

The ACS Data Handbook for Cycle 29

F814W *I*
F625W *r*
F475W *B*



STScI | SPACE TELESCOPE
SCIENCE INSTITUTE

3700 San Martin Drive
Baltimore, Maryland 21218
<http://hsthelp.stsci.edu/>

ACS Data Handbook	2
Preface	4
Handbook Structure	5
Typographic Conventions	6
Acknowledgments	7
Chapter 1: ACS Overview	8
1.1 Instrument Design and Capabilities	9
1.2 Basic Instrument Operations	14
Chapter 2: ACS Data Structure	16
2.1 Types of ACS Files	17
2.2 ACS File Structure	23
2.3 Data Storage Requirements	26
2.4 Headers and Keywords	28
Chapter 3: ACS Calibration Pipeline	45
3.1 Mikulski Archive for Space Telescopes - MAST	46
3.2 Pipeline Overview	47
3.3 Structure of calacs	51
3.4 calacs Processing Steps	56
3.5 Manual Recalibration of ACS Data	72
Chapter 4: ACS Data Processing Considerations	84
4.1 Read Noise and A-to-D Conversion	85
4.2 Bias Issues	88
4.3 Dark Current, Hot Pixels, and Cosmic Rays	96
4.4 Flat-Field Reference Files	110
4.5 Image Anomalies	120
4.6 WFC CCD Detector Charge Transfer Efficiency - CTE	128
4.7 Generic Detector and Camera Properties	132
Chapter 5: ACS Data Analysis	144
5.1 Photometry	145
5.2 Astrometry	166
5.3 Polarimetry	169
5.4 Coronagraphy	176
5.5 Ramp Filters	178
5.6 Spectroscopy with the ACS Grisms and Prisms	180

ACS Data Handbook

Version 10.0- February 2021

[PDF version](#)

Advanced Camera for Surveys Instrument Data Handbook for Cycle 29

User Support

Please contact the HST Help Desk for assistance. We encourage users to access the new web portal where you can submit your questions directly to the appropriate team of experts.

- Web: <http://hsthhelp.stsci.edu>
- E-mail: help@stsci.edu

Additional Resources

Information and other resources are available from the STScI website:

- <http://www.stsci.edu/hst/instrumentation/acs>

Revision History

Version	Date	Editor
10.0	February 2021	Ray A. Lucas, et al. Susan Rose, Sr. Technical Editor
9.0	February 2018	Ray A. Lucas, Tyler Desjardins, et al., Sharon Toolan and Susan Rose, Sr. Technical Editors
8.0	April 2016	Ray A. Lucas, et al.
7.2	June 2014	S. Gonzaga, et al.
7.1	October 2013	S. Gonzaga, et al.
7.0	August 2013	S. Gonzaga, et al.
6.0	March 2011	S. Gonzaga, et al.
5.0	March 2006	Cheryl Pavlovsky, Anton Koekemoer & Jennifer Mack
4.0	January 2005	Cheryl Pavlovsky et al.
3.0	July 2004	Cheryl Pavlovsky et al.
2.0	November 2003	Jennifer Mack et al.

1.0	January 2002	Jennifer Mack et al.
-----	--------------	----------------------

Citation

In publications, refer to this document as:

Lucas, R. A., et al., 2021, ACS Data Handbook for Cycle 29, Version 10.0 (Baltimore: STScI).

Preface

How to Use this Handbook

This handbook was created to help users process and analyze data from the Advanced Camera for Surveys (ACS) which was installed on-board the *Hubble Space Telescope* (HST) during the 2002 servicing mission (SM3B). It is presented as an independent and self-contained document, designed for users familiar with *HST* data but new to ACS. Users who wish to find more general data analysis information, including instructions for retrieving data from the *HST* Archive, a description of *HST* file formats, and details regarding data analysis software supported by STScI, are referred to a companion volume, the [Introduction to HST Data Handbooks](#).

For information about geometric distortion and drizzling ACS data, please refer to the [DrizzlePac website](#).

Since many of the instrument characteristics may be revised over a short time frame, readers are advised to consult the ACS website (<http://www.stsci.edu/hst/instrumentation/acs>) for the latest information regarding ACS performance and calibration.



STSDAS and PyRAF are not supported by STScI. The last software to use STSDAS/PyRAF was aXe, a slitless spectroscopy software, and it is now replaced by hstaxe, a PyRAF independent follow-up to aXe, available at <https://github.com/spacetelescope/hstaxe/>. All other data processing described herein now uses Python available via the STScI [AstroConda channel website](#).

Handbook Structure

The ACS Data Handbook is organized in five chapters, which discuss the following topics:

- [Chapter 1 - ACS Overview](#) describes the ACS capabilities, design, and basic instrument operations; detailed information can be found in the [ACS Instrument Handbook](#).
- [Chapter 2 - ACS Data Structure](#) describes the ACS file structures, image header keywords, and data file sizes.
- [Chapter 3 - ACS Calibration Pipeline](#) describes the processing and flow of data from the STScI Archive. There are two major components:
 - calacs**, the standard ACS calibration software, has undergone significant updates that include corrections to post-SM4 WFC image artifacts. In addition, the software generates two sets of calibrated products: one with standard calibrations and another set with pixel-based CTE corrections.
 - AstroDrizzle** creates images corrected for geometric distortion, and combines associated images with cosmic ray rejection. (Please refer to the [DrizzlePac](#) web page for details.)

This chapter also provides instructions on how to run calacs manually.

- [Chapter 4 - ACS Data Processing Considerations](#) describes how various instrument characteristics affect data processing.
- [Chapter 5 - Data Analysis](#) describes the data reduction software available to work with ACS data, and specific analysis strategies for imaging, spectroscopy, polarimetry, and coronagraphy data.

For the latest information regarding ACS performance and calibration, users are advised to consult the ACS Web pages located at:

- <http://www.stsci.edu/hst/instrumentation/acs>

Typographic Conventions

Visual Cues
Comments

To help you understand the material in this Data Handbook, we will use a few consistent typographic conventions.


Visual Cues

The following typographic cues are used:


- **bold words** identify a **Python** library or function
- `typewriter-like` words identify a file name, system command, or response that is typed or displayed.
- *italic type* indicates a new term, an important point, a mathematical -variable, or a task parameter.
- SMALLER CAPS identifies a header keyword.
- ALL CAPS identifies a table column.

Comments

Occasional side comments point out three types of information, each identified by an icon in the left margin.

 **Warning:** *You could corrupt data, produce incorrect results, or create some other kind of severe problem.*

 **Heads Up:** *Here is something that is often done incorrectly or that is not -obvious.*

 **Tip:** *No problems... just another way to do something or a suggestion that might make your life easier.*

 **Information** *especially likely to be updated on the ACS Web site is indicated by this symbol.*

Acknowledgments

The technical and operational information contained in this Handbook is the summary of the experience gained by members of the STScI ACS Team, the ACS group at the [Space Telescope European Coordinating Facility \(ST-ECF\)](#), and the ACS Instrument Definition Team (ACS IDT).

Members of the STScI ACS Team (2021) are Norman Grogin (Group Lead), Jay Anderson, Roberto Avila (Group Deputy), Ralph Bohlin, Marco Chiaberge¹, Yotam Cohen, Mees Fix, Nimish Hathi, Dean Hines, Samantha Hoffman, Vera Kozhurina-Platais, Ray Lucas, Meaghan McDonald, Melanie Olaes, and Jenna Ryon.

The ST-ECF ACS group was Martin Kuemmel, Harald Kuntschner, and Jeremy Walsh.

The ACS IDT is Holland Ford (PI), Garth Illingworth (Deputy PI), George Hartig, Mark Rafal, Frank Bartko, Tom Broadhurst, Bob Brown, Chris Burrows, Ed Cheng, Mark Clampin, Jim Crocker, Paul Feldman, Marijn Franx, David Golimowski, Randy Kimble, John Krist, Tom La Jeunesse, Mike Lesser, Doug Leviton, George Miley, Marc Postman, Piero Rosati, Bill Sparks, Pam Sullivan, Zlatan Tsvetanov, Paul Volmer, Rick White, Bob Woodruff, Terence Allen, Kenneth Anderson, David Ardila, Narciso Benitez, John Blakeslee, Rychard Bouwens, Larry Bradley, Nicholas J. G. Cross, Ricardo Demarco, Tomotsugu Goto, Caryl Gronwall, Brad Holden, Nicole Homeier, Daniel Magee, André Martel, W. Jon McCann, Simona Mei, Felipe Menanteau, Gerhardt Meurer, Veronica Motta, Alessandro Rettura, Marco Sirianni, Hien Tran, and Andrew Zirm.

The contributions of Andrea Bellini, Matthew Bourque, David Borncamp, Dan Coe, Tyler Desjardins, Susanna Deustua, David Golimowski, Warren Hack, Pey-Lian Lim, Jennifer Mack, Matthew McMaster, Nathan Miles, Max Mutchler, James Noss, Sarah Ogaz, Michele De La Pena, Nor Pirzkal¹, Blair Porterfield, Linda Smith¹, Megan Sosey, Anatoly Suchkov, Alan Welty, Tom Wheeler, and Susan Rose are greatly appreciated.

¹ European Space Agency (ESA)

Chapter 1: ACS Overview

Chapter Contents

- [1.1 Instrument Design and Capabilities](#)
- [1.2 Basic Instrument Operations](#)

1.1 Instrument Design and Capabilities

1.1.1 Detectors

1.1.2 Optical Design

1.1.3 Geometric Distortion

1.1.4 ACS Performance after Servicing Mission 4

ACS was designed for deep, visible to near-IR imaging and spectroscopic surveys using its Wide Field Channel (WFC), near-UV to near-IR imaging and coronagraphy with its now-defunct High Resolution Channel (HRC), and far-UV imaging and spectroscopy using its Solar Blind Channel (SBC). The WFC's discovery efficiency (i.e., the product of its field of view and throughput) is 10 times greater than that of WFPC2. The failure of ACS's CCD electronics in January 2007 halted its near-UV to near-IR science capabilities until Servicing Mission 4 in May 2009, during which the WFC's functionality was fully restored. Unfortunately, the HRC was not recovered. ACS comprises three channels, each optimized for specific goals:

- Wide Field Channel (WFC)
 - 202×202 arcsecond imaging field of view from 3500 Å–11,000 Å
 - 48% peak efficiency at ~7000 Å (including Optical Telescope Assembly)
 - ~0.049 arcsec/pixel with critical sampling at 11,600 Å
 - Low resolution ($R \sim 100$) wide-field spectroscopy from 5500 Å–11,000 Å
 - Polarimetric imaging with relative polarization angles of 0°, 60°, and 120°
- High Resolution Channel (HRC) [before January 2007]
 - 29×26 arcsecond imaging field of view from 1700 Å–11,000 Å
 - 29% peak efficiency at ~6500 Å
 - ~0.028 x 0.025 arcsec/pixel with critical sampling at 6300 Å
 - Low resolution ($R \sim 100$ @ 2000 Å) spectroscopy from 2000 Å–4000 Å
 - Aberrated-beam coronagraphy from 2000 Å–11,000 Å with 1.8 arcsecond and 3.0 arcsecond diameter occulting spots
 - Polarimetric imaging with relative polarization angles of 0°, 60°, and 120°
- Solar Blind Channel (SBC)
 - 35×31 arcsecond imaging field of view from 1150 Å–1700 Å
 - 8.8% peak efficiency at ~1260 Å
 - 0.034×0.030 arcsec pixels provide a good compromise between resolution and field of view
 - Low resolution ($R \sim 100$ @ 1216 Å) spectroscopy from 1150 Å–1700 Å

The WFC's high sensitivity and large field of view make it the preferred camera for deep imaging programs at red and near-infrared wavelengths. By oversampling the *HST* PSF at $\lambda > 6000$ Å, the HRC was especially useful for high precision photometry in stellar population studies before its failure in January 2007. The HRC's coronagraph was used for detection of circumstellar disks and QSO host galaxies.

1.1.1 Detectors

ACS employs different large-format detectors in each channel:

- The WFC detector is a mosaic of two 2048×4096 Scientific Imaging Technologies (SITe) charge-coupled devices (CCDs).
- The HRC detector is a single 1024×1024 SITe CCD.
- The SBC detector is a single CsI micro-channel plate (MCP) with a multi-anode microchannel array (MAMA) readout.

The WFC and HRC CCDs are thinned, backside-illuminated devices regulated by thermoelectric coolers and sealed in evacuated dewars with fused silica windows. The spectral response of the WFC CCDs is optimized for imaging at visible to near-IR wavelengths. The wavelength coverage of the HRC's CCD was similar to that of the WFC, but its spectral response was optimized for near-UV wavelengths. As with all CCD detectors, there is read noise associated with clocking and sampling the collected charge in each pixel through the output amplifiers. The dynamic range of CCD images is determined by the read noise and the depth of the pixel well (~85,000 e⁻ for WFC and ~155,000 e⁻ for HRC), which determines the saturation limit of any one pixel. Hot pixels and cosmic rays affect all ACS CCD exposures.

All ACS detectors produce a time-integrated image in the ACCUM data mode. The HRC also had a target-acquisition mode for coronagraphic observations. The minimum exposure time for WFC is 0.5 seconds and for HRC was 0.1 seconds. WFC and HRC observations are typically split into multiple dithered exposures to allow removal of hot pixels and cosmic rays in post-observation data processing.

The SBC MAMA is a photon-counting detector that provides two-dimensional imaging optimized for far-UV wavelengths. The SBC is operated only in ACCUM mode. SBC observations are subject to both scientific and absolute brightness limits. At high local (≥ 50 counts sec⁻¹ pixel⁻¹) and global ($> 200,000$ counts sec⁻¹) illumination rates, counting becomes nonlinear in a way that is not correctable. At slightly higher illumination rates, MAMA detectors can be permanently damaged. Lower absolute local and global count rate limits have been imposed that define bright object screening limits for each SBC configuration. Targets that violate these screening limits cannot be observed in the proposed configuration.

1.1.2 Optical Design

ACS has two main optical channels, one dedicated to the WFC and one shared by the HRC and the SBC. These channels are shown in [Figures 3.2 and 3.3](#) of the [ACS Instrument Handbook](#). Each channel has independent corrective optics to compensate for *HST*'s spherical aberration. The WFC has three optical elements coated with silver to optimize visible light throughput. The silver coatings absorb wavelengths shortward of 3700 Å. WFC shared two filter wheels with the HRC, which enabled internal WFC/HRC parallel observing for some filter combinations.

The HRC/SBC optical chain comprises three aluminized mirrors overcoated with MgF₂. The HRC was selected by inserting a plane fold mirror into the optical path so that the beam was imaged on the HRC's detector through the WFC/HRC filter wheels. The SBC is selected by moving the fold mirror out of the beam and allowing light to pass through the SBC filter wheel onto the SBC detector. The aberrated beam coronagraph was deployed with a mechanism that inserted a window with two occulting spots at the aberrated telescope focal plane and an apodizer at the re-imaged exit pupil. *For health and safety reasons, use of the coronagraph with SBC is forbidden.*


1.1.3 Geometric Distortion

ACS's focal planes exhibit significantly more geometric distortion than those of previous *HST* instruments. This distortion is principally caused by ACS's optical design, which has a minimal number of components for correcting the spherical aberration induced by the Optical Telescope Assembly (OTA) without introducing coma. The optics allow high throughput, but their focal surfaces are far from normal to the principal rays. The WFC detector is tilted by 22°, so its projected diagonals differ by 8%. The HRC and SBC detectors are tilted by 25°, so their projected diagonals differ by 12%. Consequently, the projected footprints of the detectors on the sky are rhomboidal rather than square, and the pixel scales are smaller along the radial direction of the OTA field of view than along the tangential direction. The angles between the projected x- and y-axis of the detectors are 84.9° for WFC1, 86.1° for WFC2, and 84.2° for HRC.

[Figure 1.1](#) shows the locations of the WFC and HRC apertures in *HST*'s V2/V3 reference frame, the rhomboidal projections of each detector, and the locations of the four readout amplifiers (A, B, C, and D) for each channel. A telescope roll angle of zero degrees corresponds to an on-sky view with the V3 and V2 axes aligned north and east, respectively. The orientations of the physical edges of the detectors are approximately parallel with the V2 and V3 coordinate axes of the telescope, but the eigenaxes of the pixel scale transformation of the WFC are along the projected diagonals of the detectors. The situation is even more irregular for the HRC and SBC because the aperture diagonals do not lie along a radius of the OTA field of view. Moreover, the scale and area of WFC pixels vary by ~10% and ~18%, respectively, from corner to corner. For HRC and SBC, the pixels scale by only ~1% from corner to corner because these detectors have smaller fields of view.

The distorted pixel scales and areas must be corrected before astrometry of ACS images is performed. To perform photometry, the pixel areas must be corrected using pixel area maps (See [Section 5.1.3](#)) or the distortion must be corrected using the [DrizzlePac](#) software.

Figure 1.1: WFC and HRC Apertures Compared with the V2/V3 Reference Frame



The readout amplifiers (A,B,C,D) are indicated on the figure. When ACS images are processed through AstroDrizzle in the DMS¹ data pipeline, the resulting drizzled images are oriented with their x,y axes corresponding approximately to the x,y axes shown in this diagram. Thus, the WFC data products are oriented so that WFC1 (which uses amplifiers A and B) is on top in the positive y -direction (also see [Section 2.2](#)), and the HRC images are oriented such that amplifiers A and B are at the top in this diagram.

1.1.4 ACS Performance after Servicing Mission 4

ACS suffered component failures in its Side 1 and Side 2 electronics in June 2006 and January 2007, respectively. Although the latter failure halted operations of the WFC and HRC cameras, the SBC was unaffected by either failure and remained operational throughout this problematic period. The WFC was recovered after the successful installation of a replacement CCD electronics box (CEB-R) and power supply during Servicing Mission 4 (SM4) in May 2009. Unfortunately, further damage to the HRC power harness in January 2007 prevented recovery of the HRC during SM4, so it remains unavailable for scientific use.

Tests conducted shortly after SM4 showed that:

- The read noise, linearity, pixel full well depth, and amplifier cross-talk of the restored WFC are as good as, or better than, the pre-failure levels ([Table 1.1](#)).

- WFC's dark current, hot pixel fraction, and charge transfer efficiency (CTE) have degraded to the levels expected after extended exposure to *HST*'s trapped radiation environment.
- All WFC images exhibit low-level stripes caused by low frequency (1 mHz to 1 Hz) 1/f noise in the bias-reference voltage generated by the CEB-R. The stripes are constant along each row and span the CCD quadrant boundaries, but they are not stable from frame to frame. The stripes' contribution to the global noise is small (about 0.9 e-), but the correlated nature of the noise may affect photometric precision for very faint sources. An algorithm for removing the stripes based on analysis of the pre-scan regions of the four WFC quadrants has been implemented in the bias-correction stage of **calacs** in the DMS data pipeline. Stripe removal is automatically performed only on full-frame (4096 × 4096 pixel) post-SM4 WFC images; it is not performed on subarray images. STScI provides a standalone routine (**acs_destripe_plus**, in the **acstools** package in the **AstroConda** distribution²) for mitigating the stripes in subarray images of sparse fields. (See [Example 5](#) in [Section 3.5.2](#). More discussion on this is also found in [Sections 3.2.1, 3.4.1, 3.4.2, 3.5.1, 3.5.2, 4.2.1](#), the [ACS web page](#), and the documentation for **acstools**.) For most observing programs, however, the stripes have a negligible effect on science results and their removal provides only cosmetic benefits.
- Bias frames obtained under default CEB-R operation show a 5–10 DN gradient spanning the rows and columns of each image quadrant. This bias gradient is stable over the time between consecutive calibration reference files, so it is effectively removed during normal image reduction and processing.
- The default CEB-R mode induces a signal-dependent bias shift. The DC level of the CEB-R's dual-slope integrator is sensitive to changes in the CCD output voltage in such a way that the bias level is shifted positively by 0.02%–0.30% (depending on the amplifier) of the signal from the previously integrated pixel. This phenomenon is well understood for full-frame WFC images, and an automatic correction for it has been implemented in the bias correction stage in **calacs**. *This correction cannot be applied to the WFC subarray images supported before Cycle 24, but it is applied to the new standard suite of subarrays available from Cycle 24 onwards.*

Table 1.1: Comparison of WFC Performance Before the Side 2 Failure and After SM4

Metric	January 2007	May 2009	January 2015
Read Noise (e- ; gain = 2)	5.5	3.9–4.7	3.9–5.0
Dark Current (e-/pix/hr)	10.7	20–25	28
Hot Pixels (%)	0.68	1.1	2.09
Full Well Depth (e-)	84,000	> 80,000	>80,000
Non-linearity (%)	< 0.1	< 0.2	<0.2
CTE (1620 ee-)	0.999949	0.99989	0.99979
Cross-talk (50Ke-source)	4×10^{-5}	$(5 \pm 4) \times 10^{-5}$	$(5 \pm 4) \times 10^{-5}$

¹The Data Management System (DMS) is the name of the pipeline software that controls the processing and archiving of data at STScI, converting telemetry into FITS data products, populating the Archive catalog, and performing housekeeping on the pipelines. See [Section 3.1](#) for more details.

² AstroConda is a Conda channel maintained by STScI that provides tools and utilities to process and analyze data from HST. Please see the [AstroConda website](#) for details.

1.2 Basic Instrument Operations

- 1.2.1 Target Acquisitions
- 1.2.2 Typical ACS Observing Strategies
- 1.2.3 Data Storage and Transfer
- 1.2.4 Parallel Operations

1.2.1 Target Acquisitions

For most ACS observations, target acquisition is simply specified by the observing aperture in the Phase II proposal. Once the telescope acquires its guide stars, the target will be within ~1–2 arcseconds of the required pointing. For observations using ramp filters, only the desired central wavelength is required to define the acquisition aperture.

For HRC coronagraphic observations, an onboard target acquisition was necessary. The nominal accuracy of the combined target acquisition and slew procedure was ~0.03 arcseconds.

1.2.2 Typical ACS Observing Strategies

WFC users (and erstwhile HRC users) must consider "packaging" their exposures to mitigate the impact of cosmic rays either by dithering their images or constructing a mosaic observing pattern to cover an extended target. SBC exposures do not suffer the effects of cosmic rays or read noise, but long integration times are often needed to obtain sufficient signal-to-noise in photon-starved bluer and/or narrow-band filters. Typical WFC and HRC observing sequences consist of a series of "CR-SPLIT" and dithered 10–20 minute exposures for each program filter. Detailed information about dither and mosaic strategies can be obtained from the [ACS Dither Web Page](#) and the [DrizzlePac website](#).

HRC coronagraphic observations required an initial target acquisition observation to permit centering of the target under the occulting mask. Observers were encouraged but not required to obtain contemporaneous images of coronagraphic PSF-reference stars in order to subtract residual scattered light and enhance the contrast of their science images.

1.2.3 Data Storage and Transfer

At the conclusion of each exposure, science data are read out from the detector and stored in ACS's internal buffer memory until they can be transferred to HST's solid state data recorder (and thereafter to the ground). The internal buffer memory is large enough to hold one full frame WFC image, or sixteen HRC or SBC images. Dumping a full ACS buffer requires 339 seconds; dumping a partially filled buffer (if, e.g., reading out the following exposure would over-fill it) requires proportionately less time. Buffer dumps can occur concurrently with a following exposure if its duration exceeds the buffer-dump duration (typically 339 seconds, corresponding to an entirely filled buffer). Otherwise, an observational overhead is imposed between exposures to clear the buffer for the next exposure.

ACS's internal buffer stores data in a 16 bit per pixel format. This structure imposes a maximum of 65,535 counts per pixel. For the SBC, this is the maximum limit on the total number of detected photons per pixel that can be accumulated in a single exposure. For WFC and HRC gain settings above 2 e-/DN, it is the depth of the pixel well (and not the 16 bit buffer format) that limits the number of photons that can be accumulated without saturating the pixel in a single exposure.

1.2.4 Parallel Operations

Before January 2007, parallel observations with the WFC and HRC were possible with ACS for certain filter combinations. ACS can be used in parallel with the other *HST* science instruments with restrictions described in detail in the [ACS Instrument Handbook](#). *No pure or coordinated parallels with the SBC are allowed.* The policy for applying for parallel observing time is described in the [HST Call for Proposals](#).

Chapter 2: ACS Data Structure

Chapter Contents

- [2.1 Types of ACS Files](#)
- [2.2 ACS File Structure](#)
- [2.3 Data Storage Requirements](#)
- [2.4 Headers and Keywords](#)

2.1 Types of ACS Files

2.1.1 Data Files and Suffixes
2.1.2 Association Tables
2.1.3 Trailer Files

2.1.1 Data Files and Suffixes

File suffixes for ACS data products are given in [Table 2.1](#), and are described below.

- Initial input files to the calibration pipeline:
 - Raw (`raw.fits`) files from Generic Conversion.
 - If applicable, an association table (`asn.fits`) for a complete "observation set."¹
- A single fully-calibrated MAMA image is given the suffix `flt.fits`. (MAMA images do not have an overscan region, and they are not affected by cosmic rays and CTE.)
- Calibration of a single CCD image:
 - After the bias image and bias level are subtracted, and the overscan regions trimmed, a temporary file, with the suffix `blv_tmp.fits`, is created. (By default, CTE-corrected images are also created for WFC data so **calacs** creates a temporary file with suffix `blc_tmp.fits`.)
 - Upon completion of additional calibration steps (dark subtraction, flat-fielding, etc.), the temporary file is renamed with the suffix `flt.fits`. For WFC data, CTE-corrected images, with suffix `flc.fits`, are also created. The `flt.fits` and `flc.fits` files will later serve as input for **AstroDrizzle**.
- A "CR-SPLIT" CCD observation:
 - Raw images from a "CR-SPLIT" observation undergo bias image and bias level subtraction. These images are then combined at the cosmic ray rejection step in **calacs** to create a temporary image with suffix `crj_tmp.fits`. For WFC images with CTE corrections, a file with suffix `crc_tmp.fits` is also created.
 - Other basic calibrations are performed on the temporary combined image. It is then renamed with the `crj.fits` suffix, and `crc.fits` suffix for WFC CTE-corrected images.
 - Individual calibrated images (`flt.fits/flc.fits`) are also created for each exposure in the "CR-SPLIT" observation.
- Multiple CCD exposures (e.g., CR-SPLIT):
 - Follows the same procedure described for "CR-SPLIT" CCD observations.
- A MAMA observation consisting of several repeated sub-exposures:
 - Calibrated `flt.fits` images are created for each sub-exposure.
 - A summed flat-fielded image is created, with suffix `sfl.fits`.

Table 2.1: ACS File Suffixes

File Suffix	File Description	Units
<code>raw.fits</code>	Raw uncalibrated image from a single exposure.	<i>DN</i>
<code>asn.fits</code>	Association table for an observation set.	–
<code>spt.fits</code>	Telemetry and engineering data.	–
<code>trl.fits</code>	Trailer file containing calacs processing comments. (This is the same as <code>.tra</code> files generated during manual calibration.)	–
<code>blv_tmp.fits</code> <code>blc_tmp.fits</code>	Overscan-trimmed individual exposure (these will be renamed <code>flt.fits</code> or <code>flc.fits</code> after all basic calibrations are completed).	<i>electrons</i>
<code>crj_tmp.fits</code> <code>crc_tmp.fits</code>	CR-rejected combined image created using <code>blv_tmp.fits</code> and <code>blc_tmp.fits</code> for WFC CTE-corrected data. (these will be renamed <code>crj.fits</code> and <code>crc.fits</code> for WFC CTE-corrected images, after all basic calibrations are completed.)	<i>electrons</i>
<code>flt.fits</code> <code>flc.fits</code>	Calibrated, flat-fielded individual exposure. CTE-corrected images with suffix <code>flc.fits</code> are also created for WFC data.	<i>electrons</i>
<code>crj.fits</code> <code>crc.fits</code>	Calibrated and combined image, with CR rejection (using images from "CR-SPLIT" sub-exposures). The CTE-corrected version, for WFC images, has suffix <code>crc.fits</code> .	<i>electrons</i>
<code>sfl.fits</code>	Calibrated and summed MAMA image (no CR rejection needed) created from sub-exposures in an observation.	<i>electrons</i>
<code>drz.fits</code> <code>drc.fits</code>	Calibrated, geometrically-corrected, dither-combined image (created by AstroDrizzle , which is not a part of calacs). The CTE-corrected version, for WFC images, has suffix <code>drc.fits</code> .	<i>electrons /sec</i>



*Intermediate calibrated products created by **calacs**, such as `sfl.fits`, `blv_tmp.fits`, `blc_tmp.fits`, `crj_tmp.fits`, `crc_tmp.fits`, `crj.fits` and `crc.fits`, are, by default, not delivered by the Archive.*

Standard calibrated files delivered from the Archive contains these extensions: `asn.fits`, `spt.fits`, `trl.fits`, `flt.fits`, and `drz.fits`. For WFC images, `flc.fits` and `drc.fits` are also delivered.

*To obtain intermediate **calacs** products from the Archive, enter specific extensions (e.g., `crj`, `crj_tmp`) in a field titled " File Extensions Requested" near the bottom of the Archive data request Web page.*

2.1.2 Association Tables

Association tables describe and track the relationship or "associations" between data products for a set of observations. Such relationships include repeated exposures in observations, "CR-SPLIT" observations, and dithered observations. ACS association tables can be used to instruct **calacs** to create different levels of calibration products. These tables are particularly useful for keeping track of complex observations, for example, an observation at a specific dither position that may be additionally split into multiple exposures. Edited association tables can also be used with **calacs** to create non-default calibration products (see [Example 3](#) in [Section 3.5](#)).

ACS data files are given the following definitions:

- A single image from an *exposure* or *sub-exposure* is the "atomic unit" of *HST* data.
- A *dataset* is a collection of files having a common rootname (first nine characters of the image name).
- A *sub-product* is created by combining a subset of the exposures in an association.
- A *product* is created by combining sub-products, or in some cases, individual exposures (before they were incorporated into a sub-product), of an association

An ACS association table has three primary columns: MEMNAME, MEMTYPE, and MEMPRSNT. (An example of an association table is shown in [Table 2.3](#).)

- The column MEMNAME lists the name of each exposure in the association and names of calacsoutput products.
- The column MEMTYPE describes the role of a file in the association. A unique set of MEMTYPES, specific to ACS, were adopted to provide descriptions for multiple products. These types are summarized in [Table 2.2](#).
- The MEMPRSNT column indicates the calibration status of each product.

Table 2.2: Exposure Types, or MEMTYPES, in ACS Associations

MEMTYPE	File Description
EXP-CRJ	An image that is part of a "CR-SPLIT" observation. ("EXP-CRJ" is used when there is only one "CR-SPLIT" observation in an association.)
EXP-CRn	Same as "EXP-CRJ," but used when there are multiple "CR-SPLIT" observations in an association. n is a numerical identification for each "CR-SPLIT" set (e.g., "EXP-CR1," "EXP-CR2").
PROD-CRJ	A calibrated and CR-rejected combined image created from a "CR-SPLIT" observation. ("PROD-CRJ" is used when there is only one "CR-SPLIT" observation in an association.)
PROD-CRn	Same as "PROD-CRJ" but used when there are multiple "CR-SPLIT" observations in an association. n is a numerical identification for each "CR-SPLIT" set and its corresponding PROD-CRn combined image (e.g., "PROD-CR1," "PROD-CR2").
EXP-RPT	An image that is part of an observation consisting of several repeated sub-exposures. ("EXP-RPT" is used when there is only one such observation in an association.)
EXP-RPn	Same as "EXP-RPT," but used when there are multiple observations in an association, each containing a series of repeated sub-exposures. n is a numerical identification for each such observation (e.g., "EXP-RP1," "EXP-RP2").

PROD-RPT	A calibrated summed MAMA image, created from an observation containing a series of repeated sub-exposures. ("PROD-RPT" is used when there is only one such observation in an association.)
PROD-RPn	Same as "PROD-RPT," but used when there are multiple observations in an association, each containing a series of repeated sub-exposures. n is a numerical identification for each such observation and its corresponding summed PROD-RPTn image (e.g., "PROD-RPT1," "PROD-RPT-2").
EXP-DTH	An image from an observation that is part of a dither pattern.
PROD-DTH	A dither-combined output product.

An example of an association table is shown in [Table 2.3](#). But first, to trace back its origins, begin by looking at the Phase II proposal commands that created the data. This example came from proposal 10605, visit 1, exposure log sheet number 1. The observations, using ACS/WFC, were taken as a two-point dither with a "CR-SPLIT=2" at each dither point. The proposal's dither pattern specification looks like this:

```

Pattern_Number: 1
Pattern_Type      Primary_Pattern      Secondary_Pattern
Pattern_Type      ACS-WFC-DITHER-LINE
Pattern_Purpose     DITHER
Number_Of_Points  2
Point_Spacing     3.011
Line_Spacing      <none>
Coordinate_Frame  POS-TARG
Pattern_Orient    85.28
Angle_Between_Sides <none>
Center_Pattern    NO

```

The exposure log sheet commands used to execute these observations were as follows:

```

-----
Exp | Target| Instr| Oper.| Aper | Spectral| Central| Optional|Num| Time| Special
Num | Name | Config| Mode| or FOV| Element| Waveln.| Parameters|Exp| Requirements
-----
1 MESSIER-081 ACS/WFC ACCUM WFC F555W CR-SPLIT=2 1 2900 S PATTERN 1 1-2
-DWARF-A
-----
(line 2 is not shown)

```

A standard calibrated data retrieval from the Archive for images taken by Visit 1, exposure 1, are listed below. (Intermediate **calacs** products from the Archive have to be specifically requested, therefore, the **crj.fits** and **crc.fits** files are not included in the delivery.)

```

j9cm01010_asn.fits  j9cm01jvq_flt.fits  j9cm01jwq_flt.fits  j9cm01k2q_flt.fits  j9cm01k4q_flt.fits
j9cm01010_drz.fits  j9cm01jvq_flc.fits  j9cm01jwq_flc.fits  j9cm01k2q_flc.fits  j9cm01k4q_flc.fits
j9cm01010_drc.fits  j9cm01jvq_spt.fits  j9cm01jwq_spt.fits  j9cm01k2q_spt.fits  j9cm01k4q_spt.fits
j9cm01010_spt.fits  j9cm01jvq_trl.fits  j9cm01jwq_trl.fits  j9cm01k2q_trl.fits  j9cm01k4q_trl.fits
j9cm01010_trl.fits

```

There are two pairs of 4 single calibrated images, one pair with suffix `flt.fits`, and the other, corrected for CTE, with suffix `flc.fits`.

- From the "CR-SPLIT" sub-exposures at dither point 1:

`j9cm01jvq_flt.fits` and `j9cm01jvq_flc.fits`
`j9cm01jwq_flt.fits` and `j9cm01jwq_flc.fits`

- From the "CR-SPLIT" sub-exposures at dither point 2:

`j9cm01k2q_flt.fits` and `j9cm01k2q_flc.fits`
`j9cm01k4q_flt.fits` and `j9cm01k4q_flc.fits`

- The combined image created by **AstroDrizzle**, in the pipeline, from four `flt.fits` and `flc.fits` images, respectively, are:
`j9cm01010_drz.fits` and `j9cm01010_drc.fits`

Table 2.3: Contents of Association Table, `j9cm01010_asn.fits`

Column	1	2	3
Label	MEMNAME	MEMTYPE	MEMPRSNT
1	J9CM01JVQ	EXP-CR1	yes
2	J9CM01JWQ	EXP-CR1	yes
3	J9CM01K2Q	EXP-CR2	yes
4	J9CM01K4Q	EXP-CR2	yes
5	J9CM01010	PROD-DTH	yes
6	J9CM01011	PROD-CR1	yes
7	J9CM01012	PROD-CR2	yes

- Rows 1 and 2 describe the association's first "CR-SPLIT" observation at dither point 1.
 - "MEMNAME" shows the image rootnames for two sub-exposures in the first "CR-SPLIT" observation.
 - "MEMTYPE" of "EXP-CR1" means that the images came from the association's first "CR-SPLIT" observation.
 - "MEMPRSNT" set to "yes" indicates that those images underwent standard calibrations in **calacs**.
- Rows 3 and 4 describe the association's second "CR-SPLIT" observation at dither point 2. The two images have rootnames "J9CM01K2Q" and "J9CM01K4Q." "EXP-CR2" means they came from the second CR-SPLIT observation, and "yes" indicates that those images underwent standard calibrations in **calacs**.
- In row 5, "PROD-DTH" and "yes" indicate that the single exposure images in the association were drizzle-combined in the pipeline by **AstroDrizzle**, and named with the rootname specified in the MEMNAME column. In this example, four `flt.fits` images were drizzled-combined to produce `j9cm01010_drz.fits`, and four `flc.fits` images were drizzle-combined to create `j9cm01010_drc.fits`. (If "MEMPRSNT" had been set to "no," the **AstroDrizzle** step would have been omitted in the pipeline.)
- In Row 6, "PROD-CR1" and "yes" indicate that the two images from the first "CR-SPLIT" (rows 1 and 2) were combined with cosmic ray rejection by **calacs** to create an image with the rootname specified in

the MEMNAME column. This was done for `flt.fits` and `flc.fits` images to create the files `J9CM01011_crj.fits` and `J9CM01011_crc.fits`, respectively.

- In Row 7, "PROD-CR2" and "yes" indicate that the two images from the second "CR-SPLIT" (rows 3 and 4) were combined with cosmic ray rejection by **calacs** to create an image with the rootname specified in the MEMNAME column. This was done for `flt.fits` and `flc.fits` images to create the files `j9cm01012_crj.fits` and `j9cm01012_crc.fits`, respectively.

2.1.3 Trailer Files

Each task in the **calacs** package creates messages that describe the progress of the calibration; these messages are directed to STDOUT (STanDard OUTput), which simply means that processing messages appear on the screen during the **calacs** run.

In pipeline processing for first and second generation *HST* instruments, where data files were calibrated one at a time, trailer files were created by simply redirecting the STDOUT contents to a file. However, the ACS pipeline was designed to calibrate several images at a time (like those described in [Section 2.1.2](#)), and create different types of output files. Therefore, each task within the **calacs** package must decide how to populate the trailer files associated with each product.

calacs will always overwrite information in existing trailer files from previous runs of **calacs** while preserving any comments generated by Generic Conversion. This ensures that the trailer files accurately reflect the most recent processing performed by **calacs**. After the Generic Conversion entries, the string "CALACSBEG" marks the first comment added to a trailer file by **calacs**. If the trailer file already exists, **calacs** searches for this string, then begins to write new processing comments from that point onwards, over-writing previous **calacs** comments. If "CALACSBEG" is not found in an existing trailer file, **calacs** will write that string at the end of the trailer file, then continue populating the trailer file with calibration processing comments.

As each image is reprocessed, an accompanying trailer text file with the suffix "tra" (without the `.fits` extension) is created. (The `trl.fits` file from the Archive has the same content, in FITS table format.)

Following the processing hierarchy specified in the association table, information in trailer files belonging to images used for creating a higher level product will be included in the trailer file of that higher level product. In other words, the trailer file for any product processed by the pipeline will contain processing comments from trailers belonging to each input file.

Linking trailer files together can result in multiple occurrences of the "CALACSBEG" string. Only the first, however, determines where **calacs** will begin overwriting comments if an observation is reprocessed.

¹ An "observation" set is a group of exposures under the umbrella of one unique ID. Each set has the same target, instrument configuration, operating mode, aperture, and spectral elements. An observation set created by the pipeline usually contains a set of dithered, repeated, and/or "CR-SPLIT" exposures taken within a visit. The data products in each calibrated observation set are described in its association table.

2.2 ACS File Structure

The ACS calibration pipeline assembles data received from *HST* into datasets, applies standard calibrations (so that calibrated image header keyword values can be entered in the Archive database), and stores uncalibrated datasets in the *HST* Data Archive. When a user requests data, it is fetched from the MAST static archive unless newer versions of the pipeline and/or calibration reference files are present. In the latter case, the requested data will be moved to the top of the reprocessing queue.

Data from the Archive arrives as multi-extension FITS files that have three main categories, as listed below. (Please see [Figure 2.1](#) for a graphical representation of this information.) Single quotes around the extension names are required in **Python** usage.

- Global header (PRIMARY): this is always in extension [0], and contains a selection of header keyword values applicable to all information in the FITS file.
- Science image ('SCI'), error array ('ERR'), and data quality image ('DQ'):

- HRC and SBC are single-detector channels. For these channels, the science image is in extension [1] or ['SCI', 1], error array in extension [2] or ['ERR', 1], and data quality array in extension [3] or ['DQ', 1].

- WFC data comes from two chips, each with its own science, error, and data quality arrays. For WFC2, the science image is in extension [1] or ['SCI', 1], error array in extension [2] or ['ERR', 1], and data quality array in extension [3] or ['DQ', 1]. For WFC1, the science image is in extension [4] or ['SCI', 2], error array in extension [5] or ['ERR', 2], and data quality array in extension [6] or ['DQ', 2]. See [Figure 2.2](#) for a visual representation of the WFC chips projected onto the V2,V3 axis.

- **AstroDrizzle** extensions:

- Drizzled data from the pipeline have the suffix `drz.fits`, and additionally `drc.fits` for WFC CTE-corrected data. These products are delivered as multi-extension FITS files with three image data extensions: science image ['SCI'], weight image ['WHT'], and context image ['CTX']. However, **AstroDrizzle** also adds new FITS extensions to the `flt.fits/flc.fits` and `drz.fits/drc.fits` files, as summarized below. For more information on drizzled data, please refer to the [The DrizzlePac](#) documentation.

- For `flt.fits/flc.fits` images: during OTFR processing, **AstroDrizzle** updates the WCS of these images with distortion corrections. Linear distortion corrections (scale, rotation, and time-dependent skew) are incorporated into the CD-Matrix. Coefficients for higher order polynomial functions that describe the distortion corrections are stored as SIP¹ header keyword values. Some distortion corrections, however, cannot be expressed as equations and have to be stored in array form as FITS extensions

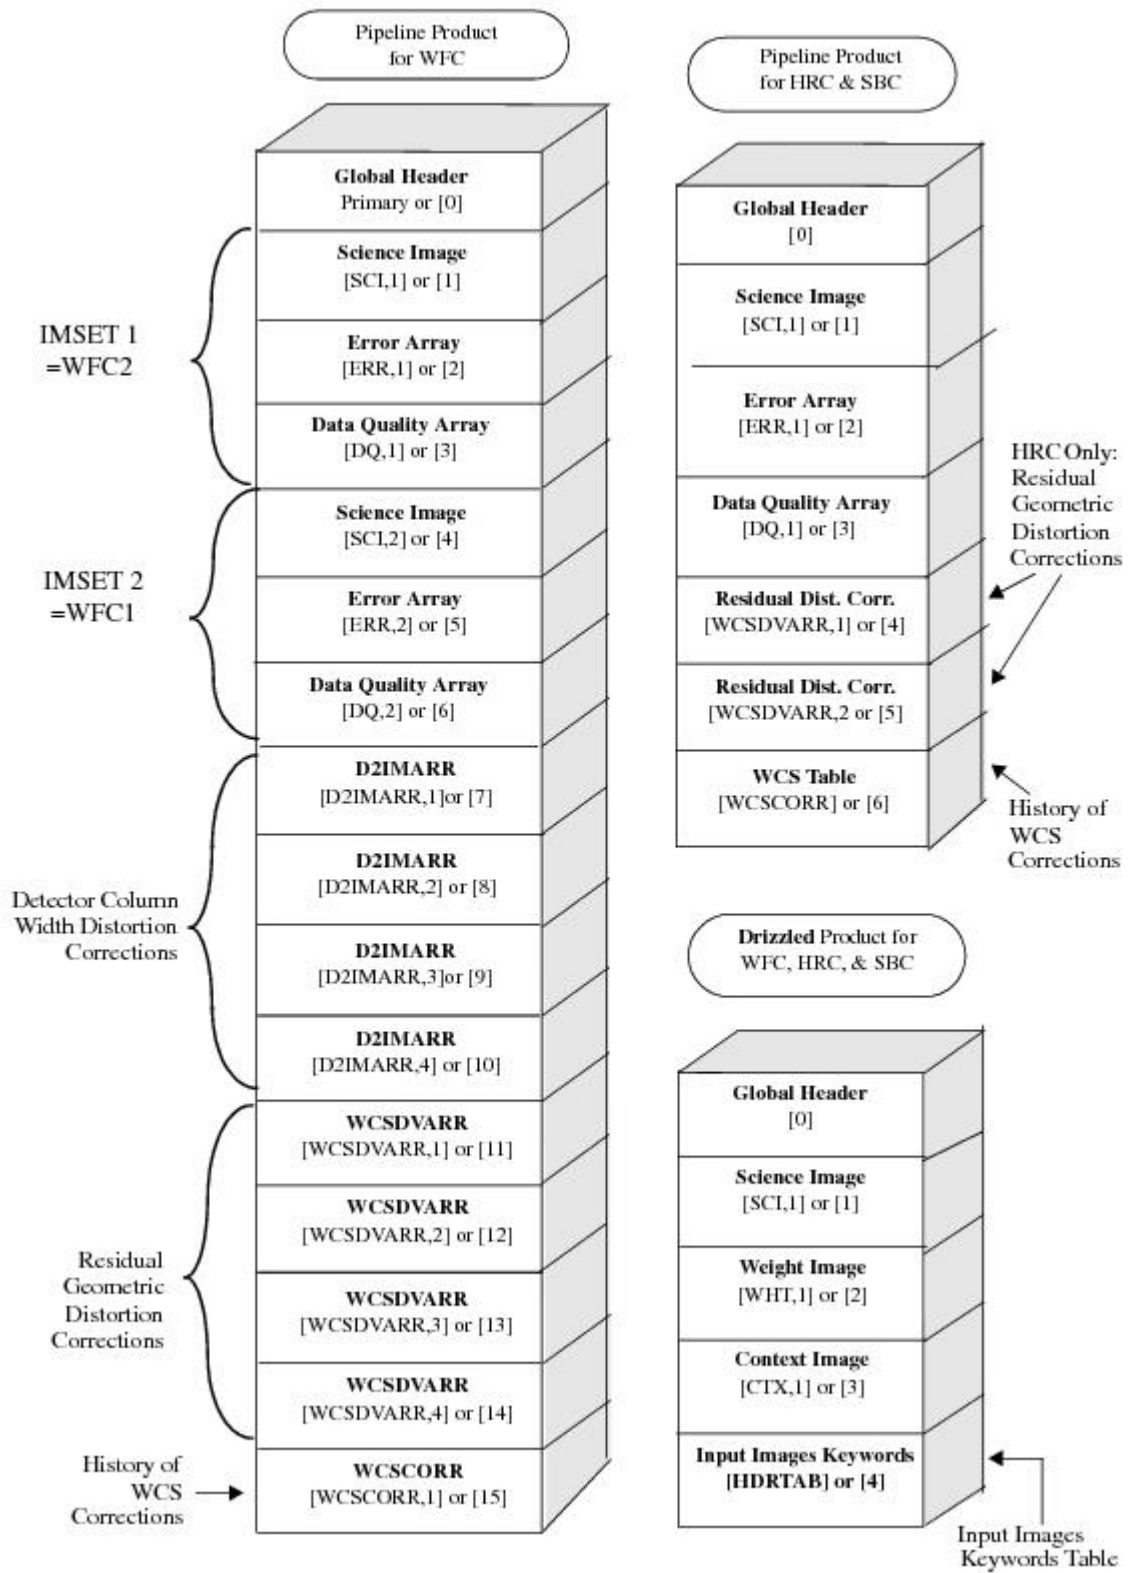
- The ['D2IMARR'] fits extensions are tabular data required only for WFC, with one extension per chip axis. They are filter-independent corrections for the CCD pixel-grid irregularities resulting from the manufacturing process and contain both X and Y corrections for each WFC CCD chip.

- The ['WCSDVARR'] fits extensions, with one extension per chip axis, hold tabular data which describe small-scale distortions due to filter-dependent non-polynomial distortion corrections which include both X and Y corrections for each WFC CCD chip. Therefore, a WFC `flt.fits/flc.fits` image has four ['WCSDVARR'] extensions while HRC and SBC only have two ['WCSDVARR'] extensions.

- If the data were reprocessed with a new distortion correction reference file, a FITS extension called ['WCSCORR'] is added that contains a history of WCS changes.

- For `drz.fits/drc.fits` images: a fourth extension has been added, called ['HDRTAB'], which is a compilation of important header keywords that have unique values for each input image.


Figure 2.1: Data Format for Calibrated and Drizzled ACS Modes



For calibrated science data, WFC1 (chip 1) corresponds to extension ['SCI',2] or [4].

Figure 2.2 shows the WFC apertures plotted with respect to the V2,V3 reference frame, and oriented such that the x-axis runs approximately towards the right and the y-axis runs approximately straight up. In pipeline data products, WFC2 is displayed below WFC1 (along the y-axis) and is therefore designated as extension 1.

Figure 2.2: WFC apertures Compared with the V2/V3 Reference Frame



In pipeline data products, WFC2 is displayed below WFC1 (along the y-axis) and is therefore designated as extension 1. To display the science image for WFC1, the user must specify the extension `fits['SCI',2]` or `fits[4]`. The WFC1 error and data quality arrays are specified as `['ERR',2]` (or [5]) and `['DQ',2]` (or [6]), respectively. The readout amplifiers (A,B,C, D) are indicated on the figure. Refer back to Figure 2.1 for the full list of extensions by detector.

¹Simple Image Polynomial (SIP) convention (Shupe, et. al, 2005). This convention has been in use for describing the geometry of Spitzer Space Telescope images, and it has become a FITS standard. Representing image distortion corrections using the SIP convention improves the handling of image combination and astrometric information.

2.3 Data Storage Requirements

Users are reminded to consider the large size of WFC exposures when allocating disk space for storing and reprocessing ACS data. Raw images serve as input to the calibration pipeline and have the file sizes (in MB) given in Table 2.4. A WFC full-frame exposure contains two SCI arrays, while a WFC subarray exposure and exposures from the HRC and SBC detectors each have a single SCI array. The raw image sizes presume that both the SCI and DQ arrays are populated with short integer values, but that the ERR arrays are NULL (all pixels have a value of zero).

During **calacs** processing, the SCI arrays are converted from integer to floating point values. The null ERR array is populated with floating point values, and the null DQ array is populated with integer values. As a result, the size of calibrated images (in MB) is much larger.

The image size in pixels is given in Table 2.5. Calibrated images taken with the WFC and HRC detectors are smaller in image dimensions than raw images because any prescan and overscan regions have been trimmed during processing.

After **calacs** calibration in the pipeline, **AstroDrizzle** combines images in an association to create a product with four extensions: science (SCI), weight (WHT), context (CTX), and header keywords (HDRTAB). Once the distortion is corrected, the size of a drizzled image will be larger (in pixel dimensions and file size) than a calibrated image because the drizzled image's projection on the sky is rhombus-shaped, rather than square. Also, the size of the image offsets, from sub-pixel dithers to image mosaics, will increase the image field of view and hence the image size. The specific dimensions of a drizzled image depends on the image orientation and on which distortion model is in use in the pipeline, and will vary slightly (about 1 to 2 pixels) due to the effects of velocity aberration.

Table 2.4: Size (in MB) of Raw, Calibrated, and Drizzled ACS Images for Each Detector

Detector	Size of FITS File (S_{raw})	Size of FITS File (S_{cal})	Size of FITS File ^a (S_{drz})
WFC (2 Chips)	34.4MB	168.0 MB	220.0 MB
HRC (1 Chip)	2.3 MB	10.6 MB	16.0 MB
SBC (1 Chip)	12.2 MB	10.6 MB	128.2 MB

^aThe drizzled image sizes are only distortion-corrected and assume no additional rotation, dither offset, or scale change.

Table 2.5: Size (in Pixels) for Raw, Calibrated, and Drizzled ACS Images for Each Detector

Detector	X_{raw}	Y_{raw}	X_{cal}	Y_{cal}	X_{drz} ^a	Y_{drz} ^a
WFC (2 Chips)	4144	2068	4096	2048	4221	4334
WFC1-1K	1046	1024	1024	1024	1063	1029
WFC1-2K	2070	2046	2048	2046	2124	2073
WFC1A-1K	2072	1024	2048	1024	2026	1082
WFC1B-1K	2072	1024	2048	1024	2079	1070

WFC2C-1K	2072	1024	2048	1024	2047	1091
WFC2D-1K	2072	1024	2048	1024	2111	1091
WFC1A-2K	2072	2068	2048	2048	2082	2075
WFC1B-2K	2072	2068	2048	2048	2124	2075
WFC2C-2K	2072	2068	2048	2048	2089	2112
WFC2D-2K	2072	2068	2048	2048	2142	2123
WFC1A-512	2072	512	2048	512	1994	594
WFC1B-512	2072	512	2048	512	2055	576
WFC2C-512	2072	512	2048	512	2027	574
WFC2D-512	2072	512	2048	512	2096	569
HRC (1 Chip)	1062	1044	1024	1024	1163	1134
SBC (1 Chip)	1024	1024	1024	1024	1423	1640

^aThe drizzled image sizes are only distortion-corrected and assume no additional rotation, dither offset, or scale change. The size of drizzled images will vary slightly due to velocity aberration effects.

While the size of calibrated, drizzled HRC and SBC images is comparable to that of STIS, WFPC2, or WFC3 IR data, the ACS WFC images are over 16 times larger and comparable to that of WFC3 UVIS. The following equation can be used to estimate the minimum amount of free storage required for processing associated ACS data:

$$D_{min} = n \cdot S_{raw} + m \cdot (1 + n) \cdot S_{cal} + m \cdot (2.1 + p^2) \cdot S_{drz}$$

where,

- D_{min} is the minimum free disk space required for processing,
- n is the number of exposures in each "CR-SPLIT" set or repeated exposures for an observation,
- m is 2 when both regular and CTE-corrected WFC products are being produced. Otherwise m is 1,
- S_{raw} is the size of the raw exposure (from [Table 2.4](#)),
- S_{cal} is the size of the calibrated exposure (from [Table 2.4](#)),
- S_{drz} is the size of the distortion-corrected exposure (from [Table 2.4](#)),
- p is the percentage shift (in pixels) across all dither positions

2.4 Headers and Keywords

ACS image headers for WFC, HRC, HRC-ACQ, and SBC observations have very similar formats. However, MAMA (SBC) headers require a slightly modified structure from the CCD format, and HRC-ACQ headers require several unique keywords.

ACS header keywords can be viewed in several ways. The example below shows how these tasks are used to get the values for group (or extension) keyword CRVAL1.

Python INPUT:

```
from astropy.io import fits

hdu = fits.open('j8cdaldsqflt.fits')
hdr = hdu['sci',1].header
hdu.close()

print(hdr['CRVAL1'])
```

Python OUTPUT: 5.952775438936009

ACS image header keywords are described in Tables 2.6 and 2.7. When printing the header, note that the keywords are grouped together by type, for example, "Target Information," "Science Instrument Configuration," "Calibration Switches," and "Calibration Reference files." There are also convenience functions for querying header keywords. For more information, see the [Astropy documentation](#).

Keywords in the image's primary header (FITS extension 0) contain parameter values that are common to the entire image dataset, and are given in Table 2.6.

Extension or group header keywords (FITS extensions > 0), as shown in Table 2.7, contain parameter values that are specific to the image group (however, many primary header keyword values are also accessible when querying an image group header).

Table 2.6: ACS primary header keywords (FITS extension 0). Some keywords are no longer in use, but are included for users who may have older data.

KEYWORD	DESCRIPTION
SIMPLE	Data conform to FITS standard (T/F)
BITPIX	Bits per data value
NAXIS	Number of data axes, always set to zero for the primary header
EXTEND	FITS data may contain extensions (T/F)
ORIGIN	FITS file originator
NEXTEND	Number of standard extensions
GROUPS	Image is in group format (T/F)

DATE	Date this file was written (YYYY-MM-DD)
FILENAME	Name of file
FILETYPE	Type of data found in data file
TELESCOP	Telescope used to acquire data
INSTRUME	Identifier for instrument used to acquire data
EQUINOX	Equinox of celestial coordinate system
DATA DESCRIPTION KEYWORDS (all detectors)	
ROOTNAME	Rootname of the observation set
IMAGETYP	Type of exposure identifier
PRIMESI	Instrument designated as prime
TARGET INFORMATION (all detectors)	
TARGNAME	Proposer's target name
RA_TARG	Right ascension of the target (deg; J2000)
DEC_TARG	Declination of the target (deg; J2000)
PROPOSAL INFORMATION (all detectors)	
PROPOSID	PEP proposal identifier
LINENUM	Proposal logsheet line number
PR_INV_L	Last name of the principal investigator
PR_INV_F	First name of the principal investigator
PR_INV_M	Middle name/initial of the principal investigator
EXPOSURE INFORMATION (all detectors)	
SUNANGLE	Angle between the Sun and V1 axis
MOONANGL	Angle between the Moon and V1 axis
SUN_ALT	Altitude of the Sun above Earth's limb
FGSLOCK	Commanded FGS lock (FINE, COARSE, GYROS, UNKNOWN)
GYROMODE	Number of gyros scheduled, T=3+OBAD (number of gyros available for use)
REFFRAME	Guide star catalog version
MTFLAG	Moving target flag (T=moving target)
DATE-OBS	UT date of observation start (YYYY-MM-DD)

TIME-OBS	UT time of observation start (HH:MM:SS)
EXPSTART	Exposure start time (modified Julian date)
EXPEND	Exposure end time (modified Julian date)
EXPTIME	Exposure duration in seconds (calculated; exposure time as executed)
TEXPTIME	Total exposure time for drizzled product
DARKTIME	Total integration time for detector dark current in seconds (calculated)
EXPFLAG	Exposure interruption indicator
QUALCOM1, QUALCOM2, QUALCOM3, QUALITY	Data quality comments. These keywords are populated if there are problems with the exposure.
POINTING INFORMATION (all detectors)	
PA_V3	Position angle of the V3-axis of <i>HST</i> (deg)
TARGET OFFSETS (all detectors)	
POSTARG1	POSTARG in axis 1 direction
POSTARG2	POSTARG in axis 2 direction
DIAGNOSTIC KEYWORDS (all detectors)	
OPUS_VER	OPUS software system version number
CSYS_VER	Calibration software system version number
CAL_VER	calacs version number
PROCTIME	Pipeline processing time (modified Julian date)
SCIENCE INSTRUMENT CONFIGURATION (all detectors except where noted)	
OBSTYPE	Observation type (imaging, spectroscopic, coronagraphic)
OBSMODE	Operating mode
CTEIMAGE	Type of charge transfer image, if applicable
SCLAMP	Lamp status: NONE or name of lamp which is on
NRPTEXP	Number of repeat exposures in set (default=1)
SUBARRAY	Data from a subarray (T) or full frame (F)

DETECTOR	Detector in use: WFC, HRC, or SBC
FILTER1	Element selected from filter wheel 1
FW1OFFST	Computed filter wheel offset
FW1ERROR	Filter wheel position error flag
FILTER2	Element selected from filter wheel 2
FW2OFFST	Computer filter wheel offset
FW2ERROR	Filter wheel position error flag
FWSOFFST	Computed filter wheel offset
FWSEERROR	Filter wheel position error flag
LRFWAVE	Proposed linear ramp filter wavelength
APERTURE	Aperture name
PROPAPER	Proposed aperture name
DIRIMAGE	Direct image for grism or prism exposure
CTEDIR	CTE measurement direction (serial or parallel; CCDs only)
CRSPLIT	Number of cosmic ray split exposures (CCDs only)
MAMA OFFSETS (SBC only)	
MOFFSET1	Axis 1 MAMA offset (low-res pixels)
MOFFSET2	Axis 2 MAMA offset (low-res pixels)
LOCAL RATE CHECK IMAGE (SBC only)	
LRC_XSTS	Local rate check image exists (T/F)
LRC_FAIL	Local rate check failed (T/F)
CALIBRATION SWITCHES: PERFORM, OMIT (all detectors except where noted)	
STATFLAG	Calculate statistics (not used; statistics are always calculated by calacs)
WRTErr	Write out error array extension (not used; always true for calacs)
DQICORR	Data quality initialization
ATODCORR	Correct for analog-to-digital conversion errors (CCDs only; not used by calacs)
BLEVCORR	Subtract bias level computed from overscan image (CCDs only)
BIASCORR	Subtract 2-D bias image (CCDs only)
FLSHCORR	Post-flash correction (CCDs only)

CRCORR	Combine observations to reject cosmic rays (CCDs only)
EXPSCORR	Process individual observations are cosmic ray rejection (CCDs only)
SHADCORR	Apply shutter shading correction (CCDs only; always set to OMIT for ACS)
GLINCORR	Correct for global detector non-linearities (SBC only)
LFLGCROR	Flag pixels for local and global non-linearities (SBC only)
PCTECORR	CTE correction (always set to OMIT for non-WFC exposures)
DARKCORR	Subtract dark image
FLATCORR	Flat-field data
PHOTCORR	Populate photometric header keywords
RPTCORR	Add individual repeat observations
DRIZCORR	Drizzle processing
SINKCORR	Flag sink pixels (affects WFC only since 2015)
CALIBRATION REFERENCE FILES (all detectors except where noted)	
BPIXTAB	Bad pixel table
MLINTAB	MAMA linearity correction table (SBC only)
CCDTAB	CCD calibration parameter table (CCDs only)
ATODTAB	Analog-to-digital correction file (CCDs only; not used by calacs)
OSCNTAB	CCD overscan table (CCDs only)
BIASFILE	2-D bias image file name (CCDs only)
FLSHFILE	Post-flash correction file (CCDs only)
CRREJTAB	Cosmic ray rejection parameters (CCDs only)
SHADFILE	Shutter shading correction file (CCDs only; not used by calacs)
PCTETAB	CTE correction table (always set to N/A for non-WFC exposures)
DRKCFE	CTE-corrected dark image file name (always set to N/A for non-WFC exposures)
DARKFILE	Dark image file name
PFLTFILE	Pixel-to-pixel flat-field name
DFLTFILE	Delta flat-field name
LFLTFILE	Low-order flat-field name
PHOTTAB	Photometric throughput table (not used by calacs)

GRAPHTAB	<i>HST</i> graph table (not used by calacs)
COMPTAB	<i>HST</i> components table (not used by calacs)
IDCTAB	Image distortion correction table
DGEOFILE	Distortion correction image (not used by AstroDrizzle for calacs)
MDRIZTAB	AstroDrizzle parameter table
CFLTFILE	Coronagraphic spot flat-field name
SPOTTAB	Coronagraphic spot offset table
IMPHTTAB	Image photometry table
D2IMFILE	Column correction reference files
NPOLFILE	Non-polynomial offsets reference files
SKNCFILE	Map of sink pixels (affects WFC only since 2015)
SATUFILE	Full-well saturation level map
COSMIC RAY REJECTION ALGORITHM PARAMETERS (WFC, HRC, HRC ACQ)	
MEANEXP	Reference exposure time for parameters
SCALENSE	Multiplicative scale factor applied to noise
INITGUES	Initial guess method (MIN or MED)
SKYSUB	Sky value subtracted (MODE or NONE)
SKYSUM	Sky level from the sum of all constituent images
CRSIGMAS	Statistical rejection criteria
CRRADIUS	Rejection propagation radius (pixels)
CRTHRESH	Rejection propagation threshold
BADINPQD	Data quality flag bits to reject
REJ_RATE	Rate at which pixels are affected by cosmic rays
CRMASK	Flag cosmic ray-rejected pixels in input files (T/F)
OTFR Keywords (all detectors)	
T_SGSTAR	OMS calculated guide star control
PATTERN KEYWORDS (all detectors)	
PATTERN1	Primary pattern type
P1_SHAPE	Primary pattern shape

P1_PURPS	Primary pattern purpose
P1_NPTS	Number of points in primary pattern
P1_PSPAC	Point spacing for primary pattern (arcsec)
P1_LSPAC	Line spacing for primary pattern (arcsec)
P1_ANGLE	Angle between sides of parallelogram pattern (deg)
P1_FRAME	Coordinate frame of primary pattern
P1_ORINT	Orientation of pattern to coordinate frame (deg)
P1_CENTR	Center pattern relative to pointing (yes/no)
PATTSTEP	Position number of this point in the pattern
POST-FLASH PARAMETERS (WFC, HRC, HRC ACQ)	
FLASHDUR	Exposure time in seconds (0.1 to 409.5)
FLASHCUR	Post-flash current (OFF, LOW, MED, HIGH)
FLASHSTA	Status (SUCCESSFUL, ABORTED, NOT PERFORMED)
SHUTRPOS	Shutter position (A or B)
ENGINEERING PARAMETERS (WFC, HRC, HRC ACQ)	
CCDAMP	CCD amplifier readout configuration
CCDGAIN	Commanded gain of CCD (electrons/DN)
CCDOFSTA	Commanded CCD bias offset for amplifier A
CCDOFSTB	Commanded CCD bias offset for amplifier B
CCDOFSTC	Commanded CCD bias offset for amplifier C
CCDOFSTD	Commanded CCD bias offset for amplifier D
JWROTYPE	ASIC WFC readout type (DS_int for dual-slope integration or CLAMP for clamp-and-sample)
CALIBRATED ENGINEERING PARAMETERS (WFC, HRC, HRC ACQ)	
ATODGNA	Calibrated gain for amplifier A (electrons/DN)
ATODGNB	Calibrated gain for amplifier B (electrons/DN)
ATODGNC	Calibrated gain for amplifier C (electrons/DN)
ATODGND	Calibrated gain for amplifier D (electrons/DN)
READNSEA	Calibrated read noise for amplifier A (electrons)
READNSEB	Calibrated read noise for amplifier B (electrons)

READNSEC	Calibrated read noise for amplifier C (electrons)
READNSED	Calibrated read noise for amplifier D (electrons)
BIASLEVA	Bias level for amplifier A (electrons)
BIASLEVB	Bias level for amplifier B (electrons)
BIASLEVC	Bias level for amplifier C (electrons)
BIASLEVD	Bias level for amplifier D (electrons)
TARGET ACQUISITION PARAMETERS (HRC ACQ only)	
AUTHOR	Headerlet created by this user
CATALOG	Astrometric catalog used for headerlet solution
DESCRIP	Short description of headerlet
DESTIM	Destination observation rootname
HDRNAME	Headerlet name
NMATCH	Number of sources used for headerlet solution
RMS_DEC	RMS in declination at reference pixel of headerlet solution
RMS_RA	RMS in right ascension at reference pixel of headerlet solution
ASSOCIATION KEYWORDS (all detectors)	
ASN_ID	Unique identifier assigned to association
ASN_TAB	Name of the association table
ASN_MTYP	Role of the member in the association
CRDS_CTX	Master calibration reference system file
CRDS_VER	Calibration reference system version number
CTE CORRECTION KEYWORDS (WFC only)	
CTE_NAME	CTE algorithm name
CTE_VER	CTE algorithm version
CTEDATE0	Date of ACS installation
PCTETRSH	CTE over subtraction threshold
CTEDATE1	Date of CTE model pinning
PCTEFRAC	CTE scaling factor
PCTELEN	Maximum length of CTE trail

PCTERNOI	Read noise amplifier clip limit
PCTENFOR	Number of iterations in forward model
PCTENPAR	Number of iterations in parallel transfer
FIXROCR	Fix readout cosmic rays
ASTRODRIZZLE KEYWORDS (all detectors)	
UPWCSVER	Version of STWCS used to update the WCS
PYWCSVER	Version of PYWCS used to update the WCS
DISTNAME	Names of all distortion model components used
SIPNAME	Specific polynomial distortion model
RULESVER	Version ID for header keyword rules file
BLENDVER	Version of blendheader software used
RULEFILE	Name of header keyword rules file
NDRIZNUM	Number of images drizzled onto output
D???OUUN	Units of output image (counts or counts/second)
D???MASK	Input weighting image
D???FVAL	Fill value for zero weight output pixels
D???KERN	Form of weight distribution kernel
D???ISCL	Default IDCTAB pixel size (arcsec)
D???COEF	Source of coefficients
D???SCAL	Pixel size of output image (arcsec)
D???GEOM	Source of geometric information
D???DATA	Input image data
D???OUWE	Output weighting image
D???OUCO	Output context image
D???DECP	Input image exposure time (seconds)
D???WKEY	Input image WCS version used
D???OUDA	Output data image
D???VER	Drizzle task version number
D???WTSC	Weighting factor for input image

D???PIXF	Linear size of drop
DEPRECATED KEYWORDS (all detectors, in alphabetical order)	
ACQTYPE	Type of acquisition
BIASLEV	CCD bias level used to process acquisition exposure
CCDBIASS	CCD bias subtracted from target acquisition image (yes/no)
CENTMETH	Target acquisition centering method
CHECKBOX	Size of checkbox for finding algorithm
CRELIM	Perform cosmic ray rejection in acquisition

Table 2.7: ACS extension header keywords (SCI, ERR, and DQ FITS extensions).

KEYWORD	DESCRIPTION
ASSOCIATION KEYWORDS (all detectors)	
XTENSION	Extension type
NAXIS1	Length of first data axis
NAXIS2	Length of second data axis
PCOUNT	Required keyword; must = 0
GCOUNT	Required keyword; must = 1
EXTNAME	Extension name
EXTVER	Extension version number
INHERIT	Inherit the primary header
EXPNAME	Exposure identifier
DATAMIN	Minimum data value
DATAMAX	Maximum data value
BUNIT	Brightness units (counts, electrons, electrons /second)
NPIX1	Length of constant array axis 1
NPIX2	Length of constant array axis 2
PIXVALUE	Values of pixels in constant array
BSCALE	Scale factor for array value to physical value
BZERO	Physical value for an array value of zero

WFC CCD CHIP IDENTIFICATION (WFC only)	
CCDCHIP	CCD chip (1 or 2)
WORLD COORDINATE SYSTEM AND RELATED PARAMETERS (all detectors)	
WCSAXES	Number of World Coordinate System axes
CRPIX1	x-coordinate of reference pixel
CRPIX2	y-coordinate of reference pixel
CRVAL1	First axis value at reference pixel
CRVAL2	Second axis value at reference pixel
CTYPE1	The coordinate type for the first axis
CTYPE2	The coordinate type for the second axis
CD1_1	Partial of first axis coordinate with respect to x
CD1_2	Partial of first axis coordinate with respect to y
CD2_1	Partial of second axis coordinate with respect to x
CD2_2	Partial of second axis coordinate with respect to y
LTV1	Offset in x to subsection start (x-coordinate of the first science pixel in the SCI array)
LTV2	Offset in y to subsection start (y-coordinate of the first science pixel in the SCI array)
RAW_LTV1	Original offset in x to subsection start (x-coordinate of the first science pixel in the SCI array)
RAW_LTV2	Original offset in y to subsection start (y-coordinate of the first science pixel in the SCI array)
LTM1_1	Reciprocal of sampling rate in x (the reference pixel size in units of current pixel size)
LTM2_2	Reciprocal of sampling rate in y (the reference pixel size in units of current pixel size)

ORIENTAT	Position angle of image y axis (degrees E of N)
RA_APER	Right ascension of aperture reference position
DEC_APER	Declination of aperture reference position
PA_APER	Position angle of reference aperture center (deg)
VAFCTOR	Velocity aberration plate scale factor
READOUT DEFINITION PARAMETERS (all detectors)	
CENTERA1	Subarray axis 1 centerpoint in unbinned detector pixels
CENTERA2	Subarray axis 2 centerpoint in unbinned detector pixels
SIZAXIS1	Subarray axis 1 size in unbinned detector pixels
SIZAXIS2	Subarray axis 2 size in unbinned detector pixels
BINAXIS1	Axis 1 data bin size in unbinned detector pixels
BINAXIS2	Axis 2 data bin size in unbinned detector pixels
PHOTOMETRY KEYWORDS (all detectors)	
PHOTMODE	Observation configuration
PHOTFLAM	Inverse sensitivity ($\text{erg}/\text{cm}^2/\text{\AA}/\text{electron}$)
PHOTZPT	ST magnitude zero points
PHOTPLAM	Pivot wavelength (Angstroms)
PHOTBW	RMS bandwidth of filter plus detector
DATA PACKET INFORMATION (all detectors)	
FILLCNT	Number of segments containing fill value for which data were lost in transmission
ERRCNT	Number of segments containing errors

PODPSFF	PODPS fill present (T/F) indicating if there are missing packets in the data stream. Missing packets are replaced with fill packets by OPUS to facilitate pipeline processing.
STDCFF	Science telemetry fill present (T/F) indicating if there were errors found in the science dataset at the Space Telescope Data Capture Facility.
STDCFFP	Science telemetry fill pattern (hex)
ENGINEERING PARAMETERS (SBC only)	
GLOBRATE	Global count rate
GLOBLIM	Was global linearity level exceeded?
MDECODT1	MAMA JMDECODT pwd temperature from eng snap 1 (deg C; temperature at the start of observation)
MDECODT2	MAMA JMDECODT pwd temperature from eng snap 2 (deg C; temperature at the end of observation)
REPEATED EXPOSURES INFORMATION (WFC, HRC, HRC ACQ)	
NCOMBINE	Number of image sets combined during cosmic ray rejection
ONBOARD COMPRESSION INFORMATION (WFC only)	
WFCMPRSD	Was WFC data compressed? (T/F)
CBLKSIZ	Size of compression block in 2 byte words
LOSTPIX	Number of pixels lost due to buffer overflow
COMPTYP	Compression type performed (Partial/Full /None)
IMAGE STATISTICS AND DATA QUALITY FLAGS (all detectors except where noted)	
NGOODPIX	Number of good pixels
SDQFLAGS	Serious data quality flags (bad pixels that are not included in calculations for header keyword image statistics, e.g., GOODMEAN, NGOODPIX, etc.)
GOODMIN	Minimum value of good pixels
GOODMAX	Maximum value of good pixels

GOODMEAN	Mean value of good pixels
SOFTERRS	Number of soft error pixels (DQF=1)
SNRMIN	Minimum signal to noise of good pixels
SNRMAX	Maximum signal to noise of good pixels
SNRMEAN	Mean value of signal to noise of good pixels
MEANDARK	Average of the dark values subtracted (electrons)
MEANBLEV	Average of all bias levels subtracted (electrons; CCDs only)
MEANFLSH	Mean number of counts in post-flash exposure (electrons; CCDs only)
MDRIZSKY	Sky value computed by AstroDrizzle
DISTORTION CORRECTION PARAMETERS (all detectors)	
A_0_2, B_0_2, A_1_1, B_1_1, A_2_0, B_2_0, A_0_3, B_0_3, A_1_2, B_1_2, A_2_1, B_2_1, A_3_0, B_3_0, A_0_4, B_0_4, A_1_3, B_1_3, A_2_2, B_2_2, A_3_1, B_3_1, A_4_0, B_4_0, A_3_2, B_3_2, B_1_4, B_2_3, B_0_5, A_5_0, A_4_1, B_4_1, B_5_0, A_1_4, A_0_5, A_2_3	Non-linear or high-order polynomial coefficients in the SIP convention describe the geometric distortion model for each image group. They are present for use by SIP-enabled code such as DS9 and for use by the pipeline software and AstroDrizzle .
A_ORDER, B_ORDER	Order of the polynomial used to describe geometric distortion corrections
IDCSCALE	Pixel scale in the IDCTAB table
IDCXREF	Reference pixel location in x as specified by the IDCTAB
IDCYREF	Reference pixel location in y as specified by the IDCTAB
IDCV2REF	Reference pixel V2 position as derived from the IDCTAB reference table
IDCV3REF	Reference pixel V3 position as derived from the IDCTAB reference table
IDCTHETA	Orientation of the detector's Y-axis relative to V3 axis, as derived from the IDCTAB reference table
OCX10	Linear distortion term without image plate scale, directly from distortion model (IDCTAB)

OCX11	Linear distortion term without image plate scale, directly from distortion model (IDCTAB)
OCY10	Linear distortion term without image plate scale, directly from distortion model (IDCTAB)
OCY11	Linear distortion term without image plate scale, directly from distortion model (IDCTAB)
D2IMEXT	Reference file used by updatewcs to create a D2IMARR extension, which holds the column or row corrections
D2IMERR1	Maximum error of NPOL correction for axis 1
D2IMDIS1	Detector-to-image correction type for axis 1
D2IMERR2	Maximum error of NPOL correction for axis 2
D2IMDIS2	Detector-to-image correction type for axis 2
D2IM1.AXIS.1, D2IM1.AXIS.2, D2IM1.EXTVER, D2IM1.NAXES, D2IM2.AXIS.1, D2IM2.AXIS.1, D2IM2.EXTVER, D2IM2.NAXES	WCS NPOL correction parameters
WCSNAMEO	WCS solution label
WCSAXESO	Number of WCS axes
CRPIX10, CRPIX20, CDELT10, CDELT20, CUNIT1, CUNIT10, CUNIT2, CUNIT20, CTYPE10, CTYPE20, CRVAL10, CRVAL20, LONPOLE, LONPOLEO, LATPOLE, LATPOLEO, CD1_10, CD1_20, CD2_10, CD2_20, CDELT1, CDELT2, RADESYS, RADESYSO	WCS solution parameters
CDERR1, CPDIS1, CDERR2, CPDIS2, CPERR1, CPERR2 DP1.AXIS.1, DP1.AXIS.2, DP1.EXTVER, DP1.NAXES, DP2.AXIS.1, DP2.AXIS.2, DP2.EXTVER, DP2.NAXES	Distortion solution parameters, if performed
NPOLEXT	NPOLFILE used, if applicable
TDD_CTB	WFC time-dependent distortion parameter (skew angle in Y-axis)
TDD_CXB	WFC time-dependent distortion parameter (scale in X-axis)
TARGET ACQUISITION KEYWORDS (HRC ACQ only)	
TG_ENAME	Target science data extension name

TG_EVER	Target science data extension version number
DEPRECATED KEYWORDS (all detectors, sorted alphabetically)	
APER1	FSW located subarray axis 1 coordinate of aperture
APER2	FSW located subarray axis 2 coordinate of aperture
APERLKA1	Axis 1 detector pixel of acquisition aperture center
APERLKA2	Axis 2 detector pixel of acquisition aperture center
AXISCORR	Axis to which the DET2IM correction is applied
BOPOFFA1	Axis 1 offset object moved off aperture (arcsec)
BOPOFFA2	Axis 2 offset object moved off aperture (arcsec)
MAXCHCNT	Counts in the brightest checkbox
OCD1_1	Partial of first axis coordinate with respect to x
OCD1_2	Partial of first axis coordinate with respect to y
OCD2_1	Partial of second axis coordinate with respect to x
OCD2_2	Partial of second axis coordinate with respect to y
OCRPIX1	x-coordinate of reference pixel
OCRPIX2	y-coordinate of reference pixel
OCRVAL1	First axis value at reference pixel
OCRVAL2	Second axis value at reference pixel
OCRTYPE1	The coordinate type for the first axis
OCRTYPE2	The coordinate type for the second axis
ONAXIS1	x-axis length
ONAXIS2	y-axis length

OORIENTA	Position angle of image y-axis (deg E of N)
TARGA1	FSW located subarray axis 1 coordinate of target
TARGA2	FSW located subarray axis 2 coordinate of target
TDDALPHA	WFC time-dependent distortion coefficient alpha
TddbBETA	WFC time-dependent distortion coefficient beta
TDD_CTA	WFC time-dependent distortion coefficient alpha
TDD_CYA	WFC time-dependent distortion coefficient alpha
TDD_CYB	WFC time-dependent distortion coefficient alpha
TDD_CXA	WFC time-dependent distortion coefficient alpha
WCSDATE	Time WCS keywords were copied

Chapter 3: ACS Calibration Pipeline

Chapter Contents

- [3.1 Mikulski Archive for Space Telescopes - MAST](#)
- [3.2 Pipeline Overview](#)
- [3.3 Structure of calacs](#)
- [3.4 calacs Processing Steps](#)
- [3.5 Manual Recalibration of ACS Data](#)

3.1 Mikulski Archive for Space Telescopes - MAST

ACS data requests are processed with the Mikulski Archive for Space Telescopes (MAST) pipeline. It provides users with calibrated data that is created using the best available reference files and latest software upgrades.

STScI's calibration pipeline consists of two main software systems: the "Data Management System" (DMS) and the "Data Archive and Distribution System" (DADS).

Raw spacecraft telemetry data from the Goddard Space Flight Center are transmitted to STScI in the form of POD files which are stored in the *HST* Archive. MAST processing begins with the POD files when DMS runs a step called *Generic Conversion* to extract and create uncalibrated "raw" data. During this step, the "Calibration Reference Database System" (CRDS), where all calibration reference files are stored, is queried to determine the most current reference files for the observation's configuration and date. DMS then calibrates the raw data using **calacs**, with the best-available reference images and tables from the CRDS. These calibrated data, along with other requested data products, are sent to the user. (See [Chapter 1](#) of the [Introduction to the HST Data Handbooks](#) for more information about retrieving data from the Archive.)

CCD dark, bias, and sink pixel reference files are frequently delivered to CRDS, or occasionally updated. Users who receive their data shortly after an observation was executed are advised to wait a few more weeks for the best-suited reference files (usually the darks) to become available in CRDS. At that point, the data could be re-retrieved from the Archive or manually recalibrated using the new reference files. To find out if reference files have been updated, compare the names of the reference files in the image header with those listed as the recommended reference files for a particular dataset in [MAST HST Data Search](#). For dated ACS products, it is often more convenient to retrieve freshly-calibrated data from the Archive before beginning data analysis. Users who choose to rerun **calacs** on raw data should consult the [ACS Reference Files web page](#) and retrieve reference files from the [CRDS web page](#) (See [Section 3.5](#) for information about running **calacs**).

3.2 Pipeline Overview

- [3.2.1 calacs: Image Calibration](#)
- [3.2.2 AstroDrizzle Processing in the Pipeline](#)
- [3.2.3 When is MAST Processing not Appropriate?](#)

Pipeline processing is carried out by two separate image processing packages: **calacs** corrects for instrumental effects to produce calibrated products. **AstroDrizzle** corrects for geometric distortion, performs cosmic ray rejection based on the individual images of the same scene, and attempts to correct for hot pixels using dithered MAST images.

3.2.1 calacs: Image Calibration

calacs controls the image calibration steps based on the type of images and/or associations:

- For CCD images, bias level removal, as well as charge transfer efficiency (CTE) corrections for WFC images, are performed on each image using the task **acsccd**, followed by the task **acscte**.
- If the association has CCD images created from "CR-SPLIT" observations, or from repeated non-dithered exposures (several sub-exposures per observation), the task **acsrej** is used to combine the images and reject cosmic rays.
- The task **acs2d** continues with routine image reductions; MAMA images are dark-subtracted (omitted by default) and flat-fielded. CCD images—single images and images combined with **acsrej**—are, as appropriate, dark-subtracted, post-flash-subtracted, and flat-fielded.
- SBC MAMA images in the association, created from repeated non-dithered exposures (several sub-exposures per observation), are summed using the task **acssum**.

Calibrated data products from **calacs** (with suffixes `flt.fits/flc.fits`, `crj.fits/crc.fits`, and `sfl.fits`) are in units of *electrons*.



calacs standard calibration final products have suffixes `flt.fits` and `crj.fits`. When run manually, if desired, calacs also creates temporary intermediate data products, such as those with the suffix `blv_tmp.fits`.

For WFC images, calacs produces counterpart data files that have undergone pixel-based CTE corrections. The CTE-corrected final data products have suffixes `flc.fits` and `crc.fits`, to complement `flt.fits` and `crj.fits` files, respectively. When calacs is run manually for WFC data, temporary intermediate data products have the letter "c" in the suffix to indicate that it has also been corrected for CTE, such as `blc_tmp.fits` as the counterpart for `blv_tmp.fits`.

Beyond calacs, the pipeline also produces two sets of drizzled data for WFC, with suffixes `drz.fits` and `drc.fits`. In this document, unless the context is specifically for one or the other, standard and CTE-corrected files mentioned by suffix will appear separated by a "/" for instance, `flt.fits/flc.fits`.

While intermediate steps in **calacs** make use of sky subtraction values to perform certain steps, such as in identifying cosmic rays, all data products created by the pipeline will *not* be sky subtracted.

Calibrated products from the pipeline may still contain some artifacts such as hot pixels, cosmic rays, and, in the case of post-SM4 WFC subarray images, bias striping. To correct for post-SM4 subarray bias striping, after bias subtraction but before the rest of the calibration steps, it is currently necessary to use a standalone **Python** routine (`acs_destripe_plus`) outside the **calacs** package (see [Section 4.2.1](#) and [Example 5](#) in [Section 3.5.2](#)). Residual hot pixels and cosmic rays may be rejected from dithered images using **AstroDrizzle** to process associations created from observations taken with the "POS TARG" or dither "PATTERN" special requirements in Phase II proposals.

calacs and Single Exposures

Each single-exposure raw image undergoes standard detector calibrations in **calacs**, such as bias subtraction, dark subtraction, and flat-fielding (see [Section 3.3](#)) to create a `flt.fits` image.

For full frame WFC images,¹ by default, a CTE-corrected image with the suffix `flc.fits` is also created. This is done regardless of whether those single images will be combined in later **calacs** steps. Data in the "SCI" (science image) and "ERR" (error image) extensions of a calibrated `flt.fits/flc.fits` image are in units of *electrons*, whereas the raw ACS images are in units of *counts*.

calacs and Combining of Sub-Exposures

Depending on how multiple sub-exposures were executed, **calacs** has two different ways to combine them.

1. If CCD images are flagged in an association table as belonging to a "CR-SPLIT" or repeated observations set,² the following steps are performed by **calacs**:
 - Bias subtraction, dark subtraction, and flat-fielding are performed on each raw image. For WFC images with PTECORR set to PERFORM, CTE correction is performed prior to dark subtraction.
 - Images are combined with cosmic ray rejection.
 - The combined image is flat-fielded to create a calibrated image file with suffix `crj.fits`. For images where CTE-correction is applicable, a CTE-corrected combined image with the suffix `crc.fits` is also created. A single fully-calibrated MAMA image is given the suffix `flt.fits`. (MAMA images do not have an overscan region, and they are not affected by cosmic rays and CTE.)
2. If SBC MAMA images are flagged in an association table as belonging to a set of repeated sub-exposures, **calacs** takes the following actions:
 - Each image is fully calibrated and flat-fielded to produce `flt.fits` files.
 - The `flt.fits` images are summed to create an image with the suffix `sfl.fits`.

Note that each single exposure image from a "CR-SPLIT" or repeated sub-exposures set will also be calibrated individually to produce a `flt.fits/flc.fits` image for later use in **AstroDrizzle** if the header value `EXPSCORR="PERFORM"` (which is currently the default).

calacs and Dithered Exposures

calacs produces a calibrated `flt.fits/flc.fits` file for each single-exposure image in an association, including those created from using dither "PATTERN" and "POS TARG" special requirements in the Phase II proposal.

If there were two or more repeated sub-exposures at a pointing, **calacs** produces a cosmic ray-rejected combined image, `crj.fits/crc.fits`, for CCD data. For SBC MAMA data, a summed image is created with the suffix `sfl.fits`.

However, **calacs** will not combine images from multiple positions within an association (like those from a dither pattern). Later in the pipeline, after **calacs** processing is completed, `flt.fits/flc.fits` images will be corrected for geometric distortion and combined, with cosmic ray and hot pixel removal, by **AstroDrizzle** (`crj.fits/crc.fits` and `sfl.fits` files are not used in **AstroDrizzle**).

Table 3.1: Input and Output Image Suffixes from calacs and AstroDrizzle for Various Observing Modes

Image Type		Image Suffixes (suffix.fits)					
		calacs Input	calacs Output	calacs Cosmic Ray Rejected?	AstroDrizzle Input	AstroDrizzle Output	AstroDrizzle Cosmic Ray Rejected?
Single		raw	flt/flc ^a	No	flt/flc	drz/drc	No
Repeated Observations	CCD	asn or raw	flt & crj flc & crc	Yes Yes	asn & flt, or flt asn & flc, or flc	drz/drc	Yes
	MAMA	asn or raw	flt & sfl	n/a ^b	asn & flt, or flt	drz	n/a ^b
CR-SPLIT		asn or raw	flt & crj flc & crc	Yes	asn & flt, or flt asn & flc, or flc	drz/drc	Yes
Dither PATTERN or POS TARG		asn or raw	flt & crj, flc & crc, flt & sfl (SBC only)	Maybe ^c	asn & flt, or flt asn & flc, or flc	drz/drc	Yes

^a CTE-corrected products from MAST only apply to full frame WFC images. Users wishing to correct old format (pre-Cycle 24, October 2016) 2K subarray images and any new format subarray images for CTE effects may do so using **acs_destripe_plus** in the **acstools** Python package. Other old format subarray sizes are not compatible with the pixel-based CTE correction.

^b SBC MAMA detectors are not sensitive to cosmic rays.

^c Depends on the image type. For "CR-SPLIT" exposures, **calacs** creates `crj.fits/crc.fits` combined images. For repeated MAMA exposures, **calacs** creates a summed `sfl.fits` file. However, combined images are not used as input to **AstroDrizzle**. Only `flt.fits/flc.fits` files are the primary input to **AstroDrizzle**; they can also be represented by an association table, if one is available.

3.2.2 AstroDrizzle Processing in the Pipeline

During pipeline processing, calibrated ACS data that belong to an association are corrected for geometric distortion and drizzle-combined with cosmic ray rejection by **AstroDrizzle**. If the associated images are dithered, they are aligned using the WCS information in their headers before being drizzle-combined. If there is no association table, each single-exposure ACS image is drizzled to correct for geometric distortion.

The resulting drizzled image, in units³ of *electrons/second*, is written to a file with the suffix `drz.fits/drc.fits`. (For WFC, data from the two chips are mosaicked together as one image.)

In the pipeline, **AstroDrizzle** and its related software rely on these reference files:

- `IDCTAB` reference table for a description of the distortion model.
- `D2IMFILE` reference file for filter independent detector pixel grid defects or irregularities in X,Y in each WFC CCD chip, only for WFC images.
- `NPOLFILE` reference file is for the non-polynomial filter dependent part of distortion, for residual distortions not accounted for by the `IDCTAB` distortion solution coefficients (nor corrected by the `D2IMFILE`, in the case of WFC images).

Information about geometric distortion from these reference files are stored as SIP header keywords and as FITS extensions in the `flt.fits/flc.fits` images. Please see [ACS ISR 15-06](#) for more information on the distortion and [Section 2.2](#) for details on the file structure.

The resulting drizzled images from the pipeline may be useful for science as-is, although subsequent manual reprocessing with **AstroDrizzle** is recommended, and sometimes required, for optimizing the data. For more information, please refer to the [DrizzlePac website](#).

3.2.3 When is MAST Processing *not* Appropriate?

The goal of the ACS pipeline is to provide data calibrated to a level suitable for initial evaluation and analysis for all users. Observers require a detailed understanding of the calibrations applied to their data and the ability to repeat, often with improved products, the calibration process at their home institutions. There are several occasions when data processed via MAST from the Archive are not ideal, requiring off-line interactive processing:

- Running **calacs** with different reference files than those specified in the image header.
- Running **calacs** with non-default calibration switch values.
- When images must be cleaned of artifacts such as hot pixels, cosmic rays, and residual artifacts such as bias striping.

Images combined by **AstroDrizzle** in the pipeline were produced using parameters that are suitable for the widest range of scientific applications (see [ACS ISR 17-02](#) for more details). Some datasets, however, could benefit significantly from manual reprocessing, for instance, by using a different pixel scale or by modifying cosmic ray rejection parameters. The same **AstroDrizzle** task used in the pipeline is also available to users in **AstroConda** for off-line processing of `flt.fits/flc.fits` images retrieved from the Archive. For more information, please refer to [DrizzlePac website](#).

¹ Since 2014, users can apply CTE corrections to old format (pre-Cycle 24, October 2016) WFC 2K subarray images and all new format subarrays using the **acs_destripe_plus** tool in **acstools**. Only full-frame images are CTE-corrected in the archive.

² The Phase II proposal's exposure log sheet line parameter "Number_of_Interations" has an integer value greater than 1.

³ The final drizzled image's unit type is set in the **AstroDrizzle** task parameter *final_units*; the choices are *cps* (counts per second, the default value) or *counts*. The unit for *counts* is specified in the image header keyword `BUNIT`. For ACS images, `BUNIT` is set to `ELECTRONS`. Therefore, ACS drizzled images are, by default, in units of electrons/second.

3.3 Structure of calacs

The **calacs** package consists of five tasks listed in [Table 3.2](#). These tasks, available in the [HSTCAL](#) package, are automatically called by calacs, but each may be run separately, via [acstools](#) in [AstroConda](#).

Table 3.2: Tasks in the calacs Pipeline

acscdd	CCD specific calibrations, except CTE corrections
acscte	CTE corrections for WFC images ¹
acsrej	Cosmic ray rejection task
acs2d	Basic MAMA and CCD calibrations
acssum	Repeated observations summing task for MAMA data

¹ CTE loss correction is available for a WFC 2K subarray of the old format or any WFC subarray of the new format. The subarray formats were switched from old to new during Cycle 24, in October 2016. However, due to uncorrected striping in all WFC subarrays, PCTECORR is set to OMIT by default in the pipeline for all subarrays. To perform CTE loss correction for an eligible WFC subarray, please set PCTECORR to PERFORM in the image header and use the '**acs_destripe_plus**' task in **acstools**. See also [Example 5](#), in [Section 3.5.2](#).

The flow of data through the ACS calibration pipeline and the decisions made while working with associated data are shown in [Figures 3.1, 3.2, and 3.3](#). They're also outlined below with the **calacs** tasks and functions in parenthesis.

1. Flag known bad pixels and A-to-D saturated pixels in the data quality (DQ) array. (**acscdd/doDQI** or **acs2d/doDQI**)
2. Subtract the bias image (CCD only). (**acscdd/doBias**)
3. Multiply by gain to convert DN to electrons. (**acscdd/toElectrons** or **acs2d/toElectrons**)
4. Subtract the bias level, which is determined from overscan regions (CCD only). For post-SM4 full frame WFC images, also correct for bias shift, cross-talk, and striping. (**acscdd/doBlev**)
5. Flag full-well saturated pixels in the data quality (DQ) array. (**acscdd/doFullWellSat**)
6. Flag sink pixels in the DQ array of WFC images. (**acscdd/doSink**)
7. Calculate a noise model for each pixel and record it in the error (ERR) array. (**acscdd/doNoise** or **acs2d/doNoise**)
8. Perform pixel-based CTE corrections for applicable WFC images. (**acscte**)
9. Combine images, with cosmic ray rejection, for "CR-SPLIT" and repeated exposures (CCD only). (**acsrej**)
10. Perform global linearity corrections (MAMA only). (**acs2d/doNonLin**)
11. Scale and subtract the dark image, and calculate the mean dark value (CCD only, by default). (**acs2d/doDark**)
12. Scale and subtract the post-flash image, if required (CCD only). (**acs2d/doFlash**)
13. Divide the image by the flat field. (**acs2d/doFlat**)
14. Apply shutter shading correction (CCD only). (**acs2d/doShad**)
15. Calculate photometry header keywords for flux conversion (except for slitless spectroscopy modes). (**acs2d/doPhot**)
16. Calculate image statistics; these values are stored in calibrated data headers. (**acs2d/doStat**)
17. Sum images from repeated sub-exposures in an ACS/SBC exposure. (**acssum**)

Calibrated `flt.fits/flc.fits` images from "CR-SPLIT" exposures, repeated sub-exposures, "POS TARG" exposures, or dither "PATTERN" exposures may be combined using **AstroDrizzle**. Please refer to the [DrizzlePac website](#) for information regarding drizzling the images.

As indicated in [Figure 3.1](#), calibration tasks that are detector-specific (like **acscdd** for WFC data only) have been separated from tasks that can process both detectors (such as **acs2d**).

The initial processing performed on CCD data alone is shown in [Figure 3.2](#). Reference files appropriate for each processing step and the calibration switches controlling them are also given beside the name of the task they control. The output (overscan-trimmed image) from **acscdd** is then sent through **acscte** (if appropriate) and **acs2d** as shown in [Figure 3.3](#).

Processing of raw MAMA data begins with **acs2d**, which initializes the error and data quality arrays (a step that was performed earlier for CCD data) and applies linearity corrections.

Figure 3.1: Flow Diagram for ACS Data, With calacs Task Names


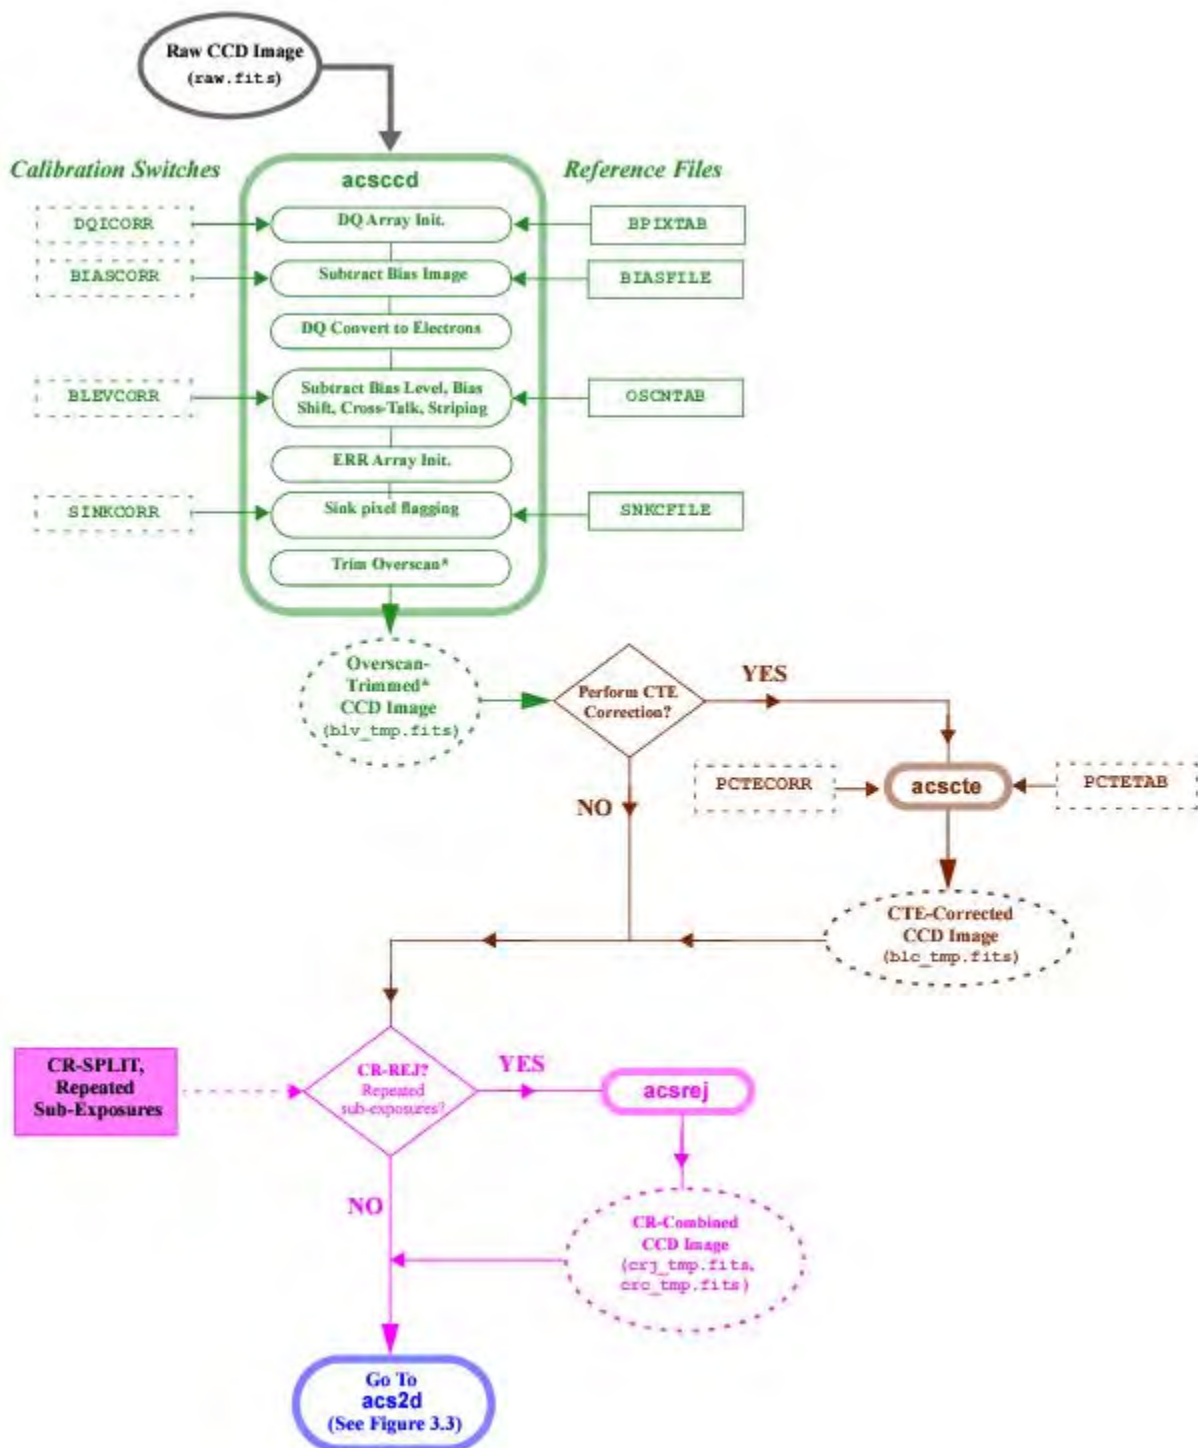
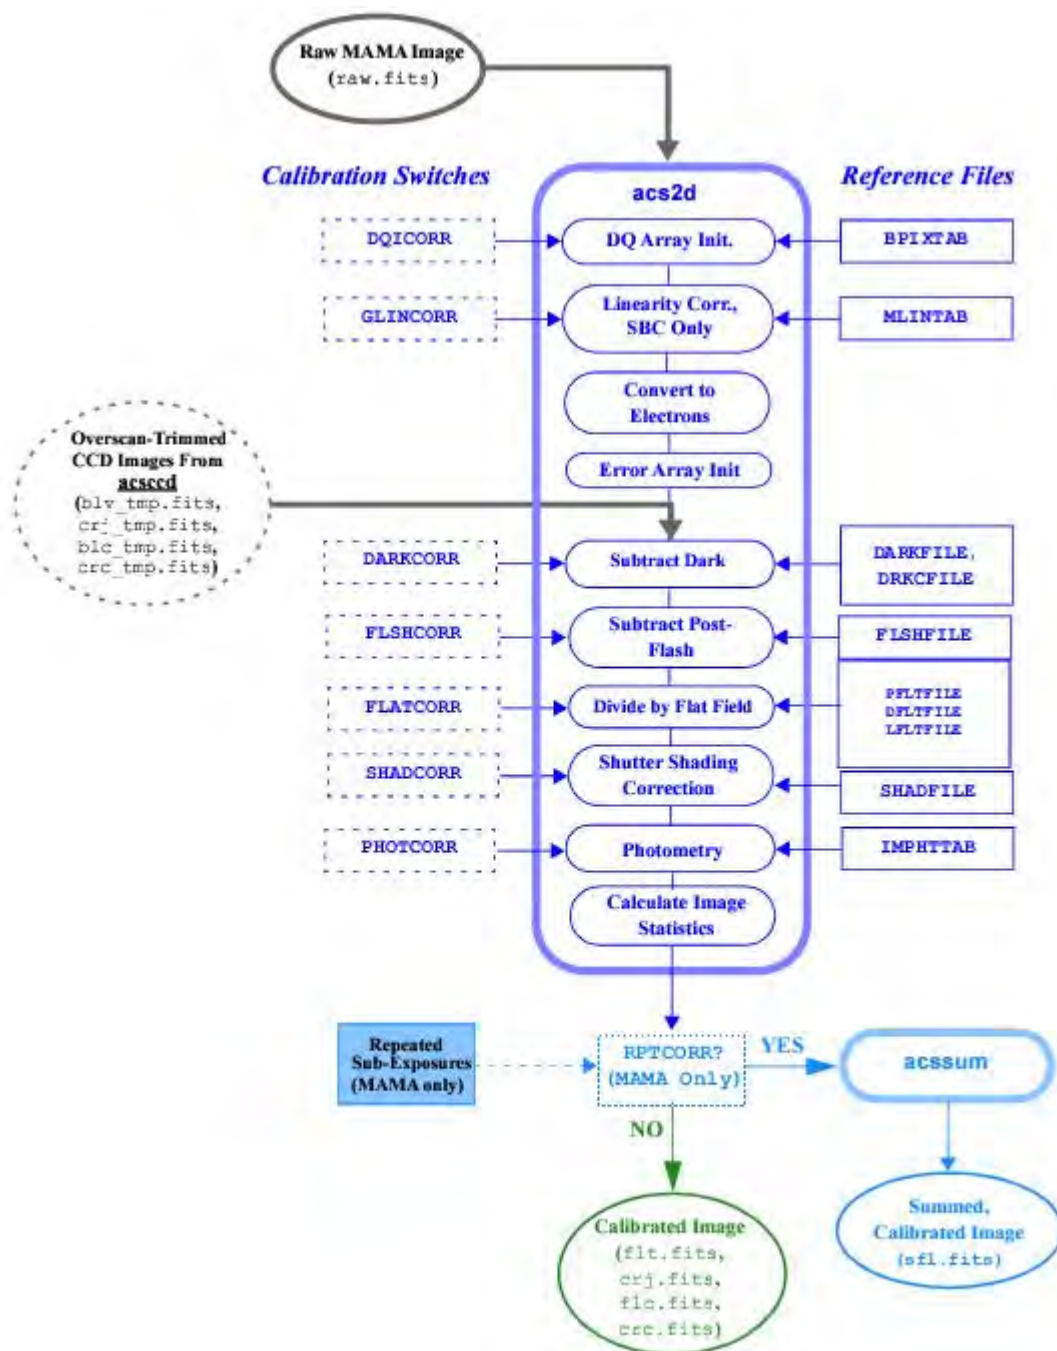


Figure 3.2: Flow Diagram for CCD Data Using `acscdd` and `acscte` in `calacs`



* The overscan areas are trimmed only if the raw image had "BLEVCORR=PERFORM." In almost all cases, this is the default value. Only in rare instances will that keyword be set to "OMIT."

Figure 3.3: Flow Diagram for MAMA and Overscan-Trimmed CCD Data Using `acs2d` in `calacs`



3.4 calacs Processing Steps

- 3.4.1 [acscdd](#)
- 3.4.2 [acscte](#) – Pixel-Based CTE Corrections
- 3.4.3 [acsrej](#)
- 3.4.4 [acs2d](#)
- 3.4.5 [acssum](#)

As mentioned earlier, the **calacs** pipeline consists of five individual calibration tasks: **acscdd**, **acscte**, **acsrej**, **acs2d**, and **acssum**. These tasks are listed in [Table 3.2](#) and diagrammed in [Figure 3.1](#). **calacs** is responsible for controlling the processing rather than the actual calibration of the data. The individual tasks within **calacs** apply the appropriate calibration steps to the data and create the output products.

Raw ACS images are in units of Data Numbers (DN). Calibrated images are in units of electrons, and drizzled images are in units of electrons/second.

The following five sections provide a detailed description of the calibration steps, and their related reference files, within each **calacs** task.

3.4.1 acscdd

This task contains the initial processing functions for all ACS CCD data. These functions are listed, in operational order, in [Table 3.3](#).

acscdd processes each image in the input list, one at a time, using the header keywords to determine which calibration steps to perform. Upon completion of **acscdd**, the overscan regions will be trimmed from the image and an output image with the file suffix "blv_tmp.fits" is created.

Table 3.3: acscdd Processing Functions

<i>doDQI</i>	Initialize data quality array
<i>doBias</i>	Subtract bias image to remove low order quasi-static structure of the bias
<i>toElectrons</i>	Convert from DN to electrons
<i>doBlev</i>	Subtract bias level as determined from pre-scan
<i>doFullWellSat</i>	Flag full-well saturated pixels
<i>doNoise</i>	Initialize error array
<i>doSink</i>	Sink pixels flagging
<i>Final Output</i>	Output overscan-trimmed or full image

doDQI - Bad Pixel Determination

- Header switch: `DQICORR`
- Header keywords updated: None

- Reference file keyword (file type, suffix): `BPIXTAB` (FITS table, `bpx.fits`), `CCDTAB` (FITS table, `ccd.fits`)

The function ***doDQI*** initializes the Data Quality (DQ) array by combining it with a table of known permanent bad pixels for the detector, stored in the bad pixel reference table (named by the image header keyword `BPIXTAB`). The type of bad pixel flags are listed in [Table 3.4](#).

The DQ array may have already been partially populated to reflect pixels which were affected by telemetry problems or compression buffer overflow. Other DQ flags will be added in further processing steps (such as cosmic ray rejection and sink pixel flagging). For CCD data, the values in the SCI extension are also checked for A-to-D saturation. See [ACS ISR 2020-03](#) for more.

For each observation, the ***doDQI*** function combines the DQ flags from prior calibration processing (obtained from the `BPIXTAB` reference table) and from saturation tests, into a single DQ result. These values are combined using a "bitwise OR operator" for each pixel. Thus, if a single pixel is affected by two DQ flags, the sum of those flag values is assigned to the corresponding pixel in the final DQ array. This array then serves as a mask so that ***calacs*** will ignore bad pixels during processing.

Table 3.4: Flags for the DQ Array

Flag Value	Definition
0	good pixel
1	Reed-Solomon decoding error
2	data replaced by fill value
4	bad detector pixel or vignetted pixel
8	masked by aperture feature
16	hot pixel (dark current > 0.14 e-/sec)
32	Pixels determined to be unstable throughout the course of an anneal cycle
64	warm pixel (dark current: 0.02–0.08 e-/sec before SM4; 0.04 to 0.08 e-/sec after SM4; 0.06–0.14 e-/sec Jan 15, 2015 to present)
128	bias structure (mostly bad columns)
256	saturation (full well saturation)
512	bad pixel in reference file
1024	Sink pixel or pixel affected by sink pixel charge traps
2048	A-to-D saturation
4096	cosmic ray rejected by AstroDrizzle (based on all <code>flt.fits/flc.fits</code> files)
8192	cosmic ray rejected by acsrej
16384	reserved (satellite trail masks, etc.)

i A more detailed description of the Data Quality flags used in ACS data products, and their history, can be found on the Web page:
http://www.stsci.edu/hst/acs/analysis/reference_files/data_quality_flags.html.

doBias - Bias Image Subtraction

- Header switch: `BIASCORR`
- Header keywords updated: None
- Reference file keyword (file type, suffix): `BIASFILE` (image, `bia.fits`)

doBias removes the low order, quasi-static structure of the bias, including the bias gradient in post-SM4 images.

Subtraction of the bias image is performed prior to cosmic ray rejection using the function **doBias**. The default reference image, named in the header keyword `BIASFILE`, is assigned based on the `DETECTOR`, `APERTURE` (for post-SM4 WFC images), `CCDAMP`, and `CCDGAIN` image header keyword values. The dimensions of the science image are used to distinguish observations which use the entire chip from subarray images.

The bias reference file data is in units of DN. Therefore, it is applied to the image before the data is converted to units of electrons. Since the overscan values were already subtracted from the bias reference image, it may have a mean pixel value of less than 1.

Due to the way the bias reference image is created, part of the dark subtraction is also included in this step. Dark counts accumulate for an additional time beyond the exposure time, primarily the time required to read out the detector, and this portion of the dark current is subtracted along with the bias. This is described in the section on [doDark - Dark Image Subtraction](#).

The bias reference files are updated regularly (for more information, see the [ACS website](#)). For full frame images, the `BIASFILE` has the same dimensions as a full-size science image, 1062×1044 for HRC and 4144×2068 for WFC, allowing for simple image arithmetic on arrays of equal size. For WFC subarray images in the post-SM4 Cycles 17–23, aperture-specific `BIASFILE`s are used to account for different bias gradients, but the concept of simple image arithmetic remains.

Only after the completion of **acsccd** are the images trimmed to 1024×1024 (HRC) and 4096×2048 (WFC). ACS CCD data are not binned, so any image which is not the full size is assumed to be a subarray.

For WFC subarray images in *HST* Cycles other than 17–23, **calacs** uses the `LTV[1, 2]` keywords to extract the appropriate region from the reference file and apply it to the subarray input image. For *HST* Cycles 17–23, the low order bias structure in WFC subarray images differs significantly from that in full frame WFC images. Therefore, bias calibration images are available for supported subarray configurations.

For users who make their own bias files, the `BIASFILE` keyword will need to be updated manually before recalibrating. This is described in detail in [Section 3.5.2](#).

toElectrons – Unit Conversion From DN to Electrons

- Header switch: None
- Header keyword updated: `BUNIT`
- Reference file keyword (file type, suffix): `CCDTAB` (FITS table, `ccd.fits`)

Data is multiplied by gain for respective amplifiers. `BUNIT` is changed from DN to electrons in SCI and ERR extensions.

***doBlev* – Bias level Correction**

- Header switch: `BLEV CORR`
- Header keywords updated: `BIASLEV[A,B,C,D]`, `MEANBLEV`, `CRPIX[1,2]`, `LTV[1,2]`
- Reference file keyword (file type, suffix): `OSCNTAB` (reference table, `osc.fits`)

There are two branches in the ***doBlev*** function, one for pre-SM4 data and another for post-SM4 data.

Pre-SM4 images

doBlev fits the bias level from the physical pre-scan¹ section (defined in the `OSCNTAB` reference table), and subtracts it from the image data.

A subset of the pixels along each row of the vertical pre-scan is used to obtain a linear fit of the bias level as a function of image row. If the measured pre-scan level for a given row is more than 3σ from the mean of all rows, that row will be excluded from the linear fit. The fitted value for each pre-scan row is then subtracted from the image pixels for that row. If the pre-scan level cannot be determined, or if the pre-scan region is not present in the input image, a default value gets derived from the CCD reference table (given by the keyword `CCDTAB`), with the value coming from the "CCDBIAS" column. That value will be subtracted and a warning message will be written to the image trailer file.

The positional image header keywords `CRPIX1`, `CRPIX2` and `LTV1`, `LTV2` are updated to reflect the offset due to removal of the overscan. In addition, the mean value of all overscan levels is computed and written to the SCI extension header keyword `MEANBLEV`. The individual bias level measured for each amplifier is written to the SCI header keywords `BIASLEVA`, `BIASLEVB`, `BIASLEVC`, and `BIASLEVD`.

The "Overscan Region Table" reference file, named in the header keyword `OSCNTAB`, describes the overscan regions to be used for measuring the bias level of the observation. Each row corresponds to a specific configuration, given by the CCD amplifier, chip, binning ratio, and dimensions. The columns `TRIMX*`² give the number of columns to trim off the beginning and end of each line (the physical overscan region), while the `TRIMY*` columns give the number of rows to trim off the top and bottom of each column (the virtual overscan region). The result of trimming (`TRIMX1 + TRIMX2`) columns and (`TRIMY1 + TRIMY2`) rows gives the final calibrated image sizes, 4096×2048 for a full WFC image and 1024×1024 for a full HRC image.

The `OSCNTAB` columns `BIASSECTA*`² and `BIASSECTB*` give the range of image columns to be used for determining the bias level in the leading and trailing regions, respectively, of the physical overscan region.

The `VX*` and `VY*` columns in the `OSCNTAB` define the virtual overscan region, but are not used in calibration.

To determine which overscan regions were actually used for determining the bias level, users are encouraged to check the `OSCNTAB` reference file. If desired for manual calibration, users may modify the bias section and virtual overscan region definitions in the reference table, but the `TRIMX*`, `TRIMY*` columns must not be changed.

Post-SM4 Full Frame WFC images

For all WFC data except subarrays, the following corrections are done by ***doBlev***:

First, a *bias shift correction* is done to remove an artifact associated with the CCDs' external pre-amplifiers and the dual-slope integrators within the CCD Electronics Box Replacement. This effect is signal-dependent, causing the pixel-to-pixel bias level to be offset by 0.02–0.3%, with an offset that decays slowly in the serial (horizontal) clocking direction. Detailed information about the bias shift can be found in [ACS ISR 2012-02](#).

The effects of *cross-talk*, negative ghost images that are mirror-symmetric to bright sources in adjacent quadrants, are also removed. To learn more about cross-talk, please refer to [ACS ISR 2010-02](#).

Corrections are simultaneously made for two other artifacts: *amplifier offsets*, slight bias level offsets in each amplifier, and bias stripes, low amplitude, horizontal striping caused by electronic $1/f$ noise by the replacement electronics (see Grogin et al., [ACS ISR 2011-05](#)).

Bias Striping Noise

For post-SM4 full frame WFC images, the bias level subtraction is more complicated, as this step also removes the $1/f$ bias striping noise. To overcome read noise fluctuations, all of the pre-scan pixels are used, not just those in columns 20–24 where the bias level has stabilized.

The procedure first computes a robust average bias using all of the pixels in columns 20–24 (inclusive) in each amplifier's pre-scan region. This constant is subtracted from each amplifier's pixels.

The next step is to remove the bias settling feature at the beginning of the pre-scan row, that occurs after the parallel shift. This trend is different for each amplifier, but is very stable and can be removed with high fidelity by subtracting, column by column, the robust average of the 2048 pixels in each pre-scan column.

With all the pre-scan pixels normalized in this way, the next step is estimation and removal of the bias striping noise. This noise is highly consistent along rows, and is also highly consistent among the four amplifiers. For each of the 2048 rows, the 24 pixels from each of the four amplifiers' pre-scans are combined for a total of 96 estimates of the systematic $1/f$ noise residual. The robust average of these 96 pixels is then subtracted from the corresponding row of all four amplifier quadrants. The striping correction has a read noise-driven uncertainty of 0.4 e⁻, compared with the highly stable 0.9 e⁻ standard deviation of the striping noise. A more accurate removal of the striping noise may be possible outside **calacs** by using the entire image region rather than only the pre-scan—please refer to the ACS Web page about [de-striping](#) for additional information, as well as the documentation for [acstools](#), and see [Example 5](#) in [Section 3.5.2](#).

In summary, the net bias level correction for each science pixel in the post-SM4 WFC full frame images is the sum of the average bias from that amplifier plus an additional correction for the striping noise in the pixel's particular row. This pre-scan-based striping correction is only possible for full frame mode; subarray modes contain too few pre-scan pixels in each row to estimate the striping correction with sufficient precision.

doFullWellSat - Full-well saturated pixels flagging

- Header switch: None
- Header keywords updated: None
- Reference file keyword (file type, suffix): SATUFILE (FITS image, `sat.fits`).

doFullWellSat flags full-well saturated pixels with the value 256 in the DQ array of WFC and HRC images. It does this by comparing the science image pixel values with the values in the full-well saturation level pixel map. See [ACS ISR 2020-02](#) for more details.

doNoise - Error Array Initialization

- Header switch: None
- Header keywords updated: None
- Reference file keyword (file type, suffix): CCDTAB (FITS table, `ccd.fits`)

Input `raw.fits` images delivered by Generic Conversion will contain a null error (ERR) array, defined by the header keywords `NPIX1`³, `NPIX2`³ and `PIXVALUE`³, where `PIXVALUE` = 0.

If this ERR array has already been expanded in prior processing and contains values other than zero, **doNoise** does nothing. Otherwise, **calacs** initializes the array and the **doNoise** function assigns a simple noise model to the ERR array.

The noise model reads the science array and calculates the error value σ (in units of electrons) for each pixel:

$$\sigma_{\text{CCD}} = \sqrt{(SCI - bias) + readnoise}^2 \quad \text{---(1a) For non-bias exposures}$$

$$\sigma_{\text{CCD}} = readnoise \quad \text{---(1b) For bias exposures}$$

$$\sigma_{\text{MAMA}} = \max(1, \sqrt{SCI}) \quad \text{---(2)}$$

The algorithms shown above are used for initializing the ERR array for the CCDs (equations 1a and 1b) and MAMA (equation 2) observations.

For CCD observations: prior to SM4, there was a negligible gradient in the *bias*, therefore a constant value was used for the *bias* variable in the noise equation. However, for post-SM4 processing, the *bias* variable now also incorporates the bias gradient as given in the corresponding bias calibration image.

Because MAMA data is not processed by **acsccd**, the **doNoise** function is later called, just for processing SBC data, as part of **acs2d**.

The "CCD Characteristics Table" reference file, recorded in the image header keyword `CCDTAB`, is used to determine the bias (for certain subarrays only), gain, and read noise values for an observation, and for use in calculating error values (σ_{CCD}) for the ERR array. The table columns are

- CCDAMP: every possible configuration of the amplifiers used for readout.
- CCDCHIP: the chip being read out.
- CCDGAIN: the commanded gain.
- CCDBIASA, CCDBIASB, CCDBIASC, CCDBIASD: the commanded bias level of the amplifiers.
- BINAXIS1, BINAXIS2: the pixel bin sizes.

The `CCDTAB` reference table uses these commanded values to determine an observation's physical readout characteristics that are written to the `flt.fits/flc.fits` calibration file. For instance, read noise values for each amplifier are populated in header keywords `READNSEA`, `READNSEB`, `READNSEC`, and `READNSED`; A-to-D gain in `ATODGNA`, `ATODGNB`, `ATODGNC`, and `ATODGND`; bias voltage offset values in `CCDOFSTA`, `CCDOFSTB`, `CCDOFSTC`, and `CCDOFSTD`.

doSink - Sink pixels flagging

- Header switch: `SINKCORR`
- Header keywords updated: None
- Reference file keyword (file type, suffix): `SNKCFIT` (FITS image, `snk.fits`).

doSink flags sink pixels and the adjacent affected pixels with the value 1024 in the DQ array of WFC images. See [Section 4.3.3](#) for a description of sink pixels. One sink pixel reference image is created for each anneal cycle after January 2015. This reference image encodes the locations of the sink pixels, trailing low pixels, and excessively high downstream pixels, as well as information relevant to flagging the correct number of trailing pixels based on the background level in a science image (see [ACS ISR 2017-01](#) for more details). The reference image in the header keyword `SNKCFIT` is assigned based on the `DETECTOR` keyword and the observation date of the science image. All WFC images taken after January 2015 will have **doSink** performed.

Final Output From `acsccd`

Upon completion of `acsccd`, the overscan regions will be trimmed from the image when it is written out, **but only if `doBlev` is performed successfully**. Otherwise, the full image array will be written out.

3.4.2 `acscte` – Pixel-Based CTE Corrections

- Header switch: `PCTECORR`
- Header keyword added or updated: `CTE_VER`, `CTE_NAME`, `CTEDATE0`, `CTEDATE1`, `PCTENFOR`, `PCTENPAR`, `PCTETLEN`, `FIXROCR`, `PCTEFRAC`, `PCTERNOI`, `PCTETRSH`, `PCTENSM`
- Reference file keyword (file type, suffix): `PCTETAB` (FITS table, `cte.fits`)

For all WFC exposures, except bias images, the pixel-based CTE correction is applied between the `acsccd` and `acs2d` steps, after `BLEVCORR`. Parameters characterizing the CTE correction are stored in a reference table, `PCTETAB`, and also in SCI image headers. As long as `PCTECORR` is set to `PERFORM` and the correct reference file or parameters are provided, CTE loss correction is also available for a WFC 2K subarray of the old format or any WFC subarray of the new format. However, due to uncorrected striping in all WFC subarrays, `PCTECORR` is set to `OMIT` by default in the pipeline for all subarrays. To perform CTE loss correction for an eligible WFC subarray, please set `PCTECORR` to `PERFORM` in the image header and use the **`acs_destripe_plus`** task in `acstools` to de-stripe the data. Then CTE correction will be performed. Upon completion, a CTE-corrected version of "`blv_tmp.fits`" (created in `acsccd`), named "`blc_tmp.fits`," is created. Both files are passed to `acs2d` for further processing. (Please see [Example 5](#) in [Section 3.5.2](#) for a worked example.)

The ACS CTE-correction algorithm has recently been updated to a "generation 2" version. Due to the time-consuming nature of this algorithm, the parallel processing option is available and enabled by default. Details regarding the earlier "generation 1" correction are available in ACS ISRs [2010-03](#) and [2012-03](#). See Anderson et al. 2017, ACS ISR 2017-07 (in prep) where there is more to read about the [pixel-based CTE correction `pctecorr`](#) and more about [command line options and the details of usage](#) of the new "gen2" CTE-correction software.

3.4.3 `acsrej`

- Header switch (in `flt.fits/flc.fits`, `crj.fits/crc.fits` images): `CRCORR`
- Header keywords updated (in `crj.fits/crc.fits` image): `TEXPTIME`, `SKYSUM`, `EXPEND`, `REJ_RATE`, `EXPTIME`, `NCOMBINE`, `ROOTNAME`
- Header keywords added or updated (in `crj.fits/crc.fits` image): `INITGUES`, `SKYSUB`, `CRSIGMAS`, `MEANEXP`, `CRRADIUS`, `CRTHRESH`, `SCALENSE`, `CRMASK`, `NEXTEND`

- Reference file keyword (file type, suffix): CRREJTAB (FITS table, `crre.fits`)

acsrej, the cosmic ray rejection task in calacs, combines "CR-SPLIT" exposures and repeated sub-exposures in an exposure, into a single image, free of cosmic rays. The task uses the same statistical detection algorithm developed for STIS data (ocrrej) and WFPC2 data (crrej), providing a well-tested and robust procedure. The parameters used by acsrej are obtained from the "[Cosmic Ray Rejection Table](#)" (named in header keyword CRREJTAB), and depend on the type of chip, number of "CR-SPLIT" or repeated sub-exposure images, and exposure time of each image.

To maintain backward compatibility (in order to use existing rejection parameter values), although input data are now in electrons and not DN, it is converted back to DN within this task prior to rejection, then back to electrons.

First, acsrej will compute the sky background using the mode of each image. Sky subtraction is performed before any statistical checks are made for cosmic rays. Next, acsrej constructs an initial comparison image from each input exposure. The comparison image can either be a median- or minimum-value sky-subtracted image constructed from all the input images, and represents the "first guess" of a cosmic ray-free image. This comparison image serves as the basis for determining the statistical deviation of each pixel from the input image.

A detection threshold is then calculated for each pixel based on the comparison image. This threshold is equal to a constant times sigma squared, given in the equation below:

$$\tau_n = \sigma^2 \times (noise + value(x, y)/gain + scale (value(x, y) - sky_n)^2)/T_n^2$$

where:

- σ is the sigma value used as the detection limit,
- *noise* is the read noise squared (in DN) and *gain* is the CCD gain (in e-/DN) for the amplifier used to read the pixel,
- *scale* is the scale factor for the noise model,
- T_n *scale is the scale factor for the noise model,*
- *value* is the pixel value (in DN) from the median- or minimum- combined comparison image.

The actual cosmic ray detection criteria at each pixel is determined as:

$$\Delta = ((pix_n(x, y) - sky_n)/T_n - median(x, y))^2$$

where:

- $pix_n(x, y)$ is the pixel value (in DN) from input image n ,
- sky_n is the sky background (in DN) of image n , and
- $median_n(x, y)$ is the median or minimum pixel value (in DN/sec.) from the comparison image.

If $\Delta > \tau_n$, the pixel is flagged as a cosmic ray in the input image's DQ array and is ignored when images are summed together. Surrounding pixels within a given expansion radius are marked as "spill" pixels and are given less stringent detection thresholds.

When all input images have been processed, the values of the acceptable pixels are summed over all input images. Each pixel in the summed output array is then scaled by the total exposure time:

$$pixout(x, y) = T \times \frac{\sum_n (pix_n(x, y) - sky_n) m_n(x, y)}{\sum_n T_n m_n(x, y)} + \sum_n sky_n$$

where:

- T_n is the exposure time for image n ,
- $m_n(x,y)$ is the mask value (0 for cr-rejected pixels, 1 for good data) for then'th image at pixel (x, y) ,
- T is the total exposure time (regardless of whether all input images were used for that particular pixel). This corresponds to the header keyword values `TEXPTIME` and `EXPTIME`.

The following `crj.fits/crc.fits` image keywords are derived from the variables in this computation:

- `TEXPTIME` = `EXPTIME` = T
- `SKYSUM` = $\sum_n sky_n$
- `REJ_RATE` = $(\sum_n T_n m_n(x,y)/T)$ averaged over all pixels
- `NCOMBINE` = n

The remaining keywords `EXPSTART` and `EXPEND` are updated based on the input image headers.

In summary, the cosmic ray rejection task sums all accepted pixel values, computes the true exposure time for that pixel, and scales the sum to correspond to the total exposure time. The final scaled, cleaned pixel is written to the comparison image to be used for the next iteration. This process is then repeated with successively less stringent detection thresholds, as specified by the `crj.fits/crc.fits` image header keyword `CRSIGMAS`. Further processing by **calacs** will scale the "pixout(x,y)" array by the gain, resulting in the summed, cosmic ray eliminated, but not sky-subtracted product (`crj.fits/crc.fits`) in units of electrons.

Cosmic Ray Rejection Table

acsrej uses the "Cosmic Ray Rejection Parameter Table" (header keyword `CRREJTAB`) to determine the number of iterations for cosmic ray rejection, the sigma levels to use for each iteration, and the spill radius to use during detection. This allows the rejection process to be tuned to each detector, with suitable defaults being applied during pipeline processing. Observers may fine-tune the cosmic ray rejection parameters when manually reprocessing data with **acsrej** by editing the `CRREJTAB` reference table.

The `CRREJTAB` reference file contains the basic parameters necessary for performing cosmic ray rejection. The column names and default values for the `CRREJTAB` are given in [Table 3.5](#). The appropriate row is selected based on the chip being processed (`CCDCHIP`), the number of images into which the exposure was split ("CR-SPLIT" Phase II optional parameter), and the exposure time of each "CR-SPLIT" image (`MEANEXP`). The sky fitting algorithm is controlled by the parameter `SKYSUB` which can have values of "mode" or "none". The "first guess" CR-combined image is then created using the median or minimum value of the input exposures, as specified by the table column `INITGUES`.

Table 3.5: Columns in Cosmic Ray Rejection Parameters Table

Column Name	Default Value	Contents
CRSPLIT	-	Number of exposures into which observation was split
MEANEXP	INDEF	Average exposure time (sec.) for each image
SCALENSE	30.0	Multiplicative term (in percent) for the noise model
INITGUES	minimum	Method for computing initial-guess image (minimum, median)

SKYSUB	mode	Sky fitting algorithm (mode, none)
CRSIGMAS	6.5, 5.5, 4.5	Rejection thresholds (sigma) for consecutive iterations
CRRADIUS	2.1	Radius (in pixels) for propagating cosmic ray
CRTHRESH	0.5555	Propagation factor
BADINPDQ	39	Data quality file bits to reject
CRMASK	yes	Flag CR-rejected pixels in input files?
CCDCHIP	-	Chip to which this conversion applies

Cosmic ray detection requires the specification of a threshold above which a pixel value is considered a cosmic ray. This threshold, defined earlier as $\tau_n = \sigma^2 \times \text{constant}$, uses the sigma rejection thresholds σ that correspond to the column CRSIGMAS in the `CRREJTAB` reference file. The table column SCALENSE is a multiplicative term (in percent) for the noise model and is given as *scale* in the threshold equation. This term can be useful when the pointing of the telescope has changed by a small fraction of a pixel between images. Under such circumstances, the undersampling of the image by the CCD will cause stars to be rejected as cosmic rays if a scale noise term is not included. This is a crude but effective step taken to satisfy the maxim of "first do no harm." However, for cases in which there have been no frame-to-frame offsets or if the image is locally well-sampled, this will unduly bias *against* rejecting cosmic rays.

Pixels within a given radius of a cosmic ray, specified in the table column CRRADIUS, will also be treated as cosmic rays. A less stringent rejection threshold, CRTHRESH, can be used for detecting pixels adjacent to a cosmic ray. As in CRSIGMAS, CRTHRESH is also given as a sigma value. If CRTHRESH is exceeded, pixels within a defined radius of the cosmic ray will also be flagged. All pixels determined to be affected by a cosmic ray will have their DQ values set to 8192, as described in [Table 3.4](#).

The pipeline adopts a conservative value for SCALENSE to avoid doing harm. In recalibrating several frames for a new cosmic ray elimination, it would be advisable to determine the full range of relative x,y offsets. An appropriate value of SCALENSE is $100 \times (\text{maximum offset in pixels})$, thus if the full offset range was 0.1 pixels an appropriate SCALENSE is 10.0. To alter this value, one may edit the `CRREJTAB` table (calibration file with the extension `crr.fits`). The number of exposures obtained in a repeated observation set can be larger than the maximum "CR-SPLIT=8" allowed in Phase II proposals.

For instance, in program 9662, 14 individual HRC exposures of duration 1.0 sec. were obtained at the same pointing. A check of relative offsets showed that a shift of about 0.2 pixels occurred between the first and last exposures. To obtain a `crj.fits` combined image of all 14 inputs at 1.0s and with better sensitivity to cosmic ray elimination than provided by the conservative default of SCALENSE = 30.0, the following should be done:

1. Create a new `asn.fits` table with one EXP-CRJ entry for each of the 14 raw 1.0s images, and an appropriately named PROD-CRJ file line.
2. Edit the `crr.fits` table adding a line with a "CR-SPLIT" value of 14, and a SCALENSE value of 20.0.
3. Given such a large number of inputs it would also make sense here to change the INITGUES value to "median."
4. Then rerun **calacs** on the new association table with 14 entries to obtain a `crj.fits` extension image based on the full stack of 1.0s images.

For a detailed discussion on manually recalibrating ACS data, please see [Section 3.5](#).

3.4.4 acs2d

Every observation, whether taken with the MAMA or CCD detectors, will be processed by **acs2d**. The primary functions of the task, listed in [Table 3.6](#), include but not limited to dark current subtraction, flat-fielding, and photometric keyword calculations. **acs2d** contains the same data quality and error array initialization functions used in **acsccd**, but **acs2d** will check to ensure that the array initialization is not performed twice on CCD data. Calibration switches in the image header control the performance of the remaining calibration functions, with MAMA-specific functions being initiated only when the relevant calibration switches are set.

Table 3.6: The Functions Performed in acs2d (in Operational Order)

<i>doDQI</i>	Initialize data quality array (if not done in acsccd)
<i>doNonLin</i>	Correct and flag non-linear data (<i>MAMA only</i>)
<i>toElectrons</i>	Convert data from DN to electrons (if not done in acsccd)
<i>doNoise</i>	Apply a simple noise model (if not done in acsccd)
<i>doDark</i>	Subtract dark image
<i>doFlash</i>	Subtract post-flash image (if required, CCD only)
<i>doFlat</i>	Divide by flat field
<i>doShad</i>	Perform CCD shutter shading correction (currently skipped)
<i>doPhot</i>	Compute photometric keyword values for header
<i>doStat</i>	Compute image statistics

***doDQI* – Bad Pixel Determination**

- Header switch: `DQICORR`
- Header keywords updated: None
- Reference file keyword (file type, suffix): `BPIXTAB` (FITS table, `bpx.fits`)

If the `DQICORR` header keyword switch is set to `COMPLETE` (e.g., for CCD data), this step is skipped. Otherwise, the same initialization is performed as described in "[doDQI – Bad Pixel Determination](#)" for **acsccd**.

***doNonLin* – Linearity Correction for MAMA Data**

- Header switch: `LFLGCROR`, `GLINCROR`
- Header keyword updated: `GLOBLIM`
- Reference file keyword (file type, suffix): `MLINTAB` (FITS table, `lin.fits`)

This routine flags global and local nonlinearity in ACS MAMA observations; the term "global" refers to the entire ACS MAMA detector, while "local" refers to an individual detector pixel. The MAMA Linearity Table, `MLINTAB`, provides the basic parameters for determining linearity.

The global limit (`GLOBAL_LIMIT`) column in this table refers to the total count rate at which the data are affected by greater than 10% non-linearity across the detector.

calacs attempts to correct for non-linearity up to the global limit using the non-linearity time constant in the column TAU. The global linearity correction is computed for every pixel below the global linearity limit specified by iteratively solving the equation $n = Ne^{(-\tau N)}$ to get the true count rate N.

The LOCAL_LIMIT can actually be much higher than the global limit and is difficult to correct using a simple algorithm. Each pixel found to exceed this limit will simply be marked as non-linear in the DQ file. This DQ flag will be extended by a fixed radius from the original pixel, given in the EXPAND column and is currently set to 2 pixels.

If the LFLGCROR switch is set to PERFORM, **acs2d** will *flag* excessive global and local nonlinearity in the DQ array. If GLINCORR is set to PERFORM, it will *correct* excessive global nonlinearity in the SCI array, if it is not too large. If the global linearity limit is exceeded, the keyword GLOBLIM in the SCI extension header will be set to EXCEEDED. Otherwise, it will have the value NOT-EXCEEDED.

doNoise - Error Array Initialization

- Header switch: None
- Header keywords updated: None
- Reference file keyword (file type, suffix): None

acs2d checks the image error array (ERR) to determine if there are non-zero values that were created from previous processing steps. If all pixel values are zero, indicating no prior calibration steps were performed, **acs2d** runs the same initializations described in "[doNoise – Error Array Initialization](#)" for **acscdd**. However, if the input image's ERR array has non-zero values, indicating processing in earlier calibration steps, the **doNoise** function in **acs2d** does not change the ERR array.

doDark – Dark Image Subtraction

- Header switch: DARKCORR
- Header keyword updated: MEANDARK
- Reference file keyword (file type, suffix): DARKFILE (image in electrons/sec., drk.fits), DRKCFIL (image in electrons/sec, dkcfits)

For the SBC, the header switch DARKCORR is set to OMIT because the SBC dark count rate is so negligible that corrections are not needed, even for long exposures.

For CCD data, the dark reference file image (in *electrons/sec*) is multiplied by the total of the image's exposure time, flash time (if post-flashed), and a 3-second idle time (post-SM4 non-BIAS WFC images only). This scaled dark reference image is subtracted from the input image in memory. The mean dark value is then computed from the scaled dark image and used to update the MEANDARK keyword in the SCI image header. Any dark accumulation during readout time is automatically included in the bias image reference file (BIASFILE), and already removed during the **doBias** step in **acscdd**.

A dark reference file is regularly created⁴ from the combination of 1000 sec. dark frames taken four times per day during several days per week. These individual dark frames are bias-corrected and combined with cosmic ray rejection. Since 2015, dark frames have been flashed to better remove CTE loss effects on warm and hot pixels, thus these darks are also flash corrected using a flash reference file.

Only hot pixels above a certain threshold (this number changes with time and is currently set at 0.08 e-/sec) are identified with the flag 16 in the data quality (DQ) array of WFC and HRC reference darks; this flag value is propagated to the DQ array of the calibrated science data. The flag 32 now marks pixels that are unstable. We now only exclude the pixels that are unstable in the subsequent processing steps since stable hot pixels, while undesirable, can be calibrated. (See [ACS ISR 2017-05](#) by Borncamp et al. for more on unstable pixels,

the reprocessing of older darks to reflect the new flag 32 scheme, and modifications to the *MDRIZTAB* associated with these changes for not only darks but also science files.) The many "warm" pixels below this threshold are assumed to be adequately corrected by the dark calibration. This produces a high signal-to-noise reference file which accurately reflects (and corrects/flags) the hot pixels present for a given observation date. The "best" dark reference file is typically not available in the pipeline for several weeks after the date of observation, because it takes a few weeks to collect enough frames to make a basedark (see [ACS ISR 2004-07](#) for more information).

The default reference files for dark subtraction, given in the header keywords `DARKFILE` and `DRKCFILE`, are assigned based on the values of the keywords `DETECTOR`, `CCDAMP`, `CCDGAIN`, and `EXPSTART` in the image header. The dark correction is applied after the overscan regions are trimmed from the input science image. As in the bias image correction (`BIASFILE`), **calacs** assumes that the images have not been binned, so any input image smaller than the full detector size will be interpreted as a subarray image. The reference file named by the `DRKCFILE` header keyword is the CTE-corrected version of the dark reference file named by the `DARKFILE` keyword, and it's only used in WFC image calibrations when pixel-based CTE correction is performed.

***doFlash* – Post-Flash Subtraction**

- Header switch: `FLSHCORR`
- Header keyword updated: `MEANFLSH`
- Reference file keyword (file type, suffix): `FLSHFILE` (image in counts, `fls.fits`)

ACS has a post-flash capability to provide the means of mitigating the effects of charge transfer efficiency degradation. The proposer controls the use of this capability via the "FLASHEXP" optional parameter in the Phase II proposal. This is only available for full frame WFC images using shutter B.

The reference file, named in the header keyword `FLSHFILE`, has the same dimensions as a full size WFC science image, 4144×2068 , allowing for simple image arithmetic on arrays of equal size. The appropriate `FLSHFILE` reference image is selected using the following keywords from the image header: `DETECTOR`, `CCDAMP`, `CCDGAIN`, `FLASHCUR`, and `SHUTRPOS`.

The function `doFlash` will subtract the post-flash in the following way:

1. The success of the post-flash exposure is verified by checking the header keyword `FLASHSTA`, which should be set to `SUCCESSFUL`. If any problems were encountered, a comment will be added to the `HISTORY` keywords in the SCI extension image header.
2. The `FLSHFILE` reference image, which is normalized to one second for the appropriate post-flash current level (`LOW`, `MED`, `HIGH`, given by the `FLASHCUR` keyword), is multiplied by the flash duration in seconds (given in the header keyword `FLASHDUR`) and subtracted from the science image.
3. Finally, the mean value of the scaled post-flash image is written to the output SCI extension image header as the keyword `MEANFLSH`.

***doFlat* – Flat-Field Image(s) Correction**

- Header switch: `FLATCORR`
- Header keywords updated: None
- Reference files keywords (file type, suffix): `PFLTFILE` (image in fractional units, `pfl.fits`), `LFLTFILE` (`lfl.fits`), `DFLTFILE` (`dfl.fits`), `CFLTFILE` (`cfl.fits`)

The ACS **doFlat** routine corrects for pixel-to-pixel and large-scale sensitivity gradients across the detector by dividing the data with the flat-field image. (Conversion to electrons no longer depends on **doFlat**, as it did in the past.)

Because of geometric distortion effects, the area of the sky seen by a given pixel is not constant, and therefore, observations of a constant surface brightness object will have count rates per pixel that vary over the detector, even if every pixel has the same sensitivity. In order to produce images that appear uniform for uniform illumination, the flat fields make an implicit correction for the geometric distortion across the field that is equivalent to dividing each pixel by the optical distortion which is normalized to unity at the center of the field. A consequence of this procedure is that two stars of equal brightness do not have the same total counts after the flat-fielding step. Thus, point source photometry extracted from a *flat-fielded* image must be multiplied by the [effective pixel area map](#), as shown in [Figure 5.1](#).

Geometric distortion corrections are also implemented in the pipeline by **AstroDrizzle** which uses the geometric distortion solution to correct all pixels to equal areas. Thus, in drizzled images, photometry is correct for both point and extended sources.

The flat-field image used to correct the data is created using up to four flat-field reference files: the pixel-to-pixel file (**PFLTFILE**), the low-order flat (**LFLTFILE**), the delta flat (**DFLTFILE**), and for HRC, the coronagraphic spot flat (**CFLTFILE**).

- The **PFLTFILE** is a pixel-to-pixel flat-field correction file containing the small scale flat-field variations.
- The **LFLTFILE** is a low-order flat which will correct for any large-scale flat-field variations across each detector. This file is stored as a binned image which is expanded when it's applied by **calacs**.
- The **DFLTFILE** is a delta-flat containing any needed changes to the small-scale **PFLTFILE**.
- The **CFLTFILE** is a spot mask which contains the vignetting patterns of the HRC occulting spots and is applied to coronagraphic observations only.

If the **LFLTFILE**, **DFLTFILE**, or **CFLTFILE** reference images are not specified in the SCI image header, only the **PFLTFILE** is used for the flat-field correction. If all four reference files are specified, they are multiplied together to form a complete flat-field correction image which is then applied to the science data.

Currently, the **LFLTFILE** and **DFLTFILE** flats are not used for ACS data. The **PFLTFILE** reference flat in the pipeline is actually a combination of the pixel-to-pixel flats taken during the ground calibration and the low-order flat correction derived in-flight. The **CFLTFILE** is applied only when the **OBSTYPE** image header keyword is equal to **CORONAGRAPHIC**.

All flat-field reference images will be chosen based on the detector, amplifier, and filters used for the observation. Any subarray science image will use the same reference file as a full-size image. **calacs** will extract the appropriate region from the reference file and apply it to the subarray input image.

doShad – Shutter Shading File Correction (skipped)

- Header switch: **SHADCORR**
- Header keywords updated: None
- Reference file keyword (file type, suffix): **SHADFILE** (image, **shd.fits**)

The **doShad** routine would apply the shutter shading correction image (named in image header keyword **SHADFILE**) to the science data if the **SHADCORR** header keyword switch were set to **PERFORM**. However, the **SHADCORR** calibration switch is currently set to **OMIT** and is unlikely to be used in future versions of **calacs**. Calibration data show that a shading correction is not needed for ACS data. For more information, refer to [ACS ISR 2003-03](#).

doPhot – Photometry Keyword Calculation

- Header switch: PHOTCORR
- Header keywords updated: PHOTMODE, PHOTFLAM, PHOTZPT, PHOTPLAM, PHOTBW
- Reference files keywords (file types, suffixes): IMPHTTAB (FITS table, `imp.fits`)

Before photometry can be performed on ACS observations, a transformation from electrons to absolute flux units must be performed. For a given instrument configuration, as described in the PHOTMODE header keyword, **calacs** uses the IMPHTTAB⁵ reference look-up table to determine the total throughput of an observing mode (linear interpolation is done using pre-computed values when necessary).

Users who wish to convert calibrated images (in units of electrons) to flux units may simply divide the image by the exposure time and then multiply by the PHOTFLAM keyword value. Drizzled images are already in units of electrons per second and may simply be multiplied by the PHOTFLAM value in the drizzled image header to create an image in flux units.

doStat – Image Statistics Determination

- Header switch: None
- Header keywords updated: NGOODPIX, GOODMIN, GOODMAX, GOODMEAN, SNRMIN, SNRMAX, SNRMEAN
- Reference files keywords (file types, suffixes): None

This routine computes the number of pixels which are flagged as "good" in the data quality array. The minimum, mean, and maximum pixel values are then calculated for data flagged as "good" in both the science and error arrays. Similarly, the minimum, mean, and maximum signal-to-noise of "good" pixels is derived for the science array. These quantities are updated in the image header.

3.4.5 acssum

- Header switch: RPTCORR
- Header keywords updated: NCOMBINE, EXPTIME, EXPEND, ROOTNAME
- Reference files keywords (file types, suffixes): None

Multiple exposures of SBC MAMA data obtained using the Phase II repeated sub-exposures specification are flagged in the association table for summing by acssum. A straight pixel-to-pixel addition of the science values is applied, and the error calculated as the square root of the sum of the squares of the errors in the individual exposures. The calibration switch RPTCORR is set to COMPLETE upon successful completion of the summation. In addition, the keywords NCOMBINE, EXPTIME, and EXPEND are adjusted to reflect the total of the summed images.

¹ In full-frame mode, each WFC CCD is read out from adjacent corners by two amplifiers, as two 2072 column × 2068 row arrays. Each row has 24 columns of physical pre-scan followed by 2048 columns of pixel data. Each column has 2048 rows of pixel data followed by 20 rows of virtual overscan. The HRC full read out is a 1062 × 1044 array: 19 columns of physical overscan followed by 1024 columns of pixel data, then another 19 columns of physical overscan. Each column consists of 1024 rows of pixel data followed by 20 rows of virtual overscan.

² Here, the asterisk is used as a wildcard, following the UNIX convention.

³ These keywords only appear in Data Quality (DQ) and Error (ERR) extensions of raw images.

⁴ The cadence of dark reference file creation has changed over time. For current information about the frequency at which darks are delivered to CRDS, as well as more in depth information about their creation, visit the [ACS website](#).

⁵ The `IMPHTTAB` reference tables, one each for ACS/HRC and ACS/WFC, contain pre-computed values of the photometry keywords for all observations.

3.5 Manual Recalibration of ACS Data

[3.5.1 Requirements for Manual Recalibration](#)
[3.5.2 calacs Examples](#)

3.5.1 Requirements for Manual Recalibration

Software Requirements

An overview of *HST* image data analysis software is available in the [Introduction to the HST Data Handbooks](#). STScI no longer supports **IRAF/PyRAF** for data analysis, and instead recommends the use of Python. The [AstroConda](#) conda channel incorporates many of the Python tools used to calibrate and analyze *HST* data including **calacs** and **AstroDrizzle**. Documentation and installation instructions are found at: <http://astroconda.readthedocs.io/>

⚠ *Before running the following examples, please make sure that your AstroConda installation is up to date. Not doing so may result in some examples breaking.*

Data Retrieval

The [Introduction to the HST Data Handbooks](#) contains an overview of data retrieval from the Archive.

Setting up "jref"

Before any recalibration can be done, the directory location for calibration reference files must be defined. For ACS, this directory is referred to as "jref", and is used as a prefix in the reference file names in the image header (i.e., jref\$qb12257gj_pfl.fits). In a Bash shell, export is used to set "jref" to a directory location. For example:

Bash INPUT

```
export jref="/mydisk/myjref/"
```

Verify your jref location using `echo $jref` prior to using these examples.

Using Non-default Reference files and Calibration Switches

By default, the Archive provides calibrated images processed with the latest available reference files at processing time. In order to use non-default reference files and calibration switch settings, manual recalibration is required. These non-default settings have to be manually updated in the uncalibrated data (the raw FITS files) before running **calacs**. The table below gives a list of all of the calibration switches in the primary header.

Table 3.7: Calibration Switch Selection Criteria

The first column shows calibration switch header keywords. The second column is a description of the keyword, and the third column shows the default values.

Switch	Description	Criteria
--------	-------------	----------

DQICORR	Data quality array initialization	DEFAULT = "PERFORM" If OBSMODE = ACQ then "OMIT" (HRC only)
ATODCORR	Analog to digital conversion	DEFAULT = "OMIT"
BLEVCORR	Overscan region subtraction	DEFAULT = "PERFORM" ("OMIT" for SBC)
BIASCORR	Bias subtraction	DEFAULT = "PERFORM" ("OMIT" for SBC)
FLSHCORR	Post-flash subtraction	DEFAULT = "OMIT" ("PERFORM" if image was post-flashed)
CRCORR	Cosmic ray rejection	If CRSPLIT >= 2 then "PERFORM" else "OMIT"
EXPSCORR	Calibrate individual exposures in an association	DEFAULT = "PERFORM"
SHADCORR	Shutter shading correction	DEFAULT = "OMIT"
PCTECORR	CTE correction	DEFAULT = "OMIT" ("PERFORM" for WFC full-frame, subarray requires manual re-calibration using acs_destripe_plus)
DARKCORR	Dark subtraction	DEFAULT = "PERFORM" ("OMIT" for SBC)
FLATCORR	Flat-field division	DEFAULT = "PERFORM"
PHOTCORR	Photometric processing	DEFAULT = "PERFORM"
RPTCORR	Repeated sub-exposure processing	If NRPTEXP > 1 then "PERFORM" else "OMIT"
DRIZCORR	Dither processing	DEFAULT = "PERFORM"
SINKCORR	Sink pixel flagging	DEFAULT = "PERFORM" for WFC only after January 15, 2015, "OMIT" for all other data
GLINCORR	Global non-linearity correction	DEFAULT = "PERFORM" (SBC only)
LFLGCORR	Local and global non-linearity flagging in DQ array	DEFAULT = "PERFORM" (SBC only)

Post-SM4 WFC Image Artifact Correction for WFC Subarray Images

Certain artifacts present in post-SM4 WFC subarray images, including bias striping, bias shift, and CTE trailing, are not fully handled by **calacs**. The amplitude of the bias shift and the CTE trailing in the 512- and 1024-pixel subarrays of *HST* Cycles 17–23 are not well characterized. However, the 2K subarray CTE trailing is near-identical to full-frame readout, and is corrected in **calacs** with the full-frame algorithm. All the new

subarray modes introduced for use in Cycle 24 onwards will have **calacs** correction of bias shift and CTE trailing. Note that CTE correction for subarray images is not performed automatically, and the image must first be de-striped (see below for more information).

Bias striping in post-SM4 subarray images is not corrected within **calacs**. Subarray stripe removal requires fitting across the entire image region, as discussed in [Section 3.4.1](#), using the stand-alone task **acs_destripe_plus** in the AstroConda **acstools** suite. (See also [Example 5](#) in [Section 3.5.2](#))

Bypassing the PHOTCORR Step

During the **doPhot** step, pixel values and units are not changed. This step only calculates the values of the calibrated image's photometric header keywords, such as the inverse sensitivity conversion factor (**PHOTFLAM**). Please refer to [Section 3.4.4 "doPhot - Photometry Keyword Calculator"](#) for more information.

When populating the photometric keywords during the **doPhot** step, **calacs** uses the **CRDS** reference file **IMPHTTAB**. Some users find it cumbersome to keep up with the updates, and prefer to simply copy the photometric keyword values from the original calibrated data into the raw image's primary header, then run **calacs** with the **PHOTCORR** switch set to **OMIT**.

3.5.2 calacs Examples

In these examples, Python is used to run **calacs**. These Python functions are simply wrappers around the C code that comprises **calacs**, and as such they may also be run outside of the Python environment.

Example 1: Reprocessing a Single Exposure Using a Different Bias File

The following example uses WFC data from the Cycle 23 CAL/ACS program 14398 (PI: Chiaberge), which monitors the ACS/WFC CTE using observations of the globular cluster NGC 104 (47 Tucanae) in the F502N filter. The dataset names are **JD1Y04RLQ** and **JD1Y04RNQ**, and are part of the association **JD1Y04010**. In our examples, we have stored the FITS files in a subdirectory called **cal_fits/** so that we can copy the files to the current working directory while preserving the original files for other examples.

Python INPUT:

```
from astropy.io import fits
from astropy.table import Table
import shutil

shutil.copy('cal_fits/jdly04010_asn.fits', '.')
shutil.copy('cal_fits/mybias.fits', '.')

#Adding the Table() function from astropy.table makes some
#tasks more convenient and makes the formatting of the table
#more readable in the notebook

asn_hdu = fits.open('jdly04010_asn.fits')
asn_table = Table(asn_hdu[1].data)
asn_hdu.close()

asn_table
```

Python OUTPUT:

MEMNAME	MEMTYPE	MEMPRSNT
str14	str14	bool

JD1Y04RLQ	EXP-CRJ	True
JD1Y04RNQ	EXP-CRJ	True
JD1Y04011	PROD-CRJ	True

For the purposes of this first example, assume that the observations *are not* part of an association. This example will illustrate the steps required to reprocess a single exposure (JD1Y04RLQ) after changing the bias reference file from the default value to a file specified by the user. ***Note that we will also CTE correct this image, but we will omit this for future examples. The CTE correction is computationally expensive and thus may take a long time to run on your machine.*** By default, it will use all of your machine's processors.

1. Make sure that the `jref` keyword is set to the ACS reference files as described earlier in this section.
2. Copy the original file into the current working directory as shown above for the association file.
3. To determine which bias reference file name was specified in the image header, open the raw FITS file and access the value for the header keyword `BIASFILE`

Python INPUT:

```
shutil.copy('cal_fits/jdly04rlq_raw.fits', '.')
hdu_raw = fits.open('jdly04rlq_raw.fits', mode = 'update')
hdr0 = hdu_raw[0].header

print(hdr0['BIASFILE'])
```

Python OUTPUT:

```
jref$09k2026tj_bia.fits
```

4. Edit the primary header of the raw image to enter the name of the new bias file called `mybias.fits` (for this example, `mybias.fits` is a copy of the default bias reference file listed in the header).

Python INPUT:

```
hdr0['BIASFILE'] = 'mybias.fits'
```

5. Set the `PHOTCORR` processing step to `OMIT` and copy the values of two useful photometric keywords from the calibrated image (retrieved from the Archive) to the raw image (see ["Bypassing the PHOTCORR"](#) for more information.) The `PHOTFLAM` keyword will be useful for photometric calibration during image analysis, and `PHOTMODE` is useful as a concise description of the observation mode. (Note: For WFC images, the keywords need to be edited for both SCI extensions, however the values in the extensions are identical.)
6. The pixel-based CTE correction algorithm was updated in July 2017, and not all files in the Archive have been reprocessed using the new method. Additionally, any files you may have downloaded from the Archive prior to August 3, 2017 will not have the new `PCTETAB` file in the header. In these cases, **calacs** will not run on files with `PCTETAB` set to an old reference file unless either `PCTECORR` is set to `OMIT` or `PCTETAB` is set to the latest reference file (alternatively, one can force **calacs** to run with the old method as well). In this example, we edit `PCTETAB` to point to the most recent `*_cte.fits` file in the `jref` directory.

Python INPUT:


```

hdr0['PHOTCORR'] = 'OMIT'

#Grab the values of the PHOTFLAM and PHOTMODE
#keywords from the FLT file in cal_fits/
#Note that we have mixed the nomenclature to use the [1]
#and ['sci',1] notations for the image extensions. Please
#see the ACS Data Handbook Section 2.2 "ACS File Structure"
#for more information.
hdu_flt = fits.open('cal_fits/jdly04rlq_flt.fits')
hdr_flt = hdu_flt[1].header
hdr1, hdr4 = hdu_raw['sci',1].header, hdu_raw['sci',2].header
hdu_flt.close()
photflam, photmode = hdr_flt['PHOTFLAM'], hdr_flt['PHOTMODE']

#Update the header keywords
hdr1['PHOTFLAM'] = photflam
hdr1['PHOTMODE'] = photmode
hdr4['PHOTFLAM'] = photflam
hdr4['PHOTMODE'] = photmode
hdr0['PCTETAB'] = 'jref$16k1747tj_cte.fits'

#Close the raw file to push the changes to the header
hdu_raw.close()

```

7. We can now run **calacs** on the raw FITS file, which will be processed with the user-specified bias reference file `mybias.fits`. The result will be two calibrated images: one with the `flt.fits` extension; and a CTE-corrected one with the `flt.fits` extension. A trailer file with the `.tra` extension describing the steps taken by **calacs** is also produced.

NOTE: Importing **calacs** will print a message (see the code block below). This is done on import only and not every time **calacs** is run.

Python INPUT:

```

from acstools.calacs import calacs

calacs('jdly04rlq_raw.fits')

```

The following tasks in the **acstools** package can be run with TEAL:

```

acs2d    acs_destripe    acs_destripe_plus
acscdd   acscte          acsrej
acssum   calacs

```

PixCteCorr is no longer supported. Please use **acscte**.

The cell below cleans up the current working directory by removing the files created by this example. **Only run this when you are finished with this example.**

Python INPUT:

```

import os, glob
files = glob.glob('jdly04*') + ['mybias.fits']
for x in files:
    if os.path.exists(x):
        os.remove(x)

```

Example 2: Reprocessing Multiple Exposures Taken with "CR-SPLIT" Within an Association

This example uses the same data from Example 1 and illustrates the steps required to reprocess an ACS association after changing the bias reference file from the default value to a file specified by the user. The steps required are similar to the previous example, with a few modifications.

1. The association table shows the images from two exposures. The MEMTYPE value "EXP-CRJ" indicate that those two images were created from a "CR-SPLIT" exposure. The cosmic ray-rejected product created by **calacs** has the rootname J1DY04011

Python INPUT:

```
from astropy.io import fits
from astropy.table import Table
from acstools.calacs import calacs
import shutil, glob

shutil.copy('cal_fits/jdly04010_asn.fits', '.')
shutil.copy('cal_fits/mybias.fits', '.')

#Adding the Table() function from astropy.table makes some
#tasks more convenient and makes the formatting of the table
#more readable in the notebook
asn_hdu = fits.open('jdly04010_asn.fits')
asn_table = Table(asn_hdu[1].data)
asn_hdu.close()

asn_table
```

Python OUTPUT:

MEMNAME	MEMTYPE	MEMPRSNT
str14	str14	bool
JD1Y04RLQ	EXP-CRJ	True
JD1Y04RNQ	EXP-CRJ	True
JD1Y04011	PROD-CRJ	True

2. Copy the images from the cal_fits/ directory to the current working directory and edit the headers to point BIASFILE to mybias.fits and update the PHOTCORR, PHOTFLAM, and PHOTMODE keywords appropriately as in Example 1. As mentioned previously, in this and future examples, we will set the PCTECORR keyword to OMIT for the sake of expediency. Then run **calacs**.

Python INPUT:

```
#Here we have hardcoded the file names, but we could also
#access the rootnames from asn_table as asn_table['MEMNAME'][0]
#and asn_table['MEMNAME'][1].
shutil.copy('cal_fits/jdly04rlq_raw.fits', '.')
shutil.copy('cal_fits/jdly04rnq_raw.fits', '.')

#Use the glob function to create a list of the raw FITS files
#that we can loop over
raw_files = glob.glob('*raw.fits')
for x in raw_files:

    #Open the primary header of the raw image and get the
```

```

#rootname for accessing the calibrated files
hdu_raw = fits.open(x, mode = 'update')
hdr0, hdr1, hdr4 = hdu_raw[0].header, hdu_raw['sci',1].header, hdu_raw['sci',2].header
rootname = hdr0['ROOTNAME'].lower()
cal_image = 'cal_fits/' + rootname + '_flt.fits'

#Get the PHOTFLAM and PHOTMODE keywords from the
#calibrated image science header
hdu_flt = fits.open(cal_image)
hdr_flt = hdu_flt['sci',1].header
hdu_flt.close()
photflam, photmode = hdr_flt['PHOTFLAM'], hdr_flt['PHOTMODE']

#Update the raw image headers
hdr0['BIASFILE'] = 'mybias.fits'
hdr0['PCTETAB'] = 'jref$16k1747tj_cte.fits'
hdr0['PHOTCORR'] = 'OMIT'
hdr0['PCTECORR'] = 'OMIT'
hdr1['PHOTMODE'] = photmode
hdr1['PHOTFLAM'] = photflam
hdr4['PHOTMODE'] = photmode
hdr4['PHOTFLAM'] = photflam

#Close the raw file to push the changes to the header
hdu_raw.close()

#Run calacs on the association file
calacs('jdly04010_asn.fits')

```

The products are three calibrated images: 1) two `flt.fits` files, one for each of the EXP-CRJ images in the association; and 2) one cosmic-ray cleaned image that is the combination of the two input images using the **acsrej** algorithm with no CTE correction applied.

The cell below cleans up the current working directory by removing the files created by this example. **Only run this when you are finished with this example.**

Python OUTPUT:

```

import os, glob
files = glob.glob('jdly04*') + ['mybias.fits']
for x in files:
    if os.path.exists(x):
        os.remove(x)

```

Example 3: Combining Exposures From Multiple Associations

This example illustrates the steps required to combine two sets of repeated observations to create a cosmic ray-rejected combined image. The data for this exercise comes from the ACS calibration program, 9662, that observed NGC 104 (47 Tucanae) using the HRC with a clear filter. The associations J8IS01020 and J8IS01040. Each association comprises of two 1-second exposures, and share the same target pointing.

1. Copy the files from the `cal_fits/` directory to the current working directory, open the association FITS files and merge the two tables into one. Only use the first two rows of each file, as the third row of each table contains a combined product that we are going to exclude.

```

from astropy.io import fits
from astropy.table import Table, vstack
from numpy import rec
from acstools.calacs import calacs
import shutil

shutil.copy('cal_fits/j8is01020_asn.fits', '.')

```

```

shutil.copy('cal_fits/j8is01040_asn.fits', '.')

asn_hdu1 = fits.open('j8is01020_asn.fits')
asn_hdu2 = fits.open('j8is01040_asn.fits')

asn_tab1 = Table(asn_hdu1[1].data)
asn_tab2 = Table(asn_hdu2[1].data)
asn_hdu1.close()
asn_hdu2.close()

#Copy one of the association files to a new file called merged_asn.fits
#for which we will overwrite the data
shutil.copy('j8is01020_asn.fits', 'merged_asn.fits')
merged_asn = fits.open('merged_asn.fits', mode = 'update')

merged_table = vstack([asn_tab1[0:2], asn_tab2[0:2]])
merged_table

```

Python OUTPUT:

MEMNAME	MEMTYPE	MEMPRSNT
str56	str56	bool
J8IS01J2Q	EXP-RPT	True
J8IS01J3Q	EXP-RPT	True
J8IS01J8Q	EXP-RPT	True
J8IS01J9Q	EXP-RPT	True

- Now that we have a table with each of the individual exposures in it, we need to add one more row to contain the cosmic-ray rejected product name (J8IS0xx1).

Python INPUT:

```

merged_table.add_row(['J8IS01xx1', 'PROD_CRJ', 'yes'])
merged_table

```

Python OUTPUT:

MEMNAME	MEMTYPE	MEMPRSNT
str56	str56	bool
J8IS01J2Q	EXP-RPT	True
J8IS01J3Q	EXP-RPT	True
J8IS01J8Q	EXP-RPT	True
J8IS01J9Q	EXP-RPT	True
J8IS01xx1	PROD_CRJ	True

Python INPUT:

```
#Replace the data in the merged_asn.fits file with the new table.
#The Astropy Table format is not the same as the FITS file table
#format, so we must convert the Astropy format to a binary FITS
#table HDU first, then replace the data array in the association
#file with the data array of the newly created table HDU.
#Then close the file to push the changes.
fits_table = fits.table_to_hdu(merged_table)
merged_asn[1].data = fits_table.data
merged_asn.close()
```

3. As in the other examples, set PHOTCORR to OMIT and copy the PHOTFLAM and PHOTMODE keywords from the calibrated files. As this is HRC data, there is only one science extension, and thus we do not need to update the headers of multiple extensions.

Python INPUT:

```
#Here we have used a more advanced syntax to compress our FOR
#statement into one line called a list comprehension. However,
#we are simply creating a list of the root names from the merged
#association if they contain the letter "Q" (i.e., they are exposures
#and not products)
rootnames = [x.rstrip().lower() for x in merged_table['MEMNAME'] if 'Q' in x]
for x in rootnames:
    shutil.copy('cal_fits/' + x + '_raw.fits', '.')
    raw_hdu = fits.open(x + '_raw.fits', mode = 'update')
    flt_hdu = fits.open('cal_fits/' + x + '_flt.fits')
    hdr0 = raw_hdu[0].header
    hdr1 = raw_hdu['sci',1].header
    flt_hdr = flt_hdu['sci',1].header
    flt_hdu.close()
    hdr0['PHOTCORR'] = 'OMIT'
    hdr1['PHOTFLAM'] = flt_hdr['PHOTFLAM']
    hdr1['PHOTMODE'] = flt_hdr['PHOTMODE']
#Run calacs on the association
calacs('merged_asn.fits')
```

The products are five images: 1) four `flt.fits` calibrated images; and 2) one `crj.fits` cosmic-ray cleaned stack of the four `flt.fits` files.

The cell below cleans up the current working directory by removing the files created by this example. **Only run this when you are finished with this example.**

Python INPUT:

```
import os, glob
files = glob.glob('j8is01*') + ['merged_asn.fits', 'merged.tra']
for x in files:
    if os.path.exists(x):
        os.remove(x)
```

Example 4: Reprocessing Images Taken as part of a Dither Pattern

The following example uses WFC data from the GOODS program 9425. These observations are from visit 54, exposure 219; the target name was "CDF-South," observed with the F606W filter. The images were part of a 2-point line dither pattern with an exposure time of 480 seconds each, with rootnames `J8E654C0Q` and `J8E654C4Q`.

This example illustrates the steps needed to reprocess data that are part of a dither pattern using a non-default dark reference file.

1. Copy the association file and images to the current working directory, then view the contents of the association file.

```
from astropy.io import fits
from astropy.table import Table
from acstools.calacs import calacs
import shutil
shutil.copy('cal_fits/j8e654010_asn.fits', '.')
shutil.copy('cal_fits/mydark.fits', '.')
asn_hdu = fits.open('j8e654010_asn.fits')
asn_table = Table(asn_hdu[1].data)
asn_hdu.close()
asn_table
```

Python OUTPUT:

MEMNAME	MEMTYPE	MEMPRSNT
str14	str14	bool
J8E654C0Q	EXP-RPT	True
J8E654C4Q	EXP-RPT	True
J8E654010	PROD-DTH	True

2. Edit the global image header for all the raw images to insert the name of the new dark reference file, `mydark.fits` (as in Example 1, we have copied the original dark file from the `jref` directory to the current working directory and renamed it `mydark.fits`).
3. Edit the `PHOTCORR`, `PHOTFLAM`, and `PHOTMODE` keywords as in previous examples.
4. Since `mydark.fits` does not have a counterpart CTE-corrected dark reference file, set `PCTECORR` to `OMIT` so that CTE-corrected images are not generated.

```
#Get the rootnames of the exposures directly from the association
#table and copy the raw files from the cal_fits/ directory to the
#current working directory
raw_files = [asn_table['MEMNAME'][0].rstrip().lower(),
              asn_table['MEMNAME'][1].rstrip().lower()]

for x in raw_files:
    shutil.copy('cal_fits/' + x + '_raw.fits', '.')
    #Get the raw and FLT image headers
    hdu_raw = fits.open(x + '_raw.fits', mode = 'update')
    hdu_flt = fits.open('cal_fits/' + x + '_flt.fits')
    hdr0, hdr1, hdr4 = hdu_raw[0].header, hdu_raw['sci',1].header, hdu_raw['sci',2].header
    hdr_flt = hdu_flt['sci',1].header
    hdu_flt.close()
    #Grab the photometry keywords from the FLT header
    photmode, photflam = hdr_flt['PHOTMODE'], hdr_flt['PHOTFLAM']
    #Set all of the necessary keywords in the raw FITS file
    hdr0['DARKFILE'] = 'mydark.fits'
    hdr0['PHOTCORR'] = 'OMIT'
    hdr0['PCTECORR'] = 'OMIT'
    hdr1['PHOTMODE'] = photmode
    hdr1['PHOTFLAM'] = photflam
    hdr4['PHOTMODE'] = photmode
    hdr4['PHOTFLAM'] = photflam
    #Close the raw FITS file to push the changes to the header
    hdu_raw.close()
#Run calacs on the association file
calacs('j8e654010_asn.fits')
```

The result of this **calacs** run will be two calibrated images, one for each exposure in the association. This time, there will only `beflt.fits` files, not CTE-corrected `flc.fits` files, as we have set `PCTECORR` to `OMIT`. These files can then be used in **AstroDrizzle** to create a drizzled image.

The cell below cleans up the current working directory by removing the files created by this example. **Only run this when you are finished with this example.**

Python INPUT:

```
import os, glob
files = glob.glob('j8e654*') + ['mydark.fits']
for x in files:
    if os.path.exists(x):
        os.remove(x)
```

Example 5: CTE Correction On Sub-array Images

The ACS Instrument Team implemented new flight software in Cycle 24 that changed the way CCD sub-array images are read out. This change makes it possible to apply the pixel based CTE correction to sub-array images. Here we illustrate the steps necessary to implement the correction in post processing.

The correction is implemented by using a tool within the **acstools** python package called **acs_destripe_plus**. This tool is useful for when built-in **calacs** de-stripping algorithm using overscans is insufficient or unavailable. The use of this tool to apply CTE correction is limited to the following cases:

1. All 2K sub-arrays taken after SM4.
2. Smaller sub-arrays taken after October 2016.

As a reminder, in order to get the latest CTE correction you will need to install the AstroConda environment. This will include **acstools** and other python packages that you will need. Instructions on how to install AstroConda can be found in the [AstroConda website](#).

For this example we will use the observation of the 47 Tuc calibration field with rootname `JD5702JPQ`, which we have copied to a file called `filename.fits` in the following example. These data were taken using the WFC1B-1K aperture. This is a 1024 x 2048 pix sub-array.

1. Update the image header to turn on the `PCTECORR` switch and update reference file information as necessary.

```
PCTECORR = "PERFORM"
PCTETAB = "jref$19i16323j_cte.fits"
DRKCFIL = "jref$name_of_appropriate_dark_file_dkc.fits"
```

You can do this with some python commands:

Python INPUT:

```
from astropy.io import fits

with fits.open('filename.fits',mode='update') as hdu:
    hdu[0].header['PCTECORR'] = 'PERFORM'
    hdu[0].header['PCTETAB'] = 'jref$19i16323j_cte.fits'
    hdu[0].header['DRKCFIL'] = 'jref$19j14351j_dkc.fits'
```

Run `acs_destripe_plus`

Finally, to run the correction run the following commands in python:

Python INPUT:

```
from acstools import acs_destripe_plus

acs_destripe_plus.destripe_plus('jd5702jpq_raw.fits',cte_correct=True)
```

The tool will take some time to run. In an 8 core machine it takes 3 minutes for the CTE correction to complete.

This tool will give you and FLT and an FLC file. The FLC file is the one that has been CTE corrected.

Correct `acs_destripe_plus` WCS

The subarray products produced by this process do not have the proper WCS information in the header. This is normally done by the pipeline via an additional call to **AstroDrizzle**. You can do this yourself with two python commands:

Python INPUT:

```
from stwcs import updatewcs

updatewcs.updatewcs('*flc.fits')
```

With that, you should have calibrated files ready for scientific analysis or to process using **Astrodrizzle**.

Chapter 4: ACS Data Processing Considerations

Chapter Contents

- [4.1 Read Noise and A-to-D Conversion](#)
- [4.2 Bias Issues](#)
- [4.3 Dark Current, Hot Pixels, and Cosmic Rays](#)
- [4.4 Flat-Field Reference Files](#)
- [4.5 Image Anomalies](#)
- [4.6 WFC CCD Detector Charge Transfer Efficiency - CTE](#)
- [4.7 Generic Detector and Camera Properties](#)

4.1 Read Noise and A-to-D Conversion

[4.1.1 Read Noise](#)

[4.1.2 A-to-D Conversion](#)

4.1.1 Read Noise

Read noise is an irreducible contribution to the total error budget. It is linked to the readout process and there are no reduction steps that can minimize or remove it. The influence of read noise is minimized by reducing the number of pixels in the photometry aperture, and by subdividing the total exposure into the minimum number of component readouts while still permitting an adequate cosmic ray and hot pixel removal, as well as pixel-phase sampling. Read noise is independent of position on the CCD. Tables [4.1](#), [4.2](#), and [4.3](#) show gain and read noise values of the four WFC amplifiers for different time periods. [Figure 4.1](#) illustrates the history of the amplifier-based read noise (in electrons) since 2003.

Figure 4.1: History of the amplifier-based read noise (in electrons).




Table 4.1: WFC Amplifier Gain and Read Noise After Installation of the CEB-R (Valid After May 2009)

Values apply to dual-slope integrator mode of pixel sampling.

CCD	Amp	GAIN (e-/DN)				READ NOISE (e-)			
		0.5	1.0	1.4	2.0 ^a	0.5	1.0	1.4	2.0 ^a
WFC1	A	0.53	1.03	1.45	2.020	3.92	4.05	4.28	4.57
WFC1	B	0.50	0.96	1.36	1.886	3.24	3.36	3.61	3.91
WFC2	C	0.53	1.03	1.45	2.017	3.54	3.69	3.95	4.25
WFC2	D	0.53	1.02	1.45	2.011	3.31	3.43	3.71	5.3 ^b

^a Default Gain.

^b Valid after 20 January 2013; read noise was 4.04 e- prior to this date.

Table 4.2: CCD Gain and Read Noise Operated Under Side 1 of Original CEB (March 2002 to June 2006)

CCD	Amp	GAIN (e-/DN)			READ NOISE (e-)		
		1	2 ^a	4	1	2 ^a	4
WFC1	A	1.000	2.002	4.01	5.57	5.84	–
WFC1	B	0.972	1.945	3.90	4.70	4.98	–

WFC2	C	1.011	2.028	4.07	5.18	5.35	–
WFC2	D	1.018	1.994	4.00	4.80	5.27	–
HRC	C	1.163	2.216	4.235	4.46	4.80	5.86

^a Since Cycle 14 (under Side 1 operations), the default gain has been GAIN = 2. Prior to this, GAIN = 1 was the default.

Table 4.3: CCD Gain and Read Noise Operated Under Side 2 of Original CEB (July 2006 to January 2007)

CCD	Amp	GAIN (e-/DN)			READ NOISE (e-)		
		1	2 ^a	4	1	2 ^a	4
WFC1	A	1.000	2.002	4.01	5.29	5.62	–
WFC1	B	0.972	1.945	3.90	4.45	4.74	–
WFC2	C	1.011	2.028	4.07	5.03	5.34	–
WFC2	D	1.018	1.994	4.00	4.55	4.89	–
HRC	C	1.163	2.216	4.235	4.36	4.82	5.44

^a Default Gain.

4.1.2 A-to-D Conversion

The analog information (electrons) accumulated in the CCD is converted into data numbers (DNs) by the analog-to-digital converter (ADC). Both ACS CCD cameras employ 16-bit ADCs, which can produce a maximum of $2^{16} - 1 = 65,535$ DN. If the gain conversion factor g is expressed in e-/DN, the largest number of electrons representable by these ADCs is given by $g \times 2^{16}$. Any charge packet containing more than $g \times 2^{16}$ electrons will reach the digital saturation limit (see [Section 4.7.4](#)). The analog-to-digital converter produces only discrete output levels. This means that a range of analog inputs can produce the same digital output. This round-off error is called quantizing noise. It can be shown (Janesick, J. R., 2001, *Scientific Charge-Coupled Devices*, SPIE Press) that quantizing noise QN is constant for a given gain setting when expressed in DN: $QN(DN) = 12^{-1/2} = 0.288675$. Quantizing noise can be converted into noise electrons as: $QN(e-) = 0.288675 \times g$.

The measured read noise (RN; reported in [Tables 4.1](#), [4.2](#), and [4.3](#)) is the quadrature sum of this quantizing noise and the intrinsic read noise (IRN) associated with a particular readout amplifier:

$$RN = (IRN^2 + (0.288675 \times g)^2)^{1/2}$$

4.2 Bias Issues

- [4.2.1 Bias Calibration](#)
- [4.2.2 Bias Jump](#)
- [4.2.3 Post-SM4 WFC Signal-dependent Bias Shift](#)
- [4.2.4 Bias Subarrays for WFC and HRC](#)

4.2.1 Bias Calibration

calacs performs the bias correction in two steps (see [Section 3.4.1](#)):

- (1) **doBias**: a "superbias" reference image is subtracted from the science image to remove the fixed bias structure and during-readout dark current.¹ The superbias, constructed from individual bias frames obtained three times a week, samples the fixed bias structure at a high signal-to-noise ratio and is free of cosmic ray artifacts. This calibration step is applied to the image before its conversion to units of electrons, so the superbias is in units of DN.
- (2) **doBlev**: this step subtracts the bias level from the image, after the image has been converted from DNs to electrons. For pre-SM4 data, it fits the bias level from the physical overscan, and subtracts it from the science data. For post-SM4 data, additional steps are required to remove the fixed bias structure and "readout dark," the dark current that accrues during readout ([ACS ISR 2014-02](#)).

WFC Bias Level Determination

Each quadrant of the WFC focal plane (A & B for WFC1, C & D for WFC2) has two overscan regions: a 24 pixel-wide leading physical pre-scan at columns 1–24 for amplifiers A & C, and columns 4121–4144 for amplifiers B & D, and a 20 row-wide virtual overscan at rows 2049–2068. The physical pre-scan is produced by 24 extra pixels in the CCDs' serial registers between the readout amplifiers and the imaging region of the CCD. The virtual overscans are obtained by over-clocking the last rows in the imaging regions of each CCD 20 times.


After each vertical row shift, the bias level requires some time to reach its nominal level. The bias levels in the first 18 columns of the physical pre-scans associated with each WFC amplifier decay quasi-exponentially to their nominal levels. The bias levels in the imaging areas of the CCDs can be safely measured using the six columns of the physical pre-scans adjacent to the imaging areas, i.e., columns 19–24 for amplifiers A & C, and columns 4121–4126 for amplifiers B & D.

The virtual overscan is not used to estimate the bias level. This region exhibits large scale structure that is quadrant- and gain-dependent. Moreover, this region can be contaminated by deferred charge due to degradation in the parallel charge transfer efficiency.

The WFC bias frames show small differences between the bias levels of the physical pre-scans and the imaging region of the CCD ([Sirianni, et al., 2002, HST Calibration Workshop, STScI, page 82](#)). These bias offsets vary from amplifier to amplifier and they can be as large as 3.5 DN. If these offsets were constant, a full frame bias subtraction (doBias) would remove any differences between the pre-scans and the imaging region. Unfortunately, the offsets show random variations of about 0.3 DN that may be caused by interference between the WFC integrated electronics module and the telescope and/or other science instruments. The accuracy of the bias level subtraction in a single quadrant is limited by this random effect. Consequently, sky background levels often appear discontinuous across the boundaries of adjacent image

quadrants after **calacs** processing ([Figure 4.2](#)). Automated photometry of point or extended sources that span the quadrant boundaries should therefore be considered suspect. In such cases, measuring and subtracting the sky background levels in each quadrant separately is recommended.

Figure 4.2: Calibrated WFC1 Image Showing the Quadrant-to-quadrant Jump



Post-SM4 WFC Bias Gradient

Since Servicing Mission 4 (SM4), the WFC bias frames exhibit two-dimensional spatial gradients of 5 DN to 10 DN within each image quadrant ([Figure 4.3](#)). These gradients are stable within the time spanned by each superbias reference image, and so they are completely removed (along with other fixed pattern noise) in the ***doBias*** step of ***calacs***. These gradients are characteristics of the dual-slope integrator (DSI) implemented in the replacement CCD electronics to reduce the noise incurred during pixel sampling. The gradients are mainly caused by slow drifts of the bias reference voltages during and after the readout of each row of pixels. A small fraction of the gradient is also due to accumulated dark counts during the readout process. Since the time for a given pixel to read out depends on its vertical distance from the amplifier, there is a vertical gradient in accumulated readout dark counts, which adds to the bias gradient.

These gradients were not produced by the pre-SM4 CCD electronics, which used the clamp-and-sample technique of pixel sampling. The replacement electronics also offer the clamp-and-sample option and, consequently, gradient-free biases, but this option increases the read noise of the images by about 0.5 e- to 1.0 e-. Therefore, the DSI is the default mode for post-SM4 WFC operations, while the clamp-and-sample is unsupported.

Figure 4.3: Bias Gradients Seen Within Each Image Quadrant in the WFC Superbias (In Units of Counts) after SM4





Figure 4.4: A side-by-side comparison of the same image section in a single Bias (left) and a de-striping-uncorrected Superbias (right), from late August 2009




The linear stretch is identical, and is equivalent to ± 3 of the single bias. The local pixel-to-pixel noise (including both read noise and striping noise) of the superbias, which comprises 16 biases (including the one shown on the left) is a factor of ~ 3.6 times lower than the noise level of the single bias.

Post-SM4 WFC Bias "Striping"

The ACS CCD Electronics Box Replacement includes a SIDE CAR Application- Specific Integrated Circuit (ASIC) that exhibits a low frequency noise (1 mHz to 1 Hz) on the bias and reference voltages it generates for the WFC CCDs. This noise contribution does not matter for bias voltages going to the CCD since it is canceled out by correlated double sampling (CDS). However, there is one reference voltage from the ASIC that is used to offset the signal applied after the CDS stage (see [ACS ISR 2011-05](#)). Here, the noise does not cancel out, and manifests as a slow moving variation of the baseline. In practice, "striping" is observed in all post-SM4 WFC images that is virtually uniform across both amplifier readouts (the entire 4096 columns) of each WFC CCD ([Figure 4.4](#)).

Because of the uniformity of the striping across WFC rows, it is straightforward to characterize and remove this low-level $1/f$ noise from WFC bias frames. The amplitude distribution is well fit by a Gaussian of $\sigma_G = 0.74$ e- with an enhanced negative tail, giving an overall $\sigma = 0.9$ e-. (See [Figure 4.5](#)). This is under 25% of the WFC read noise. Averaging N bias images reduces the $1/f$ noise by nearly a factor of $N^{1/2}$, so the total noise in the post-SM4 WFC superbias reference images approaches pre-2007 levels. The striping is not well estimated by the limited WFC overscan regions alone, complicating stripe removal in non-bias frames. A WFC superdark reference image (average of ~ 24 darks) is effectively stripe-free, but science targets rarely comprise so many exposures.

Figure 4.5: Characterization of the Low-level 1/f Noise from WFC Bias Frames




HRC Bias Level Determination

HRC images were read out using only one amplifier. Each image has three overscan regions: a physical pre-scan and overscan of width 19 pixels at columns 1–19 and 1044–1062, and virtual overscan of width 20 pixels at rows 1025–1044. The first 10 columns of the physical pre-scan exhibit bias-settling behavior similar to that described for the WFC quadrants. The bias level is therefore measured in the last six columns of the pre-scan. There is not a significant difference between the bias levels in the pre-scan and imaging regions of the CCD.

4.2.2 Bias Jump

WFC images obtained before SM4 showed intermittent bias variations of a few tenths of a DN during readout. Bias frames occasionally exhibited horizontal bias jumps in one or more quadrants that lasted for several hundreds of rows ([Figure 4.6](#)). The probable cause of these jumps was electronic interference from other scientific instruments and/or spacecraft activities. There is no automatic detection of these bias jumps within the calibration pipeline. Bias jumps at the sub-DN level are not important for most science applications, but users should be aware of their possible existence in their **calacs** image products.

Figure 4.6: Bias Jump in the WFC1 Quadrant B Only



The vertical stripes in the data are hot columns and are unrelated to bias.

4.2.3 Post-SM4 WFC Signal-dependent Bias Shift

The DSI mode (Dual-slope Integrator, implemented in the replacement CCD electronics during SM4) of WFC operation induces a signal-dependent bias shift, the cause of which is closely related to that of the bias gradient described in [Section 4.2.1](#). The DC level of the DSI mode is sensitive to changes in the CCD output voltage in such a way that the pixel bias level is shifted positively by 0.02%–0.30% (depending on the amplifier) of the signal from the previously integrated pixel. This phenomenon is well characterized for ACS/WFC full frame images and can be analytically removed using a parametric algorithm described in [ACS-ISR 2012-02](#). The **calacs** pipeline now performs this correction, only for ACS/WFC post-SM4 full frame images, during the **doBlev** stage.

4.2.4 Bias Subarrays for WFC and HRC

Before SM4, the bias reference frames for science images obtained with WFC and HRC subarray readout modes were simply extracted from the appropriate locations of the full frame bias reference images, provided the subarrays did not cross quadrant boundaries. Tests showed that subarray science images that were calibrated with the relevant extracted regions from a full frame reference image were just as good as those calibrated using reference images with the same subarray readout patterns. Users were advised not to

use subarrays that spanned amplifier quadrants because doing so required the procurement of single amplifier subarray bias images at the expense of the users' observing time.

Unfortunately, this convenient use of full frame reference images became unsuitable for post-SM4 subarray science images because the two-dimensional bias gradients imposed by the DSI (see [Section 4.2.1](#)) are dependent on the timing patterns used to read out the CCD. The bias gradients seen in the standard 512×512 and 1024×1024 subarray images are significantly different from the gradients seen in the full-quadrant and full frame readout modes.

For the post-SM4 Cycles 17–23, subarray bias reference images have been obtained depending on the number and nature of the science programs that use subarrays during any given cycle:

- For GO programs (with well-defined scheduling windows) the subarray bias reference frames are inserted directly into the GO proposal and linked to the science visits so that they are obtained within two weeks of the science exposures.
- For SNAP programs (with no scheduling windows) subarray bias reference files are created at bi-weekly intervals, as is done for the full frame biases in the Daily Monitor calibration program.
- User-defined subarrays are no longer supported.

In May 2016, partway through *HST* Cycle 23, the ACS flight software was changed to introduce a new set of WFC subarray modes that make obsolete the modes used during post-SM4 Cycles 17–23 ([ACS ISR 2017-03](#)). These subarray modes were designed to have identical readout timing pattern as the WFC full-frame mode, except with readout from only one amplifier (always 2048 columns plus pre-scan) and potentially with fewer rows (512 or 1024 rows; 2048-row subarray also available). On-orbit tests in November 2015 indicated that these revised subarray modes successfully reproduce the full-frame bias gradient structure. Bias reference files for these subarray modes are extracted from the appropriate locations of the full-frame bias reference files once again. An initial, pedestal-like offset was found in the new, **calacs**-processed subarray images between amplifiers A and C and amplifiers B and D. The offset was due to an incorrect mapping of the overscan regions in the OSCNTAB reference calibration file for amplifiers B and D. An updated version of the OSCNTAB reference file was released in February 2017 to remove the pedestal-like offset ([ACS ISR 2017-06](#)).

calacs is designed to perform bias subtraction of both pre-SM4 and post-SM4 WFC subarray images; no special directions for the calibration pipeline are needed. The pipeline initially performs a search for contemporary superbias images with the appropriate subarray dimensions and, if unsuccessful, reverts to the pre-SM4 procedure of extracting the corresponding region from a contemporary full frame superbias.

¹ Due to the way the bias reference image is created, part of the dark subtraction is also included in this step. Dark counts accumulate for an additional time beyond the exposure time, primarily the time required to read out the detector, and this portion of the dark current is subtracted along with the bias.

4.3 Dark Current, Hot Pixels, and Cosmic Rays

- 4.3.1 Dark Current
- 4.3.2 Hot Pixels
- 4.3.3 Pixel Stability
- 4.3.4 Sink Pixels
- 4.3.5 Cosmic Rays

4.3.1 Dark Current

The procedure for creating ACS dark reference files and applying dark subtraction to ACS science data is described in detail in [Section 3.4.4](#). Because it takes around one month to collect enough frames to create the dark reference files, the "best" dark reference file is typically not available in the pipeline for several weeks after the date of observation. Users may verify whether the dark reference file most appropriate to their observations has been delivered for pipeline use by checking the ACS reference file Web pages: <https://hst-crds.stsci.edu>


Using an old dark reference file will produce a poor dark correction, either leaving too many hot or unstable pixels uncorrected and unflagged, or creating many negative "holes" caused by the correction of hot pixels which were not actually hot in the science data (i.e., if the detectors were annealed in the interim).

WFC and HRC

The dark current is not constant across the CCDs. [Figure 4.7](#) shows dark current features in the WFC1 (above) and WFC2 (below) chips. These features were observed in pre-flight tests, and have generally remained stable in orbit.

There is a gradient, most noticeable on the WFC1 chip, going from a dark edge in the amplifier A quadrant (upper left) to a bright corner in the amplifier B quadrant (upper right). There are two horizontal bright bands of elevated dark current in the center of the WFC2 chip. Many faint rings are also visible, all concentric with the center of both chips. These features are likely intrinsic to the chips themselves, artifacts embedded in (or on) the silicon during various stages of the CCD manufacturing process.

Figure 4.7: High S/N Combination of WFC Dark Frames Illustrating Dark Current Structure




The WFC1 and HRC histograms in [Figure 4.8](#) and [Figure 4.9](#) show the growth of hot pixels (for more information, please refer to [Section 4.3.2](#)). A less obvious result is that the peak of the normal pixel (Poisson) distribution (i.e., the mean dark current, excluding the hot pixels) has also increased.

The increase in mean dark current for WFC has gone from 6.8 e-/pixel/hour at launch in March 2002 to 11.1 e-/pixel/hour (an average of 11.4 for WFC1 and 10.8 for WFC2) in April 2004. Following SM4, a dark current of 20 to 25 e-/pixel/hour was measured and, as of May 2017, the average dark current measured is 67 e-/pixel/hour.


For HRC the change in dark current has been from 9.3 e-/pixel/hour at launch to 13.4 e-/pixel/hour in April 2004. The dark calibration tracked the mean dark current very closely at 2 week intervals (see [Section 3.4.4](#)).

Figure 4.8: WFC Dark Current Histogram for WFC1



Growth of hot pixels over time, as measured for WFC1 in CTE-corrected superdark files. Statistics for WFC2 are nearly identical. Increases in the dark current and number of hot pixels over time were mitigated somewhat by the WFC temperature change from -77°C to -81°C on July 4, 2006. There was a substantial increase in hot and warm pixels starting in January 2015, with the introduction of post-flash to the dark calibration images, which allowed more accurate measurements of the dark frame despite CTE losses.

Figure 4.9: HRC Dark Current Histogram



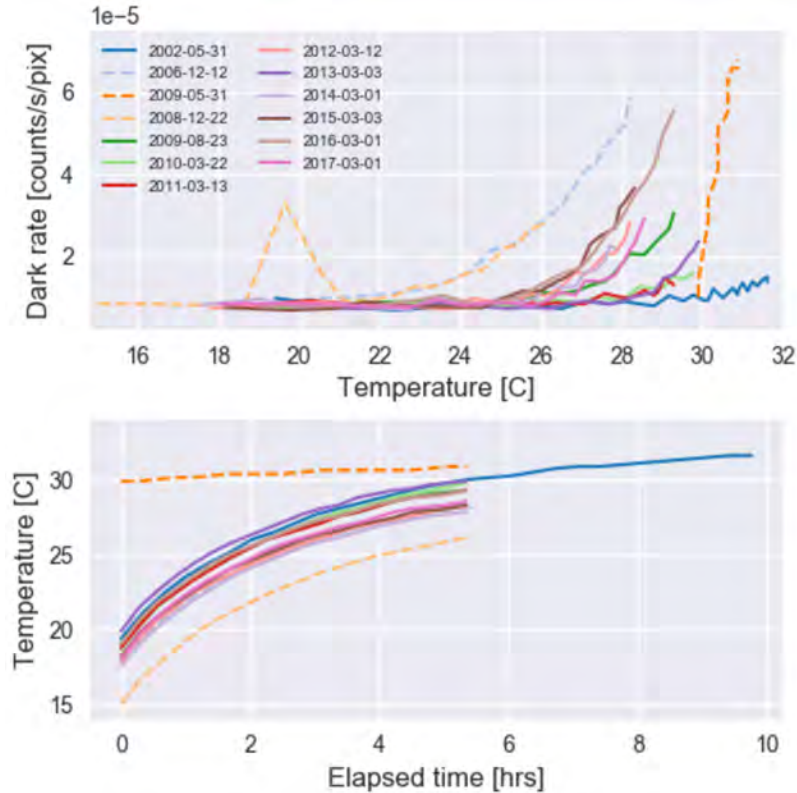
Data from the first (March 2002) and the last (January 2007) superdarks are shown. Dark current and hot pixels increased over time, as expected. Unlike [Figure 4.8](#), HRC was not affected by the temperature change in WFC.

SBC

The SBC, or ACS MAMA detector, intrinsically has no read noise and very low detector noise levels, which are normally negligible compared to statistical fluctuations. Dark frames for SBC are typically taken once per year to monitor their levels. The mean dark rate when the instrument is cooler than 25 °C is 8.11×10^{-6} counts /pixel/second ([ACS ISR 2017-04](#)).

[Summed SBC dark images](#) are delivered to [CRDS](#) from time to time. These are only valid for a period within 2 hours of the SBC turn-on, because the temperature increases with time and causes a subsequent rise in the dark current (see [Figure 4.10](#)). Measurements of the dark rate after the instrument has been on for 5 hours show a slight increase in the dark rate over time, at a rate of 2.0×10^{-6} counts/pixel/ second/year ([ACS ISR 2017-04](#)). This is driven by an elevated dark rate in the central region of the detector at warmer temperatures. However, dark correction of SBC images remains unnecessary and is not used in the calibration pipeline because the correction is negligible.

Figure 4.10: SBC Dark Rate and Operating Temperature



Top panel shows the dark rates vs. temperature measured from all dark rate monitoring programs. The bottom panel shows how the temperature changed from the time the instrument was turned on until the end of the observations. Dashed lines correspond to observations that were excluded from further analysis. (See [ACS ISR 2017-04](#) for details.)

4.3.2 Hot Pixels

When pixels are damaged by radiation or other causes, they can suffer enhanced dark current. Such pixels are called hot pixels. Although the increase in the mean dark current with proton irradiation is important, of greater consequence is the large increase in dark current non-uniformity.

Field-enhanced pixels have been classified into two categories: warm and hot pixels. The definition of "warm" and "hot" pixel is somewhat arbitrary, and there have been several changes to the definition of warm and hot pixels throughout the lifetime of ACS. In January 2015 the ranges of hot and warm pixels were adjusted as follows. A pixel above 0.14 e-/pixel/second is considered a "hot" pixel. A pixel below the hot pixel range but above 0.06 e-/pixel/second is considered a "warm" pixel. This change in definition was made after the addition of post-flash to the dark calibration images, which partially alleviates effects of CTE ([ACS ISR 2015-03](#)). The new values were chosen by comparing the hot and warm pixel percentages found in the years following SM4. In this way the hot and warm pixel flag will continue to keep track of the pixels that have the worst levels of artificial signal. Previous post-SM4 limits were set to 0.08 e-/pixel/second for hot pixels, and 0.04 e-/pixel/second for warm pixels.

Warm and hot pixels accumulate as a function of time on orbit. Defects responsible for elevated dark rate are created continuously as a result of the ongoing displacement damage on orbit. The number of new pixels with a dark current higher than the mean dark rate increases every day by few to several hundreds depending on the threshold. The reduction of the operating temperature of the WFC CCDs in 2006

dramatically reduced the dark current of the hot pixels, but over time the values have continued to rise. The smoothing effects of CTE have also increased the difficulty of accurately measuring dark current as CTE trail profiles change with time.

Table 4.4: Creation Rate of New Hot Pixels (Pixel/Day)

Threshold (e-/pixel/second)	WFC (-77°C)	WFC (-81°C)	HRC (-80°C)
0.02	815 ± 56	n/a	125 ± 12
0.04	616 ± 22	427 ± 34	96 ± 2
0.06	480 ± 13	292 ± 8	66 ± 1
0.08	390 ± 9	188 ± 5	48 ± 1
0.10	328 ± 8	143 ± 12	35 ± 1
1.00	16 ± 1	10 ± 1	1 ± 0.5

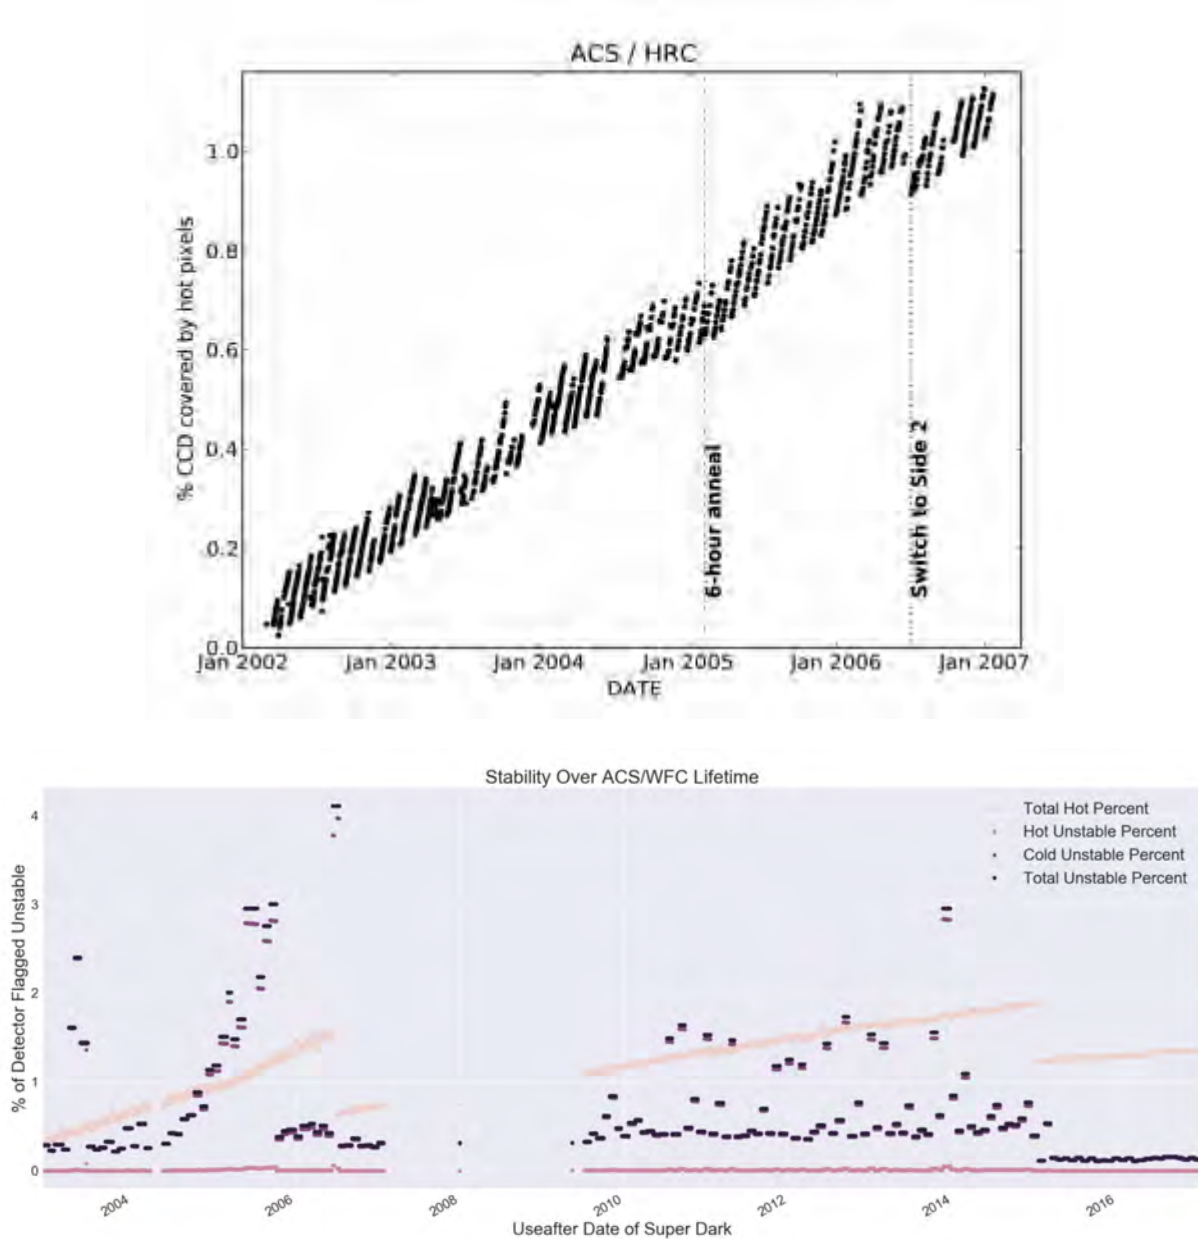
Table 4.5: Annual Permanent Hot Pixel Growth (%)

Threshold (e-/pixel/second)	WFC (-77°C)	WFC (-81°C)	HRC (-80°C)
> 0.02	1.6	n/a	1.54
> 0.04	0.78	0.32	0.52
> 0.06	0.46	0.18	0.29
> 0.08	0.30	0.16	0.21
> 0.10	0.23	0.13	0.17
> 1.00	0.03	0.02	0.02

Like other CCDs on *HST*, the ACS devices undergo a monthly annealing process. (The CCDs and the thermal electric coolers are turned off and the heaters are turned on to warm the CCDs to ~19°C.) Although the annealing mechanism at such low temperatures is not yet understood, after this "thermal cycle" the population of hot pixels is reduced (see [Figure 4.11](#)). The anneal rate depends on the dark current rate; very hot pixels are annealed more easily than warmer pixels. For pixels classified as "hot" (those with dark rate > 0.14 e-/pix/sec.) the anneal heals ~3% for WFC and ~14% for HRC.

Annealing has no effect on the normal pixels that are responsible for the increase in the mean dark current rate. Such behavior was also seen with STIS and WFC3 CCDs during ground radiation testing. Since the anneal cycles do not repair 100% of the hot pixels, there is a growing population of permanent hot pixels (see [Figure 4.11](#)).

Figure 4.11: Hot Pixel Growth Rate for HRC and WFC



These figures show hot pixel growth rates (DQ flag 16) in the WFC and HRC. Top plot: The sawtooth patterns correspond to anneal cycles. For HRC, the growth rate increased slightly when the anneal duration was reduced from 12 hours to 6 hours—a slight drop coincided with the switch to Side 2 electronics. Bottom plot: Percent of WFC detector flagged unstable as a function of superdark date. The total percentage of hot pixels, both stable and unstable, is plotted in light pink. The percentages of hot and cold unstable pixels, and their sum, are also plotted. The sharp decline in percentage of hot pixels in 2006 is due to the changing of operating temperature of the instrument. The decline in 2015 is when the darks began to be post-flashed. Only hot unstable pixels will be flagged before 2015, whereas all unstable pixels will be flagged starting in 2015.

The dark current in field-enhanced hot pixels can be dependent on the signal level, so the noise is much higher than the normal shot noise. As a consequence, since the locations of warm and hot pixels are known

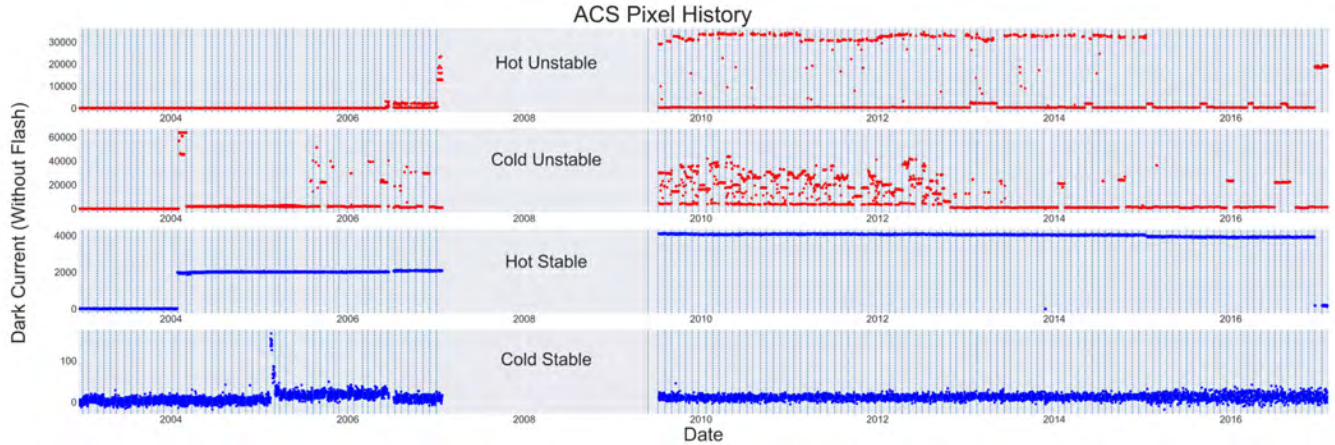
from dark frames, they are flagged in the data quality array. The hot pixels can be discarded during image combination if multiple exposures have been dithered. The standard "CR-SPLIT" approach allows rejection of cosmic rays, but hot pixels cannot be eliminated in post-observation processing without dithering between exposures.

Given the transient nature of hot pixels, users are reminded that some hot pixels may not be properly flagged in the data quality array (because they spontaneously "healed" or because their status changed in the period spanning the reference file and science frame acquisition), and therefore could create false positive detections in some science programs.

4.3.3 Pixel Stability

In principle, warm and hot pixels should be eliminated by the superdark subtraction. However, some pixels show a dark current that is not stable with time but switches between well-defined levels (see [Figure 4.12](#), see also [ACS ISR 2017-05](#)). These fluctuations may have timescales of a few minutes and have the characteristics of random telegraph signal (RTS) noise.

Figure 4.12: ACS Pixel History




Examples of the four classifications of pixel stability in plots of dark current over the lifetime of ACS. Top panel: A pixel whose high average dark current (in electrons) which varies strongly between exposures, now identified as a hot, unstable pixel. Top middle panel: An unstable pixel with normal average dark current, now identified as a cold, unstable pixel. Bottom middle panel: A hot pixel which varies very little between exposures, now identified as a hot, stable pixel. Bottom panel: A cold, stable pixel. Vertical lines are anneal boundaries. The gap in the center is when ACS was inoperable due to electronics failure. (See [ACS ISR 2017-05](#) for details.)

An analysis of the stability of every pixel in the WFC detector was performed using every dark image from the lifetime of ACS ([ACS ISR 2017-05](#)). For each anneal cycle, the stability of each pixel was determined by $f = (\text{Var}(\text{SCI}) - \text{Mean}(\text{ERR}^2)) / \text{Mean}(\text{SCI}) + 1$, where SCI and ERR are the values of that pixel in the SCI and ERR extensions, respectively, in all of the dark images taken during the anneal cycle. Pixels with f values above the threshold $f_{\text{threshold}} = 75 * \text{Mean}(\text{SCI})^{-0.75} + 2$ are considered unstable. [Figure 4.13](#) shows the f values of each pixel as a function of mean pixel value in electrons as measured from the dark frames from the Nov 18, 2015 anneal. The dotted green line marks the threshold above which pixels are considered unstable.

This analysis has shown that the vast majority of hot pixels are stable over an anneal cycle. In one example anneal from 2015, 213,999 pixels, or 1.27% of the detector, were flagged as hot, but only 257 of them, or 0.002% of the detector, are considered unstable. The total number of pixels considered unstable in this anneal is 20,113, 0.12% of the detector, so there are 19,491 normal unstable pixels.

The superdarks for each anneal will be updated to include flags of value 32 in their DQ extensions at the locations of unstable pixels. **calacs** will propagate this information into the science images it processes. Therefore, instead of discarding all hot pixels during image combination, as was done previously, stable hot pixels are retained and their noise is propagated into the ERR extensions of each image. An appropriate dither strategy between exposures is nevertheless recommended.

Figure 4.13: Density plot of mean pixel intensity versus stability for the Nov 18, 2015 Anneal




Vertical green line is the hot pixel threshold. Horizontal blue line is a stability of 1. Dotted green line is the stability threshold, everything above will be marked unstable. Note: the axes and colormap are in log space. Mean pixel value is in total electrons in a 1000.5 second dark and includes flash. (See [ACS ISR 2017-05](#) for details.)

4.3.4 Sink Pixels

Sink pixels are certain pixels in a CCD detector that are anomalously low compared to the background. This is likely due to the presence of extra charge traps in the pixel. A study of the sink pixels present in the WFC detector is presented in [ACS ISR 2017-01](#). Charge traps prevent electrons in the sink pixel itself from being read out, and they can also trap electrons from pixels that are transferred through them during the readout process, giving rise to a low-valued trail following the sink pixel. The apparent length of the trail depends on the background level of the image, as shown in [Figure 4.14](#) ([ACS ISR 2017-01](#)). In addition, about 30% of the time in WFC, a charge excess is found in the pixel immediately closer to the amplifier.

Sink pixels and the pixels they affect, both above and below the sink pixel, are identified in the average post-flashed 0.5-second WFC dark image from each monthly anneal cycle, which have been available since January 2015. One sink pixel reference image (`snk.fits`) is produced for each anneal and is used by **calacs** to flag sink pixels and the pixels they affect with the value 1024 in the DQ extension of WFC science images. Since August 3, 2017, when version 9.2.0 of **calacs** was implemented, this flagging is performed for all WFC images processed by **calacs** observed after January 1, 2015. About 0.3 to 0.5% of pixels in the WFC detector are considered sink pixels in a given anneal. Depending on the background level of an image, 1–2% of pixels will be flagged with the value 1024 in its DQ extension ([ACS ISR 2017-01](#)).

Figure 4.14: Flash-subtracted Short Dark



A 100×100 pixel² region in the flash-subtracted short dark for the 2016-09-25 anneal cycle centered on a deep SP with a trail extending towards the top of the image. Trails are visible following some of the other SPs, but many appear to be individual pixels.

4.3.5 Cosmic Rays

Like all *HST* cameras before it, the ACS HRC and WFC images are heavily peppered by cosmic rays in even the shortest of exposures. For full orbit integrations, approximately 5% of the pixels receive significant charge from cosmic rays via direct deposition or from diffusion from nearby pixels. Great care must be taken in planning and analyzing *HST* ACS observations to minimize the impact of cosmic rays on science images.

Many science observations require a careful consideration of individual cosmic ray events. To either remove cosmic rays or distinguish them from astrophysical sources, users might consider the distributions of observed cosmic ray fluxes, sizes, anisotropies, and the number of attached pixels per event.

Fractional Coverage

For most users of the HRC and WFC, the most important characteristic of cosmic rays is simply the fraction of pixels they impact. This number provides the basis for assessing the risk that the target(s) in any set of exposures will be compromised. For ACS the observed rate of cosmic ray impacts on an individual frame varies by a factor of two depending on the proximity of the spacecraft to the confluence of the Earth's magnetic field lines (e.g., the South Atlantic Anomaly). For a 1000 second exposure, the fraction of pixels affected by cosmic rays (in non-SAA passages) varies between 1.5% and 3%. This fraction is the same for the

WFC and HRC despite their factor of two difference in pixel areas because the census of affected pixels is dominated by charge diffusion, not direct impacts. This fraction is also consistent with what was observed for WFPC2.

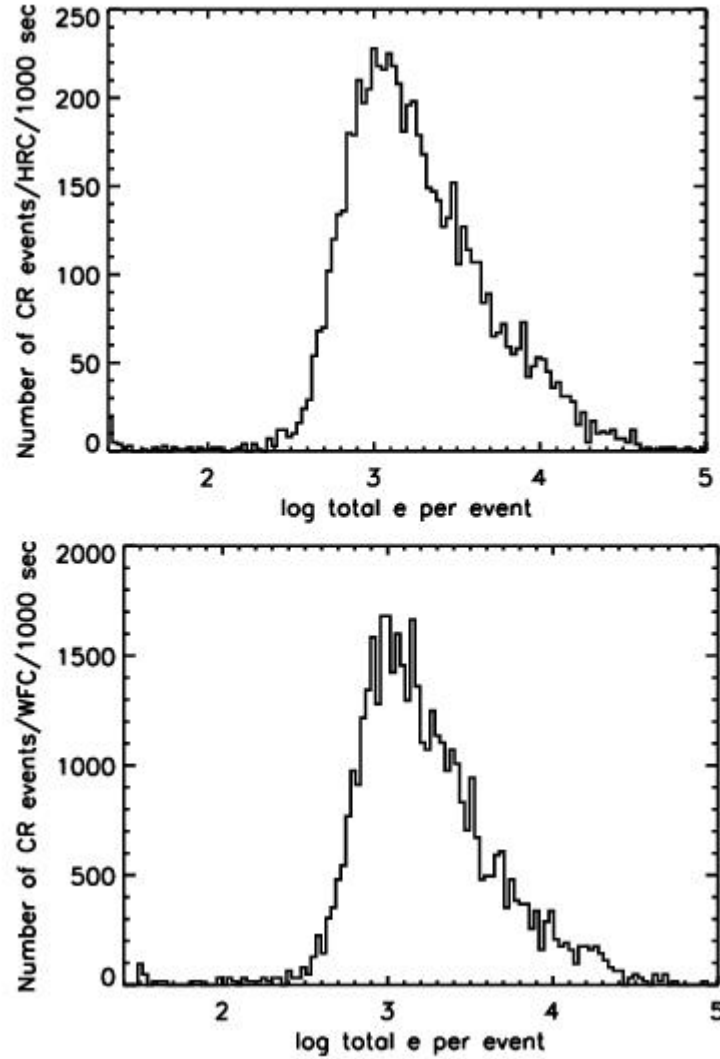
For most science observations, a single "CR-SPLIT" or dither (i.e., two exposures) is sufficient to ensure that measurements of the targets are not compromised by cosmic rays. Due to detector artifacts like hot pixels, dithered exposures are strongly recommended over a "CR-SPLIT." More consideration is required for survey-type observations with WFC, a bonafide survey instrument. Observers seeking rare or serendipitous objects as well as transients may require that every single WFC pixel in at least one exposure among a set of exposures is free from cosmic ray impacts. For the cosmic ray fractions of 1.5% to 3% in 1000 seconds, a single ~2400 second orbit must be broken into 4 dithered exposures of 500 to 600 seconds each to reduce the number of un-cleanable pixels to 1 or less.

Electron Deposition

The flux deposited on the CCD from an individual cosmic ray does not depend on the energy of the cosmic ray but rather the length it travels in the silicon substrate. As a result, the deposition distribution has a well-defined minimum with few events of less than 500 electrons (where such low-electron events correspond to cosmic rays which pass through the CCD at a normal angle of incidence). As seen in [Figure 4.15](#), the median charge deposited for WFC and HRC is about 1000 electrons, the same as for WFPC2.

A useful characteristic of the deposition distribution is its well-defined minimum; e.g., multi-pixel events which have an apparent magnitude of 25th or fainter, in a 500 second broad-band exposure, are unlikely to be caused by cosmic rays. Such information can help with the removal of false positives from searches for faint transients (e.g., high-redshift SNe).

Figure 4.15: Electron Deposition by Cosmic Rays on HRC (top) and WFC (bottom)




A minimum deposition of $\sim 500 e^-$ is seen corresponding to cosmic rays with normal incidence. The median deposition is $\sim 1000 e^-$.

Attached Pixels

As seen in [Figure 4.16](#), for HRC and WFC, the salient features of electron deposition are a strong peak in the distribution function at 4 to 5 pixels. On the smaller side there is a sharp decline in events. Although a few events are seen that encompass only one pixel, examination of these events indicate that at least some and maybe all of these sources are actually transient hot pixels or unstable pixels which can appear hot in one exposure (with no charge diffusion) and normal in the next. There is a long tail in the direction towards increasing numbers of attached pixels. Some of these are likely due to two individual events associated by their chance superposition, but more are from oblique incidence cosmic rays that skim the surface of the CCD leaving a long trail (which is wider near the surface). Unfortunately the number of attached pixels is not a very useful characteristic to distinguish cosmic rays from unresolved astrophysical sources.

Figure 4.16: Distribution of the Number of Pixels Associated with a Single Cosmic Ray Event for the HRC (top) and WFC (bottom)



Some bias exists for events > 6 pixels, which may be composed of two events with chance superposition. This distribution does not account for possible charge trails left by the CTE deterioration of WFC.

4.4 Flat-Field Reference Files

- [4.4.1 Ground Flats \(P-flats\)](#)
- [4.4.2 L-flats](#)
- [4.4.3 Pipeline Flat-Field Files](#)
- [4.4.4 Earth Flat Verification for the HRC](#)
- [4.4.5 Sky Flats for the WFC](#)

The flat-field reference files currently in use by the ACS calibration pipeline were derived via different methods, depending on the detector and the filter used. These "LP-flats" are a combination of a "P-flat," which accounts for the pixel-to-pixel variations in sensitivity, and of an "L-flat," which models the low-frequency variations in sensitivity over the detector field of view. CCD LP-flats were created in the laboratory using an external illumination source and corrected in-flight from dithered stellar observations of 47 Tucanae. For the SBC, P-flats were derived from in-flight internal lamp exposures, while the SBC L-flats were derived from dithered NGC 6681 exposures. (Flat field files in CRDS have the suffix `pf1.fits` despite being combinations of L and P flats. See *doFlat* in [Section 3.4.4](#).)

4.4.1 Ground Flats (P-flats)

CCDs

In early 2001, flat-field images for the ACS were produced in the laboratory (see [ACS ISR 2001-11](#)) using the Refractive Aberrated Simulator/Hubble Opto-Mechanical Simulator (RAS/HOMS). The RAS/HOMS is a *HST* simulator capable of delivering OTA-like external monochromatic point source and broad-band full field illumination above its refractive cutoff wavelength of ~ 3500 Å.

Because the RAS/HOMS optics are opaque below 3500 Å, flats for the UV filters F330W and F344N (see [ACS ISR 2003-02](#) and [ACS ISR 2005-12](#)) were created using in-flight observations of the bright Earth (see [Section 4.4.4](#)). Unfortunately, red leaks in F220W and F250W are so large that the out-of-band light dominates, and the lab flats made with the deuterium lamp illumination are superior to the Earth flats for these two filters.

A total signal of about 100,000 electrons per pixel is required for each flat field to avoid degrading the intrinsic pixel-to-pixel rms response of $< 1\%$ for the ACS CCD detectors. The flats are normalized by dividing by the average number of counts in the central 1% of the frame. In the case of full WFC frames, the WFC2 images are divided by the WFC1 central value in order to preserve the overall sensitivity difference between the two CCD chips across the ~ 50 pixel gap that separates the two independent pieces of the WFC detector. For small filters that cover just part of one quadrant of the WFC (i.e., F892N, polarizers), flats were masked to unity below 90% of the central value of the data, so that no flat-field correction is done on the scattered light outside the physical edge of the filter.

More information on the HRC and WFC ground flats may be obtained from [ACS ISR 2001-11](#) for the standard filters, polarizers, and coronagraph, from [ACS ISR 2002-01](#) for the ramp filters, and from [ACS ISR 2002-04](#) for the prism and grism. The stability of P-flats is tested in each cycle. Please see [ACS ISR 2007-01](#) for more information about the stability of P-flats after the cooldown to -81°C in 2006. Observations made during SMOV SM4 show that WFC P-flats are stable.

MAMA

Flat fields for the full set of SBC filters were also taken in the laboratory ([ACS ISR 1999-02](#)), but were later replaced with in-flight observations using the internal deuterium lamp ([ACS ISR 2005-04](#) and [ACS ISR 2016-02](#)). Analysis of original laboratory flats indicates that the P-flat response is independent of wavelength, so the

F125LP lamp flat was used for all filters. The internal lamp illumination does not simulate the OTA optics and, therefore, is useful only for correcting the pixel-to-pixel detector response. In order to accurately model the low-frequency variations in sensitivity, which may depend on wavelength, dithered star field observations were required. (See [Section 4.4.2](#) for a discussion of the SBC L-flats.) A study of the changes in the SBC P-flat finds that random pixel-to-pixel fluctuations have been small, but there are changes in the larger-scale fringing patterns that warrant the production of a new P-flat in the near term. In addition, the internal deuterium lamp is degrading with time, and currently produces about 65% of the original brightness ([ACS ISR 2016-02](#)).

4.4.2 L-flats

CCDs

The large scale uniformity of the WFC and HRC detector response, as provided by the CCD laboratory flats, has been improved in-flight by using multiple dithered pointings of stars in 47 Tucanae. By placing the same stars over different portions of the detector and measuring relative changes in brightness, low frequency spatial variations in the response of each detector have been measured (including the pixel-area corrections discussed in [Section 3.4.4](#)). Photometric errors of $\pm 3\%$ to $\pm 9\%$, corner-to-corner, have been found in the original WFC and HRC laboratory flat fields (see [ACS ISR 2002-08](#)). The derived L-flats are based on a 4th-order polynomial fit and are shown in [Figure 4.17](#) and [Figure 4.18](#) for WFC and HRC, where white indicates that the photometry produced using the laboratory flats is too faint with respect to the true stellar magnitude, and black indicates that the photometry is too bright. There is a continuous gradient in the L-flat correction along the diagonal of the detector which corresponds to the axis of maximum geometric distortion.

L-flats were determined from in-flight observations using filters F435W, F555W, F606W, F775W, F814W, and F850LP for both the WFC and HRC. The HRC study included two additional filters: F475W and F625W. The L-flat correction for the remaining filters was derived by using linear interpolation as a function of wavelength. The pivot wavelength of each filter was used for the interpolation, where the resulting L-flat is equal to the weighted average of the L-flat for the two filters nearest in wavelength. Due to red leaks in the F220W and F250W HRC filters, no L-flat correction has been applied, and errors in the flats of $\pm 2\%$ to $\pm 4\%$, corner-to-corner, are expected for these filters. For a detailed discussion of ACS L-flat corrections, refer to [ACS ISR 2002-08](#). For a discussion of the mathematical algorithm used to derive the L-flats, refer to [ACS ISR 2003-10](#).


Following the recovery of ACS with Side 2 electronics in July 2006, the temperature set-point was lowered from a nominal value of -77°C to -81°C for the WFC in order to minimize impacts of the continuously growing hot pixels.

The required L-flat changes range from $\sim 0.6\%$ peak-to-peak at F435W to 0.15% at F814W (see [ACS ISR 2006-06](#)). Revised L-flats were created and delivered, post-cooldown, in July 2006.

MAMA


As was done for the CCDs, the SBC L-flats were derived using dithered star field observations. Instead of 47 Tucanae, however, the UV-bright globular cluster NGC 6681, which is rich in blue horizontal branch stars, was selected. This work is summarized in [ACS ISR 2005-13](#) and [ACS ISR 2019-04](#). The required corrections to the in-flight lamp flats are given in [Figure 4.19](#) and range by $\pm 15\%$, depending on wavelength. Since insufficient observations with the F122M filter exist, this L-flat is simply a copy of the F115LP filter correction. Six new flat fields have been delivered for use in the calibration pipeline, and the resulting photometric accuracy is now $\pm 2.5\%$ for F115LP, F122M, F125LP, F140LP, and F150LP, and $\pm 3.3\%$ for F165LP.

Figure 4.17: WFC Low Frequency (L-flat) Flat-Field Corrections Required for the Laboratory Data.




While the pixel-to-pixel (P-flat) structure of the laboratory flats is robust, a low frequency correction is required to achieve uniform detector response. This correction ranges from $\pm 5\%$ for the F555W filter to $\pm 9\%$ for the F850LP filter.

Figure 4.18: HRC Low Frequency (L-flat) Flat-Field Corrections Required for the Laboratory Flats



This correction ranges from $\pm 3\%$ for the F555W filter to $\pm 6\%$ for the F850LP filter.

Figure 4.19: SBC Low Frequency (L-flat) Flat-Field Corrections Required for the In-flight Deuterium Lamp Data



4.4.3 Pipeline Flat-Field Files

The inferred L-flat corrections were multiplied by the corresponding P-flats, and the resulting flat fields have been in use in the calibration pipeline. Observers are encouraged to check the ACS website for the most recent flat fields available for recalibration at:

<https://hst-crds.stsci.edu>

[Figure 4.20](#) shows the corrected WFC laboratory flats for several broad-band filters. Note that on the sky, a gap of ~50 pixels exists between the top and bottom halves that is not shown here. The central donut-like structure is due to variations in chip thickness (see [ACS ISR 2003-06](#)) and is dependent on wavelength. Pixels in the central region, for example, are less sensitive than surrounding pixels in the blue F435W filter, and more sensitive in the red F850LP filter.

For the HRC, corrected laboratory flats for the same broad-band filters are shown in [Figure 4.21](#). The donut-like structure seen in the WFC response is not found in the HRC flats. The small dark rings are shadows of dust on the CCD window (see [Section 4.5.1](#)). The large dust mote seen in the WFC F606W flat is due to dust on the F606W filter. That portion of the filter is not "seen" by the HRC detector.

Because of the observed changes in P-flats discussed in the previous section, two versions of the LP-flat exist in the pipeline for each SBC filter. The set used for observations after 2007 are presented in [Figure 4.22](#).

Because of geometric distortion effects, the area of the sky seen by a given pixel is not constant; therefore, observations of a constant surface brightness object will have count rates per pixel that vary over the detector, even if every pixel has the same sensitivity. In order to produce images that appear uniform for uniform illumination, the observed flat fields include the effect of the variable pixel area across the field. A consequence of dividing by these flat fields is that two stars of equal brightness do not have the same total counts after the flat-fielding step. Therefore, point source photometry extracted from a flat-fielded image (`flt.fits/flc.fits` and `crj.fits/crc.fits`) must be multiplied by the effective pixel area map (see [Section 5.1.3](#)). Alternately, users may use images processed by **AstroDrizzle** (with suffix `drz.fits/drc.fits`) where a geometric distortion solution has been applied to correct all pixels to equal areas. In drizzled images, photometry is correct for both point and extended sources.

Accuracy of pipeline flats can be verified using a variety of complementary methods. [Section 4.4.2](#) explained how follow-up observations of the same stellar field can be used to verify the derived L-flat corrections. Alternately, observations of the bright Earth can provide a uniform flat-field source for the complete OTA optical complement and incorporate both the low frequency L-flat and the high frequency pixel-to-pixel P-flat response. Earth flats are described in [Section 4.4.4](#). For most filters, the flats agree to within ~1%, except for the interpolated L-flat filters which usually agree to within ~2%.

The third method for verifying the ACS pipeline flats is with sky flats. These can be made by filtering and summing many observations of a sparse field. Sky flats have been created for a few of the most frequently used broad-band WFC filters and are discussed in detail in [Section 4.4.5](#). The sky flats are generally similar to corrected ground flats at the 1% level, in accordance with the results of the previous two methods. While the WFC can show residuals in the central donut-region which are as large as 2%, these are most likely due to differences in the color of the spectrum of the sky from that of the bright globular cluster stars used for the L-flat determination.

To summarize, the pipeline flats were created by correcting the pixel-to-pixel flats by low-frequency corrections derived from dithered stellar observations. For the most used modes, the flats are accurate to better than 1% across the detector.

Figure 4.20: WFC Flat Fields




Figure 4.21: HRC Flat Fields





Figure 4.22: SBC Flat Fields



4.4.4 Earth Flat Verification for the HRC

Because the RAS/HOMS optics are opaque below 3500 Å (see [Section 4.4.1](#)), flat fields for the HRC UV filters were created using in-flight observations of the bright Earth. The Earth is a poor flat-field source at optical wavelengths because structure in the cloud cover can cause streaking in the flat. However, HRC modes that utilized the F220W, F250W, F330W, and F344N were immune to streaks because of the large optical depth down to the tropospheric cloud layers. The bright Earth then provided a uniform flat-field source for the complete OTA+HRC optical complement.

The required calibration flats, which incorporates both the low frequency L-flat and the high frequency pixel-to-pixel P-flat response, can be easily produced from these observations. Unfortunately, the red leaks in F220W and F250W are so large that the out-of-band light dominated, and the lab flats made with the deuterium lamp illumination (see [ACS ISR 2001-11](#) for details) are superior to the observed Earth flats for the modes that included these two filters. Because no L-flat correction has been applied, errors in the flats of $\pm 2\%$ to $\pm 4\%$, corner-to-corner, are expected for these filters. HRC F330W and F344N pipeline flats, on the other hand, are defined entirely by Earth flat data (see [ACS ISR 2003-02](#) and [ACS ISR 2005-12](#)). With ~ 20 – 30 observations, each over the course of 3 years, these flats have very high signal-to-noise and showed repeatability to much better than the required 1% accuracy.

HRC Earth flats at wavelengths longward of 4000 \AA often showed streaking and non-uniformities from clouds or the terminator. However, a number of the images are free from these defects and provided an independent verification of the stellar L-flat correction technique at visible wavelengths. Unfortunately, WFC Earth flats suffer from a shutter light leak, and the Earth limb is too bright for SBC observations.

Several hundred observations of the bright Earth were obtained using the full set of HRC standard filters ([ACS ISR 2005-12](#)). In general, the pipeline flat fields are confirmed to a precision of $\sim 1\%$, validating the stellar L-flat corrections. The "interpolated" L-flats are not significantly worse than the L-flats derived from the Earth observations (see [Section 4.4.2](#)). One exception is the F550M filter which shows a total deviation of more than 2%. Other exceptions are the four longest wavelength HRC filters which show large systematic differences with the pipeline flats. These differences appear to be caused by stray light originating from the detector surface, where most of the long wavelength photons were reflected and then scattered back from nearby focal plane structures. Any filter transmitting at these long wavelengths would have seen the extra pattern from this light, though the strength of the additional stray light is proportional to the total flux of the source. Thus, for large diffuse objects that fully illuminated the detector, these Earth flats are more appropriate for calibration than the existing pipeline flats, which are appropriate for point sources.

4.4.5 Sky Flats for the WFC

The ACS team has made use of the extensive imagery from the Cycle 13 GOODS survey (Programs 9425 and 9583; PI M. Giavalisco) to construct high signal-to-noise sky flats. These sparsely populated, high Galactic latitude exposures have relatively uniform sky that can be stacked to further quantify the pixel-to-pixel variation of the instrumental response in the WFC F606W, F775W, and F850LP filters with some 50–70 images apiece.

The sky flats were created by median-combining the pipeline-reduced `flt.fits` files after removing cosmic rays and masking all of the sources. Because the GOODS data contained 2–4 dithers at each pointing, masking was necessary to eliminate the high values at each pixel. In addition, before combining, each image was corrected for the pedestal bias signature and inspected for scattered light.

Because sky flats are created from the pipeline calibrated `flt.fits` files, any flat-field signatures that are not in the pipeline flats should appear in the ratio. The combined sky flat in each filter was box-medianaed for comparison to the corrected ground flat. The resulting ratios of pipeline to sky flats show variations across the field of view of $< 2\%$ for each of the three filters. The sky flats independently created from the two separate GOODS fields, showed excellent agreement ($< 1\%$) for all three filters.

4.5 Image Anomalies

- [4.5.1 Dust Motes](#)
- [4.5.2 Optical Ghosts and Scattered Light](#)
- [4.5.3 Cross-Talk](#)
- [4.5.4 Scattered Earth Light](#)
- [4.5.5 HRC Polarimetry](#)

This section presents an overview of image anomalies in ACS images. Additional information is available in [HLA ISR 08-01](#) and [ACS ISR 2016-06](#). Examples of scattered light anomalies characterized in [ACS ISR 2016-06](#) may be viewed using the [dragon's breath](#) and [guidestars](#) interactive web tools.

4.5.1 Dust Motes

Several circular patterns consisting of a dark ring with a bright center, are visible in the ACS flat fields, with typical diameters of ~30 pixels on HRC and ~100 pixels on WFC. These artifacts are shadows of dust on the CCD windows and are weaker on the $f/25$ WFC than on the $f/68$ HRC. The motes can be seen in WFC and HRC flats in [Figure 4.20](#) and [Figure 4.21](#).

Since the shapes and depths of these motes are almost independent of wavelength, their effects will be removed by the flats to $\ll 1\%$, unless any of these particulate contaminants move to different positions on the CCD windows. In case of particulate migrations, the internal lamp flats have a lower f -ratio with a wider angular distribution and cannot be used to patch the flat fields because they wash out the mote shadows. To correct for new motes, patches to the pipeline flats must be made using the original laboratory flats, corrected for the low-frequency flats derived in-flight or, for short wavelengths, using observations of the bright Earth.

Larger motes are sometimes present due to dust and blemishes on several ACS filters, including F606W, POL0V, and POL60V. Because the filters are located farther from the detector windows in a converging light beam, imperfections on the filters produce an out-of-focus image at the detector, where typical mote diameters are about 350–400 pixels on the WFC and about 250–300 pixels on the HRC. One of these large motes can be seen on chip 2 of the WFC F606W flat in [Figure 4.20](#).

Until April 2004, the positioning accuracy of the filter wheels has been within ± 1 motor steps of the nominal position. This delta corresponds to a distance on the detectors of ~18 HRC pixels and ~20 WFC pixels. Features with sharp transmission gradients at the filter wheels cause a corresponding flat-field instability, where errors are 1%–3% for a few pixels near the blemishes. If the filter wheel lands in a different place the dust mote will move. Blemish mis-registration is an error in the pixel-to-pixel high frequency component of the flat fields and is not related to the low-frequency L-flat correction, which has been applied for all standard and polarizing filters. For details on the L-flat correction, see [Section 4.4.2](#).

This problem was recognized and addressed before launch by a laboratory calibration campaign to obtain flat fields at the nominal position and at plus and minus one step for the F606W + CLEAR on the WFC and for the two POLV filters in combination with the highest priority F475W, F606W, and F775W filters on both HRC and WFC. Since the resolver position uniquely determines the filter wheel step, the ACS pipeline data processing has been enhanced to automatically apply the proper flat for the wheel step position. The keyword `FWOFFSET` has been added to the ACS image headers to indicate the position of the filter wheel. In April 2004, an update was made to the ACS flight software and the filter wheel is now always positioned at its nominal position. For more information on flat fields for filter wheel offset positions, refer to [ACS ISR 2003-11](#).

4.5.2 Optical Ghosts and Scattered Light


ACS was designed with a requirement that no single stray light feature may contain more than 0.1% of the detected energy in the object that produced it. This goal has generally been met, but during extensive ground and SMOV test programs, a few exceptions have been identified ([Hartig et al. 2003](#)) such as the WFC elliptical haloes and the F660N ghosts. While these ghosts exceed the specified intensity, their origin and characteristics are well defined and they should have minimal impact on the ACS science program. Details about some optical ghosts relevant for ACS are described below.

- WFC elliptical haloes: these ghosts are caused by reflection from the CCD surface (which lies at an $\sim 20^\circ$ angle to the chief ray) up to the detector windows and back to the CCD. They show up as pairs of elliptical annuli, aligned along the negative diagonal of the FOV (see [Figure 4.23](#)), and are observed when bright sources are placed on the lower right (D amplifier) quadrant of the WFC detector. The surface brightness of the annuli increases and size decreases with proximity to the corner. Two pairs of ghosts are seen, produced by reflection from the four window surfaces. The total energy fraction in each ghost may exceed 0.2% of the target signal.
- F660N ghosts: the F660N narrow band filter produces pairs of relatively bright circular annuli stationed near to (but radially outward from) the target image (see [Figure 4.24](#)). This is due to reflection from the two surfaces of the second "blocker" substrate back to the many-layer dielectric stack on the first substrate, which in turn reflects at high efficiency at the filter wavelength range. These haloes contain $\sim 2\%$ of the detected target energy and are always about 10 and 20 pixels in diameter.
- Annular ghosts: large annular ghosts¹ near their parent images are caused by reflection from the detector windows, back to the filters, then returning to the CCD (see [Figure 4.23](#)). Another type of annular ghost arises from reflections between the inner and outer window surfaces; these are much smaller in diameter, relatively low in intensity (well within the specification) and are displaced radially from the parent image by a small amount.
- Glint: a star that falls close to the gap between the two WFC CCDs can sometimes create a reflection known as a "glint," thin rays of light extending from the chip gap, as shown in [Figure 4.25](#). These rays are the reflection of starlight off the residual indium solder used for chip attachment.
- "Dragon's breath" is caused by light reflections involving the knife-edged mask in front of the CCD detector. This anomaly typically appears as a narrow spike, appearing like a tongue of flame, extending from the edge of the frame (see #4 in [Figure 4.23](#)). Edge glow is caused by a similar effect, however produces somewhat round and diffuse bright area on the edge of the frame. See [ACS ISR 2016-06](#) for more details of these anomalies.

In general, little can be done about these anomalies in the post-observation data processing phase. Instead, some judicious planning of the actual observations, particularly if bright sources are expected in, or near, the field of view, is recommended. For instance, the impact of diffraction spikes (which for ACS lie along X and Y axes) and of CCD blooming (which occurs along the Y direction) due to a bright star, can be reduced by choosing an ORIENT² which prevents the source of interest from being connected to the bright star along either of these axes. Similarly, the impact of WFC elliptical haloes can be minimized by avoiding a bright star in the quadrant associated with amplifier D.

(See the [dragon's breath](#) and [guidestars](#) interactive web tools and [ACS ISR 2016-06](#).)


Figure 4.23: Optical Ghosts, Diffraction Spikes, and Blooming



This image was obtained during ground calibration at Ball Aerospace using a HeNe laser (633 nm) through F625W ([Hartig, et al](#)). There are nine highly saturated point sources in and just off of the field of view. Note the diffraction spikes along the X and Y axes of the WFC, and the severe blooming of the charge along the Y -axis. Several optical ghosts with different origins and intensities (see text below for details) are visible.


- (1) WFC elliptical haloes show up in the lower left as pairs of elliptical annuli aligned along the negative diagonal of the FOV;
- (2) large annular ghosts are seen near their parent images;
- (3) to the lower right, there exist smaller fainter annular ghosts which tend to be displaced radially from the parent image by a small amount;
- (4) "Dragon's Breath" is a shower of scattered light from a very bright star that is just off the edge of the CCD. This rare anomaly occurs when a star falls at the edge of the mask in front of the chip; the starlight reflects off the CCD, then off the mask, and back to the detector.

Figure 4.24: F660N Optical Ghosts



This 700 second F660N exposure illustrates the F660N ghosts, namely pairs of relatively bright circular annuli stationed near to (but radially outward from) the target image. These haloes contain ~2% of the detected target energy and are always about 10 and 20 pixels in diameter.

Figure 4.25: Glint



Thin rays of light extending from the chip gap are sometimes caused by the reflection of light from a star that falls close to the inter-chip gap. See [Hartig \(2002\)](#) and [HLA ISR 2008-01](#) for more on glints from stars on edges of the ACS/WFC chip gap.

4.5.3 Cross-Talk

The ACS/WFC detector has four amplifiers (A, B, C, D; see [Figure 1.1](#)) through which the four quadrants of the detector are read separately and simultaneously. As the quadrants are read out, electronic cross-talk between the amplifiers can be induced. As a result, an imaged source in one quadrant may appear as a faint, mirror-symmetric ghost image in the other quadrants. The ghost image is often negative; therefore, bright features on the "offending" quadrant show up as dark depressions on the "victim" quadrants. See [Figure 4.26](#) and [Figure 4.27](#) for actual examples.


A complete study of this issue can be found in Suchkov, A., et al. 2010, [ACS ISR 2010-02](#).

The latest **calacs** corrects for cross-talk in post-SM4 full frame WFC images as part of the **doBlev** stage (see [Section 3.4.1](#)).




*In general, the conclusions from the ACS/WFC cross-talk characterization are as follows:
After SM4, cross-talk due to low signal sources is much weaker than before SM4, and can be safely ignored for both GAIN = 1 and GAIN = 2.
For the current default GAIN = 2, the post-SM4 cross-talk is insignificant ($< \sim 0.0001\%$ (between amplifier pairs on different WFC CCDs, but intra-CCD amplifier cross-talk is relatively uniform at $\sim 0.009\%$ and is removed by calacs in the 'doBlev' stage.*

Figure 4.26: Cross-talk in Pre-SM4 ACS Image of NGC 4710



This image (from program ID 10594, image j9ew02w1q, GAIN = 2) shows the galaxy in quadrants C and D. Its cross-talk ghosts, seen as dark oval shapes in quadrants A and B, are due to low signal offending sources of $\sim 100\text{ e}^-$ to 1000 e^- in the area of the reddish rim of the galaxy. Also the ghosts of three largest galaxies in quadrant B are easily identifiable in quadrant A. The image in quadrants A and B is stretched within a narrow signal range centered at the sky background level, which makes the ghosts distinctly stand out against the background.

Figure 4.27: Post-SM4 Image of NGC 6217



The image shows NGC6217 centered in quadrant A, obtained with GAIN = 2 (program ID 11371, frame ja7z03unq). The images in quadrants B, C, and D are stretched similarly to that in quadrants A and B in [Figure 4.26](#). No ghosts from low signal areas in quadrant A are seen in the other three quadrants.

4.5.4 Scattered Earth Light

Most observers will not experience significant issues with scattered Earth light in their observations. Normally observations are scheduled only when the bright Earth limb is more than 20° from the *HST* pointing direction. This is sufficient to eliminate serious impacts from scattered Earth light—the most severe impact will be for observers with targets in the CVZ³ who may notice the sky background increased by a factor of 2 or 3.

It is possible to make arrangements for observations at smaller bright Earth limb angles, and these images have a potential for serious impacts from scattered light. There are two types of impact: elevated background and non-uniformity in the background. For example, at a bright Earth avoidance angle of 14° , it is possible for the sky level to be increased by a factor of 100 compared to normal pointings away from the Earth; this will of course have a serious impact on the background noise and detection of faint targets. Also, non-uniformity can arise since the scattered light is taking an increasingly non-standard path through the *HST* optics, and hence the flat-fielding becomes corrupted. At this same angle of 14° , it is possible to have both large scale gradients across the field of view (up to $\sim 20\%$ amplitude in the WFC) and small scale features in the background (up to $\sim 12\%$ in WFC and $\sim 30\%$ in the HRC). See [ACS ISR 2003-05](#) for more details.

4.5.5 HRC Polarimetry

Each ACS filter is designed to maintain confocality with both the WFC and HRC when paired with a clear aperture in the other filter wheel. To maintain this confocality when filters are paired with polarizers, the polarizers were fabricated with additional lensing power which alters the pixel scale and geometric distortion

of each camera. The additional distortion is further complicated by localized optical defects (bubbles and wrinkles) in the polaroid materials (See [ACS ISR 2004-09](#)).

Polynomial solutions of the geometric distortion for each filter are used by the **AstroDrizzle** stage of the ACS calibration pipeline to produce geometrically rectified and resampled WFC and HRC images for photometric and astrometric use. These rectified images are provided as FITS images with the `drz.fits/drc.fits` extension.

Presently, distortion solutions have not been derived for HRC images obtained with the UV or visible polarizers. Such images constitute about 4% of the HRC datasets in the *HST* archive. Consequently, the `drz.fits` files produced by **AstroDrizzle** for polarized HRC images have a pixel scale that differs by ~3% from the correct pixel scale obtained for non-polarized images. A correct distortion solution will be generated before the planned creation of a static HRC image Archive. Until then, users must exercise caution when performing astrometry or surface brightness measurements with polarized HRC images.

¹ In [Figure 4.23](#), the fringes visible in the annulus arise from interference between HeNe laser light reflections from the two surfaces of the windows. This image was obtained during ground testing.

² Spacecraft roll angle specified for an observation in the Phase II proposal exposure logsheet.

³ Continuous Viewing Zone: a declination band at $\pm 61.5^\circ$ where targets may be viewed without occultations at some time during the 56-day precessional cycle of the *HST* orbit.

4.6 WFC CCD Detector Charge Transfer Efficiency - CTE

4.6.1 The Issue

4.6.2 Improving CTE: Considerations Before Making the Observations

4.6.3 Improving CTE: Post-Observation Image Restoration


4.6.1 The Issue

The ACS/WFC CCD detectors operate by the simple process of converting incoming photons into electron/hole pairs, collecting the electrons in each pixel, and then transferring those electrons across the detector array during the device readout. The transfer process moves each pixel's electrons down along the columns and then across in a transfer register to the amplifier located in the corner of the detector array.

When the detectors were manufactured, these transfers were extremely efficient (typically 0.999996 of each charge packet was transferred successfully from one pixel into the next), which means that slightly over 99% of the charge collected in a pixel would be delivered to the transfer register. Once in space, the flux of energetic particles such as relativistic protons and electrons damages the silicon lattice of the CCD detectors. This creates both "hot" pixels and charge traps. This radiation damage is cumulative and was unavoidable given current technologies for detector construction and shielding.

The charge traps degrade the efficiency with which charge is transferred from pixel to pixel during the readout of the CCD array. This is seen directly, as shown in [Figure 4.28](#), in the "charge trails" that follow hot pixels, cosmic rays, and bright stars that can extend to over 50 pixels in length.

Figure 4.28: CTE Trails



A section (800×800) of an ACS frame of 47 Tucanae. Note the presence of trails extending from the stars indicating the effect of CTE on the detector.

4.6.2 Improving CTE: Considerations Before Making the Observations

The simplest mitigation of the imperfect CTE is to reduce the number of charge transfers required for a given source to reach the readout amplifier on the CCD. If the source of interest is small ($\sim 10''$ across or less), placing it close to the corner of the detector will result in greatly enhanced net CTE. APT now has a pre-defined aperture setting for this purpose called `WFCL-CTE`.

CTE is a strong function of the signal level in the pixels through which a charge packet must pass on its way to the transfer register. Observations with very low background ($< 20 \text{ e}^-$ for ACS) will suffer large losses for very faint sources. This is likely to be problematic for narrow band filters and extremely short exposures. In these cases, raising the background will greatly improve the CTE and thus the S/N of these sources. For users planning to stack multiple images to reach very faint limits, they should plan to achieve a background level of $\sim 20 \text{ e}^-$ for ACS.

The background can be increased in several ways:

1. Longer exposure times
2. Selection of a broader filter

3. Addition of internally generated photons (i.e., "post-flash").

ACS/WFC contains LED lamps configured to illuminate the side of the shutter blade that faces the CCD detector. Designed to provide fairly even illumination at low signal level, these lamps provide a "post-flash" capability. While the post-flash lamp can be used to increase the background in an image, it is recommended with reservations due to a 50% variation in the signal across both chips. See [ACS ISR 2014-01](#) and the [ACS website](#) for more details.

4.6.3 Improving CTE: Post-Observation Image Restoration

The ACS team has developed and implemented a post-observation correction algorithm based upon the Anderson and Bedin (2010, [PASP](#), 122, 1035) methodology. This empirical algorithm is based on a model for charge-transfer loss and release that reproduced the observed trails behind warm pixels. The correction software then uses an iterative forward-modeling process to estimate the source image from the observed trailed image.


The original version of the code worked very well for intermediate to high flux levels (> 200 electrons), but data were not available at the time to test it at lower flux levels. Several ISRs describe the original correction: [ACS ISR 2011-01](#), [ACS ISR 2012-03](#), and [ACS ISR 2012-04](#).

Recently, we have taken calibration data that will allow a higher fidelity model at all flux levels. These data have been used to re-parameterize the model and make additional improvements (such as read-noise mitigation). The new model is now available as a part of the standalone pipeline and will soon be available via MAST. An ISR (Anderson 2017, in prep.) describing the new correction and evaluating its performance with on-sky tests is in progress and should be available before the end of 2017.

While pixel-reconstruction algorithms may do a good job removing trails behind stars, cosmic rays, and hot pixels, they have one serious and fundamental limitation: they cannot restore the lost S/N in the image. This limitation notwithstanding, the reconstruction algorithm provides the best understanding of the "original" image before the transfer, and also helps understand how the value of each pixel may have been modified by the transfer process. This algorithm is available in the ACS pipeline; standard calibrated products are now available both with and without this correction.

In general, we find that the correction is good to about 25% for stars with moderate signal to noise, so one can get a sense of the reconstruction error by determining the amplitude of the correction and taking 25% of that as the error. The upcoming ISR ([ACS ISR 2018-04](#)) should provide some empirical demonstrations of the efficacy of the correction.

Figure 4.29: An Example of the Pixel-Based CCD Corrections



(Left) A 1000 x 1000 pixel region at the top of the chip 1 extension in image `jbmncOakq_flt.fits`. CTE vertical trails are clearly visible. **(Right)** The reconstructed CTE-corrected `flc.fits` image after the execution of `calacs`. An alternate method for post-observation restoration involves a simple recalibration of the photometry using correction curves that have been provided in Chiaberge, M. ([ACS ISR 2012-05](#)). This can be effective for isolated point sources on flat backgrounds, but is less effective for extended sources or sources in crowded regions. Please refer to [Section 5.1.5](#) for more details

The expected losses should be taken into consideration when one is deciding on the best CTE-mitigation strategy, which may involve taking fewer longer exposures to preserve S/N even with the increased cosmic-ray contamination.

4.7 Generic Detector and Camera Properties

- [4.7.1 Gain Calibration](#)
- [4.7.2 Full Well Depth](#)
- [4.7.3 Linearity at Low to Moderate Intensity](#)
- [4.7.4 Linearity Beyond Saturation](#)
- [4.7.5 Shutter Stability](#)

4.7.1 Gain Calibration

The ACS CCDs have default gain values of approximately 2 e-/DN for both the WFC and HRC (prior to Cycle 14 the WFC default was 1 e-/DN). The relative values of gains 1 and 2 for the WFC are determined to an accuracy of about 1 part in 10,000. Although the original default settings were used to establish adjustments to basic quantum efficiency curves, current gains have consistent values. Minor errors in the default gain values would therefore be absorbed into revisions of the QE curves. There is no compelling reason to seek refinement in the quantitative values of the default used for primary calibrations and the bulk of the science program.

The default gain for the WFC does sample the full well depth of the CCDs. However, for HRC, the default fell short by 22%. Use of the next higher gain value of approximately 4 e-/DN for the HRC provided full sampling of the 165,000 e- full well depth; since the read noise was only marginally higher than with the default gain, and the read noise for the HRC is nearly critically sampled, even at the higher gain value, (much more so than any WFPC2 data ever taken), many science and calibration programs may have logically chosen the higher gain.

ACS gain values in use between on-orbit installation (in March 2002) and January 2004 can be traced to ground test results. The primary tool for measuring gain values is the photon-transfer method which is described by [Janesick, J.R. \(2001\)](#) in Scientific charge-coupled devices, Bellingham, WA: SPIE Optical Engineering Press, 2001, xvi, 906 p. SPIE Press monograph, PM 83.

This technique relies on analysis of pairs of identical exposure flat fields taken at a range of intensity levels. Over an ensemble of pixels at a given exposure level, the relation between differences of intensity values (noise) and the direct signal level depends on the read noise and the gain. This relation can be fit at a range of intensity levels to uniquely determine these two quantities. This technique can also provide limited information on linearity¹ and the saturation count level² by determining where the photon-transfer curve departs from linearity at high signal levels. Errors of about 0.6% in the gain values are quoted for the WFC determinations, and similar values hold for the HRC.

On the WFC, with two CCDs and the default use of two readout amplifiers per CCD, errors of 0.6% in the normalization of gains, quadrant-to-quadrant, would leave easily visible steps in observations of spatially flat sources. A re-determination of the WFC gains ([ACS ISR 2002-03](#)), maintaining the same mean over all amplifiers at a given gain setting but using a continuity constraint across quadrant boundaries, provided an improvement in amplifier-to-amplifier gains at better than 0.1% for the default gains. Any steps of intensity seen in data across quadrant boundaries are likely to reflect minor errors in bias and overscan values (see [Section 4.2](#)), rather than errors in relative gains.

As documented in [ACS ISR 2004-01](#), on-orbit data was used to redetermine the mean absolute gain values relative to the standard adopted for the default gains on each camera, e.g., WFC GAIN = 1 and HRC GAIN = 2 values were retained and other gains adjusted relative to these. Through analysis of observations of the same stellar field, in the same filter and at the two gain settings, it is possible to obtain accurate adjustments

of these gain values relative to the default gain levels. The results discussed in the above-mentioned ISR supply these to better than 0.1%, removing errors that averaged about 1% for the two cameras prior to updates made in January 2004.

Absolute errors of about 0.6% could remain in the default gain values, but these are of no real consequence since redetermination of the CCD quantum efficiency would compensate for this, since the QE adjustments were based on data acquired with the default gains. Basic photometric calibrations apply with equal accuracy to data acquired at all supported gains (assuming for the non-default gains that reprocessing was done after January 6, 2004 to include the revised values).

Tables 4.6 and 4.7 show the CCD gain values for pre-SM4 ACS operations, based on pre-launch and on-orbit calibrations.

Table 4.6: Pre-SM4 HRC Gain Values based on Pre-launch and On-orbit Calibrations

Amplifier	GAIN	Prelaunch Calibration (e ⁻)	On-orbit Calibration (e ⁻)
C	4	4.289	4.235
C	1 ¹	1.185	1.163

¹ GAIN = 1 was "available but unsupported"

Table 4.7: Pre-SM4 WFC Gain Values based on Pre-launch and On-orbit Calibrations

Amplifier	GAIN	Prelaunch Calibration (e ⁻)	On-orbit Calibration (e ⁻)
A	2	2.018	2.002
A	4 ¹	4.005	4.011
B	2	1.960	1.945
B	4 ¹	3.897	3.902
C	2	2.044	2.028
C	4 ¹	4.068	4.074
D	2	2.010	1.994
D	4 ¹	3.990	3.996

¹ GAIN = 4 was "available but unsupported"

Following SM4, the gains available with the new electronics are, in units of e⁻/DN, 0.5, 1.0, 1.4, and 2.0. Only GAIN = 2.0 is currently supported. New relative gain values were measured and are discussed in [ACS ISR 2009-03](#). The average absolute gain remains unchanged.

4.7.2 Full Well Depth

Conceptually, full well depths can be derived by analyzing images of a rich star field taken at two significantly different exposure times, identifying bright but still unsaturated stars in the short exposure image, calculating which stars will saturate in the longer exposure and then simply recording the peak value reached for each star in electrons (using a gain that samples the full well depth, of course). In practice, as discussed in [ACS ISR 2004-01](#), it is also necessary to correct for a ~10% "piling-up " effect of higher values being reached at extreme levels of over-saturation relative to the value at which saturation and bleeding begins to neighboring pixels in the column.

Since the full well depth may vary over the CCDs, it is desired to have a rich star field observed at a gain that samples the full well depth, and for which a large number of stars saturate. Calibration programs have serendipitously supplied the requisite data of rich fields observed at two different exposure times.

High Resolution Channel

The HRC shows large scale variation of about 20% over the CCD. The smallest full well depth values are at about 155,000 e- and the largest at about 185,000 e-, with 165,000 e- representing a rough estimate at an area-weighted average value. See Figure 3 of [ACS ISR 2004-01](#) for details.

Wide Field Channel

As with the HRC, there is a real and significant large scale variation of the full well depth on the WFC CCDs. The variation over the WFC CCDs is from about 72,000 e- to 82,000 e- with a typical value of about 77,000 e-. There is a significant offset between the two CCDs. The spatial variation may be seen as Figure 4 in [ACS ISR 2020-02](#) and Figure 4 in [ACS ISR 2004-01](#).

4.7.3 Linearity at Low to Moderate Intensity


High Resolution Channel

Rich star fields observed at quite different exposure times provide a simple, direct test for linearity. In [Figure 4.30](#) the results of such tests are shown in two ways. The first is a simple plot of aperture sum values in the long exposures versus the same stars on the same pixels in the short exposures—no deviations from linearity are evident.

For a more sensitive test the bottom plot shows the results of summing counts over all stars within a defined magnitude range in the short and long exposure cases, separately, before then taking the ratio and normalizing to the relative exposure times.

The linearity of the HRC at low and moderate intensity levels, as evidenced by comparing stars observed with exposures differing by a factor of 60, appears to be excellent.

Figure 4.30: Two Representations of Linearity Test Results for Low to Moderate Count Levels, in the HRC




The upper panel shows aperture sums over 9 pixels for all stars used over a range emphasizing results at low to moderate count levels. The plotted line has a slope set by the relative exposure time. In the lower panel the ratio of counts in ensembles of stars divided into factor of two intensity bins, and further normalized by the relative exposure time are shown. One sigma error bars are derived based on the ensemble signal-to-noise of the short exposure case. The lowest bin has 81 stars, with typical values of 400 stars per bin above this.

Wide Field Channel


In April 2002, one month after ACS's installation, small-aperture stellar photometry with WFC was very linear down to flux levels of about 5 e- in the central pixel (Figure 4.32). At this level, stars are not recognizable in single exposures. Since then, the degradation of CTE by radiation damage has significantly affected the photometric linearity of the WFC, especially for faint sources and low sky background. Figure 4.31 shows the measured and predicted CTE losses for point sources measured with small apertures and different sky background levels. As early as 2003, up to 5% of the signal from a faint source (30 sec. exposure of 20th mag. star through narrow band filters) was lost from the 3-pixel photometric aperture. Thus, any intrinsic non-linearity in the WFC CCDs at low signal levels became negligible compared with the normal losses expected from degraded CTE only a few months after ACS had been placed aboard *HST*.

Figure 4.31: Predicted Impact of CTE on Science Images for WFC



The magnitude loss is estimated for a star located at the chip middle point along the y-axis ($Y = 1024$) and refers to counts measured in the 3 pixel aperture. Two vertical dashed lines refer to the epoch of the ACS failure in January 2007 and to the date of SM4 (May 2009).

Figure 4.32: Two Representations of Linearity Test Results for Low to Moderate Count Levels, in the WFC



The upper panel shows simple aperture sums, for long and short exposures, for all stars of low to moderate intensity; the line has a slope set by the relative exposure times. The lower panel shows ratio of counts summed over all stars within intensity bins (in factor of two steps) for the 1200 sec. exposures after normalization to the relative exposure time. To account for minor encircled energy differences for the very small 9 pixel (total, not radius) apertures used, all points have been normalized by 1.006, the initial value for the brightest bin. The error bars show ± 1 deviations based on the total signal-to-noise of the short exposure sums. The number of stars per bin is typically about 400, although the lowest bin contains only 33 stars.

4.7.4 Linearity Beyond Saturation

The response of the ACS CCDs remains linear not only up to, but well beyond, the point of saturation when using a gain³ value that samples the full well depth. [ACS ISR 2004-01](#) shows the well-behaved response of the ACS: electrons are clearly conserved after saturation. This result is similar to that of the STIS CCD ([Gilliland et al., 1999, PASP, 111, 1009](#)) and the WFPC2 camera ([Gilliland, R. L. 1994, ApJ, 435, L63–66](#)). It is possible to perform photometry on point sources that remain isolated simply by summing over all of the pixels affected by bleeding if the gain value samples the full well depth. Given the larger dynamic range afforded before saturation at GAIN = 2⁴ for WFC, and the only modestly-increased read noise coupled with the potentially beneficial aspect of being able to recover photometry on saturated objects, a gain value of 2 for WFC, currently the default, remains an optimal choice for saturated star photometry.


The extent to which accurate photometry can be extracted for point sources in which one or more pixels have exceeded the physical full well depth is explored in this section. Only the cases of $\text{GAIN} = 4^4$ for the HRC and $\text{GAIN} = 2^4$ for the WFC are considered, which provide direct sampling of the count levels independent of whether saturation has occurred. Ideal data for these tests consist of multiple exposures taken back-to-back on a moderate to rich star field with a broad range of exposure times resulting in both unsaturated and saturated data for many stars.

High Resolution Channel

[Figure 4.33](#) illustrates linearity beyond saturation. In the lower panel, peak values in the long exposure are plotted against peak values from the short exposure. Over the expected linear domain, points fall within a narrow cone centered on a line that has a slope equal to the exposure time ratio while above this the values in the long exposure saturates as expected. Deviations from lying perfectly along the line here result primarily from a 0.1 pixel offset between the images used for this test leading to different relative fractions of light falling on the central pixel.

The upper panel of [Figure 4.33](#) shows the same stars but now using identical extraction apertures in the two exposures. The vertical line, near 205,000 e- on the x-axis is a separation point where stars below have a central pixel that remained unsaturated, while stars above had a saturated central pixel. There is no difference between the aperture photometry of point sources that are saturated up to 5 times the pixel-well depth, and those that are not. Within the domain sampled here the accuracy of saturated star photometry is much better than 1%.

Figure 4.33: Linearity Beyond Saturation in the HRC




The lower panel shows peak pixel values for many stars observed for 350 seconds (long) and 60 seconds (short) with the HRC. Over the range for which stars are unsaturated in both exposures, there is a linear relation with slope equal to exposure time ratio. Brighter stars in the long exposure saturate at values near the full well depth. The upper panel shows the same stars using photometry sums from identical apertures in both exposures; stars above the vertical mark near 205,000 e^- , on the x-axis, separates stars that are saturated in the long exposure. The photometric response when summing over saturated pixels that have been bled into remains perfectly linear far beyond saturation of the central pixel at GAIN = 4.

Wide Field Channel


Figure 4.34, shows that the pixel response of the WFC detector is highly linear up to the full well level of ~80,000 electrons. Beyond this level, accurate aperture photometry is still possible. In Figure 4.35, results for the WFC show that over a range of nearly 4 magnitudes beyond saturation, photometry remains linear to < 1%.

Figure 4.34: Photon Transfer Curve for WFC Obtained During SM4 Optimization Campaign



Each data point shows the mean of signal (in DN) and variance from an identical image pair measured with 25 equidistant 40×40 pixel boxes on the Amplifier B quadrant at commanded gain of 2. The dashed line shows that the detector is linear up to full well. The dotted line marks the full well, which is approximately 40,000 DN (with about 10% variation across the detector).

Figure 4.35: Linearity Beyond Saturation in the WFC



The lower panel shows peak pixel values for many stars observed for 340 seconds (long) and 10 seconds (short) with the WFC. Over the range for which stars are unsaturated in both exposures, there is a linear relation with slope equal to the exposure time ratio. Brighter stars in the long exposure saturate at values near the full well depth. The middle panel shows the same stars using photometry sums from identical apertures in both exposures; stars above the vertical mark near $12,350e^-$, on the x-axis, separates stars that are saturated in the long exposure. The upper panel shows the ratio of long to short aperture sums normalized by the relative exposure time plotted against the degree of over-saturation in the central pixel of the long exposure. The lines in the top and middle panels represent the ideal case of perfect linearity in the pixel response. The fact that the data falls on the lines (or are symmetrically scattered about the lines) indicates that the WFC CCDs are close to perfectly linear. The photometric response when summing over saturated pixels that have been bled into remains perfectly linear far beyond the saturation of the central pixel at GAIN = 2.

These linearity results are based upon comparisons of `flt.fits` images. A comparison of long and short drizzled images show that the conservation of flux property by Drizzle leads to equally good results for linearity beyond saturation. An analysis of the drizzled data sets corresponding to [Figure 4.35](#) showed equally impressive results.

4.7.5 Shutter Stability

For each detector, the shutter consists of two blades located in front of the CCD entrance window. The optical path is blocked by one blade. When a command is sent to begin an exposure, the blade sweeps uniformly across the detector by 90° to open the aperture, exposing the CCD for the commanded integration time. When the exposure is complete, the shutter rotates by another 90° in the same direction so the second blade covers the aperture. A single exposure, therefore, rotates the shutter mechanism by 180°, i.e., one blade opens the aperture while the other closes it. If the blades sweep at a uniform speed, all pixels will be exposed for an identical integration time. The shortest possible exposure time for WFC is 0.5 seconds, where the blade rotates continuously through 180°. (A 0.6 sec integration time is not allowed.) For more information about the shutter, please refer to [ACS ISR 2003-03](#) which followed after the earlier ACS IDT report "WFC and HRC Shutter Shading and Accuracy (Feb. 2001)," by Martel, G. Hartig, and M. Sirianni which may no longer be available on-line.

For short exposure times, field dependent timing (shutter shading), "A" versus "B" blade shutter⁵ control dependence, stability, and timing accuracy were assessed for the HRC and WFC during ground-based testing and through utilizing on-orbit data. The ground results using time variable flat-field sources suggested deviations in actual versus nominal exposure times up to, but no larger than ~1% for short exposures of less than 1.0 sec.

For a full discussion of on-orbit data used to assess shutter stability see [ACS ISR 2003-03](#) (Gilliland and Hartig 2003).

Even at the shortest exposures, shutter shading measurements did not exceed ~0.5% center-to-edge for either camera. This supports a previous decision to not actively invoke a shading correction in **calacs**.

Stability of shutter timing was a bit more problematic. For the HRC, stability appeared excellent when measured as rms across several exposures, and the errors remain well under 0.5%. At exposure times of 1.0 second, peak-to-peak fluctuations on the HRC were at the 0.1% level. Only at 0.1 seconds did the HRC exposures exhibit peak-to-peak fluctuations up to ~1.0% (and part of this may have been measurement error).

The WFC shutter timing stability was not as good, but was also more difficult to quantify given a higher level of systematic errors. Again, measured as an rms across several exposures, errors were well under 0.5%. However, out of 7 pairs of back-to-back 1.0 second exposures, two of the pairs had individual components differing by > 1.3%. Where such differences existed on WFC, it was always the case that shutter timing under A control was shorter than the following B exposure. At 0.5 seconds on WFC (where shutter operation is one continuous rotation), the greatest discrepancy between A to B remained < 0.3%. These results suggest that if generally short WFC exposures are required (where short is taken to be < 2 seconds), then 0.5 second exposures appear to be stable and would likely support 1% accuracy, but exposures in the range of 0.7 seconds to less than 2.0 seconds may experience timing fluctuations that could compromise such accuracy. This conclusion regarding WFC exposures is not regarded as robust, but is offered as that most consistent with a simple and conservative interpretation of the test data.

In adopting a threshold of 0.5% in absolute timing, only 4 exposure times over HRC and WFC required revisions from pre-launch values. The 0.1 second HRC exposure is actually larger by 4.1%, while WFC exposures at 0.5 seconds, 0.7 seconds and 0.8 seconds in reality differ by +1.6%, -1.0% and -0.6% respectively. For these four exposures, using revised exposure times in **calacs** was recommended in order to support accurate photometry; retrievals made after March 11, 2004 will have invoked use of these corrected values.

- ¹ Linearity depends on the constancy of the lamp source, or on experimental techniques to control for drifts.
- ² The saturation count level may be different for a uniform illumination pattern than for point sources.
- ³ The commanded gain value for a CCD image is recorded in the image header keyword `CCDGAIN`. For WFC, the calibrated gain values for amplifiers A, B, C, and D are held by the image header keywords `ATODGNA`, `ATODGNB`, `ATODGNC`, and `ATODGND`, respectively. For HRC, the calibrated gain value for amplifier C is held by the keyword `ATODGNC`.
- ⁴ The use of "GAIN" in this context refers to the GAIN optional parameter in the Phase II proposal exposure logsheet. This commanded gain value is identical to the `CCDGAIN` header keyword value.
- ⁵ The shutter used for an exposure is recorded in the header keyword `SHUTRPOS`.

Chapter 5: ACS Data Analysis

Chapter Contents

- [5.1 Photometry](#)
- [5.2 Astrometry](#)
- [5.3 Polarimetry](#)
- [5.4 Coronagraphy](#)
- [5.5 Ramp Filters](#)
- [5.6 Spectroscopy with the ACS Grisms and Prisms](#)

5.1 Photometry

- 5.1.1 Photometric Systems, Units, and Zeropoints
- 5.1.2 Aperture and Color Corrections
- 5.1.3 Pixel Area Maps
- 5.1.4 PSF
- 5.1.5 CTE
- 5.1.6 Red Leak
- 5.1.7 UV Sensitivity

5.1.1 Photometric Systems, Units, and Zeropoints

It is strongly recommend that, whenever practical, ACS photometric results be referred to a system based on its own filters. Transformations to other photometric systems are possible (see [Ryon et al. ACS ISR 2019-10](#), [Sirianni et al., 2005, PASP, 117, 1049](#)) but such transformations have limited precision and strongly depend on the color range, surface gravity, and metallicity of the stars.

For ACS filters, three magnitude systems are commonly used: VEGAMAG, STMAG, and ABMAG; all of which are based on absolute flux.

The absolute effective flux for any ACS filter and for any stellar flux distribution (SED) can be computed from the throughput for the entire system (OTA + ACS CAMERA + FILTER + DETECTOR) according to Equation 3 of [Bohlin et al. 2014, PASP, 126, 711](#).

VEGAMAG is a widely used standard star system, defined as relative photometry to the actual star Vega; Vega magnitudes can be converted to absolute effective fluxes using the flux distribution of this *HST* secondary standard star¹ in the CALSPEC database at: <http://www.stsci.edu/hst/instrumentation/reference-data-for-calibration-and-tools/astronomical-catalogs/calspec>

The commonly used photometric systems ABMAG ([Oke, J. B. 1964, ApJ 140, 689](#)) or STMAG ([Koorneef, J. et. al. 1985, in Highlights of Astronomy \(IAU\), Vol. 7, ed. J.-P. Swings, 833](#)) are directly related to physical units. The choice between observational and flux-based systems is mostly a matter of personal preference. Any new determination of ACS's absolute efficiency will result in revised magnitudes for these three photometric systems that are based on absolute physical flux.

VEGAMAG

The VEGAMAG system uses Vega (α Lyr) as the standard star. The spectrum of Vega used to define this system is a composite spectrum of empirical and synthetic spectra ([Bohlin & Gilliland 2004, AJ, 127, 3508](#)). The "Vega magnitude" of a star with flux F is

$$\text{VEGAMAG} = -2.5 \times \log\left(\frac{F}{F_{\text{Vega}}}\right)$$

where F_{vega} is the calibrated spectrum of Vega in **pysynphot**. In the VEGAMAG system, by definition, Vega has zero magnitude at all wavelengths.

STMAG and ABMAG

These two similar photometric systems are flux-based systems. The conversion is chosen such that the magnitude in V corresponds roughly to that in the Johnson system.

In the STMAG system, the flux density is expressed per unit wavelength, while in the ABMAG system, the flux density is expressed per unit frequency. The magnitude definitions are:

- **STMAG** = $-2.5 \times \log F_\lambda - 21.1$
- **ABMAG** = $-2.5 \times \log F_\nu - 48.60$

where F_ν is expressed in $\text{erg cm}^{-2} \text{sec}^{-1} \text{Hz}^{-1}$, and F_λ in $\text{erg cm}^{-2} \text{sec}^{-1} \text{\AA}^{-1}$. Another way to express these STMAG (zero point given by *HST* header keyword PHOTZPT) and ABMAG zero points of -21.1 and -48.6 is to say that an object with a constant $F_\nu = 3.63 \times 10^{-20} \text{ erg cm}^{-2} \text{sec}^{-1} \text{Hz}^{-1}$ will have magnitude AB = 0 in every filter, and an object with $F_\lambda = 3.63 \times 10^{-9} \text{ erg cm}^{-2} \text{sec}^{-1} \text{\AA}^{-1}$ will have magnitude STMAG = 0 in every filter.

Zeropoints

Another sort of zeropoint is the "instrumental zeropoint," which is the magnitude of an object that produces one count per second.

Each zeropoint refers to a count rate measured in a specific aperture. For point source photometry, the measurement of counts in a large aperture is not possible for faint targets in a crowded field. Therefore, counts are measured in a small aperture, then an aperture correction is applied to transform the result to an "infinite" aperture. ***For ACS, all zeropoints refer to a nominal "infinite" aperture of radius 5".5.***

By definition, the magnitude in the passband P in any of the ACS systems is given by:

$$\text{ACSmag}(P) = -2.5 \log(\text{total electrons/sec}) + \text{zeropoint}$$

The choice of the zeropoint determines the magnitude system of ACSmag(P). There are several ways to determine the instrumental zeropoints:

- Use **pysynphot** to renormalize a spectrum to 1 count/second in the appropriate ACS passband and specify an output zeropoint value based on a selected magnitude system. See <http://www.stsci.edu/hst/instrumentation/acs/data-analysis/zeropoints> for more. (Be sure to verify that the most updated throughput tables are being used.) In the following example, a 10,000 K blackbody is renormalized to 1 count/second and the zeropoint for the ACS/WFC F555W filter on the WFC1 CCD is computed on the MJD 57754 (January 1, 2017). More information about how to use **pysynphot** may be found at: <http://pysynphot.readthedocs.io>.

Python INPUT:

```
import pysynphot as S

#Note that the string in ObsBandpass() must not contain spaces
band = S.ObsBandpass('acs,wfcl,f555w,mjd#57754')
spec = S.BlackBody(10000)
spec_norm = spec.renorm(1, 'counts', band)
obs = S.Observation(spec_norm, band)

print(obs.effstim('stmag'))
```

Python OUTPUT: 25.667076345950036

- Use values of photometric header keywords (shown below), in the SCI extension(s) of the images, to calculate the STMAG or ABMAG zeropoint:
 - PHOTFLAM is the inverse sensitivity ($\text{erg cm}^{-2} \text{sec}^{-1} \text{\AA}^{-1}$); it represents the flux of a source with constant F_{λ} which produces a count rate of 1 electron per second.
 - PHOTPLAM is the pivot wavelength.
 - PHOTZPT is the STMAG zeropoint, permanently set to -21.1.

The header keywords PHOTFLAM, PHOTZPT, and PHOTPLAM relate to the STMAG and ABMAG zeropoints through these formulae (See also [Bohlin, et al. 2011, AJ, 141, 173](#)).

$$\text{instr_STMAG_ZPT} = -2.5 \times \log(\text{PHOTFLAM}) - \text{PHOTZPT} = -2.5 \times \log(\text{PHOTFLAM}) - 21.1$$

$$\text{instr_ABMAG_ZPT} = -2.5 \times \log(\text{PHOTFLAM}) - 21.1 - 5 \times \log(\text{PHOTFLAM}) + 18.6921$$

- Use the [ACS Zeropoints Calculator](http://www.stsci.edu/hst/instrumentation/acs/data-analysis/zeropoints), or to calculate them yourself follow the instructions at <http://www.stsci.edu/hst/instrumentation/acs/data-analysis/zeropoints>.

The WFC zeropoints for -81° C for post-SM4 data are updated by [Bohlin 2016, AJ, 152, 60](#).

5.1.2 Aperture and Color Corrections

In order to reduce errors due to background variations and to increase the signal-to-noise ratio, aperture photometry and PSF-fitting photometry are usually performed by measuring the flux within a small radius around the center of the source. (For a discussion on the optimal aperture size, see [Sirianni et al. 2005](#)). However, a small aperture measurement needs to be adjusted to a "total count rate" by applying an aperture correction.

For point sources, ACS zeropoints (available in STMAG, ABMAG, and VEGAMAG) are applied to a measured magnitude that is aperture-corrected to a nominal "infinite" aperture of radius 5.5". For surface photometry, these zeropoints can be directly added to values in units of magnitude/arcsec². The Sirianni paper has been superseded in a number of aspects by [Bohlin \(2016\)](#) and the revised encircled energy fractions for both a 0.5" and 1.0" aperture from Table 8-9 of [Bohlin \(2016\)](#) are reproduced here as [Table 5.1](#). Aperture corrections for the SBC are reported in Table 2 of [ACS ISR 2016-05](#) and use 4" for the "infinite" aperture.

Table 5.1: Encircled Energy Fractions for K Type and Hotter Stars in 0.5" and 1.0" Apertures
Typical formal 1 sigma uncertainties on the 0.5" values are 0.004 for WFC and 0.004 for HRC, while uncertainties for the 1" aperture are 0.001 and 0.003 ([Bohlin, 2016](#)).

FILTER	WFC (0.5")	HRC (0.5")	WFC (1.0")	HRC (1.0")
F220W	--	0.868	--	0.948
F250W	--	0.884	--	0.946
F330W	--	0.898	--	0.943
F344N	--	0.899	--	0.943
F435W	0.907	0.910	0.941	0.944
F475W	0.912	0.914	0.944	0.946

F502N	0.914	0.916	0.945	0.947
F555W	0.915	0.919	0.946	0.949
F550M	0.915	0.920	0.947	0.949
F606W	0.915	0.920	0.947	0.950
F625W	0.916	0.919	0.948	0.950
F658N	0.916	0.917	0.948	0.949
F660N	0.916	0.917	0.948	0.949
F775W	0.916	0.884	0.950	0.927
F814W	0.915	0.862	0.949	0.910
F892N	0.898	0.773	0.942	0.844
F850LP	0.893	0.756	0.940	0.831

To determine the aperture corrections, users should measure a few bright stars in an uncrowded region of the field of view with both the smaller measurement aperture and the 0.5" radius aperture. Users should determine the aperture correction between their own small-aperture photometry and aperture photometry with a 0.5" radius aperture; this is done by measuring a few bright stars in an uncrowded region of the field of view with both the smaller measurement aperture and the 0.5" radius aperture. The difference between the two apertures should then be applied to all the small aperture measurements. If such stars are not available, small aperture encircled energies have been tabulated by [Bohlin \(2016\)](#). However, accurate aperture corrections are a function of time and location on the chip and also depend on the kernel used by **AstroDrizzle**². Blind application of tabulated encircled energies for small radii should be avoided.

Aperture corrections for near-IR filters present further complications because the ACS CCD detectors suffer from scattered light at long wavelengths. These thinned backside-illuminated devices are relatively transparent to near-IR photons; the transmitted long wavelength light illuminates and scatters in the CCD soda glass substrate, is reflected back from the header's metallized rear surface, then re-illuminates the CCDs frontside photosensitive surface ([Sirianni et al. 1998, *proc SPIE* vol. 3555, 608, ed. S. D'Odorico](#)). The fraction of the integrated light in the scattered light halo increases as a function of wavelength. As a consequence, the PSF becomes increasingly broad with increasing wavelengths. WFC CCDs incorporate a special anti-halation aluminum layer between the frontside of the CCD and its glass substrate. While this layer is effective at reducing the IR halo, it appears to give rise to a relatively strong scatter along one of the four diffraction spikes at wavelengths greater than 9000 Å ([Hartig et al. 2003, *proc SPIE* vol. 4854, 532, ed. J. C. Blades, H. W. Siegmund](#)), see [Section 5.1.4](#) and [Bohlin \(2012\)](#) for more details.

The same mechanism responsible for the variation of the intensity and extension of the halo as a function of wavelength is also responsible for the variation of the shape of the PSF as a function of color of the source. As a consequence, in the same near-IR filter, the PSF for a red star is broader than the PSF of a blue star. [Gilliland & Riess \(2002 *HST Calibration Workshop Proceedings*, page 61\)](#) and [Sirianni et al. 2005](#) provide assessments of the scientific impact of these PSF artifacts in the red. The presence of the halo has the obvious effect of reducing the signal-to-noise and the limiting magnitude of the camera in the red. It also impacts the photometry in very crowded fields. The effects of the long wavelength halo should also be taken into account when performing morphological studies and performing surface photometry of extended objects (see [Sirianni et al. 2005](#) for more details).

The aperture correction for red objects should be determined using an isolated, same-color star in the field of view, or by using the effective wavelength versus aperture correction relation ([Sirianni et al., 2005](#); [Bohlin 2012, 2016](#)). If the object's spectral energy distribution (SED) is available, an estimate of the aperture correction is also possible with **pysynphot**; the parameter *aper* has been implemented to call the encircled energy tables in the *obsmode* **pysynphot** field for ACS. A typical *obsmode* for an aperture of 0.5" would be specified as "acs,wfc1,f850lp,aper#0.5". A comparison with the infinite aperture magnitude using the standard *obsmode* "acs,wfc,f850lp" would give an estimate of the aperture correction to apply. Please refer to the [PySynphot documentation](#) for more details.

Color Correction

In some cases, ACS photometric results must be compared with existing datasets in different photometric systems (e.g., WFPC2, SDSS, or Johnson-Cousins). Because the ACS filters do not have exact counterparts in any other standard filter sets, the accuracy of these transformations is limited. Moreover, if the transformations are applied to objects whose spectral type (e.g., color, metallicity, surface gravity) do not match the spectral type of the calibration observation, systematic effects could be introduced. The transformations can be determined by using **pysynphot**, or by using the published transformation coefficients ([ACS ISR 2019-10](#), [Sirianni et al. 2005](#)). In any case, users should not expect to preserve the 1%–2% accuracy of ACS photometry on the transformed data.

5.1.3 Pixel Area Maps

When ACS images are flat-fielded by the **calacs** pipeline; the resultant *flt.fits/flc.fits* files are "flat" if the original sky intensity was also "flat." However, there is still very significant geometric distortion remaining in these images. The pixel area on the sky varies across the field, and as a result, relative point source photometry measurements in the *flt.fits/flc.fits* images are incorrect.

One option is to drizzle the data; this will remove geometric distortion while keeping the sky-flat. Therefore, both surface and point source relative photometry can be performed correctly on the resulting *drz.fits/drc.fits* files. The inverse sensitivity (in units of $\text{erg cm}^{-2} \text{sec}^{-1} \text{\AA}^{-1}$), given by the header keyword **PHOTFLAM**, can be used to compute the STMAG or ABMAG zeropoint, and to convert flux in electrons/seconds to absolute flux units:

$$\begin{aligned}\text{STMAG_ZEROPOINT} &= -2.5 \text{ Log (PHOTFLAM) } - \text{PHOTZPT} \\ \text{ABMAG_ZEROPOINT} &= -2.5 \text{ Log (PHOTFLAM) } - 21.10 - 5 \text{ Log (PHOTPLAM) } + 18.6921\end{aligned}$$

where,

- STMAG_ZEROPOINT is the ST magnitude zeropoint for the observing configuration (given in the header keyword **PHOTMODE**).
- ABMAG_ZEROPOINT is the AB magnitude zeropoint for the observing configuration.
- PHOTFLAM³ is the mean flux density (in $\text{erg cm}^{-2} \text{sec}^{-1} \text{\AA}^{-1}$) that produces 1 count per second in the *HST* observing mode (**PHOTMODE**) used for the observation.
- PHOTZPT³ is the ST magnitude zeropoint (= 21.10).
- PHOTPLAM³ is the pivot wavelength.

Remember that for point source photometry, both these zeropoints should be applied to measurements corrected to the standard "infinite" aperture; for ACS, this is a radius of 5.5".

Additional information about zeropoints is available at the [ACS Zeropoints Web page](#).

Users who wish to perform photometry directly on the distorted `flt.fits/flc.fits` and `crj.fits/crc.fits` files, rather than the drizzled (`drz.fits/drc.fits`) data products, will require a field-dependent correction to match their photometry with that obtained from drizzled data. Only then can the PHOTFLAM and PHOTPLAM values in the `flt.fits/flc.fits` and `crj.fits/crc.fits` images be used to obtain calibrated STMAG or ABMAG photometry. (Note: the corresponding `drz.fits/drc.fits` image has identical PHOTFLAM and PHOTPLAM values.)

The correction to the `flt.fits/flc.fits` images may be made by multiplying the measured flux in the `flt.fits/flc.fits` image by the pixel area at the corresponding position using a pixel area map (PAM), and then dividing by the exposure time. The easiest way to do it is to simply multiply the `flt.fits/flc.fits` image with its corresponding pixel area map.

$$\text{DRZ_flux} = \text{FLT_flux} * \text{PAM} / \text{exposure time.}$$

For example, in Python:

```
from astropy.io import fits
import shutil

shutil.copy('jdly04rlq_flt.fits', 'jdly04rlq_flt_pam.fits')
flt_hdu = fits.open('jdly04rlq_flt_pam.fits', mode = 'update')
pam1_hdu = fits.open('wfc1_pam.fits')
pam2_hdu = fits.open('wfc2_pam.fits')


flt_hdu['sci', 1].data *= pam2_hdu[0].data
flt_hdu['sci', 2].data *= pam1_hdu[0].data


flt_hdu.close()
pam1_hdu.close()
pam2_hdu.close()
```

`drz.fits/drc.fits`, `crj.fits/crc.fits` and `flt.fits/flc.fits` images have units of electrons /seconds, and have the same PHOTFLAM values.

The PAM for the WFC is approximately unity at the center of the WFC2 chip, ~0.95 near the center of the WFC1 chip and ~1.12 near the center of the HRC. PAM images, which are very large, may be downloaded directly from this page: <http://www.stsci.edu/hst/instrumentation/acs/data-analysis/pixel-area-maps>.

Figure 5.1: Variation of the WFC and HRC Effective Pixel Area with Position in Detector Coordinates





Pixel Area Map Concept Illustration


To illustrate the concepts of extended source and point source photometry on `flt.fits/flc.fits` and `drz.fits/drc.fits` images, consider a simple idealized example of a 3×3 pixel section of the detector, assuming that the bias and dark corrections are zero and that the quantum efficiency is unity everywhere.

Example #1. Illustration of Geometric Distortion on a Constant Surface Brightness Object

For an extended object with a surface brightness of $2e^-/\text{pixel}$ in the undistorted case, an image without geometric distortion is:

Actual scene on the sky	2	2	2
	2	2	2
	2	2	2


In reality, ACS suffers from geometric distortion. As a consequence, the pixel scale varies across the detector. The result is that the sky area coverage per pixel is not identical.




The pixel area map (PAM), shown below, illustrates the differences in area for each pixel, represented by a fractional value, due to geometric distortion.

PAM	1.4	1.2	1.0
	1.2	1.0	0.9
	1.0	0.9	0.8

As a result, the raw data shows an apparent variation in surface brightness because some pixels detect flux from a larger sky area than others.




The geometrical area of each pixel is imprinted in the flat field, along with its photometric sensitivity. In this example, the quantum efficiency is unity everywhere, so the flat field is the equivalent of the PAM:



ACS flat fields are designed to produce a flat image when the instrument is uniformly illuminated. This, however, means that pixels which are smaller than average on the sky are boosted, while pixels with relatively large areas are suppressed. Application of the PAM removes this effect—pixels now show the true relative illumination they receive from a uniform source. However, the image remains geometrically distorted. Thus, when doing aperture photometry on the field, the user should take into account that aperture sizes defined in pixels are not uniform in size across the field of view.

If **AstroDrizzle** is run on a `flt.fits/flc.fits` image, the output image is free of geometric distortion and is photometrically accurate.



When drizzling a single image, the user may want to use the Lanczos kernel, which provides the best image fidelity for the single image case. However, this kernel does not properly handle missing data and causes ringing around cosmic rays. Thus, the Lanczos kernel should not be used for combining multiple images where sections of the image lost to defects on one chip can be filled in by dithering.

For additional information about the inner workings of **AstroDrizzle**, please refer to the [DrizzlePac website](#).


drz/drc. fits image	2.0	2.0	2.0
	2.0	2.0	2.0
	2.0	2.0	2.0

Example #2 Illustration of Geometric Distortion and Integrated Photometry of a Point Source


This example considers observing a point source and that all the flux is included in the 3 × 3 grid. Let the counts distribution be:

Actual scene on the sky	2.0	10.5	2.0
	10.5	50	10.5
	2.0	10.5	2.0


The total counts are 100. Due to geometric distortion, the PSF, as seen in the raw image, is distorted. The total counts are conserved, but they are redistributed on the CCD, as shown in the fractional area values below.



After the flat-field correction, however, the total counts are no longer conserved:



In this example the counts now add up to 89.08, instead of 100. In order to perform integrated photometry, the pixel area variation needs to be taken into account. This can be done by multiplying the `flt.fits/flc.fits` image by the PAM or by running **AstroDrizzle**. *Only by running **AstroDrizzle** can the geometric distortion be removed*, but both approaches correctly recover the count total as 100.



This graphical depiction of geometric distortion and drizzling is just an idealized example. In reality, the PSF of the star extends to a much bigger radius. If a user decides to work on the flat-fielded image after

correcting by the pixel area map, he or she needs to calculate a new aperture correction to the total flux of the star. The aperture corrections discussed in [Section 5.1.2](#) are for `flt` or `crj` images at the WFC1-1K reference position, as corrected by the PAMs. However, the drz drizzled corrections are the same down to an aperture radius of 3 pixels, where the encircled energy differs by 0.5%.

5.1.4 PSF

PSF Field Dependence

Point spread functions in the ACS cameras are relatively stable over the field of view, especially when compared to previous generation cameras such as WFPC2. Variations in the HRC are very small and probably negligible when using apertures greater than $r = 1.5$ pixels or using PSF fitting. However, the WFC PSF varies enough in shape and width that significant photometric errors may be introduced when using small apertures or fixed-width PSF fitting. These effects are described in detail in [ACS ISR 2003-06](#).


The WFC PSF width variation is mostly due to changes in CCD charge diffusion. Charge diffusion, and thus the resulting image blur, is greater in thicker regions of the detector (the WFC CCD thickness ranges from 12.6 to 17.1 microns, see [Figure 5.2](#)). At 500 nm, the PSF FWHM varies by 25% across the field. Because charge diffusion in backside-illuminated devices like the ACS CCDs decreases with increasing wavelength, the blurring and variations in PSF width will increase towards shorter wavelengths. At 500 nm, photometric errors as much as 15% may result when using small ($r < 1.5$ pixel) apertures. At $r = 4$ pixels, the errors are reduced to $< 1\%$. Significant errors may also be introduced when using fixed-width PSF fitting (see [ACS ISR 2003-06](#)).

PSF shape also changes over the WFC field due to the combined effects of aberrations like astigmatism, coma, and defocus. Astigmatism noticeably elongates the PSF cores along the edges and in the corners of the field. This may potentially alter ellipticity measurements of the bright, compact cores of small galaxies at the field edges. Coma is largely stable over most of the field and is only significant in the upper left corner, and centroid errors of ~ 0.15 pixels may be expected there.

Observers may use **TinyTim** to predict the variations in the PSF over the field of view for their particular observation. **TinyTim** accounts for wavelength and field-dependent charge diffusion and aberrations. The software is available at: <http://www.stsci.edu/software/tinytim/>

For point source relative astrometry, procedures for obtaining the best results are described in [ACS ISR 2006-01](#).

Figure 5.2: WFC Chip Thickness (left) and PSF FWHM (right)




PSF Long Wavelength Artifacts

Long wavelength ($\lambda > 700$ nm) photons can pass entirely through a CCD without being detected and enter the substrate on which the detector is mounted. In the case of the ACS CCDs, the photons can be scattered to large distances (many arcseconds) within the soda glass substrate before reentering the CCD and being detected—it creates a large, diffuse halo of light surrounding an object, called the "red halo." This problem was largely solved in the WFC by applying a metal coating between the CCD and the mounting substrate that reflects photons back into the detector. Except at wavelengths longer than 900 nm (where the metal layer becomes transparent), the WFC PSF is unaffected by the red halo. The HRC CCD, however, does not have this fix and is significantly impacted by the effect.

The red halo begins to appear in the HRC at around 700 nm. It exponentially decreases in intensity with increasing radius from the source. The halo is featureless but slightly asymmetrical, with more light scattered towards the lower half of the image. By 1000 nm, it accounts for nearly 30% of the light from the source and dominates the wings of the PSF, washing out the diffraction structure. Because of its wavelength dependence, the red halo can result in different PSF light distributions within the same filter for red and blue objects. The red halo complicates photometry in red filters. In broad-band filters like F814W and especially F850LP (in the WFC as well as the HRC), aperture corrections will depend on the color of the star, see [Section 5.1.2](#) for more discussion of this. Also, in high-contrast imaging where the PSF of one star is subtracted from another (including coronagraphic imaging), color differences between the objects may lead to a significant residual over- or under-subtracted halo.

In addition to the halo, two diffraction spike-like streaks can be seen in both HRC and WFC data beyond 1000 nm (including F850LP). In the WFC, one streak is aligned over the left diffraction spike while the other is seen above the right spike. For HRC, the streak is aligned over the right diffraction spike while the other is seen below the left spike. These seem to be due to scattering by the electrodes on the back sides of the detectors. They are about five times brighter than the diffraction spikes and result in a fractional decrease in encircled energy. They may also produce artifacts in sharp-edged extended sources.

Figure 5.3: CCD Scatter at Red Wavelengths.



WFC images of the standard star GD71 through filters F775W (9 sec., left), F850LP (24 sec., middle), and FR1016N at 996 nm (600 sec., right). The CCD scatter, undetected below ~800 nm, grows rapidly with longer wavelength. In addition to the asymmetrical, horizontal feature, a weaker diagonal streak also becomes apparent near 1 micron.

HRC and SBC UV PSFs

Below 3500 Å, the low- and mid-spatial frequency aberrations in *HST* result in highly asymmetric PSF cores surrounded by a considerable halo of scattered light extending 1 to 2 arcseconds from the star. The asymmetries may adversely affect PSF-fitting photometry if idealized PSF profiles are assumed. Also, charge scattering within the SBC MAMA detector creates a prominent halo of light extending about 1" from the star that contains roughly 20% of the light. This washes out most of the diffraction structure in the SBC PSF wings. An updated study of the SBC PSF can be found in [Avila and Chiaberge, ACS ISR 2016-05](#).

5.1.5 CTE

ACS's WFC and HRC cameras have CCD detectors that shift charge during readout and therefore suffer photometric losses due to imperfect charge transfer efficiency (CTE). Such losses are particularly significant along the parallel direction (y-axis on the detector).

Two methods are currently available to correct photometry for CTE losses. The first method uses a pixel-based correction. The second method utilizes a photometric correction formula and may be used to correct photometry of point sources.

Pixel-Based CTE Correction

[Anderson and Bedin \(2010, PASP, 122, 1035\)](#) have developed an empirical approach based on the profiles of warm pixels to characterize the effects of CTE losses for WFC. Such an algorithm first develops a model that reproduces the observed trails, and then inverts the model to convert the observed pixel values in an image into an estimate of the original pixel values. The pixel-based CTE correction, applicable only to WFC images, has been implemented in the ACS calibration pipeline (**calacs**). Data products that are corrected for CTE losses have the suffix `flc.fits` and `drz.fits`, which correspond to the (uncorrected) `flt.fits` and `drz.fits` images. Additional information on the pixel-based CTE correction can be found in [Section 4.6](#).

Photometric CTE Correction

The CTE-correcting formulae described in [ACS ISR 2009-01](#) and [ACS ISR 2012-05](#) can also estimate lost flux as a function of source brightness, sky brightness, *x* and *y* position, and time. The pixel-based CTE correction on stellar fields is in general agreement with these photometric correction formulae. Statistically significant deviations are observed only at low stellar fluxes (~300 e- and lower) and for background levels close to 0 e-.

For low stellar fluxes, the correction formulae may be more accurate; however, these formulae fail for short exposures with sky values near or below zero, because of the $\log(\text{sky})$ or sky to a negative power term.

The formulae are calibrated for conventional aperture photometry with an aperture radius of 3 pixels (3 and 5 pixels for WFC). CTE loss calculations are based on the number of transfers in the y direction that a source undergoes; charge losses that depend on pixel transfers in the x direction are negligible. More details can be found in [ACS ISR 2009-01](#) and [ACS ISR 2012-05](#) for pre- and post-SM4 data, respectively.

CTE Correction Cookbook

This section briefly describes the procedure that users should follow in order to apply the CTE correction formula.

1. Obtain `flt.fits` images using **calacs**, then use **AstroDrizzle** to create `drz.fits` data. Alternatively, `drz.fits` files from the *HST* Archive may be used if those images are acceptable.
2. Multiply the `drz.fits` images by the exposure time.
 - For a single exposure `drz.fits` image, multiply the image by its exposure time.
 - If the `drz.fits` image was created from combining several `flt.fits` images with the same exposure time, multiply the `drz.fits` image by the exposure time of a single `flt.fits` image.
 - If the `drz.fits` file was made from input `flt.fits` images with different exposure times, do not use that drizzle-combined image. Instead, run **AstroDrizzle** to create single exposure `drz.fits` files for each `flt.fits` file. In these situations, the recommended method for photometry is to perform the measurements on each single exposure `drz.fits` image. Those results can be corrected for CTE losses, then averaged to obtain a more accurate measurement of the stellar flux.
3. Perform aperture photometry with any preferred software (e.g., **daophot**). For images drizzled at the native scale, set the photometry aperture radius, r , to 3 pixels. Measure the background for each star locally (e.g., in an annulus of $r_{min} = 13$ pixels and $r_{max} = 18$ pixels, centered on the star).
4. Obtain the epoch of the observation, [Table 2.6](#), in modified Julian days (MJD) from the header keyword EXPSTART.
5. To measure the number of Y-transfers for each star, measure the coordinates of the stars directly on the corresponding `flt.fits` image. For stars that fall in WFC2 (extension 1 of a multi-extension `flt.fits` file), the number of transfers, Y_{tran} , is simply the y coordinate of the star. For WFC1, in extension 4 of the same multi-extension FITS file, the number of transfers is $Y_{tran} = 2049 - y_{star}$ where y_{star} is the y coordinate of the star on that frame.

Note the need for particular care in dithered observations. If the position of the star changes only by a few (< 10) pixels, the correction is still well within the error even for large losses. For HRC, $Y_{tran} = y_{star}$ for readouts with amps C or D, $Y_{tran} = 1024 - y_{star}$ for readouts with amps A or B.

If the dithers have larger shifts, carefully check the original position of each star in the `flt.fits` files, then derive the correction using the average value of Y_{tran} . While doing so, it is also recommended to verify that parameters used in **AstroDrizzle** do not generate any biases when the cosmic ray rejection step is performed—the flux of a star located at different positions on the chip might significantly differ, and **AstroDrizzle** might interpret such objects as cosmic rays.

6. Apply the appropriate correction formula (see below) where
 - FLUX* is the stellar flux measured within the aperture radius (in electrons).
 - SKY* is the local background for each star (in electrons).
 - t* is the observation date in modified Julian days.
 - Y_{tran} is defined in step 5.

7. Apply the magnitude correction to the results of the performed photometry (in magnitudes: $mag_{corrected} = mag_{measured} - \Delta mag$).
8. Perform aperture correction (see [Bohlin 2016](#)).

Any additional corrections (i.e., transforming the flux from e- into e-/sec and applying the zeropoints to transform the flux into AB or Vega magnitudes) should be performed after all of the steps above.

- For WFC and pre-SM4 data (MJD < 54129), use the following formula:

$$\Delta mag = 10^A \times SKY^B \times FLUX^C \times (Y_{tran}/2000) \times (MJD-52333) / 365$$

with the following coefficients (and associated 1 σ uncertainties)

$A = -0.14$ (0.04), $B = -0.25$ (0.01), $C = -0.44$ (0.02).

- For WFC and post-SM4 data users should apply the following formula

$$\Delta mag(Y, t, SKY, FLUX) = [p_1 \text{Log}(SKY) \text{Log}(FLUX) t + p_2 \text{Log}(SKY) \text{Log}(FLUX) + p'_1 \text{Log}(SKY) t + q_1 \text{Log}(Flux) t + p'_2 \text{Log}(SKY) + q_2 \text{Log}(FLUX) + q'_1 t + q'_2] * Y_{tran} / 2000$$

The original coefficients are derived in [ACS ISR 2012-05](#). Updated coefficients are reported on the ACS Website at: <http://www.stsci.edu/hst/instrumentation/acs/performance/cte-information>

Users should be reminded that the formula for WFC was calibrated using stellar fluxes between ~50 e- and ~80,000 e- (measured within the 3 pixel aperture radius), and for background levels between ~0.1 e- and ~50 e-. Therefore, to ensure the highest level of accuracy, the formula should be used to correct photometry of stars that are within the range specified above. However, note that for stellar fluxes and background levels higher than the above limits, the amount of the correction is < 2%, even for stars located at the edge of the chip, far from the amplifiers. For very bright stars, the CTE formula currently overestimates CTE losses. For the specific case of very bright stars, the use of `flc.fits` and `drc.fits` files (i.e., those obtained with the pixel-based CTE correction included in the ACS pipeline) is a better option.

Note that for post-SM4 WFC data, users can utilize the online CTE correction calculator available at the following URL: <https://acsphotometriccte.stsci.edu/>

A calculator-specific version of this walkthrough, which differs slightly from the instructions above, is available at that URL as well.

- For HRC, the CTE correction can be performed using

$$\Delta mag = 10^A \times SKY^B \times FLUX^C \times (Y_{tran}/1000) \times (MJD-52333) / 365$$

with the following coefficients (and associated 1sig uncertainties)

$A = -0.44$ (0.05), $B = -0.15$ (0.02), $C = -0.36$ (0.01).

The Relationship Between Field-Dependent Charge Diffusion and CTE

Because the WFC and HRC CCDs were thinned during the manufacturing process, there are large-scale variations in thicknesses of their pixels. Charge diffusion in CCDs depend on the pixel thickness (thicker

pixels suffer greater diffusion), so the width of the PSF and, consequently, the corresponding aperture corrections are field-dependent (see [Section 5.1.4](#)). [ACS ISR 2003-06](#) characterizes the spatial variation of charge diffusion as well as its impact on fixed aperture photometry. For intermediate and large apertures ($r > 4$ pixels), the spatial variation of photometry is less than 1%, but it becomes significant for small apertures ($r < 3$ pixels).

It is unnecessary to decouple the effects of imperfect CTE and charge diffusion on the field dependence of photometry. The CTE correction formulae account for both effects, as long as the user seeks to correct photometry to a "perfect CTE" in the aperture used to obtain measurements (e.g., the recommended aperture radius of 3 pixels). An aperture correction from the measuring aperture to a nominal "infinite" (5.5" radius) aperture is still necessary (see [Section 5.1.2](#)).

Internal CTE Measurements

CTE measurements from internal WFC and HRC calibration images have been obtained since the cameras were integrated into ACS. Extended Pixel Edge Response (EPER) and First Pixel Response (FPR) images are routinely collected to monitor the relative degradation of CTE over time. These images confirm a linear degradation, but are otherwise not applicable to the scientific calibration of ACS data. Results of the EPER and FPR tests were published in [ACS ISR 2005-03](#) and routinely updated at: <http://www.stsci.edu/hst/instrumentation/acs/performance/cte-information>

5.1.6 Red Leak

HRC

When designing a UV filter, a high suppression of off-band transmission, particularly in the red, had to be traded with overall in-band transmission. The very high blue quantum efficiency of the HRC, compared to WFPC2, made it possible to obtain an overall red leak suppression comparable to that of the WFPC2 while using much higher transmission filters. The ratio of in-band versus total flux, determined using in-flight calibration observations, is given in [Table 5.2](#) for the UV and blue HRC filters, where the cutoff point between in-band and out-of-band flux is defined as the filter's 1% transmission points. This is described in [ACS ISR 2007-03](#). This ISR also reports on the percentage of in-band flux for seven stellar spectral types, elliptical galaxies spectrum (Ell. G), a reddened ($0.61 < E(B-V) < 0.70$) starburst galaxy (SB), and four different power-law spectral slopes $F_{\lambda} \propto \lambda^{\alpha}$.

Clearly, red leaks are not a problem for F330W, F435W, and F475W and are important for F250W and F220W. In particular, accurate UV photometry of objects with the spectrum of an M star will require correction for the red leak in F250W and will be essentially impossible in F220W. For the latter filter, a red leak correction will also be necessary for K and G types.

Table 5.2: In-band Flux as a Percentage of the Total Flux

	O5V	B0V	A0V	F0V	G0V	K0V	M2V	Ell. G	SB	$\alpha = -1$	$\alpha = 0$	$\alpha = 1$	$\alpha = 2$
F220W	99.97	99.96	99.73	99.24	96.1	79.52	5.02	95.48	99.67	99.86	99.69	99.24	97.85
F250W	99.98	99.98	99.8	99.71	99.42	98.64	76.2	98.15	99.82	99.93	99.87	99.74	99.45
F330W	99.99	99.99	99.99	99.99	99.99	99.99	99.95	99.99	99.99	99.99	99.99	99.99	99.99

SBC

The visible light rejection of the SBC is excellent, but users should be aware that stars of solar type or later will have a significant fraction of the detected flux coming from outside the nominal bandpass of the detector. Details are given below in [Table 5.3](#).

Table 5.3: Visible-Light Rejection of the SBC F115LP Imaging Mode

Stellar Type	Percentage of all Detected Photons which have $\lambda < 1800 \text{ \AA}$	Percentage of all Detected Photons which have $\lambda < 3000 \text{ \AA}$
O5V	99.5	100
B1V	99.4	100
A0V	98.1	100
G0V	72.7	99.8
K0V	35.1	94.4

5.1.7 UV Sensitivity

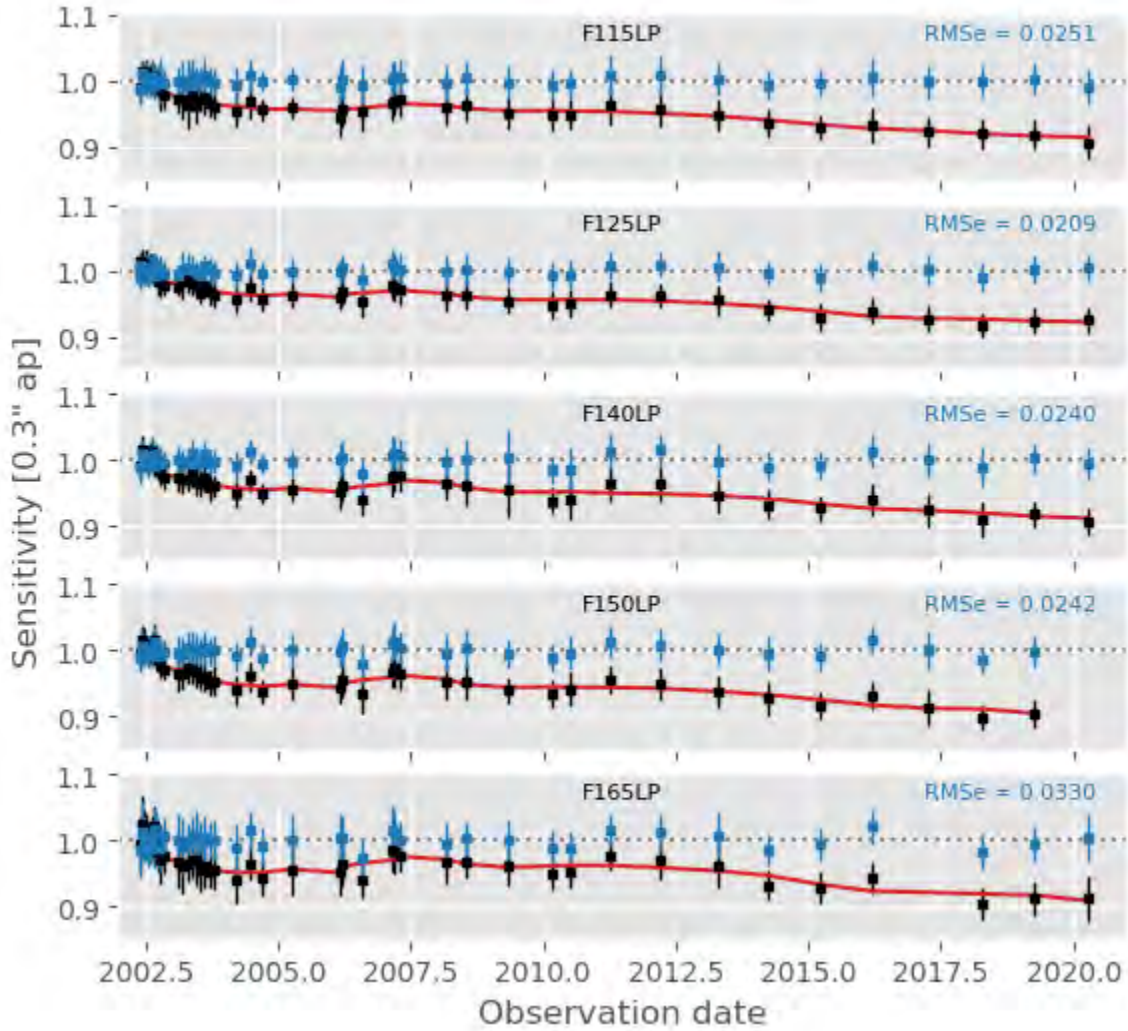
In an ongoing calibration effort, the star cluster NGC 6681 has been observed since the launch of ACS to monitor the UV performance of the HRC and SBC detectors.

Results for the HRC detector for the first year following launch were published in [ACS ISR 2004-05](#). For the three filters, F220W, F250W, and F330W, eight standard stars in the field were routinely measured, indicating a sensitivity loss of not more than ~1% to 2% per year.

Results for the SBC detector have been obtained using repeated measurements of ~50 stars in the same cluster to provide an estimate of the UV contamination in addition to the accuracy of the existing flat fields. The uniformity of the detector response must be corrected prior to investigating any temporal loss in sensitivity. See [Section 4.4.2](#) for a discussion of the SBC flat corrections.

The sensitivity of MAMA detectors declines with time ([ACS ISR 2019-04](#)). The sensitivity of the SBC has declined by up to ~10% since launch, with a rate of about 0.5% per year since 2007 (see [Figure 5.4](#)). This time-dependent sensitivity requires observation date-dependent zeropoints. The necessary files have been delivered to the calibration pipeline so that the `PHOTFLAM` header keyword in SBC images is populated with the correct value. Appropriate zeropoints can also be obtained from the [ACS Zeropoints Calculator](#).

Figure 5.4: Sensitivity of the SBC imaging modes.



Sensitivity measurements for five SBC imaging modes. Black boxes are the mean and standard deviation of the sensitivity before time-dependent sensitivity (TDS) correction. The red line is the extracted TDS. Blue boxes are the sensitivity after applying TDS correction.

¹ The instrumental flux calibrations are defined by observations of the three primary *HST*WD standards GD71, GD153, and G191B2B, which then defines the instrumental sensitivity function for absolute flux. For example, the STIS total throughput sensitivity is used to derive flux distributions for any star observed, including Vega (Bohlin & Gilliland 2004). In other words, the STIS flux for Vega depends only on (1) the absolute fluxes for the three WD primaries, (2) the precision of establishing the STIS calibration from the observations of the three primaries, and (3) the quality of the STIS observations of Vega. Because the *HST* /STIS Vega flux depends on the flux of other stars, Vega is, by definition, a *HST* secondary standard except for its monochromatic flux at the one wavelength of 5556 Å which was measured originally relative to tungsten filament lamps. The 5556 Å Vega flux is used to normalize the WD model fluxes distributions via their brightness relative to Vega, as measured by the STIS data. Thus, Vega remains an *HST* primary star at the one monochromatic wavelength of 5556 Å, making it unique as a sort of hybrid *HST* secondary/primary star.

² The final PSF depends slightly on the kernel and pixel scale used to create the output image. This is because the final image is a convolution of the optical/electronic PSF with the final "final_scale" and "final_pixfrac"

values. The effect on the PSF is small when the "final_pixfrac" and "final_scale" values are small, and when the measuring aperture is just a few original pixels in radius.

³ Header keyword in calibrated images (`flt.fits/flc.fits` and `drz.fits/drc.fits`).

⁴ STIS ceased science operations in August 2004 due to a power supply failure within its Side-2 electronics (which had powered the instrument since May 2001 after the failure of the Side-1 electronics). In May 2009, during SM4, repairs to the electronics fully restored STIS operation.

5.2 Astrometry

- 5.2.1 Coordinate Transformations
- 5.2.2 Absolute and Relative Astrometry
- 5.2.3 Impact of Guide Star Failure

5.2.1 Coordinate Transformations

There are three coordinate systems applicable to ACS images.

- The position of a pixel on the geometrically distorted raw image (`raw.fits`) or, identically, the position on the flat-fielded images (`flt.fits/flc.fits`) after pipeline processing through **calacs**.
- The pixel position on the drizzled images (`drz.fits/drc.fits`) created by **AstroDrizzle** which corresponds to an undistorted pixel position on a tangent plane projection of the sky.
- The corresponding position (RA, Dec) on the sky

A pixel position on a drizzled image (`drz.fits/drc.fits`) may be transformed to a position on the celestial sphere (RA, Dec) using the task **pixtosky** (in the **DrizzlePac** package, see the [DrizzlePac website](#) for details). There is a corresponding task, **skytopix**, also in the **DrizzlePac** package, that transforms a RA, Dec position to a pixel position on a drizzled image.

 For more information about drizzling HST images, please refer to the [DrizzlePac website](#).

5.2.2 Absolute and Relative Astrometry

The astrometric information in the header of an ACS image comes indirectly from the positions of the guide stars used during the observations. As a result, the absolute astrometry attainable by using the image header world coordinate system directly is limited by two sources of error. First, the positions of guide stars are not known to better than about 0.3 arcseconds. Second, the mapping from the guide star to the instrument aperture introduces a smaller, but significant error.

Although absolute astrometry cannot be done to high accuracy without additional knowledge, relative astrometry with ACS is possible to a much higher accuracy. In this case, the limitations are primarily the accuracy with which the geometric distortion of the camera has been characterized, see the [DrizzlePac website](#) for details. With the inclusion of a ¹time-dependent skew in the ACS distortion model used by **AstroDrizzle**, the accuracy of alignment between ACS/WFC images is ~0.05 pixels or better. (Please see [ACS ISR 2015-02](#) and [ACS ISR 2015-06](#)).

Accurate astrometric measurements, especially for faint sources, should take into account the effects of CTE, as described in [ACS ISR 2007-04](#). The Institute is monitoring the variations of the linear skew terms and will continue updating the corresponding astrometric reference files described in the above-mentioned ISR.

5.2.3 Impact of Guide Star Failure

The normal guiding mode uses two guide stars that are tracked by two of HST's Fine Guidance Sensors (FGSs). On some occasions, when two suitable guide stars are not available, single-star guiding is used with

the telescope roll controlled by the gyros. These observations will suffer from small drift rates. To determine the quality of tracking during these observations, please refer to the [Introduction to the HST Data Handbooks](#) for information about jitter files.

In single guide star guiding, typical gyro drift rates produce a roll about the guide star of 1.0–1.5 mas per 1 second exposure, which in turn introduces a translational drift of the target on the detector. This roll is not reset and continues to build over multiple orbits and reacquisitions, until the next full guide star acquisition.

The exact size of the drift depends on the exact roll drift rate and distance from the single guide star to the target in the *HST* field of view. Additional information is available in [ISR TEL 2005-02](#). For ACS with single-star guiding, the typical and maximum drift rate of the target on the detector are shown in [Table 5.5](#).

Table 5.5: Drift Rates for Single-Star Guiding with ACS

TYPICAL	per 1000 sec. exposure	per orbit (96 min.)
WFC	0.0041 arcsec (0.08 pix)	0.024 arcsec (0.47 pix)
HRC	0.0048 arcsec (0.19 pix)	0.028 arcsec (1.11 pix)

MAXIMUM	per 1000 sec. exposure	per orbit (96 min.)
WFC	0.0080 arcsec (0.16 pix)	0.046 arcsec (0.92 pix)
HRC	0.0092 arcsec (0.37 pix)	0.053 arcsec (2.12 pix)

The drift over an orbital visibility period can be calculated from the values in [Table 5.5](#). The typical visibility period in an orbit (outside the Continuous Viewing Zone [CVZ]) ranges from 52 to 60 minutes, depending on target declination. The drifts inherent to single-star guiding are **not** represented in the image header astrometric information, and have two important consequences:

- There will be a slight drift of the target on the detector within a given exposure. For the majority of observations and scientific applications this will not degrade the data (especially if the exposures are not very long). The drift is smaller than the FWHM of the point spread function (PSF). Also, the typical jitter of the telescope during an *HST* observation is 0.003–0.005 arcsec, even when two guide stars are used.
- There will be small shifts between consecutive exposures. These shifts can build up between orbits in the same visit, and will affect the **AstroDrizzle** products from the pipeline because it depends on the header WCS (predicted) positions to determine image offsets when combining dithered images. Therefore, the structure of sources in the image will be degraded during the cosmic ray rejection routine. This problem can, however, be addressed during post-processing by manually running **AstroDrizzle**, after the images have been aligned using **TweakReg**.

Even when two guide stars are used, there is often a slow drift of the telescope up to 0.01 arcsec/orbit due to thermal effects. So, it is generally advisable to check the image shifts, and if necessary, align the images using the **TweakReg** task in **DrizzlePac** to improve the alignment of the exposures before running **AstroDrizzle**.

In summary, for most scientific applications, single-star guiding will not degrade the usefulness of ACS data, provided that the images are aligned (if necessary, use **TweakReg** to update the image alignment) before running **AstroDrizzle**. However, single-star guiding is not recommended for the following applications:

- Programs that require very accurate knowledge of the PSF, including coronagraphic programs and astrometric programs.
- Programs that rely critically on achieving a dithering pattern that is accurate on the sub-pixel scale. (However, note that even with two-star guiding this can often not be achieved).

Observers who are particularly concerned about the effect of pointing accuracy on the PSF can obtain quantitative insight using the [TinyTim](#) software package. While this does not have an option to simulate the effect of a linear drift, it can calculate the effect of jitter of a specified RMS value.

¹ Please check the [ACS Distortion page](#) for updated information about the skew component in the ACS distortion model.

5.3 Polarimetry

- [5.3.1 Absolute Calibration](#)
- [5.3.2 Instrumental Issues](#)
- [5.3.3 Flats](#)
- [5.3.4 Polarization Calibration](#)

5.3.1 Absolute Calibration

ACS contains a set of six filters that are sensitive to linear polarization; there are three visible polarizer filters with their polarization directions set at nominal 60° angles to each other, and three UV polarizer filters arranged in a similar manner. These filters are typically used in combination with a spectral filter which largely defines the spectral bandpass. In most cases observers will obtain images of the target in each of the three filters. The initial calibration steps for polarization data are identical to that for data taken in any other filter—the data are bias-corrected, dark-subtracted, and flat-fielded in the normal manner. The polarization calibration itself is accomplished by combining the set of images (or the resulting counts measured on the images) in the three filter rotations to produce a set of I, Q, and U images, or equivalently, a set of images giving the total intensity, fractional polarization, and polarization position angle.

i ACS/WFC polarization data are taken with a 2048 x 2048 subarray. Post-SM4 WFC subarray observations are not de-striped by default, and thus are also not corrected for CTE loss by the data pipeline for storage in MAST. Users wishing to work on de-striped, CTE-corrected data will need to set the PCTECORR flag to "PERFORM" in the primary header, update the PCTETAB keyword to the correct reference file, and run `acs_destripe_plus` from the `acstools` Python package. See the [subarray data processing example notebook](#) for assistance.

5.3.2 Instrumental Issues

The design of ACS is far from ideal for polarimetry. Both the HRC and WFC optical chains contain three tilted mirrors and utilize tilted CCD detectors. These tilted components will produce significant polarization effects within the instrument that must be calibrated out for accurate results. There are two primary effects in the tilted components—diattenuation and phase retardance. Diattenuation refers to the fact that a tilted component will likely have different reflectivities (or transmissions) for light which is polarized parallel and perpendicular to the plane of the tilt. This can be an important source of instrumental polarization, and can also alter the position angle of the polarization E-vector. The second effect, phase retardance, will tend to convert incident linear polarized light into elliptically polarized light. These effects will have complex dependencies on position angle of the polarization E-vector, and hence will be difficult to fully calibrate. Additional discussion of these effects can be found in [WFPC2 ISR 1997-11](#), [ACS ISR 2004-10](#), and [ACS ISR 2007-10](#).

The instrumental polarization, defined as the instrument's response to an unpolarized target, provides a simple measure of some of these effects. [Figure 5.5](#) shows the instrumental polarization derived for the HRC through on-orbit observations of unpolarized stars (*HST* programs [9586](#) and [9661](#)). The instrumental polarization is approximately 5% at the red end of the spectrum, but rises in the UV to about 14% at the shortest wavelengths. Also shown is a rough model for the effects of the M3 mirror together with a very crude model of the CCD. The mirror is aluminum with a 606 Å thick overcoat of Magnesium Fluoride and has an incidence angle of 47° . Since details of the CCD are proprietary, it has been simply modeled as Silicon at an incidence angle of 31° ; no doubt this is a serious over-simplification. [Figure 5.6](#) shows the same plot for the WFC, which has an instrumental polarization around 2%. Here the IM3 mirror is a proprietary Denton

enhanced Silver Coating with an incidence angle of 49° , and the CCD has an incidence angle of 20° . While the lower instrumental polarization of the WFC seems attractive, users are cautioned that the phase retardance effects are not known for the Denton coating, and have some potential to cause serious problems—if sufficiently large, the retardance could produce a large component of elliptical polarization which will be difficult to analyze with the linear polarizers downstream.

Figure 5.5: Instrumental Polarization for the HRC





Figure 5.6: Instrumental Polarization for the WFC



The ACS polarizer filters were characterized prior to installation in ACS by Leviton and the results as summarized in [Figure 6.1](#) of the *ACS Instrument Handbook*. The cross-polarized transmissions are essentially zero for the POLV set, but are significant in the UV (5% to 10%) and far-red (20%) for the POLUV set.

One further issue for polarizer data is added geometric distortion. The polarizers contain a weak lens which corrects the optical focus for the presence of two filters in the light path. The lens causes a large scale distortion which appears to be well-corrected by the drizzle software. There is also, however, a weak (± 0.3 pixel) small scale distortion in the images caused by slight ripples in the polarizing material. There is presently no correction available for this. There is also the possibility of polarimetric field dependencies; while there has been study of intensity flats for the polarizers, the polarization field dependencies are not known.

5.3.3 Flats

Flat fields for the ACS polarizers were obtained in the laboratory and corrected for low frequency variations using the in-flight L-flat corrections which were derived for the standard (non-polarizer) filters. The pivot wavelength of the combined optical components is typically within 1% when the standard filters are used in combination with the polarizers instead of with the clear filters. To assess the accuracy of this approximation, in-flight observations of the bright Earth using the F435W+POLUV filters were compared with the F475W+POLV filters with the corrected laboratory flats. The HRC Earth flats agreed with the corrected lab flats to better than 1%, where the largest deviations occurred near the edges of the detector.

5.3.4 Polarization Calibration

An extensive series of on-orbit polarization calibration observations were carried out in Cycles 11 and 12 (programs [9586](#), [9661](#), and [10055](#)). These included observations of unpolarized and polarized standard stars, the star cluster 47 Tuc, and an extended reflection nebula. Additional observations of polarized standards were taken over a wide and well-sampled range of *HST* roll angles to help quantify the angular dependences which are expected as the wavefront interacts with the diattenuation and phase retardation in the mirrors and CCD.

Preliminary calibrations, based largely on data in programs [9586](#) and [9661](#), are available for use by polarization observers. The number of polarimetric observations obtained with ACS is very small compared to other modes. As a result of this, the polarimetric mode has not been calibrated as precisely as other modes because of limited resources. For details, please see Cracraft & Sparks 2007 ([ACS ISR 2007-10](#)).

The ACS Team is currently working to improve the calibration of the polarizers using data from two programs (13964 and 14407). These proposals follow up earlier polarimetry calibration programs. Observations will be taken of the polarized star, Vela1-81 (GSC 08169-01120), and the unpolarized source G191B2B (EGGR-247). The stars will be observed through the F435W, F475W, F606W, F658N, F660N, and F775W filters with their respective polarizers. Each of the filter/polarizer combinations will be used at three different roll angles, in order to better quantify the instrumental linear polarization (such as from mirrors in the camera), and errors in the polarization angles of the filters. The F606W observations were done at a single roll angle as a sanity check of earlier observations. Updates will be posted on the [ACS website](#).

The strategy was to create calibrations based upon the small amount of data that have been analyzed. This calibration can be applied to either aperture photometry results, or to the images themselves (i.e., for an extended target). The calibration process began with the polarization "zeropoint" using corrections which were derived from observations of unpolarized standard stars.

i An update to the **acstools** Python package in December 2020 adds the **polarization_tools** module that provides conveniences for the following equations and calibration coefficients. Astropy table versions of the efficiency corrections ([Table 5.6](#)) and cross-polarization leak terms ([Table 5.7](#)) are included for convenience. An [example notebook](#) demonstrating the basic functionality of the **polarization_tools** module is available.

Efficiency corrections $C(\text{CCD}, \text{POLnXX}, \text{spectral filter}, n)$ were applied to the observed count rate r_{obs} in each of the three polarizers (POLnUV or POLnV, where $n = 0, 60, 120$). These corrections are tabulated in [Table 5.6](#), and have been scaled such that Stokes I will approximate the count rate seen with no polarizing filter.

$$r_n = C(\text{CCD}, \text{POLnXX}, \text{spectral filter}, n)r_{\text{obs},n}$$

Next, an "instrumental" Stokes vector is computed for the target.

$$\begin{aligned} I &= \frac{2}{3}(r_0 + r_{60} + r_{120}) \\ Q &= \frac{2}{3}(2r_0 - r_{60} - r_{120}) \\ U &= \frac{2}{\sqrt{3}}(r_{60} - r_{120}) \end{aligned}$$

Next, the fractional polarization of the target is computed. Also included is a factor which corrects for cross-polarization leakage in the polarizing filters (see Table 5.7 for the average correction factor per filter for a spectrum flat in wavelength).

$$P = \frac{\sqrt{Q^2 + U^2}}{I} \times \left[\frac{T_{\text{par}} + T_{\text{perp}}}{T_{\text{par}} - T_{\text{perp}}} \right]$$

Finally, the position angle on the sky of the polarization E-vector is computed. The parameter PAV3 is the roll angle of the HST spacecraft, and is called PA_V3 in the data headers. The parameter χ contains information about the camera geometry which is derived from the design specifications; for HRC, $\chi = -69.4^\circ$, and for the WFC, $\chi = -38.2^\circ$. Note that the arc tangent function must be properly defined; here, the result is defined as positive in quadrants I and II, and negative in III and IV.

$$\theta = \frac{1}{2} \tan^{-1} \left(\frac{U}{Q} \right) + \text{PAV3} + \chi$$

For example, a target that gives 65192, 71686, and 66296 counts per second in the HRC with F606W and POL0V, POL60V, and POL120V, respectively, is found to be 5.9% polarized at PA = 96.9°.

The full instrumental effects and the above calibration have been modeled together in an effort to determine the impacts of the remaining uncalibrated systematic errors. These will cause the fractional polarizations to be uncertain at the one-part-in-ten level (e.g., a 20% polarization has an uncertainty of 2%) for highly polarized sources; and at about the 1% level for weakly polarized targets. The position angles will have an uncertainty of about 3°. (This is in addition to uncertainties which arise from photon statistics in the observer's data.) This calibration has been checked against polarized standard stars (~5% polarized) and found it to be reliable within the stated errors. Better accuracy will require improved models for the mirror and detector properties as well as additional on-orbit data. No calibration has been provided for F220W, F250W, or F814W, as they are believed to be too unreliable at this time. There is also some evidence of a polarization pathology in the F625W filter, and observers should be cautious of it until the situation is better understood. In addition, one incidence of a 5° PA error for F775W has been observed, suggesting this waveband is not calibrated as well as the others.

Table 5.6: Efficiency Correction Factors C(CCD, POLnXX, specfilt, n) for Polarization Zeropoint

CCD	POL Filter	Spectral Filter	n = 0	n = 60	n = 120
HRC	POLUV	F330W	1.7302	1.5302	1.6451
	POLV	F435W	1.6378	1.4113	1.4762
	POLV	F475W	1.5651	1.4326	1.3943
	POLV	F606W	1.4324	1.3067	1.2902
	POLV	F625W ¹	1.0443	0.9788	0.9797
	POLV	F658N ¹	1.0614	0.9708	0.9730

	POLV	F775W	1.0867	1.0106	1.0442
WFC	POLV	F475W	1.4303	1.4717	1.4269
	POLV	F606W	1.3314	1.3607	1.3094
	POLV	F775W	0.9965	1.0255	1.0071

¹ Not scaled for Stokes I

Table 5.7: Flat-Spectrum (with Wavelength) Average Cross-Polarization Leak Correction Factors

CCD	POL Filter	Spectral Filter	T_{parallel}	$T_{\text{perpendicular}}$	Leak Correction
HRC	POLUV	F330W	0.4810	0.0470	1.2167
	POLUV	F435W	0.5247	0.0416	1.1724
	POLV	F606W	0.5158	5.591×10^{-5}	1.0002
	POLV	F625W	0.5147	2.874×10^{-5}	1.0001
	POLV	F658N	0.5174	2.355×10^{-5}	1.0001
	POLV	F775W	0.6043	0.0732	1.2758
WFC	POLV	F475W	0.4239	1.524×10^{-4}	1.0001
	POLV	F606W	0.5157	5.591×10^{-5}	1.0002
	POLV	F775W	0.6041	0.0737	1.2778

5.4 Coronagraphy

- [5.4.1 Flat Fielding](#)
- [5.4.2 PSF Subtraction](#)
- [5.4.3 Photometry](#)
- [5.4.4 Ghosting](#)

The HRC coronagraph allowed high contrast imaging of faint point or extended sources around bright stars or AGN by preventing saturation of the detector and suppressing the diffraction pattern of the bright central source.

Pipeline calibration of coronagraphic HRC images match that of direct HRC images except in the flat-fielding stage. The coronagraphic images require division by a pixel-to-pixel flat and an extra flat that contains the vignetting shadows of the occulting spots appropriately shifted to their locations on the observation date.

5.4.1 Flat Fielding

The coronagraphic flat fields differ from normal HRC flats in two ways. The Lyot stop alters both the large scale flat-field illumination pattern and the diffraction patterns from dust specks and filter pinholes. The occulting spots also cast shadows in the center and the upper left corner of the field of view. Because the spots are in the aberrated *HST* beam, the shadows are vignettted up to 0.5" beyond the nominal edges of the spots. This vignetting makes the spots appear diffuse. The occulting spots wander by several pixels over weekly time scales, so their effects must be treated separately from the static features in the flat field. Consequently, there are two distinct flat-field reference images that must be applied during calibration of coronagraphic science images. See Krist et al. ([ACS ISR 2004-16](#)) for detailed information.

The first reference flat is a static pixel-to-pixel flat (header keyword `PFLTFILE`, `pfl.fits`) that contains the dust specks, detector response patterns, etc. For the five supported coronagraphic filters, static flats have been created using either ground-based images (F606W, F814W) or on-orbit "Earth flats" (F330W, F435W, F475W). In the static flats, the occulting spot shadows have been replaced with the corresponding regions from the standard direct-imaging flats for the same filters. These modified static flats are accurate to better than 2% over large spatial scales, but they may be less accurate for pixel-to-pixel response. For all other filters, the standard direct-imaging HRC flats are used instead of static flats, which may cause local errors of up to 10% at the locations of dust spots. Observers should examine the flat fields used to calibrate their data to determine whether any features in their science images are caused by these local flat-field errors.

The second reference flat is a spot flat field (header keyword `CFLTFILE`, `cfl.fits`) which contains the vignettted shadows of the occulting spots. There is a different spot flat for each fully supported filter (F330W, F435W, F475W, F606W, F814W). These spot flats must be shifted so that the locations of the shadows match those at the time of the science observations. STScI regularly monitored the location of the small occulting spot using Earth flats. A table of spot position versus time is available as a calibration reference files (header keyword `SPOTTAB`, `csp.fits`). The ACS calibration pipeline finds the spot location in the table that is closest to the observation date and then shifts the spot flat by the required offset. The coronagraphic science image is then divided by the product of the static reference flat and the shifted spot flat.

5.4.2 PSF Subtraction

Although the coronagraph suppresses diffracted light from the occulted source, it does reduce the halo of scattered light created by the polishing errors in *HST*'s optics. Scattered light must be subtracted using an image of an isolated star whose color is similar to that of the science target. Observers should observe a PSF reference star immediately before or after their science observations.

The PSF reference image must be scaled and aligned to match the intensity and position of the occulted science target. The scale factor and alignment offset can be derived from non-coronagraphic imaging of the target and reference sources or by normalizing and aligning the coronagraphic PSF reference image with the science image. Sub-pixel registration is typically required and is best accomplished by iteratively shifting the PSF reference image until the subtraction residuals are minimized and symmetrical about the source (notwithstanding any potential circumstellar material).

Higher-order interpolation should be used because of the high frequency structure in the scattered light halo (Bi-linear interpolation is usually insufficient.) The **astropy.photutils.psf** module contains tools for fitting and subtracting PSFs.

5.4.3 Photometry

The coronagraph's Lyot stop alters the PSF of the field sources and reduces the system throughput by 52.5% relative to the normal HRC imaging mode. The stop broadens the field PSF and places more light into the diffraction spikes and Airy rings (see [Chapter 5](#) in the *ACS Instrument Handbook*). These effects must be considered when performing photometric measurements of field sources. Note that the reduction in throughput is included in **pysynphot** when the `coron` specifier is included in the observing mode (e.g., "`acs, hrc, f606w, coron`").

5.4.4 Ghosting

A point-like ghost appears about 30 pixels to the lower right of every field star in HRC coronagraph images. Each ghost is about 8 magnitudes fainter than its associated field star. There are no known ghosts associated with the occulted target source.

5.5 Ramp Filters

5.5.1 Introduction

5.5.2 Existing Ground Calibrations

5.5.3 In-Flight Calibrations

5.5.1 Introduction

Five slots on ACS filter wheel #2 each contain a ramp filter unit that is made up of an inner, an outer, and a middle segment. For the sake of simplicity, these three segments will be referred to as individual inner (IRAMP), outer (ORAMP), and middle (MRAMP) ramp "filters" although strictly speaking, they form part of the same physical ramp filter unit.

All fifteen ramp filters can be used with the WFC over specified regions of the WFC1 and WFC2 chips. IRAMP and ORAMP filters can only be placed on the WFC1 or WFC2 chips, respectively, over regions which define the WFC1-IRAMP and WFC2-ORAMP apertures. The MRAMP filter is, by default, placed on the WFC1 chip over the WFC1-MRAMP aperture, but it can, in principle, also be placed on the WFC2 chip over the WFC2-MRAMP aperture. The different ramp apertures and their reference points on the WFC chip are shown in [Table 7.6](#) and [Figure 7.4](#) of the *ACS Instrument Handbook*. In practice, the observer specifies a ramp filter and a central wavelength, and the filter wheel is appropriately rotated in order to place the central wavelength at the reference point of the relevant aperture.

With the HRC, only the five middle ramp filters were available and they cover the region over the HRC chip defined by the HRC aperture ([Table 7.7](#), *ACS Instrument Handbook*).

The fifteen ramp filters provided a total wavelength coverage of 3700 Å–10,700 Å. Twelve of the filters have a narrow bandpass ($\Delta\lambda/\lambda$) of 2% and three have a medium bandpass of 9% ([Table 5.2](#), *ACS Instrument Handbook*).

5.5.2 Existing Ground Calibrations

For each of the 15 ramp filters and different specified central wavelengths, the transmission was measured as a function of wavelength over the bandpass. For each ramp, analytical fits to these measurements generated eleven transmission profiles spaced by 10% of the wavelength coverage of the ramp ([ACS ISR 2000-05](#)). Using these delivered profiles, **pysynphot** interpolates the transmission curve for any wavelength within the bandpass of each ramp filter. The interpolation errors are below 1% which is better than the measurement accuracy.

WFC and HRC flat fields have been derived from pre-flight data by illuminating the ramps with continuum and monochromatic light using the Refractive Aberrated Simulator/Hubble Opto-Mechanical Simulator (RAS/HOMS) ([ACS ISR 2002-01](#)). Each of the five WFC ramp filter units associated with the five wheel #2 slots is made up of three segments (inner, outer, middle) which cover three different adjacent wavelength ranges. A composite WFC flat was made for each of the five ramp filter units by combining three pieces of full field broad-band surrogate LP-flats (flats that incorporate both low frequency and pixel-to-pixel variations) nearest in central wavelength. For instance, for the ramp filter unit made up of FR505N (middle), FR551N (inner), and FR601N (outer) segments, surrogate broad-band LP-flats of F475W, F555W, and F606W were used.

HRC ramp flats were made from surrogate HRC flats appropriate for the middle segment. The broad-band flats reduce the rms structure in monochromatic P-flats from 1% to 0.2%, except at the shortest measured wavelength of 3880 Å.

After flat-fielding some test ramp images with the relevant composite L-flats, correcting for anamorphic geometric distortion, rotating, and median filtering to remove residual noise, the response for FR505N is found to be uniform to below 1% over a 30×80 arcsec field of view for the WFC and 11.5×24 arcsec FOV for the HRC. Out of the five middle ramp filters, the FR388N shows the least uniform response ($< 10\%$) over such regions ([ACS ISR 2002-01](#)).

Values for the dispersion of the middle ramp filters FR388N, FR459M, FR505N, FR656N, and FR914M are, respectively, 0.192, 0.0399, 0.146, 0.120, 0.0219 mm/Å at the ramp filters' wavelength.

Due to the extra divergence of the beam from the filter to the detector, the dispersion at the WFC and HRC detector is estimated to be a factor of 1.2 and 1.3 lower respectively.

5.5.3 In-Flight Calibrations

An attempt has been made to check the wavelength calibration of the ramp filters by crossing them with the grism ([ACS ISR 2007-11](#)). In these tests, the wavelength setting of HRC ramps agreed well with the calibration measurements. In contrast, some of the WFC ramp filters showed notable offsets from their expected centers when tested in this fashion.

There is reason to believe, however, that the offsets seen in the WFC ramps may be caused by a "filter wedge" produced by using the grism and ramp filters in combination. One of the filters that showed the largest offset with respect to its width was FR782N. This filter was used to observe red-shifted H α sources, and showed that the *HST* fluxes agreed with those measured from the ground to 20%, which is within their continuum subtraction error (Ovezier and Heckman, personal communication). For now, observers are advised to use the standard settings of the ramps for their observations when using the WFC.

5.6 Spectroscopy with the ACS Grisms and Prisms

5.6.1 What to Expect from ACS Slitless Spectroscopy Data

5.6.2 Pipeline Calibration

5.6.3 Slitless Spectroscopy Data and Dithering

5.6.4 Extracting and Calibrating Slitless Spectra

5.6.5 Accuracy of Slitless Spectra Wavelength and Flux Calibration

The ACS WFC has a grism (G800L) which provides slitless spectra at a dispersion of $\sim 40 \text{ \AA/pixel}$ in the first order over the whole field of view. The same grism, used with the HRC, produced higher dispersion spectra ($\sim 23 \text{ \AA}$ in first order).

A prism (PR200L) available with the HRC provided slitless spectra from the UV cut-off to $\sim 3500 \text{ \AA}$.

The SBC is fitted with two prisms (PR110L and PR130L) which provide far-UV spectra, the latter blocked below 1230 \AA to prevent transmission of Lyman α .

All of the ACS spectroscopic modes present the well-known advantages and disadvantages of slitless spectroscopy. The chief advantage is area coverage enabling spectroscopic surveys. Among the disadvantages are overlap of spectra, high background from the integrated signal over the passband and modulation of the resolving power by the sizes of dispersed objects. For ACS, the low background from space and the narrow PSF make slitless spectra competitively sensitive to ground-based spectra from larger telescopes in the presence of sky background.

The primary aim of reduction of slitless spectra is to provide one-dimensional wavelength and flux calibrated spectra of all the objects with detectable spectra. The reduction presents special problems on account of the dependence of the wavelength zero point on the position of the dispersing object, the blending of spectra in crowded fields, and the necessity of flat fields over the whole available wavelength range. For ACS slitless modes, a dedicated package was developed by the ST-ECF to enable automatic and reliable extraction of large numbers of spectra. This package, called **aXe**, is sufficiently general to be used for other slitless spectra applications. A separate software package, **hstaxe** (available at <https://github.com/spacetelescope/hstaxe>), which is a follow-up to **aXe**, can now be used to extract and calibrate spectra from ACS grism exposures. Additional **aXe**-related information is available at: <https://www.stsci.edu/scientific-community/software/axe>.

5.6.1 What to Expect from ACS Slitless Spectroscopy Data

The default method for taking ACS slitless spectra is to take a pair of images, one direct and one dispersed, in the same orbit (using special requirement AUTOIMAGE=YES¹, see the [STScI Phase II Proposal Instructions](#) for details). Or, at the least, to take them with the same set of guide stars and without a shift in-between imaging and spectroscopic observations. The filter for the direct image is set by the AUTOIMAGE=YES parameter. Alternately, a direct image using a filter in the range of the spectral sensitivity can be explicitly specified in the exposure logsheet; in that case, AUTOIMAGE should be set to "NO."

The direct image provides the reference position for the spectrum and thus sets the pixel coordinates of the wavelength zero point on the dispersed image. The adequacy of the direct image to provide the wavelength zero point can be assessed by checking the WCS of the direct and dispersed images. In general this can be trusted at the level of a pixel or better. Detailed information about the pointing during an observation can be obtained from the jitter files.

In the case of a direct image not being available, such as when the pairing of a dispersed image with a direct image was suppressed (AUTOIMAGE=NO), then it is possible, for G800L grism observations, to use the zeroth


order to define the wavelength zero point. However the detected flux in the zeroth order is a few percent of that in the first order and it is dispersed by about 5 pixels, so the quality of the wavelength information will be degraded.

As an example, [Figure 5.7](#) shows an ACS image (one chip only) taken with F775W and [Figure 5.8](#) shows the companion G800L dispersed image. This illustrates the general characteristics of slitless spectra and the characteristics of ACS spectra in particular.

Figure 5.7: Direct Image of WFC1 in F775W



Figure 5.8: WFC1 G800L Slitless Image Corresponding to the Direct Image



Bright Stars

The brightest objects produce spectra which can extend far across the chip. For the G800L grism, only about 10% of the flux is in orders other than the first, but for bright objects, orders up to +4 and -4 can be detected. These spectra, while in principle can be analyzed², are a strong source of contamination for fainter spectra. In addition, for the grism, higher order spectra are increasingly out of focus and thus spread in the cross-dispersion direction.

In the case of the prisms, the strongly non-linear dispersion can result in heavily saturated spectra to longer wavelengths. An additional effect for bright stars is the spatially extended spectra formed by the wings of the PSF.

Resolution and Object Size

In slitless spectroscopy, the object itself provides the "slit." The ACS PSF has a high Strehl ratio for most channels and over most of the accessible wavelength range of the dispersing elements. Therefore, the degradation of point sources beyond the theoretical resolution is minimal. However, for an extended object, the spectral resolution will be degraded depending on the size and light distribution in the object, and spectral features will be diluted (see [ACS ISR 2001-02](#)).

Zeroth Order

The grism zeroth order, only detectable for brighter objects since it contains about 3% of the total flux, can be mistaken for an emission line. The direct image can be used to determine the position of the zeroth order, and for brighter sources, distinguish unequivocally between the zeroth order and an emission line.

Background

The background in a single pixel is the result of the transmission across the whole spectral range of the disperser and can thus be high depending on the spectrum of the sky background. The SBC PR130L prism, for example, provides lower detected background than the PR110L since it excludes the geocoronal Lyman α

line. Bright objects in the field elevate the local background both in the dispersion direction and in the cross-dispersion direction. This background needs to be carefully removed before extracting the spectra of targets.

Crowding

Although [Figure 5.7](#) and [Figure 5.8](#) show a relatively uncrowded field, close examination shows many instances where spectra overlap, either in the dispersion direction or between adjacent spectra. It is important to know if a given spectrum is contaminated by a neighbor. This can be done by obtaining slitless spectra at a number of different roll angles, which improves the chances of cleanly extracting spectra that overlap in the dispersion direction.

Extra Field Objects

There will inevitably be cases of objects outside the field which produce detectable spectra. This is more serious for the HRC and SBC where the fields are small and the spectra are long relative to the size of the detector. In such cases reliable wavelengths cannot be assigned since the zero point of the wavelength scale cannot be determined unless the zeroth order is also present. Brighter sources are a source of contamination, and correct flagging is hampered by the lack of knowledge of the precise image position.

5.6.2 Pipeline Calibration

The direct image obtained before a grism-dispersed observation is fully reduced by **calacs**. However, the only pipeline steps applied to a dispersed image involve noise and data quality array initializations, linearity corrections for the SBC, and for the CCDs, bias subtraction, dark correction, removal of the image overscan areas, and sink pixel flagging (if the data were taken after Jan 1, 2015) will also take place. The grism-dispersed images are also flat-fielded using a dummy flat field during this step, and the data units are converted to electrons (which includes gain conversions for the CCDs³). [Table 5.7](#) shows the calibration switches appropriate for a WFC G800L frame.

A dummy flat, filled with the value "1" for each pixel, is used because no single flat-field image can be correctly applied to slitless spectroscopy data, since each pixel cannot be associated with a unique wavelength. Rather, flat-fielding is applied later during the extraction of spectra using an **aXe/hstaxe** task called **aXe_PETFF** (see [Section 5.6.4](#)). Each pixel receives a flat-field correction dependent on the wavelength falling on that pixel as specified by the position of the direct image and the dispersion solution. No unique photometric keywords can be attached to all the spectra in a slitless image, so the photometric header keywords are left as default values as shown in [Table 5.8](#).

For all subsequent reductions of slitless data using **aXe/hstaxe**, the pipeline flt or `flt.fits` files should be the starting point.

Table 5.7: Calibration Switch Settings for ACS WFC Slitless Spectrometry Frames

Keyword	Switch	Description
STATFLAG	F	calculate statistics? (T/F)
WRterr	T	write out error array extension? (T/F)
DQICORR	PERFORM	data quality initialization?
ATODCORR	OMIT	correct for A to D conversion errors?
BLEVCORR	PERFORM	subtract bias level computed from overscan image?

BIASCORR	PERFORM	subtract bias image?
FLSHCORR	OMIT	subtract post-flash image?
CRCORR	OMIT	combine observations to reject cosmic rays?
EXPSCORR	PERFORM	process individual observations after cr-reject?
SHADCORR	OMIT	apply shutter shading correction?
DARKCORR	PERFORM	subtract dark image?
FLATCORR	PERFORM	apply flat field correction?
PHOTCORR	OMIT	populate photometric header keywords?
RPTCORR	OMIT	add individual repeat observations?
DRIZCORR	PERFORM	process dithered images?
SINKCORR	OMIT	flag sink pixels? (PERFORM IF WFC data with date > Jan 1 2015)

Table 5.8: ACS Slitless Spectra (Grism or Prism) Photometric Keywords

Keyword	Value	Description
PHOTMODE	'ACS WFC2 G800L'	observation con
PHOTFLAM	0.00000E+00	inverse sensitivity, ergs/cm ² /Ang/electron
PHOTZPT	0.00000	ST magnitude zero point
PHOTPLAM	0.00000	Pivot wavelength (Angstroms)
PHOTBW	0.00000	RMS bandwidth of filter plus detector

5.6.3 Slitless Spectroscopy Data and Dithering

The common desire to dither ACS imaging data in order to improve the spatial resolution, and to facilitate removal of hot pixels and cosmic rays, applies equally well to slitless spectroscopy data. Typically, for long exposures, and especially for parallel observations, the data is obtained as several sub-orbit dithered exposures.

AstroDrizzle (see the [DrizzlePac website](#)) corrects for the large geometrical distortion of the ACS and provides a very convenient tool for combining ACS datasets. However, slitless grism data have both a spatial and a spectral dimension, but the correction for geometric distortion is only applicable to the cross-dispersion direction.

aXe/hstaxe ACS grism dispersion solutions are computed for images that have not been corrected for geometric distortion (see [ISR 2003-01](#), [2003-07](#)). Tracing the light path for slitless spectroscopic data shows that the distortions apply to the position and shape of images as they impinge on the grism, but not from the grism/prism to the detector. Extraction of drizzled (geometrically-corrected) slitless data for WFC, the ACS channel with the largest geometric distortion, have shown that the correction has the effect of decreasing

the resolution (increasing the dispersion in Å/pixel) across the whole field by around 10%. An additional complexity is that the dispersion solution would have to match the set of **drizzle** parameters used (such as **pixfrac** and **scale**). This is clearly disadvantageous, therefore, *users are advised against extracting spectra from **MultiDrizzle**- or **AstroDrizzle**-combined slitless data. Individual `flt.fits` or `flc.fits` images should be used for extraction of spectra.* However, the **MultiDrizzle** or **AstroDrizzle** combination of many grism images is useful for visual assessment of the spectra.

When images are combined using **AstroDrizzle**, cosmic rays and hot pixels are detected and flagged in mask files for later use in creating the final drizzle-combined image. **AstroDrizzle** also updates the data quality array for each input `flt.fits` image with flags to denote the cosmic rays and hot pixels that were detected in them. Therefore, running **AstroDrizzle** is recommended for the sole purpose of updating each `flt.fits` file to flag cosmic rays and hot pixels. The **aXe** Spectral Extraction package (see [Section 5.6.4](#)) can then be run on these updated `flt.fits` images to extract the spectra from each image separately. A separate software package, **hstaxe** (available at <https://github.com/spacetelescope/hstaxe>), which is a follow-up to **aXe**, can now be used to extract and calibrate spectra from ACS grism exposures. The individual extracted one-dimensional (1D) wavelength vs. flux spectra, with the bad pixels excluded from the spectrum, can then be converted to individual FITS images.

However, for data taken with the G800L grism for both WFC and HRC, **aXe/hstaxe** offers, with **aXedrizzle** (see [Section 5.6.4](#)), a method to combine the individual images of the first order slitless spectra in two-dimensions (2D) into a pseudo long-slit spectrum. **aXedrizzle** derives, for each first order spectrum on each image, the transformation coefficients to co-add the pixels to deep drizzled 2D spectra. Then, the final, deep 1D spectra are extracted from the drizzled 2D spectra.

5.6.4 Extracting and Calibrating Slitless Spectra

The software package **aXe** provides a streamlined method for extracting spectra from the ACS slitless spectroscopy data. A separate software package, **hstaxe** (available at <https://github.com/spacetelescope/hstaxe>), which is a follow-up to **aXe**, can now be used to extract and calibrate spectra from ACS grism exposures.

This section provides an overview of the **aXe/hstaxe** software capabilities. For details about using **aXe**, please refer to the [aXe manual](#) and demonstration package available at: <https://github.com/spacetelescope/hstaxe/>

Preparing the aXe/hstaxe Reduction

Input data for **aXe/hstaxe** are a set of dispersed slitless images, and catalogs of the objects for the slitless images. (The object catalog, which may be displaced from the slitless images, may alternatively refer to a set of direct images.)

For each slitless mode (e.g., WFC G800L, HRC G800L, HRC PR200L, SBC PR110L and SBC PR130L), a configuration file is required which contains the parameters of the spectra. Information about the location of the spectra relative to the position of the direct image, the tilt of the spectra on the detector, the dispersion solution for various orders, the name of the flat-field image, and the sensitivity (flux per Å/e-/sec) table, is held in the configuration file that enables the full calibration of extracted spectra.

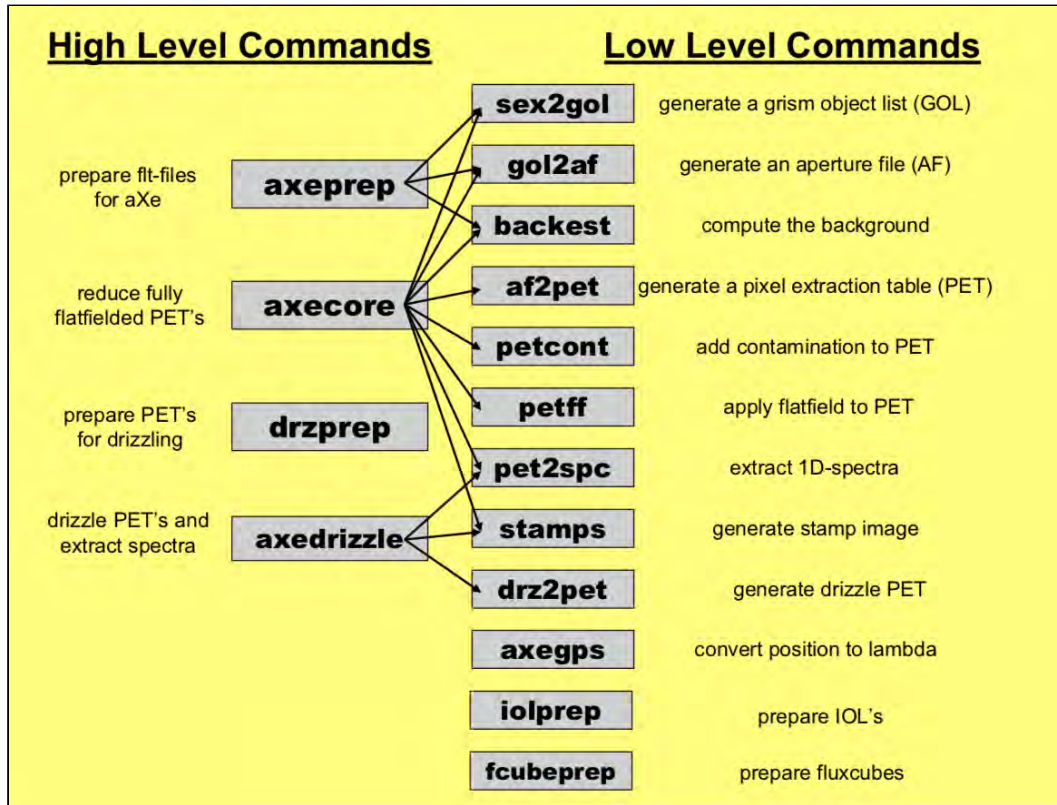
For each instrument mode, the configuration files and all necessary calibration files for flat-fielding and flux calibration can be downloaded following the links given on the [aXe Web page](#).

aXe/hstaxe Tasks

The **aXe/hstaxe** software has several tasks which can be successively used to produce extracted spectra.

There are two classes of **aXe/hstaxe** tasks. The so-called "Low Level Tasks" work on individual grism images. The "High Level Tasks" work on datasets to perform certain processing steps for a set of images. Often, the High Level Tasks use Low Level Tasks to perform reduction steps on each frame. A set of four High Level Tasks, as shown in [Figure 5.9](#), cover all steps of the **aXe/hstaxe** reduction.

Figure 5.9: Summary of aXe/hstaxe High and Low Level Commands



This figure is reproduced from the [aXe User Manual version 2.3](#).

aXedrizzle, for Processing G800L Data

The **aXe/hstaxe** data reduction is based on the individual extraction of object beams on each `flt.fits` science image. ACS datasets usually consist of several images, with small position shifts or dithers between them. The **aXedrizzle** technique offers the possibility to combine the first order beams of each object extracted from a set of dithered G800L slitless spectra images, taken with the WFC and HRC.


The input for **aXedrizzle** consists of flat-fielded and wavelength-calibrated pixels extracted for all beams on each science image. For each beam, those pixels are stored as stamp images in multi-extension FITS files.

For each beam, **aXe/hstaxe** computes transformation coefficients which are required to drizzle the single stamp images of each object onto a single deep, combined 2D spectral image. These transformation coefficients are computed such that the combined drizzled image resembles an ideal long slit spectrum with the dispersion direction parallel to the x-axis and cross-dispersion direction parallel to the y-axis. The wavelength scale and the pixel scale in the cross-dispersion direction can be set by the user with keyword settings in the **aXe/hstaxe** configuration file.

If the set of slitless images to drizzle has a range of position angles, then caution is required in the use of **aXedrizzle**. Since the shape of the dispersing object acts as the "slit" for the spectrum, the object's orientation and length may differ for extended objects which are not circular. Naively combining the slitless spectra at different position angles with **aXedrizzle** will result in an output spectrum which suffers distortions—the resultant 1D extracted spectrum would be a poor representation of the real spectrum.

To extract the final 1D spectrum from the deep 2D spectral image, **aXe/hstaxe** uses an (automatically created) adapted configuration file that takes into account the modified spectrum of the drizzled images (i.e., orthogonal wavelength and cross-dispersion, and the Å/pixel and arcsec/pixel scales).

Figure 5.10: Illustration of the aXedrizzle Process for One Object



Panel *a* shows one individual grism image with an object marked. Panel *b* displays the stamp image for this object out of the grism image *a*. Panel *c* shows the drizzled grism stamp image derived from *a*, and the final co-added 2D spectrum for this object is given in panel *d*.

Background Subtraction

aXe/hstaxe has two different strategies for removal of the sky background from the spectra.

The first strategy is to perform a global subtraction of a scaled "master sky" frame from each input spectrum image at the beginning of the reduction process. This removes the background signature from the images; the remaining signal can be assumed to originate only from the sources, and is extracted without further background correction in the **aXe/hstaxe** reduction. Master sky frames for the HRC and WFC G800L configurations are available for download from links given on the [ACS Prism/Grism web page](http://www.stsci.edu/hst/instrumentation/acs/performance/prism--grism/wfc-g800l).

The second strategy is to make a local estimate of the sky background for each beam by interpolating between the adjacent pixels on either side of the beam. In this case, an individual sky estimate is made for every beam in each science image. This individual sky estimate is processed (flat-fielded, wavelength-calibrated) parallel to the original beam. Subtracting the 1D spectrum extracted from the sky estimate from the 1D spectrum derived from the original beam results finally in the pure object spectrum.

Output Products: Extracted Spectra

The primary output of **aXe/hstaxe** is the file of extracted spectra (SPC). This is a multi-extension FITS binary table with as many table extensions as there are extracted beams. The beams are designated by an alphabetic postscript that correspond to the order defined by the configuration file. For example the G800L WFC 1st order is BEAMA in the configuration file `ACS.WFC*.conf`, see: <http://www.stsci.edu/hst/instrumentation/acs/performance/prism--grism/wfc-g800l>.

For instance, a first-order spectrum of an object designated as number 8 in an input *Sextractor* catalog, for WFC1 of the slitless image `grism.fits`, would be accessed as `grism_5.SPC.fits[BEAM_8A]`. That table contains 12 columns: wavelength; total, extracted, and background counts along with their associated errors; calibrated flux with its error and weight; a contamination flag. The header of the primary extension of the SPC table is a copy of the header of the frame from which the spectrum was extracted.

aXe/hstaxe can also create, for each beam, a stamp image for the individual inspection of single beams. The stamp images of all beams extracted from a grism image are stored as a multi-extension FITS file with each

extension containing the image of a single extracted beam. The stamp image of the first order spectrum of object number 8 from the above example would be stored as `grism_5.STP.fits[BEAM_8A]`. It is also possible to create stamp images for 2D drizzled grism images.

Other output products such as Pixel Extraction Tables, contamination maps, or Object Aperture Files are intermediate data products and their contents are not normally inspected.

Handling the aXe/hstaxe Output

The output products from **aXe/hstaxe** consist of ASCII files, FITS images, and FITS binary tables. The FITS binary tables containing wavelength, flux and flux uncertainties can be plotted as shown in [Figure 5.11](#).

When there are many detected spectra on a single image, such as is usually the case for data taken with the WFC G800L, then a dedicated task **aXe2web** is available at the [aXe Web page](#). **aXe2web** creates html pages consisting of direct image cut-outs, stamp images and 1D spectra (see [Figure 5.12](#)) for each extracted beam. This enables convenient browsing of large numbers of spectra or the publishing of **aXe/hstaxe** spectra on the web with minimal interaction.

Figure 5.11: Example WFC G800L Extracted Spectrum





Figure 5.12: Web page Produced by *aXe2web* for Browsing Extracted Spectra



Tuning the Extraction Process

All *aXe*/*hstaxe* tasks have several parameters which may be tuned in various ways to alter or improve the extraction. Please see the [aXe User Manual](#) for further details.

5.6.5 Accuracy of Slitless Spectra Wavelength and Flux Calibration

Wavelength Calibration

The dispersion solution was established by observing astronomical sources with known emission lines (e.g., for the WFC G800L and the HRC G800L, Wolf-Rayet stars were observed; see [ACS ISR 2003-01](#) and [ACS ISR 2003-07](#)). The field variation of the dispersion solution was mapped by observing the same star at different positions over the field. The internal accuracy of these dispersion solutions is good (see the ISR's for details) with an rms generally less than 0.2 pixels.

For a given object, the accuracy of the assigned wavelengths depends most sensitively on the accuracy of the zeropoint and the transfer of the zeropoint from the direct image to the slitless spectrum image. Provided that both direct and slitless images were taken with the same set of guide stars, systematic pointing offsets less than 0.2 pixels can be expected. For faint sources the error on the determination of the object centroid for the direct image will also contribute to wavelength error. Realistic zeropoint errors of up to 0.3 pixels are

representative. For the highest wavelength accuracy it is advised to perform a number of pixel offsets of about ten percent of the field size for each of several direct and slitless image pairs and average the extracted spectra.

Flux Calibration

The sensitivity of the dispersers was established by observing a spectrophotometric standard star at several positions across the field. The sensitivity (**aXe/hstaxe** uses a sensitivity tabulated in ergs/cm/cm/sec/Å per detected Å) was derived using data flat-fielded by the flat-field cube (see [Section 5.6.2](#)). In fact, the adequacy of the WFC flat field was established by comparing the detected counts in the standard star spectra at several positions across the field. These tests showed that the sensitivity at the peak of the WFC G800L varied by less than 5% across the whole field. However at the lower sensitivity edges of the spectra, to the blue (<5,800Å) and to the red (>10,000Å), the counts in the standards are low, and the errors in flux calibration approach 20%. In addition, small errors in wavelength assignment have a large effect in the blue and red where the sensitivity changes rapidly with wavelength. This often leads to strong upturns at the blue and red ends of extracted flux calibrated spectra, whose reality may be considered suspect. See [ST-ECF ISR 2008-01](#).

¹ Controls the automatic scheduling of image exposures for the purpose of spectra zero point determination of grism observations. By default, a single short image through a standard filter will be taken in conjunction with each Exposure Specification using the grism for external science observations.

² (dispersion solutions have been derived for orders -3, -2, -1, 0, 1, 2 and 3 for the WFC G800L for example; see [ACS ISR 2003-01](#))

³ This dummy flat corrects the gain offsets between the four WFC quadrants. For the HRC, which has only one amplifier, a single gain value is used. There is no gain correction for the SBC.