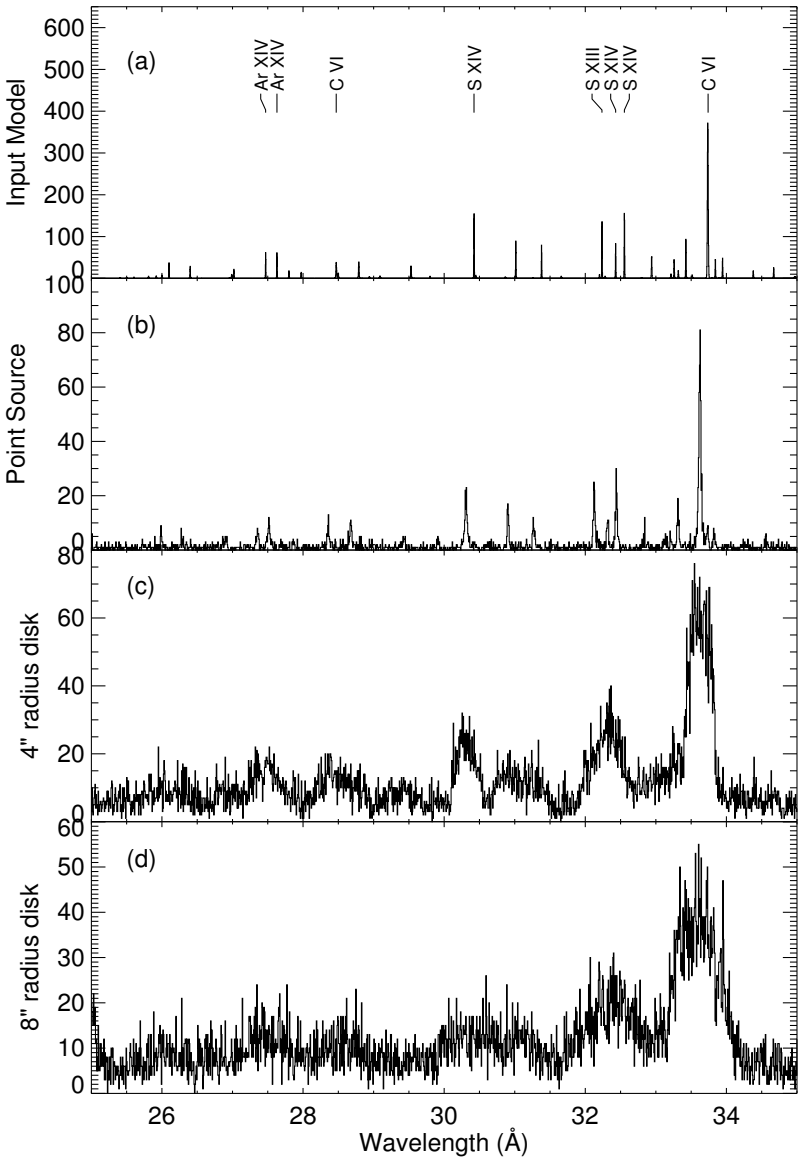

Figure 9.30: *MARX* simulation of spectra showing the effect of source extent. The panels show (a) computed input spectrum, (b) a *MARX* output of LETG spectrum of a point source, (c) the same as (b) except that the source is a disc of uniform brightness with radius of 4 arcsec, and (d) the same but with radius of 8 arcsec. See Figure 9.11 and Section 9.3.3 for a discussion of extended sources.

9.6 References

Further information on LETGS performance and calibration can be found on https://cxc.harvard.edu/cal/letg/ and https://cxc.harvard.edu/cal/letg/detailed_info.html, and in yearly *Chandra* Newsletter articles at https://cxc.harvard.edu/newsletters/.

Part I

Appendices

Appendix A

Contact Information

A.1 Contact Information

The Proposal Review is organized for NASA by the *CXC* Director's Office (*CDO*) under the direction of Dr. Rodolfo Montez Jr. Questions should be submitted to the *CDO* via our HelpDesk (http://cxc.harvard.edu/helpdesk/) (preferable) or by email: (cxchelp@cfa.harvard.edu). We strongly recommend the use of this HelpDesk as opposed to contacting individual staff members, as it is routinely monitored and not affected by the schedules of individuals.

A.2 *CDO* Staff

- Dr. Pat Slane, Director
- Dr. Thomas Connor
- Dr. Antonella Fruscione
- Dr. Amruta Jaodand
- Dr. Rodolfo Montez Jr.
- Ms. Tara Gokas
- Mr. Evan Tingle

Appendix B

Acronym List

This list is a superset of acronyms and abbreviations used in this document.

2MASS 2 Micron All-Sky Survey

ACA Aspect Camera Assembly

ACIS Advanced CCD Imaging Spectrometer

ACIS-I ACIS Imaging array

ACIS-S ACIS Spectroscopic array

ADU Analog-to-Digital Unit

AGASC AXAF Guide and Aspect Star Catalog

AGN Active Galactic Nucleus or Nuclei

ARF Ancillary Response File

ASC AXAF Science Center

ASCA Advanced Satellite for Cosmology and Astrophysics

AXAF Advanced X-Ray Astrophysics Facility

BI Back-illuminated

CALDB Calibration Data Base

CAP Central Aperture Plate

CC Continuous Clocking

CCD Charge Coupled Device

CCDM Communication, Command, and Data Management

CCT Chandra Cool Targets

CDO *CXC* Director's Office

CfP Call for Proposals

CGCD Crossed Grid Charge Detector

CGRO Compton Gamma-Ray Observatory

CHaRT *Chandra* Ray Tracer

CIAO Chandra Interactive Analysis of Observations

COLDEN Calculate Neutral Hydrogen Column Density

COUP Chandra Orion Ultradeep Project

CPS Chandra Proposal System or Chandra Proposal Software

CSS Coarse Sun Sensor

CTI Charge Transfer Inefficiency

CXC Chandra X-ray Center

CXCDS Chandra X-ray Center Data Systems

CXO Chandra X-ray Observatory

DEA Detector Electronics Assembly

DPA Digital Processing Assembly

DSN Deep Space Network

EA Effective Area

ECS External Calibration Source

EDT Eastern Daylight Time

EPHIN Electron Proton Helium Instrument

ESA Earth Sensor Assembly

ESC Engineering Support Center

EUV Extreme Ultraviolet

EXOSAT European X-ray Observatory Satellite

FI Front-Illuminated

FITS Flexible Image Transport System

FLA Fiducial Light Assemblies

FOT Flight Operations Team

FOV Field of View

FP Focal Plane

FPC Flow Proportional Counter

FSS Fine Sun Sensor

FTS Fiducial Transfer System

FUV Far-UV

FWHM Full Width at Half Maximum

GESS Grating Element Support Structure

GO General Observer

GSFC Goddard Space Flight Center

GTO Guaranteed Time Observer

HEASARC High Energy Astrophysics Science Archive Research Center

HEG High Energy Grating

HESS HETG Support Structure

HETG High Energy Transmission Grating

HETGS HETG Spectrometer

HPD Half Power Diameter

HRC High Resolution Camera

HRC-I HRC Imaging Detector

HRC-S HRC Spectroscopic Detector

HRI High Resolution Imager

HRMA High Resolution Mirror Assembly

HXDS HRMA X-ray Detection System

IDL Interactive Data Language

IPI Instrument Principal Investigator

IPS Integral Propulsion System

IRU Inertial Reference Unit

ISIS Interactive Spectral Interpretation System

ISM Interstellar Medium

IUS Inertial Upper Stage

JPL Jet Propulsion Laboratory

LESF Low-energy Suppression Filter

LETG Low Energy Transmission Grating

LETGS LETG Spectrometer

LMXB Low-Mass X-ray Binary

LRF Line Response Function

LSF Line Spread Function

LTS Long Term Schedule

MARX Model of AXAF Response to X-rays

MCP Micro Channel Plate

MEG Medium Energy Grating

MIT Massachusetts Institute of Technology

MPE Max-Planck-Institut für Extraterrestrische Physik

MSC Mirror Shell Coordinates

MSFC Marshall Space Flight Center

MTA Monitoring and Trends Analysis

MUPS Momentum Unloading Propulsion System

NASA National Aeronautics and Space Administration

NASCOM NASA Communication system

NGAS Northrop Grumman Aerospace Systems

NMM Normal Maneuver Mode

NPM Normal Pointing Mode

OAC Orbital Activation and Checkout

OBA Optical Bench Assembly

OBC On-Board Computer

OBF Optical Blocking Filter

OBSCAT Observation Catalog

OBSVIS Observation Visualizer

OCC Operations Control Center

ODB Optical Data Base

OFLS Off-line System

ONLS On-line System

OR Observation Request

OSIP Order Sorting Integrated Probability

PCA Principal Component Analysis

PCAD Pointing Control and Aspect Determination

PDA Permanent Default Aimpoint

PDF Portable Document Format

PHA Pulse Height Amplifier (or Amplitude)

PI Principal Investigator

PIMMS Portable Interactive Multi-Mission Software

PRoVis Pitch, Roll and Visibility tool

PSF Point Spread Function

PSMC Power Supply and Mechanism Controller (for ACIS)

QE Quantum Efficiency

QEU Quantum Efficiency Uniformity

RCS Reaction Control System

RCTU Remote Command and Telemetry Unit

RMF Redistribution Matrix File

RMS Root Mean Squared

ROSAT Roentgen Satellite

RRC Retroreflector Collimator

RWA Reaction Wheel Assembly

SADA Solar Array Drive Assembly

SAO Smithsonian Astrophysical Observatory

SI Science Instrument

SIM Science Instrument Module

SKIRU Space-Qualified Kearfott Inertial Reference Unit

SMF Software Maintenance Facility

S/N Signal-to-Noise ratio

SNR Supernova Remnant

SOHO Solar and Heliospheric Observatory

SPIE Society of Photo-Optical Instrumentation Engineers

SRON Laboratory for Space Research

SSD Solid-State Detector

STS Short Term Schedule

TE Timed Exposure

TOO Target of Opportunity

UTC Coordinated Universal Time

UV Ultraviolet

UVIS UV/Ion Shield

VCDU Virtual Channel Data Unit

VETA Validation Engineering Test Article (mirror)

VF Very Faint

XMM X-ray Multi-Mirror Mission

XRCF X-ray Calibration Facility

XSPEC X-ray Spectral fitting package