ACS Data Handbook

SPACE
TELESCOPE
SCIENCE
INSTITUTE

User Support

For prompt answers to any question, please contact the STScI Help Desk.

- **E-mail:** help@stsci.edu
- **Phone:** (410) 338-1082 or 1-800-544-8125

World Wide Web

Information and other resources are available on the ACS World Wide Web page:

- URL: http://www.stsci.edu/hst/acs/

ACS Revision History

Version	Date	Editor
3.0	July 2004	Cheryl Pavlovsky, Adam Riess, Jennifer Mack and Ron Gilliland, ACS Editors
2.0	November 2003	Jennifer Mack and Ron Gilliland
1.0	January 2002	Jennifer Mack, ACS Editor Bahram Mobasher, Chief Editor

Authorship

This document is written and maintained by the ACS+WFPC2 Branch in the Instruments Division of STScI which consists of the following individuals: Biretta, J., Boffi, F., Bohlin, R., Cox, C., Fruchter, A., Giavalisco, M., Gilliland, R., Gonzaga, S., Heyer, I., Hook, R., Jogee, S., Koekemoer, A., Mack, J., Mutchler, M., Pavlovsky, C., Platais, V., Riess, A., Sirianni, M., van der Marel, R., and van Orsow, D. Other scientists who contributed to the writing of this handbook include Clampin, M., de Marchi, G., Hack, W., Hartig, G., Krist, J., Sparks, W., and Welty, A. at STScI or formerly at STScI, Blakeslee, J. and Meurer, G., of the ACS Science Team, and Pasquali, A. and Walsh, J. at the ST-ECF.

In publications, refer to this document as:

Pavlovsky, C., et al. 2004, "ACS Data Handbook", Version 3.0, (Baltimore: STScI).

Table of Contents

Chapter 4: Distortion Correction and Drizzling of ACS Images 4-1

Preface

How to Use this Handbook

This handbook is designed to help users manipulate, 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 and is designed for users familiar with HST data but new to ACS. Users who wish to find more general information including instructions for retrieving data from the HST archive, a description of HST file formats, and a discussion of IRAF/STSDAS software for displaying and processing these data, are referred to a companion volume, the Introduction to the *HST Data Handbook* at:

http://www.stsci.edu/hst/HST_overview/documents/datahandbook

While the purpose of this handbook is to provide comprehensive information for the treatment of ACS data and present the in-flight data calibration program results, a brief discussion of the ACS capabilities, design and basic instrument operations (Chapter 1) is also included. For a detailed discussion please refer to the ACS Instrument Handbook. One important concern for users is the large volume of ACS data compared to other HST instruments, addressed in Chapter 2, with a discussion of the ACS file structure. Chapter 3 describes the ACS calibration pipeline software, the ACS reference files, and instructions for manually recalibrating data. Given the large ACS field of view, and hence the importance of geometric distortion, Chapter 4 is devoted to a thorough discussion of the ACS distortion. This includes the new Python package **PyDrizzle** which corrects for distortion in ACS images and drizzles multiple dithered images within **CALACS**. We also recommend the new **MultiDrizzle** software for users who wish to improve or fine-tune the drizzling process, choosing parameters specific to their scientific goals. Chapter 5 describes the sources of uncertainty in ACS data and chapter 6 describes the data reduction issues relevant for work with ACS data.

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

Cheryl Pavlovsky, Adam Riess, Jennifer Mack & Ron Gilliland (Editors, ACS Data Handbook)

ACS Overview

In this chapter. . .

This chapter provides an overview of the capabilities and design of the Advanced Camera for Surveys (ACS) and describes the basic instrument operations. For more information we refer you to the *ACS Instrument Handbook* which gives a technical description of the instrument's properties, expected performance, operations, and calibration.

1.1 Instrument Design and Capabilities

The ACS camera is designed to provide HST with a deep, wide-field survey capability from the visible to near-IR, high resolution imaging from the near-UV to the near-IR, and solar blind far-UV imaging. The primary design goal of the ACS Wide-Field Channel is to achieve a factor of 10 improvement in discovery efficiency compared to WFPC2, where discovery efficiency is defined as the product of imaging field of view (FOV) and instrument throughput.

ACS comprises three channels, each optimized for a specific goal:

- Wide field channel (WFC): ~202 × 202 arcsecond field of view from 3700–11,000 Å, and peak efficiency of 48% (including the Optical Telescope Assembly). The plate scale of ~0.049 arcsec/pixel provides critical sampling at 11,600 Å.

- High resolution channel (HRC): 29 × 25 arcsecond field of view from 2000–11,000 Å, and peak efficiency of 29%. The plate scale of ~0.027 arcsec/pixel provides critical sampling at 6300 Å.

- Solar Blind Channel (SBC): 35 × 31 arcsecond field of view from 1150-1700 Å, and peak efficiency of 7.5%. The plate scale of ~0.032 arcsec/pixel provides a good compromise between resolution and field of view.

In addition to the these three prime capabilities, ACS also provides:

- Grism spectroscopy: Low resolution (R~100) wide field spectroscopy from 5500–11,000 Å available in both the WFC and the HRC.

- Objective prism spectroscopy: Low resolution (R~100 @ 2000 Å) near-UV spectroscopy from 2000–4000 Å available in the HRC.

- Objective prism spectroscopy: Low resolution (R~100 @ 1216 Å) far-UV spectroscopy from 1150–1700 Å available in the SBC.

- Coronagraphy: Aberrated beam coronagraphy in the HRC from 2000–11,000 Å with 1.8 arcsecond and 3.0 arcsecond diameter occulting spots.

- Imaging Polarimetry: Polarimetric imaging in the HRC and WFC with relative polarization angles of 0°, 60° and 120°.

ACS is a versatile instrument that can be applied to a broad range of scientific programs. For example, the high sensitivity and wide field of the WFC in the visible and near-infrared will make it the instrument of choice for deep imaging programs in this wavelength region. The HRC, with its excellent spatial resolution, provides full sampling of the HST PSF at $\lambda > 6000$ Å and can be used for high precision photometry in stellar population programs. The HRC coronagraph can be used for the detection of circumstellar disks and QSO host galaxies.

1.1.1 Detectors

ACS uses one or more large-format detectors in each channel:

- The WFC detector, called `ACS/WFC,` employs a mosaic of two 2048 × 4096 Scientific Imaging Technologies (SITe) CCDs, with ~0.049 arcsecond pixels, covering a nominal ~202 × 202 arcsecond field of view, and a spectral response from ~3700 to 11,000 Å.

- The HRC detector, called `ACS/HRC,` is a 1024 × 1024 SITe CCD, with ~0.028 × 0.025 arcsecond pixels, covering a nominal 29 × 25 arcsecond field of view, and spectral response from ~2000 to 11,000 Å.

- The SBC detector, called the `ACS/SBC,` is a solar-blind CsI Multi-Anode Microchannel Array (MAMA), with 1024 × 1024 ~0.034 × 0.030 arcsecond pixels, and a nominal 35 × 31 arcsecond field of view, with far-UV spectral response from ~1150 to 1700 Å.

The WFC & HRC CCDs

The ACS CCDs are thinned, backside-illuminated devices cooled by thermo-electric cooler (TEC) stacks and housed in sealed, evacuated dewars with fused silica windows. The spectral response of the WFC CCDs is optimized for imaging at visible to near-IR wavelengths. The HRC CCD covers wavelengths similar to the WFC but the spectral response has been additionally optimized for the near-UV. Both CCD cameras produce a time-integrated image in the `ACCUM` data-taking mode. The HRC can also be operated in target acquisition (ACQ) mode for coronagraphic observations. As with all CCD detectors, there is noise (*readout noise*) and time (*read time*) associated with reading out the detector following an exposure. The minimum exposure time is 0.1 *sec* for HRC, and 0.5 *sec* for WFC, and the minimum time between successive identical exposures is 45 *sec* (HRC) or ~135 *sec* (WFC) for full-frame and can be reduced to ~36 *sec* for subarray readouts. The dynamic range for a single exposure is ultimately limited by the depth of the CCD full well (~85,000 e^- for the WFC and 155,000 e^- for the HRC), which determines the total amount of charge that can accumulate in any one pixel during an exposure without saturation. Cosmic rays will affect all CCD exposures: CCD observations should be broken into multiple exposures whenever possible, to allow removal of cosmic rays in post-observation data processing.

The SBC MAMA

The SBC MAMA is a *photon-counting* detector which provides a two-dimensional ultraviolet capability, optimized for the far-UV. It can only be operated in `ACCUM` mode. The ACS MAMA detector is subject to both *scientific* and *absolute* brightness limits. At high local (≥50 *counts sec^{-1} pixel^{-1}*) and global (>285,000 *counts sec^{-1}*) illumination rates, counting becomes nonlinear in a way that is not correctable. At only slightly higher illumination rates, the MAMA detectors are subject to damage. We have therefore defined absolute local and global count-rate limits, which translate to a set of configuration-dependent bright-object screening limits. Sources which violate the absolute count rate limits in a given configuration *cannot be observed in those configurations*.

1.1.2 ACS Optical Design

The ACS design incorporates two main optical channels: one for the WFC and one which is shared by the HRC and SBC. These channels are illustrated in figures 3.1 and 3.2 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 instrument throughput in the visible. The silver coatings cut off at wavelengths shortward of 3700 Å. The WFC has two filter wheels which it shares with the HRC, offering the possibility of internal WFC/HRC parallel observing for some filter combinations. The HRC/SBC optical chain comprises three aluminized mirrors, overcoated with MgF_2. The HRC or SBC channels are selected by means of a plane fold mirror. The HRC is selected by inserting the fold mirror into the optical chain so that the beam is imaged onto the HRC detector through the WFC/HRC filter wheels. The SBC channel is selected by moving the fold mirror out of the beam to yield a two mirror optical chain which images through the SBC filter wheel onto the SBC detector. The aberrated beam coronagraph is accessed by inserting a mechanism into the HRC optical chain. This mechanism positions a substrate with two occulting spots at the aberrated telescope focal plane and an apodizer at the re-imaged exit pupil.

While there is no mechanical reason why the coronagraph could not be used with the SBC, for health and safety reasons (due to brightness limits of the MAMA detector) **use of the coronagraph is forbidden with the SBC**.

1.1.3 ACS Geometric Distortion

The ACS detectors exhibit significantly more distortion than previous HST instruments. All ACS observations must be corrected for distortion before any photometry or astrometry is derived. For a thorough discussion of ACS Geometric Distortion, we refer the reader to Chapter 4.

The principal cause of the ACS distortion is that the optics have been designed with a minimum number of components, consistent with correcting for the spherical aberration induced by the Optical Telescope Assembly (OTA), without introducing coma. The result is a high throughput, but focal surfaces far from normal to the principal rays. The WFC detector is tilted at 22 degrees giving an elongation of 8% along the diagonal. The HRC and SBC detectors have a 25 degree tilt giving an elongation of 12%. In each case, the scale in arcseconds per pixel is smaller along the radial direction of the OTA field of view than along the tangential direction. When projected on the sky, this causes each detector to appear "rhombus-shaped" rather than square. The angle on the sky between the x and y axes is 84.9 degrees for the WFC1, 86.1 for the WFC2 and 84.2 degrees for the HRC.

The orientations of the ACS detector edges are approximately in line with the V2 and V3 coordinate axes of the telescope. Consequently, the eigenaxes of the scale transformation are along the diagonals for WFC and the apertures and pixels appear non-rectangular in the sky projection. For the HRC and SBC the situation is even more irregular because the aperture diagonals do not lie along a radius of the HST field of view. Figure 1.1 shows the ACS apertures in the telescope's V2/V3 reference frame and illustrates the "rhombus" shape of each detector. A telescope roll angle of zero degrees would correspond to an on-sky view with the V3 axis aligned with North and the V2 with East. The readout amplifiers are also indicated for each detector.

If these were the only distortions present, their impact on photometry and mosaicing/dithering could be simply computed. A more problematic effect is the variation of scale and pixel area across each detector. For the WFC this amounts to a change of ~10% in scale from corner to corner. For the HRC and SBC this variation is only about 1%, since these detectors cover much smaller fields of view. The area on the sky covered by a WFC pixel varies by ~18% from corner to corner, allowance for which *must* be made in photometry.

Dithering and mosaicing are complicated by the fact that an integral pixel shift near the center of the detector will translate into a non-integral displacement for pixels near the edges. This is not a fundamental limitation, but will imply some computational complexity in registering images and will depend on an accurate measurement of distortions.

Figure 1.1: ACS apertures compared with the V2/V3 reference frame. The read-out amplifiers (A,B,C,D) are indicated on the figure. The WFC data products from the calibration pipeline are oriented so that WFC1 (chip 1, which uses amps A,B) is on top. The HRC data products are also oriented such that amps A,B are on top, but they are inverted from the WFC images with respect to the sky.
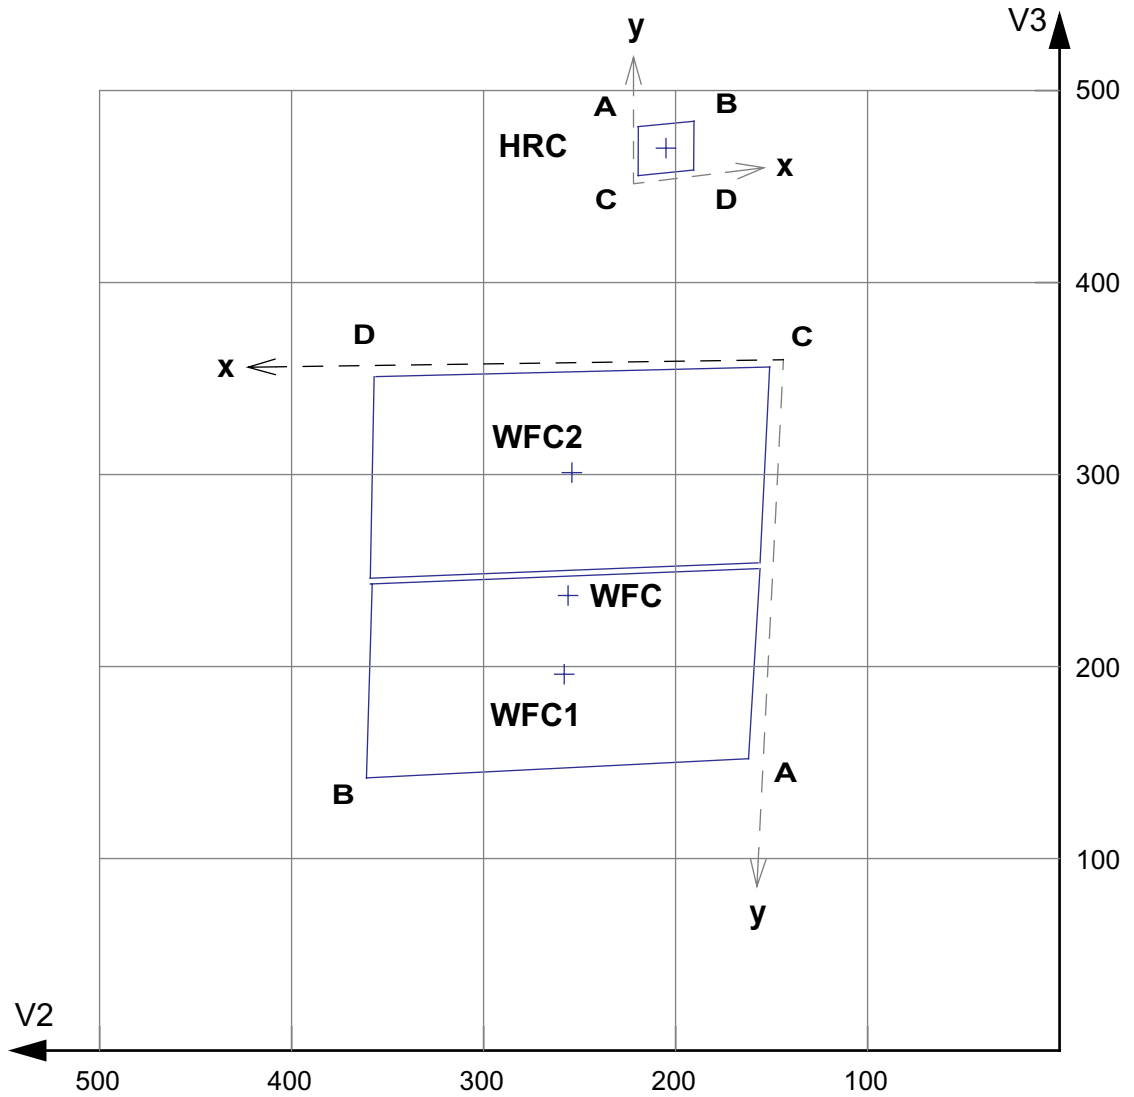
1.2 Basic Instrument Operations

1.2.1 Target Acquisitions

For the majority of ACS observations target acquisition is simply a matter of defining the appropriate aperture for the observation. Once the telescope acquires its guide stars, the target will be within ~1–2 arcseconds of the specified pointing. For observations with the ramp filters, one must specify the desired central wavelength for the observation. For the special case of coronagraphic observations, an onboard target acquisition will need to be specified. The nominal accuracy of the combined target acquisition and slew procedure is ~0.03 arcseconds, comparable to that achieved by STIS.

1.2.2 Typical ACS Observing Sequence

An important issue for observers to consider is the "packaging" of their observations, i.e. how observations are CR-SPLIT to mitigate the impact of cosmic rays, whether sub-stepping or "dithering" of images is required, and how, if necessary, to construct a mosaic pattern to map the target. For an online library of pointing patterns, refer to the ACS dither web page:

http://www.stsci.edu/hst/acs/proposing/dither.

HRC observations and narrowband observations with the WFC are more likely to be read-noise limited, requiring consideration of the optimum CR-SPLIT times. Observations with the MAMA detectors do not suffer from cosmic rays or read noise, but long integration times are often needed to obtain sufficient signal-to-noise in the photon-starved ultraviolet.

A typical ACS observing sequence consists of a series of CR-SPLIT and dithered ~10–20 minute exposures for each program filter. Coronagraphic observations require an initial target acquisition observation to permit centering of the target under the occulting mask. Observers generally need not take their own calibration exposures.

1.2.3 Data Storage and Transfer

At the conclusion of each exposure, the science data are read out from the detector and placed in ACS's internal buffer memory, where they are stored until it can be transferred to the HST solid state data recorder (and thereafter to the ground). The internal buffer memory is large enough to hold one WFC image, or sixteen HRC or SBC images, and so the buffer will typically need to be dumped during the following WFC exposure, assuming it is longer than 340 seconds. For shorter exposures an extra overhead of this length is imposed.

ACS's internal buffer stores the data in a 16 bit-per-pixel format. This structure imposes a maximum of 65,535 counts per pixel. For the MAMA detectors this maximum is equivalent to a limit on the total number of detected photons per pixel which can be accumulated in a single exposure. For the WFC and HRC, the full well (and not the 16 bit buffer format) limits the photons per pixel which can be accumulated without saturating in a single exposure when GAIN > 1 and GAIN > 2, respectively, are selected.

1.2.4 Parallel Operations

Parallel observations with the WFC and HRC are possible with ACS for certain filter combinations. ACS can be used in parallel with any of the other science instruments on HST within restrictions described in detail in the *ACS Instrument Handbook*. There are significant constraints on the use of the MAMA detectors in parallel. The policy for applying for parallel observing time is described in the *HST Call for Proposals*.

ACS Data Structure

In this chapter. . .

This chapter describes the ACS data products, including a discussion of file suffixes, association tables, and trailer files. It describes the FITS file structure and data storage requirements for ACS images. Finally, a description of the FITS header keywords is provided for both primary and extension headers.

2.1 Types of ACS Files

2.1.1 Data Files and Suffixes

The file suffixes given to ACS data products are described in Table 2.1 and closely mimic the suffixes used by STIS. The initial input files to the calibration pipeline are the raw (RAW) files from Generic Conversion and the association (ASN) table, if applicable, for the complete observation set.

For CCD images, a temporary file, with the suffix BLV_TMP, is created once bias levels are subtracted and the overscan regions are trimmed. This file is renamed with the FLT extension after the standard calibrations (flat fielding, dark subtraction, etc.) are complete. The FLT files will serve as input for cosmic ray rejection (if required). For CR-SPLIT exposures, a temporary 'cr-combined image' (CRJ_TMP) is created and then renamed with the CRJ extension once basic calibrations are complete. Single MAMA images are given the FLT suffix once calibrations are complete. By definition, these images do not have an overscan region and are not affected by cosmic rays. The calibrated products of a REPEAT-OBS association will be several individually calibrated FLT exposures and a summed flat-fielded (but not cosmic-ray cleaned) SFL image.

Table 2.1: ACS File Suffixes:

File Suffix	Description	Units
_RAW	Raw data	*DN*
_ASN	Association file for observation set	
_SPT	Telemetry and engineering data	
_TRL	Trailer File with processing comments	
_BLV_TMP	Overscan-trimmed individual exposure (renamed to _FLT)	*DN*
_CRJ_TMP	Uncalibrated, CR-rejected Combined image (renamed to _CRJ)	*DN*
_FLT	Calibrated, Flat fielded individual exposure	*electrons*
_CRJ	Calibrated, CR-rejected, Combined image	*electrons*
_SFL	Calibrated, Repeat-Obs, Combined image	*electrons*
_DRZ	Calibrated, Geometrically Corrected, Dither-Combined image	*electrons/sec*

2.1.2 Association Tables

Association tables are useful for keeping track of the complex set of relationships that can exist between exposures taken with ACS, especially with REPEAT-OBS, CR-SPLIT, and dithered exposures. Images taken at a given dither position may be additionally CR-SPLIT into multiple

exposures. In these cases, associations are built to describe how each exposure relates to the desired final product. As a result, ACS associations will create one or more science products from the input exposures, unlike NICMOS or STIS associations. The relationships defined in the association table determine how far through the calibration pipeline the exposures are processed and when the calibrated exposures are combined into sub-products for further calibration.

ACS data files are given the following definitions:

- An *exposure* is a single image, the "atomic unit" of HST data.

- A *dataset* is a collection of files having a common root name (first 9 characters).

- A *sub-product* is a dataset created by combining a subset of the exposures in an association.

- A *product* is a dataset created by combining sub-products of an association.

ACS association tables were designed to closely resemble the NICMOS association format, with three primary columns: `MEMNAME`, `MEMTYPE`, and `MEMPRSNT`. The column `MEMNAME` gives the name of each exposure making up the association and output product name(s). The column `MEMTYPE` specifies the role the file has in the association. A unique set of `MEMTYPES` specific to ACS were adopted to provide the support for multiple products. These types are summarized in Table 2.2.

Table 2.2: Exposure types in ACS associations. The suffix "n" is appended to the `MEMTYPE` when multiple sets are present within a single association.

MEMTYPE	Description
EXP-CRJ	Input CR-SPLIT exposure (single set)
EXP-CRn	Input CR-SPLIT exposure for CR-combined image n (multiple sets)
PROD-CRJ	CR-combined output product (single set)
PROD-CRn	CR-combined output product n (multiple sets)
EXP-RPT	Input REPEAT-OBS exposure (single set)
EXP-RPn	Input REPEAT-OBS exposure for repeated image n (multiple sets)
PROD-RPT	REPEAT-OBS combined output product (single set)
PROD-RPn	REPEAT-OBS combined output product n (multiple sets)
EXP-DTH	Input dither exposure
PROD-DTH	Dither-combined output product

A sample association table for a two-position dithered observation with `CR-SPLIT=2` is presented in Table 2.3. This example shows how both `MEMNAME` and `MEMTYPE` are used to associate input and output products. The `MEMTYPE` for each component of the first `CR-SPLIT` exposure, `JxxxxxECM` and `JxxxxxEGM`, are given the type `EXP-CR1`. The sub-product `Jxxxxx011` is designated in the table with a `MEMTYPE` of `PROD-CR1`. The last digit of the product filename corresponds to the output product number in the `MEMTYPE`. A designation of zero for the last digit in the filename is reserved for the dither-combined product.

The column `MEMPRSNT` indicates whether a given file already exists. For example, if cosmic ray rejection has not yet been performed by **CALACS**, the `PROD-CRn` files will have a `MEMPRSNT` value of "no". The sample association table in Table 2.3 shows the values of `MEMPRSNT` prior to **CALACS** processing.

Table 2.3: Sample Association Table `Jxxxxx010_ASN`

MEMNAME	MEMTYPE	MEMPRSNT
JxxxxxECM	EXP-CR1	yes
JxxxxxEGM	EXP-CR1	yes
Jxxxxx01<u>1</u>	PROD-CR1	no
JxxxxxEMM	EXP-CR2	yes
JxxxxxEOM	EXP-CR2	yes
Jxxxxx01<u>2</u>	PROD-CR2	no
Jxxxxx010	PROD-DTH	no

2.1.3 Trailer Files

Each task in the **CALACS** package creates messages during processing which describe the progress of the calibration and which are sent to STDOUT. In the pipeline processing for other HST instruments, trailer files were created by simply redirecting the STDOUT to a file. Because multiple output files can be processed in a single run of **CALACS**, creating trailer files presents a unique challenge. Each task within the **CALACS** package must decide which trailer file should be appended with comments and automatically open, populate, and close each trailer file.

CALACS will *always overwrite* information in 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. The string *CALACSBEG* will mark the first comment added to the trailer file. If a trailer file already exists, **CALACS** will search for this string to determine where to append

processing comments. If it is not found, the string will be written to the end of the file and all comments will follow. Thus any comments from previous processing are overwritten and only the most current calibrations are recorded.

As each image is processed, an accompanying trailer file with the '.trl' ending will be created. Further processing will concatenate all trailer files associated with the output product into a single file. Additional messages will then be appended to this concatenated file. Thus, some information is duplicated across multiple trailer files, but for any product processed within the pipeline, the trailer file is ensured to contain processing comments from 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.

2.2 ACS File Structure

The ACS calibration pipeline assembles data received from HST into datasets, applies the standard calibrations, and stores the uncalibrated datasets in the HST Data Archive. The structure of these data products is based on the STIS and NICMOS file format and consists of multi-extension FITS files which store science (SCI), data quality (DQ) and error (ERR) arrays, as shown in Figure 2.1.

ACS WFC data come from two CCD chips and are treated as separate observations, with a SCI, DQ and ERR array for each chip, and with both chips being stored in the same FITS file. The result is a FITS file with 6 data extensions plus a global header for a single WFC exposure. The WFC apertures are plotted with respect to the V2/V3 reference frame in Figure 2.2 and are oriented such that the x-axis is approximately to the right and the y-axis is approximately straight up. For pipeline data products, chip 2 is below chip 1 in y-pixel coordinates and was therefore defined as extension 1. Thus, the chip/extension notation is counterintuitive. To display the science image for chip 1, the user must specify the extension 'file.fits[sci,2]'. Similarly, the data quality and error array for chip 1 are specified as [dq,2] and [err,2].

A single HRC or SBC exposure comes from a single chip and has only 3 data extensions plus a global header. While the raw and calibrated WFC images contain 6 data extensions, the drizzled product will always contain 3 data extensions, no matter which detector was used. When WFC data is drizzled, both chips are included in a single FITS extension.

Figure 2.1 illustrates the basic format for storing ACS images. An uncalibrated (RAW) exposure contains a primary header plus, for each chip, a SCI extension (in 16-bit integer format), an empty ERR extension, and an empty DQ array. The calibrated product from **CALACS** contains a

primary header plus, for each chip, a SCI extension (in 32-bit float format), an ERR extension (32-bit float), and a DQ extension (16-bit integer).

PyDrizzle removes the effects of the geometric distortion and produces a drizzled (DRZ) product which has the multi-extension FITS file shown in Figure 2.1. The SCI extension contains the distortion-corrected data (32-bit float), the WHT extension contains the weight mask (32-bit float), and the CTX extension contains the context image (32-bit integer). For more information on these DRZ product extensions, refer to "Data Products" in Section 4.5.

Figure 2.1: Data format for calibrated and drizzled ACS modes. Note that for calibrated science data, WFC1 (chip 1) corresponds to extension [sci,2].
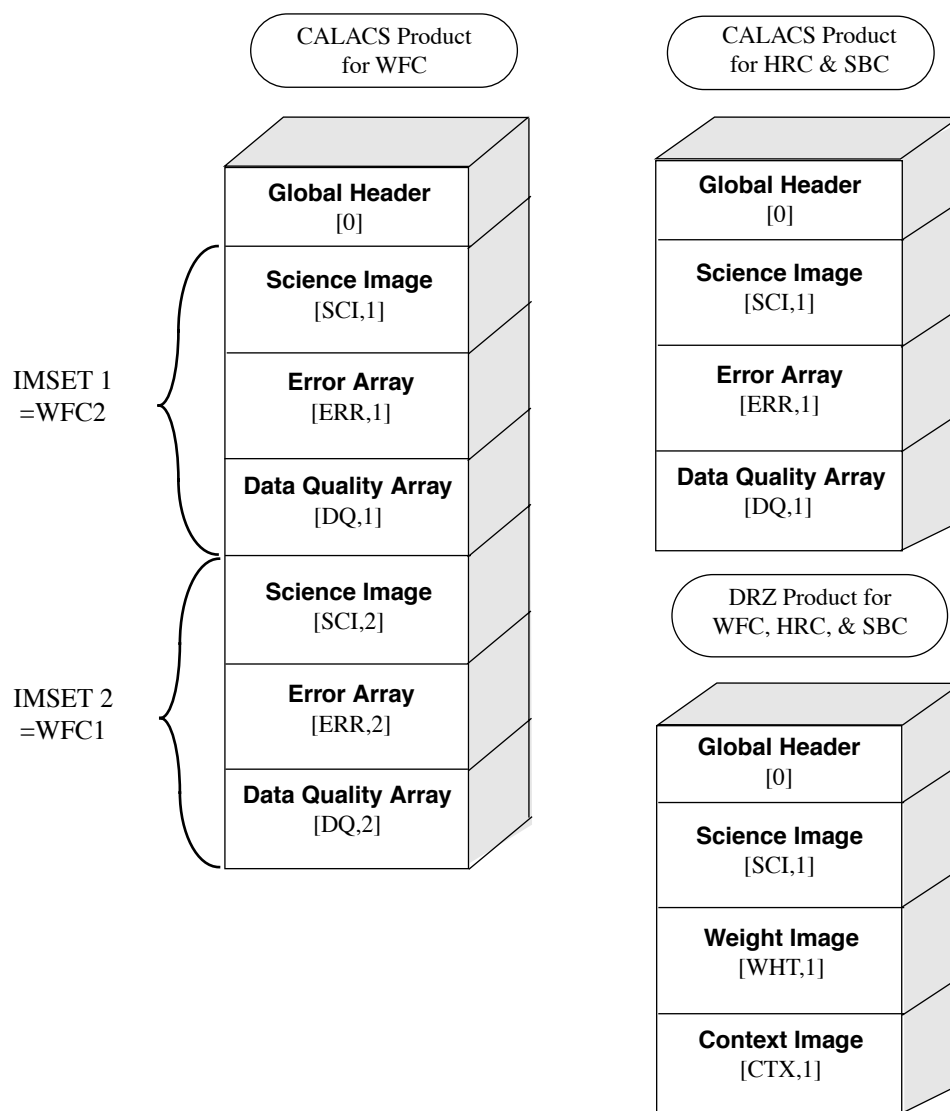
Figure 2.2: WFC apertures compared with the V2/V3 reference frame. The read-out amplifiers (A,B,C,D) are indicated on the figure. The WFC data products from the calibration pipeline are oriented so that WFC1 (chip 1, which uses amps A,B) is on top. The science image for chip 1 may be specified as [sci,2].
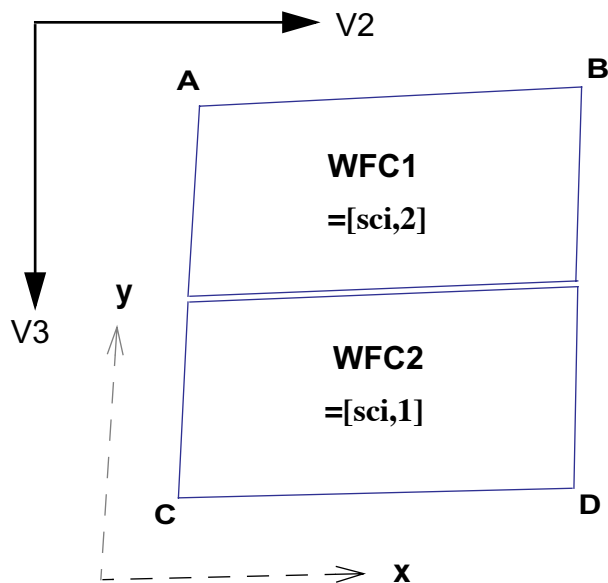
2.3 Data Storage Requirements

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. The WFC image includes two SCI arrays, while the HRC and SBC detectors have a single 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 than raw images because the overscan regions have been trimmed during processing.

PyDrizzle combines the science, data quality, and error arrays for each detector/array into a single science, weight, and context extension. Once the distortion is corrected, the size of the drizzled images (in pixels and therefore also in MB) will be larger than the calibrated images due to the fact that their projection on the sky is rhombus-shaped, rather than square. The specific dimensions of the drizzled image will depend on which

distortion model is currently in use in the pipeline and will vary slightly (~1-2 pixels) due to the effects of velocity aberration.

Table 2.4: Size (in MB) of raw, calibrated, and drizzled ACS images for each detector. The drizzled image sizes are distortion-corrected only and assume no dither offset or scale change.

Detector	Size of FITS file (S_{raw})	Size of FITS file (S_{cal})	Size of FITS file (S_{drz})
WFC (2 Chips)	68.7 MB	168 MB	215 MB
HRC (1 Chip)	4.4 MB	10.5 MB	16 MB
SBC (1 Chip)	4.2 MB	10.5 MB	16 MB

Table 2.5: Size (in pixels) for raw, calibrated, and drizzled ACS images for each detector. The drizzled image sizes are distortion-corrected only and assume no dither offset or scale change. The size of drizzled images will vary slightly due to velocity aberration effects.

Detector	X_{raw}	Y_{raw}	X_{cal}	Y_{cal}	X_{drz}	Y_{drz}
WFC (2 Chips)	4144	2068	4096	2048	4213	4240
HRC (1 Chip)	1062	1044	1024	1024	1164	1138
SBC (1 Chip)	1024	1024	1024	1024	1397	1341

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

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

where:

- D_{min} is the minimum free disk space required for processing,
- n is the number of exposures in each `CR-SPLIT` set or `REPEAT-OBS` set,
- 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

While the ACS image headers for the WFC, HRC, HRC-ACQ, and SBC observations have very similar formats, two main differences are apparent. The MAMA (SBC) headers require a slightly modified structure from the CCD (WFC, HRC) headers, and the HRC-ACQ headers require several unique keywords.

In Table 2.6 and Table 2.7, we describe the ACS header keywords, with a note indicating to which detector those keywords refer. The primary header keywords (FITS extension 0) are given in Table 2.6, and the extension header keywords (FITS extensions > 0) are given in Table 2.7. For more information about primary and extension FITS headers, refer to Section 2.1.1 of the HST Data Handbook Introduction.

Table 2.6: ACS Primary Header Keywords (FITS extension 0)

KEYWORD	DESCRIPTION
SIMPLE	data conform to FITS standard (T/F)
BITPIX	bits per data value
NAXIS	number of data axes
EXTEND	file may contain standard extensions (T/F)
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 log line number
PR_INV_L	last name of principal investigator
PR_INV_F	first name of principal investigator

EXPOSURE INFORMATION (All Detectors)	
SUNANGLE	angle between sun and V1 axis
MOONANGL	angle between moon and V1 axis
SUN_ALT	altitude of the sun above Earth's limb
FGSLOCK	commanded FGS lock (fine, coarse, gyros, unknown)
DATE-OBS	UT date of start of observation (yyyy-mm-dd)
TIME-OBS	UT time of start of observation (hh:mm:ss)
EXPSTART	exposure start time (modified Julian date)
EXPEND	exposure end time (modified Julian date)
EXPTIME	exposure duration (seconds)--calculated
EXPFLAG	exposure interruption indicator
POINTING INFORMATION (All Detectors)	
PA_V3	position angle of V3-axis of HST (deg)
TARGET OFFSETS (POSTARGS) (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
CAL_VER	CALACS code version
PROCTIME	pipeline processing time (MJD)
SCIENCE INSTRUMENT CONFIGURATION (All Detectors)	
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
FILTER2	element selected from filter wheel 2
FWOFFSET	computed filter wheel offset (0, +1, -1)
FWERROR	filter wheel position error flag (T, F)
LRFWAVE	proposed linear ramp filter wavelength
APERTURE	aperture name
PROPAPER	proposed aperture name
CRSPLIT	number of cosmic ray split exposures
DIRIMAGE	direct image for grism or prism exposure
MAMA OFFSETS (SBC Only)	
MOFFSET1	axis 1 MAMA offset
MOFFSET2	axis 2 MAMA offset

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)**
STATFLAG	calculate statistics? (T/F)
WRTERR	write out error array extension? (T/F)
DQICORR	data quality initialization?
ATODCORR	correct for A to D conversion errors?
BLEVCORR	subtract bias level computed from overscan image?
BIASCORR	subtract bias image?
FLSHCORR	subtract post flash image?
CRCORR	combine observations to reject cosmic-rays?
EXPSCORR	process individual observations after cr-reject?
SHADCORR	apply shutter shading correction?
GLINCORR	correct for global detector non-linearities? (SBC only)
LFLGCORR	flag pixels for local and global nonlinearities? (SBC only)
DARKCORR	subtract dark image?
FLATCORR	apply flat field correction?
PHOTCORR	populate photometric header keywords?
RPTCORR	add individual repeat observations?
DRIZCORR	process dithered images?
CALIBRATION REFERENCE FILES	**(All Detectors)**
BPIXTAB	bad pixel table
CCDTAB	CCD calibration parameter table
ATODTAB	analog-to-digital correction file
OSCNTAB	CCD overscan table
BIASFILE	bias image file name
FLSHFILE	post flash image file name
CRREJTAB	cosmic-ray rejection parameter table
SHADFILE	shutter shading correction file
MLINTAB	MAMA linearity correction table (SBC only)
DARKFILE	dark image file name
PFLTFILE	pixel-to-pixel flat field file name
DFLTFILE	delta-flat field file name
LFLTFILE	low order flat
PHOTTAB	photometric throughput table (not used)
GRAPHTAB	HST graph table
COMPTAB	HST components table
IDCTAB	image distortion correction table

COSMIC RAY REJECTION ALGORITHM PARAMETERS (WFC, HRC, HRC ACQ)	
MEANEXP	reference exposure time
SCALENSE	multiplicative scale factor applied to noise
INITGUES	initial guess method (min, med)
SKYSUB	sky value subtracted (mode, none)
SKYSUM	sky level from the sum of all constituent images
CRSIGMAS	statistical rejection criteria
CRRADIUS	rejection propagation radius (pixels)
CRTHRESH	rejection propagation threshold
BADINPDQ	data quality flag used for rejection
REJ_RATE	rate at which pixels are affected by cosmic rays
CRMASK	flag CR-rejected pixels in input files? (T/F)
PATTERN KEYWORDS (WFC, HRC, SBC)	
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 patt (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 read out configuration
CCDGAIN	commanded gain of CCD (electrons/DN)
CCDOFSTA	commanded CCD bias offset for amplifier A (electrons)
CCDOFSTB	commanded CCD bias offset for amplifier B (electrons)
CCDOFSTC	commanded CCD bias offset for amplifier C (electrons)
CCDOFSTD	commanded CCD bias offset for amplifier D (electrons)

CALIBRATED ENGINEERING PARAMETERS (WFC, HRC, HRC ACQ)	
ATODGNA	calibrated gain for CCD amplifier A (electrons/DN)
ATODGNB	calibrated gain for CCD amplifier B (electrons/DN)
ATODGNC	calibrated gain for CCD amplifier C (electrons/DN)
ATODGND	calibrated gain for CCD 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 only)	
ACQTYPE	type of acquisition
ACQNAME	rootname of acquisition exposure
TARGET ACQUISITION PARAMETERS (HRC ACQ only)	
ACQTYPE	type of acquisition
CENTMETH	target acquisition centering method
CRELIM	perform cosmic ray rejection in acquisition
CCDBIASS	CCD bias subtracted from target acq image? (y/n)
BIASLEV	CCD bias level used to process acquisition exp.
CHECKBOX	size of checkbox for finding algorithms
TARGAREA	area of target (detector pixels)
ASSOCIATION KEYWORDS (All Detectors)	
ASN_ID	unique identifier assigned to association
ASN_TAB	name of association table
ASN_MTYP	role of exposure in association

Table 2.7: ACS Extension Header Keywords (FITS extensions > 0)

KEYWORD	DESCRIPTION
XTENSION	extension type
BITPIX	bits per data value
NAXIS	number of data axes
NAXIS1	length of first data axis
NAXIS2	length of second data axis
PCOUNT	number of group parameters
GCOUNT	number of groups
INHERIT	inherit the primary header
EXTNAME	extension name
EXTVER	extension version number
ROOTNAME	rootname of the observation set
EXPNAME	9 character exposure identifier
DATAMIN	minimum data value (electrons)
DATAMAX	maximum data value (electrons)
BUNIT	brightness units (counts, electrons, electrons/sec)
BSCALE	scale factor for array value to physical value
BZERO	physical value for an array value of zero
WORLD COORDINATE SYSTEM AND RELATED PARAMETERS (All Detectors)	
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 w.r.t. x
CD1_2	partial of first axis coordinate w.r.t. y
CD2_1	partial of second axis coordinate w.r.t. x
CD2_2	partial of second axis coordinate w.r.t. y
LTV1	offset in X to subsection start
LTV2	offset in Y to subsection start
LTM1_1	reciprocal of sampling rate in X
LTM2_2	reciprocal of sampling rate in Y
RA_APER	RA of aperture reference position
DEC_APER	Dec of aperture reference position
PA_APER	position angle of aperture reference position (deg)
ORIENTAT	position angle of image y axis (degrees E of N)
VAFACTOR	velocity aberration plate scale factor

WFC CCD CHIP IDENTIFICATION (WFC only)	
CCDCHIP	CCD chip (1 or 2)

READOUT DEFINITION PARAMETERS (All Detectors)	
CENTERA1	subarray axis1 center point in unbinned detector pixels
CENTERA2	subarray axis2 center point in unbinned detector pixels
SIZAXIS1	subarray axis1 size in unbinned detector pixels
SIZAXIS2	subarray axis2 size in unbinned detector pixels
BINAXIS1	axis1 data bin size in unbinned detector pixels
BINAXIS2	axis2 data bin size in unbinned detector pixels

PHOTOMETRY KEYWORDS (All Detectors)	
PHOTMODE	observation configuration for photometric calibration
PHOTFLAM	inverse sensitivity (erg s^{-1} cm^{-2} A^{-1} for 1 electron/sec)
PHOTZPT	ST magnitude zero point
PHOTPLAM	pivot wavelength
PHOTBW	rms bandwidth of filter plus detector

REPEATED EXPOSURES INFO (WFC, HRC, HRC ACQ)	
NCOMBINE	number of image sets combined during CR rejection

DATA PACKET INFORMATION (All Detectors)	
FILLCNT	number of segments containing fill
ERRCNT	number of segments containing errors
PODPSFF	podps fill present? (T/F)
STDCFFF	ST DDF fill present? (T/F)
STDCFFP	ST DDF fill pattern (hex)

ON-BOARD COMPRESSION INFORMATION (WFC Only)	
WFCMPRSD	WFC data compressed? (T/F)
COMPTYP	compression type performed? (partial/full)
CBLKSIZ	size of compression block in bytes
LOSTPIX	number of pixels lost due to buffer overflow

TARGET ACQUISITION PARAMETERS (HRC ACQ Only)	
TARGA1	fsw located subarray axis1 coord. of target
TARGA2	fsw located subarray axis2 coord. of target
APERA1	fsw located subarray axis 1 coord. of aperture
APERA2	fsw located subarray axis 2 coord. of aperture
APERLKA1	axis 1 detector pixel of acq aperture center
APERLKA2	axis 2 detector pixel of acq aperture center
MAXCHCNT	counts in the brightest checkbox
BOPOFFA1	axis 1 offset (arcs) object moved off aperture
BOPOFFA2	axis 2 offset (arcs) object moved off aperture

ENGINEERING PARAMETERS (SBC Only)	
GLOBRATE	global count rate
GLOBLIM	was global linearity level exceeded?
IMAGE STATISTICS AND DATA QUALITY FLAGS (All Detectors)	
NGOODPIX	number of good pixels
SDQFLAGS	serious data quality flags considered 'bad' by CALACS
GOODMIN	minimum value of good pixels (electrons)
GOODMAX	maximum value of good pixels (electrons)
GOODMEAN	mean value of good pixels (electrons)
SNRMIN	minimum signal to noise of good pixels
SNRMAX	maximum signal to noise of good pixels
SNRMEAN	mean signal to noise of good pixels
SOFTERRS	number of soft error pixels (DQF=1)
MEANBLEV	average of the subtracted bias levels (electrons)
MEANDARK	average of the subtracted dark values (electrons)
MEANFLSH	average of the subtracted post flash values (electrons)

CHAPTER 3:

ACS Calibration Pipeline

In this chapter. . .

This chapter describes the ACS calibration pipeline and the On-the-Fly-Reprocessing, which is standard for all HST data requested from the STScI archive. The pipeline software consists of two separate packages: **CALACS** and **PyDrizzle**. **CALACS** removes various instrumental signatures, combines multiple exposures, and updates header keyword values. **PyDrizzle** corrects images for geometric distortion and combines observations which were taken as part of a dither pattern.

The data products from OTFR are described, as well as considerations and requirements for manual recalibration. The remaining bulk of the chapter is devoted to a detailed description of **CALACS**. While **PyDrizzle** is also used during pipeline processing, a detailed discussion of this software is deferred to Chapter 4.

3.1 On The Fly Reprocessing (OTFR)

OTFR is now the standard way of processing data which has been requested by the user from the STScI archive. It provides the best calibrated products by reprocessing the raw telemetry files "on-the-fly" for distribution each time data is requested. Each observation requires calibrations suited to the particular mode used for taking data. Calibration data specific to the various modes of ACS operation are prepared and archived in the Calibration Data Base System (CDBS).

The STScI calibration pipeline consists of two main software systems: the Operations Pipeline Unified System (OPUS) and the Data Archive and Distribution System (DADS). Raw spacecraft telemetry data from Goddard Spaceflight Center are transmitted to STScI in the form of POD files. When a user requests data from the HST Archive, OTFR uses the POD files as input to the OPUS step called Generic Conversion and generates the uncalibrated "raw" data. At the same time, Generic Conversion queries the CDBS to determine the most up-to-date reference files for the observation configuration. OPUS then runs **CALACS** to process the uncalibrated data, using specific ACS reference images and tables from the CDBS. With these data, DADS populates a database which is accessible to users via StarView. DADS then distributes any data (both calibrated and uncalibrated) requested for download to the user. (See Chapter 1 of the HST Introduction for more information.)

The most current ACS reference files are used by **CALACS** each time OTFR is run. The calibration reference files (e.g. flat fields, bad pixel tables) are also available from the HST Archive. Since reference files such as CCD biases and darks are frequently updated, OTFR may use different reference files depending on the date of reprocessing. Previously, the archived data from DADS had to be manually recalibrated by running the calibration software on the user's home workstation with the updated reference files. OTFR replaces this step by automatically recalibrating with the most current reference data available. The user simply waits until the contemporaneous reference files are in place and requests the data via StarView or the HST web-based archive request form. The image header keywords are then updated with the names of the reference files used during that OTFR run. The PROCTIME keyword is updated to reflect the pipeline processing date and time in MJD.

OTFR allows the user to avoid downloading archived data which is outdated because of software changes made for bug fixes, improved algorithms, new capabilities, or header keyword changes. An archive request for calibrated data will result in OTFR reprocessing and distributing data which uses the latest software versions available.

Currently OTFR will distribute *all* files associated with an ACS observation, including both the uncalibrated and calibrated files. Future

versions could enable users to select only certain parts of the dataset, for example the final calibrated image. Future versions could also enable users to set calibration parameters for a particular OTFR run. Until then, OTFR will process with default values. Of course, the option of recalibrating aside from OTFR always exists. (See "Manual Recalibration of ACS Data" in Section 3.5.)

3.2 Pipeline Overview

Pipeline processing is carried out by 2 separate packages: **CALACS**, which corrects for instrumental effects and produces calibrated products, and **PyDrizzle**, which corrects for distortion and combines associated dithered data.

3.2.1 CALACS: Image Calibration

The **CALACS** software consists of a series of individual tasks that:

- Orchestrate the flow of data through the pipeline.

- Perform the initial tasks of basic two-dimensional image reduction (e.g., overscan subtraction, bias subtraction) for CCD data.

- Reject cosmic rays from CR-SPLIT CCD data.

- Perform the remaining tasks of basic two-dimensional image reduction (e.g., dark subtraction, flat fielding) for CCD and MAMA data.

- Sum any REPEAT-OBS exposures.

The calibrated products from the pipeline may still contain artifacts such as hot pixels and cosmic-rays. If hot pixels are to be detected and flagged by **CALACS**, dithered exposures are required. Similarly, cosmic ray rejection by the pipeline requires the specification of CR-SPLIT exposures in the observing program. In the future, improved pipeline software will allow all associated data (including REPEAT-OBS exposures and dithered exposures) to be automatically corrected for cosmic-rays.

While intermediate steps in **CALACS** make use of sky subtraction, for example for identifying cosmic-rays, all data products created by the pipeline will *not* be sky subtracted. Fully calibrated data products (with suffixes FLT, CRJ, SFL) will be in units of *electrons*.

CALACS and Single exposures

A single raw ACS exposure is processed by **CALACS** following the steps outlined in Section 3.3. These steps include the standard detector calibrations: bias subtraction, dark subtraction, flat fielding, etc. When

CALACS processes a single raw exposure, the calibrated product is given the FLT suffix. The data in the calibrated SCI and ERR extensions are in units of *electrons*.

CALACS and Multiple exposures

The same processing steps performed on single images are performed for multiple images which are part of an association. The result is a single calibrated product in units of *electrons*.

CALACS will recognize and correctly process `CR-SPLIT` or `REPEAT-OBS` exposures by interpreting the association table and determining which exposures should be combined during calibration. `REPEAT-OBS` exposures are individually run through all the calibration steps to create calibrated, flat-fielded (FLT) product for each input file. The FLT images are then summed to create the summed flat-fielded (SFL) product. No cosmic-ray rejection is performed for `REPEAT-OBS` images.

Observations taken as part of a CR-SPLIT association, on the other hand, are combined during cosmic-ray rejection. First, the bias and dark subtraction is performed. Then, the images are combined while at the same time rejecting cosmic rays. The product is then flat fielded to create a single calibrated, cosmic-ray rejected, flat-fielded (CRJ) image.

CALACS and Dithered exposures

Observations which use the dither patterns provided in the proposal instructions will be automatically associated for combining in the pipeline. Any pattern designed with POS-TARGs will not be associated. When processing observations which are part of a dither pattern, **CALACS** will produce a calibrated FLT, CRJ, or SFL file at each dither position. It will not, however, combine images from multiple positions. Further pipeline processing by **PyDrizzle** will correct for geometric distortion and combine dithered images.

Table 3.1: The input and output image suffixes from **CALACS** and **PyDrizzle** are given for various observing modes.

Image Type	CALACS Input	CALACS Output	PyDrizzle Input	PyDrizzle Output	Cosmic Ray Rejected?
Single	RAW	FLT	FLT	DRZ	No
RPT-OBS	ASN+RAW	FLT+SFL	ASN+SFL	DRZ	No
CR-SPLIT	ASN+RAW	FLT+CRJ	ASN+CRJ	DRZ	Yes
DITH-PATTERN	ASN+RAW	FLT	ASN+FLT	DRZ	No

3.2.2 PyDrizzle: Distortion Correction and Dither Combination

All ACS data will be automatically corrected for distortion during pipeline processing using a task called **PyDrizzle**. This task relies on the IDCTAB reference file for a description of the distortion model. **PyDrizzle** understands the ACS association tables which allow the pipeline to combine dithered observations. Only observations which use the dither patterns provided in the proposal instructions will be automatically associated for combining in the pipeline. Programs which rely on explicit POS-TARG commands will NOT be associated, resulting in separately calibrated images for each position. **PyDrizzle** will automatically produce images which are both astrometrically and photometrically accurate regardless of whether they were taken as a single exposure or a set of dithered exposures.

The flat-fielded, calibrated products from **CALACS** are listed in Table 3.1. The calibrated product of a single ACS exposure is an FLT file. The product of a CR-SPLIT or REPEAT-OBS association is a CRJ or SFL file, respectively. The product of a dither pattern is an FLT, CRJ or SFL file at *each* dither position. Using these files as input, **PyDrizzle** performs the geometric correction on *all* data (dithered or not) and combines multiple dithered images into a single output image with the DRZ suffix. For WFC observations, both chips are combined into a single extension. **PyDrizzle** then converts the data to units of count rate (*electrons/sec*).

It is important to note that **PyDrizzle**, when combining dithered images, does *not* remove cosmic-rays. This can be done offline with a new task called **MultiDrizzle** using the individual FLT files as input. Subsequent reprocessing with **PyDrizzle** or **MultiDrizzle** offline may be required to obtain the desired scientific result. For information on obtaining this software, we refer the reader to Section 3.5.1.

3.2.3 When is OTFR *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 frequently 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 institution. There are several occasions when OTFR is not ideal and when off-line interactive processing by the user is required:

- when running **CALACS** with personal versions of reference files,

- when running **CALACS** with non-default calibration switch values,

- when images must be cleaned of artifacts such as new hot pixels or cosmic-rays.

Many ACS datasets do not contain multiple exposures at the same pointing and hence will not be cleaned of cosmic rays by OTFR

processing. In addition, the increasing prevalence of hot pixels, particularly in the WFC mode, makes dithered observations highly desirable. In these cases the pipeline will provide products which are geometrically corrected but which still contain cosmic rays.

To address this problem, the **MultiDrizzle** script has been developed to be run interactively outside the pipeline. **MultiDrizzle** provides a single-step interface to the complex suite of tasks in the STSDAS **dither** package. It makes use of the **PyDrizzle** software and can only be executed in the PyRAF environment. For more information, we refer the reader to Chapter 4.

3.3 Structure of CALACS

The **CALACS** package consists of four individual tasks listed in Table 3.2. These tasks are called automatically by **CALACS**, but each may be run separately should some variation in the normal processing be desired by the user.

Table 3.2: Tasks in the **CALACS** pipeline

ACSCCD	CCD specific calibrations
ACSREJ	Cosmic-ray rejection task
ACS2D	Basic MAMA and CCD calibrations
ACSSUM	Repeat-obs summing task

The flow of data through the ACS calibration pipeline and the decisions made while working with associated data are diagrammed in Figure 3.1. These calibration steps are also outlined below:

- Calculate a noise model for each pixel and record in the error array
- Flag known bad pixels and saturated pixels in the data quality array
- Subtract bias-level determined from overscan regions (CCD only)
- Subtract bias image (CCD only)
- Subtract post-flash image, if required (CCD only)
- Perform cosmic-ray rejection and combining of CR–SPLIT data (CCD only)
- Perform global linearity corrections (MAMA only)
- Scale and subtract dark image and calculate mean dark value
- Divide by flat field and multiply by gain
- Calculate photometry header keywords for flux conversion
- Calculate image statistics

Additional calibration steps are performed by **PyDrizzle**, including:

- Correct the geometric distortion of non-dithered images

- Convert the data units from counts to count rates

- Combine calibrated dithered observations into a single product

As indicated in Figure 3.1, the calibration tasks have been split to handle CCD-specific calibrations (**ACSCCD**) separately from those steps which can be applied to any ACS data (**ACS2D**). For example, the MAMA data obtained with the SBC do not have the overscan regions found in CCD data.

The initial processing performed on CCD data alone is shown in Figure 3.2. The 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 **ACSCCD** is then sent through **ACS2D** as shown in Figure 3.3. Raw MAMA data are processed from the start with **ACS2D**, which initializes the error and DQ arrays (already performed for CCD data) and which performs linearity corrections.

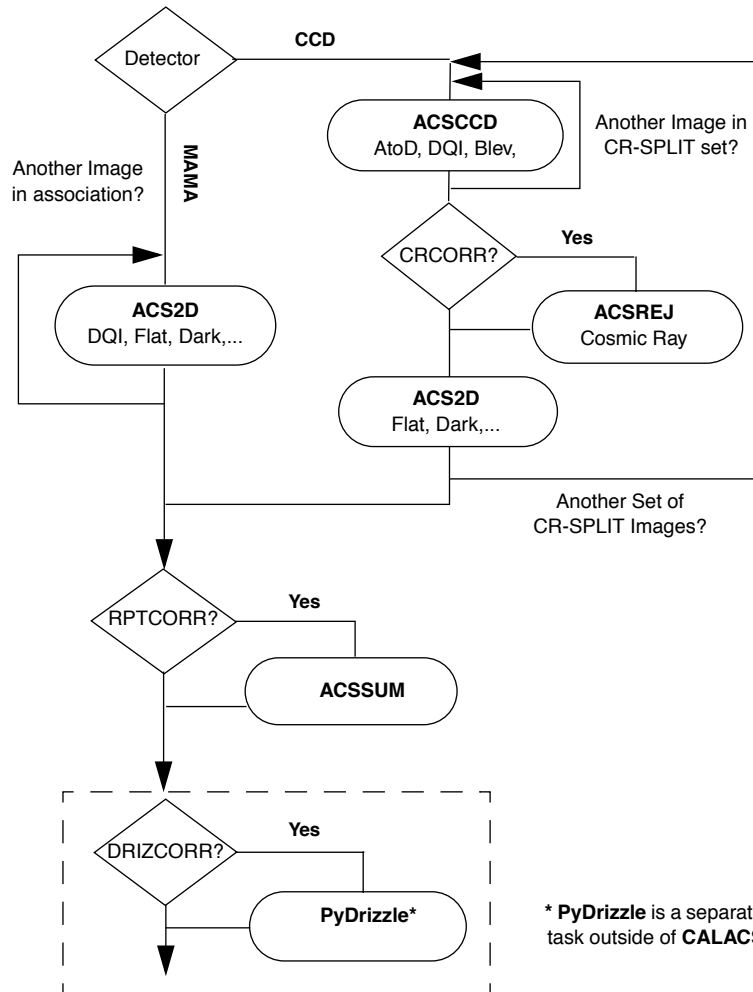Figure 3.1: Flow Diagram for ACS data with **CALACS** task names in bold.

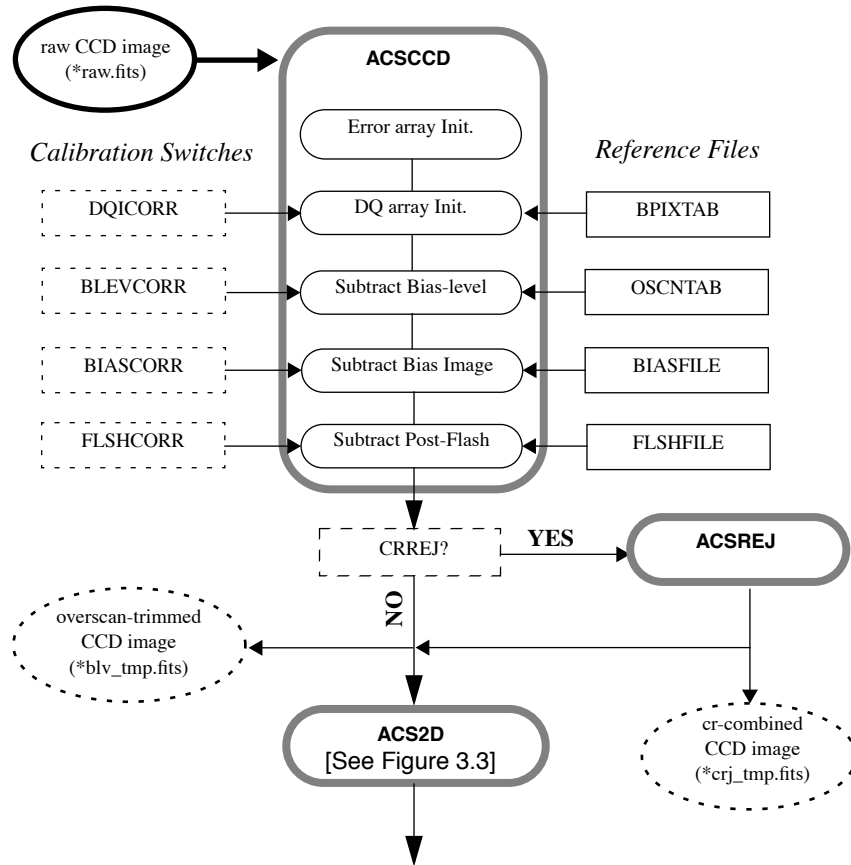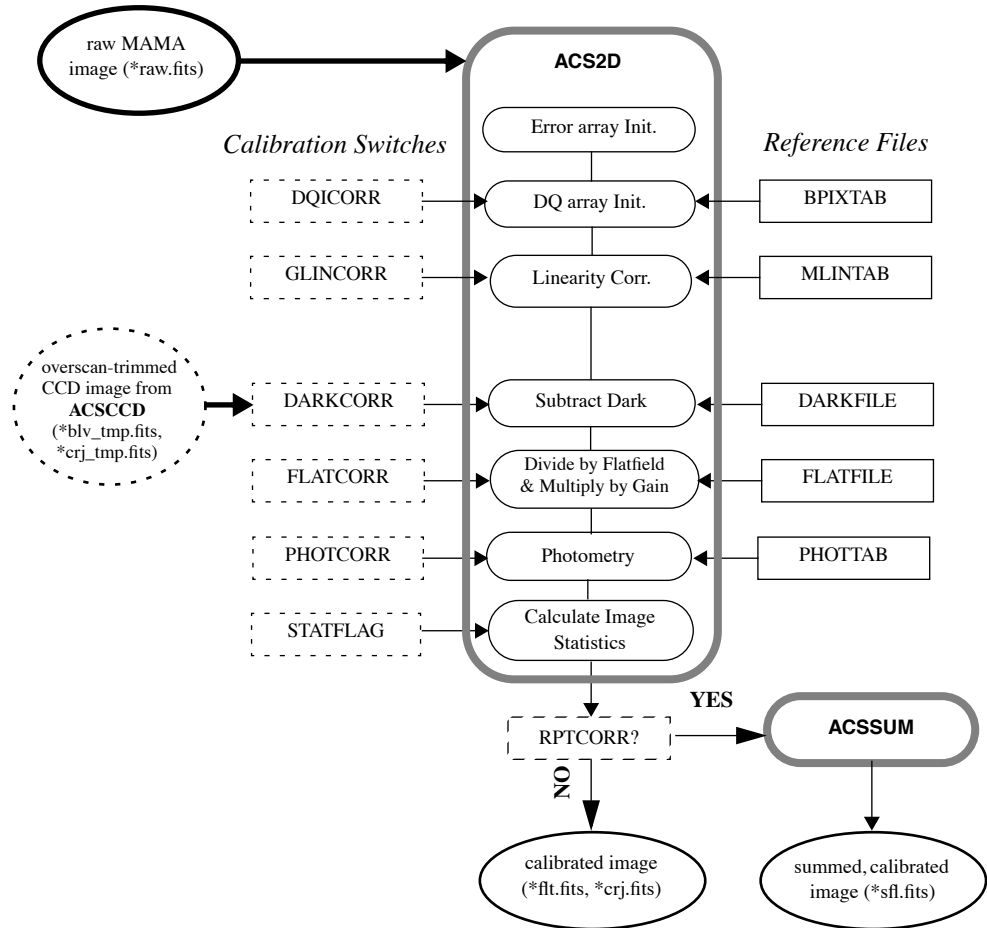Figure 3.2: Flow diagram for CCD data using **ACSCCD** in **CALACS**.

Figure 3.3: Flow diagram for MAMA and overscan-trimmed CCD data using **ACS2D** in **CALACS**.

3.4 CALACS Processing Steps

The **CALACS** pipeline consists of four individual calibration tasks: **ACSCCD**, **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 actually calibrating the data. The individual tasks apply the desired calibration steps to the data and create the output products, including the trailer files.

In the following four sections, we describe each **CALACS** task, give a detailed description of the calibrations steps performed within each task, and give a brief description of the reference files used for each step.

3.4.1 ACSCCD

This routine contains the initial processing tasks for all ACS CCD data. These tasks are listed in operational order in Table 3.3. The bulk of this routine was based on the code for **CALSTIS1**. **ACSCCD** processes each image in the input list one at a time, using the header keywords to determine which calibrations are performed. Upon completion of **ACSCCD**, the overscan regions will be trimmed from the image and an output image with the file extension '_blv_tmp.fits' is created.

Table 3.3: **ACSCCD** processing tasks

doNoise	Initialize error array
doDQI	Initialize data quality array
doAtoD	Perform A to D correction (currently skipped)
doBlev	Subtract bias level from overscan region
doBias	Subtract bias image
doFlash	Subtract post-flash image (if required)
Final Output	Output overscan-trimmed or full image

doNoise - Error Array Initialization
- Header Switch: None
- Header Keywords Updated: None
- Reference File: CCDTAB (*_ccd.fits)

The first processing step is initializing the image error array. The **doNoise** function examines the ERR extension of the input data to determine the state of the array. Input '_raw.fits' images delivered by Generic Conversion will contain a NULL ERR array, defined by the keywords NPIX1, NPIX2 and PIXVALUE, where PIXVALUE=0. If the ERR array has already been expanded and contains values other than zero, then this function does nothing. Otherwise, **CALACS** will initialize the array and **doNoise** will assign a simple noise model.

The noise model reads the science array and for each pixel calculates the error value σ (in DN for the CCD, in counts for the MAMA):

$$\sigma_{\text{CCD}} = \sqrt{(DN - bias)/(gain) + (readnoise/gain)^2}$$

$$\sigma_{\text{MAMA}} = max(1, \sqrt{counts})$$

Because MAMA data is not processed with **ACSCCD**, the **doNoise** function is later called as part of **ACS2D**.

The CCDTAB reference file, the CCD Characteristics Table, is used to initialize the ERR array by determining the bias, gain, and readnoise for each observation. The table contains one row for each amp configuration which can be used during readout. Each configuration is uniquely identified by the list of amps possible (CCDAMP), the particular chip being read out (CCDCHIP), the commanded gain (CCDGAIN), the commanded bias level

(CCDBIAS), and the pixel bin size (BINAXIS). CCDTAB uses these commanded values to determine each amp's physical read-out characteristics, including readnoise (READNSE), A-to-D gain (ATODGN), and bias level (CCDOFST). Note that by using the bias level from the CCDTAB, instead of from a full bias image, to create the ERR array, slight errors may be introduced which could be relevant for faint sources.

doDQI - Bad Pixel Determination

- Header Switch: DQICORR
- Header Keywords Updated: None
- Reference File: BPIXTAB (*_bpx.fits), CCDTAB (*_ccd.fits)

The function **doDQI** initializes the data quality (DQ) array by combining it with a table of known bad pixels for the detector, stored in the Bad-Pixel reference table (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). For CCD data, the values in the SCI extension are also checked for saturation by comparison with the value of the SATURATE column in the CCD parameters table CCDTAB.

The **doDQI** function combines the DQ flags from pre-processing, from the BPIXTAB, and from saturation tests into a single result for the particular observation. These values are combined using a "bitwise/logical OR operator" for each pixel. Thus, if a single pixel is affected by two DQ flags, those flag values will be added in the final DQ array. This array then serves as a mask so that **CALACS** can ignore bad pixels during processing.

The BPIXTAB reference file maintains a record of the (x,y) position and DQ value for all known bad pixels in each CCD. These pixels may change as some hot pixels are annealed and others appear over time, for example. Permanently bad pixels due to chip defects may be flagged during Generic Conversion, and it is the job of the BPIXTAB to maintain the list of bad pixels applicable for a given time period.

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 beyond aperture
8	masked by aperture feature
16	new hot pixel
32	"blob" at end of hot column
64	pre-existing hot pixel
128	bias level pixel
256	saturation (full well or a-to-d)
512	bad pixel in reference file
1024	weak trap
2048	a-to-d saturation*
4096	*reserved*
8192	cosmic ray rejected during image combination

*Note: The ability to flag pixels affected by a-to-d situation as 2048 will be implemented in the pipeline in Spring 2004.

doAtoD - A-to-D Correction (skipped)

- Header Switch: ATODCORR
- Header Keywords Updated: None
- Reference File: ATODTAB (*_a2d.fits)

The **doAtoD** function converts pixel values in the raw input image from units of DN (16-bit integer) to counts (32-bit float). An analog-to-digital correction would be applied if the CCD electronic circuitry, which performs the analog-to-digital conversion, were biased toward the assignment of certain DN (data number) values. Ground test results show that this correction is not currently needed, so the ATODCORR switch will always be set to OMIT.

doBlev - Bias-level Correction

- Header Switch: BLEVCORR
- Header Keywords Updated: BIASLEV[A,B,C,D], MEANBLEV, CRPIX[1,2], LTV[1,2]
- Reference File: OSCNTAB (*_osc.fits)

The **doBlev** function fits the bias level from the overscan region and subtracts it from the image data. The definition of the overscan region is

taken from the OSCNTAB reference file. With these regions defined, the overscan level for each row of the input image is measured and a straight line is fit as a function of image line number. The fit will then be evaluated for each row of the input image, and that value will be subtracted from each pixel of that row. If the measured overscan level for a given row is greater than 3 sigma times the mean of all rows, that row will be excluded from the linear fit. If the overscan level cannot be determined, or if the overscan region is not present in the input image, a default value (the CCDBIAS from the CCD parameters table) will be subtracted instead. In this case, an error message will be written to the image trailer file.

In **CALSTIS**, the overscan regions were removed from the image just after the bias-level subtraction, since no subsequent processing was required. **CALACS**, on the other hand, retains the output image in memory until the completion of all the processing steps in **ACSCCD**. The overscan regions will not be trimmed until the image is written to disk at the completion of **ACSCCD**.

The positional keywords CRPIX* and LTV* will be updated to reflect the offset due to removing the overscan. In addition, the mean value of all overscan levels will be computed and written to the SCI extension header as MEANBLEV. The individual bias level measured for each amplifier will be written to the SCI header as BIASLEV*.

The OSCNTAB reference file (Overscan Region Table) 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 amp and chip used. The columns TRIMX* 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, 4096x2048 for a full WFC image and 1024x1024 for a full HRC image. The OSCNTAB columns BIASSECTA* 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 virtual overscan region is also defined in the OSCNTAB, where (VX1, VY1) indicates the beginning of the upper or lower overscan region, extending to the position (VX2, VY2).

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.

doBias - Bias Image Subtraction
- Header Switch: BIASCORR
- Header Keywords Updated: None
- Reference File: BIASFILE (*_bia.fits)

Subtraction of the bias image is performed prior to cosmic-ray rejection using the function **doBias**. The appropriate reference image, BIASFILE, is selected using the DETECTOR, CCDAMP, and CCDGAIN keywords. The dimensions of the science image are used to distinguish observations which use the entire chip from sub-array images. Because the bias image has already been scaled by the gain and had the overscan values subtracted, the bias image may have a mean pixel value which is less than one.

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 further in the section on Dark Image Subtraction.

The bias reference files are updated frequently (~weekly). The BIASFILE has the same dimensions as a full-size science image, 1062x1044 for HRC and 4144x2068 for WFC, allowing for simple image arithmetic on arrays of equal size. Only after the completion of **ACSCCD** are the images trimmed to 1024x1024 (HRC) and 4096x2048 (WFC). **CALACS** assumes that CCD data will not be read-out in binned mode, so any image which is not the full size is assumed to be a sub-array.

For sub-array images, **CALACS** uses the LTV[1,2] keywords to extract the appropriate region from the reference file and apply it to the sub-array input image. For users who take 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.

doFlash - Post-Flash Subtraction

- Header Switch: FLSHCORR
- Header Keywords Updated: MEANFLSH
- Reference File: FLSHFILE (*_fls.fits)

ACS has a post-flash capability to provide the means of mitigating the effects of Charge Transfer Efficiency degradation. The proposer controls whether to use this capability via the FLASHEXP optional parameter in the Phase II proposal.

The function **doFlash** will subtract the Post-Flash reference file, FLSHFILE, from the science image. This file has the same dimensions as a full-size science image, 1062x1044 for HRC and 4144x2068 for WFC, allowing for simple image arithmetic on arrays of equal size. The appropriate FLSHFILE is selected using the following keywords from the image header: DETECTOR, CCDAMP, CCDGAIN, FLASHCUR and SHUTRPOS.

First, the success of the Post-Flash exposure is verified by checking the keyword FLASHSTA. If any problems were encountered, a comment will be added to the history comments in the SCI extension header. The FLSHFILE is then normalized to one second for the appropriate post-flash current level (LOW, MED, HIGH), given by the FLASHCUR keyword.

The FLSHFILE is then multiplied by the flash duration in seconds (FLASHDUR) and is subtracted from the science image. Finally, the mean value of the scaled Post-Flash image is written to the output SCI extension header as the keyword MEANFLSH.

Final Output

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 ACSREJ

- Header Switch: CRCORR

- Header Keywords Updated: TEXPTIME, SKYSUM, EXPEND, REJ_RATE, EXPTIME, NCOMBINE, ROOTNAME

- Header Keywords added as 'HISTORY' Comments: INITGUES, SKYSUB, CRSIGMAS, MEANEXP, CRRADIUS, CRTHRESH, SCALENSE, CRMASK, NEXTEND

- Reference File: CRREJTAB (*_crr.fits)

ACSREJ, the cosmic-ray rejection task in **CALACS**, combines CR-SPLIT exposures into a single image, first detecting and then replacing flagged pixels. The task uses the same statistical detection algorithm developed for STIS data (**OCRREJ**) and WFPC2 data (**CRREJ**), providing a well-tested and robust procedure.

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. The 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 readnoise squared (in DN) and *gain* is the CCD gain (in e⁻/DN) for the amp used to read the pixel,
- *scale* is the scale factor for the noise model,
- T_n is the exposure time (in seconds) for the input image, and
- *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:

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

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 the n'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 keywords TEXPTIME, EXPTIME.

The following keywords are derived from the variables in this computation:

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

The remaining keywords EXPSTART, 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 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) in units of *electrons*.

Cosmic-ray Rejection Table

ACSREJ uses the Cosmic-ray Rejection Parameter Table (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.

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), 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 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 is defined above as $\tau_n = \sigma^2 \times constant$ and uses the sigma rejection thresholds σ. These sigmas correspond to the column CRSIGMAS in the CRREJTAB reference file. 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 the image is locally well-sampled, this will unduly bias *against* rejecting cosmic rays.

Pixels within a given radius, CRRADIUS, of a cosmic-ray 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 for 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.

A much more detailed discussion of cosmic ray elimination in repeated ACS exposures may be found in ACS ISR 2002-08, Section 5.1. Of particular note is the appropriate value of SCALENSE to use; the pipeline adopts a conservative value 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*(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 use **tedit** on the CRREJTAB file (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 example, for program 9662 used in Examples 3 and 4 in Section 3.5.2, there were 14 individual 1.0s HRC exposures all at nominally the same pointing. A check of relative offsets shows that a shift of about 0.2 pixels occurred between the first and last exposures. To obtain a crj 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. Use **tedit** on 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 to 'median' as the INITGUES value. Then rerun **CALACS** on the new association table with 14 entries to obtain a CRJ extension image based on the full stack of 1.0s images.

For a detailed discussion on manually recalibrating ACS data, refer to Section 3.5.

3.4.3 ACS2D

Every observation, whether taken with the MAMA or CCD detectors, will be processed by **ACS2D**. The **ACS2D** primary functions are listed in Table 3.6 and include the 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)

doNoise	Apply a simple noise model (if not done in ACSCCD)
doDQI	Initialize data quality array (if not done in ACSCCD)
doNonLin	Correct and flag non-linear data (**MAMA only**)
doDark	Subtract dark image
doFlat	Divide by flat field and multiply by gain
doShad	Perform CCD shutter shading correction (currently skipped)
doPhot	Compute photometric keyword values for header
doStat	Compute image statistics

doNoise - Error Array Initialization
- Header Switch: None
- Header Keywords Updated: None
- Reference File: None

ACS2D will first check that the image error array has not been expanded, indicating that no previous processing has been performed. In this case, the value of all pixels in the ERR array will have a value of zero. **ACS2D** will then perform the same initialization as described in "doNoise - Error Array Initialization" for **ACSCCD**. If, however, the input image has already been processed, no changes to the ERR array will be made.

doDQI - Bad Pixel Determination
- Header Switch: DQICORR
- Header Keywords Updated: None
- Reference File: BPIXTAB (*_bpx.fits)

If the DQICORR header keyword switch is set to COMPLETE (e.g. CCD data), this step will be skipped. Otherwise, the same initialization will be performed as described in "doDQI - Bad Pixel Determination" for **ACSCCD**.

doNonLin - Linearity Correction for MAMA Data
- Header Switch: LFLGCORR, GLINCORR
- Header Keywords Updated: GLOBLIM
- Reference File: MLINTAB (*_lin.fits)

This routine is capable of flagging global and local nonlinearity in ACS MAMA observations, where the term global refers to the entire ACS detector and 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 from this table refers to the total count rate at which the data are affected by greater than 10% non-linearity across the detector.

CALACS will attempt 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 LFLGCORR 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".

doDark - Dark Image Subtraction
- Header Switch: DARKCORR
- Header Keywords Updated: MEANDARK
- Reference File: DARKFILE (*_drk.fits)

This task is responsible for subtracting the dark image from the input image. For CCD data, the dark image (in *electron s^{-1}*) is multiplied by the exposure time and then divided by the gain before subtracting. (MAMA data do not require a gain correction.) The dark reference file, DARKFILE, is read in line-by-line and 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. The "dark time" is simply the exposure time; it does NOT include the idle time since the last flushing of the chip or the readout time. Any dark accumulation during readout time or post-flash idle will be included automatically into the BIASFILE.

A unique dark reference file is created (approximately) daily using four 1000 sec frames for both the WFC and HRC. The four darks are combined (with cosmic ray rejection) and bias-corrected to form a "daydark". Every two weeks, a high signal-to-noise "basedark" is created from all the frames (~56) in that period. It does not contain the newest hot pixels present in the daydarks since they are rejected with the cosmic rays during image combination. A copy of the basedark is made for each day in the two-week period, and daily hot pixels are skimmed from the respective daydarks and added to the basedark copies. Only the hot pixels above 0.08 e⁻/sec are identified with flag 16 in the data quality (DQ) array of WFC and HRC reference darks, which propogates to the DQ array of the calibrated science data. The flagged hot pixels are typically excluded in subsequent processing steps, as they are likely too noisy to be correctable. But the

many "warm" pixels below this threshold are assumed to be adequately corrected by the dark calibration. This produces a reference file with high signal-to-noise, which accurately reflects (and corrects/flags) the hot pixels present for a given observation date. The "best" dark file is typically not available in the pipeline until 2-3 weeks after the date of observation, because it takes a few weeks to collect enough frames to make a basedark (see ACS ISR 04-07 for more information).

The reference file for dark subtraction, DARKFILE, is selected based on the values of the keywords DETECTOR, CCDAMP, and CCDGAIN in the image header. The dark correction is applied after the overscan regions are trimmed from the input science image. As for the 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 sub-array image.

Sub-array science images will use the same reference file as a full-sized image. **CALACS** simply extracts the appropriate region from the reference file and applies it to the sub-array input image.

doFlat - Flat field image(s) Correction
- Header Switch: FLATCORR
- Header Keywords Updated: None
- Reference File: PFLTFILE (*_pfl.fits), LFLTFILE (*_lfl.fits), DFLTFILE (*_dfl.fits)

As done for STIS, the ACS **doFlat** routine corrects for pixel-to-pixel and large-scale sensitivity gradients across the detector by dividing the data by the flat field image. This contrasts with WFPC2 and NICMOS who deliver an "inverse" flat to the pipeline and then multiply the data by the flat field. When performing the flat field correction, **CALACS** multiplies by the gain so that the calibrated data will be in units of electrons.

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 6.1. This correction is accounted for in pipeline processing by **PyDrizzle** which uses the geometric distortion solution to correct all pixels to equal areas. Thus, in the 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 three flat field reference files: the pixel-to-pixel file (PFLTFILE), the low-order

flat (LFLTFILE), and the delta flat (DFLTFILE). 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 being applied by **CALACS.** The DFLTFILE is a delta-flat containing any needed changes to the small-scale PFLTFILE.

If the LFLTFILE and DFLTFILE are not specified in the SCI header, only the PFLTFILE will be used for the flat field correction. If all three reference files are specified, they will be read in line-by-line and multiplied together to form a complete flat field correction image. 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.

All flat field reference images will be chosen based on the detector, amplifier, and filters used for the observation. Any sub-array 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 sub-array input image.

doShad - Shutter Shading File Correction (skipped)
- Header Switch: SHADCORR
- Header Keywords Updated: None
- Reference File: SHADFILE (*_shd.fits)

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 03-03, "Stability and Accuracy of HRC and WFC Shutters" and to the report "WFC and HRC Shutter Shading and Accuracy" on the ACS Science and Engineering Team web site at:

http://acs.pha.jhu.edu/instrument/calibration/results/by_it
em/shutter/shading_accuracy_feb2001

The **doShad** routine applies the shutter shading correction image (SHADFILE) to the science data if SHADCORR is set to PERFORM. The input image is corrected for the differential exposure time across the detector caused by the shutter opening. This correction only has a significant effect on images with short exposure times (less than 10 seconds).

Pixels are corrected based on the exposure time using the relation:

$$corrected = uncorrected \times EXPTIME/(EXPTIME + SHADFILE).$$

The SHADFILE is selected using the DETECTOR keyword in the input science image. At present, this reference file is an over scanned-trimmed

image, binned by a factor of 8 for each dimension, resulting in a 128x128 HRC and a 256x512 WFC SHADFILE.

The shutter shading correction can either be applied during **ACS2D** processing for single or `REPEAT-OBS` exposures or during cosmic-ray rejection in **ACSREJ** for `CRSPLIT` exposures.

doPhot - Photometry Keyword Calculation

- Header Switch: PHOTCORR
- Header Keywords Updated: PHOTMODE, OBSMODE, PHOT-FLAM, PHOTZPT, PHOTPLAM, PHOTBW
- Reference Files: GRAPHTAB (*_tmg.fits), COMPTAB (*_tmc.fits)

Before photometry can be performed on ACS observations, a transformation from electrons to absolute flux units must be performed. **CALACS** follows the WFPC2 methodology for calculating the photometry keywords in the calibration pipeline. The calibration reference files, GRAPHTAB and COMPTAB, point to the SYNPHOT tables containing the latest ACS component tables. These tables contain the throughput as a function of wavelength for the various ACS detector and filter combinations. This reliance on SYNPHOT allows the ACS team to maintain the latest throughput files in SYNPHOT to keep **CALACS** up to date. For further discussion of SYNPHOT, refer to Chapter 3 of the HST Data Handbook Introduction.

The keywords PHOTMODE and OBSMODE are built to reflect the configuration of the instrument during the exposure. **CALACS** then uses the OBSMODE string with SYNPHOT to compute the total throughput array for the observation, based on the aperture and filter throughputs, and the photometry keywords: PHOTFLAM (the inverse sensitivity in units of *erg s^{-1} cm^{-2} A^{-1}* for 1 *electron s^{-1}*), PHOTZPT (the Space Telescope magnitude zero point), PHOTPLAM (the bandpass pivot wavelength), and PHOTBW (the bandpass RMS width). 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. Drizzled DRZ images are already in units of *electrons per second* and may simply be multiplied by the PHOTFLAM value to attain flux units.

doStat - Image Statistics Determination

- Header Switch: STATFLAG
- Header Keywords Updated: NGOODPIX, GOODMIN, GOOD-MAX, GOODMEAN, SNRMIN, SNRMAX, SNRMEAN
- Reference File: 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.

3.4.4 ACSSUM

- Header Switch: RPTCORR
- Header Keywords Updated: NCOMBINE, EXPTIME, EXPEND, ROOTNAME
- Reference File: None

ACSSUM sums a list of images one line at a time, using the association file to define the input. 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. Note that for multiple exposures taken via REPEAT-OBS instead of the observationally equivalent CR-SPLIT, the final processed image is *not* cosmic-ray rejected.

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.

3.5 Manual Recalibration of ACS Data

3.5.1 Requirements for Manual Recalibration

Retrieving Software and Input Data Files

If the observer decides to recalibrate their ACS data, the following must be available on their system:

- **CALACS** software from the STSDAS **hst_calib.acs** package
- **PyDrizzle** task from the STSDAS **dither** package
- PyRAF, to run **PyDrizzle**, obtained from STScI
- Reference files obtained from STScI Archive, or personal versions
- Uncalibrated data (*_raw.fits) from OTFR
- Association table, if applicable.

Uncalibrated data and calibration reference files can be obtained from OTFR via StarView or the web-based archive request form:

http://archive.stsci.edu/cgi-bin/hst.

StarView is available for users to download from:

http://starview.stsci.edu/html.

The pages include instructions on installing and using StarView and on how to become a registered archive user.

The most recent version of STSDAS, which includes the **CALACS** package required for reprocessing, can be downloaded from:

```
http://www.stsci.edu/resources/software_hardware/stsdas
```

The most recent versions of PyRAF, **PyDrizzle**, and **MultiDrizzle** can be downloaded from:

```
http://www.stsci.edu/resources/software_hardware/stsdas/pyd
rizzle
```

Setting up 'jref'

Before any recalibration may be done, the user's local directory containing the calibration reference files must be defined. For ACS, this directory is referred to as 'jref'. The image header already contains the appropriate keywords defining the reference files. The user must simply define the location of the 'jref' directory in the Unix environment.

```
setenv jref /mydisk/myjref/
```

Once this path is set, IRAF must be restarted in the same window in which the setup was performed.

Selecting Calibration Switches

OTFR will use the latest available calibration reference files by default. In order to use non-default reference files, manual recalibration is required. The calibration reference file keywords will need to be updated manually in the uncalibrated data with the non-default filenames before running **CALACS**. The selection criteria in Table 3.7 are used to set the values for the calibration switch header keywords in uncalibrated ACS data.

Table 3.7: Calibration Switch Selection Criteria.

Switch	Description	Criteria
DQICORR	Data Quality Array Initialization	DEFAULT = **"PERFORM"** If OBSMODE = ACQ then **"OMIT"**
ATODCORR	Analog to Digital Conversion	DEFAULT = **"OMIT"**
BLEVCORR	Overscan Region Subtraction	DEFAULT = **"PERFORM"**
BIASCORR	Bias Subtraction	DEFAULT = **"PERFORM"**
FLSHCORR	Post Flash Subtraction	DEFAULT = **"OMIT"**
CRCORR	Cosmic Ray Rejection	If CRSPLIT >= 2 then **"PERFORM"** If CRSPLIT < 2 then **"OMIT"**
EXPSCORR	Calibrate individual exposures in an association	DEFAULT = **"PERFORM"**
SHADCORR	Shutter Shading Correction	DEFAULT = **"OMIT"**
DARKCORR	Dark Subtraction	DEFAULT = **"PERFORM"**
FLATCORR	Flat field Division	DEFAULT = **"PERFORM"**
PHOTCORR	Photometric Processing	DEFAULT = **"PERFORM"**
RPTCORR	Repeat Processing	DEFAULT = **"OMIT"** If NRPTEXP > 1 then **"PERFORM"**
DRIZCORR	Dither Processing	DEFAULT = **"PERFORM"**
GLINCORR	Global Non-Linearity Correction	SBC DEFAULT = **"PERFORM"**
LFLGCORR	Local and Global Non-Linearity Flagging in DQ Array	SBC DEFAULT = **"PERFORM"**

Bypassing the PHOTCORR Step

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 pipeline calibration tasks do not alter the units of the pixels in the image. Instead they calculate and write the inverse sensitivity conversion factor (PHOTFLAM) and the ST magnitude scale zero point (PHOTZPT) into header keywords in the calibrated data. Refer to the section "doPhot - Photometry Keyword Calculation" in Section 3.4.3 for more information.

When populating the photometric keywords during the PHOTCORR step, **CALACS** uses the STSDAS synthetic photometry package, **SYNPHOT**, which contains the throughput curves of all HST optical components. (For information on retrieving the **SYNPHOT** dataset, see Appendix A.3.2 of the HST Data Handbook Introduction.) The **SYNPHOT** dataset contains numerous files which may be updated on a regular basis, making it cumbersome for the user to set up. A simple alternative is to set the PHOTCORR calibration switch to OMIT in the primary header of the uncalibrated image. The user may simply copy the photometric keywords from the calibrated data into the raw image header and then run **CALACS** with the PHOTCORR switch set to OMIT.

Speed of Pipeline Processing

Reprocessing HRC or SBC data will not put a burden on most computing systems since the image sizes, both science data and reference files, are relatively small. Processing WFC observations, on the other hand, will require more computing power, including both CPU run time and disk space.

Great care has been taken to minimize the memory requirements of the pipeline software to accommodate most computing configurations. Line-by-line I/O is used during pipeline processing and is particularly useful when more than one image is operated on at a time, for example during flat fielding or co-adding images. This, unfortunately, places an extra burden on the I/O capabilities of the computing system. **CALACS** requires up to 130MB of memory to process WFC data, while **PyDrizzle** requires up to 400MB.

An Ultra-10 class workstation (300 MHz, 512MB RAM) using **CALACS** to process WFC data will take ~13 minutes for a CR-SPLIT=2 association or ~4.5 minutes for a single WFC exposure. **PyDrizzle** requires ~5 minutes to process a single WFC exposure. These times can be scaled by the processor speed of the computing system, and result in much faster run times for newer workstations.

3.5.2 CALACS Examples

We present several examples of **CALACS** reprocessing. The steps required for general **PyDrizzle** reprocessing are summarized in the "PyDrizzle Examples" on page 4-28.

Example 1: Reprocessing a Single Exposure

The following example uses HRC data from the flat field calibration program 9019 which observed the stellar cluster 47 Tucanae. The observations are from visit 07, exposure 12 and utilize the F814W filter. The exposures are CR-SPLIT into two exposures of 20 sec each. The association table for this exposure is j8bt07020_asn.fits. Typing 'tprint j8bt07020_asn.fits' reveals the rootnames of the individual exposures: j8bt07oyq and j8bt07ozq.

For the purposes of this first example, let us assume that the observations are *not* part of an association. This example will illustrate the steps required to reprocess a single exposure after changing the bias reference file from the default value to a file specified by the user.

1. In the Unix environment, we must specify the location of the 'jref' directory where the calibration reference files reside and then start IRAF. We assume here that the user has requested calibration files as part of the data archive request and stored these in their directory '/mydisk/myjref/', where '/mydisk/' specifies the directory containing the pipeline data.

(Note: This setup must be done in the same window in which IRAF will be used. Note also that setting 'jref' from within IRAF will not work even though typing 'show jref' would suggest it might.)

```
unix> setenv jref /mydisk/myjref/
unix> cl
```

2. To see what bias reference file is currently in use, we use the hedit task. Note that the period at the end of the line is required. The lines below beginning with the 'iraf>' symbol indicate commands typed into IRAF. Lines with no symbol indicate output from IRAF.

```
iraf> hedit j8bt07oyq_raw.fits[0] BIASFILE .
j8bt07oyq_raw.fits[0], BIASFILE = jref$m4r1753rj_bia.fits
```

3. Next, we edit the global image header of the raw image to reflect the name of the new bias file.

```
iraf> hedit j8bt07oyq_raw.fits[0] BIASFILE jref$mybias.fits verify-
j8bt07oyq_raw.fits[0], BIASFILE: jref$m4r1753rj_bia.fits->
jref$mybias.fits
j8bt07oyq_raw.fits[0] updated
```

4. We then set the PHOTCORR processing step to 'OMIT' and copy the photometric keywords from the calibrated to the raw image. (Alternately, the user may keep track of these numbers in any other preferred manner. Most users will only require knowledge of the PHOTFLAM keyword for photmetric calibration.) These steps allow users to bypass this SYNPHOT-based calibration step, due to the cumbersome requirements for setup. (See "Bypassing the PHOTCORR Step" on page 3-27 for more information.)

```
iraf> hedit j8bt07oyq_raw.fits[0] PHOTCORR omit verify-
j8bt07oyq_raw.fits[0], PHOTCORR: PERFORM -> omit
j8bt07oyq_raw.fits[0] updated
iraf> hedit j8bt07oyq_flt.fits[1] phot* .
j8bt07oyq_flt.fits[1], PHOTMODE = "ACS HRC F814W"
j8bt07oyq_flt.fits[1], PHOTFLAM = 1.3260861E-19
j8bt07oyq_flt.fits[1], PHOTZPT = -2.1100000E+01
j8bt07oyq_flt.fits[1], PHOTPLAM = 8.1209805E+03
j8bt07oyq_flt.fits[1], PHOTBW = 7.0114880E+02
iraf> hedit j8bt07oyq_raw.fits[1] PHOTFLAM 1.3260861E-19 verify-
j8bt07oyq_raw.fits[1],PHOTFLAM: 0.000000E+00 -> 1.3260861E-19
j8bt07oyq_raw.fits[1] updated
```

5. Within the directory containing the pipeline products ("/mydisk/"), we create the subdirectory "recal", copy the RAW and ASN files to this subdirectory, and 'cd' to this directory.

6. Finally, we load the packages **stsdas.hst_calib.acs** and run **CALACS** on the single raw exposure.

(Note that invoking this package will result in the message: "Warning: package 'acs' includes Python tasks that require PyRAF". This applies only for **calacsdriz** which calls the **CALACS** and **PyDrizzle** tasks in succession.)

```
iraf> stsdas.hst_calib.acs
iraf> calacs j8bt07oyq_raw.fits
```

The product will be a single calibrated image with the FLT extension.

Example 2: Reprocessing Multiple Exposures 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. IRAF output comments which are similar to Example 1 have been omitted.

1. First, we look at the contents of the image association.

```
iraf> tprint j8bt07020_asn.fits

# MEMNAME        MEMTYPE          MEMPRSNT

j8bt07oyq        EXP-CRJ          yes

j8bt07ozq        EXP-CRJ          yes

j8bt07021        PROD-CRJ         yes
```

2. To see what bias reference file is currently in use, we use **hedit**.

```
iraf> hedit j8bt07*raw.fits[0] BIASFILE .
```

3. Next, we edit the global image header of all the raw images to reflect the name of the new bias file.

```
iraf> hedit j8bt07*raw.fits[0] BIASFILE jref$mybias.fits verify-
```

4. We then set the PHOTCORR processing step to 'OMIT' and copy the photometric keywords from the calibrated to the raw image.

```
iraf> hedit j8bt07*raw.fits[0] PHOTCORR omit verify-
iraf> hedit j8bt07*flt.fits[1] phot* .
iraf> hedit j8bt07*raw.fits[1] PHOTFLAM 1.3260861E-19 verify-
```

5. Within the directory containing the pipeline products ("/mydisk/"), we create the subdirectory "recal", copy the RAW and ASN files to this subdirectory, and 'cd' to this directory.

6. Finally, we run **CALACS** on the image association.

```
iraf> stsdas.hst_calib.acs
iraf> calacs j8bt07020_asn.fits
```

(If executing this command in the same directory in which you have run the previous example, then one of the FLT files will already exist and **CALACS** will not overwrite existing images. To execute in this case, first delete the existing FLT files or move them to a separate directory.)

The product will be two separate calibrated images with the FLT extension (j8bt07oyq_flt.fits, j8bt07oyq_flt.fits) and a single cr-combined image with the CRJ extension (j8bt07021_crj.fits).

Example 3: Removing Cosmic-rays from Repeated (RPT-OBS) Exposures

The following example uses HRC data from the shutter stability calibration program 9662 which observed the stellar cluster 47 Tucanae. The observations are from visit 01, exposure 20 and utilize the CLEAR filter. The two repeated (RPT-OBS) exposures are 1.0 sec each.

In this example, we illustrate the steps required to remove cosmic-rays from repeated observations which use the RPT-OBS syntax. By definition, the standard calibration product of multiple repeated exposures is a summed image. We can fool the software into rejecting cosmic-rays from the images, but this requires modification of the image association table.

1. First, we look at the contents of the image association.

```
iraf> tprint j8is01020_asn.fits

# MEMNAME        MEMTYPE         MEMPRSNT

j8is01j2q        EXP-RPT         yes

j8is01j3q        EXP-RPT         yes

j8is01021        PROD-RPT        yes
```

2. Update calibration switches in the global header using the **hedit** task.

```
iraf> hedit *raw.fits[0] CRCORR perform add-
iraf> hedit *raw.fits[0] RPTCORR omit add-
```

3. Update association table MEMTYPE column, by changing each instance of RPT to CR. This may be done using the text editor functionality of **tedit**.

```
iraf> tedit j8is01020_asn.fits

# MEMNAME        MEMTYPE         MEMPRSNT

j8is01j2q        EXP-CRJ         yes

j8is01j3q        EXP-CRJ         yes

j8is01021        PROD-CRJ        yes
```

4. We then set the PHOTCORR processing step to 'OMIT' and copy the photometric keywords from the calibrated to the raw image.

```
iraf> hedit j8is01*raw.fits[0] PHOTCORR omit verify-
iraf> hedit j8is01*flt.fits[1] phot* .
iraf> hedit j8is01*raw.fits[1] PHOTFLAM 4.5803514E-20 verify-
```

5. Within the directory containing the pipeline products ("/mydisk/"), we create the subdirectory "recal", copy the RAW and ASN files to this subdirectory, and 'cd' to this directory.

6. Finally, we run **CALACS** on the image association.

```
iraf> stsdas.hst_calib.acs
iraf> calacs j8is01020_asn.fits
```

The product will be two calibrated images with the FLT extension (j8is01j2q_flt.fits, j8is01j3q_flt.fits) and a cr-combined image with the CRJ extension (j8is01021_crj.fits) instead of the usual SFL extension.

Example 4: Combining Exposures from Multiple Associations

This example uses the same data from Example 3 plus an additional exposure from the same program: visit 01, exposure 40. These additional data also use the CLEAR filter and contain two repeated 1.0 sec exposures. The association table for this exposure is j8is01040_asn.fits. In this example, we illustrate the steps required to combine two separate sets of repeated observations.

1. First we merge the two association tables and edit them appropriately (using **tmerge** and **tedit**, respectively).

```
iraf> tmerge j8is01020_asn.fits,j8is01040_asn.fits \
>>>   outtable="merged_asn.fits" option=append
iraf> tprint merged_asn.fits

# MEMNAME        MEMTYPE        MEMPRSNT

j8is01j2q       EXP-RPT        yes

j8is01j3q       EXP-RPT        yes

j8is01021       PROD-RPT       yes      <---remove this line

j8is01j8q       EXP-RPT        yes

j8is01j9q       EXP-RPT        yes

j8is01041       PROD-RPT       yes      <---edit this line

iraf> tedit merged_asn.fits

# MEMNAME        MEMTYPE        MEMPRSNT

j8is01j2q       EXP-RPT        yes

j8is01j3q       EXP-RPT        yes

j8is01j8q       EXP-RPT        yes

j8is01j9q       EXP-RPT        yes

j8is01xx1       PROD-RPT       yes      <---new product name
```

2. We then set the PHOTCORR processing step to 'OMIT' and copy the photometric keywords from the calibrated to the raw image using **hedit**.

```
iraf> hedit j8is01*raw.fits[0] PHOTCORR omit verify-
iraf> hedit j8is01*flt.fits[1] phot* .
iraf> hedit j8is01*raw.fits[1] PHOTFLAM 4.5803514E-20 verify-
```

3. Within the directory containing the pipeline products ("/mydisk/"), we create the subdirectory "recal", copy the RAW and ASN files to this subdirectory, and 'cd' to this directory.

4. Finally, we run **CALACS** on the new image association.

```
iraf> stsdas.hst_calib.acs
iraf> calacs merged_asn.fits
```

The product will be 4 calibrated images with the FLT extension and a single summed image with the SFL extension (j8is01xx1_sfl.fits).

Example 5: 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 and target the CDF-South with the F606W filter. The images are part of a 2-point line dither pattern and are 480 sec each. In this example, we illustrate the steps required to reprocess data which is part of a dither pattern, using a non-default dark reference file. The steps are similar to Example 2, but the format of the association and the data products are unique.

1. First, we look at the contents of the image association.

```
iraf> tprint j8e654010_asn.fits
# MEMNAME        MEMTYPE         MEMPRSNT
j8e654c0q        EXP-DTH         yes
j8e654c4q        EXP-DTH         yes
j8e654011        PROD-DTH        yes
```

2. To see what dark reference file is currently in use, we use **hedit**.

```
iraf> hedit j8e654*raw.fits[0] DARKFILE .
```

3. Next, we edit the global image header of all the raw images to reflect the name of the new dark reference file.

```
iraf> hedit j8e654*raw.fits[0] DARKFILE jref$mydark.fits verify-
```

4. We then set the PHOTCORR processing step to 'OMIT' and copy the photometric keywords from the calibrated to the raw image.

```
iraf> hedit j8e654*raw.fits[0] PHOTCORR omit verify-
iraf> hedit j8e654*flt.fits[1] phot* .
iraf> hedit j8e654*raw.fits[1] PHOTFLAM 7.7792272E-20 verify-
```

5. Within the directory containing the pipeline products ("/mydisk/"), we create the subdirectory "recal", copy the RAW and ASN files to this subdirectory, and 'cd' to this directory.

6. Finally, we run **CALACS** on the image association.

```
iraf> stsdas.hst_calib.acs
iraf> calacs j8e654010_asn.fits
```

The product of the DITH-PATTERN will be two separate calibrated images with the FLT extension. In subsequent processing (see Chapter 4 for details), **PyDrizzle** will combine these FLT files into a single DRZ image (j8e654011_drz.fits).

CHAPTER 4:

Distortion Correction and Drizzling of ACS Images

In this chapter. . .

This chapter describes the geometric distortion of the ACS and presents the available tools and methods for removing the distortion and combining dithered datasets into clean, undistorted, and photometrically faithful output products.

The first section describes how the distortion has been quantified and how this information is held in reference files which are used by the calibration pipeline. Plots of the distortion field are given for the three ACS detectors.

The remaining sections give an overview of the drizzle method and discuss the higher level **PyDrizzle** and **MultiDrizzle** software tools available for distortion correction and image combination.

4.1 ACS Geometric Distortions

ACS produces the largest single images ever taken with HST. However, because the optics were designed with a minimum number of components, ACS focal surfaces are far from normal to the principal rays. This results in an image of the sky which is distorted in a large, but predictable manner. The distortion consists of two effects, discussed previously in Section 1.1.3. The first is the elongation of the ACS apertures, causing the pixel scale to be smaller along the radial direction of the Optical Telescope Assembly (OTA) field of view than along the tangential direction. The second effect is the variation of pixel area across the detector.

The results presented in the discussion below are derived from multiple observations of dense star fields with the images displaced from each other by amounts comparable with the size of the field of view. The following discussion is based on reports from the HST Calibration Workshop Proceedings: *Calibration of Geometric Distortion in the ACS Detectors* (Meurer et al. 2002) and *Astrometry with the Advanced Camera: PSFs and Distortion in the WFC and HRC* (Anderson 2002) at:

http://www.stsci.edu/hst/HST_overview/documents/calworkshop
/workshop2002/CW2002_TableOfContents

4.1.1 Distortion Model

The heart of the distortion model relates detector pixel position (x, y) to sky position (see ACS ISR 2000-11) using a polynomial transformation given by:

$$x_c = \sum_{i=0}^{k} \sum_{j=0}^{i} a_{i,j}(x - x_r)^j (y - y_r)^{i-j}$$

$$y_c = \sum_{i=0}^{k} \sum_{j=0}^{i} b_{i,j}(x - x_r)^j (y - y_r)^{i-j}$$

where k is the polynomial order of the fit, (x_r, y_r) is the reference pixel position taken to be the center of each detector or WFC chip, and (x_c, y_c) is the undistorted position in arcseconds. The coefficients to the fits, $a_{i,j}$ and $b_{i,j}$, are free parameters. A similar form is used for the inverse relation which provides the conversion from units of arcseconds to distorted input pixels.

For the WFC, an offset is applied to put the two CCD chips on the same coordinate system:

$$X = x_c + \Delta x(chip), \quad Y = y_c + \Delta y(chip).$$

The offsets $\Delta x(chip), \Delta y(chip)$ are 0,0 for chip 1 and correspond to the separation between chips 1 and 2 for chip 2. The chip 2 offsets are free parameters in the fit. (X, Y) correspond to tangential plane positions in

arcseconds which are tied to the HST $(V2, V3)$ coordinate system. Finally, all frames are transformed to the same coordinate grid on the sky:

$$X_{sky} = \Delta X_i + \cos(\Delta\theta_i)X' - \sin(\Delta\theta_i)Y',$$

$$Y_{sky} = \Delta Y_i + \sin(\Delta\theta_i)X' + \cos(\Delta\theta_i)Y',$$

where the free parameters $\Delta X_i, \Delta Y_i, \Delta\theta_i$ are the position and rotation offsets of frame i.

There is one added complication in the scale term of the geometric distortion due to the effect of differential velocity aberration. The apparent position on the sky of any targets is displaced by the effect of velocity aberration. The speed of the earth in its orbit is about 30 km/sec. Figure 4.1 illustrates a target which is viewed at an angle α with respect to the velocity of the telescope. The apparent position of the target is at angle α' where:

$$\tan\alpha' = \frac{\sqrt{1-\beta^2} \cdot \sin\alpha}{\cos\alpha + \beta}$$

and β is the net telescope velocity divided by the speed of light. The change in angle can amount to tens of arcseconds, but the HST flight software continually accounts for this and keeps the apparent position of the target at a fixed aperture reference point. However a secondary differential effect is a small change in the scale of magnitude equal to:

$$\frac{\sqrt{1-\beta^2}}{1 + \beta\cos\alpha}$$

The scale change is of order 1 part in 10^4. For targets near the ecliptic, this can cause a misalignment between different exposures of up to 1.4 pixels in the ACS field of view in WFC images taken 6 months apart. This scale change, illustrated in Figure 4.2, can complicate the removal of cosmic rays from images and can cause artifacts when constructing mosaics.

The effective scale change in arcsec per pixel is actually the reciprocal of the factor given above, i.e., if the apparent positions appear closer together than a single pixel sees a larger portion of the sky. The ACS science images now include a keyword VAFACTOR equal to the approximation $(1+\beta \cos\alpha)$. This keyword is used by the **PyDrizzle** software when correcting for geometric distortion.

For a further discussion of the effect of velocity aberration, see *The Effect of Velocity Aberration Correction on ACS Image Processing*: HST Calibration Workshop Proceedings (Cox & Gilliland 2002) at:
http://www.stsci.edu/hst/HST_overview/documents/calworkshop
/workshop2002/CW2002_Papers/CW02_cox

Figure 4.1: Schematic illustration of the angular change due to velocity aberration.
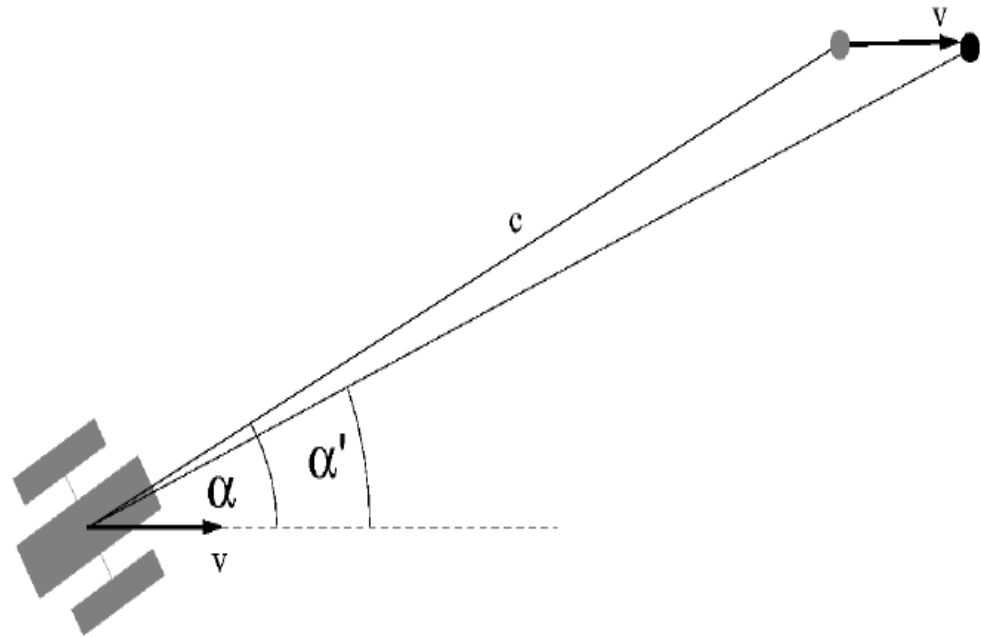
Figure 4.2: Pixel shifts on WFC caused by scale changes from differential velocity abberation.

PIXEL SHIFTS

Maximum shift due to HST orbit

Maximum variation due to Earth's orbit
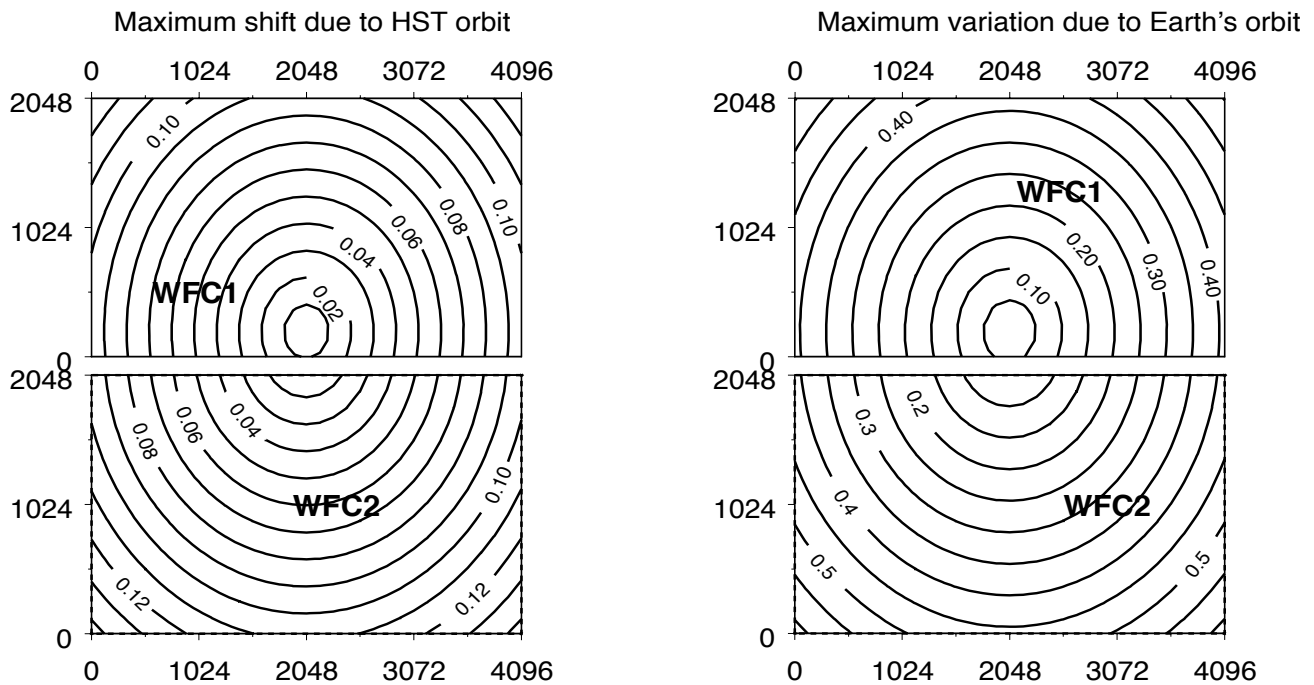
4.1.2 On-orbit Calibration Programs

The positions of stars observed multiple times in the dithered star fields are used to iteratively solve for the free parameters in the distortion

solution: fit coefficients $a_{i,j}$, $b_{i,j}$; chip 2 offsets Δx, Δy (WFC only); frame offsets ΔX_i, ΔY_i, $\Delta\theta_i$; and tangential plane position X_{sky}, Y_{sky} of each star used in the fit.

The original ACS geometric distortion campaign observed the core of 47 Tucanae with the WFC and HRC (program 9028, Meurer). The exposures were designed to detect stars on the main sequence turn-off, allowing a high density of stars with relatively short exposures. The F475W filter (Sloan g') was used to minimize the number of saturated red giant branch stars in the field. For calibrating the distortion in the SBC, exposures of NGC6681 (program 9027, Hartig) were chosen for the relatively high density of UV emitters (hot horizontal branch stars).

Initially the distortion analysis was performed only with the F475W filter for both WFC and HRC detectors. Further measurements and analysis have provided solutions for 10 HRC filters and 12 WFC filters. Scale variations of a few parts in 10^4 are found between regular filters but changes of a percent or more are evident when polarizing filters are included.

Originally, only images taken with a single roll angle were used to define the distortion solutions. The solution using only these data is degenerate in the zeroth (absolute pointing) and linear terms (scale, skewness). Therefore the largest commanded offsets with a given guide star pair were used to set the linear terms. However, comparison of corrected coordinates to astrometric positions showed that residual skewness in the solution remained. Hence, as of November 2002, the results for the WFC and HRC are based on data from multiple roll angles, incorporating data obtained from the calibration outsourcing program 9443 (King & Anderson). The overall plate scale is set by the largest commanded offset. For the HRC, the linear scale is set by matching HRC and WFC coordinates, since the same field was used in the SMOV observations. The zeroth order terms (position of ACS apertures in the V2,V3 frame) were determined from observations of an astrometric field.

4.1.3 Results

The most obvious component of the ACS distortion is a marked skew which results in the X/Y detector axes lying at ~85 degrees to each other on the sky. In addition, the distortion in all ACS detectors is highly non-linear (see Figure 4.3-4.6). A quartic fit ($k=4$) is adequate for characterizing the distortion to an accuracy much better than 0.2 pixels over the entire field of view.

WFC

The WFC detector is tilted at 22 degrees giving an elongation of 8% along the diagonal. The non-square projected aperture shape is evident from Figure 1.1 in Chapter 1, where the x axis is approximately in the V2 direction and the y axis is in the -V3 direction.

The WFC distortion is also illustrated in Figure 4.3, a vector displacement diagram which shows the contribution of the non-linear part

of a quartic fit to the data. The vectors represent the degree of distortion to be expected in the WFC beyond the directional dependence of the plate scale. For display, the vectors are magnified by a factor of 5 compared to the scale of the x and y axes. The largest displacement indicated at the top left corner of the figure is ~82 pixels or about 4 arcseconds. While a quartic solution is adequate for most purposes, binned residual maps (Figure 4.6) show that there are significant coherent residuals in the WFC solutions which have amplitudes up to ~0.1 pixels. These residuals are not currently accounted for in pipeline calibration, although efforts are underway to do so in the future (see Section 4.1.4).

The pixel scales and array sizes are summarized in Table 4.1 for each ACS detector. A contour plot of relative pixel area across the WFC is shown in Figure 6.1. The range of area is from 0.89 to 1.08 times the central value.

When deriving the ACS flat fields, an implicit correction is made for the distortion which is equivalent to dividing each pixel by the relative pixel area. Thus, point source photometry extracted from a flat fielded image must be multiplied by the effective pixel area. This correction is accounted for in pipeline processing by **PyDrizzle** which uses the distortion solution to correct all pixels to equal areas. When working with FLT files this correction must be applied manually (see Section 6.1.3 for details).

Table 4.1: Pixel scale *(arcsec/pixel)* and array size (*pixels*) for each ACS detector.

Detector	Pixel Scale at Center X_{scale}, Y_{scale}	Average Pixel Scale X_{scale}, Y_{scale}	Array size (pixels)	Array size (arcsec)
WFC1	0.04927, 0.04864	0.04917, 0.04861	4096x2048	201"x100"
WFC2	0.04985, 0.05031	0.04976, 0.05029	4096x2048	204"x103"
HRC	0.02844, 0.02484	0.02843, 0.02484	1024x1024	29"x25"
SBC	0.03374, 0.03009	0.03377, 0.03010	1024x1024	35"x31"

Figure 4.3: Non-linear component of the ACS distortion for the WFC detector.
Note that this figure is rotated 180 degrees with respect to the pipeline calibration
products, where WFC2 is the lower half of the detector.

WFC non−linear distortion vectors (x5)
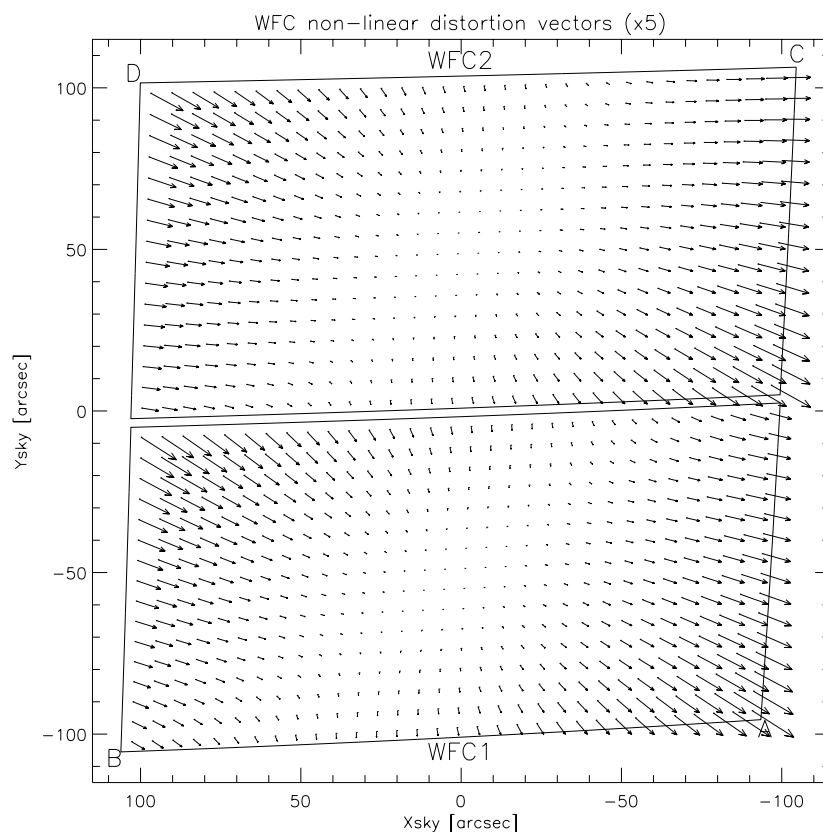
HRC

The High Resolution Channel has its edges aligned approximately along the V2 and V3 axes (see Figure 1.1). In this case, the center of the aperture lies on a line passing through the V2/V3 origin and making an angle of 22 degrees with the V3 axis. Because the aperture diagonal does not correspond to a radius of the HST field of view, the distortion has no particular symmetry with respect to the detector axes. Because the focal plane, and therefore the detector plane is 25 degrees away from the plane normal to the light path, the scales along the axes differ by ~14%. However, since the HRC is less than 30 arcsec across, the *total* variation in scale across the detector is much less than for the WFC, being only about 1%.

A vector plot of the deviation from linearity is given in Figure 4.4 in which the deviations have been magnified by a factor of 10 for illustrative purposes. The largest displacement is 4.9 pixels in the top left corner and corresponds to about 0.1 arcsec. While a quartic solution is adequate for most purposes, binned residual maps (Figure 4.6) show that there are significant coherent residuals in the HRC solutions which have amplitudes up to ~0.1 pixels. These residuals are not currently accounted for in pipeline calibration, although efforts are underway to do so in the future (see Section 4.1.4).

The pixel scales and array sizes are summarized in Table 4.1 for the HRC. The variation of pixel area across the detector is shown in Figure 6.1. The maximum deviation from the central value is just over 2%. As explained for WFC, this must be taken into account when doing photometry, either by using **PyDrizzle** output products (DRZ files) or by manually correcting FLT files with a pixel area map (see Section 6.1.3 for details).

Figure 4.4: Non-linear component of the ACS distortion for the HRC detector.
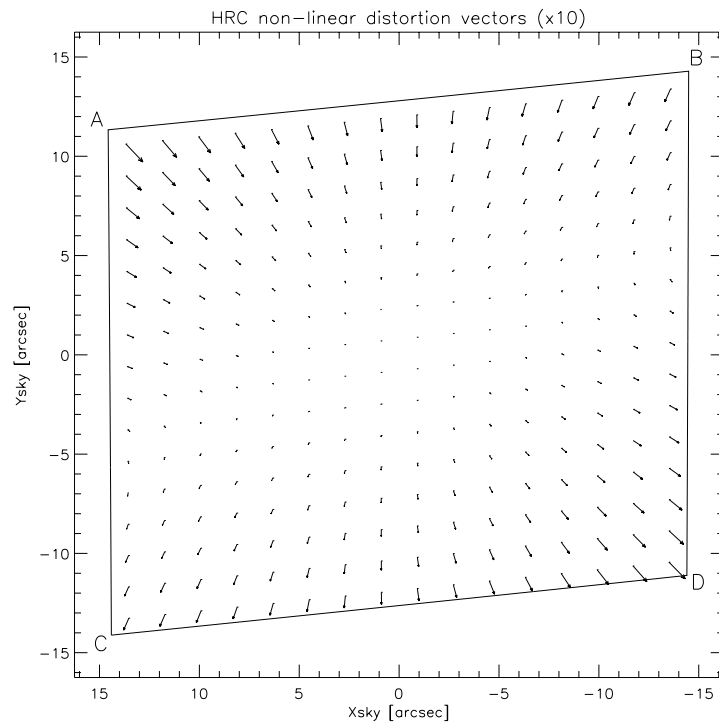
SBC

The Solar Blind Channel contains the MAMA detector. It is centered at approximately the same place as the HRC in the V2/V3 plane and is slightly larger, about 35 x 30 arcsec, versus 29 x 25 arcsec for the HRC. The average scales in the x and y directions are given in Table 4.1 and are calculated by matching observations of the same sky area by HRC and SBC.

The optical distortions are close to those displayed by the HRC. The only difference in their light paths is the presence of a plane M3 fold mirror which reflects light away from the SBC onto the HRC. Because there is one less mirror in the SBC light path, the distortion corrected image (in sky coordinates) in Figure 4.5 appears flipped with respect to the HRC image in Figure 4.4. The MAMA has a small amount of extra distortion in the detector itself, arising from irregularities in the multi-channel plate. The largest difference from a square pattern is 2 pixels, or 0.06 arcsec. As explained for WFC, pixel area variations must be taken into account when doing photometry, either by using **PyDrizzle** output products (DRZ files) or by manually correcting FLT files with a pixel area map (see Section 6.1.3 for details).

Figure 4.5: Non-linear component of the ACS distortion for the SBC detector.
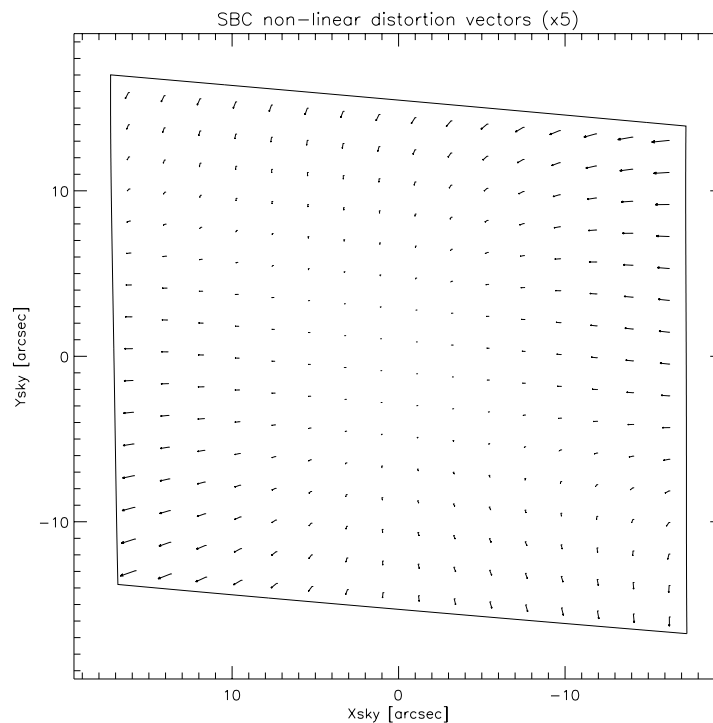
Figure 4.6: Binned residuals to quartic distortion model fits for the ACS WFC and HRC detectors. Note that the WFC figure is rotated 180 degrees with respect to the pipeline calibration products, where WFC2 is the lower half of the detector. The large residuals close to the upper edge of the HRC map, left of center, are due to the presence of the Fastie occulting finger. There is insufficient data to make the equivalent plot for the SBC.
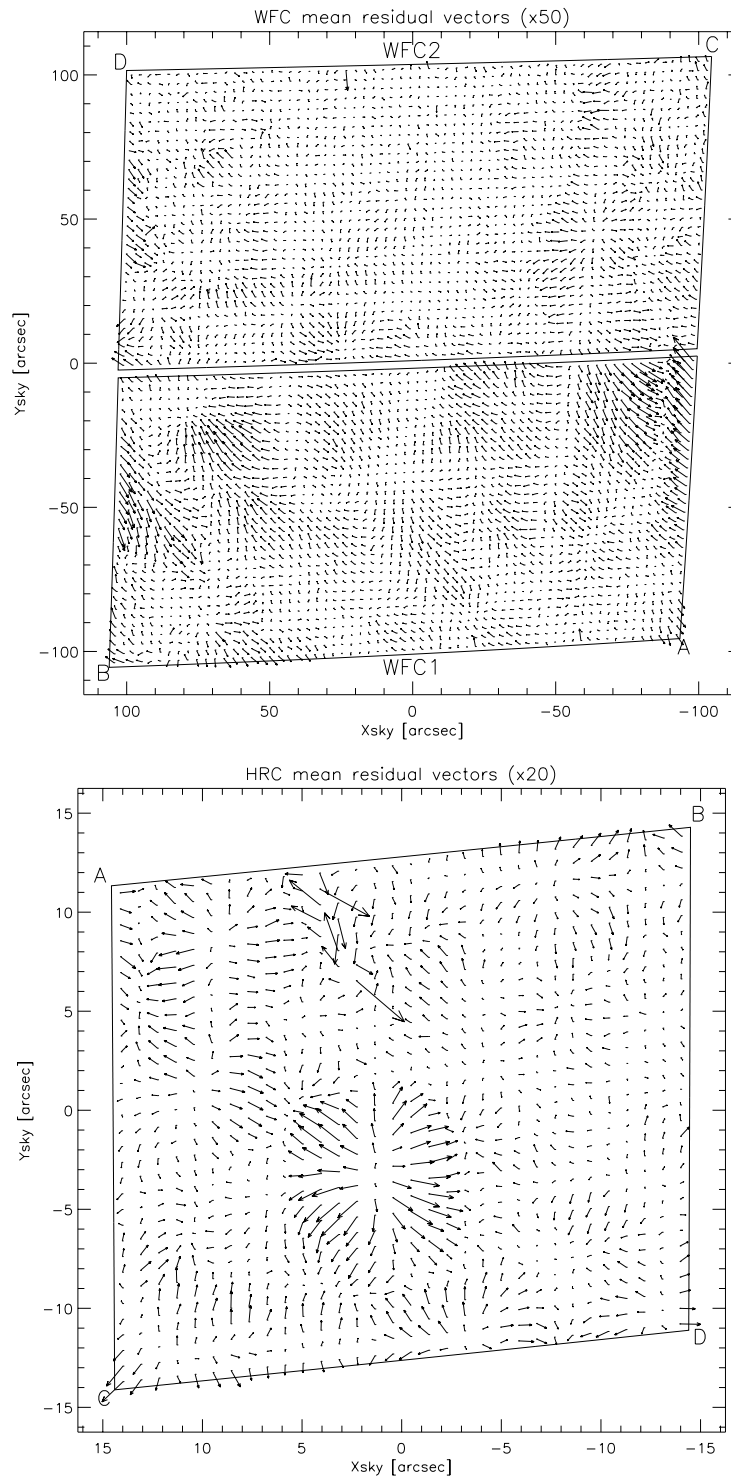
4.1.4 Distortion Table (IDCTAB)

The Science Instrument Aperture File (SIAF) contains position and scale information for every aperture of each of HST's science instruments. Thus, it supports accurate target acquisitions, image processing, and photometry. The conversion between distorted and undistorted positions in the SIAF is controlled by a polynomial expansion which is derived by the individual instrument teams. A simple, common reference file can therefore be created for each instrument which describes the geometric distortion of each detector. In this section, we describe the format of this common reference file which has been developed to be as general as possible so that it can be used for many different instruments or detectors.

For ACS, the coefficients to the polynomial fit are found in the IDCTAB reference file. The distortion correction is not a part of **CALACS**. A separate task, **PyDrizzle** (see Section 4.3), converts the distortion model in the IDCTAB into a form which is used by **drizzle** to produce geometrically corrected images.

Application of the Model

The steps performed in the application of the IDCTAB can be summarized in the following steps:

- First, the offset in position of each pixel relative to the reference position is computed.

- These deltas are used to compute the undistorted positions in arcseconds with the coefficients of the polynomial fit given in the IDCTAB reference file.

- The undistorted positions are divided by the desired output plate scale, where the default value is taken from the SCALE column in the IDCTAB unless specified by the user.

- Finally, both chips (WFC only) are combined into a single image by adding the appropriate offset to each pixel to get the final position.

The fit for each chip ensures that the distorted and corrected Y axes for a chip remain at the same angle with respect to the telescope V3 axis. This allows each chip to have a different rotation relative to the other, though this effect has been empirically shown to be less than 0.001 degree.

Format of the IDCTAB Reference file

The IDCTAB is stored as a FITS binary table (calibration file with '_idc.fits' extension), where the primary header contains keywords INSTRUME and DETECTOR which define the model appropriate to a specific instrument and detector. The NORDER keyword gives the order of the polynomial stored in the table. While the original laboratory data provided a cubic solution to the distortion, more accurate in-flight data now supports a fourth-order solution (NORDER=4).

The first extension of the IDCTAB contains the geometric distortion solution, where the table columns are listed in Table 4.2. The table contains distinct rows for each detector chip and filter combination specified by the DETCHIP and FILTER columns. If the DIRECTION column is 'INVERSE' rather than 'FORWARD', the coefficients specify the conversion from undistorted to distorted pixel positions. The SCALE keyword gives the default pixel scale of the drizzled image, where the default output scale for drizzled WFC images is 0.05 arcsec/pixel and for HRC/SBC images is 0.025 arcsec/pixel.

Table 4.2: Columns in the IDCTAB

Column	Description
DETCHIP	ID of chip/detector used for observation
DIRECTION	Application direction of coefficients (FORWARD,INVERSE)
FILTER1	Name of filter in filter wheel 1
FILTER2	Name of filter in filter wheel 2
XSIZE	Raw image size in X direction (pixels)
YSIZE	Raw image size in Y direction (pixels)
XREF	X position of reference pixel
YREF	Y position of reference pixel
V2REF	V2 position of reference point (arcsec)
V3REF	V3 position of reference point (arcsec)
SCALE	Scale of square corrected pixel (arcsec/pix)
CX10, CX11,...	Distortion Coefficients for X position
CY10, CY11,...	Distortion Coefficients for Y position

The distortion coefficients provide a correction for each pixel relative to the detector chip reference point specified in the XREF and YREF columns. The relationship between each chip must be factored in separately. HST relies on a V2/V3 coordinate system to define each detector's position (in arcseconds) relative to the center of HST's field of view. The V2REF and V3REF columns contain the reference position of each detector in the V2/V3 system. The absolute position is secondary to the relative position of this reference point, because only the offset between V2REF/V3REF positions is used for combining multiple chips into a single output image. The XREF/YREF and V2REF/V3REF positions represent the same point in different coordinate systems.

The columns CX*ij* and CY*ij* contain the coefficients of the polynomial fit for each chip and convert an input pixel position to an undistorted

position. The column names contain the indices for the coefficients from the polynomial fit, where CX11 and CY11 correspond to the a_{11} and b_{11} coefficients. For subarray data, the input pixel position must be corrected for the subarray offset in order to put it in the coordinate frame of the full chip.

The presently implemented IDC tables include separate rows of coefficients for all the measured filter combinations. Uncalibrated combinations default to the F475W filter solution with scale allowances made in the case of polarizing filters.

The solutions described by the IDC tables are correct to about 0.1 pixels. However the distortion can be specified to 0.01 pixels by the inclusion of extra corrections (such as those shown in Figure 4.6) which can not be expressed by the fourth order polynomials. These further corrections are available in the form of two images for each detector for the x and y directions. We have given these images the extension DXY. These correction images are planned to become part of the pipeline process in late 2004. For filter combinations for which these images do not exist, the final DXY image correction step will be omitted. However note that for most purposes the polynomial solutions provide adequate precision.

4.2 Combination of Dithered Data

4.2.1 Introduction to Dithering

Many ACS datasets are dithered - they have shifts between successive images which may be small or large. Dithering is a well established technique for HST imaging and has many advantages. For a general introduction to the concepts behind dithering, please consult the HST Dither Handbook (Koekemoer et al. 2002).

The primary motivations for dithering are:

- To allow the detection and flagging of image artifacts. Cosmic rays may be detected by taking multiple CR-SPLIT images at the same position, but other artifacts such as hot pixels require exposures with small dithers if they are to be detected and flagged. In addition, small-scale dithers are useful for filling in regions not covered by the detector, like the WFC chip gap or the HRC occulting finger.

- To improve the sampling of the PSF. In many ACS observation modes the pixel grid spacing is inadequate to fully sample the optical PSF falling on it. In these cases sub-pixel dithering will allow some of this lost information to be regained and the subsequent reconstruction of an improved final image.

- Larger dithers are needed to create mosaics of larger regions of sky or to cover the gap between the two chips of the ACS/WFC.

Because the non-linear component of the ACS optical distortion is large, dither shifts will vary across the field of view. This must be taken into account when planning observations. For more information on ACS dither patterns, see the ACS Patterns section of the HST Phase II proposal instructions and additional variations on the ACS "dither" webpage:
http://www.stsci.edu/hst/acs/proposing/dither.

4.2.2 Introduction to Drizzling

The methods for combining dithered ACS data which are described in this chapter are all based on the drizzle method (Fruchter & Hook, 2002 PASP, 114, 144) and are described in detail in the *HST Dither Handbook* (Koekemoer et al. 2002). Here we give a brief summary of the method and refer the reader to the aforementioned paper for further details and quantitative information about its properties.

Drizzle maps each pixel of the input image onto pixels in a sub-sampled output image, taking into account shifts and rotations between images and the optical distortion of the camera. However, in order to avoid re-convolving the image with the large pixel of the camera, the user may shrink the pixel before it is averaged into the output image. This shrinking factor is known as the "pixfrac" and is normally between 0.0 (corresponding to interlacing input pixels onto single output ones) and 1.0 (which corresponds to shift-and-add). The degree of sub-sampling (the size of the output pixels relative to that of the original ones) is also a user controllable parameter called "scale" which is typically in the range 1-0.5 for the ACS.

This algorithm has a number of advantages. It preserves both absolute surface and point-source photometry, hence flux can be measured using an aperture whose size is independent of position on the chip. And because the method anticipates that a given output pixel may receive no information from a given input pixel, missing data (due for instance to cosmic rays or detector defects) do not cause a substantial problem, so long as there are enough dithered images to fill in the gaps caused by these zero-weight input pixels. Finally, the linear weighting scheme is statistically optimum when inverse variance maps are used as the input weights.

The original drizzle method mapped a shrunken version of each input pixel onto the output and distributed weight according to the overlaps with the output pixel grid. For the case of multiple sub-pixel dithers of significantly undersampled images, this was close to optimum. However, for some purposes this "kernel" is not ideal and the latest implementations of drizzle allow the user to select from other options. For example, if the original method is applied to a single image, where pixfrac and scale are

both normally unity, the method becomes equivalent to bilinear interpolation. This results in an output image having strongly correlated noise and also degraded resolution. For images from the ACS in the red, even the WFC channel is close to critically sampled, rather than strongly undersampled, and it would appear to be useful to consider more optimal interpolators as drizzle kernel options. Some of these are now implemented in **drizzle** and may also be used via the higher-level scripts described. The **drizzle** help file lists all the options. The "Lanczos3" kernel is particularly recommended as it minimizes apparent noise correlation in the output image and maximizes resolution. It is, however, significantly slower.

The basic drizzle method is implemented as an IRAF task within the STSDAS **dither** package which is invoked by **PyDrizzle** or **MultiDrizzle**.

4.2.3 Overview of Drizzling Software Tools

The tools for the combination of dithered HST images have evolved through the years, starting with those developed for the Hubble Deep Field in late 1995. First came low level tools which required considerable effort to use. Subsequently higher level wrapper scripts, written in the Python language, have been developed to ease such combination. Some of these capacities are now available within the pipeline. The main three levels applicable to current ACS imaging are as follows:

- The low level tools for mapping pixels from input to output images and support tasks for measuring shifts, detecting cosmic ray events, etc. are available within the dither package of STSDAS. These include **drizzle** itself and **blot** (reverse **drizzle**). The user will now probably prefer to access these tools via one of the wrappers rather than use them directly.

- **PyDrizzle** and associated Python scripts are layered on these and automate some of the processing. They convert the IDCTAB information into a form suitable for drizzle, create appropriate association tables and drizzle images into an appropriate output frame. These functions are implemented in the pipeline and are described in detail below. The user will again probably not use these tools directly. **PyDrizzle** is available within the STSDAS system and requires the PyRAF environment.

- The third level is the **MultiDrizzle** script which is again layered on top of the above tools. It aims to provide automated image combination, including the detection and flagging of cosmic rays. Although **MultiDrizzle** can be run using its default settings to do a good combination for many typical ACS dithered datasets, it also allows the

user considerable control and flexibility. **MultiDrizzle** is available as a beta-test version but will be implemented in the STSDAS system in 2004. Like **PyDrizzle**, it requires the PyRAF environment.

We now briefly describe the pipeline software **PyDrizzle** and several associated tools. A detailed description of **MultiDrizzle** follows, as this is the recommended interactive tool for drizzling ACS images.

4.3 PyDrizzle

4.3.1 Application by the Pipeline

The ACS calibration pipeline currently consists of 2 processing stages, **CALACS** and **PyDrizzle**, which are run one after the other. Standard **CALACS** processing includes bias and dark subtraction, removal of the overscan regions, and flat fielding. It will produce, among other products, a calibrated image with the FLT suffix for each input file. The **CALACS** portion of the pipeline will calibrate and combine CR-SPLIT images. It will not, however, correct for geometric distortion or combine images from multiple dither positions.

Geometric distortion is removed for all images (dithered or not) in the second stage of processing via **PyDrizzle**. This software understands the ACS association tables, allowing the pipeline to combine associated dithered observations into a single image. These images will still contain cosmic-rays unless CR-SPLIT observations were specified within the dither pattern. For WFC observations, both chips are combined into a single extension. Finally, **PyDrizzle** converts the data to units of count rate (*electrons/sec*).

The IRAF **dither** package has provided users the basic tools for linear image combination via the task **drizzle**. This task can produce mosaics, correct for geometric distortion, and even remove cosmic rays, creating a single, undistorted output image with uniform photometric and astrometric properties. A significant amount of work, however, is required to determine the necessary input parameters, causing **drizzle** to be fairly user-intensive. These steps include:

- requiring the offsets in x and y from the reference image as input.
- requiring the output image size and pixel scale as input.
- only working with a single image extension at a time.
- requiring the rotation of each image extension as input.

These requirements made it impossible to use **drizzle** in a pipeline environment to automatically process images.

PyDrizzle was written to automate the use of **drizzle** in the pipeline environment. It can also simplify using **drizzle** for manual image

reprocessing. Pipeline use relies on default parameter values which are taken from the IDCTAB reference file and include the default plate scale and the distortion model. This file provides information about the position of each ACS detector on the sky, allowing **PyDrizzle** to combine images from multiple detectors at once. **PyDrizzle**, like **drizzle**, works with count rates, allowing images with differing exposure times to be combined easily. Regardless of whether observations were taken as a single image or a set of dithered exposures, **PyDrizzle** produces images which are both astrometrically and photometrically accurate.

Processing comments are recorded in the trailer file for the DRZ image, including which version of **drizzle** was used, what parameters were used, and which images were combined (if dithered). The same default parameters are used for all observations in the pipeline. These parameters were chosen to avoid introducing any scale changes, shifts, or rotations relative to the original pointing, aside from those corrections incorporated in the distortion model itself. While the pipeline products will be properly corrected for distortion, the combination of dithered observations may not be ideal, and small pointing errors or defects, such as hot pixels or cosmic-rays, may still be present.

4.3.2 Association Tables

Before the **drizzle** task could be used in an automated manner, two obstacles had to be overcome: calculating the required input parameters and working with a wide range of input image formats. **PyDrizzle** relies on several features to address these problems:

- An association table will relate multiple input images to an output product.
- Python "structures" will be created in memory for each input image and for the output field.
- Output field parameters (RA/Dec of center, pixel size, field of view) will be specified using a 'SkyField' function.

The calibration pipeline uses association tables to define how multiple images, taken at different pointings, relate to a single output image. **PyDrizzle** relies on these association tables to automatically combine and correct ACS dithered observations or any other related set of observations.

Each input/output exposure within **PyDrizzle** knows about:

- its own world coordinate system,
- the distortion model for its detector,
- the range of output pixels covered by the image, and
- any instrument specific characteristics.

The **PyDrizzle** design provides the flexibility to support multiple instruments simultaneously, and to combine observations from several instruments into a single output field with arbitrary specification. The virtual output image is built using the World Coordinate System (WCS) information from each input image and the distortion model. This provides all the information necessary to calculate the input parameters for **drizzle**. These parameters are then passed to **drizzle** which actually performs the image combination.

4.3.3 PyDrizzle Data Products

The output from **PyDrizzle** is a single multi-extension FITS file with the suffix '_drz.fits'. The first extension contains the science (SCI) image which is corrected for distortion and which is dither-combined (or mosaiced), if applicable. The drizzled SCI image is in units of *electrons per second*. All pixels have equal area on the sky and equal photometric normalization across the field of view, giving an image which is both photometrically and astrometrically accurate for both point and extended sources. The dimensions of the output image will be computed on-the-fly by **PyDrizzle** and the default output plate scale will be read from the IDCTAB. These parameters, however, may be chosen by the user during reprocessing to best suit the actual data.

The second extension of the output image contains the weight (WHT) image. It gives the relative weight of the output pixels and can be considered an effective exposure time map. The data quality map created by **CALACS** is used by **PyDrizzle** in creating the weight image.

The third extension of the **PyDrizzle** output image contains the so-called context (CTX) image which encodes information about which input image contributes to a specific output pixel. This is done using a bitmask for each output pixel, where 'bit set' means that the image, in the order it was combined, contributed with non-zero weight to that output pixel. The context image starts as a single 32-bit integer image but is extended as a cube with additional 32-bit deep planes as required to handle all the input images.

4.4 Reprocessing with PyDrizzle

The goal of the ACS pipeline is to provide data calibrated to a level suitable for initial evaluation and analysis for all users. Observers frequently 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 institution. There are several occasions when off-line interactive processing with **PyDrizzle** is required:

- Observations which have been manually reprocessed with **CALACS** will still require a correction for geometric distortion.

- Users may wish to modify the default **PyDrizzle** parameters, specifying an alternate output image orientation, pixel "shrinking factor" pixfrac, or pixel scale.

- When images have been taken across multiple visits, the observations are not automatically associated. Creating custom association tables will allow data from multiple visits to be combined.

- For optimal image registration, the user may wish to fine-tune the WCS information by supplying additional shifts and/or rotations.

- During pipeline processing, **PyDrizzle** makes assumptions about which DQ flags should be considered bad. These pixels will be excluded when creating the drizzled product. Users may wish to select which flagged pixels should actually be considered good and included in the drizzled product.

The **PyDrizzle** software uses the Python language and can only be executed in the PyRAF environment. PyRAF, developed at STScI, allows users to run tasks using the standard IRAF parameter-based interface or by

using separate Python programs. The PyRAF interface and the most recent version of **PyDrizzle** can be downloaded from:

http://www.stsci.edu/resources/software_hardware/stsdas/pyd
rizzle.

The calibrated products from **CALACS** and the corresponding input to **PyDrizzle** are described for various observing modes in Table 3.1. For example, a single image will produce an FLT file, a RPT-OBS exposure will produce an SFL file, a CR-SPLIT exposure will produce a CRJ file, and a DITH-PATTERN will produce a series of FLT, CRJ, or SFL files at *each* dither position. **PyDrizzle** may be executed using any of the calibrated data products or the association table itself. When executed on an association table, **PyDrizzle** looks only for the existing product file and geometrically corrects that image. **PyDrizzle** will create a single drizzled image with the DRZ suffix for each calibration product. In the case of dithered observations which are part of a pattern, the individual FLT files will be combined to create a single DRZ image.

PyDrizzle was not developed with the capability of detecting cosmic-rays, so unless cosmic-ray rejection is performed in the first stage by **CALACS**, the calibrated, drizzled product will still contain cosmic-rays. A new package, **MultiDrizzle**, addresses this deficiency and is discussed in detail in Section 4.5.

4.4.1 Specifying the 'Bits' Parameter

The data quality flags which were set during **CALACS** processing can be used as bit masks when drizzling, and the user may specify which bit values should actually be considered 'good' and included in image combination. This is done via the parameter 'bits'. Any pixel flagged otherwise will be considered 'bad' and will be given zero weight. If a pixel is flagged as bad but has non-zero weight in any other image, the pixel will be replaced with the value from that image during image combination. Otherwise, the pixel will be replaced with the 'fill value'. The default 'fill value' used by **PyDrizzle** during pipeline processing is zero. If the 'fill value' is set to INDEF during reprocessing with **PyDrizzle** or **MultiDrizzle**, and there are no pixels with non-zero weight, the pixel will retain its original value.

The flags for the DQ array were presented in Chapter 3 and are summarized in Table 3.4. The bits value used during pipeline processing is 8578 which is the sum of 8192 (corrected cosmic rays) + 256 (saturated data) + 128 (bad pixel in bias file) + 2 (pixel rejected in second iteration of cr-rejection). Note that these pixels were flagged during **CALACS** processing as "bad" or "suspect", but may have been corrected in later processing steps. For example, cosmic rays are flagged 8192, but these pixels are replaced by unaffected pixels from other images in the

CR-SPLIT association. The bits parameter indicates which "suspect" pixels to keep. When the total counts in a given pixel have exceeded the full well, for example, the DQ flag is 256. However, testing shows that counts still accumulate in a highly linear manner and that photometry of saturated stars is quite practical if using a gain that samples the full well depth.

The default bits value in the **PyDrizzle** and **MultiDrizzle** software is zero, but should be reset by the user prior to reprocessing. The value chosen for this parameter is completely up to the user and should be selected based on the specific calibration needs. For HRC processing, the user may want to add 8 (masked by aperture) to the 8578 value so that pixels which lie behind the occulting finger mask are not given zero weight. (In Figure 4.9 of Section 4.4.4 "PyDrizzle Examples", bit 8 was not included in the bits parameter prior to drizzling. Pixels behind the occulting finger were therefore given zero weight and replaced with the default fill value=zero.) Of course, the photometry in this region will be inaccurate due to improper flat fielding, so this bit would likely only be set for aesthetic reasons.

4.4.2 Creating Custom Association Tables

Association tables for dithered, REPEAT-OBS, or CR-SPLIT observations are generated by the **CALACS** pipeline. These tables keep track of the input exposure filenames and the output product filenames. Some types of observations, however, will not have association tables. Others will have multiple tables from different visits which need to be combined into one. In the following discussion, we present the methodology for creating custom association tables, either by merging association tables or creating them from scratch. Users also have the capability to manually edit association tables to include any known image offsets. This can be done using the TTOOLS task **tedit** in IRAF, where the user adds the columns XOFFSET, YOFFSET and/or ROTATION to the association table. Alternately, an ascii file with the known image offsets may be specified and the association automatically updated.

Merging Association Tables

Observing programs which cover a large portion of the sky will generally be made up of multiple pointings. Each of these pointings may be dithered or split for cosmic-ray rejection and will possess their own association table. In order for **PyDrizzle** to produce a single mosaic product, a composite association table must be built from the separate association tables. Users can easily create this composite by merging the individual tables using the TTOOLS task **tmerge** with "option = append".

The default product rootname will be taken from the first listed DTH-PROD in the composite association table. This rootname can be modified by the user (with **tedit**) to suit the desired naming convention.

Generally, this rootname should match the rootname of the composite association table. A detailed example of merging association tables is given in step 1 of Example 4 from Section 3.5.2.

Creating Association Tables from Scratch

Some images will not have association tables. These may be single exposures of a given target which were taken at different times and orientations. They may even be exposures taken with different instruments. Observations taken using POS-TARG offsets, which are specified manually in the proposal, will not have association tables either.

A Python tool called **buildAsn** has been developed to create association tables which may then be used by **PyDrizzle**. This task can be run under PyRAF using an IRAF parameter-based "epar" interface or directly using a Python command-line interface. Since it resides alongside the **PyDrizzle** code, no extra setup is required to use this task.

A built-in help facility provides details on additional parameters which can be used with **buildAsn** when using the Python syntax. (For those needing a step-by-step introduction to starting PyRAF, see the examples in Section 4.4.4. Building an association table will be part of the **PyDrizzle** Examples.) It can be accessed using:

```
pyraf> from pydrizzle import buildasn
pyraf> buildasn.help()
```

The following illustrates how to use **buildAsn** to create an association table from all calibrated input files with the FLT suffix found in a single directory. Within PyRAF and using the IRAF syntax, we type:

```
pyraf> buildasn mymosaic suffix='flt'
```

Alternately, the same result may be obtained using the Python syntax:

```
pyraf> buildasn.buildAsnTable('mymosaic',suffix='flt')
```

In this example, the association table will have the name 'mymosaic_asn.fits', the product name will be 'mymosaic_drz.fits', and all files with the FLT suffix found in the working directory will be included in the association. To specify a subset of images within a given directory, the user may specify the 'suffix' parameter to be a filelist ('@filename'), a wild-card list ('*8cd*flt*'), or any portion of the filename ('crj' or 'f555w').

If user determined offsets are available, **buildAsn** has the capability of incorporating them into the association table. These offsets (XOFFSET,

YOFFSET, and/or ROTATION) are specified by the file listed in the 'shiftfile' parameter. Using the IRAF syntax:

```
pyraf> buildasn mymosaic suffix='flt' shiftfile='shift.txt'
```

This option allows users to fine-tune the final image combination by providing corrections to the header WCS information. For more information on the format of user-defined shifts, see Section 4.4.3. For an example illustrating the steps required to apply a shiftfile to an existing association, see Example 3 in Section 4.4.4.

4.4.3 Applying User-defined Shifts

The ability of **PyDrizzle** to properly align dithered observations relies on the accuracy of the header WCS information. Unfortunately, there are times when those WCS values may be slightly inaccurate, resulting in a mis-alignment of the final drizzled product. For example, one might wish to combine observations taken in different visits. If the visits were separated by a substantial time interval, then the telescope probably used different guide stars. Due to limitations in the intrinsic accuracy of the Guide Star Catalog, this can lead to WCS coordinates that differ by up to 1-2" between visits. Also when all exposures that one wishes to combine were taken in a single visit, it might still be important to correct for relative shifts between images. Even in two-star fine-guiding mode, thermal effects on the telescope can cause slow drifts that can build up to ~50 mas over a few orbits (or in rarer occasions within a single orbit). This corresponds approximately to 1 WFC pixel or 2 HRC pixels. Such shifts between exposures can degrade image quality and corrupt cosmic-ray rejection when left uncorrected.

By deriving shifts based on the position of objects in the data, the user may refine the shifts computed from the WCS header information to create a precisely-aligned drizzled product. Many observers have developed their own methods of comparing images and computing offsets. The conventions described here should support the majority of users, making it a simple matter to incorporate shift computations into **MultiDrizzle** or **PyDrizzle.**

Delta Shifts

A *delta shift* is defined as the residual shift required to align sources, *after* applying the offsets implied by the header WCS. Delta shifts may be determined by separately drizzling images onto a common WCS frame with the same central RA/Dec. This is performed in the '*driz_separate*' step of **MultiDrizzle** (see section 4.5 for more information) or by setting the **PyDrizzle** parameter 'single=yes'. Object lists derived for each drizzled image may then be matched and fit to derive a single delta shift for each image.

Delta shifts are defined to be in the 'input' frame of reference when the distortion-corrected, drizzled image has the same orientation and scale as the distorted image. Shifts will be in the 'output' frame of reference if any rotation or scaling was applied while drizzling. When the shifts are given in the 'output' frame of reference, **PyDrizzle** requires the specification of the reference image to properly account for any orientation and scale changes.

Absolute Shifts

An *absolute shift*, on the other hand, is the total shift between two images and includes offsets implied by the header WCS. Absolute shifts may be determined by separately drizzling images to their own unique WCS frame. **PyDrizzle** will use the header information to optimally place the drizzled image in the frame.

To derive absolute shifts in the 'input' frame of reference, each distorted image is drizzled separately, and **PyDrizzle** uses the unique WCS information from each frame to choose the central RA/Dec, the image orientation and final image size. No additional rotation or scaling is applied while drizzling. (Alternately, the user may catalogue sources in each distorted image and then apply the distortion model to the X/Y positions. However, this requires tasks to transform coordinates from distorted to undistorted space. These tasks are still being developed for release in the STSDAS dither package.)

Absolute shifts in the 'output' frame of reference may be computed by separately drizzling each distorted image, with the same rotation and scaling, onto a unique WCS frame. Target lists in each drizzled image may then be matched to derive the absolute shifts.

Figure 4.7: Delta Shifts in the 'input' and 'output' frame of reference.
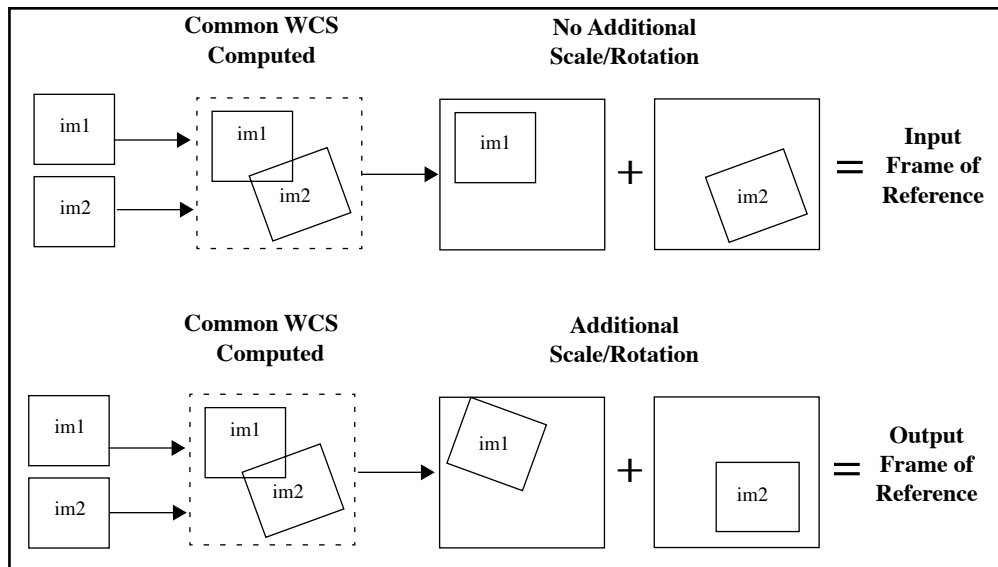
Figure 4.8: Absolute Shifts in the 'input' and 'output' frame of reference. The direction of north is indicated by the bold arrows.
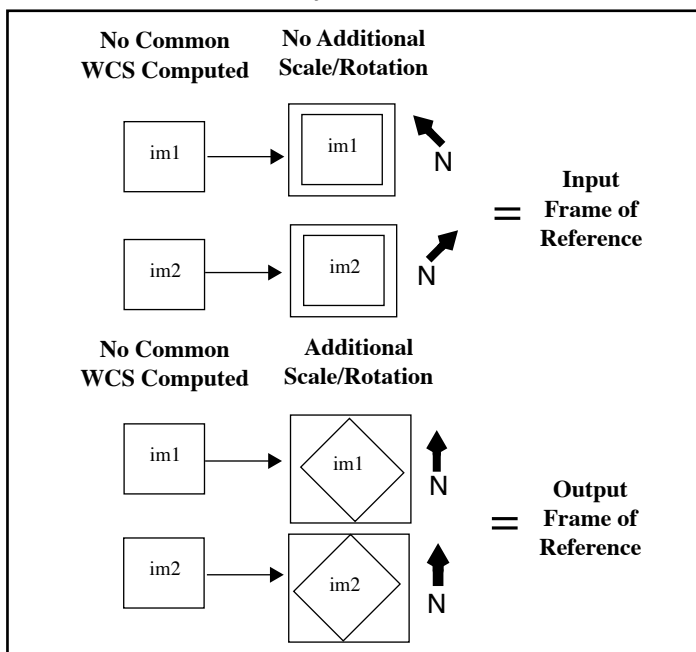
Shift Units

PyDrizzle is capable of interpreting shifts in units of *pixels* or *arcseconds*. A *pixel shift* is the difference in the X/Y pixel position of a source in a given image compared to the reference image. The shifts are interpreted with respect to each image's X/Y pixel reference position and require the specification of 'input' or 'output' frame of reference.

An *arcsecond shift* is the difference in the RA/Dec coordinates of a source in a given image compared to the reference image. Object coordinates must be derived *after* correcting for distortion. These shifts are *not* angular offsets on the sky, rather they are changes in the RA/Dec of the reference pixel and do not require any 'cos(Dec)' adjustment. Arcsecond shifts are a natural result of using the RA/Dec coordinates of image sources. They eliminate any ambiguity with the chosen reference frame since RA/Dec coordinates are fixed on the sky. Arcsecond shifts are interpreted as offsets in the direction of increasing RA and Dec.

Sign Conventions

The shifts are defined in terms of how the targets in the image need to move, rather than on how the telescope would need to move, to allow precise registration. **PyDrizzle** assumes that the shifts were computed as 'image - reference'. For example, when using **geomap** to cross-correlate two source lists, the input image coordinates will be in the first two columns and the reference image coordinates will be in the last two columns.

Use of Shifts in Association Tables

The *shift file* is an ASCII file containing the user-computed shifts for a set of images and is used for updating the offsets in the association table. Observers may use either the ASN table delivered by the pipeline or create/modify their own table as described in Section 4.4.2. Both **PyDrizzle** and **MultiDrizzle** use the **buildasn** task to incorporate information from the shift file into the association table. The *shift file* uses the following format:

```
# units: pixels              Values: pixels (default), arcseconds
# frame: input               Values: input (default), output
# form: absolute             Values: absolute (default), delta
# reference: <filename>      Only required when 'frame = output'
filename1  xshift1 yshift1  [rotation1]
filename2  xshift2 yshift2  [rotation2]
```

The first three lines specify the units, the frame of reference, and the form of the shifts, and only those lines with non-default values are necessary. While the '#' symbol is required, the 'Values:' comments are not. When shifts are given in the 'output' frame of reference, it is necessary to specify the name of the reference image used.

The following is an example of a shift file in units of pixels, measured in the 'input' frame of reference, with the form specified as 'delta' shifts. An x and y shift and a rotation (in degrees) is specified for each dataset. If only a simple shift is required, then the rotation column need not be specified.

```
#units: pixels
#frame: input
#form: delta
j8c0b1skq_flt  -2.4 -1.2 -0.002
j8c0b1snq_flt  -4.3 -2.3  0.001
j8c0b1sqq_flt  -6.9 -3.5  0.003
```

Using this shiftfile as input, the **buildasn** task will update the ASN table as indicated in Table 4.3.

Table 4.3: Association table updated with user-defined shifts.

MEMNAME	MEMTYPE	XOFF-SET	YOFF-SET	XDELTA	YDELTA	ROTATION
j8c0b1skq	EXP-DTH	0	0	-2.4	-1.2	-0.002
j8c0b1smq	EXP-DTH	0	0	-4.3	-2.3	0.001
j8c0b1sqq	EXP-DTH	0	0	-6.9	-3.5	0.003

The shift file can contain shifts for files other than those contained in the ASN table. In that case, only those entries found in the ASN table will be used to update the table. Thus, a single shift file can be generated for a whole set of observations which are represented by multiple association tables.

Association (ASN) tables can support either delta shifts or absolute shifts in units of pixels or arcseconds. Absolute shifts are stored in columns 'XOFFSET' and 'YOFFSET', while delta shifts are stored in columns 'XDELTA' and 'YDELTA'. The 'ROTATION' column is also available if a simple shift does not produce ideal registration and indicates how much should be added to the image's ORIENTAT to get proper alignment.

If *absolute offsets* are provided in the shift file, **PyDrizzle** will populate the OFFSET columns in the association and fill the DELTA values with zeroes. Similarly, if *delta shifts* are provided, those values will be used to populate the DELTA columns and the OFFSET columns will be zero. If shifts are given in both the OFFSET and DELTA columns, the OFFSET column will be applied and the DELTA columns will be ignored. This convention used by **buildasn** eliminates any ambiguity as to which values are applied to the data. Editing the ASN table directly will then allow the user to further update the offsets using either *delta* or *absolute* shifts.

4.4.4 PyDrizzle Examples

These examples further describe the steps required for pipeline reprocessing and serve as a continuation of the **CALACS** examples described in Section 3.5.2. We advise the reader to work through the **CALACS** examples first to develop a solid understanding of the calibrated data products and their potential limitations.

Example 1: Drizzling a Single Exposure

The following is a continuation of Example 1 from Section 3.5.2 where manual recalibration of a single raw exposure with **CALACS** produced a single image with the FLT suffix.

1. In the Unix environment, the 'jref' environment must be established prior to loading PyRAF. This directory indicates the location of the geometric distortion reference file (IDCTAB) specified in the image header. For more information, see "Setting up jref" in Section 3.5.1. If running IRAF and PyRAF simultaneously, 'jref' will need to be defined in each xterm window before loading the software. Alternately, this path may be initialized in the user's .setenv file. PyRAF

should then be started in its own xterm window and the **stsdas**, **analysis**, and **dither** packages loaded.

```
unix> setenv jref /mydisk/myjref/
unix> pyraf
pyraf> stsdas
pyraf> analysis
pyraf> dither
```

2. Next, we run **PyDrizzle** on the calibrated FLT image (j8bt07oyq_flt.fits). By default, the drizzled output product will be named j8bt07oyq_drz.fits. To specify an output file which is different than the default, the user may set the 'outfile' parameter to 'myfile_drz.fits', for example. Section 4.4.1 defined the 'bits' parameter; here we adopt the recommended value of 8578. The reader may also wish to add 8 to this value to see the effect of not dropping pixels behind the occulting finger. Prior to running, we recommended clearing any preset values using the 'unlearn' command.

```
pyraf> unlearn pydrizzle
pyraf> pydrizzle j8bt07oyq_flt.fits bits=8578
```

The distortion corrected HRC science image is shown in Figure 4.9. The image on the left is the flat fielded FLT product from **CALACS** and still contains distortion. The image on the right is the calibrated, distortion-corrected DRZ product from **PyDrizzle** using 'bits=8578'.

Figure 4.9: The effect of applying the distortion correction to a single HRC image. The image on the left is the flat-fielded FLT product from **CALACS**. The image on the right is the distortion-corrected DRZ product from **PyDrizzle**.
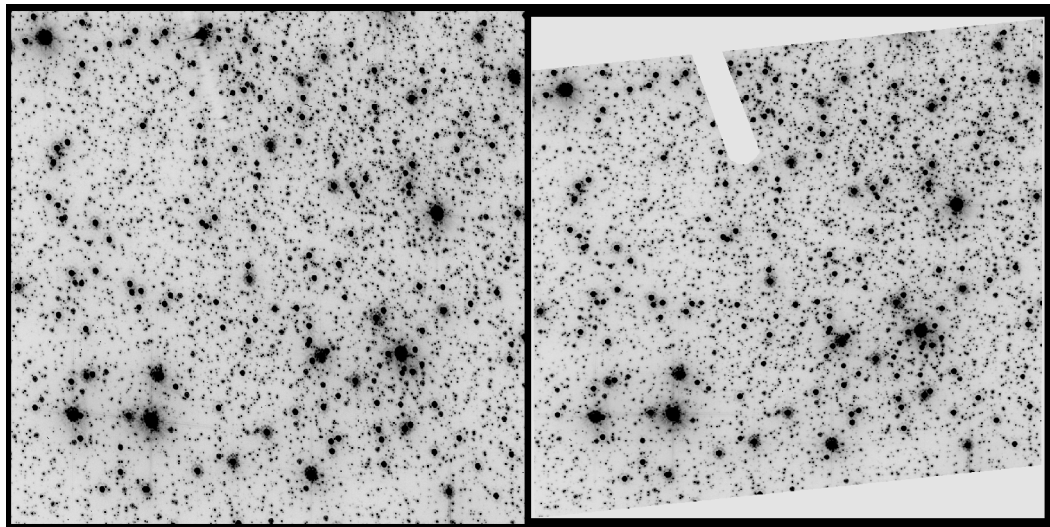
Example 2: Drizzling Images taken as part of a Dither Pattern

The following is a continuation of Example 5 from Section 3.5.2 where manual recalibration of a 2-point line dither pattern produced two FLT images (j8e654c0q_flt.fits, j8e654c4q_flt.fits).

1. Using **tprint**, we can view the contents of the image association, including the rootnames of the individual dithered images and the default rootname of the drizzled product.

```
pyraf> tprint j8e654010_asn.fits

# MEMNAME       MEMTYPE          MEMPRSNT

j8e654c0q       EXP-DTH          yes

j8e654c4q       EXP-DTH          yes

j8e654011       PROD-DTH         yes
```

2. Next we run **PyDrizzle** on the image association. The two FLT images are combined to create a single DRZ image called 'j8e654011_drz.fits'.

```
pyraf> pydrizzle j8e654010_asn.fits bits=8578
```

Note that **PyDrizzle** is not designed to remove cosmic rays from these images. This step requires the new software **MultiDrizzle**, which is described in Section 4.5.

Example 3a: Improving the Image Registration - 'Delta Shifts' in the 'Input Frame'

For observations with large dithers or for images which were taken at nominally the same pointing but during different visits, users may wish to fine-tune the image registration, improving on the world coordinate system in the image header. In this case, the images must be separately drizzled and the position of stars cross-correlated. The derived shifts may then be used to update the image association prior to drizzling the final product. More generally, if the drizzle combined product of multiple exposures have PSFs that are not round or as sharp as expected, a likely culprit is the need for improved shifts.

The following example uses WFC observations from separate visits in program 9018: visit 61, exposure 02 (j8c061020_asn.fits) on 19April 2002 and visit C1, exposure 01 (j8c0c1010_asn.fits) on 09May2002. Each observation is CR-SPLIT with an exposure time of 45 seconds in F606W. The images are at the same pointing, use the same guide stars, and are at the same ORIENT. We describe the steps required to register the calibrated CRJ products from separate visits. However, the individual FLT files could be alternatively used if improved registration of each CRJ image is desired, although use of FLT images would result in cosmic-ray retention.

1. First, we use the CRJ images from visits 61 and c1 to create a new association table with the suffix 'v61c1'. We advise using the **tprint** task to verify the contents of the association.

```
pyraf> import buildasn
pyraf> buildasn f606w_v61c1 suffix='crj.fits'
pyraf> tprint f606w_v61c1_asn.fits
# MEMNAME     MEMTYPE   MEMPRSNT  XOFF YOFF XDELTA YDELTA ROTATION
j8c061021    EXP-DTH    yes        0.   0.    0.     0.     0.
j8c0c1011    EXP-DTH    yes        0.   0.    0.     0.     0.
f606w_v61c1 PROD-DTH    no         0.   0.    0.     0.     0.
```

2. Next, we separately drizzle the images onto a common output WCS. This is achieved by setting the **PyDrizzle** parameter 'single'=yes. Any differences in the aperture RA and Dec, any POS-TARG offsets, or any orientation differences will be automatically accounted for by **PyDrizzle** which uses the header WCS to transform the images. The drizzled CRJ products are appended with the suffix '_single_sci.fits'.

```
pyraf> pydrizzle f606w_v61c1_asn.fits bits=8578 single=yes
```

3. Using point sources in each image, we create and cross correlate coordinate lists from each visit. For example, we used **daofind** to create star lists for each DRZ product separately. Then we used **xyxymatch** to match the coordinate lists within a user specified tolerance. Finally, we use **geomap** with the matched coordinate list to compute the transformation required to map the reference coordinate system to the input coordinate system, allowing for a shift and/or rotation. We create the following "shiftfile", a simple ascii table called 'shift.txt', using the derived offsets. The shifts were derived in units of *pixels* and are in the 'input' frame of reference since no additional rotation or scale terms were introduced while drizzling. The derived shifts using the above methods are *delta* shifts because the images were drizzled onto a common WCS.

```
#units: pixels
#frame: input
#form:  delta
j8c061021_crj    0.00    0.00  0.000
j8c0c1011_crj   24.46  -56.54  0.001
```

4. Next, we delete and rebuild the previous association table, this time specifying the derived offsets via the shiftfile. To verify that the association has been correctly populated, we use the **tprint** command.

Note that the *delta* shifts appear in the DELTA columns in the association table.

```
pyraf> delete f606w_v61c1_asn.fits
pyraf> buildasn f606w_v61c1 suffix='crj.fits'
shiftfile='shift.txt'
pyraf> tprint f606w_v61c1_asn.fits
# MEMNAME    MEMTYPE   MEMPRSNT  XOFF YOFF XDELTA YDELTA ROTATION
j8c061021   EXP-DTH    yes        0.  0.    0.      0.     0.
j8c0c1011   EXP-DTH    yes        0.  0.    24.46  -56.54  0.001
f606w_v61c1 PROD-DTH   no         0.  0.    0.      0.     0.
```

5. Finally we run **PyDrizzle** on the new image association. The two CRJ images are combined to create a single DRZ image called 'f606w_v61c1_drz.fits'.

```
pyraf> pydrizzle f606w_v61c1_asn.fits bits=8578 single=no
```

Example 3b: Improving the Image Registration - 'Absolute Shifts' in the 'Output Frame'

As a continuation of the previous example, we illustrate how to specify *absolute* shifts in the 'output' frame of reference.

1. First, the CRJ images from each visit are drizzled to separate output images. Absolute shifts are determined by separately drizzling images onto their own unique WCS frame. In this example, an additional rotation is applied to the drizzled images so that north is oriented up.

```
pyraf> pydrizzle j8c061021_crj.fits output=j8c06102X_drz.fits
bits=8578 rotate+ orient=0
pyraf> pydrizzle j8c0c1011_crj.fits output=j8c0c101X_drz.fits
bits=8578 rotate+ orient=0
```

2. As in the previous example, we create and cross correlate coordinates lists from each visit using point sources in each image. We create the following "shiftfile", a simple ascii table called 'shift1.txt', using the

derived offsets. Note that because we are using absolute shifts, we must specify the name of the reference image.

```
#units: pixels
#frame: output
#form:  absolute
#reference: j8c06102X_drz.fits
j8c061021_crj    0.00   0.00  0.000
j8c0c1011_crj  -41.64  45.51  0.001
```

3. Next, we build a new image association using the two CRJ files and the derived shift file.

```
pyraf> import buildasn
pyraf> buildasn f606w_v61c1rot suffix='crj.fits'
shiftfile='shift1.txt'
```

4. To verify that the association has been correctly populated, we use the **tprint** command. Notice that the *absolute* shifts have been correctly added to the OFFSET columns, not the DELTA columns.

```
pyraf> tprint f606w_v61c1rot_asn.fits
# MEMNAME      MEMTYPE  MEMPRSNT  XOFF   YOFF XDELTA YDELTA ROTATION
j8c061021     EXP-DTH    yes       0.    0.    0.    0.     0.
j8c0c1011     EXP-DTH    yes     -41.64 45.51  0.    0.     0.001
f606w_v61c1rot PROD-DTH  no        0.    0.    0.    0.     0.
```

5. Finally we run **PyDrizzle** on the new image association, specifying a rotation such that north is up. The two CRJ images are combined using the updated shift information to create a single DRZ image called 'f606w_v61c1rot_drz.fits'.

```
pyraf> pydrizzle f606w_v61c1rot_asn.fits bits=8578 rotate+
orient=0
```

4.5 MultiDrizzle

MultiDrizzle, developed in its original form by A. Koekemoer (Koekemoer et al. 2002, HST Calibration Workshop 2002, p.381), is a general purpose tool for drizzle combination of images. It provides a single-step interface to the complex suite of tasks in the STSDAS **dither** package. These tasks include, in order: initial image registration, creation of a cleaned median image, transformation back to the input image plane, creation of cosmic ray masks, and final drizzling. (For a general introduction, refer to the HST Dither Handbook.) The goal of **MultiDrizzle** is to provide a high-level task which has an extensive suite of user-adjustable parameters which, if left at their default values, will allow the task to perform all steps in a single operation with no user intervention. At the same time, the parameters allow the user a large amount of flexibility in controlling the relevant aspects of these steps, in case the default parameters are not sufficient for specific scientific applications.

The current version of **MultiDrizzle** uses the image header world coordinate system (WCS) to deduce the image-to-image offsets. However, user-supplied offsets may also be applied in a flexible manner. In the future, tasks to ease the determination of such shifts will be made available. **MultiDrizzle** processing with default parameters may eventually be part of pipeline processing, if a detailed assessment shows that this is feasible.

MultiDrizzle and **PyDrizzle** are both Python scripts which can only be run within the PyRAF environment. **MultiDrizzle** calls **PyDrizzle** to manage the images, to compute the **drizzle** parameters, and to control the operation of both the **drizzle** and **blot** tasks. While the latest release of STSDAS includes **PyDrizzle**, **MultiDrizzle** must be obtained separately (see below). It is recommended that the complete current STSDAS package is installed to avoid potential missing components or mismatched versions. STSDAS and PyRAF are available from:

http://www.stsci.edu/resources/software_hardware.

Because many of these components are evolving rapidly, subsequent updates will be available from the same site.

While the software has been well tested and is fairly stable, it is still periodically being improved. We therefore refer readers to the **MultiDrizzle** web site for the most up-to-date software:

http://stsdas.stsci.edu/pydrizzle/multidrizzle

and to the ACS Drizzling web site for additional documentation:

http://www.stsci.edu/hst/acs/analysis/drizzle.

The following sections describe in detail the tasks performed by **MultiDrizzle** and give an example of its use. A help document describing the numerous **MultiDrizzle** parameters can be accessed by typing 'help

multidrizzle' from within PyRAF. HST Dither Handbook is currently under revision to include documentation for **MultiDrizzle** and several detailed ACS drizzling examples. Version 3.0 is expected to be released by mid-2004.

4.5.1 Reprocessing with MultiDrizzle

In Section 4.4, we described several occasions when off-line interactive processing with **PyDrizzle** is required, for example, when images are taken across multiple visits, when fine-tuned image registration is required, or when the user wishes to specify which DQ flags should be considered *good*. **MultiDrizzle** will address these same issues, while at the same time providing additional capabilities such as cosmic ray rejection and sky subtraction. Because it allows access to an extensive suite of user-adjustable parameters, **MultiDrizzle** gives users more flexibility in selecting the parameters that best suit their observations. This is all made possible within a single, easy-to-use interface.

MultiDrizzle creates many large files and does a considerable amount of processing. It is hence recommended that a modern fast machine with several GBytes of free space be used if possible. In the example to follow, each of the input FLT files is ~168MB and the final combined output product is 310MB. Many large intermediate files are also created but may be cleaned up if required.

The only calibration reference file required by **MultiDrizzle** is the distortion table (IDCTAB) discussed earlier in this chapter. This table describes the ACS distortion in terms of polynomial coefficients. The appropriate reference file is recorded in the image header keyword IDCTAB. This file should be obtained from the STScI archive (via Starview) and placed in some local directory. Before starting **MultiDrizzle** the environment variable "jref" must be set appropriately to point to this directory (see Example 1 in Section 4.4.4 for more information). Once the FLT files have been collected and the 'jref' environment established, PyRAF should be started and the **stsdas**, **analysis**, and **dither** packages loaded.

The IDCTAB reference file in the image header must NOT be modified prior to running MultiDrizzle or PyDrizzle. To make use of an updated distortion solution, users are advised to re-retrieve their data from the archive such that the CD-matrix keywords in the image header are correctly populated.

To avoid potential problems, it is recommended that the parameters for both **MultiDrizzle** and **PyDrizzle** be reset to the default values:

```
pyraf> unlearn pydrizzle
pyraf> unlearn multidrizzle
```

4.5.2 Tasks Performed by MultiDrizzle

MultiDrizzle starts with the flat fielded, but still distorted, FLT files created by **CALACS** in the pipeline. While these are the only mandatory input, there are many additional optional inputs including lists of improved registration information. In default mode **MultiDrizzle** performs the following steps in order, although the user has full control via the parameter file, of which steps are to be performed or omitted. Each step is described in detail in the example which follows in Section 4.6.

1. StaticMask: Create a mask of bad pixels, using information from the images themselves and associated data quality files,

2. SkySub: Carry out sky subtraction on each individual input image,

3. Driz_Separate: Calculate relative image shifts based on header information in the world coordinate system keywords (optionally supplemented by user-supplied shift information) and drizzle the input images onto a series of separate output images,

4. Median: Combine these images to create a clean median image,

5. Blot: Transform, or blot, the median image back to the frame of the input images,

6. Driz_cr: Create a cosmic ray mask for each input image based on comparison between the blotted median and its derivative and the original input image,

7. Driz_Combine: Perform the final **drizzle** combination step, applying the cosmic ray masks.

A warning to users: MultiDrizzle modifies the original calibrated images, for example, when updating the image DQ array with the static pixel mask and when performing the sky subtraction. When reprocessing, we advise retaining a copy of the original files in a separate directory.

4.5.3 Specifying Shifts Between Images

By default, **MultiDrizzle** uses the world coordinate system (WCS) information in the image headers to align the images. While dithers applied to a target within a single visit of HST are usually accurately reflected by the WCS, this is not the case for multiple visits which normally require guide star re-acquisitions and may utilize different guide stars. Even for nominally back-to-back exposures that are part of a CR-SPLIT association, offsets large enough to degrade final combinations do sometimes occur. As a result, it is essential to accurately determine image-to-image shifts, and possibly rotations, before running **MultiDrizzle**.

These (delta) shifts may be determined by separately drizzling each image onto a common WCS frame. This is performed by the 'driz_separate' task in step 3. Object lists derived for each drizzled image may then be matched and fit to derive a single shift for each image. More details on this topic and a description of the shift file formats are given in Section 4.4.3.

Once the shifts have been determined, **MultiDrizzle** must be rerun from the beginning with the shift file specified in the 'shiftfile' parameter. **MultiDrizzle** will update the association table to reflect these 'delta' offsets and will now use both the header WCS information plus the shift file to appropriately align images.

4.6 MultiDrizzle Example

The MultiDrizzle software is under continuing development. The following example was written in December 2003 should be used as a guide to facilitate reprocessing. While the algorithms and the overall processing steps will remain stable, specific details of the software may change.

This example describes the combination of six ACS/WFC images of the "Tadpole Galaxy" UGC 10214 which formed part of the ACS Early Release Observations (EROs) from program 8992. These images were acquired in two visits, with a significant shift between visits, as noted by the target RA and Dec, and with small dithers within each visit. The images were taken with the F606W filter and have the exposure times listed in Table 4.4. Prior to attempting this example, we assume that users have worked through the **PyDrizzle** examples in Section 4.4.4 and are familiar with using association tables and shift files.

The ERO images are available from the HST archive for anyone wishing to repeat this example. A new capability has been developed in the StarView archive interface, allowing the user to request FLT files only. The FLT files are the output from **CALACS** and have been corrected for bias, dark current and flat fielding. They are not corrected for distortion and still contain numerous cosmic rays. They form the input for **MultiDrizzle**.

Figure 4.10 shows chip 1 for one of these FLT files, j8cw54ovq_flt[sci,2], which still contains numerous cosmic ray events, hot pixels, and other artifacts. The rectangular, rather than 'rhombus' shaped image is a clear indicator that the geometric distortion has not yet been corrected.

Table 4.4: ERO datasets used in the accompanying **MultiDrizzle** example. The two sets of images were taken at different epochs and have a significant shift between epochs. Within a given epoch, small POS-TARG offsets are specified in accordance with the ACS-WFC-DITHER-LINE pattern which spans the WFC interchip gap.

Dataset	Date-Obs	RA (deg), Dec (deg)	POS-TARG1	POS-TARG2	Exposure Time (sec)
j8cw04abq_flt	01 Apr 2002	241.553625,+55.431094	0.000	0.000	150.0
j8cw04c2q_flt	01 Apr 2002		0.248	3.001	580.0
j8cw04c7q_flt	01 Apr 2002		0.248	3.001	840.0
j8cw54oiq_flt	09 Apr 2002	241.537833,+55.427594	0.000	0.000	150.0
j8cw54orq_flt	09 Apr 2002		0.248	3.001	580.0
j8cw54ovq_flt	09 Apr 2002		0.248	3.001	840.0

Figure 4.10:   A single, calibrated FLT file, j8cw54orq_flt.fits[sci,2], which still contains numerous cosmic rays, hot pixels, and other artifacts.
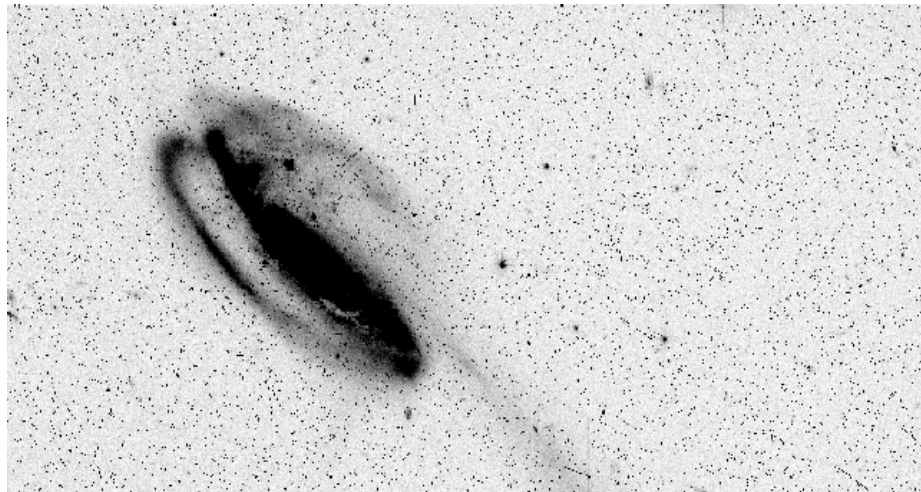
4.6.1 Detailed Processing Steps & Parameters

First, the appropriate calibrated FLT files and the geometric distortion reference table (IDCTAB) should be placed on the user's local disk. Next the 'jref' directory should be defined, PyRAF started, and the **dither** package loaded. To run **MultiDrizzle**, the parameters may be edited in the standard way using the "epar" facility.

The **MultiDrizzle** software has an extensive set of parameters, but the default values should allow the task to process nearly any set of images for an initial review. The parameters are separated according to the processing step they control, making it easier to interpret them.

In this example, we describe the parameters for each step in succession, though in practice, the user would set all relevant parameters at once. While the majority of relevant parameters are discussed here, a help document describing all parameters and tasks can be accessed by typing 'help multidrizzle' from within PyRAF.

Initial Setup

Parameters	Default	Description
output		Rootname for output drizzled products
suffix	flt	Suffix of input files in current directory
filelist		List of input files
refimage		Reference image with desired output WCS
runfile	multi.run	File for logging the script commands
restart		Parameter allowing processing to resume at a given step
coeffs	header	Use header-based distortion coefficients? (header, none)
context	no	Create context image during final drizzle?
clean	no	Remove temporary files?
section		Extension or group to be drizzled
bits	0	Integer mask bit values considered good
ra		right ascension output frame center
dec		declination output frame center
build	yes	Create multi-extension output file?
shiftfile		Shift file for improving WCS registration

The only required parameter is the rootname for the 'output' drizzled product. By default, **MultiDrizzle** will look for all files in the working directory with the FLT extension. The user may modify the 'suffix' parameter to include some other extension or may specify a subset of

images via the 'filelist' parameter. Using the images specified, **MultiDrizzle** calls **PyDrizzle** which uses the '**buildAsn**' task (see Section 4.4.2) to create an association table named 'output_asn.fits'. This association will be used to define the data set. The header WCS information from the entire set is used to define a common WCS output frame, and **MultiDrizzle** sets the **drizzle** parameters appropriately. If user defined shifts are available, these may be specified in the 'shiftfile' parameter, and the association table will be updated accordingly. However, a shift file is not usually available until after step 3, separately drizzling the images onto a common WCS, has been performed and objects matched.

A reference image which has the desired output WCS may be specified, and the input images will be drizzled to match the WCS of this image. Alternately, the central RA and Dec (ra, dec) of the reference pixel and the dimensions of the output frame (outnx, outny) may be specified, if desired, though reasonable values will be automatically determined from the images' WCS if these parameters are left blank. While the central RA and Dec are specified in the initial setup parameters listed above, the output image dimensions are specified in both the 'driz_separate' and 'driz_combine' parameters in steps 3 and 7, respectively.

The 'bits' parameter is defined as the integer sum of all bit values from the input images' DQ array that should be considered 'good' when building the weighting mask. Because **MultiDrizzle** was designed for use with multiple instruments, the default value is set to zero. For ACS data, the recommended default value is 8578. (For more information on selecting the appropriate bits for your data, refer to Section 4.4.1.) Information from the DQ array for each chip is used in combination with the 'bits' parameter to create temporary mask files called '*_inmask?.fits'. Pixels which were flagged in the DQ array *and* which were not specified as good via the bits parameter are assigned the value 0. All other pixels are set to 1 in the mask.

The distortion reference file is read from the image header via the IDCTAB keyword which specifies the name and location of the appropriate file. The distortion coefficients for each chip are written to temporary ascii files named '*_coeffs?.dat'. If the user wishes to retain these files, the parameter 'clean' should be set to 'no'.

In general, the default **MultiDrizzle** parameters will work well for most data. When setting up the Tadpole images, for example, we have specified *only* the following non-default parameters: output='example', 'bits=8578', 'combine_nhigh=2', and 'driz_cr_snr=4.0 3.0'. The choice of these last two parameters is explained in the sections that follow. In default mode, **MultiDrizzle** performs each of its 7 steps in order. In this example, however, we perform some of the steps and examine the intermediate products before final drizzle combination is performed. Approximately 5GB of free disk space is required for this example when intermediate products are not removed. An outline of the entire process is described below:

1. Run only steps 1 through 3 of Section 4.5.2 to create sky-subtracted, separately drizzled images which are based on a common WCS.

```
pyraf> multidrizzle output='example' bits=8578 static+ skysub+
driz_separate+ median- blot- driz_cr- driz_comb-
```

2. Measure the positions of stars in the separately drizzled images and derive a shift file which defines the residual offsets.

3. Rerun step 3 using the derived 'shiftfile' to create new separately drizzled images. Also turn on step 4 to create a well-aligned median image.

```
pyraf> multidrizzle output='example' bits=8578 shiftfile='shifts'
static- skysub- driz_separate+ median+ combine_nhigh=2 blot-
driz_cr- driz_comb-
```

4. Examine the median image to ensure that the PSF is 'round' and 'narrow' and that cosmic-rays and other artifacts are appropriately rejected.

5. Run steps 5 through 7 to transform the median image back to the reference frame of each of the original input images and to derive cosmic ray masks. Using these new masks, perform the final drizzle combination.

```
pyraf> multidrizzle output='example' bits=8578 shiftfile='shifts'
static- skysub- driz_sep- median- blot+ driz_cr+ driz_cr_snr='4.0
3.0' driz_comb+
```

Once the optimal set of parameters and the optional shift file is derived, as described in the outline above, **MultiDrizzle** may be executed in a single pass by turning all steps on and by specifying any desired non-default parameters:

```
pyraf> multidrizzle output='example' bits=8578 shiftfile='shifts'
static+ skysub+ driz_sep+ median+ combine_nhigh=2 blot+ driz_cr+
driz_cr_snr='4.0 3.0' driz_comb+
```

While **MultiDrizzle** may be executed from the command line, as shown in the above examples, it may also be run from the 'epar' facility which allows the user to see all parameters at once and to turn particular steps on and off. We recommend that beginners use the 'epar facility' to become familiar with all steps and parameters before running **MultiDrizzle** from the command line.

1. Creating the Static Mask

Parameters	Default	Description
static	yes	Create static bad-pixel mask from data?
staticfile		Name of (optional) input static bad-pixel mask
static_goodval	1.0	Value of good pixels in the input static mask

Output Files	Modified DQ arrays in the original input files

When 'static=yes', this step goes through each of the input images, calculates the rms value for each chip, and identifies pixels that are below the median value by more than 5 times the rms. It is aimed at identifying pixels that may have high values in the dark frame, which is subtracted during calibration, but may not necessarily have high values in the images, and thus subtraction gives them strongly negative values. Such pixels are not always flagged in the DQ file, and this step allows them to be identified. Sometimes such pixels fall on bright objects so instead of being negative, they would be positive but lower than surrounding pixels. If the images are dithered, then they should land on blank sky at least some of the time, in which case they will appear negative and will be flagged.

For the Tadpole Example, we have left the 'StaticMask' parameters to their default values. After this step is performed, the image DQ arrays are updated with the static mask, and new flagged pixels are set with bit 64. The '*_inmask?.fits' files are subsequently updated during further processing.

2. Performing Sky Subtraction

Parameters	Default	Description
skysub	yes	Perform sky subtraction?
skytype	"single"	Type of sky subtraction (single, quadrants)
skyname	"SKYSUM"	Header keyword containing sky value (SKYSUM = ACS, BACKGRND = WFPC2)
skywidth	50.0	Interval width for sky statistics
skystat	"median"	Sky correction statistics parameter (median, mode, mean)
skylower	-50.0	Lower limit of usable data for sky (always in DN)
skyupper	200.0	Upper limit of usable data for sky (always in DN)

Output Files	Modified science array in the original input files

When 'skysub=yes', this task will subtract the sky from each chip ('skytype=single') or from each of the four individual amplifiers on the WFC chips separately ('skytype=quadrants'). The other parameters correspond directly to those in the **sky** task in the **dither** package, and are passed to it exactly as they are specified here.

The main parameter that users must set is the histogram width, though the default value is good for most cases. The width determines the region of the pixel histogram used to determine the image statistics. It should be set to include most pixels in the sky (so substantially more than the FWHM of the sky distribution) but not so large as to include a substantial amount of power from objects or cosmic rays.

The 'SkySub' task will update the header keyword defined by the parameter 'skyname' with the derived sky value for each chip and will subtract the sky from the original FLT images. Sky subtraction is recommended for effective cosmic-ray flagging and removal, but only if sufficient blank sky is available to perform an accurate determination.

Great care must be taken when choosing to implement sky subtraction, because the original calibrated images will be modified, and because if sufficiently blank sky is not available, sky subtraction will produce erroneous results. In the case of the Tadpole images, adequate blank sky is available for each WFC chip to allow an accurate sky determination when 'skytype' is set to 'single'. Choosing 'skytype=quadrants', on the other hand, results in an over-subtraction of the 'A' amplifier (upper-left quadrant) since the galaxy light dominates the background in this region of the detector.

3. Drizzling to Separate Output Images

Parameters	Default	Description
driz_separate	yes	Drizzle to separate output images?
driz_sep_outnx		Size of x-axis for separate output frame
driz_sep_outny		Size of y-axis for separate output frame
driz_sep_kernel	"turbo"	Shape of kernel function (square, point, gaussian, turbo, tophat, lanczos3)
driz_sep_scale	1.0	Linear size of output pixels (relative to input)
driz_sep_pixfrac	1.0	Linear size of drop in input pixels
driz_sep_rot	0.	Rotation of input image to be applied (degrees anti-clockwise)
driz_sep_fillval	INDEF	Value assigned to undefined output pixels

Output Files	Images= '*_single_sci.fits', '*_single_wht.fits'

When 'driz_separate=yes', the input images are corrected for geometric distortion and drizzled onto separate output frames which have a common WCS. Any shifts, rotations, or scale changes are calculated from the image headers by **PyDrizzle**. The output image dimensions are calculated on-the-fly and the pixel scale is taken from the column 'scale' from the IDCTAB, where the default values are 0.05 *arcsec*/*pix* for the WFC and 0.025 *arcsec*/*pix* for the HRC. The drizzled images are in units of *electrons*/*sec*.

By default, the 'driz_separate' task uses the 'turbo' drizzle kernel and **drizzle** parameters 'pixfrac=1' and 'scale=1'. For more information on setting these parameters, refer to the HST Dither Handbook. These parameters can be changed; for example, masks can be substantially improved by specifying a smaller value of scale (e.g., 0.5 or 0.66), with the primary trade-off being much larger images (their size increases as the inverse square of the value of 'scale') and increased computation time.

In the Tadpole example, we have left the 'Driz_Separate' parameters to their default values. While **PyDrizzle** produces a final drizzled image containing 3 extensions (the science, weight, and context images), the 'driz_separate' products are separate science and weight images named '*_single_sci.fits' and '*_single_wht.fits'. No context image is created.

One of the singly drizzled FLT images, 'j8cw54ovq_flt_single_sci.fits', is shown in Figure 4.11. This image still contains numerous cosmic ray events, hot pixels, and other artifacts. The 'rhombus' shape is a result of correcting the geometric distortion. The corresponding weight image, 'j8cw54ovq_flt_single_wht.fits', is shown in Figure 4.12, where white indicates pixels with zero weight. Due to the effects of distortion and

varying pixel area in the FLT images, the weight image changes gradually across the detector. Because the association table was used to define a common WCS for all images, the drizzled image is 'padded' with zeros outside the boundary of the original array. The weight image is set to zero in these regions, allowing these pixels to be rejected during median combination.

The separately drizzled science images may be used to improve the image registration prior to final drizzle combination. In this example, the images form two groups of three. While the WCS information for images within a single group (visit) are adequate to align them, there is a small residual offset between visits. Shifts which are determined from separately drizzling images onto a common WCS are by definition 'delta' shifts (see Section 4.4.3 for details), and will be applied in addition to any offsets from the WCS when a 'shiftfile' is provided. Because no additional scale or rotation was applied to the singly drizzled images, the shifts are in the 'input' frame of reference.

In this example, unsaturated stars in the short exposures (j8cw04abq and j8cw54oiq) were used to derive the offsets between the two groups of images. Shifts were measured to the nearest 0.1 pixel and are listed below in the form of a shift file. Refined shifts within a given visit may also be determined, but these are typically less than a tenth of a pixel, and accuracy at this level is not as critical for extended sources as it may be for point sources.

```
#units: pixels
#frame: input
#form: delta
j8cw04abq_flt  0.0  0.0
j8cw04c2q_flt  0.0  0.0
j8cw04c7q_flt  0.0  0.0
j8cw54oiq_flt -5.6 -0.3
j8cw54orq_flt -5.6 -0.3
j8cw54ovq_flt -5.6 -0.3
```

When the 'driz_separate' step is run for the second time, but with a 'shiftfile' specified, the association table for the data set will be automatically updated. To confirm that the new separately drizzled images are appropriately registered, the position of stars should again be examined. It is also useful to create a median image and examine the width and shape of the PSF over the entire FOV to look for any effects of mis-registration. Median combination is performed in Step 4 of **MultiDrizzle.**

Figure 4.11: The singly drizzled FLT image 'j8cw54ovq_flt_single_sci.fits' from the **MultiDrizzle** example.
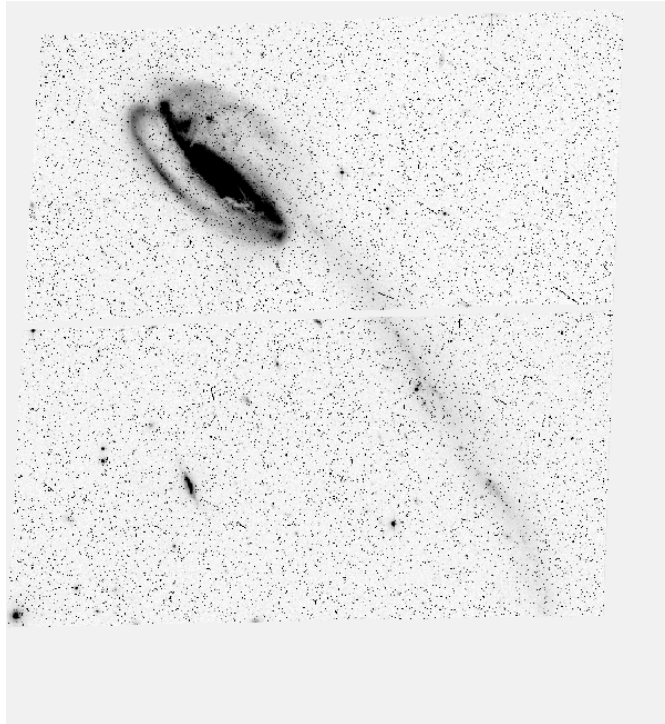
Figure 4.12: The weight image 'j8cw54ovq_flt_single_wht.fits' corresponding to the singly drizzled image in Figure 4.11 where white indicates zero weight.
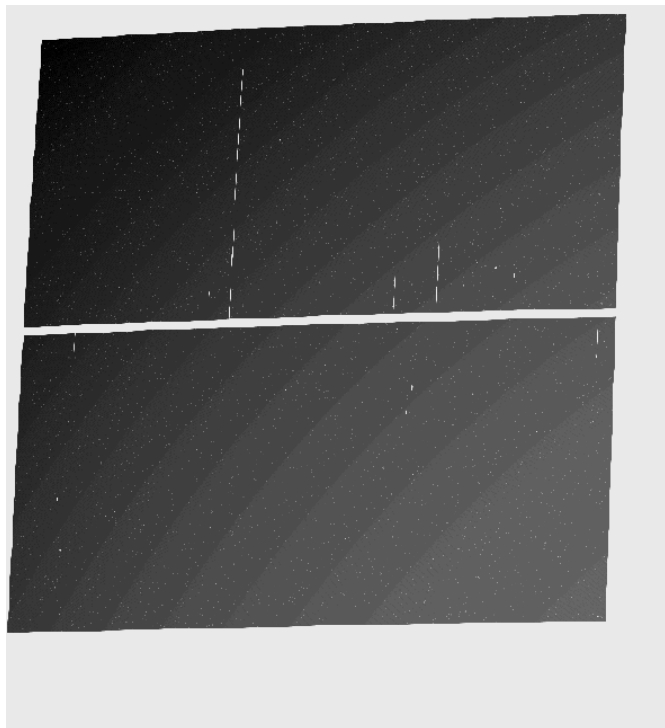
4. Creating the Median Image

Parameters	Default	Description
median	yes	Create median image?
median_newmasks	yes	Create new masks when doing the median?
median_preclean	no	Remove intermediate files prior to doing median?
combine_type	"median"	Type of combine operation (average, median, sum, minmed)
combine_reject	"minmax"	Type of rejection (none, minmax, ccdclip, crreject, sigclip, avsigclip, pclip)
combine_nsigma	6 3	Significance for accepting min instead of median
combine_nlow	0	minmax: Number of low pixels to reject
combine_nhigh	1	minmax: Number of high pixels to reject
combine_lthresh	INDEF	Lower threshold for clipping input pixels
combine_hthresh	INDEF	Upper threshold for clipping input pixels
combine_grow	1.0	Radius (pixels) for neighbor rejection

Output Files	Images= 'output_med.fits', '*_single_wht_maskhead.pl'

When 'median=yes', this step creates a median image from the separate drizzled input images, allowing a variety of combination and rejection schemes. If 'combine_type' is set to 'median' or 'average', then the routine calls the IRAF task **imcombine**, passing to it the values of 'combine_reject' (the rejection algorithm chosen, usually expected to be 'minmax'), 'combine_nlow' and 'combine_nhigh' (the number of low and high pixels to reject), and 'combine_grow' (the amount by which flagged pixels can grow). All **imcombine** parameters other than those specified above are reset to their default values. While the 'median' is recommended, a slightly more sophisticated algorithm than those available in **imcombine** will be used when 'combine_type=minmed'. This algorithm is described at the end of this section.

If 'median_newmasks=yes', then the singly drizzled weight maps ('*_single_wht.fits') are used to create pixel masks for each image (with values 0 and 1) which are named '*_single_wht_maskhead.pl'. The IRAF task **mask_head** prepares the singly drizzled images for use with **imcombine** by populating the header bad pixel mask keyword 'BPM' for each image. These masks will be used by **imcombine** when combining images, where the assumed mask parameters are 'masktype=goodvalue' and 'maskvalue=1', indicating that pixels assigned a value of 1 are considered good.

If 'median_newmasks=no', this task will use whatever masks are specified by the user (and which are created offline) in the 'BPM' header keyword of each image. In general, however, it is recommended that the pixel masks which are generated by default are used instead.

Selecting the best parameters for the median step can be an iterative process and should always involve examination of the clean, combined product to verify that the majority of cosmic-rays and other artifacts are successfully removed. The rejection algorithm which is ultimately chosen depends largely on the number of datasets being combined and the amount of overlap between dithered images.

In this example, we have chosen the default parameters 'combine_type=median', 'combine_reject=minmax', and 'combine_nlow=0'. Instead of the default value, we have set 'combine_nhigh=2' so that hot pixels are rejected in the outer portions of the image, which have only 3 input images contributing to the median calculation, and for which 2 of the 3 images are at the same dither position. Thus, the **imcombine** parameter 'nkeep' (the minimum number of pixels retained) is 1 in the outer portion of the median image. In the central regions of the image, six images have contributed to the median, but two high images are rejected, so 'nkeep' is 4. The six separately drizzled images are combined using the bad pixel masks and the rejection parameters specified above to create a single clean median image named 'example_med.fits'. This median image is shown in Figure 4.13.

Figure 4.13: The cleaned median image created using the 6 separately drizzled Tadpole images and their bad pixel masks.
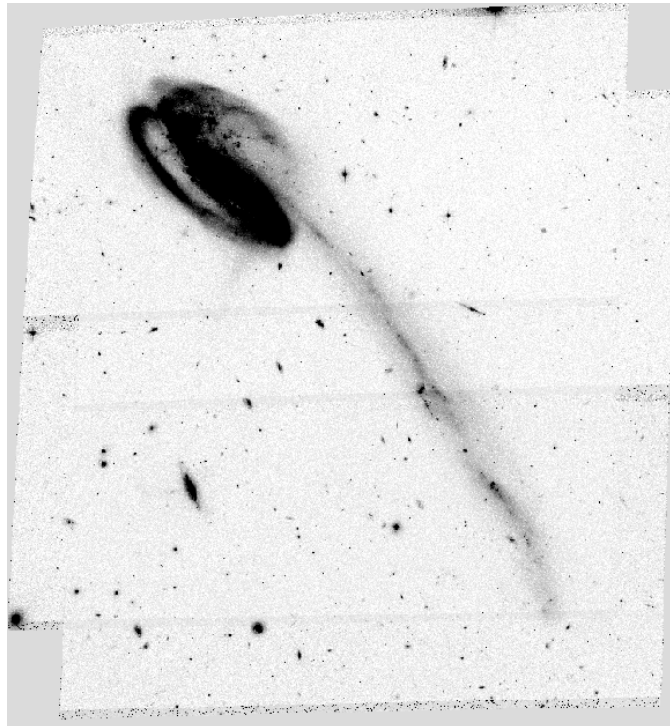
When 'combine_type=minmed', a slightly more sophisticated algorithm will be used to combine images. This algorithm requires significantly more computation time and disk space (>1GB) compared to the other available algorithms. The basic concept is that each pixel in the output combined image will be either the median or the minimum of the input pixel values, depending on whether the median is above the minimum by more than n times sigma. An estimate of the "true" counts is obtained from the median image (after rejecting the highest-valued pixel), while the minimum is actually the minimum unmasked ("good") pixel. This algorithm is designed to perform optimally in the case of combining only a few images (3 or 4), where triple-incidence cosmic rays often pose a serious problem for more simplified median combination strategies. The algorithm performs the following steps:

1. Create median image, rejecting the highest pixel and applying masks.

2. Use this median to estimate the true counts, and thus derive an rms.

3. If the median is above the lowest pixel value by less than the first value in 'combine_nsigma', then use the median value, otherwise use the lowest value.

4. If 'combine_grow' > 0, repeat the above 3 steps for all pixels around those that have already been chosen as the minimum, this time using a lower significance threshold specified as the second value in 'combine_nsigma'.

The last step is very successful at flagging the lower signal-to-noise "halos" around bright cosmic rays which were flagged in the first pass.

5. Blotting Back the Median Image

Parameters	Default	Description
blot	yes	Blot the median back to the input frame?
Output Files	Images= '*_sci?_blt.fits'	

When 'blot=yes', this task takes the median image and uses the **dither** package **blot** task to apply the geometric distortion and to transform ('reverse drizzle') it back to the reference frame of each of the original individual input images. This involves reversing the shifts and reapplying the geometric distortion that had been removed in step 3. In addition, the median image is resampled to the pixel scale of the original images and is trimmed to match the dimensions of each input image. This step is done in preparation for subsequent cosmic-ray rejection in step 6. The blotted frames are named '*_sci?_blt.fits'.

If desired, the user may wish to display the input images and blink them with their 'blotted' counterparts. The 'blotted' images should align perfectly with their respective input images and should be reasonably similar in appearance, except for the fact that they should be cleaned of cosmic rays and other defects.

6. Creating Cosmic Ray Masks

Parameters	Default	Description
driz_cr	yes	Perform CR rejection with deriv and driz_cr?
driz_cr_snr	"3.0 2.5"	Driz_cr.SNR parameter
driz_cr_scale	"1.2 0.7"	Driz_cr.scale parameter
Output Files		Images= '*_sci?_blt_deriv.fits', '*sci?.fits', '*_sci?_crderiv.pl', '*_sci?_cor.fits', '*_sci?_mask.pl'

When 'driz_cr=yes', this step uses the original input images, the blotted median images, and the derivative of the blotted images (which it creates using the **deriv** task) to create a cosmic ray mask for each input image (using the **driz_cr** task).

First, the **deriv** task uses the blotted median images ('*_sci?_blt.fits') from step 5 to calculate the absolute value of the difference between each pixel and its four surrounding neighbors. For each pixel, the largest of these four values is saved in an output image, '*_sci?_blt_deriv.fits', which represents an effective gradient or spatial derivative.

These derivative images are used by the task **driz_cr** when comparing the original and blotted images. First, the original FLT images for each chip are copied to files named '*_sci?.fits' which are the required input for the **driz_cr** task. These images are compared with the corresponding blotted median image '*_sci?_blt.fits' and its absolute derivative '*_sci?_blt_deriv.fits' to create a mask of cosmic rays (and other blemishes, like satellite trails). Where the difference is larger than can be explained by noise statistics, or the flattening effect of taking the median, or perhaps an error in the shift (the latter two effects are estimated using the image derivative), the suspect pixel is masked. Cosmic rays are flagged using the following rule:

|data_image - blotted_image| > scale*deriv_image + SNR*noise

where 'scale' is the user supplied **driz_cr** parameter listed above and is defined as the multiplication factor applied to the derivative before determining if the difference between the data image and the blotted image is sufficiently great to require masking. 'Noise' is calculated using a combination of the detector read noise and the poisson noise of the blotted median image plus the sky background.

The user must specify a cut-off signal-to-noise (SNR) value for determining whether a pixel should be masked. Actually, two cut-off signal-to-noise ratios are needed, one for detecting the primary cosmic ray, and a second for masking lower-level bad pixels adjacent to those found in

the first pass. After the first pass through the image, the procedure is thus repeated on pixels that are adjacent to previously masked pixels using a lower SNR threshold, since cosmic rays often extend across several pixels.

The final output is a cosmic-ray mask file named '*_sci?_crderiv.pl'. One of the resulting masks for chip 1 is shown in Figure 4.14 and should be blinked with the original image j8cw54orq_flt[sci,2] from Figure 4.10 (or the equivalent j8cw54orq_flt_sci2.fits file) to visually ascertain that all cosmic rays were flagged. If it appears that the central pixels of some stars were unnecessarily masked, the 'driz_cr_scale' parameter should be increased. If not enough cosmic rays were masked out, this parameter should be decreased. In this example, the default 'driz_cr_snr' values "3.0 2.5" were too stringent and resulted in flagging the centers of stars and the core of the Tadpole galaxy. Instead, we have increased the default SNR values to "4.0 3.0" to create ideal cosmic ray masks for this data set.

The **driz_cr** task also creates a '*_sci?_cor.fits' image, where flagged pixels are replaced with pixels from the blotted median image. The cosmic ray mask files are then multiplied by the bad pixel masks (which are a combination of the image DQ array and the static masks) to create a final mask file for each input image, '*_sci?_mask.pl', which will be used during final drizzle combination.

Figure 4.14: A single cosmic ray mask 'j8cw54orq_flt_sci2_crderiv.pl'. This mask should be blinked with the original image 'j8cw54orq_flt[sci,2]' from Figure 4.10 or the equivalent 'j8cw54orq_flt_sci2.fits' output file to assure that optimal parameters were chosen in the **driz_cr** task.
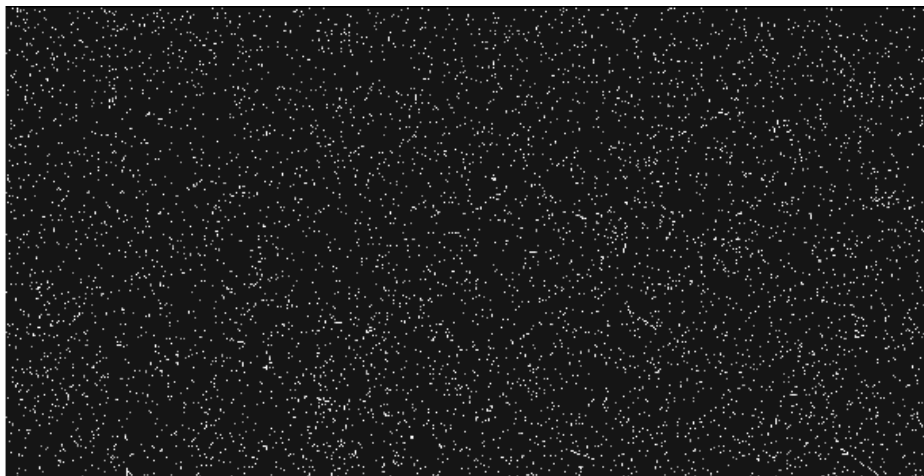
7. Performing the Final Drizzle Combination

Parameters	Default	Description
driz_combine	yes	Perform final drizzle image combination?
final_outnx		Size of FINAL output frame x-axis
final_outny		Size of FINAL output frame y-axis
final_kernel	"square"	Shape of kernel function (square, point, gauss-ian, turbo, tophat, lanczos3)
final_scale	1.0	Linear size of output pixels (relative to input)
final_pixfrac	1.0	Linear size of drop in input pixels
final_rot	0.	Rotation of input image to be applied (degrees anti-clockwise)
final_fillval	INDEF	Value given to undefined output pixels

Output Files	Images= output_sci.fits, output_wht.fits

When 'driz_combine=yes', this step takes the original input images, together with the final cosmic ray masks, and drizzles them all onto a single output image. The standard **drizzle** parameters kernel, scale, pixfrac, and rot can be specified by the user, if desired. By default the scale of the output image is 1.0, but the user is encouraged to experiment with other options (e.g. scale=0.5 and pixfrac=0.7 yields a sharper output PSF).

When the following initial setup parameters are set: 'build=yes' (default) and 'context=yes' (non-default), the final **MultiDrizzle** output image will be a single multi-extension FITS file named 'example_drz.fits'. The format of this file is identical to the DRZ product from **PyDrizzle** which is delivered from the archive (see Section 4.3.3) and which contains the science image in extension 1, the weight image in extension 2, and the context image in extension 3. When 'build=no', these files will be written to separate output files. When the default value 'context=no' is used, no context image is created.

The first extension of the drizzled product contains the science (SCI) image which is corrected for distortion and which is dither-combined (or mosaiced), if applicable. The drizzled SCI image derived from the Tadpole example is presented in Figure 4.15 and is in units of *electrons/sec*. All pixels have equal area on the sky and equal photometric normalization across the field of view, giving an image which is both photometrically and astrometrically accurate for both point and extended sources. The dimensions of the output image are computed on-the-fly by **MultiDrizzle** and the default output plate scale is read from the 'scale' column in the IDCTAB. These parameters, however, may be chosen by the user to best suit the actual data.

The second extension of the output image contains the weight (WHT) image. This image gives the relative weight of the output pixels and, in standard processing using the **MultiDrizzle** defaults, it can be considered an effective exposure time map. The weight image from the example is shown in Figure 4.16, where darker areas have higher weight. The chip edges and gaps are clearly visible, as are column defects and cosmic ray features. The bulk of the image is "dark gray" corresponding to the overlap of all six inputs. In this area the weight value is ~3140, equal to the sum of the exposure times of the six images which contribute. There is also a smooth variation across the image due to the variation of the pixel area on the sky caused by the distortion.

The third extension of the **MultiDrizzle** output image contains the context (CTX) image which encodes information about which input image contributes to a specific output pixel. This is done using a bitmask for each output pixel, where 'bit set' means that the image, in the order it was combined, contributed with non-zero weight to that output pixel. The context image starts as a single 32-bit integer image but is extended as a cube with additional 32-bit deep planes as required to handle all the input images. The context image for the Tadpole example is shown in Figure 4.17. As there are six input images, each with two chips which are treated separately, this image has 12 bit planes which may be set. The darkest area shown corresponds to the [sci,2] chip from all six inputs and hence has all the following even bits set: 2+8+32+128+512+2048=2730. Cosmic ray hits or other defective pixels contribute to the appropriate bit plane with zero weight and hence appear as lighter spots.

Figure 4.15: The science (SCI) extension of the drizzled product from the **MultiDrizzle** example. This image has been corrected for distortion and drizzled onto a single mosaic using the six images in the dither pattern.
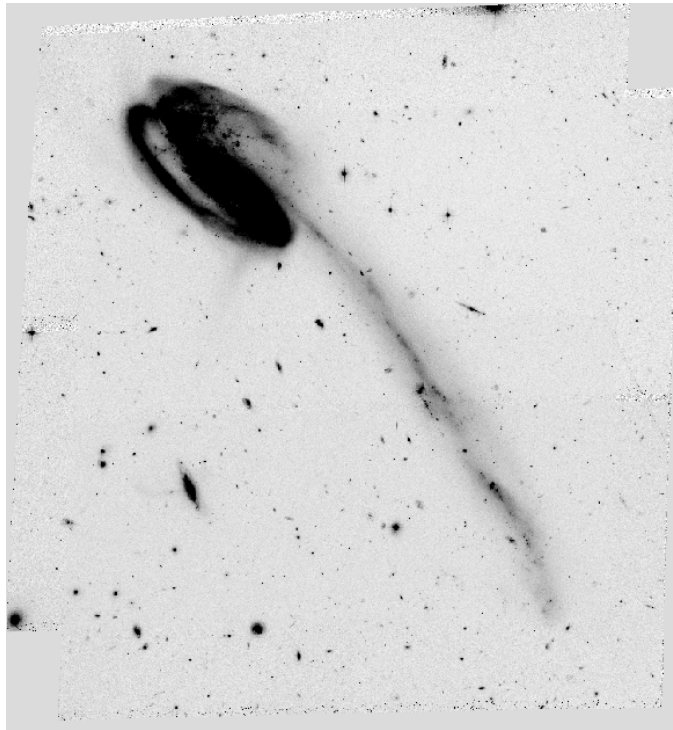
Figure 4.16: The corresponding weight (WHT) extension of the drizzled product from the example.
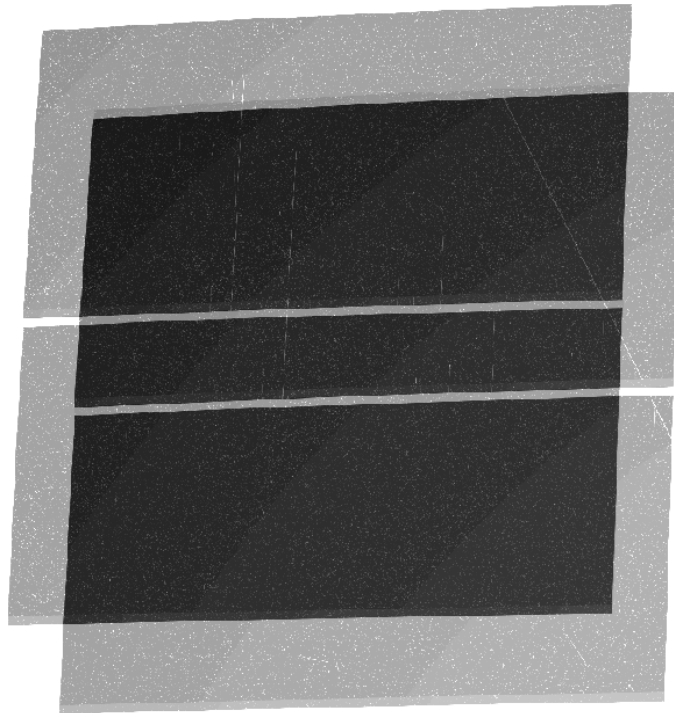
Figure 4.17: The corresponding context (CTX) extension of the drizzled product from the example.
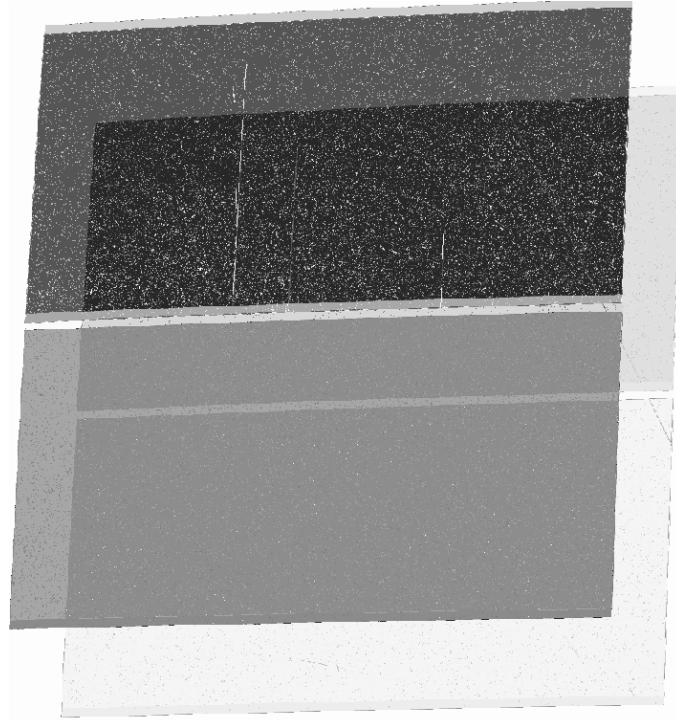
ACS Error Sources

In this chapter. . .

This chapter describes the sources of uncertainty in ACS data, and provides pointers to resources beyond this manual for the latest descriptions of our evolving understanding of these issues.

5.1 Readout Noise and A-to-D Conversion

5.1.1 Readout Noise

The readout 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 effect of the read noise can however be minimized by keeping the number of pixels where signal is measured small and by minimizing the number of CCD readouts (while still allowing for good cosmic ray and hot pixel removal). The read noise is independent of position on the CCD. Table 5.1 lists the read noise (R) expressed in electrons (rms), which is associated with the readout from each amplifier for the main gain settings.

Table 5.1: Readout Noise (R) and Total Noise (Readout+Quantization, TRN) for ACS CCDs

Chip	Amp	Gain 1			Gain 2		
		Gain	R	TRN	Gain	R	TRN
WFC1	A	1.00	4.99	5.00	2.002	5.24	5.27
WFC1	B	0.972	4.89	4.90	1.945	5.08	5.11
WFC2	C	1.011	5.16	5.17	2.028	5.37	5.40
WFC2	D	1.018	4.82	4.83	1.994	5.11	5.14
HRC	C	1.185	4.46	4.47	2.216	4.71	4.78

5.1.2 A-to-D Conversion

The analog information (electrons) accumulated in the CCD are converted into data numbers (DN) by the analog-to-digital converter (ADC). Both ACS CCD cameras employ 16-bit ADCs, which can produce a maximum of $2^{16} = 65,535$ DN. If the gain conversion factor (g) is expressed in (e⁻/DN) the largest number of electrons representable by these ADCs is given by $g*2^{16}$. Any charge packet containing more than $g*2^{16}$ electrons will reach the digital saturation limit (see Section 5.6.4). 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 quantization noise. It can be shown (Janesick J.R., 2001, Scientific Charge-Coupled Devices, SPIE Press.) that the quantization noise QN is constant for a given gain setting when expressed in DN: $QN(DN)=12^{-1/2}$. The quantizing noise can be converted into noise electrons as: $QN(e^-)=0.288675*g$. The total noise (TRN) associated with CCD readout noise (R) and quantizing noise is obtained by adding the two figures in quadrature (see Table 5.1), $TRN = (R^2 +(0.288675*g)^2)^{1/2}$.

5.2 Bias Subtraction

5.2.1 Bias Calibration Issues

Bias reference frames are acquired daily for scientific calibration purposes and for monitoring the detector performance. Every week multiple bias frames are combined together into a reference superbias image. The combination removes the cosmic rays accumulated during the readout time and enhances the signal-to-noise ratio of the fixed pattern noise.

CALACS performs the bias correction in two steps (see Section 3.4.1): **doBlev** subtracts the bias level from the overscan region and **doBias** subtracts a reference bias image.

Each quadrant of the WFC focal plane (A&B for WFC1 and C&D for WFC2) has two overscan regions: a 24-pixel-wide leading physical overscan at columns 1-24 and a 20-row-wide virtual overscan at rows 2049-2068 obtained by overclocking after the readout of the last CCD row.

All four WFC amplifiers produce a horizontal ramp in the leading overscan which extends up to 18 columns toward the active area (see Sirianni, M., Martel, A.R., & Hartig, G.F., 2001, JHU-ACS internal report, available on-line at:

http://acs.pha.jhu.edu/instrument/calibration/results/by_it
em/detector/wfc/build4/overscan/).

The amplitude of the ramp is different for each amplifier, but it is generally larger in the A&D quadrants than in the B&C quadrants. The perturbation associated with the ramp fades gradually and disappears within approximately 15 columns. The bias level from the physical overscan can be safely measured using the last six columns adjacent to the active area [columns 19-24 for A&D and columns 4121-4126 for B&C].

The virtual overscan is not used to estimate the bias level. This region shows large scale structure which is quadrant and gain dependent. This structure, due to the bias fixed pattern noise, makes using the virtual overscan in an automatic routine problematic. Moreover, this region can be contaminated by deferred charge due to degradation in the parallel charge transfer efficiency.

Analysis of the active area in the bias frames shows small differences in the bias level between the leading physical overscan and the active area (Sirianni, et al., 2002, "HST Calibration Workshop" STScI, page 82). This offset varies from amplifier to amplifier and it can be as large as 3.5 DN. Ideally, no offset would be seen. In principle, if the residual offset between the imaging area and the overscan region were always of the same amplitude, a full frame bias subtraction (**doBias**) would remove any residual differences. Unfortunately, the offset is not constant but shows random variations on the order of few tenths of a DN. These variations are likely due to 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.

HRC images are read out using only one amplifier. Each image has three overscan regions: the 19-pixel-wide leading and trailing physical overscans at columns 1-19 and 1044-1062, respectively, and the 20-row-wide virtual overscan at rows 1025-1044. The readout amplifier produces a horizontal ramp in the leading overscan which extends up to 10 columns toward the active area. As in the case of WFC images, the bias level is therefore measured in the last six columns of the leading physical overscan. There is not a significant difference between the bias level in the overscan area and in the active area.

Even after the subtraction of the overscan bias level, some structure remains in the bias frames. This structure includes real bias pattern noise

and also some dark current accumulated during the time required to read out the detector. To remove this structure a superbias is subtracted from each science image.

In principle, a perfect bias subtraction would give an image, that once flat fielded and converted to electrons should not present any discontinuity at the boundaries between quadrants.

However, in the case of WFC images the amplitude of the residual offset level varies between each quadrant. This results from the bias level being measured in the leading physical overscan and then subtracted from the active area of each quadrant, causing the center of the resulting image to show a jump between the two adjacent quadrants, A&B or C&D (Sirianni, et al., 2002, "HST Calibration Workshop" STScI, page 82). This jump is not removed by the superbias subtraction. Indeed, since this offset (usually with a maximum amplitude of 0.3 DN) is present in all calibration and scientific frames, the final product of **CALACS**, after bias and dark subtraction, may show a quadrant-to-quadrant jump as large as few DNs (Figure 5.1).

Figure 5.1: Calibrated WFC1 image showing the quadrant-to-quadrant jump
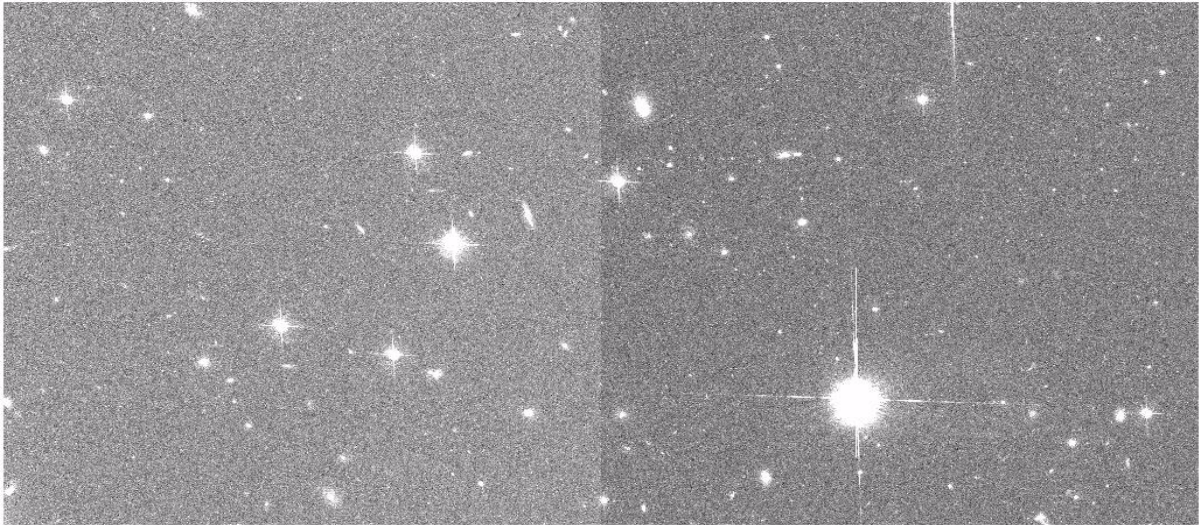
For point source photometry, the local background in an annular region is typically subtracted so that the residual offset has essentially no impact on the integrated magnitudes. However, if the object falls across the quadrant jump the photometry should be considered suspect. Similarly, extended objects may spread over two or more quadrants, so their surface brightness profile will suffer from the residual offset. More studies are in progress to better characterize this problem and to develop a correction that can be applied directly to the calibrated images. At the moment we suggest fitting the sky level in each quadrant separately.

A detailed study of the on-orbit characteristics of bias frames, the overscan regions and superbiases used in the pipeline will be published in a future ISR.

5.2.2 Bias Jump

There is some evidence for intermittent variations at the sub-DN level in the WFC bias level during the readout. On a few occasions some frames show horizontal bias jumps in one or more quadrants that can run for several hundreds of rows (Figure 5.2). These effects occur at a few tenths of a DN level and they might appear in one or more quadrants simultaneously. The probable source is interference from other activities on the spacecraft and/or scientific instruments with the electronics. At the moment there is no automatic detection of bias jumps within the calibration pipeline. Though bias jumps at the sub-DN level are not important for most applications, users should be cautioned that they might exist in processed data and we recommended a careful inspection of the data for their presence.

Figure 5.2: Bias Jump in the WFC1 Quadrant B only. The vertical stripes in the data are hot columns and are unrelated to bias.
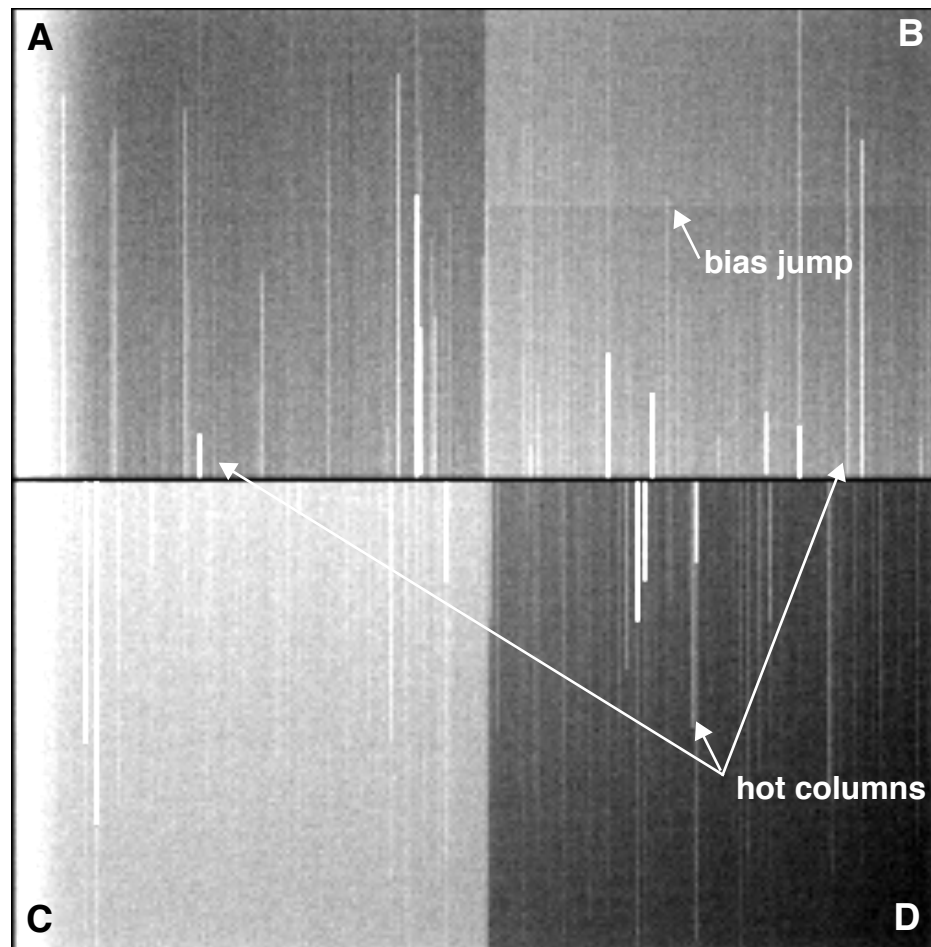
5.2.3 Bias Subarrays for the WFC and HRC

When science data are obtained with use of subarrays, dark and flat field images will be extracted from full-frame images. Bias images will also be extracted from full-frame images, if the subarray is contained in a single quadrant. Tests have shown that this does not degrade the quality of the dark, flat field or bias corrections as compared to full-frame data. It is strongly advised that subarrays not be defined such that they span amplifier quadrants. In the case that the subarray extends into the adjacent amplifier quadrant, users must obtain their own subarray bias images (with one-amp readout), otherwise the bias level difference between adjacent amplifier quadrants would be applied inappropriately.

5.3 Dark Current, Hot Pixels, and Cosmic Rays

5.3.1 Dark Current

The procedure for creating ACS dark reference files and applying dark subtraction to ACS science data is described in detail in Section 3.4.3. Because it takes a few weeks to collect enough frames to create the dark reference file, the "best" darkfile is typically not available in the pipeline until 2-3 weeks after the date of observation. Users may verify whether the darkfile most appropriate to their observations has been installed for pipeline use by checking the ACS reference file web pages:

http://www.stsci.edu/hst/acs/analysis/reference_files/drkim age_list.html.

Using an old reference file will produce a poor dark correction: either leaving too many hot 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/HRC

The dark current is not constant across the CCDs. Figure 5.3 shows the 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 5.3: High S/N combination of WFC dark frames illustrating dark current structure
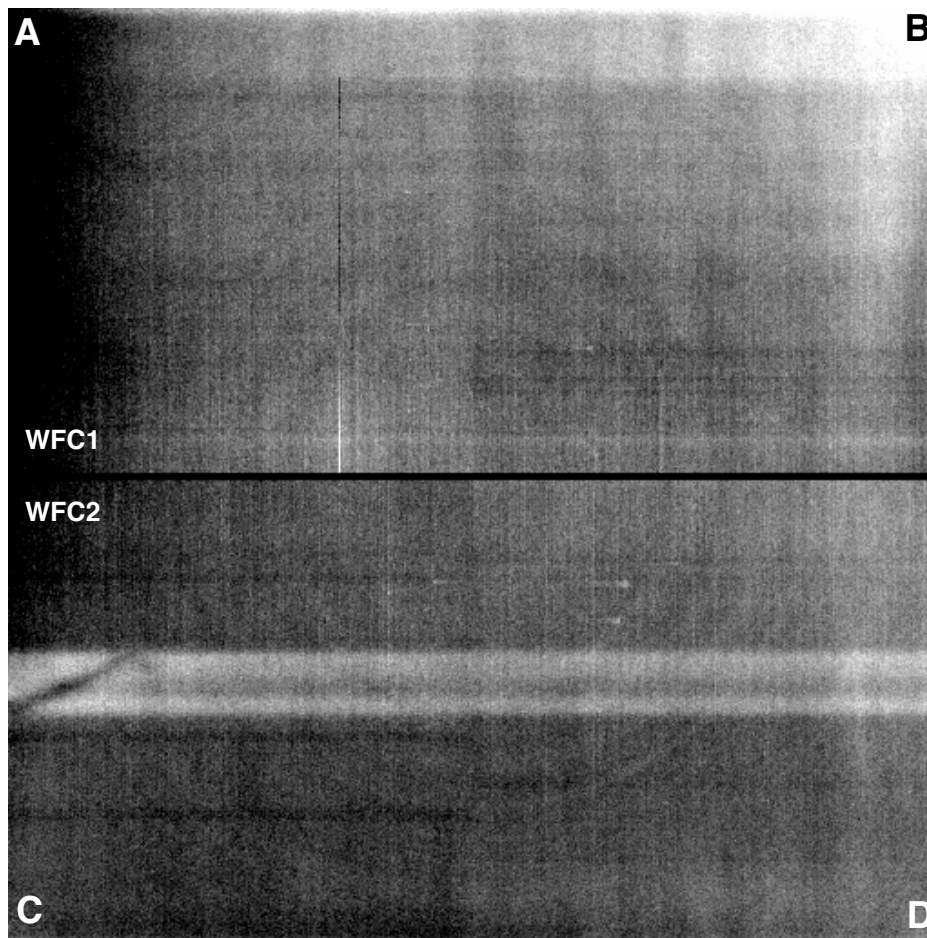
The WFC1 and HRC histograms in Figure 5.4 and Figure 5.5 show the growth of hot pixels (for more information, refer to Section 5.3.2). A less obvious result is that the peak of the normal pixel (Poisson) distribution (i.e., the mean dark current, excluding the hot pixels) has also increased.

The increase in the mean dark current for WFC has been from 6.8 e⁻/pixel/hr at launch to 11.1 e⁻/pixel/hr (an average of 11.4 for WFC1 and 10.8 for WFC2) in April 2004. For HRC the change in dark current has been from 9.3 e⁻/pixel/hr at launch to 13.4 e⁻/pixel/hr in April 2004. The dark calibration tracks the mean dark current very closely at 2 week intervals (see Section 3.4.3).

Figure 5.4: WFC dark current histogram for chip 1. Statistics for chip 2 are nearly identical. The figures illustrate the growth of hot pixels over time.

ACS / WFC1 dark current histograms
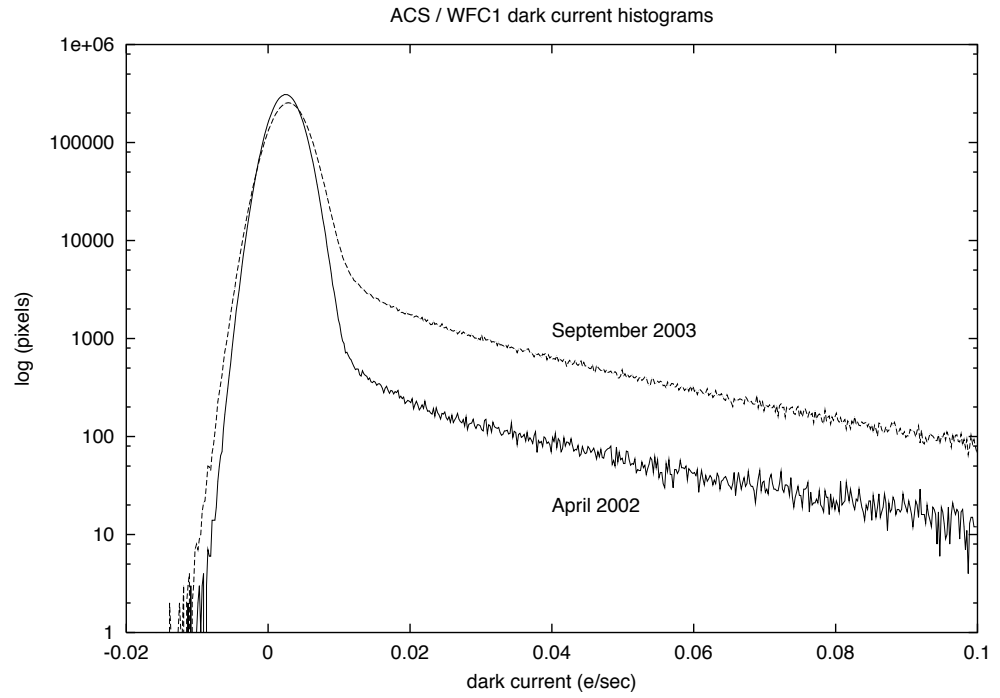
Figure 5.5: HRC dark current histogram

ACS / HRC dark current histograms
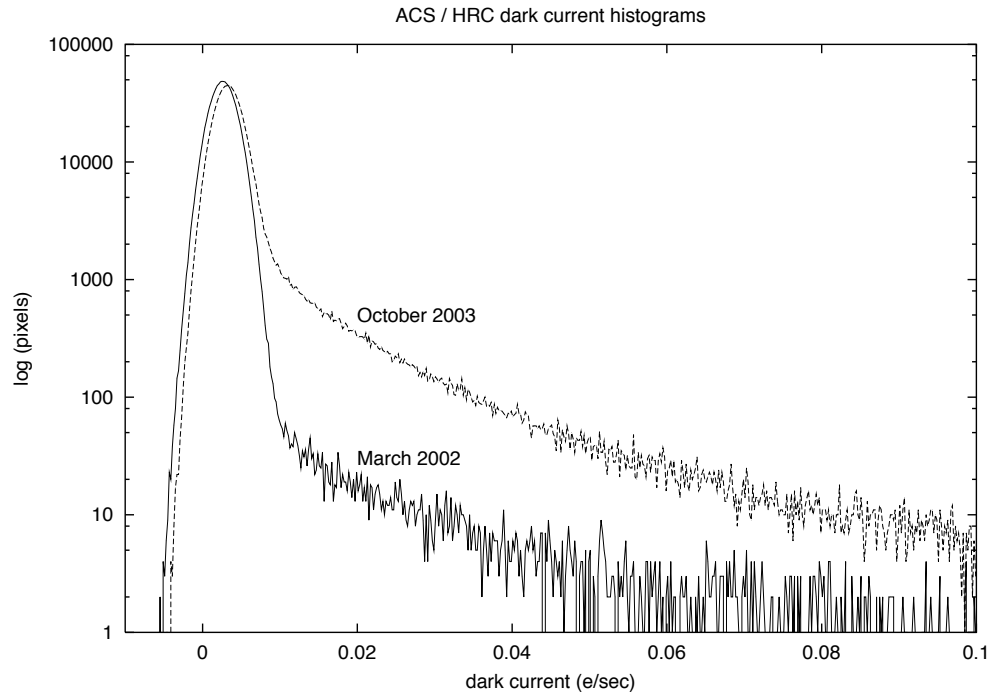
SBC

The SBC, or ACS MAMA detector, intrinsically has no readout noise and very low detector noise levels which normally will be negligible compared to statistical fluctuations. A dark current image was compiled during ground tests with a total exposure time of 14,400 seconds and remains as the calibration dark image. This has a mean dark count of 3.8×10^{-5} counts/pixel/second. Approximately once per month since ACS was installed, 1800 second SBC darks have been accumulated. The dark rate for these is typically 2.0×10^{-5} counts/pixel/second. The principal reason for taking these images is to provide a measure of the health of the SBC. No major fluctuation has been noted. An in-flight super-dark will be constructed from the in-flight measurements.

5.3.2 Hot Pixels

A "hot pixel" is a state of radiation damage affecting the dark current production of an individual pixel. Such pixels likely result from the close interaction of an incident heavy nucleon with Si nuclei in a pixel, creating new Si-SiO2 interface states. Once produced, such pixels almost always remain hot before the next anneal producing a continuum of excess dark current rates more than 1 to 4 orders of magnitude greater than the mean value. Although such pixels have temporarily lost their ability to yield precise measurements of faint photon flux, the majority of such pixels are "healed" in monthly anneals when the CCD temperature is raised by ~100 degrees over operating temperature for ~16 hours.

As seen in Figure 5.4 and Figure 5.5, the dark current distribution is well described by a Gaussian. As expected from experience with earlier HST cameras, very significant tails (all pixels with dark current above the 5 sigma limit [~0.02 e$^-$/sec]) are seen in these distributions. These tails consist of the pixels that have suffered radiation damage. The pixels with values between 0.02 and 0.08 e$^-$/sec are referred to as "warm" pixels and those above 0.08 e$^-$/sec as hot pixels. Figure 5.6 and Figure 5.7 show these tails for WFC and HRC for two different dates. As shown, many thousands (HRC) or tens of thousands (WFC) of pixels have greatly elevated dark current with the number markedly increasing with time.

Figure 5.6: Dark Current Distribution Tails in WFC for March 15, 2002 and March 28, 2002. All pixels with dark current above the 5 sigma limit (~0.02 e⁻/sec) are due to so-called "warm" and "hot" pixels which have suffered radiation damage. The number of these pixels can be seen to increase rapidly with time.
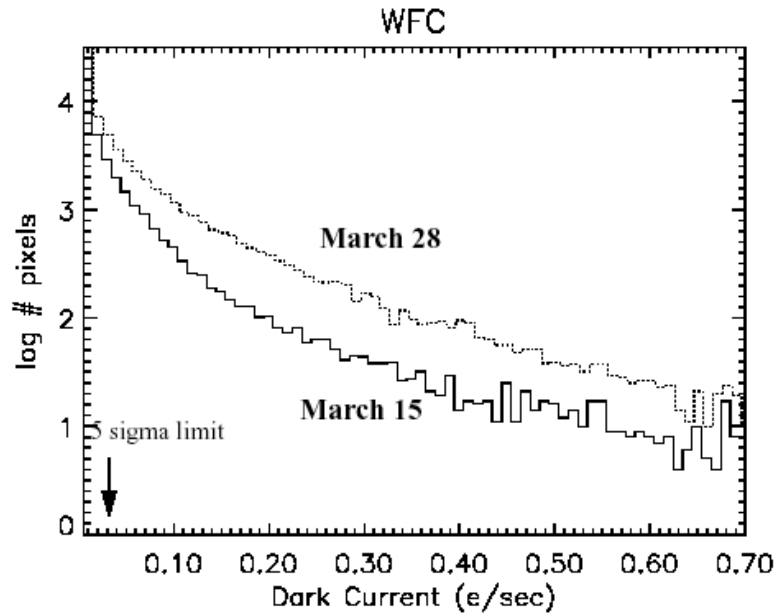
Figure 5.7: Dark Current Distribution Tails in HRC for March 15, 2002 and March 28, 2002.
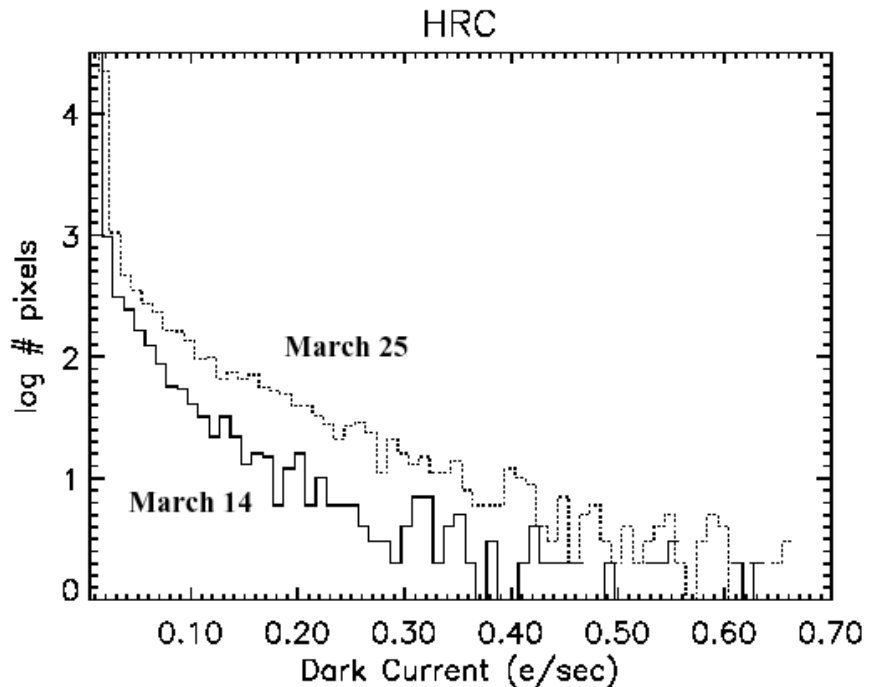
Trending

Even with 16 million pixels, the WFC would not be expected to have any pixels with dark current more than 6 standard deviations beyond the mean (i.e., with dark current greater than 0.02 e$^-$/sec) were it not for the effects of radiation damage. To analyze hot pixel trending, we have chosen a conservative limit of twice that, or 0.04 e$^-$/sec for WFC and 0.08 e$^-$/sec for HRC as a threshold above which we consider a pixel to be "hot" and not part of the normal distribution of pixel dark current. Note that even though trending is analyzed at these limits, only pixels above 0.08 e$^-$/sec are flagged in the DQ array.

For WFC we find a growth rate of approximately 1200 new hot pixels per day with dark current greater than 0.04 e$^-$/sec. For HRC the number of new hot pixels per day above the threshold is approximately 90. Because the distribution of dark current in hot pixels is strongly concentrated near the threshold, the specific number of such pixels is necessarily a strong function of the chosen threshold.

Subtraction of a superdark frame from a science image can remove the dark current from hot pixels just as it does for normal pixels. However, hot pixels are often orders of magnitude noisier than normal pixels, which in many cases limits their ability to provide useful measurements of flux. For the hottest pixels, the observed noise exceeds even the Poisson noise by an order of magnitude. Observations which do not employ contemporaneous superdarks will suffer from uncorrected and unflagged hot pixels, a particular concern to programs whose aim is to search for astrophysical transients or broad-band drop-outs. The problems associated with hot pixels can be mitigated by using the best reference files (usually necessitating recalibration with OTFR several weeks after the data were obtained), flagging and discarding data from hot pixels, and dithering observations to provide additional sampling of pixel positions on the sky and allow rejection of hot pixels.

Figure 5.8: Hot pixel growth rate for HRC (top) and WFC (bottom).

WFC, hot pix > 0.04 e/s
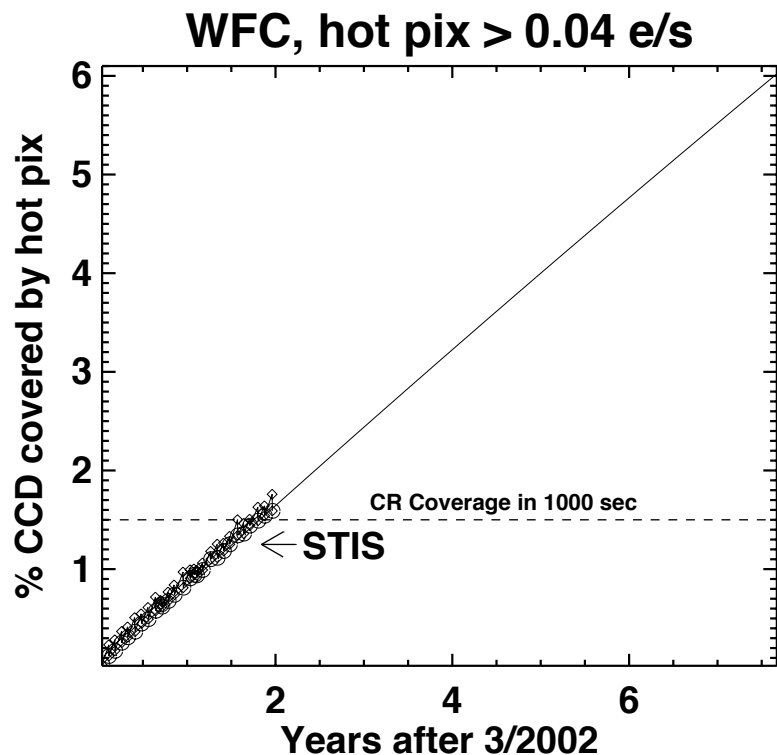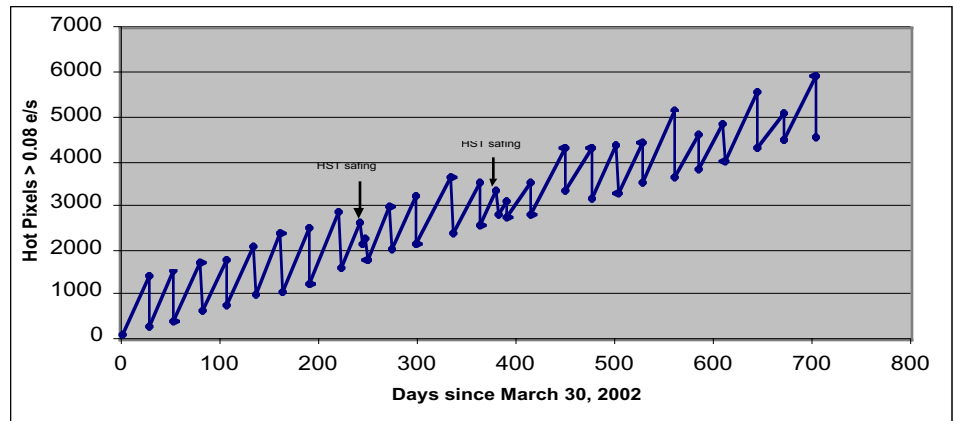
The anneal rate of new hot pixels (dark current > 0.04 e⁻/sec) on the ACS WFC has been disappointingly low: ~60% in the first eight monthly anneal cycles of the instrument. This rate is significantly lower than the observed and characteristic value for other HST CCDs: 80-85% for WFPC2, STIS, and ACS HRC. The anneal rate is significantly less than 60% for pixels which are much hotter than 0.04 e⁻/sec. The consequence of poor annealing is a greater fractional coverage of the camera by pixels with elevated dark current than was the experience of other HST CCDs. Figure 5.8 shows the observed growth of hot pixels on both HRC and WFC. The saw-tooth appearance is due to regular anneals. For WFC hot pixel coverage is more than cosmic ray coverage in a typical exposure. By the

end of the decade one in every 16 pixels is predicted to be hot on WFC. For more details see ACS ISR 02-09.

5.3.3 Cosmic Rays

Like all HST cameras before it, the ACS HRC and WFC fields of view 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.

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 is observed for WFPC2. For most science observations, a single CR-split (i.e., two exposures) is sufficient to insure that measurements of the targets are not compromised by cosmic rays.

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 sec, a single ~2400 second orbit must be broken into 4 exposures (4 CR-splits of 500 to 600 sec each) to reduce the number of uncleanable pixels to 1 or less. (It is also recommended that users dither these exposures to remove hot pixels as well.)

Cosmic Ray Events

Many science observations require a careful consideration of individual cosmic rays events. To either remove cosmic rays events 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.

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 5.9, 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 5.9: Electron deposition by cosmic rays on HRC (top) and WFC (bottom). A minimum deposition of ~500 e is seen corresponding to CR's with normal incidence. The median deposition is ~1000 e.
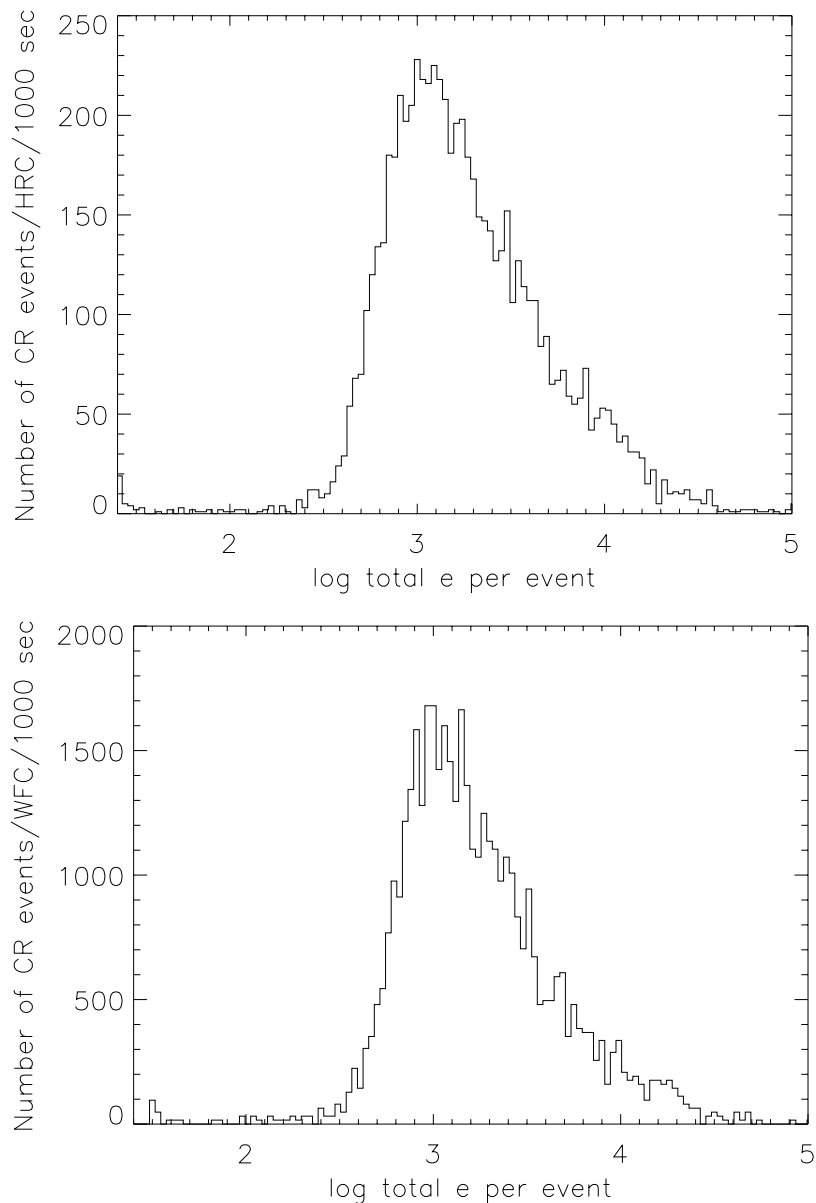
Attached Pixels

As seen in Figure 5.10 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 which 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. Such pixels are very rare but do exist. 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 which 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 5.10: Distribution of the number of pixels associated with a single CR 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.
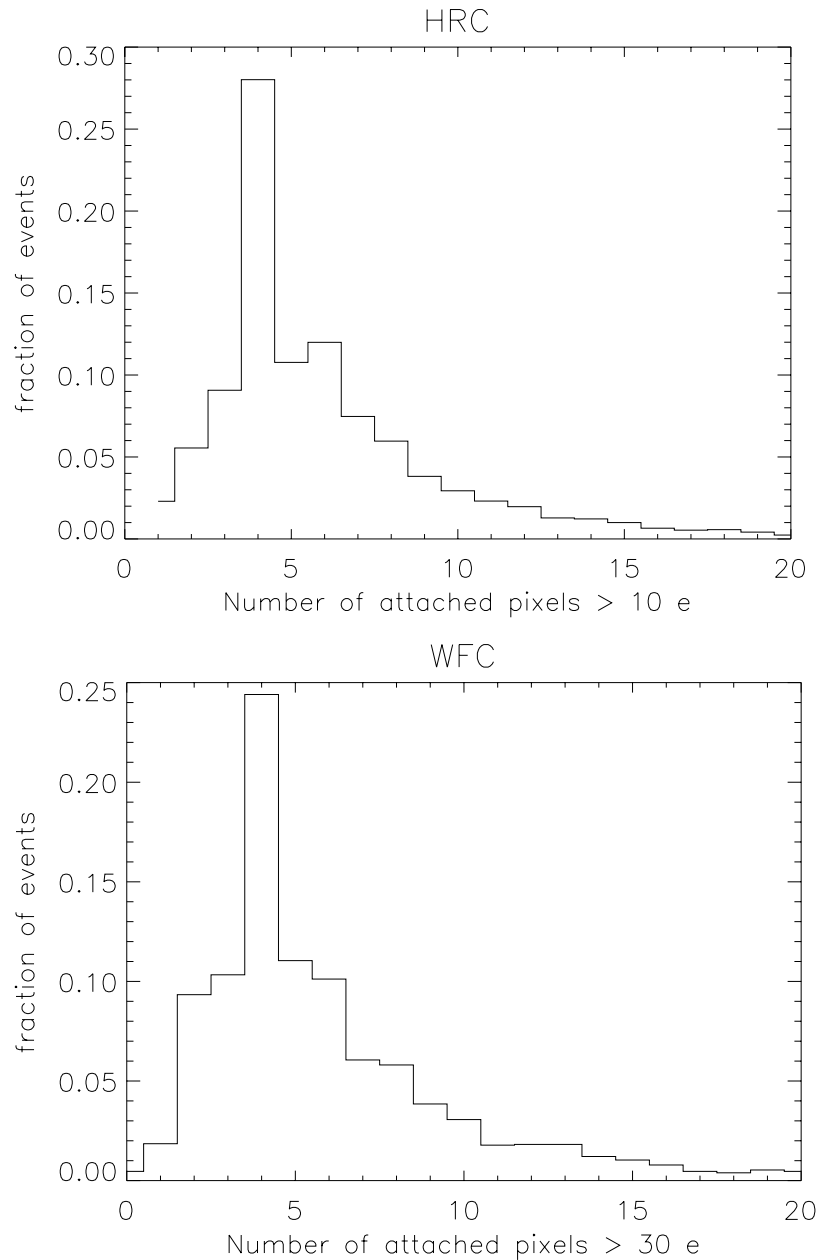
5.4 Flat Fields

5.4.1 Ground Flats (P-flats)

In early 2001, flat field images for the ACS were produced in the laboratory using the Refractive Aberrated Simulator/Hubble Opto-Mechanical Simulator (RAS/HOMS). The RAS/HOMS is an 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 03-02) were created using inflight observations of the bright earth. (See Section 5.4.4 for more information on the ACS Earth flat program.)

The resulting flat fields include both the low frequency (L-flat) and the high frequency pixel-to-pixel (P-flat) structure. 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 chip 2 images are divided by the chip 1 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.

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. Thus, point source photometry extracted from a flat fielded image must be multiplied by the effective pixel area map (see Section 6.1.3). This correction is accounted for in pipeline processing by **PyDrizzle** (see Section 3.4.3), which uses the geometric distortion solution to correct all pixels to equal areas. In the drizzled images, photometry is correct for both point and extended sources.

More information on the HRC and WFC ground flats may be obtained from ACS ISR 01-11 for the standard filters, polarizers, and coronagraph, from ACS ISR 02-01 for the ramp filters, and from ACS ISR 02-04 for the prism and grism. Ground flats for the SBC detector are described in ACS ISR 99-02.

5.4.2 Inflight L-flat Correction

The large scale uniformity of the WFC and HRC detector response, as provided by the ACS laboratory flats, has been improved inflight 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. Photometric errors of up to +/-9% corner-to-corner have been found in the original WFC and HRC laboratory flat fields (see ACS ISR 02-08). The derived L-flats are based on a 4th-order polynomial fit and are shown in Figure 5.11 and Figure 5.12 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 inflight 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. For a detailed discussion of ACS L-flat corrections, refer to ACS ISR 02-08. For a discussion of the mathematical algorithm used to derive the L-flats, refer to ACS ISR 03-10.

To verify the interpolated L-flat corrections, additional observations of the same stellar field have been acquired and are undergoing analysis. The L-flat calibration program revisits the same target several times per year; therefore, observations are obtained for a range of orientations due to the roll of the telescope. Differential photometry of stars falling on different portions of the detector as the telescope rolls will provide a verification of the L-flats for the interpolated filters and will also provide an independent test of the absolute sensitivity dependence with time for the full suite of ACS filters. Initial testing indicates that the photometric response for a given star is now the same to ~1% for any position in the field of view for filters which were observed during the inflight L-flat campaign. For the interpolated L-flats, the expected accuracy is ~2-3%. Once the analysis for the additional stellar observations is complete, new flat fields will be delivered to the pipeline and the errors are expected to be reduced to <1%.

Figure 5.11: 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 +/- 5% for the F555W filter to +/- 9% for the F850LP filter.
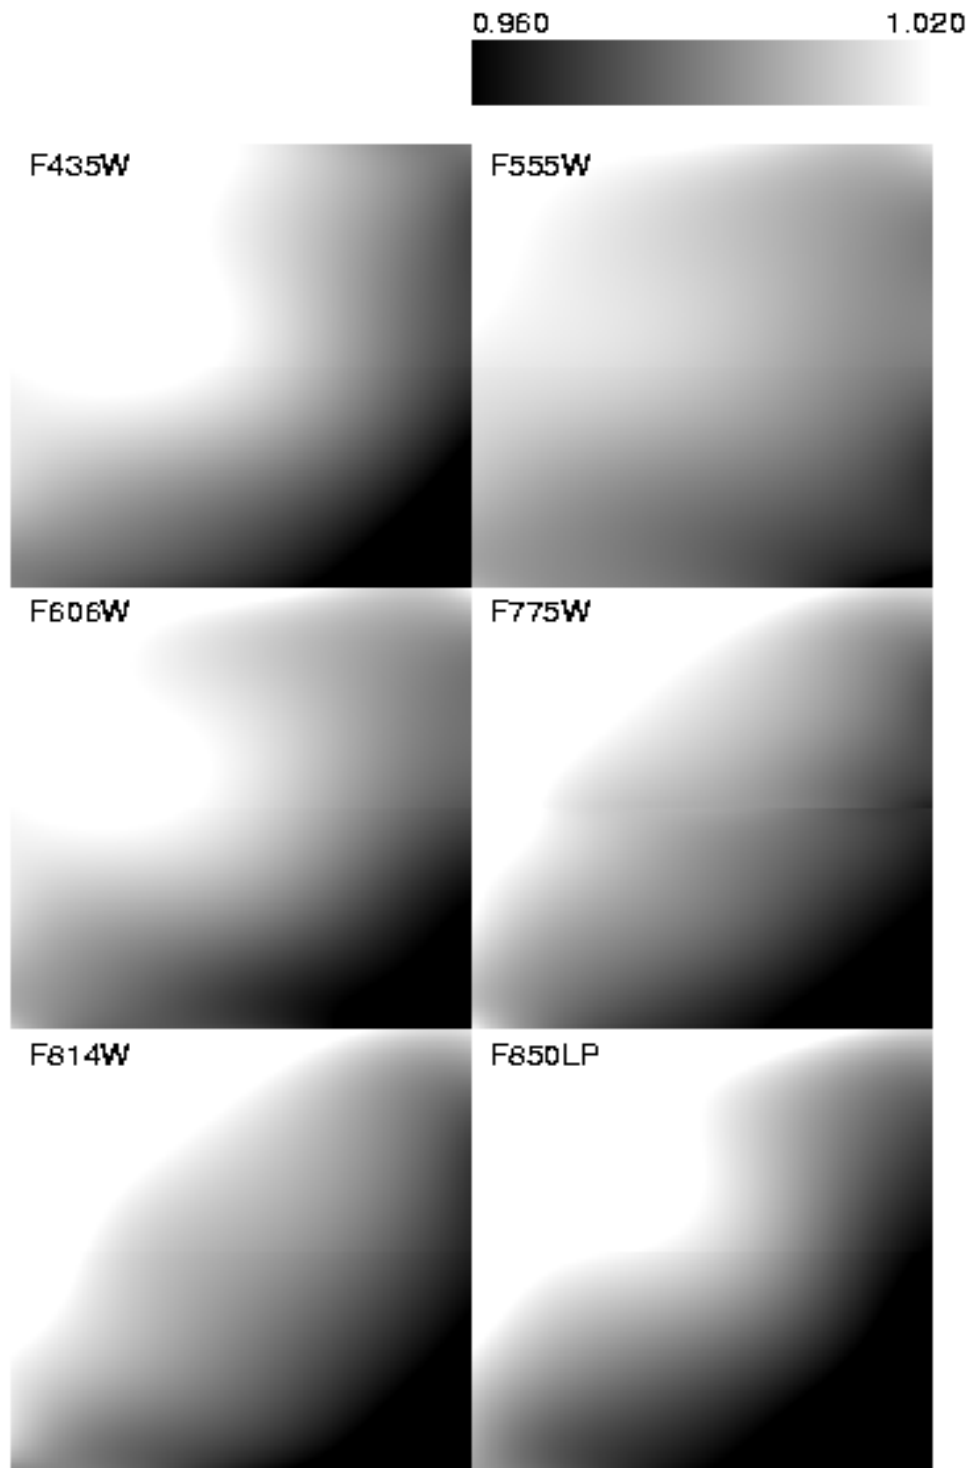
Figure 5.12: HRC low frequency (L-flat) flat-field corrections required for the laboratory flats. This correction ranges from +/- 3% for the F555W filter to +/- 6% for the F850LP filter.
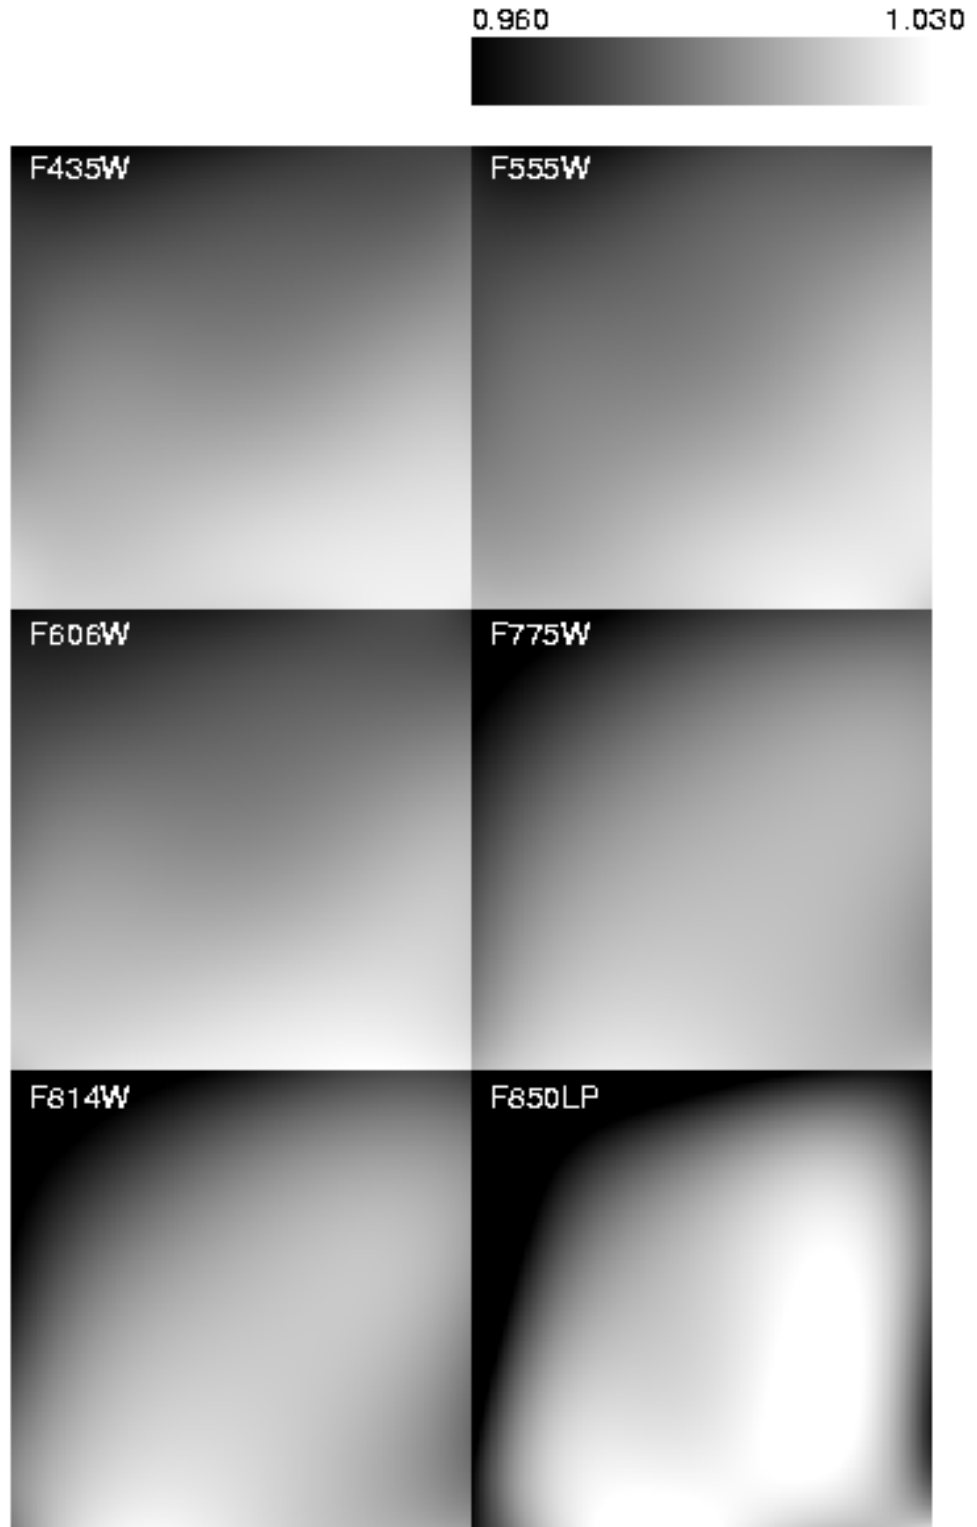
5.4.3 Pipeline Flats

The inferred low-frequency (L-flat) corrections were applied to the laboratory flats, and the resulting flat fields have been in use in the calibration pipeline. Observers are encouraged to check the ACS web site for the most recent flat fields available for recalibration at http://www.stsci.edu/hst/acs/analysis/reference_files/flatimage_list.html

Figure 5.13 shows the corrected WFC laboratory flats for several broadband 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 03-06) 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.

The HRC corrected laboratory flats for the same broadband filters are shown in Figure 5.14. 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 5.5.1). 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.

The accuracy of pipeline flats can be verified using a variety of complementary methods. In Section 5.4.2, we discussed 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 5.4.4. 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 broadband WFC filters and are discussed in detail in Section 5.4.5. The sky flats are generally similar to the 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 laboratory 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. For a few filters, improved L-flats will be derived from additional stellar observations to improve the flat field errors to <1% for all modes. Additional verification of the pipeline flats will be provided by both earth flats and sky flats, thus confirming the derived L-flats and setting limits to their dependence on the color of the source spectrum.

Figure 5.13: WFC Flat Fields
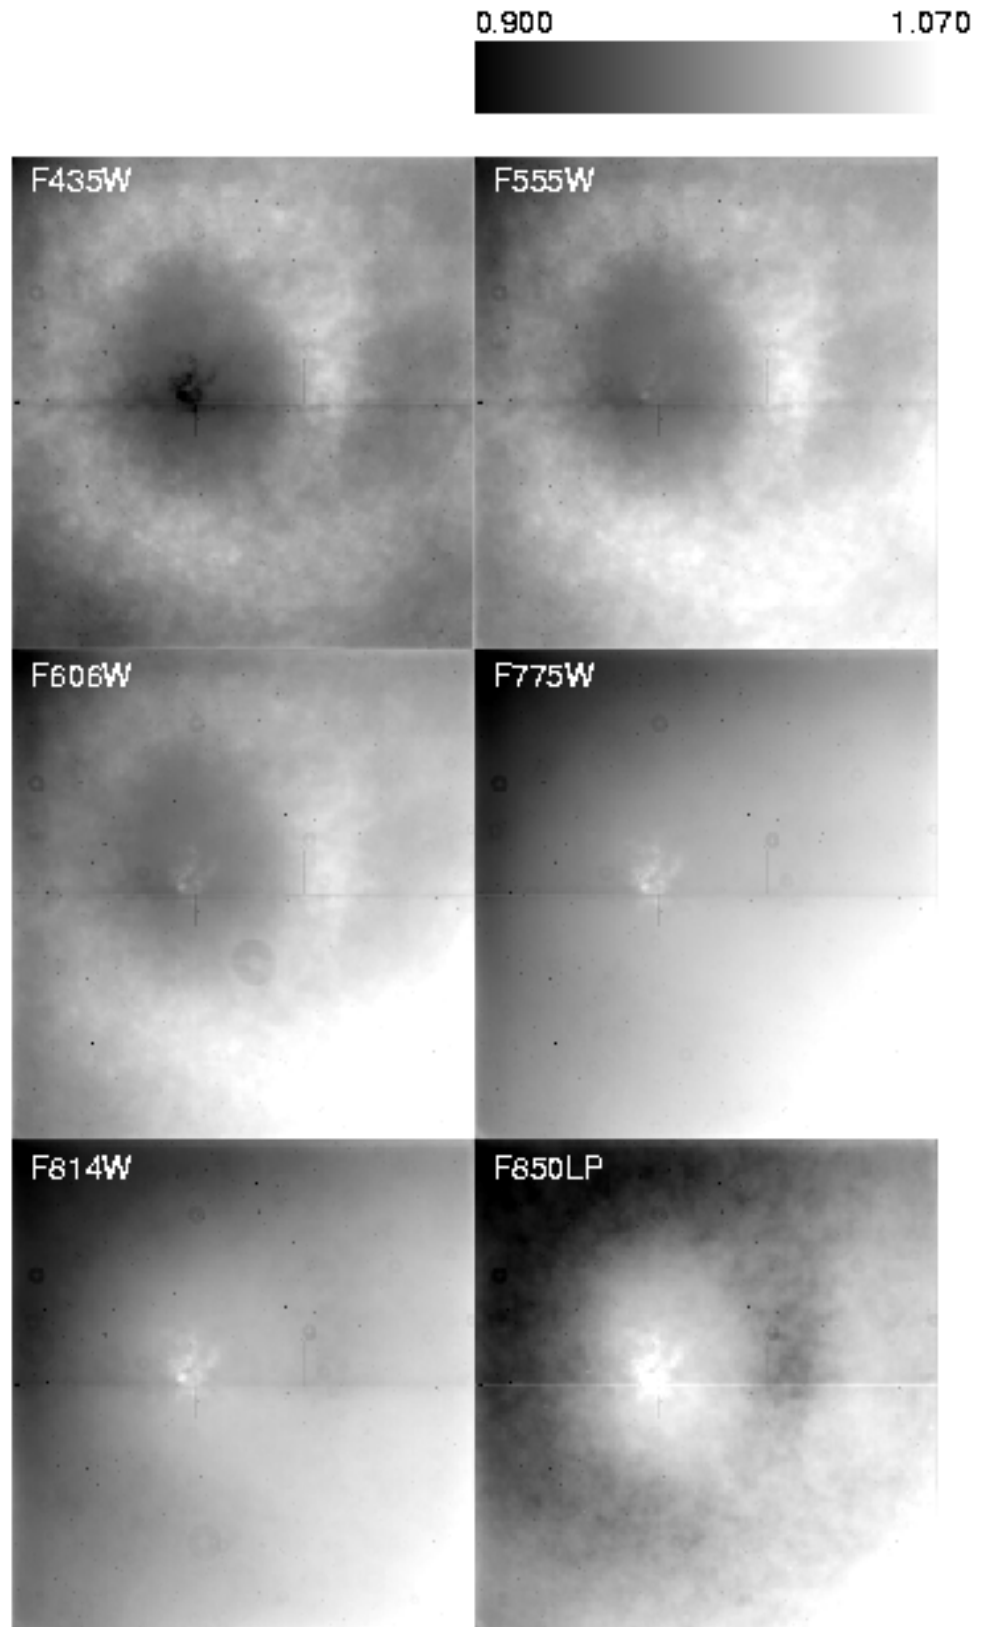
Figure 5.14: HRC Flat Fields
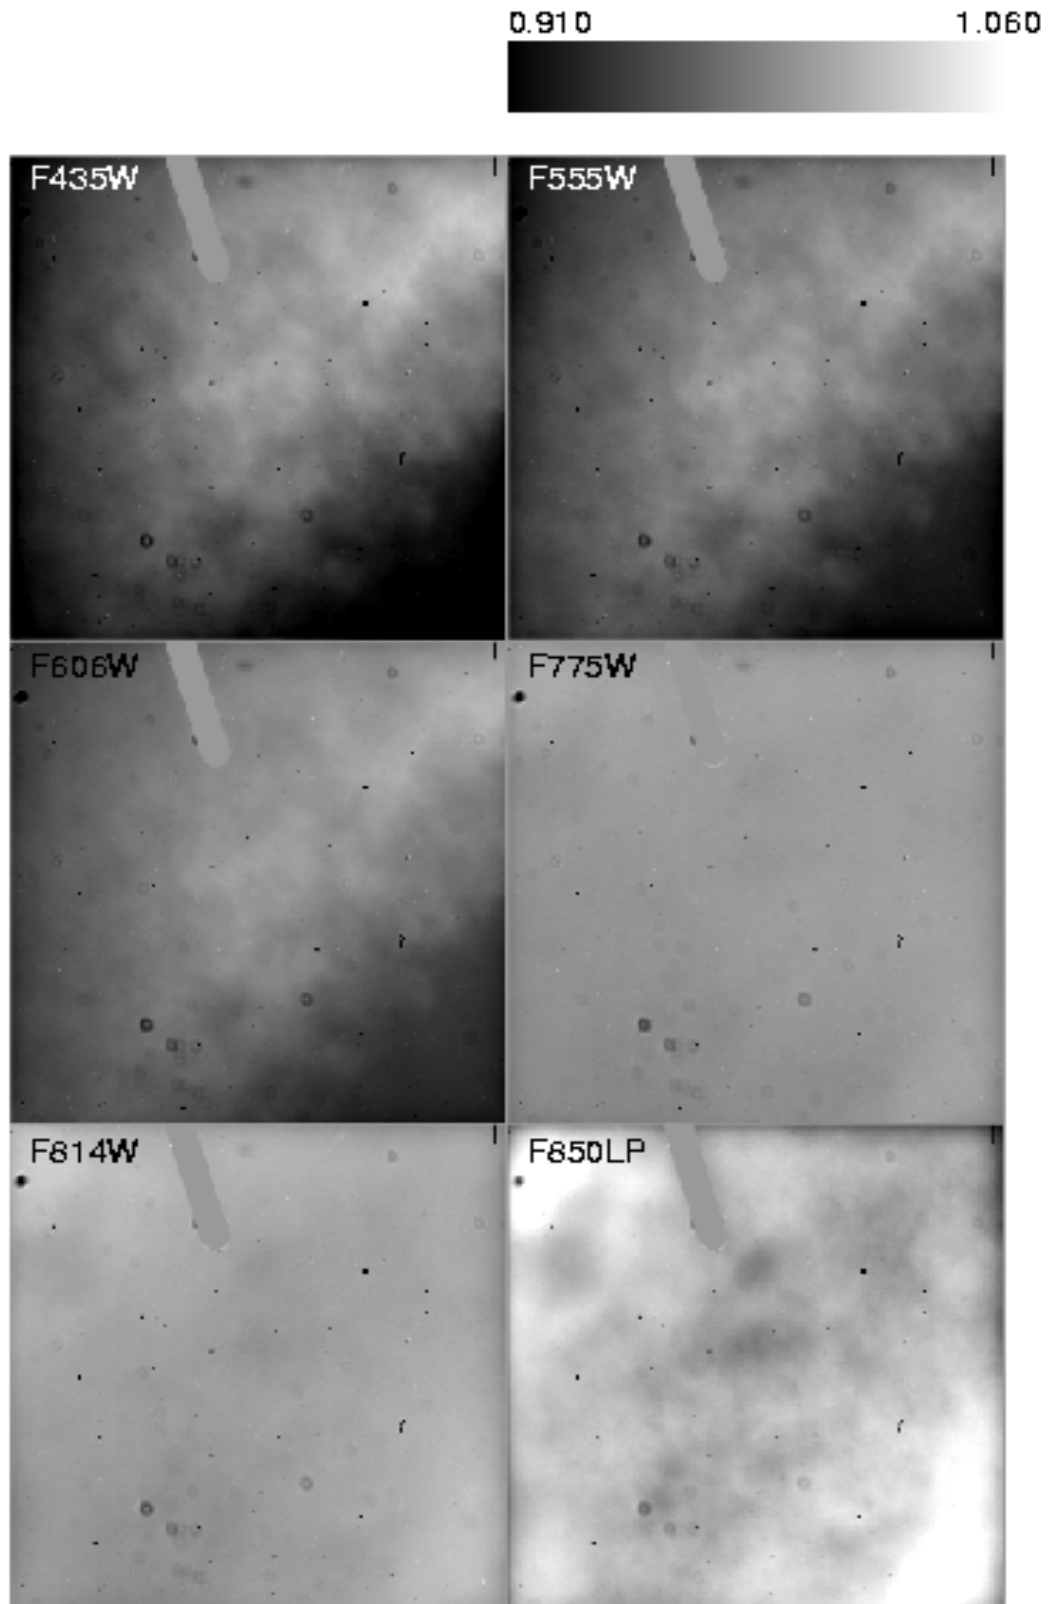
5.4.4 Earth Flat Verification

Because the RAS/HOMS optics are opaque below 3500 Å (see Section 5.4.1), flat fields for the UV filters were created using inflight 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 utilize the F220W, F250W, F330W, and F344N are immune to streaks because of the large optical depth down to the tropospheric cloud layers. The bright Earth then provides a uniform flat field source for the complete OTA+HRC optical complement. The required calibration flat, 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 dominates, and the lab flats made with the deuterium lamp illumination (see ACS ISR 01-11 for details) are superior to the observed earth flats for the modes that include these two filters. Because no L-flat correction has been applied, errors in the flats of +/- 2 to 4%corner-to-corner are expected for these two filters.

The HRC earth flats at wavelengths longward of 4000 Å often show streaking and non-uniformities from clouds or the terminator. However, a number of the images are free from these defects and provide a promising alternative to sky flats as well as an independent verification of the stellar L-flats. A comparison of the current pipeline flats to the new HRC earth flats shows differences of typically less than ~1% for the broadband filters and ~2% for the interpolated L-flat filters. An ISR describing these results is in preparation and will be posted on the ACS site:

http://www.stsci.edu/hst/acs/documents/isrs.

At the shortest exposure time of 0.5 sec, most WFC modes saturate, except for a few narrow band filters (i.e., F502N, F658N, F660N, and F892N) which are only saturated approximately half the time. Earth flats in these narrow band filters would nicely compliment the stellar L-flat program and sky flat results, which use only broadband filters. More observations in the WFC are required to produce good enough earth flats for comparison with the pipeline flats.

5.4.5 Sky Flats

The ACS sky flat program takes advantage of the many survey programs that have executed on ACS, such as the GOODS program, to build high signal-to-noise flats. These non-proprietary data consist mostly of sparsely populated images with relatively uniform sky that can be stacked to further quantify the pixel-to-pixel variation of the instrumental response.

The GOODS program consists of either 15 or 16 unique pointings over 5 epochs for both the Chandra Deep Field South (CDFS) and the Hubble Deep Field North (HDFN) using the most commonly used F606W, F775W

and F850LP filters in the WFC. In total between 50-70 images were stacked in each filter.

The sky flats were created by median combining the pipeline reduced FLT files after removing cosmic rays and masking all of the sources. Because the GOODS data contained 2-4 dithers at each pointing, depending on the filter, 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.

The combined sky flat in each filter was box medianed down to a 32x32 image, in order to directly compare it to the corrected ground flat. Since the sky flats were created starting from the pipeline calibrated FLT files, it will show any flat field signatures that were not accounted for by the pipeline.The resulting sky flats show variations across the field of view of <2% for each of the 3 filters. The CDFS and HDFN fields were analyzed separately and compared. The two fields showed excellent agreement (<1%) for all 3 filters. The inferred sky flats will be made available on the ACS web site and a future ISR will describe the sky flats in more detail. Please check the ACS web site for updates on this issue.

5.5 Image Anomalies

5.5.1 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. These motes can be seen in the WFC and HRC flats in Figure 5.13 and Figure 5.14. Since the shapes and depths of these motes are almost independent of wavelength, their effects will be removed by the flats to <<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 inflight 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 ~350-400 pixels on the WFC and ~250-300 pixels on the HRC. One of these large motes can be seen on chip 2 of the WFC F606W flat in Figure 5.13.

Until April 2004, the positioning accuracy of the filter wheels has been within +/- one motor step 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 misregistration 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 5.4.2.

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 03-11.

5.5.2 Optical Ghosts and Scattered Light

The ACS was designed with a requirement that no single straylight feature may contain more than 0.1% of the detected energy in the object producing it. This goal has generally been met, but during the extensive ground and SMOV test programs a few exceptions have been identified (Hartig et al. 2002, Proc SPIE 4854) 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. We describe below in more detail some of the optical ghosts relevant for ACS.

- WFC elliptical haloes: These ghosts are caused by reflection from the CCD surface (which lies at an ~20 degree 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 5.15), 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 5.16). 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 ~2% of the detected target energy and are always ~10 and 20 pixels in diameter.

• Annular ghosts: Large annular ghosts near their parent images are caused by reflection from the detector windows, back to the filters, then returning to the CCD. The fringes arise from interference between the HeNe laser light reflections from the two surfaces of the windows (Figure 5.15). 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.

In general, little can be done about ghosts and blooming in the post-observation data processing phase. Instead, some judicious planning of the actual observations, particularly if bright sources are expected in 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 ORIENT which prevents the source of interest from being connected to the bright star along either of these axes. Alternatively, a suitable ORIENT could move the bright star(s) into the interchip gap or off the field of view altogether. Similarly, the impact of WFC elliptical haloes can be minimized by avoiding a bright star in the quadrant associated with amplifier D.

Figure 5.15: Optical ghosts, diffraction spikes, and blooming. This image was obtained during ground calibration at Ball Aerospace using a HeNe laser (633nm) through F625W. It is a factor of ~50,000 saturated with 9 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 for details) are visible: (a) WFC elliptical haloes show up in the lower left as pairs of elliptical annuli aligned along the negative diagonal of the FOV; (b) Large annular ghosts are seen near their parent images; (c) To the lower right, there exist smaller fainter annular ghosts which tend to be displaced radially from the parent image by a small amount.
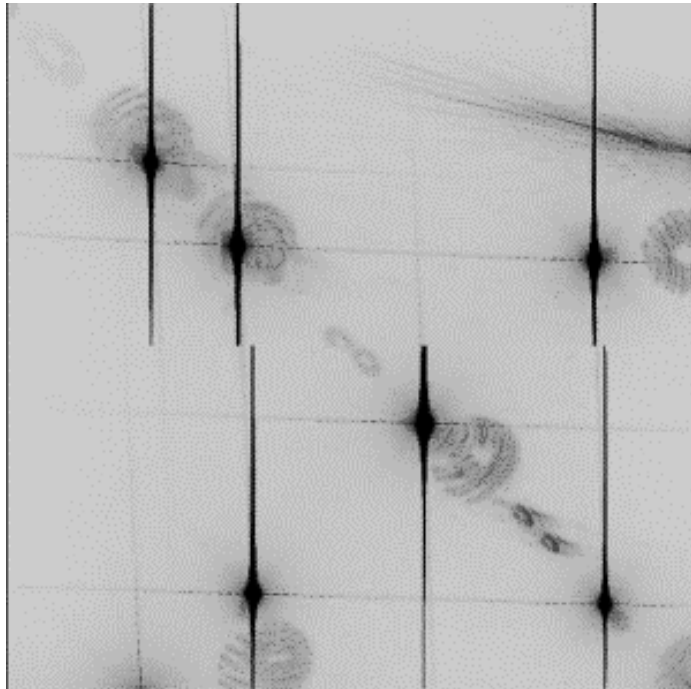
Figure 5.16: F660N optical ghosts. This 700 s F660N exposure illustrates the F660N ghosts, namely pairs of relatively bright circular annuli stationed near to (but radially outward from) the target image. These haloes contain ~2% of the detected target energy and are always ~ 10 and 20 px in diameter.
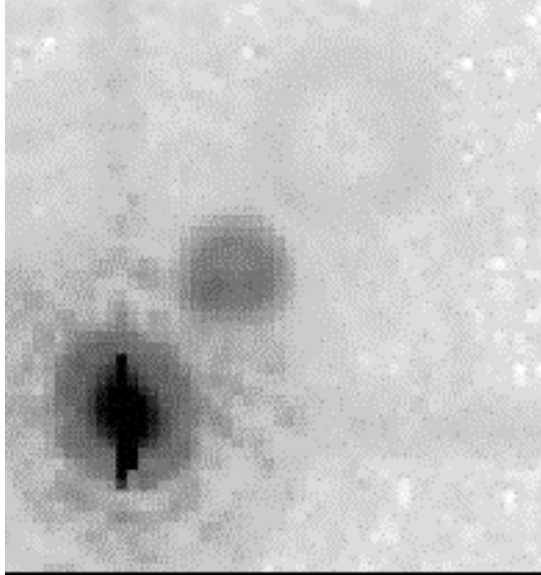
5.5.3 Cross Talk

Images obtained with the WFC are affected by a small amount of electronic cross talk between the four CCD quadrants that correspond to the four amplifiers of the two detectors. The effect produces electronic "ghost" (i.e., negative) images in a given quadrant that mirror real images recorded on other quadrants. Figure 5.17 shows an example of the cross talk in one of the images obtained during the Great Observatories Origins Deep Survey (GOODS) program, a pointing in the GOODS-S field taken through the F850LP passband. Source images a, b, c and d, which are recorded on quadrant C, produce negative ghost images a', b', c', d', a'', b'', c'', d'', and a''', b''', c''' and d''' in quadrants A, B and D, respectively.

Inspection of images affected by cross talk reveals a number of interesting characteristics of this effect. The cross talk ghosts have a surface brightness that is a few electron lower than those of the average background and there does not appear to be a simple proportionality between the light profile of the sources and the ghosts. Rather, as Figure 5.18 shows, the strength of the cross talk, measured as the average surface brightness of ghosts, seems to be roughly proportional to the logarithm of the sources.

The intensity of the cross talk ghost also appears to vary depending on which quadrant hosts the source and which hosts a ghost. For example, sources in quadrant D of Figure 5.17 produce clearly visible ghosts in quadrant C, but much fainter ghosts (if any at all) in quadrant A and B. Also, sources in quadrant A and B do not seem to produce any obvious ghosts in the other quadrants.

Figure 5.17: Cross Talk in the WFC Detectors. Due to the cross talk, galaxies a, b, c and d recorded in quadrant C have mirror ghost images in quadrants A, B, and D. Note the effect appears stronger in quadrants C and B than in quadrant A, and that sources in quadrant D produce ghosts that follow a similar trend. However, sources in quadrants A and B do not seem as effective in producing ghost images.
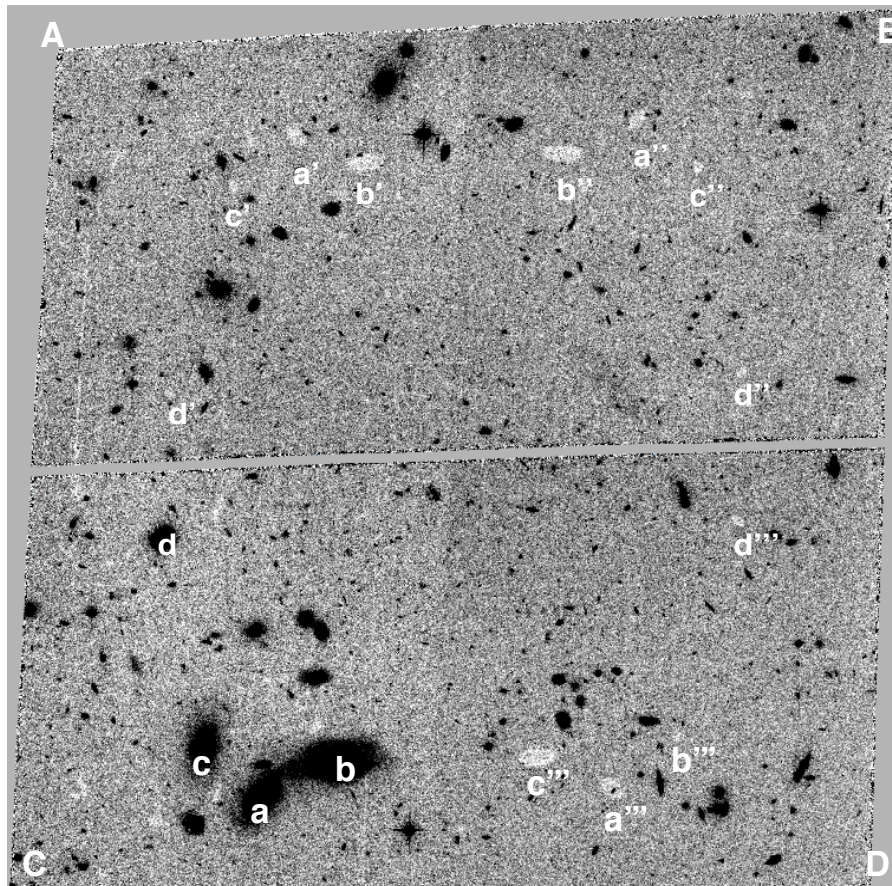
Figure 5.18: The average surface brightness of ghost images plotted as a function of the flux of the sources. All sources are in quadrant C. There is a correlation between the surface brightness of the ghosts and the flux of the sources.
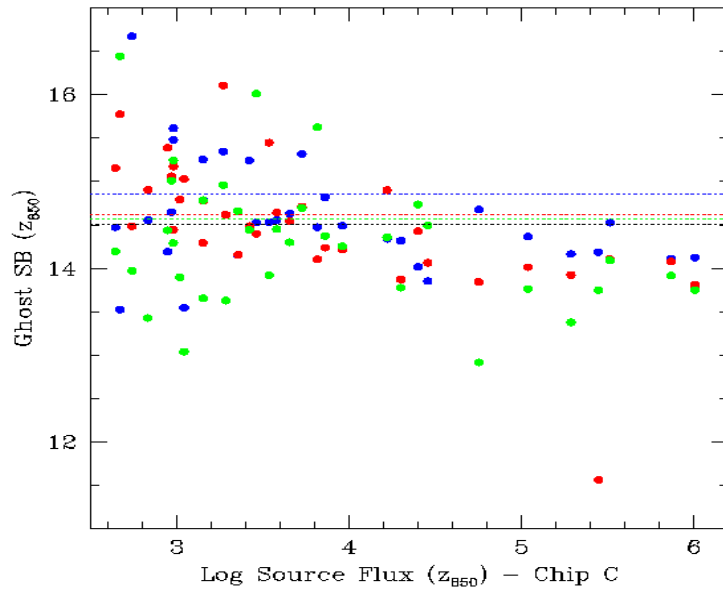
The magnitude of the cross talk depends on the background of the victim quadrant, the sense of the correlation being that the cross talk is stronger when the background is lower. Figure 5.19 shows that in images with low background the cross talk creates ghosts whose surface brightness is ~40% lower than the average background. In images with high background the ghost surface brightness is only ~3% lower.

Tests on existing images show that the cross talk has minimal impact on flux measurements of sources superimposed on ghost images and thus mostly presents a cosmetic issue. The bottom panel of Figure 5.20 shows six growth curves of the same star derived from six individual frames, all with identical exposure time and filter. In three of the frames, due to the adopted dithering strategy, the star is located within a major ghost image (the annulus around the photometric aperture is also within the ghost image), while in the remaining three frames the star is outside the ghost image. The observed photometric scatter is typical for a star with this brightness of post-pipeline processed data (FLT files) with no correction for geometrical distortions and no cosmic ray removal. As a comparison, the upper panel shows the growth curves of a star of comparable apparent magnitude in the same frames but located in an region unaffected by cross talk. The photometry of the latter star has the same amount of scatter as the former one, showing in particular that the cross talk acts as an additive effect and not as a multiplicative one.

Figure 5.19: The dependence of the cross talk, measured as the difference of the sky background in affected and non affected areas, as a function of the sky background. The strength of the cross talk diminishes as the background increases.
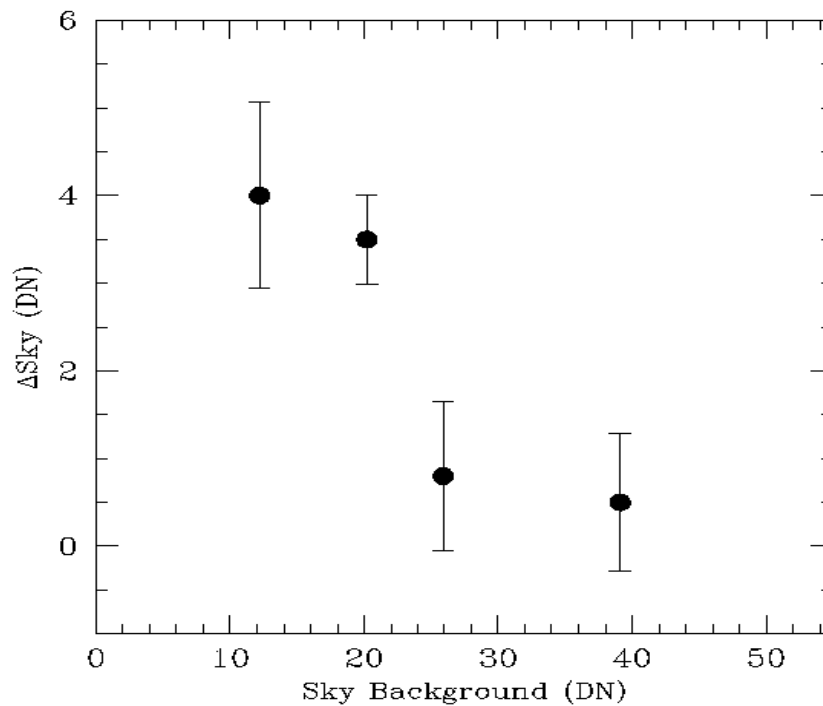
Figure 5.20: BOTTOM - Six growth curves of the same star derived from six individual frames, all with identical exposure time and filter. In three of the frames, due to the adopted dithering strategy for the observations, the star is located within a major ghost image (the annulus around the photometric aperture is also within the ghost image), while in the remaining three frames the star is outside the ghost image. The observed photometric scatter is typical, for a star with this brightness, of post-pipeline processed data (FLT files) with no correction for geometrical distortions and no CR removal. TOP - As above but for a star of comparable apparent magnitude and in the same frame that is located in a region unaffected by cross talk. The photometric scatter of the two stars is very similar.
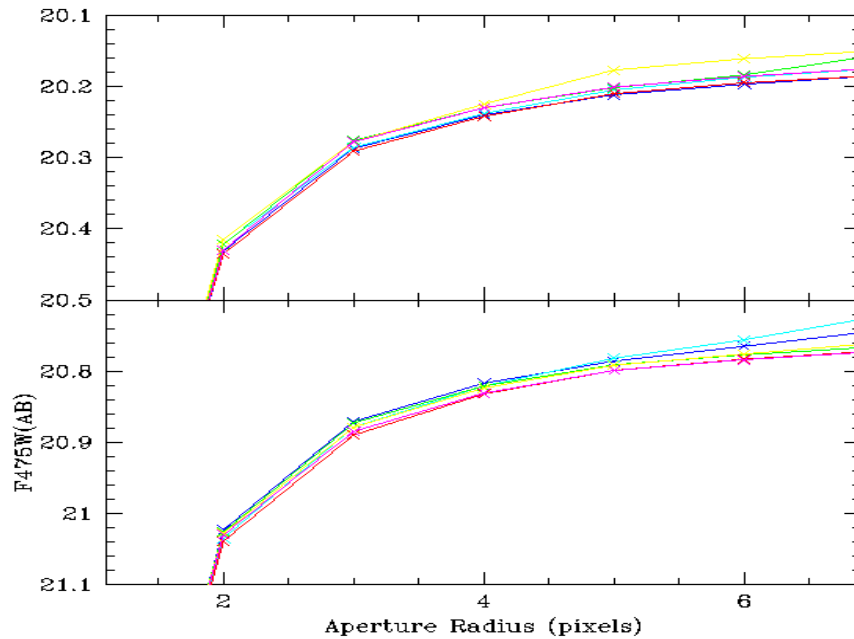
There is evidence that the strength of the cross talk is significantly weaker in images acquired with GAIN=2, than in images taken with GAIN=1. Figure 5.21 shows two exposures through the same filter and with the same exposure time, the former acquired with GAIN=1 (upper panel), the second with GAIN=2 (lower panel). Cross talk ghosts are clearly observed in the former image, but not in the latter one. A more quantitative comparison supports the visual inspection. Figure 5.22 shows the median value of pixels in ghost images as a function of the source pixel flux, i.e., of pixels located in mirror symmetric locations relative to the pixels in the source quadrant.

Observers who do not need to sample the readout noise of the CCD with the accuracy afforded by the GAIN=1 setting, might want to consider to use GAIN=2 in order to minimize the cross talk.

At this time we have not yet determined a robust way to predict the size of a cross talk residual and the exposure parameters on which it depends. Efforts are ongoing to identify algorithms to remove the effect of cross talk from the WFC images and we encourage GO's to look for updates on the ACS STScI web page.

Figure 5.21: Two exposures through the same filter (F814W) and with the same exposure time (22.5 sec), the former acquired with GAIN=1 (TOP), the second with GAIN=2 (BOTTOM). Cross talk ghost images are observed in the former image, but not in the latter one.
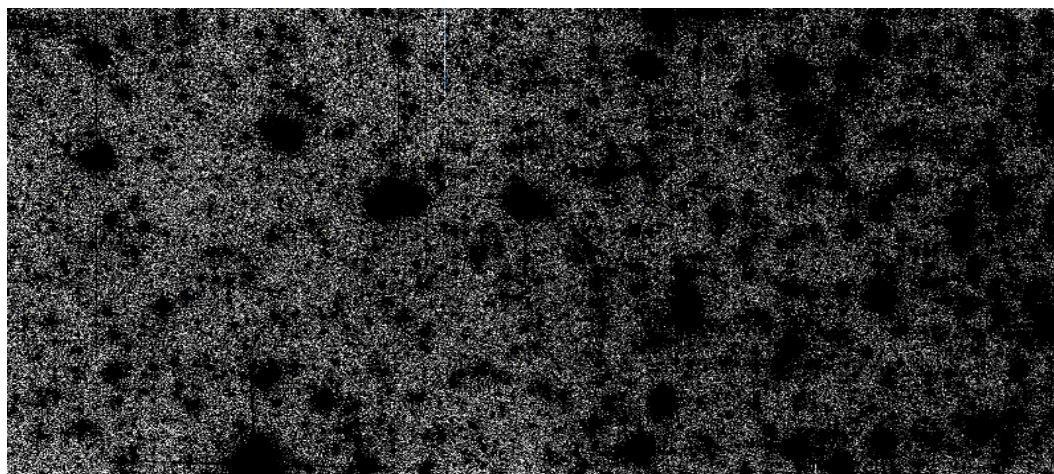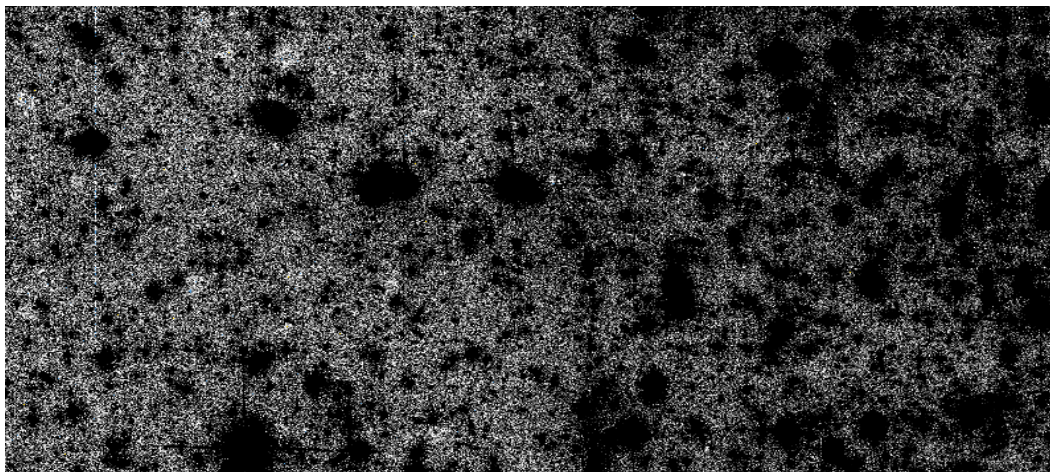
Figure 5.22: The median value of pixels in ghost images as a function of the source pixel flux, i.e. of pixels located in mirror symmetric locations relative to the pixels in the source quadrant, derived from the CCD frames shown in Figure 5.17. The images are sky subtracted. Source pixel flux is binned in logarithmic bins 0.2 in size. Using the median ensures that mostly background pixels, as opposed to pixels affected by sources, are being considered in the analysis. Filled symbols are for GAIN=2, hollow one for GAIN=1. The pixels in the image with GAIN=1 have systematically lower values than those of the image with GAIN=2, an indication that the cross talk is stronger in this case. Their values also decrease as a function of the source pixel value more steeply than in the GAIN=2 case.
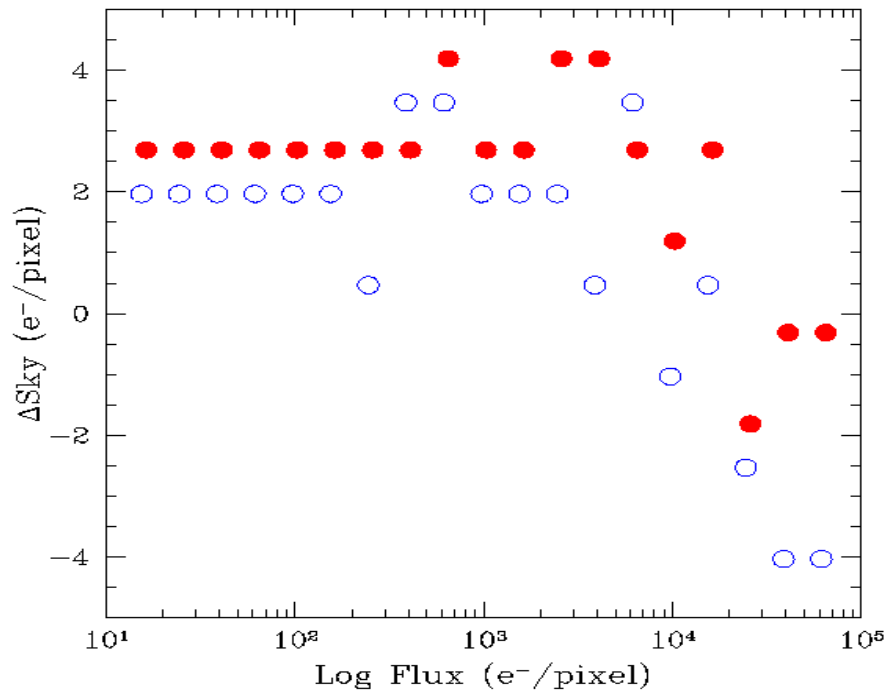
5.5.4 Scattered Earth Light

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 degrees 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 CVZ 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 degrees 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 degrees, it is possible to have both large scale gradients across the field of view (up to ~20% amplitude in the WFC) and small scale features in the background (up to ~12% in WFC and ~30% in the HRC). See ACS ISR 2003-05 for more details.

5.6 Generic Detector and Camera Properties

5.6.1 Gain Calibration

The ACS CCDs have default gain values of approximately 1 e-/DN (WFC) and 2 e-/DN (HRC) -- these settings have, in particular, been used to establish adjustments to basic quantum efficiency curves. 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 gains do not sample the full well depths of the CCDs, falling short by about 35% (WFC), and 22% (HRC). Use of the next higher gain values of approximately 2 e-/DN (WFC) and 4 e-/DN (HRC) do provide full sampling of the roughly 85,000 e- and 165,000 e- full well depths, respectively. Since the readout noise is only marginally higher than with the default gains, and is well-sampled even at these higher values (true for WFC, not quite for HRC, but even here its much better than for any WFPC2 data ever taken), many science programs and calibration programs may logically choose these gains.

ACS gain values in use between on-orbit installation 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 in the Appendix at

http://acs.pha.jhu.edu/instrument/calibration/results/by_it
em/detector/wfc/build4/gain/

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 readout 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

linearity (which depends on the constancy of the lamp source, or experimental techniques to control for drifts) and the saturation count level (which may be different for a uniform illumination pattern than for point sources) 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 as given at

http://acs.pha.jhu.edu/instrument/calibration/results/by_it em/detector/hrc/build1/gain/

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 redetermination of the WFC gains (ACS ISR 02-03), 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 5.2), rather than errors in relative gains.

In ACS ISR 04-01 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 here 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).

High Resolution Camera

The prelaunch GAIN = 4 value using the default AMP C is 4.289. The revised value based on on-orbit calibrations is 4.235. For GAIN = 1 (less important since this is available-but-unsupported), the prelaunch value of 1.185 has been revised to 1.163.

Wide Field Camera

Revisions of WFC GAIN = 2 values based on on-orbit testing are for prelaunch (new): AMP A 2.018 (2.002), AMP B 1.960 (1.945), AMP C

2.044 (2.028), and AMP D 2.010 (1.994). And for the available-but-unsupported GAIN = 4 the prelaunch (new) values are: AMP A 4.005 (4.011), AMP B 3.897 (3.902), AMP C 4.068 (4.074), and AMP D 3.990 (3.996).

5.6.2 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 04-01, 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 to neighboring pixels in the column begins.

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 Camera

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 representing a rough estimate at an area weighted average value. See Figure 3 of ACS ISR 04-01 for details.

Wide Field Camera

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 80,000 e- to 88,000 e- with a typical value of about 84,000 e-. There is a significant offset between the two CCDs. The spatial variation may be seen as Figure 4 in Gilliland ACS ISR 04-01.

5.6.3 Linearity at Low to Moderate Intensity

High Resolution Camera

Rich star fields observed at quite different exposure times provide a simple, direct test for linearity. In Figure 5.23 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 exposure -- 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 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 5.23: 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.
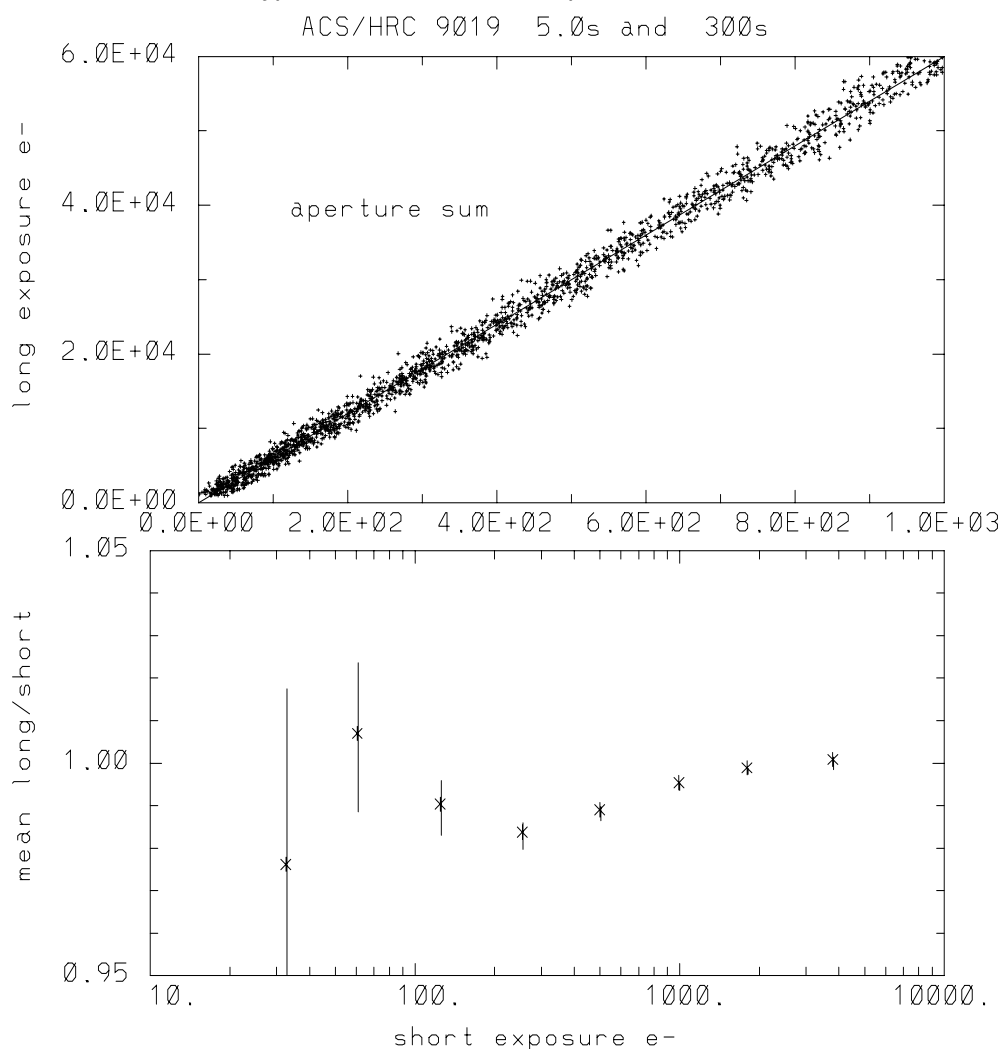
Wide Field Camera

As shown in Figure 5.24 counts within small apertures suggest excellent linearity down to quite low exposure levels of about 5 e- in the central pixel for stellar sources, i.e., down to a level where individual stars

cannot be recognized in single exposures. The data used for this test were taken in April, 2002, only one month after on-orbit installation of ACS. CTE effects are of course a source of low-intensity non-linearity (that correlates with y-positioned detector) and should have been very small at this early date (tests verified this). By contrast, for the approximately one year after launch observations used in ISR 2003-09 (quantifying CTE more generally), the CTE impact (at middle of CCDs) at the three lowest intensity levels in Figure 5.27 are at 24, 16, and 10.3% respectively. Thus, any intrinsic non-linearities that may exist at low intensity levels are small compared to the normal losses expected from CTE growth after ACS had been on orbit for only a few months. Within the mutual error bars results for the GAIN = 2 comparisons are entirely consistent -- the low-level linearity of the WFC channel is obviously excellent.

Figure 5.24: The upper panel shows simple aperture sums for the 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 second to mean 22.5 second exposures after normalization to the relative exposure time. To account for minor encirled energy differences for the very small 9 pixel (total, not radius) aper-tures used all points have been normalized by 1.006, the initial value for the brightest bin. The error bars show plus/minus one sigma 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.
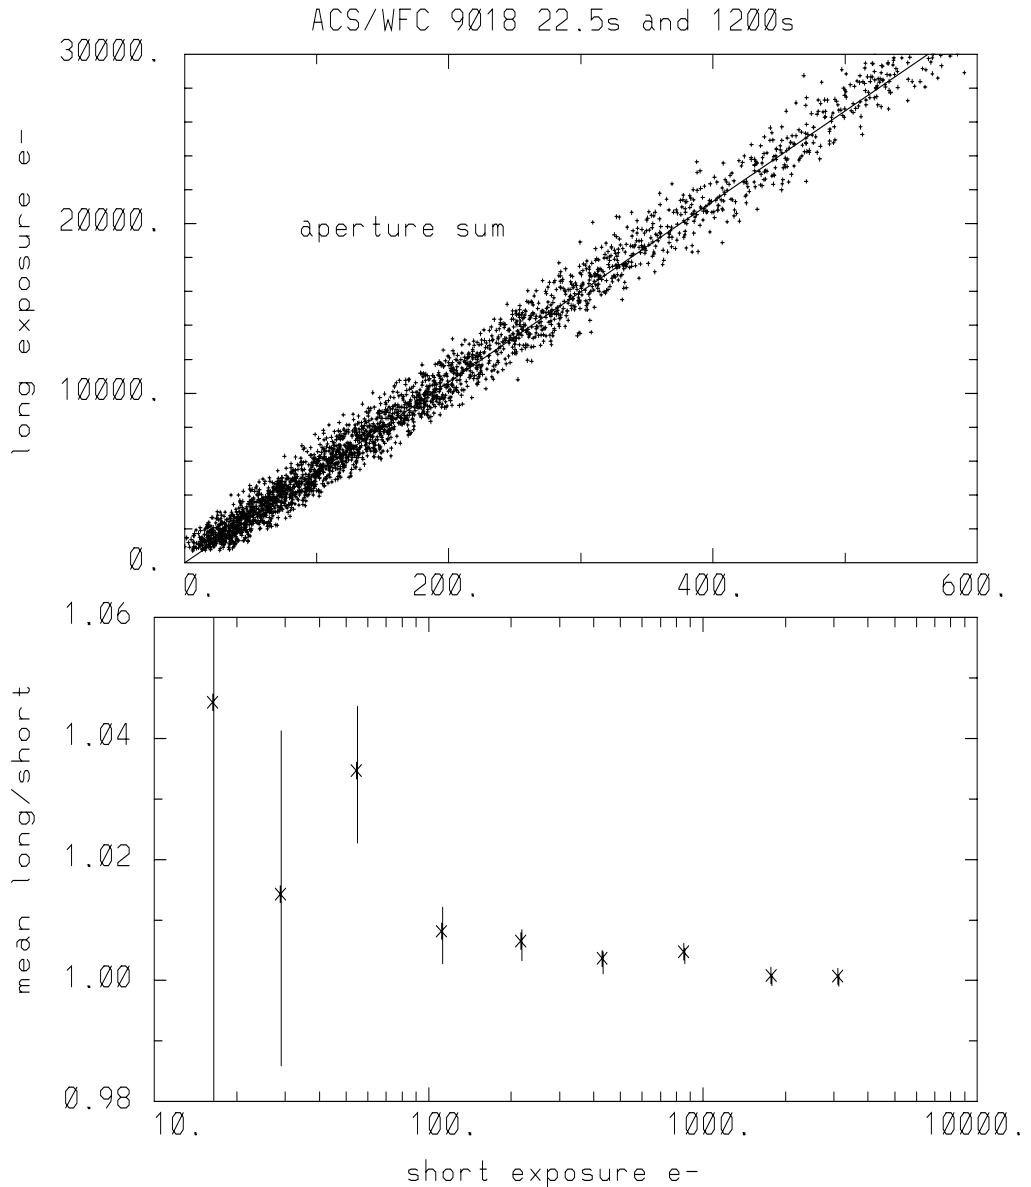
5.6.4 Linearity Beyond Saturation

The response of the ACS CCDs remains linear not only up to, but well beyond, the point of saturation when using a GAIN value that samples the full well depth. ISR 04-01 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 1009-1020) and the WFPC2 camera (Gilliland, R.L. 1994, ApJ, 435, L63-66). It is possible to easily perform photometry on point sources that remain isolated simply by summing over all of the pixels bled into with a GAIN value that samples the full well depth. Given the extra dynamic range afforded before saturation at GAIN=2 (WFC) and GAIN=4 (HRC), and the only modestly increased readout noise coupled with the potentially beneficial aspect of being able to recover photometry on saturated objects, more frequent use of these GAIN values may be appropriate.

Here 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. Only the cases of GAIN = 4 for the HRC and GAIN = 2 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 Camera

Figure 5.25 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 saturate 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 5.25 shows the same stars but now using identical extraction apertures in the two exposures, the vertical tick mark near 205,000 e- on the x-axis flags stars below which the central pixel remained unsaturated, above which the central pixel experienced saturation. In the aperture sum data there is no difference in the long exposure photometric quality between stars that saturate, even up to a factor of 5, and those that do not. Within the domain sampled here the accuracy of saturated star photometry is much better than 1%.

Figure 5.25: 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 on the x-axis separates stars that are saturated in the long expo- sure. 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.
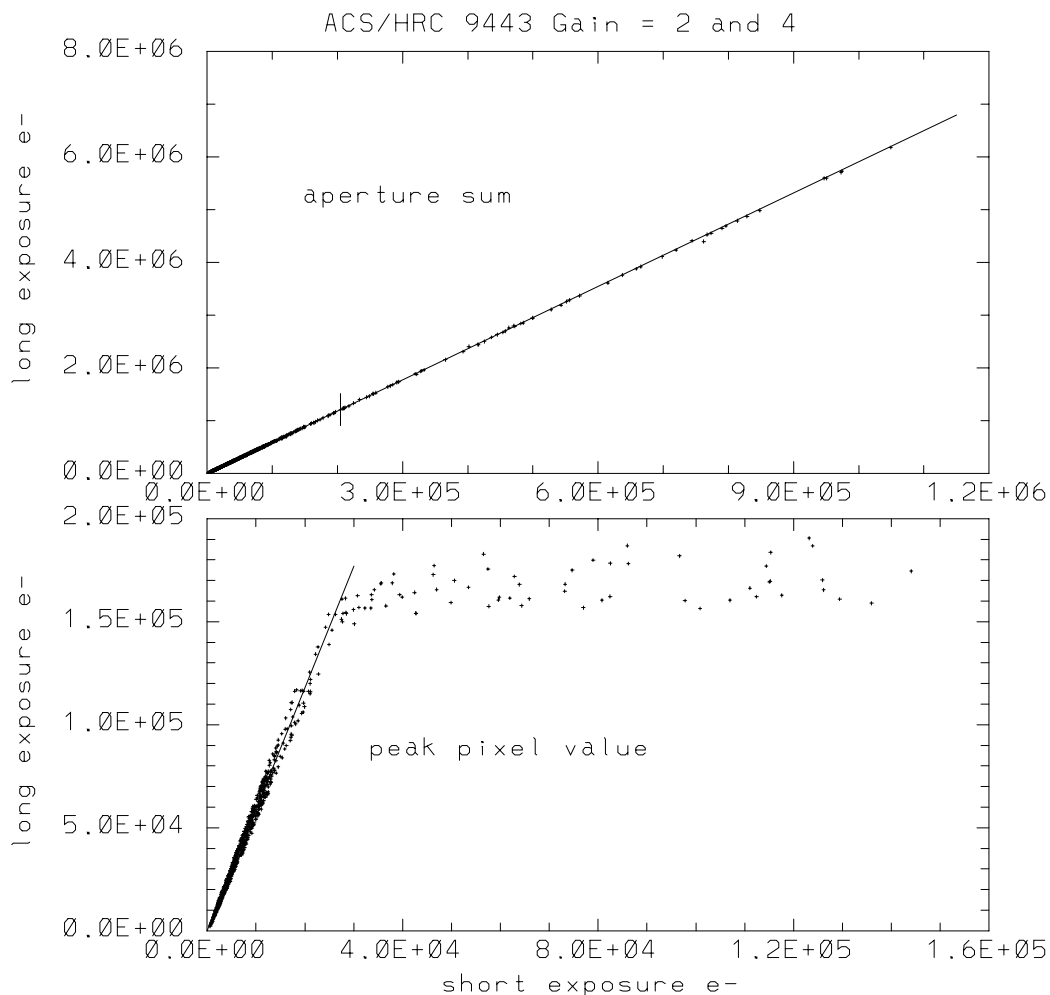
Wide Field Camera

Figure 5.26 shows results for the WFC. Over a range of nearly 4 magnitudes beyond saturation, photometry remains linear to < 1%.

All of the above linearity results are based upon comparisons of FLT images. The conservation of flux property of drizzle leads to equally good results for linearity beyond saturation comparing long and short DRZ images. An analysis of the drizzled data sets corresponding to Figure 5.26 show equally impressive results.

Figure 5.26: 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 12,350 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 plotted line in each panel represents perfect linearity, rather than fits to the data. 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=2.
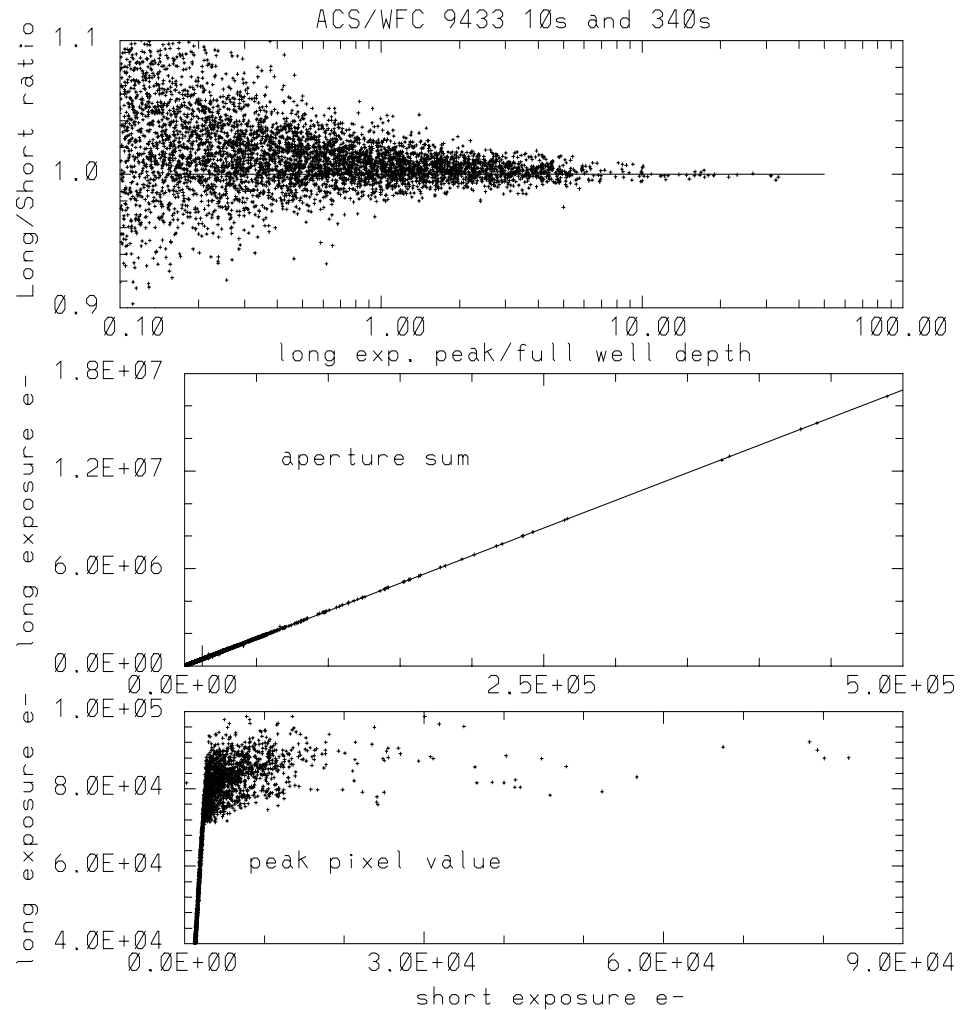
5.6.5 Shutter Stability

The shutter mechanism operation varies for different commanded exposure times. For example, the WFC shutter management varies for the shortest exposures, e.g., at 0.5 sec for WFC it is a continuous motion, 0.6

sec is not allowed and 0.7 sec begins the standard mode of operation. Similar differences exist for the shortest HRC options. Ground-based testing is discussed in Martel, Hartig and Sirianni (2001c) and should be referred to for background on operation of the shutters.

For short exposure times, field dependent timing (shutter shading), A versus B blade shutter 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 were suggestive of 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 03-03 (Gilliland and Hartig 2003).

Even at the shortest exposures, shutter shading measures 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**.

No systematic biases in timing reach the 0.5% level for any average over multiple pairs at any exposure time when switching between A and B shutter control. This supports not actively utilizing a record of which shutter was used.

Stability of shutter timing is a bit more problematic. For the HRC stability appears excellent, measured as rms across several exposures, the errors remain well under 0.5%. At exposure times of 1.0 sec peak-to-peak fluctuations on the HRC are at only the 0.1% level. Only at 0.1 sec do the HRC exposures exhibit peak-to-peak fluctuations up to ~1.0% (and part of this may be measurement error). The WFC shutter timing stability is 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 are well under 0.5%. However, out of 7 pairs of back-to-back 1.0 sec exposures, two of the pairs have individual components differing by > 1.3%. Where such differences exist on WFC, it is always the case that shutter timing under A control is shorter than the following B exposure. At 0.5 sec on WFC (where shutter operation is one continuous rotation) the greatest discrepancy A to B remains < 0.3%. These results suggest that if generally short WFC exposures are required (where short is taken to be < 2 sec), then 0.5 sec exposures seem stable and likely to support 1% accuracy, but exposures over the range 0.7 to <2.0 sec 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.

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 sec HRC exposure is actually larger than this by 4.1% WFC exposures at 0.5 sec, 0.7 sec and 0.8 sec in reality differ by +1.6%, -1.0% and -0.6%

respectively. For these 4 exposures, using revised times in **CALACS** is recommended in order to support accurate photometry; retrievals made after March 11, 2004 will have invoked use of these corrected values.

CHAPTER 6:

ACS Data Analysis

In this chapter. . .

This chapter describes the data reduction applications available to work with ACS data, and the specific analyses you may want to apply to your imaging, spectroscopy, polarimetry and coronagraphy data.

6.1 Photometry

6.1.1 Photometric Systems, Units, and Zeropoints

We strongly recommend that, whenever practical, ACS photometric results be referred to a system based on its own filters. Transformation to other photometric systems is possible (see Sirianni et al., 2004, PASP submitted) but such transformations have limited precision and strongly depend on the color range, surface gravity and metallicity of the stars considered.

Three magnitude systems, based on ACS filters are commonly used, VEGAMAG, STMAG and ABMAG. They are all based on the knowledge of the throughput of the entire system (OTA + ACS CAMERA + FILTER +

DETECTOR). These synthetic systems are tied to the observed standard stars. VEGAMAG is a standard-based system and therefore is subject to systematic errors tied to the calibration of the standard star. In the last decades it has become more and more common to use photometric systems such as ABMAG (Oke, J.B. 1964, ApJ 140,689) or STMAG (Koorneef, J. et. al. 1986, in Highlights of Astronomy (IAU), Vol.7, ed. J.-P. Swings, 833) which are directly related to physical units, and they are therefore much simpler and cleaner. The choice between standard-based and flux-based systems is mostly a matter of personal preference. Any new determination of the absolute efficiency of the instrument results in revised magnitudes for all three systems.

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 constructed by assembling empirical and synthetic spectra together (Bohlin & Gilliland, 2004, AJ, in press). The vega magnitude of a star with flux F:

$$VEGAMAG = -2.5 \log 10\left(\frac{F}{F_{vega}}\right)$$

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

STMAG and ABMAG

These two similar photometric systems are flux-based. The conversion is chosen so 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 Log F_{λ} - 21.10
- ABMAG = -2.5 Log F_{ν} - 48.60

where F_{ν} is expressed in erg cm^{-2}sec^{-1}Hz^{-1}, and F_{λ} in erg cm^{-2}sec^{-1}Å^{-1}. Another way to express these zeropoints is to say that an object with a constant $F_{\nu} = 3.63 \times 10^{-20}$ erg cm^{-2}sec^{-1}Hz^{-1} will have magnitude AB=0 in every filter, and an object with $F_{\lambda} = 3.63 \times 10^{-9}$ erg cm^{-2}sec^{-1}Å^{-1} will have magnitude ST=0 in every filter.

Zeropoints

By definition the zeropoint of an instrument is the magnitude of an object that produces one count per second. Each zeropoint refers to a count rate measured in a specific aperture. Since for point source photometry it is not practical to measure counts in a very large aperture, counts are instead measured in a small aperture and 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 arcsec. Independently from the different definitions, the magnitude in the passband P in any of the ACS systems is given by:

ACSmag(P)=-2.5 log(total electrons/sec) + zeropoint.

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

- Using SYNPHOT you can renormalize a spectrum to 1 count/sec in the appropriate ACS passband and output the zeropoint in the selected magnitude system (provided that the most updated through-put tables are included in your SYNPHOT distribution). This example renormalizes a 10,000 K blackbody:

```
iraf> calcphot obsmode "rn(bb(10000),band(obsmode),1,counts)" form
```

where `obsmode` is the passband, for example "acs,hrc,f555w", and `form` is the output magnitude system: ABMAG, VEGAMAG or STMAG.

- Using the header keywords. The simplest way to calculate the most updated STMAG and ABMAG zeropoints for your data is to use the photometric keywords in the SCI extension of ACS images:
  - PHOTFLAM is the inverse sensitivity (erg $cm^{-2}sec^{-1}Å^{-1}$); 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.

The header keywords PHOTFLAM and PHOTPLAM relate to the STMAG and ABMAG zeropoints through the formulae:

STMAG_ZPT = -2.5 Log (PHOTFLAM) - PHOTZPT = -2.5 Log (PHOTFLAM) - 21.10

ABMAG_ZPT = -2.5 Log (PHOTFLAM) - 21.10 - 5 Log (PHOTPLAM) + 18.6921

- Using a list of published zeropoints such as those published on the ACS web page:

  http://www.stsci.edu/hst/acs/analysis/zeropoints

This web page will be maintained with the most current zeropoints. More information about the zeropoint determinations can be found in Sirianni et al. (2004, PASP submitted) and ACS ISR 04-08.

6.1.2 Aperture and Color Corrections

In order to reduce errors due to the residual flat fielding error and background variations and to increase the signal-to-noise ratio, both of the most popular photometric techniques, 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. 2004 PASP submitted). As a consequence, this measurement must be tied to the total count rate by applying an aperture correction.

For historic consistency, ACS zeropoints refer to the total flux in a nominal "infinite" aperture and are ready to use for surface photometry. When performing point source photometry, users should convert their count rates within any given aperture to the count rates in the infinite aperture. In order to facilitate this conversion Sirianni et al. (2004, PASP submitted) provides the correction derived from the encircled energy profiles of an intermediate aperture (0.5" radius) to an infinite aperture. This correction is filter dependent and varies between ~0.09 and ~0.12 for WFC and between ~0.09 and ~0.31 for the HRC.

Users should determine the offset between their own photometry and aperture photometry with a 0.5" radius aperture. This is usually done by measuring a few bright stars in an uncrowded region of the field of view and applying the offset to all photometric measurements. If such stars are not available, encircled energies have been tabulated (Sirianni et al. 2004, PASP submitted). However users should be reminded that accurate aperture corrections are a function of time and location on the chip and also depend on the kernel used in Drizzle. They should avoid the blind application of tabulated encircled energies especially at small radii.

The aperture correction for the near-IR filters presents further complications. All 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 header, a soda glass substrate. It is then reflected back from the header's metallized rear surface and re-illuminates the CCDs front-side 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 broader at increasing wavelengths. WFC CCDs incorporate a special anti-halation aluminum layer between the frontside of the CCD and its glass substrate (Sirianni et al. 1998, proc SPIE vol.3555, 608, ed. S. D'Odorico). 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 wavelength > 9000 Å (Hartig et al. 2003, proc SPIE vol.4854, 532, ed. J.C. Blades, H.W. Siegmund), see Section 6.1.4 for more details.

The same mechanism responsible for the variation of the intensity and extension of the halo as a function of the 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 et al. (2002, proc SPIE vol.4854, 532, ed. J.C. Blades, H.W. Siegmund.) and Sirianni et al. (2004, PASP submitted) 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., 2004, PASP submitted, for more details).

The aperture correction for red objects should be determined using an isolated same-color star in the field of view, or using the effective wavelength versus aperture correction relation (Sirianni et al. 2004, PASP submitted). If the object's SED is available an estimate of the aperture correction is also possible with SYNPHOT. In fact, a new keyword "aper" has been implemented to call the encircled energy tables in the OBSMODE for ACS. A typical OBSMODE for an aperture of 0.5" would read like "acs,wfc,f850lp,aper$0.5". A comparison with the infinite aperture magnitude using the standard OBSMODE "acs,wfc,f850lp" would give the estimate of the aperture correction to apply. See the SYNPHOT handbook for more details.

Color Correction

In some cases it may be necessary to compare ACS photometric results with existing datasets in different photometric systems (e.g., WFPC2, SDSS or ground Johnson-Cousins). Since the ACS filters do not have exact counterparts in any other "standard" filter set, 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, serious systematic effects could be introduced. The transformations can be determined by using SYNPHOT, or by using the published transformation coefficients (Sirianni et al., 2004, PASP submitted). In any case users should not expect to preserve the 1-2% accuracy of ACS photometry on the transformed data.

6.1.3 Pixel Area Maps

When ACS images are flat fielded by the CALACS pipeline the resultant FLT files are flat if the original sky intensity was also flat. However, there is still very significant geometric distortion remaining in these images and the

pixel area on the sky will vary around the field. Consequently, the relative photometry in the FLT images must be corrected.

One option is to drizzle the data to remove the geometric distortion and allow the sky to remain flat. Therefore both surface and point-source relative photometry are correct in the resulting DRZ files and the photometric zeropoint in the header will allow conversion from electrons/s to absolute flux units.

Unfortunately users who wish to perform photometry directly on the distorted FLT files, rather than the drizzled (DRZ) data products, will require a field-dependent correction to match their photometry with the DRZ zeropoint.

The correction to the FLT images may be made by multiplying the measured flux on the FLT image by the pixel area at the corresponding position using a pixel area map (PAM), and then dividing by the exposure time (DRZ images have units of electrons/s). The pixel area must be expressed in units of the reference scale (50mas/pixel for WFC and 25mas/pix for HRC). Once this correction has been applied the same zeropoint as applied to the drizzled data products may be used to convert to absolute flux units.

$$DRZ\_flux = FLT\_flux * PAM / exposure\ time.$$

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.

Users may construct PAM images for themselves by downloading the script example files as well as the coefficients files which are needed from the web page:

http://www.stsci.edu/hst/acs/analysis/PAMS

and executing the scripts in IRAF or PyRAF. Alternatively PAM images may be downloaded directly from the same page, but they are very large.

Figure 6.1: Variation of the WFC and HRC effective pixel area with position in detector coordinates.

WFC PIXEL AREA MAP

Normalized to 0.05 arcsecond square pixel
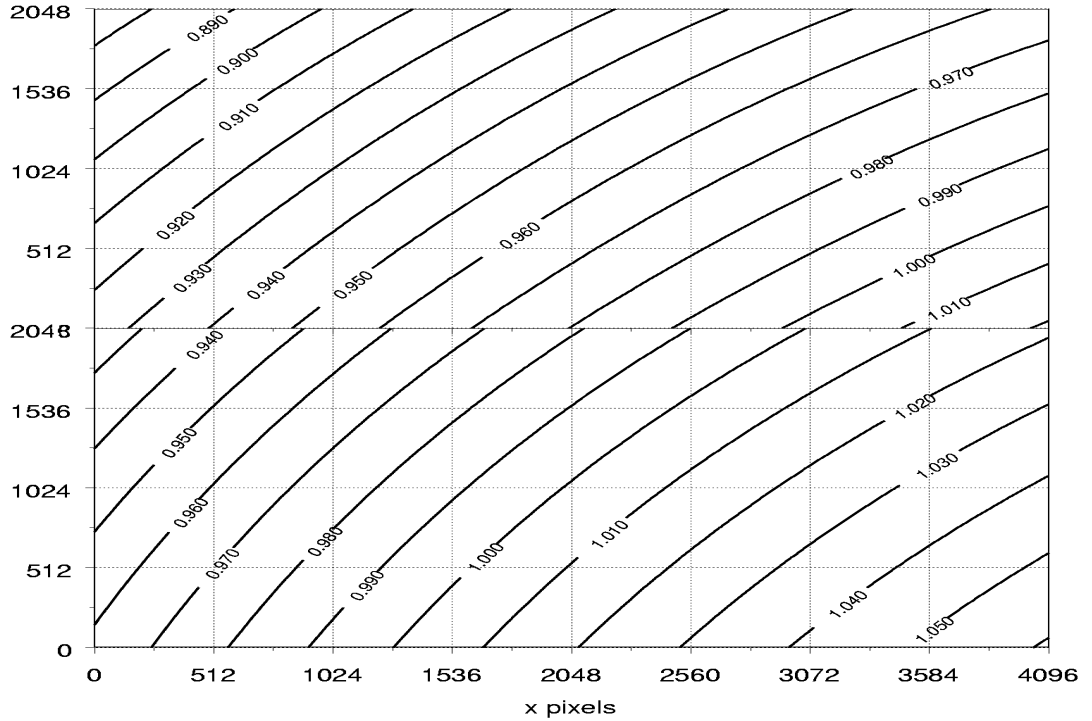
x pixels

HRC PIXEL AREA MAP

Normalized to 0.025 arcsecond square
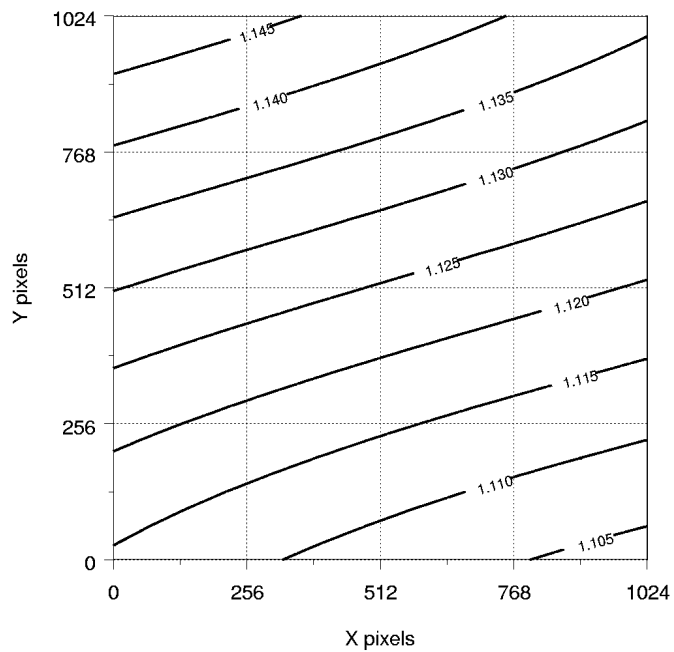
X pixels

PAM Concept Illustration

To illustrate the concepts of extended source and point source photometry on FLT and DRZ images we consider a simple idealized example of a 3x3 pixel section of the detector. We assume that the bias and dark corrections are zero and that the quantum efficiency is unity everywhere.

Example #1 Constant Surface Brightness Object

Let's suppose we are observing an extended object with a surface brightness of 2 e⁻/pixel in the undistorted case. With no geometric distortion the image is:

Actual scene on the sky:

2	2	2
2	2	2
2	2	2

In reality ACS suffers from geometric distortion and as a consequence pixels are not square and the pixel area varies across the detector:

Detector pixel layout on the sky:

Let's suppose the pixel area map (PAM) is:

PAM:

1.4	1.2	1.0
1.2	1.0	0.9
1.0	0.9	0.8

As a result in the raw data there is an apparent variation in surface brightness.

RAW image:

2.8	2.4	2.0
2.4	2.0	1.8
2.0	1.8	1.6

The geometrical area of each pixel is imprinted in the flat field as well as the photometric sensitivity. In this example, since we assumed that the quantum efficiency is unity everywhere, the flat field is just the equivalent of the PAM:

Flat field:

1.4	1.2	1.0
1.2	1.0	0.9
1.0	0.9	0.8

ACS flat fields are designed to level out a uniformly illuminated source and not to conserve total integrated counts, so after the flat field correction the FLT image has the correct surface brightness and can be used to perform surface photometry. However the image morphology is distorted.

FLT image:

2.0	2.0	2.0
2.0	2.0	2.0
2.0	2.0	2.0

PyDrizzle can be run on the FLT image. The result is that each pixel is free of geometric distortion and is photometrically accurate.

DRZ image:

2.0	2.0	2.0
2.0	2.0	2.0
2.0	2.0	2.0

Example #2 Integrated photometry of a point source

Now let's suppose we are observing a point source and that all the flux is included in the 3x3 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 the geometric distortion, the PSF as seen in the raw image is distorted. The total counts are conserved, but they are redistributed on the CCD.

RAW image:

1.9	10	1.9
10.1	50	10.1
2.5	11	2.5

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

FLT image:

1.36	8.33	1.9
8.42	50	11.22
2.5	12.22	3.13

In this example the counts now add up to 99.08, instead of 100.

In order to perform integrated photometry the pixel area variation need to be taken into account. this can be done by multiplying the FLT image by the PAM or by running **PyDrizzle**.

FLTxPAM				DRZ image			
1.9	10	1.9		2.0	10.5	2.0
10.1	50	10.1		10.5	50	10.5
2.5	11	2.5		2.0	10.5	2.0

Only running of **PyDrizzle** removes the geometric distortion, but both approaches correctly recover the count total as 100. Users should be cautioned that this is just an idealized example. In reality the PSF of the star extends to a much bigger radius. If the user decides to work on the flat fielded image after correcting by the pixel area map, they need to calculate a new aperture correction to the total flux of the star. The aperture corrections discussed in Section 6.1.2 are only for **PyDrizzle** output images. In most cases the aperture correction for distorted images will be quite different from the same star measured in the DRZ image. This is particularly true for small radius apertures.

6.1.4 PSF

PSF Field Dependence

The point spread functions in the ACS cameras are more stable over the field of view than in any other HST camera, especially when compared to WFPC2. The PSF does vary slightly in both shape and width across the field to a degree that may affect photometric measurements. This effect is described in detail in ACS ISR 03-06. The 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.

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 6.2). 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 wavelength, the blurring and variations in PSF width will increase towards shorter wavelengths. At 500 nm, photometric errors of 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 03-06).

The shape of the PSF 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, and can be found at:

http://www.stsci.edu/software/tinytim/

Figure 6.2: WFC chip thickness (left) and PSF FWHM (right).

WFC CCD Epitaxial Layer Thickness

12.6 μm   17.1 μm

WFC F550M PSF FWHM (2.12 to 2.49 pixels)
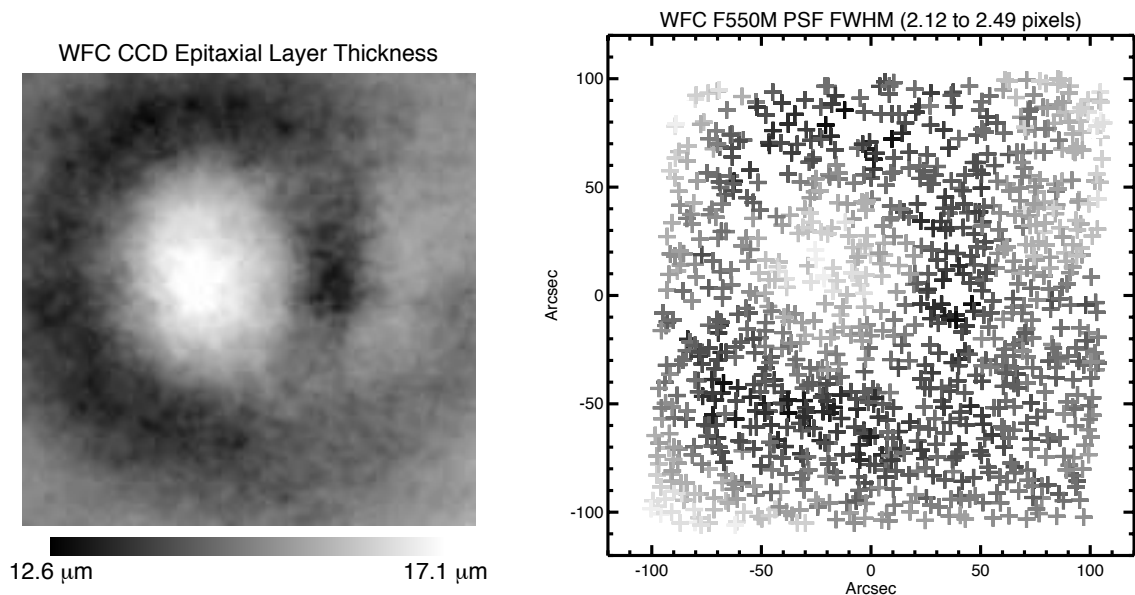
PSF Long-Wavelength Artifacts

Long wavelength (λ > 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. This 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 broadband 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 6.1.2. 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 (right spike in the HRC) while the other is seen above the right spike (below the left spike in the HRC). 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 6.3: CCD scatter at red wavelengths. WFC images of the standard star GD71 through filters F775W (9s, left), F850LP (24s, middle), and FR1016N at 996 nm (600s, 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.
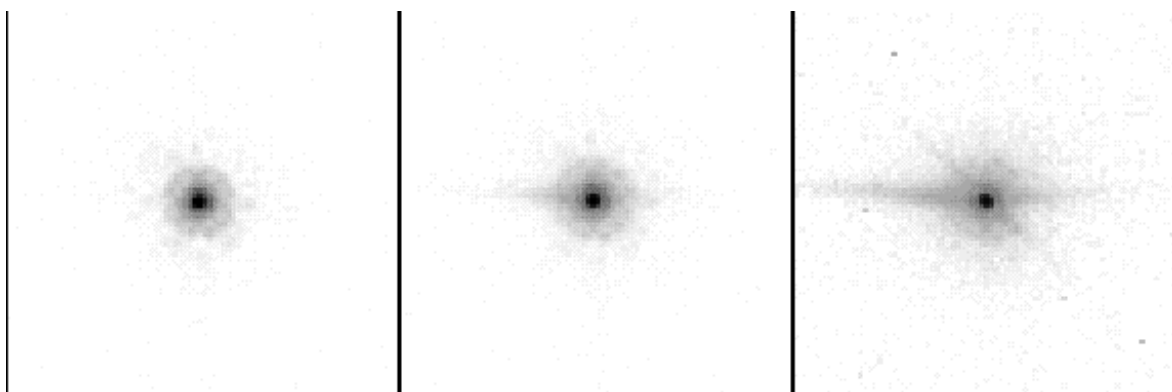
HRC and SBC UV PSFs

Below 3500Å the low and mid-spatial-frequency aberrations in HST result in highly asymmetric PSF cores surrounded by a considerable halo of scattered light extending 1-2 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.

6.1.5 CTE

To date, all CCDs flown in the harsh radiation environment of HST suffer degradation of their charge transfer efficiency (CTE). The effect of CTE degradation is to reduce the apparent brightness of sources, requiring the application of photometric corrections to restore measured integrated counts to their "true" value.

The principle of correcting for CTE losses to achieve unbiased measurements is straightforward, but in practice it can be quite difficult to utilize. The charge loss due to imperfect CTE depends on combinations of scene characteristics such as the total counts and physical extent of a source on the detector, its background rate (global and local), the number of parallel and serial transfers, and elapsed time in the life of the detector. It is difficult to quantify and calibrate the degree to which each of these factors affects the measurements of a source over time. The general solution has been to use repeated observations of cluster fields to quantify the impact of CTE on point sources, to use internal measures of CTE degradation to track changes in between costly, external calibrations and, in the future, to do additional studies of the impact to more complicated sources (e.g., extended sources) when the CTE losses become large.

ACS has two CCDs, WFC and HRC, which clock charge and are therefore subject to CTE losses. Here we describe the external calibration of their photometric losses due to imperfect CTE. All CTE calibrations were derived from matched sets of observations of 47 Tucanae using a wide range of position, filter, and exposure time. All CTE measurements were performed by comparing the brightness of individual stars as a function of pixel transfers. The present discussion summarizes the results presented in ACS ISR 03-09 and ACS ISR 04-06.

Figure 6.4 gives a clear and consistent snapshot of the impact of imperfect CTE on photometry in early 2003 for WFC. The statistics are sufficient to detect non-trivial CTE losses from individual bins (~a few hundred stars) at the ~5-10 sigma confidence level for faint stars.

Figure 6.4: An example of the linear relationship determined for WFC for one filter and exposure time (F606W for 30 sec providing a sky level of <s>=3.6 electrons) at one brightness range (stars with an average of <f>~400 electrons and a full range of +/-1 mag) contained within a r=3 pixel aperture. The upper panel shows a parallel transfer relation, the lower panel illustrates the serial transfer dependence.

F606W,30.00s, <f>=402.616(e), <s>=3.5602(e), ap=3

Parallel

% loss=4.076
% error=0.578

F606W,30.00s, <f>=391.664(e), <s>=3.7983(e), ap=3

Serial

% loss=0.328
% error=0.590
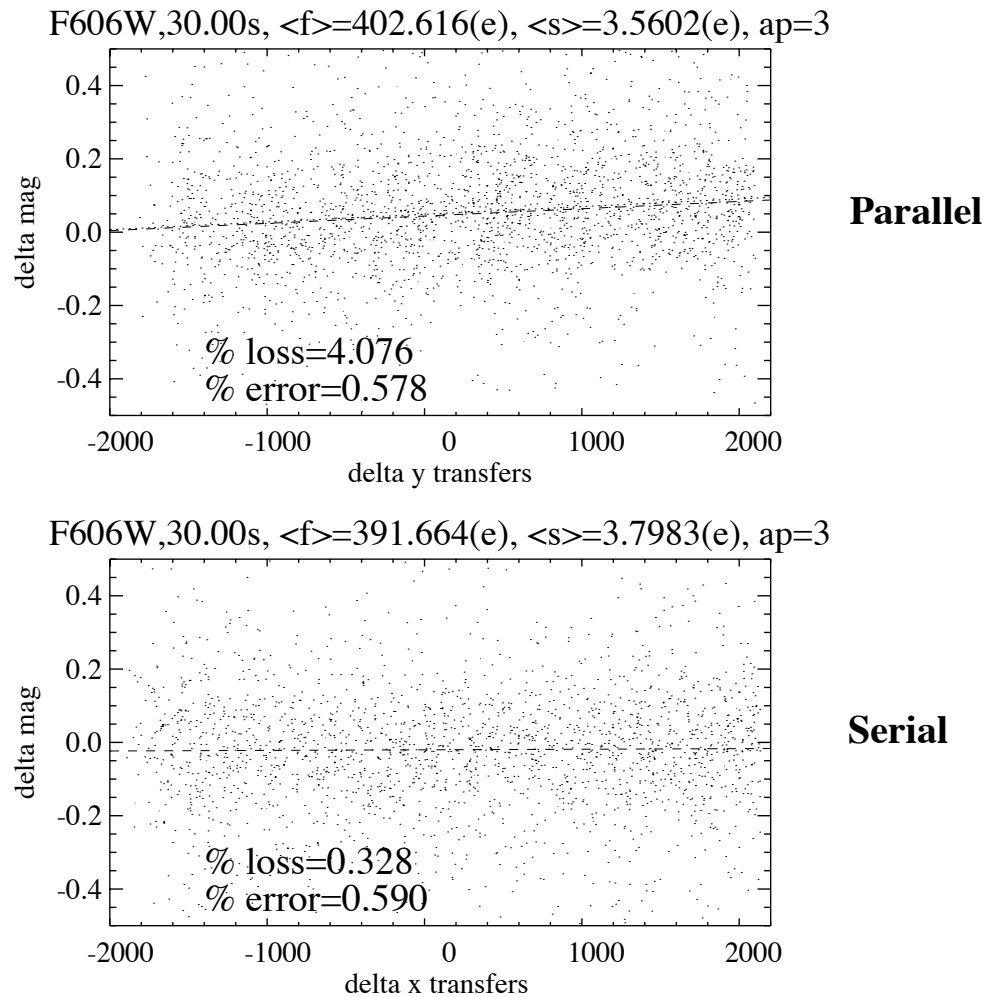
Figure 6.5: The dependence of parallel transfer CTE loss for WFC (over 2048 transfers) versus stellar flux contained within an r=3 pixel aperture (all images at low sky level and taken in March 2003). As shown, the relation is strong and suggests a power-law relation which was utilized in the correction formulae. The panel on the right shows the same in log-log space.

Low Background (2-4 e) Parallel CTE
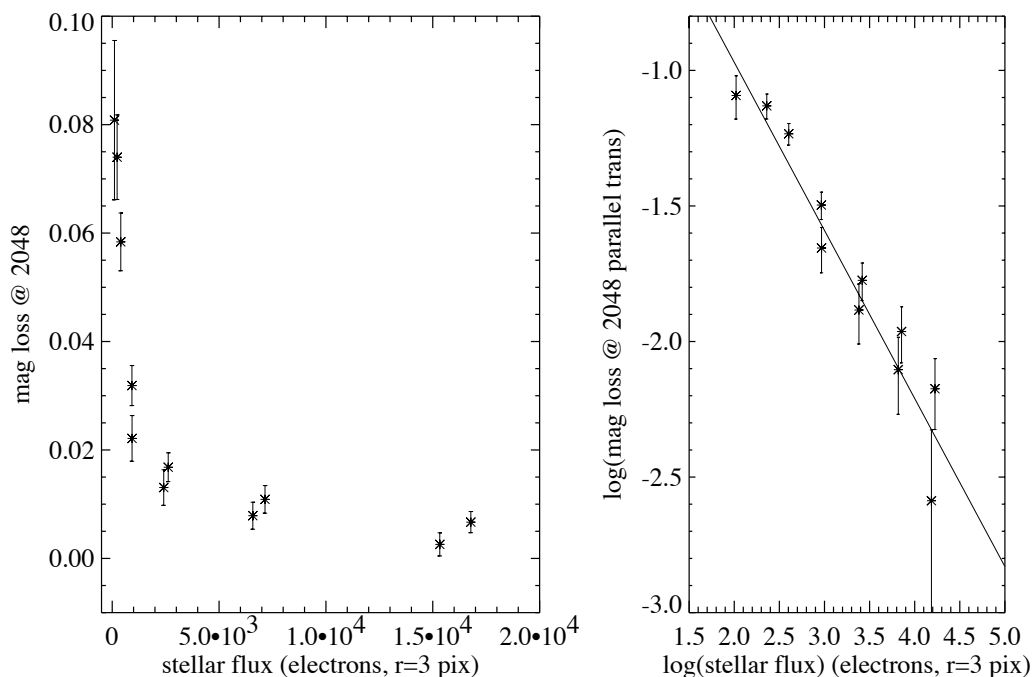
WFC Parallel and Serial CTE

As seen in Figure 6.5, the parallel WFC CTE loss appears to be a strong function of the stellar flux. This is not surprising given similar experience with WFPC2 and STIS. Charge traps may be present at varying depths of the silicon substrate of WFC. Larger charge packets (from brighter stars) may be subject to greater absolute flux losses by accessing deeper traps, but smaller charge packets (from fainter stars) appear to lose a larger fraction of their flux. As seen in Figure 6.5, the data suggests a power-law relationship between CTE loss and stellar flux which we have utilized in the following correction formula.

In Figure 6.6 we show the relationship between CTE loss and background. Increased background can mitigate the CTE losses (presumably by filling traps in advance of the arrival of the stellar charge packet during read-out), and this phenomenon has been seen for WFPC2 and STIS. For the dependence on background we also use a power-law in our correction formula.

Figure 6.6: The relationship between parallel transfer CTE loss and sky background for WFC in March 2003. As seen, the correlation is weak (and is insignificant for apertures of r=5 and 7) and in the sense of reduced loss with higher background. The measurement was made for stars with a flux of ~1000 electrons to remove the dependence on stellar flux.

Parallel CTE Loss @ Stellar Flux ~1000 e, r=3
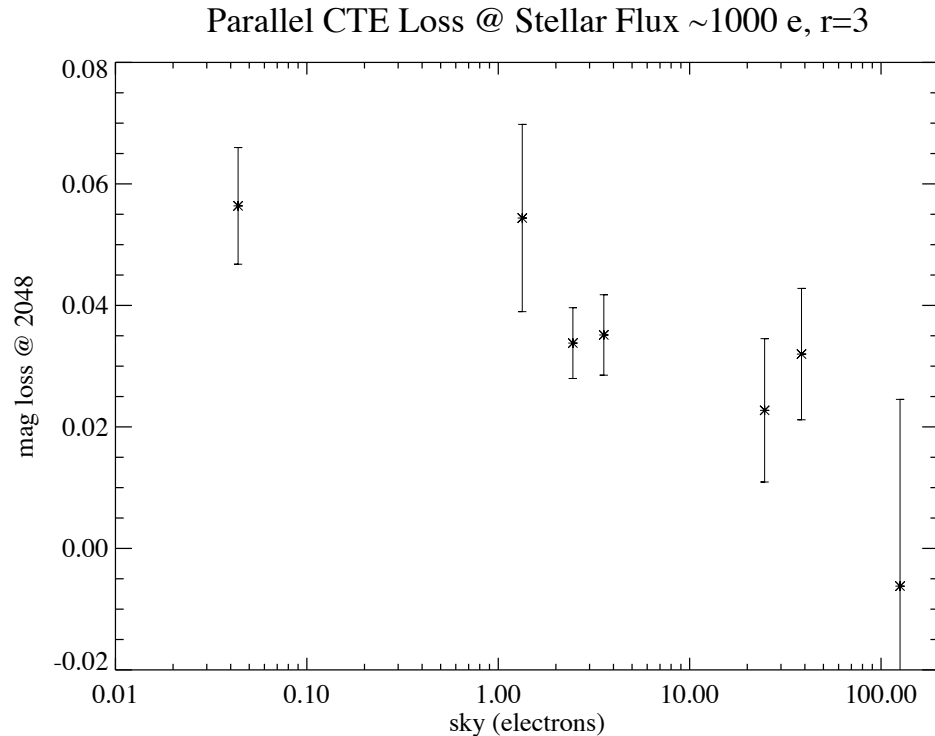
Our CTE parameterization is:

$$YCTE = 10^A \times SKY^B \times FLUX^C \times \frac{Y}{2048} \times \frac{(MJD - 52333)}{365}$$

The term, (Y/2048) reflects the linear relationship of the CTE loss with pixel transfer as seen in Figure 6.4. The last term involves the modified Julian date (MJD) from the header and reflects a linear degradation with time. The full range of sky values (SKY) sampled extended from ~0.1 e⁻ (F502N for 30 sec) to 125 e⁻ (F606W for 1100 sec). The full range of aperture flux values (FLUX) sampled ranged from ~100 e⁻ to 300,000 e⁻.

Table 6.1 contains the best fit values of the parameters A, B, and C and their uncertainties for the equation given above, for 3 different aperture sizes. The parameters were determined by simultaneously fitting data obtained in March 2003, August 2003 and February 2004. The fit to the data is shown in Figure 6.7. To determine the CTE correction (YCTE) expressed in magnitudes for a particular image observers should:

• select the appropriate aperture size,

- measure the flux in the aperture and the sky value (if working with drizzled data in e⁻/sec one must multiply by the exposure time),

- determine the number of parallel transfers, Y (Y=y coordinate for 1<y<2048, Y=4096-y for y>2048),

- determine the MJD,

- and evaluate the correction.

The correction must be derived for the individual exposures, not for a "stack". Figure 6.7 shows the corrections derived from the formula for different observing parameters.

Table 6.1: CTE Correction Coefficients and Uncertainties for WFC Parallel Transfer.

aperture radius (pixels)	A(σ)	B(σ)	C(σ)
3	0.14(0.14)	-0.31(0.02)	-0.64(0.05)
5	-0.55(0.15)	-0.32(0.03)	-0.40(0.04)
7	-0.70(0.24)	-0.34(0.04)	-0.36(0.07)

In principle, first fitting the CTE trends in bins of stellar flux (and individual sky levels), followed by the step of constraining the multi-parameter fitting formula could result in information loss. However, the combination of contaminants (cosmic rays) and the large range of photon statistics (requiring different sigmas for sigma-clipping) dictated the adopted approach.

As of 2004 we have seen no evidence of photometric losses due to imperfect serial CTE on the WFC.

Figure 6.7: Parallel CTE trends. The rows and columns show the photometric losses (at y=2048) for different sky levels and dates, respectively for aperture radii of 3 pixels. The fitted line uses the global, time-dependent correction formula in "WFC Parallel and Serial CTE" with the appropriate stellar flux and sky levels.
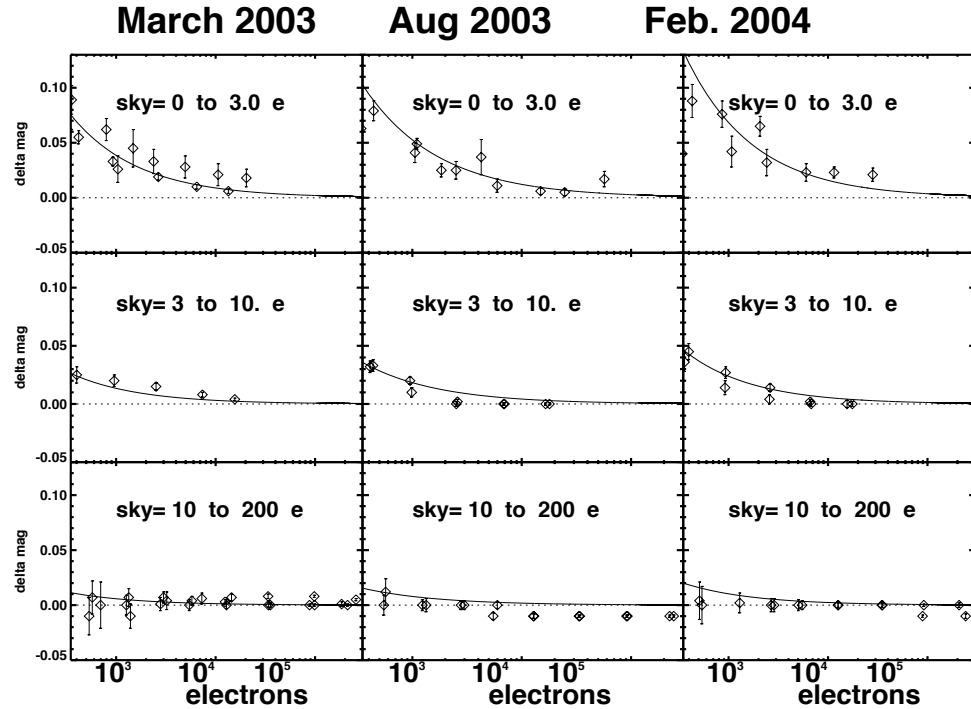
HRC Parallel and Serial CTE

The same procedure described above for WFC was utilized to analyze HRC data and to quantify the photometric losses due to imperfect CTE. By utilizing all 4 read-out amplifiers, we obtained 2 independent measures of the CTE loss trend for all combinations of filter and exposure time and stellar flux.

Overall, the quantification of photometric losses due to imperfect CTE on HRC is of much lower accuracy than for WFC due to the small field of view of HRC (45 times less area) and the resulting paucity of stars. The adopted parallel CTE correction parameterization in magnitudes is:

$$YCTE = 10^A \times SKY^B \times FLUX^C \times \frac{Y}{1024} \times \frac{(MJD - 52333)}{(52714 - 52333)}$$

The terms in the equation are the same as for the WFC equation except for the modified Julian date (MJD) term, which characterizes the time evolution of the degradation of CTE away from the calibration date (MJD=52714). Table 6.2 contains the best fit values of the parameters A, B and C and their uncertainties for the YCTE equation for HRC. At present, our calibration of CTE on HRC is based on a single epoch (March 2003).

The method to correct HRC data is the same as for WFC. Observers should:

- select the appropriate aperture size,

- measure the flux in the aperture and the sky value (if working with drizzled data in electrons/sec one must multiply by the exposure time),

- determine the number of parallel transfers, Y (Y=y coordinate for read-outs with amps C or D, Y=1024-y for read-outs with amps A or B),

- determine the MJD,

- and solve for the correction.

For most observations the default amplifier, C, is used but users can determine which amplifier was used by consulting the header keyword CCDAMP. Similar to the case for WFC, at present we see no significant photometric loss in HRC due to serial transfers.

Table 6.2: CTE Correction Coefficients and Uncertainties for HRC Parallel Transfer.

aperture (pixels)	A(σ)	B(σ)	C(σ)
3	-0.89(0.26)	-0.24(0.13)	-0.21(0.07)
5	-0.87(0.36)	-0.27(0.17)	-0.21(0.09)
7	-2.22(1.10)	-0.58(0.46)	0.12(0.28)

Similar to the case for WFC, at present we see no significant photometric loss in HRC due to serial transfers.

CTE Trending and Evolution

The YCTE correction formulae include a linear time-dependence term which is bounded by the condition that the photometric losses due to imperfect CTE were zero at launch (MJD=52333). While future data will help refine the time-dependence, we believe the linear time-dependence is well motivated by other, on-orbit data such as internal CTE measurements (see "Internal CTE") and measurements of cosmic ray charge deferred tails.

The Relationship between field-dependent Charge Diffusion and CTE

The CCDs in HRC and WFC are both thinned and the thinning process inevitably results in large-scale variations in the CCD thickness. Because the charge diffusion of the CCD depends on thickness (thicker regions suffer greater diffusion), the breadth of the PSF and aperture corrections can also exhibit a field-dependence (see Section 6.1.4.). ACS

ISR 03-06 has characterized 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 very small (less than 1%), but becomes significant for small apertures (r<3 pixels).

It may be of interest to decouple the impact of imperfect CTE and charge diffusion on the field dependence of photometry but it is not necessary for observers to consider these two effects separately. The correction formulae given in Section 6.1.5 include both effects and hence corrects photometry for both effects. This is true as long as the user seeks to correct photometry to a perfect CTE and fiducial PSF width standard. An aperture correction from the fiducial aperture to a nominal "infinite" aperture is still necessary (see Section 6.1.2).

Internal CTE

Internal CTE measurements have been obtained since before launch for both the WFC and HRC. Extended Pixel Edge Response (EPER) and First Pixel Response (FPR) data is routinely collected to monitor the relative degradation of the CTE over time. These data indicate a linear degradation but are otherwise not directly useful for the scientific calibration of ACS data. These results will be published in an ISR that is presently in preparation and will be posted on the web site:

http://www.stsci.edu/hst/acs/documents/isrs/.

6.1.6 Red Leak

HRC

When designing a UV filter, a high suppression of off-band transmission, particularly in the red, has to be traded with overall in-band transmission. The very high blue quantum efficiency of the HRC compared to WFPC2 makes 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 with SYNPHOT, is given in Table 6.3 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. The same ratio is also listed for the equivalent filters in WFPC2. Clearly, red leaks are not a problem for F330W, F435W, and F475W and are more 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 6.3: In-band Flux as a Percentage of the Total Flux.

Stellar Type	WFPC2 F218W	HRC F220W	WFPC2 F255W	HRC F250W	WFPC2 F300W	HRC F330W	WFPC2 F439W	HRC F435W	WFPC2 F450W	HRC F475W
O5V	99.8	99.8	99.6	99.7	99.9	99.9	99.9	99.9	99.9	99.9
B1V	99.7	99.7	99.6	99.7	99.9	99.9	99.9	99.9	99.9	99.9
A1V	99.4	99.1	99.2	99.3	99.2	99.9	99.9	99.9	99.9	99.9
F0V	98.5	97.8	98.8	99.0	98.8	99.9	99.9	99.9	99.9	99.9
G2V	92.5	90.2	97.4	98.4	97.4	99.9	99.9	99.9	99.8	99.9
K0V	71.7	69.6	95.0	97.3	95.0	99.9	99.9	99.9	99.8	99.9
M2V	0.03	2.5	45.5	71.9	45.4	99.9	99.9	99.9	99.6	99.9

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 6.4.

Table 6.4: Visible-Light Rejection of the SBC F115LP Imaging Mode.

Stellar Type	Percentage of all Detected Photons which have $\lambda < 1800\text{Å}$	Percentage of all Detected Photons which have $\lambda < 3000\text{Å}$
O5V	99.5	100
B1V	99.4	100
A0V	98.1	100
G0V	72.7	99.8
K0V	35.1	94.4

6.1.7 UV Contamination

In an ongoing calibration effort we are regularly observing the star cluster NGC 6881 to monitor the UV performance of the HRC and SBC detectors in ACS. Our results indicate that ACS does not suffer from UV contamination (see ACS ISR 04-05 for details). Over time both the HRC and the SBC have a constant sensitivity within 1% (except the SBC filter F165LP which is within 2%). The three HRC filters F220W, F250W and F330W, do not change sensitivity by more than 1%. Presently, for the SBC the only conclusion is that any change in sensitivity is less than 2% because the precision of the current analysis is limited by the lack of an accurate L-flat field for the SBC.

6.2 Astrometry

6.2.1 Coordinate Transformations

There are three coordinate systems applicable to ACS images. First, there is the position of a pixel on the geometrically distorted raw image (RAW) or, identically, the position on the flat fielded images (FLT) after pipeline processing through **CALACS**. Second, there is the pixel position on the drizzled images (DRZ) created by **PyDrizzle** which corresponds to an undistorted pixel position on a tangent plane projection of the sky. Third, there is the corresponding position (RA, Dec) on the sky.

A position measured on a drizzled output image (DRZ) from **PyDrizzle** or MultiDrizzle may be transformed to a position on the celestial sphere using the utility **xy2rd** in the **stsdas.toolbox.imgtools** package within STSDAS. There is also a corresponding **rd2xy** task to go from the sky to the pixel position on a drizzled frame. These tools cannot be used for RAW/FLT files as they do not include the very large effects of geometric distortion.

The transformation between the pixel position on a distorted (RAW or FLT) and drizzled (DRZ) image may be performed using utilities in the dither package in STSDAS. The tasks **traxy** and **tranback** implement the same geometric mapping as **drizzle** and **blot** except applied to X,Y pixel positions rather than images. For more details on how to use these the IRAF help files should be consulted. As a more convenient high-level wrapper the **tran** task, which will be released as part of the late-2004 version of STSDAS, allows the mapping of X,Y positions between FLT and DRZ images. This task uses information in the image headers to define the necessary geometrical transformation and requires that the coefficients files are also available. These files are not provided by the STSDAS archive and must be recreated by running **PyDrizzle** on the FLT file on the local machine. Please note that to run **PyDrizzle** you also need to have access to the IDCTAB.

For example, if an object was found to be at (123,234) on an FLT image (test_flt.fits[sci,1]) the position on the drizzled DRZ product (test_drz.fits) can be found as follows:

```
pyraf> tran test_flt.fits[sci,1] test_drz.fits[1] forward 123 234
Running Tran Version  0.11 (May 2004)
-Reading drizzle keywords from header...
-Image  test_flt.fits[sci,1]  was # 1
-to be drizzled onto output  test_drz.fits[1]
-Transforming position...
Xin,Yin: 123.00000 234.00000  Xout,Yout: 128.20980  239.63962
PAM:      0.96218
```

The last number refers to the pixel-area map value at this point (see Section 6.1.3).

The reverse operation, DRZ to FLT, can be applied as follows:

```
pyraf> tran test_flt.fits[sci,1] test_drz.fits[1] backward
       128.20980 239.63962
Running Tran Version  0.11 (May 2004)
-Reading drizzle keywords from header...
-Image  test_flt.fits[sci,1]  was # 1
-to be drizzled onto output  test_drz.fits[1]
-Transforming position...
Xin,Yin:    123.00000   234.00000 Xout,Yout:    128.20980
239.63962
```

Finally the **xytosky** task, which forms part of the current dither package, will convert a pixel position on a distorted FLT file directly to a sky position, by applying the distortion correction from the IDCTAB and using the world coordinate information from the header. It may be used as follows:

```
pyraf> xytosky test_flt.fits[sci,1] x=123 y=234 linear=no
- IDCTAB: Distortion model from row 22 for chip 2 : CLEAR and F814W
RA (deg.) =  5.53751673945 , Dec (deg.)=  -72.0652050903
```

This task doesn't require the coefficient files but, like **PyDrizzle**, requires a copy of the IDCTAB to be available. Both **tran** and **xytosky** have options for lists of positions to be supplied as text files to allow multiple positions to be efficiently transformed.

6.2.2 Absolute and Relative Astrometry

The astrometric information in the header of an ACS image comes indirectly from the positions of the guide stars used. 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 1 arcsecond. 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. This is discussed in detail in Section 4.1. Typical accuracy of the distortion correction in the pipeline with fourth order polynomial solutions is 0.1 pixels (5mas for the WFC). The application of the work of Anderson and

King will allow even better distortion correction by the inclusion of additional small corrections held in image files (e.g., Figure 4.6) which can model effects not included in the polynomial solutions. We plan to include these corrections in the pipeline in late 2004.

6.2.3 Impact of Guide Star Failure

The guiding performance and pointing stability of HST are described in the HST Primer. The normal guiding mode uses two guide stars that are tracked by two of HST's Fine Guidance Sensors (FGSs). However, sometimes two suitable guide stars are not available and single-star guiding is used instead 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 review Appendix C of the Introduction to the HST Data Handbook.

The gyros have a typical drift rate of 1.0 mas/sec, although on rare occasions it can be as large as 1.5 mas/sec. This causes a rotation of the target around the single guide star, which in turn introduces a small translational drift of the target on the detector. 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. For ACS with single-star guiding, the typical and maximum drift rate of the target on the detector are in Table 6.5.

Table 6.5: Drift Rates for single-star guiding with ACS

TYPICAL	per 1000 sec exposure	per orbit(96 min)
WFC	0.0041 arcsec (0.08pix)	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 these numbers; the typical visibility period in an orbit (outside the Continuous Viewing Zone [CVZ]) is in the range 52-60 minutes, depending on target declination (see Section 6.3 of the Cycle 12 HST Primer). The drifts inherent to single-star guiding have two 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 func-

tion (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. This is not a problem if, during data reduction, different exposures are aligned before co-addition or comparison. The **MultiDrizzle** software provides a way of doing this. Also, there is always a slow drift of the telescope up to 0.01 arcsec/orbit due to thermal effects, even when two guide stars are used. So even when two guide stars are used it is often necessary to align exposures before co-addition.

In summary, for most scientific applications single-star guiding will not degrade the usefulness of ACS data. However, we do not recommend single-star guiding 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 with 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.

6.3 Polarimetry

6.3.1 Absolute Calibration

ACS contains a set of six filters which are sensitive to linear polarization; there are three visible polarizer filters with their polarization directions set at nominal 60 degree angles to each other, and three UV polarizer filters arrayed 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.

6.3.2 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 an 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 97-11) and Biretta, et al., 2004 (ACS ISR in preparation).

The instrumental polarization, defined as the instrument's response to an unpolarized target, provides a simple measure of some of these effects. Figure 6.8 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 degrees. The details of the CCD are proprietary so we have simply modeled it as Silicon at an incidence angle of 31 degrees; no doubt this is a serious over-simplification. Figure 6.9 shows the same plot for the WFC, which has an instrumental polarization around 2%. Here the IM3 mirror is a proprietary Denton enhance Silver Coating with an incidence angle of 49 degrees, and the CCD has an incidence angle of 20 degrees. While the lower instrumental polarization of the WFC seems attractive, we caution 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 6.8: Instrumental Polarization for the HRC
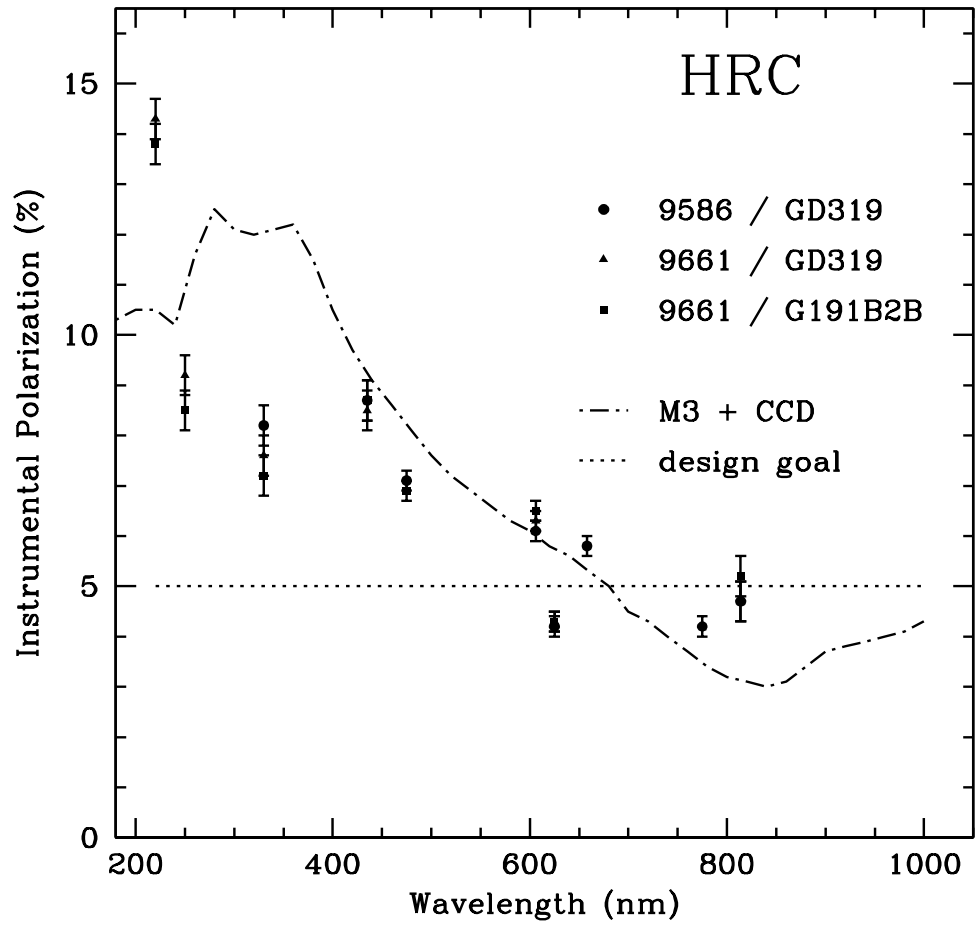
Figure 6.9: instrumental Polarization for the WFC
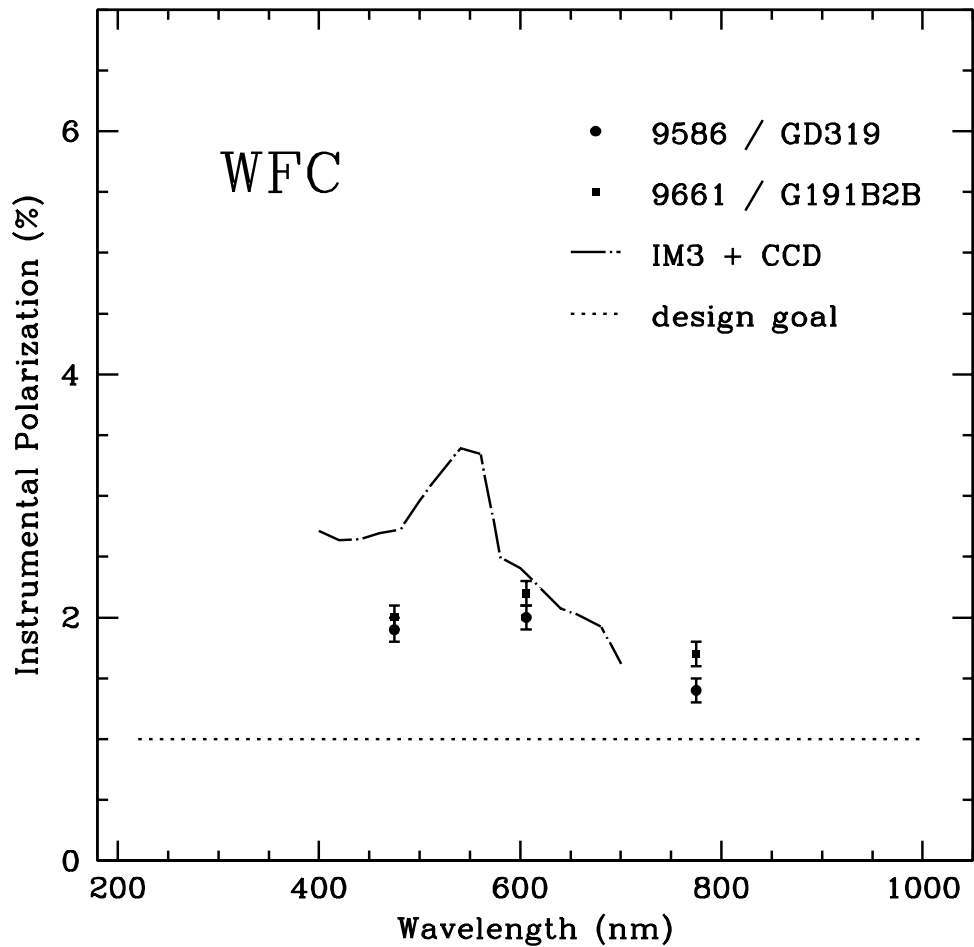
The ACS polarizer filters were characterized prior to installation in ACS by Leviton and the results as summarized in Figure 5.1 of the ACS Instrument Handbook. The cross-polarized transmissions are essentially zero for the POLV set, but are significant in the UV (5% - 10%) and far-red (20%) for the 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, although we hope to have one soon. There is also the possibility of polametric field dependences; while there has been study of intensity flats for the polarizers, the polarization field dependencies are not known.

6.3.3 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, we compared in-flight observations of the bright earth using the F435W+POLUV filters and 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. Improved polarizer L-flats are currently being derived from stellar observations, and these should bring the flat field uncertainties down to <1%.

6.3.4 Polarization Calibration

An extensive series of on-orbit polarization calibration observations were carried out at the end of Cycle 11 (programs 9586 and 9661). These included observations of unpolarized and polarized standard stars, the star cluster 47 Tuc, and an extended reflection nebula. Additional observations of polarized standards are needed over a wide and well-sampled range of HST roll angles; this will help quantify the angular dependences which are expected as the wavefront interacts with the diattenuation and phase retardation in the mirrors and CCD. Some of these observations are included in program 10055, which is underway (May 2004).

We now present a highly-preliminary calibration for use by polarization observers which is based largely on data in programs 9586 and 9661. (Observers are urged to check the ACS web pages for updates and improvements, since the material below will almost certainly be obsolete before this appears in print.) The strategy here is to create a calibration based upon the small amount of data which are available as of this writing. This calibration can be applied to either aperture photometry results, or to the images themselves (i.e., for an extended target). We begin by calibrating the polarization "zeropoint" using corrections which are derived from observations of unpolarized standard stars. Corrections C(CCD, POLnXX, spectral filter, n) are 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 6.6, and have been scaled such that Stokes I will approximate the count rate seen with no polarizing filter.

$$r(n) = C(CCD, POLnXX, \text{spectral filter}, n)\, r_{obs}(n)$$

Next, an "instrumental" Stokes vector is computed for the target. We note our definition of U is slightly different from that in the ACS Instrument Handbook:

$$I = \left(\frac{2}{3}\right)[r(0) + r(60) + r(120)]$$

$$Q = \left(\frac{2}{3}\right)[2r(0) - r(60) - r(120)]$$

$$U = \left(\frac{2}{\sqrt{3}}\right)[r(60) - r(120)]$$

Next we compute the fractional polarization of the target. We include a factor which corrects for cross-polarization leakage in the polarizing filters. This correction is useful for the POLUV filters, and values of T_{par} and T_{perp} can be found in Figure 5.4 of the ACS Instrument Handbook:

$$P = \frac{\sqrt{Q^2 + U^2}}{I} \times \left[\frac{T_{par} + T_{perp}}{T_{par} - T_{perp}}\right]$$

Finally we compute the position angle on the sky of the polarization E-vector. 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 we have χ=-69.4 degrees, and for the WFC we have χ=-38.2 degrees. Note that the arc tangent function must be properly defined; here we use a definition where the result is positive in quadrants I and II, and negative in III and IV.

$$PA = \left(\frac{1}{2}\right)Tan^{-1}\left(\frac{U}{Q}\right) + (PAV3) + \chi$$

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 degrees.

We have modeled the full instrumental effects and the above calibration together in an effort to determine the impacts of the remaining uncalibrated systematic errors. We believe 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 degrees. (This is in addition to uncertainties which arise from photon statistics in the observer's data.) We have checked this calibration 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. We have not given a calibration for F220W, F250W, or F814W, as we believe

these are 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, we have seen one incidence of a 5 degree PA error for F775W, suggesting this waveband is not calibrated as well as the others.

Table 6.6: Correction factors C(CCD, POLnXX, specfilt, n) for polarization zeropoint.

CCD	POLnXX	Spectral Filter	n=0	n=60	n=120
HRC	POLnUV	F330W	1.7302	1.5302	1.6451
		F435W	1.6378	1.4113	1.4762
	POLnV	F475W	1.5651	1.4326	1.3943
		F606W	1.4324	1.3067	1.2902
		F625W[1]	1.0443	0.9788	0.9797
		F658N[1]	1.0614	0.9708	0.9730
		F775W	1.0867	1.0106	1.0442
WFC	POLnV	F475W	1.4303	1.4717	1.4269
		F606W	1.3314	1.3607	1.3094
		F775W	0.9965	1.0255	1.0071

1. Not scaled for Stokes I

6.4 Coronagraphy

The coronagraphic mode of ACS HRC provides improved imaging of low-contrast features, such as circumstellar disks, compared to what is possible with direct, normal imaging. It does this by preventing saturation of the detector and suppressing the diffraction pattern of the central source.

While the coronagraphic data appears at first to be similar to normal ACS images, there are some important differences. The calibration steps for both types of data are the same, except that the coronagraphic flat fields are different and for the moment must be manually applied to the images, The coronagraphic Lyot stop reduces throughput and must be accounted for when performing photometric measurements, and the stop also alters the PSF of field sources.

6.4.1 Flat Fielding

The coronagraphic flat fields are different from normal HRC flats. The most obvious differences are the shadows of the two occulting spots located at the center and the upper left corner. Because the spots are in the aberrated HST beam, there is vignetting out to 0.5 arcsec past the spot edges, causing them to appear somewhat diffuse. In addition, the Lyot stop alters the large scale flat field variation pattern as well as the diffraction pattern caused by dust specks and filter pinholes. Because the occulting spots wander by several pixels over weekly time scales, their effects must be treated separately from the static features in the flat field. Therefore, there are two separate flats that must be applied to coronagraphic images.

The first is the static flat that is essentially like a normal flat, with dust specks, detector response patterns, etc. For certain filters, static flats have been created using either coronagraphic ground test flats (F606W, F814W) or earth flats (F330W, F435W, F475W). These have been modified by replacing the occulting spot patterns with corresponding regions from normal flats. These should be accurate to better than 2% over multi-pixel spatial scales, but they may be less accurate for pixel-to-pixel response. For all other filters, normal HRC flat fields are used as the static flat in the pipeline. These may create some visible artifacts where dust spots occur, causing localized errors of up to 10%. Observers should examine the flat fields used to calibrate their data to verify that any features are not caused by these errors. The STScI pipeline divides all coronagraphic data by the corresponding static flat when calibrated images are requested from the archive.

The second flat that applies to coronographic data is the spot flat that only contains the vignetting patterns of the occulting spots. This flat must be shifted so that the patterns match the positions of the spots at the time a given set of science images was taken. STScI regularly monitors the position of the small occulting spot using earth flats. A table of the positions over time is available on the ACS web site. As of the time of writing, the pipeline does not create nor uses spot flats, so they must be manually applied by the observer. They are available for download from the ACS web site in the HRC Coronagraph Flat Image List of the reference files section at:

http://www.stsci.edu/hst/acs/analysis/reference_files/hrc_c
oron_spotflat_list.html

There is a different spot flat for each fully-supported filter (F330W, F435W, F475W, F606W, F814W). The observer must shift the flat by the amount listed in the spot position table (also in the Reference Files section of the web site). Shifting via bilinear interpolation (an option in imshift in IRAF) is sufficient. Then, the pipeline-produced (or **CALACS**) FLT or CRJ files should be divided by the shifted spot flat. The images must be

divided by the spot flat before geometric correction (the pipeline-produced DRZ files do not account for the spot flat).

6.4.2 PSF Subtraction

While the coronagraph suppresses the central source's diffracted light, it does not affect the halo of scattered light created by the polishing errors in the telescope's optics. That must be subtracted out by using the image of an isolated star with a color similar to that of the source. Observers should observe a PSF reference star immediately before or after their science observations.

The reference PSF must be scaled to match the intensity of the central source in the science observation. The normalization constant can be derived from photometry of non-coronagraphic images of the stars (if such were taken), or by iterative adjustment until the subtraction residuals are visually minimized. Experience shows that the by-eye value is typically accurate to within 5% or less.

The reference PSF must also be aligned with the science image for optimal subtraction. Alignment errors of about 0.05 pixels can be seen by eye in the subtracted image. Subpixel registration is required and is best accomplished by iteratively adjusting the shift of the reference PSF until the subtraction residuals are minimized and symmetrical about the source (discounting any potential circumstellar material). To achieve the best results, higher-order interpolation methods should be used for shifting -- due to the high frequency structure in the scattered light halo, bilinear interpolation is insufficient. In IDL, for instance, the INTERPOLATE routine with cubic convolution interpolation will provide good results. In IRAF imshift using sinc interpolation may also suffice. It may be necessary to iteratively interleave shift and normalization adjustments.

6.4.3 Photometry

Observers with ACS coronagraphic data must account for the effects of the Lyot stop when making photometric measurements. The stop decreases field source intensity by 52.5% relative to the normal HRC imaging mode. Note that this reduction is included in SYNPHOT when the coron specifier is included in the observing mode (e.g., "acs,hrc,f606w,coron"). The stop also broadens the field PSF and places more light into the diffraction spikes and Airy rings (see Chapter 5 in the ACS Instrument Handbook).

6.4.4 Ghosting

A point-source-like ghost appears about 30 pixels to the lower right of every field star and is about 8 magnitudes fainter than the star. There are no ghosts known to be associated with the occulted central source.

6.5 Ramp Filters

6.5.1 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, we will refer to these three segments 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 8.1 and Figure 8.3 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 are available and they cover the region over the HRC chip defined by the HRC aperture (Table 8.2; ACS Instrument Handbook).

The fifteen ramp filters provide a total wavelength coverage of 3700-10,700 Å. Twelve of the filters have a narrow bandpass ($\Delta\lambda/\lambda$) of 2% and three have a medium bandpass of 9% (Table 4.2; ACS Instrument Handbook).

6.5.2 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 (ISR ACS 00-05). Using these delivered profiles, SYNPHOT 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) (ISR ACS 02-01). 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 broadband 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 broadband LP-flats of F475W, F555W, and F606W were used. HRC ramp flats were made from surrogate HRC flats appropriate for the middle segment. These flats have been delivered to the ACS pipeline. The broadband 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 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 30x80 arcsec field of view for the WFC and 11.5x24 arcsec FOV for the HRC. Out of the five middle ramp filters, the FR388N shows the least uniform response (< 10%) over such regions (ISR ACS 02-01).

Preliminary 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

(http://acs.pha.jhu.edu/instrument/filters/ramp/coefficients/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.

6.5.3 Future In-flight Calibrations

While the existing ground-based surrogate broadband L-flats provide some preliminary calibrations, improved monochromatic and continuum ramp L-flats are planned in order to take into account the fact that the transmitted wavelengths and sensitivity varies as a function of position within the ramp aperture. Two calibration steps are planned in Cycle 12:

1. A well-calibrated region in the Orion Nebula is being observed with the primary goal of deriving monochromatic L-flats at specified central wavelengths for (a) the WFC (inner+outer ramp) H-alpha

(FR716N, FR782N) and [OIII] (FR551N, FR601N) ramp filters, (b) WFC middle ramp filters (FR388N, FR459M, FR505N, FR656N, FR914M), and (c) HRC middle ramp filters (FR388N, FR459M, FR505N, FR656N, FR914M).

2. Two well-calibrated standard star fields in 47 Tucanae are being observed with the primary goal of deriving WFC and HRC continuum L-flats for three middle ramp filters (FR656N, FR388N, and FR505N). These filters are selected as they are frequently used to map the H-alpha 6563 Å, [OII] 3727 Å and [OIII] 5007 Å lines, which are of great astrophysical interest, and they include the filters which seem to be the least (FR388N) and most (FR505N) uniform.

6.6 Spectroscopy with the ACS Grisms and Prisms

The ACS WFC has a grism (G800L) which provides slitless spectra at a dispersion of ~40 Å/pixel in the first order over the whole field of view. The same grism used with the HRC produces higher dispersion spectra (~23 Å in first order). A prism (PR200L) available with the HRC provides slitless spectra from the UV cut-off to ~3500 Å. The SBC is fitted with two prisms (PR110L and PR130L) which provide far-UV spectra, the latter blocked below 1230 Å to prevent transmission of Lyman alpha.

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 has been developed by the ST-ECF to enable automatic and reliable extraction of large numbers of spectra. This package, called aXe

(`http://www.stecf.org/software/axe`)

is sufficiently general to be used for other slitless spectra applications. Nevertheless, for the more exacting extractions of small numbers of objects requiring a high degree of interaction, then the relevant packages in noao.twodspec

(`http://www.stecf.org/scripts/irafhelp?twodspec`)

and noao.onedspec, may be more appropriate.

(`http://www.stecf.org/scripts/irafhelp?onedspec`)

6.6.1 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=YES, 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. The filter to be used for the

direct image is set by the AUTOIMAGE=YES parameter or can be user selected in the range of the spectral sensitivity by setting AUTOIMAGE=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. Some information on the jitter during an exposure is available in the ACS Quicklook products (see ACS ISR 99-10). More 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 with the 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 6.10 shows a WFC2 image (one chip only) taken with F775W and Figure 6.11 shows the companion G800L dispersed image (for one chip). This illustrates well the general characteristics of slitless spectra and the characteristics of ACS spectra in particular.

Figure 6.10: Direct image of WFC chip2 in F775W.
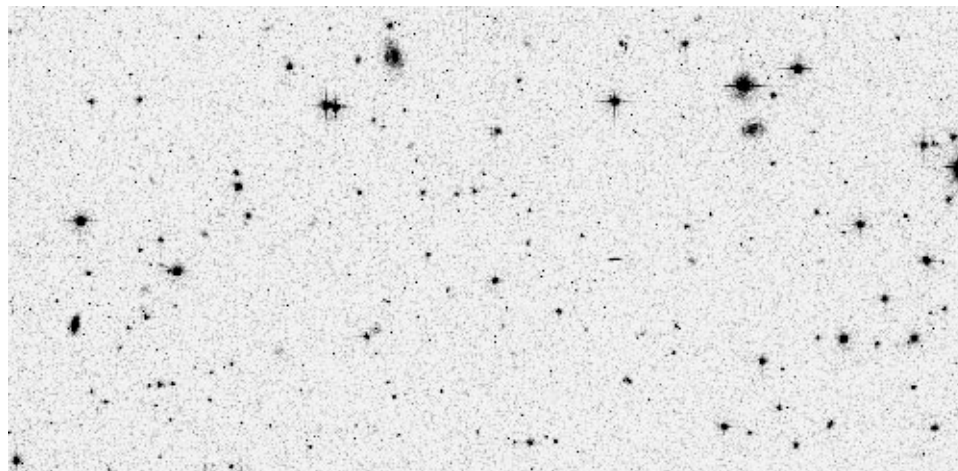
Figure 6.11: WFC G800L slitless image corresponding to the direct image.
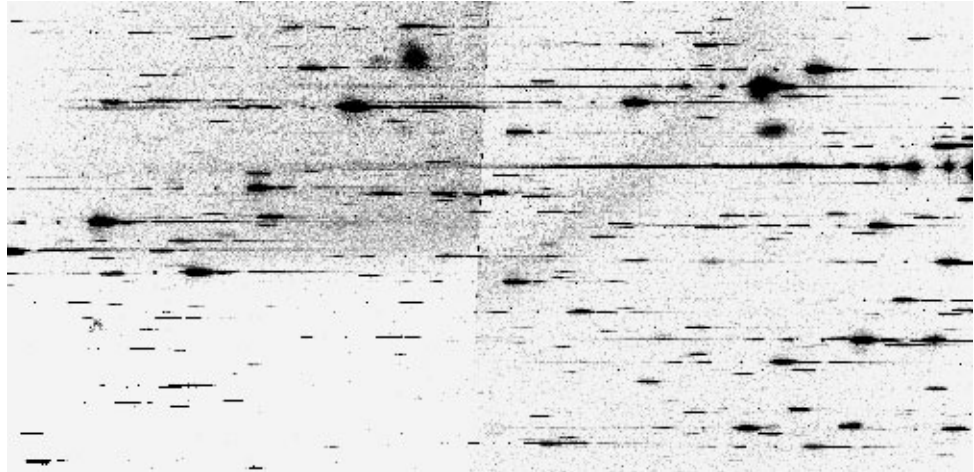
Bright Stars

The brightest objects produce spectra which can extend far across the chip. Although for the G800L grism only about 10% of the flux is in orders other than the first, for bright objects, orders up to +4 and -4 can be detected. These spectra, while in principle can be analyzed (dispersion solutions have been derived for orders -3, -2, -1, 0, 1, 2 and 3 for the WFC G800L for example; see ACS ISR 03-01), provide a strong source of contamination for fainter spectra. In addition for the grism, the 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 spectrometry 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 01-02).

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 alpha 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 6.10 and Figure 6.11 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 that of 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, which 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 will be a source of contamination but correct flagging is hampered by the lack of knowledge of the precise image position.

6.6.2 Pipeline Calibration

The direct image half of the grism dispersed and direct image pair can be fully reduced by **CALACS** including bias subtraction, dark subtraction, flat fielding and inclusion of photometric zero point in the header (processing by ACS2D, see Section 3.4.3). However on account of the individual wavelength coverage of each object in the dispersed image, the pipeline reduction of slitless images is necessarily more restricted. Only the steps of the ACSCCD routine, involving noise and data quality initialization, bias subtraction and trimming of the overscan are applied. In addition, division by a dummy flat field (unity filled) is also applied in order to correct for the CCD gain, providing an image in electrons. This "flat" image (unity) is supplied as a pixel-to-pixel flat field file. Application of this dummy flat corrects very well the offsets in level between both quadrants of the WFC1 and WFC2 chips which have differing gain values. For the single quadrant detectors (HRC and SBC) this step merely corrects

the ADU values by a single gain value. Table 6.7 shows the calibration switches appropriate for a WFC G800L frame.

Table 6.7: 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?

No unique photometric keywords can be attached to all the spectra in a slitless image, so the photometric header keywords are left as default - see Table 6.7.

Table 6.8: ACS Slitless spectra (grism or prism) photometric keywords

Keyword	Value	Description
PHOTMODE	'ACS WFC2 G800L'	observation con
PHOTFLAM	0.00000E+00	inverse sensitivity, ergs/cm2/Ang/electron
PHOTZPT	0.00000	ST magnitude zero point
PHOTPLAM	0.00000	Pivot wavelength (Angstroms)
PHOTBW	0.00000	RMS bandwidth of filter plus detector

In contrast to direct imaging data, 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 during

the extraction of spectra in the aXe task aXe_PETFF (see "Main Steps of an aXe Extraction" in Section 6.6.3). 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.

For all subsequent reductions of slitless data, the pipeline product FLT files should be the starting point. The units of the SCI and ERR extensions are electrons.

6.6.3 Extracting and Calibrating Slitless Spectra

The set of C programs under the package name of aXe (also available in **pyraf.stsdas.acs**) provides a streamlined method for extracting spectra from the ACS slitless spectroscopy data. The starting point is always a direct image, a dispersed slitless image and a catalogue of the objects on the direct image. 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 and enables the full calibration of extracted spectra. Since the configuration file is quite general, it can be adapted to other HST instruments with slitless spectroscopy facilities, such as NICMOS.

There is a detailed aXe manual and so a only a brief outline of its use will be presented here. A demonstration package with direct and slitless spectroscopy images and a script to drive the extraction can be downloaded from the aXe web page:

(http://www.stecf.org/software/aXe).

All catalogues are ASCII files and all image files are in FITS format. The data [SCI], data quality [DQ] and statistical error [ERR] data are accessed through the extension names.

The input slitless images for aXe are assumed to be in e⁻/sec. Thus the FLT file must be normalized by the exposure time. Use of files in these units ensures that errors are correctly handled and propagated through the extraction process.

Main Steps of an aXe Extraction

1. Produce a deep catalogue (ASCII) of the objects on the direct image. The aXe package was designed to use a catalogue generated by Sextractor (http://terapix.iap.fr/cplt/oldSite/soft/releases.html). A minimum number of catalogue keyword parameters is required (see http://www.stecf.org/software/aXe/sex/aXe.param) to run aXe

from this catalogue, but other catalogue formats could be converted to hold columns with these names.

2. Run aXe_SEX2GOL to convert the direct image Sextractor catalogue into a list of slitless image objects (called a Grism Object List, GOL). This task reads the WCS of both direct and slitless images and so can allow for shifts (usually small) in pointing between both images.

3. Run aXe_GOL2AF to convert the output catalogue of aXe_SEX2GOL into a list of BEAM specifications where slitless spectra are predicted, called an Object Aperture File (OAF). A BEAM is defined as a single order of the spectrum of a single cata-logued object. This task requires the configuration file for the slitless mode. aXe_GOL2AF can be used to specify the extraction boxes for objects (OAF) and for the corresponding region for the background (Background Aperture File, BAF).

4. Run aXe_AF2PET to copy the pixel values for each aperture defined by an OAF or BAF to a Pixel Extraction Table (PET). For the object apertures the PET is referred to an Object PET (OPET) and for background as a Background PET (BPET). A PET is a FITS binary table which for each pixel within an aperture holds trace information (such as wavelength and distance from the predicted position of the spectrum trace) as well as data value, statistical error and data quality.

5. Run aXe_BE to generate a FITS file of the background used in sub-tracting the object spectra. There are various interpolation options for the background including the option to use the whole column for background estimation. The resulting FITS file is not used for further processing but provides the user with an image of the background used. This is very useful for checking purposes or could be used for subsequent processing, outside aXe.

6. Run aXe_PET2CONT to check whether any pixels within a BEAM are contaminated by pixels from another BEAM. If a pixel within one beam is contaminated by one or more BEAM of other targets, the contamination column of the PET file is incremented.

7. Run aXe_PETFF to flat field a PET table by a flat field cube. The flat field cube is formed by a polynomial fit to the flat field variation of each pixel with wavelength over the wavelength range of the detector. Both the OPET and BPET must be identically flat fielded.

8. Run aXe_PET2SPC to extract each BEAM of each spectrum from its pixels, collapsing the 2D spectrum to a 1D spectrum, subtracting the background and applying wavelength and flux calibration. The result is a single FITS binary table containing all the spectra extracted from the slitless image.

9. Run aXe_STAMP to produce a "stamp" image of the region of each BEAM of each extracted object. This is also a single FITS binary table for the slitless image. This tool is primarily used for checking the veracity of an extraction. The stamp images can be rectified to give an orthogonal wavelength cross-dispersion image, similar to a long slit image.

Output Products: Extracted Spectra

The primary output of aXe 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; the correspondence with order being defined by the configuration file. For example the G800L WFC 1st order is BEAMA in the configuration file ACS.WFC.CHIP2.Cycle11.4.conf see:

http://www.stecf.org/instruments/acs/calib/Cycle11/WFC/ACS.WFC.CHIP1.Cycle11.4.conf.

The first order spectrum of the object with number 8 in the input Sextractor catalogue for the slitless image grism.fits, would therefore be accessed as grism.SPC.fits[BEAM_8A]. The table contains 12 columns listing wavelength, total and extracted and background counts and their errors, the calibrated flux and error, the weight and 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 spectra were extracted.

The stamp image is also a multi-extension fits file with each extension containing the image of a single extracted beam. The other output products -- background image and contamination map are simple FITS files. The PET files are multi-extension FITS binary tables organized similarly to the SPC file, but are files for intermediate products and their contents are not normally inspected.

Tuning the Extraction Process

All these tasks have input and output parameters so they may be tuned in various ways. See the aXe manual for further details.

aXe was planned as a set of tools which can be used to build up a full reduction of a slitless spectroscopy image. For example, several of the tools are run repeatedly as part of a full reduction. It proves most convenient to run aXe as a script which successively calls the required tasks. An example script is contained in the demo and is described in detail in the aXe manual. Often a reduction will need to be repeated, for example, changing the cut between the magnitude of objects and BEAMs to be extracted and those "marked" (i.e., considered to contribute to contamination but are not extracted, such as the zeroth order of grism spectra or saturated spectra).

The output products from aXe consist of ASCII files, FITS images and FITS binary tables. The FITS binary tables can be assessed using the tasks in the **stsdas.ttools** package and wavelength flux plots, with error bars, can

be plotted using **stsdas.graphics.stplot.sgraph** (see Figure 6.12 for an example). When there are many detected spectra on a single image, such as for WFC G800L, then a dedicated task aXe2web is available (http://www.stecf.org/software/aXe) to make html pages consisting of direct image cut-outs, stamp images and 1D spectra (see Figure 6.13). This enables convenient browsing of large numbers of spectra.

Two facets of the reduction of ACS slitless data are not yet handled. First, fringing occurs in the WFC and HRC detectors arising from multiple reflections within the CCD layers, with an amplitude highest in the wavelength range 8000-9500Å (see ACS ISR 03-12). Fringing modulates the signal of the slitless spectra and should be corrected for the highest fidelity applications. Second, for optimal extraction of faint spectra, weights should be applied to the data points during the extraction process. While the PET files have a column for weight, which is applied during the extraction process (aXe_PET2SPC), the current version sets the value of the weights to unity. A task to reset the weights based on the cross-dispersion profile shape is required.

Figure 6.12: Example WFC G800L extracted spectrum.
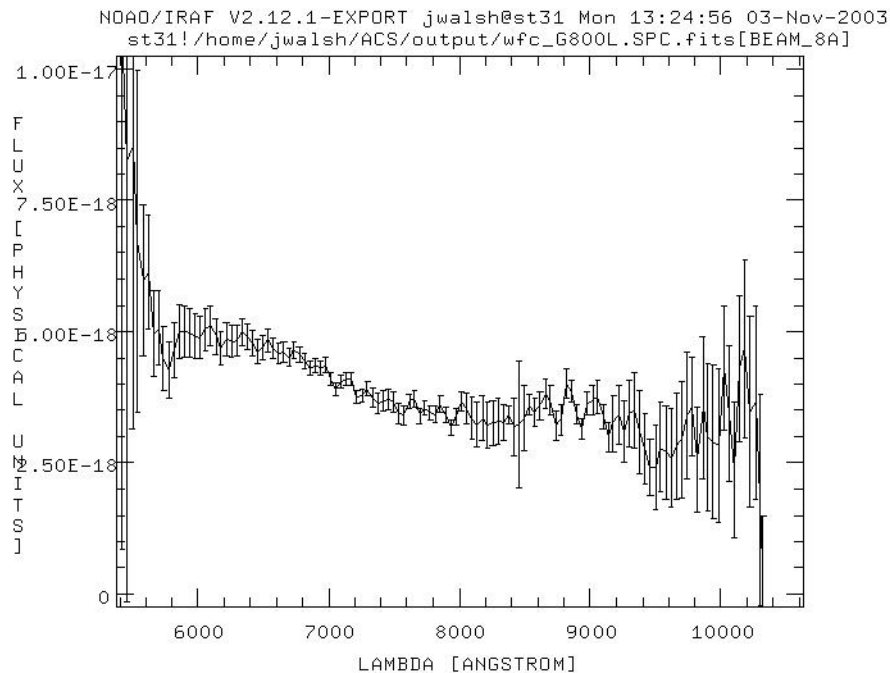
Figure 6.13: Web page produced by axe2Web for browsing extracted spectra.
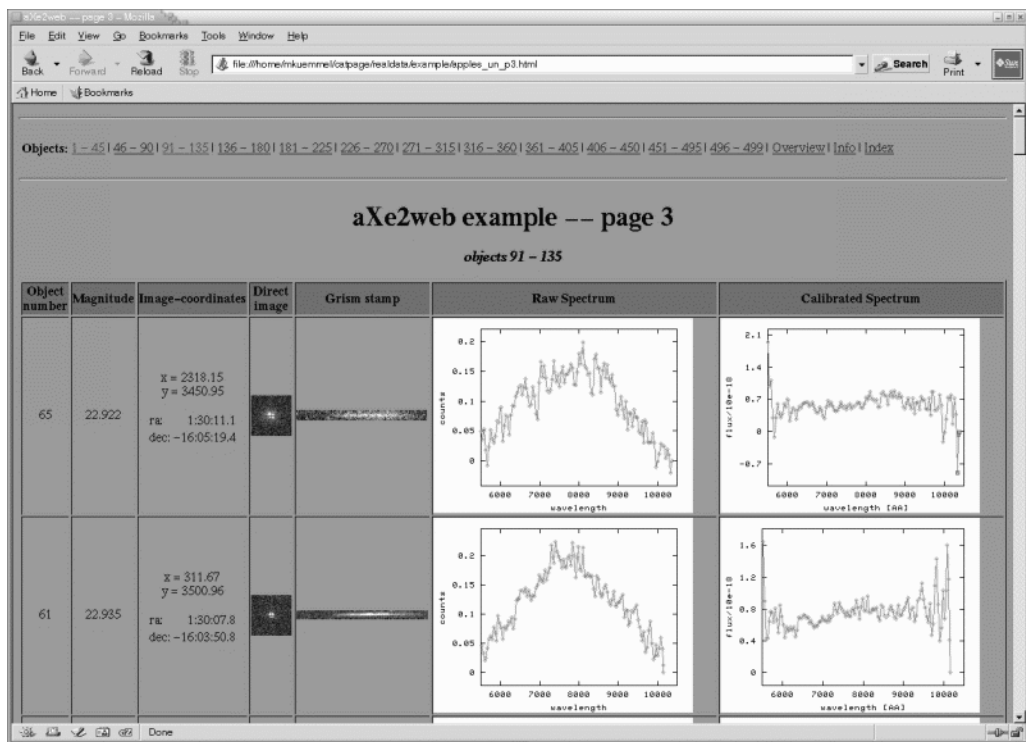
6.6.4 How to Handle Dithered Slitless Spectroscopy Data

The necessity to dither ACS imaging data in order to improve the spatial resolution and to facilitate removal of hot pixels, applies equally well to slitless spectroscopy data. Typically for long slitless spectroscopic exposures, and especially for parallel slitless data, the data taking is broken into many sub-orbit dithered exposures. During the combination of ACS imaging data by **PyDrizzle**, the correction for the large geometrical distortion has to be taken into account. However slitless data has both a spatial and a spectral dimension and correction for the geometrical distortion, while applicable to the cross-dispersion direction, will distort the wavelength dispersion. Tracing the light path for slitless 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 and geometrically corrected slitless data for WFC (the ACS channel in which the geometrical distortion is largest) 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%. This is clearly disadvantageous but a configuration file for extracting dithered WFC G800L images is available at

http://www.stecf.org/instruments/acs/calib/Cycle11/WFC/ACS.
WFC.drizzle.2.conf

although this configuration file only handles first and zeroth order spectra. No flat field can be applied to drizzled slitless combined data since each pixel in the drizzled data is a combination of multiple input pixels, each with different wavelength response.

The **MultiDrizzle** script (see Section 4.5.2) however provides a very convenient tool for combining many dithered slitless spectral images in order to provide a deep slitless image, free of cosmic rays, for examination. The companion direct images can also be combined with **MultiDrizzle** to derive a deep catalogue. However for extraction of the spectra, it is recommended to extract the spectra from each image separately and combine the 1D spectra for each object. **MultiDrizzle** is then applied to provide data quality images with the cosmic rays and hot pixels flagged, so that these affected pixels are not included in the extracted spectra. Thus **MultiDrizzle** is run on all the slitless data files in a set, the data quality arrays in the FLT files are updated, and then aXe is run on this set of data files. The individual extracted wavelength, flux spectra can then be converted to FITS images and noao.onedspec.scombine used, for example, to merge all the spectra of each object. An alternative, preferable, approach under investigation is to combine the images of the slitless spectra in 2D. The pixels from the PET files from different exposures could then be drizzled and cosmic rays flagged using **MultiDrizzle**. Spectra could be extracted using axe_PET2SPC to produce the final combined wavelength and flux calibrated 1D data.

6.6.5 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 03-01 and ACS ISR 03-07). 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 zero point and the transfer of the zero point from the direct 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 zero point 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 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 6.6.2). 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 (<5,800 Å) and to the red (>10,000 Å), 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 end of extracted flux calibrated spectra, whose reality may be considered suspect.

Index

W

Warm pixels ACS:5—9
Weight Image ACS:4—20, ACS:4—53
WFC
    capabilities ACS:1—1
    chip offsets ACS:4—2
    description ACS:1—2, ACS:1—3
    distortion ACS:1—4, ACS:4—5

Z

Zeropoint
    determining ACS:6—2
    photometric ACS:6—3