

Version 6.0

Galaxy Cluster MACS J1149.5+223
HST ACS/WFC and WFC3/IR

WFC3 Data Handbook

June 2024



STScI | SPACE TELESCOPE
SCIENCE INSTITUTE

3700 San Martin Drive
Baltimore, MD 21218
<https://hsthelp.stsci.edu>

WFC3 Data Handbook	3
Acknowledgments	5
What's New in This Revision	6
Preface	8
Handbook Structure	9
Typographic Conventions	10
Chapter 1: WFC3 Instruments	13
1.1 Instrument Overview	14
1.2 The UVIS Channel	16
1.3 The IR Channel	18
Chapter 2: WFC3 Data Structure	20
2.1 Types of WFC3 Files	21
2.2 WFC3 File Structure	31
2.3 Data Storage Requirements	40
2.4 Headers and Keywords	41
Chapter 3: WFC3 Data Calibration	64
3.1 The calwf3 Data Processing Pipeline	65
3.2 UVIS Data Calibration Steps	72
3.3 IR Data Calibration Steps	86
3.4 Pipeline Tasks	98
3.5 Manual Recalibration of WFC3 Data	114
Chapter 4: WFC3 Images: Distortion Correction and AstroDrizzle	118
4.1 WFC3 Geometric Distortion	119
4.2 Distortion Corrections and Image Combination	123
4.3 References	126
Chapter 5: WFC3 UVIS Sources of Error	127
5.1 Gain and Read Noise	128
5.2 Bias Subtraction	130
5.3 Dark Current and Hot Pixels	131
5.4 UVIS Flat Fields	137
5.5 Image Anomalies	147
5.6 Generic Detector and Camera Properties	152
5.7 UVIS Photometry Errors	165
5.8 References	166
Chapter 6: WFC3 UVIS Charge Transfer Efficiency - CTE	169
6.1 Overview	170
6.2 CTE Losses And Background	171
6.3 The Nature Of CTE Losses	172
6.4 The Pixel-Based Model	175
6.5 Empirical Corrections	180
6.6 Dealing With CTE Losses in WFC3 UVIS Images	181
6.7 Sink Pixels	182
6.8 References	183
Chapter 7: WFC3 IR Sources of Error	184
7.1 WFC3 IR Error Source Overview	185
7.2 Gain	186
7.3 WFC3 IR Bias Correction	187
7.4 WFC3 Dark Current and Banding	188
7.5 Blobs	193
7.6 Detector Nonlinearity Issues	198
7.7 Count Rate Non-Linearity	200
7.8 IR Flat Fields	201
7.9 Pixel Defects and Bad Imaging Regions	209

7.10 Time-Variable Background	213
7.11 IR Photometry Errors	219
7.12 References	222
Chapter 8: Persistence in WFC3 IR	225
8.1 Persistence in WFC3 IR	226
8.2 Evaluating the Amount of Persistence in Science Images	230
8.3 Mitigating the Effects of Persistence	233
8.4 References	236
Chapter 9: WFC3 Data Analysis	237
9.1 Photometry	238
9.2 Astrometry	262
9.3 Spectroscopy	265
9.4 STSDAS, STSCI_PYTHON and Astropy	276
9.5 Specific Tools for the Analysis of WFC3	277
9.6 References	283
Chapter 10: WFC3 Spatial Scan Data	288
10.1 Analysis of Scanned Data	289
10.2 IR Scanned Data	290
10.3 UVIS Scanned Data	294
10.4 References	295

WFC3 Data Handbook

Version 6.0 - June 2024

Wide Field Camera 3 Data Handbook

Table of Contents

[Expand all](#) [Collapse all](#)

User Support

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

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

Additional Resources

Information and other resources are available from the WFC3 website:

- <http://www.stsci.edu/hst/instrumentation/wfc3>

Revision History

Version	Date	Editor
6.0	June 2024	Amanda Pagul and Isabel Rivera, et al.
5.0	Sep 2021	Kailash Sahu, et al.
4.0	May 2018	Mario Gennaro, et al.
3.0	January 2016	Susana Deustua, et al.
2.1	May 2011	Abhijith Rajan, et al.
2.0	October 2010	Abhijith Rajan, et al.
1.0	January 2009	Jessica Kim Quijano, Howard Bushouse, and Susana Deustua

Contributors

This document is written and maintained by the WFC3 Team in the Instruments Division of STScI. The WFC3 Team, at the time of this writing, consists of M. Alam, J. Anderson, S. Baggett, V. Bajaj, A. Calamida, F. Dauphin, M. de la Pena, L. Dressel, J. D. Green, K. Huynh, H. Khandrika, B. Kuhn, J. Mack, M. Marinelli, C. Martlin, P. R. McCullough, A. O'Connor, A. Pagul, A. Pidgeon, M. Revalski, I. Rivera, S. Shenoy, D. Som.

Citation

In publications, refer to this document as:

- Pagul, A. and Rivera, I., et al., 2024, "WFC3 Data Handbook", Version 6.0, (Baltimore: STScI).

Acknowledgments

Current Members of the WFC3 Team

M. Alam, J. Anderson, S. Baggett, V. Bajaj, A. Calamida, F. Dauphin, M. de la Pena, L. Dressel, J. D. Green, K. Huynh, H. Khandrika, B. Kuhn, J. Mack, M. Marinelli, C. Martlin, P. R. McCullough, A. O'Connor, A. Pagul, A. Pidgeon, M. Revalski, I. Rivera, S. Shenoy, D. Som.

Former Members of the WFC3 Team

W. Baggett, H. Bond, T. Borders, M. Bourque, A. Bowers, G. Brammer, T. Brown, H. Bushouse, L. Cawley, T. Dahlen, S. Deustua, M. Dulude, M. Durbin, M. Fall, D. Figer, J. Fowler, M. Gennaro, M. Giavalisco, S. Gonzaga, C. Gosmeyer, H. Gunning, D. Hammer, C. Hanley, G. Hartig, B. Hilbert, M. Jones, J. Kalirai, P. Knezek, A. Kovacs, V. Kozhurina-Platais, H. Kurtz, R. Kutina, J. Lee, C. Lisse, K. Long, O. Lupie, J. MacKenty, M. McKay, J. Medina, I. Momcheva, M. Montes, N. Nikolov, K. Noeske, C. Pavlovski, N. Pirzkal, J. Kim Quijano, M. Rafelski, A. Rajan, N. Reid, A. G. Riess, M. Robberto, R. Ryan, E. Sabbi, C. Shanahan, K. C. Sahu, M. Sosey, K. Stevenson, M. Stiavelli, B. Sunnquist, D. Taylor, A. Viana, M. Wong

Thanks

The entire WFC3 team helped and actively participated in updating this Data Handbook. We thank Hsing-Ying Liu, Mariarosa Marinelli, Paul Mulgrew, and Jenna Ryon for their invaluable contributions to the editing and production of this Handbook.

What's New in This Revision

The entire Data Handbook has been revised and updated to include the recent publications and results from various calibration programs. We list below the main changes made in different sections.

- Information on Hubble Advanced Products added; additional auxiliary products described ([Section 2.1](#))
- **calwf3** procedure updated ([Chapter 3](#))
- UVIS calibration procedure and flowchart updated ([Section 3.2](#))
- IR calibration procedure and flowchart updated ([Section 3.3](#))
- Astrometric alignment updates described ([Section 4.2.1](#))
- generation of Hubble Advanced Products described ([Section 4.2.1](#))
- inclusion of updates to **Drizzlepac** ([Section 4.2.2](#))
- updated UVIS gain values ([Section 5.1.1](#))
- 2021 and 2022 superbias reference file updates ([Section 5.2.1](#))
- updated UVIS hot pixel characterization ([Section 5.3](#))
- updated description of UVIS full well saturation flagging ([Section 5.6.1](#))
- updated saturation map reference file ([Section 5.6.1](#))
- updated analysis of UVIS shutter timing jitter ([Section 5.6.4](#))
- new figure detailing photometric noise sources ([Section 5.6.4](#))
- new figures comparing UVIS CCD detection layer thickness and fringing in quad filter flats ([Section 5.6.5](#))
- updated fringing comparison in different bands ([Section 5.6.5](#))
- discussion of x-CTE contribution to photometry ([Section 6.1](#))
- new figure visualizing the "trails" that result from CTE taking charge away from downstream pixels, depositing it into upstream pixels ([Section 6.3](#))
- updated characterization of post-flash needed to reduce CTE losses for faint sources ([Section 6.4](#))
- updated discussion of **hst1pass** ([Section 6.4](#))
- updated information of where to find an analysis of long term UVIS CTE behavior ([Section 6.5](#))
- updated IR gain values ([Section 7.2](#))
- description of new Jupyter notebook to flag blobs using a convolutional neural network (CNN) ([Section 7.5.1](#))
- updated detector map of new blobs ([Section 7.5.2](#))
- updated dithering strategies for blob correction ([Section 7.5.3](#))
- updated collection of time-dependent IR bad pixel tables ([Section 7.9.1](#))
- updated IR Bad Pixel Mask ([Section 7.9.1](#))
- description of 4 new notebooks that are designed to aid users in working with IR file structures and visualization, as well as identifying and correcting issues with data ([Section 7.10](#))
- updated demonstration of background subtraction in IR images affected by scattered light ([Section 7.10](#))
- updated figure detailing deviation of measurements of GF-153 (a spectrophotometric star) versus signal-to-noise ([Section 7.11](#))
- updated photometric repeatability figure and description ([Section 7.11](#))
- figure updated showing pixels affected by persistence subtracted ([Section 8.3](#))
- Photometry section completely rewritten, including a checklist of steps for measuring photometry of sources on images collected with the WFC3 UVIS and IR detectors, and a description of each ([Section 9.1](#))
- discussion of **slitlessutils** and additional resources ([Section 9.3](#))
- new sky image for G280 ([Section 9.3.4](#))
- updated discussion and guidance on usage of **HSTaXe** ([Section 9.3.6](#))

- updated discussion of G102 and G141 background sky signal and calibration procedure ([Section 9.3.6](#))
- updated flux monitor figure for G102 and G141 ([Section 9.3.7](#))
- updated discussion of WFC3 analysis tools and corresponding Jupyter notebooks available for users ([Section 9.5](#))
- updated PSF modeling notebook ([Section 9.5.2](#))
- updated Jupyter notebook on using **stsynphot** to compute photometric keyword values ([Section 9.5.3](#))
- updated discussion of WFC3/UVIS zeropoints ([Section 9.5.3](#))
- updated Jupyter notebook for UVIS time-dependent photometry ([Section 9.5.3](#))
- updated discussion of Jupyter notebooks to help mitigate variable IR background ([Section 9.5.4](#))
- updated discussion of **HSTaXe** and **slitlessutils** ([Section 9.5.6](#))
- updated discussion of new tools developed by ACS for satellite trail flagging ([Section 9.5.9](#))
- updated figure of charge accumulation in a pixel sampled with RAPID mode for staring and scanned modes ([Section 10.2](#))
- updated sensitivity losses in IR spatially scanned spectra ([Section 10.2](#))
- updated sensitivity losses in UVIS spatially scanned spectra ([Section 10.3](#))

Preface

How to Use this Handbook

This handbook is designed to provide users with an introduction on how to understand, manipulate, and analyze data from the Wide Field Camera 3 (WFC3), which was installed on-board the Hubble Space Telescope (HST) during the 2009 Servicing Mission 4 (SM4). It is presented as an independent and self-contained document, referred to as the “WFC3 Data Handbook.”

Information about HST not specific to WFC3 is not discussed here. Users are referred to a companion volume, [Introduction to the HST Data Handbooks](#), for more general information about the details of acquiring data from the HST archive, HST file formats, and general purpose software for displaying and processing HST data.

For detailed information about the instrument’s capabilities and design, and how to plan observations, users should refer to the [WFC3 Instrument Handbook](#).

Handbook Structure

The WFC3 Data Handbook is organized in 10 chapters, which discuss the following topics:

- [Chapter 1 - WFC3 Instruments](#) - provides a brief summary of the capabilities and design of the Wide Field Camera 3.
- [Chapter 2 - WFC3 Data Structure](#) - describes WFC3 data products as would be obtained from the [MAST](#), file name conventions, and auxiliary files. The structure of WFC3 images is explained, and an extensive list of all the header keywords found in typical WFC3 data files is included.
- [Chapter 3 - WFC3 Data Calibration](#) - explains the flow of calibration steps performed on WFC3 data with the STScI software [calwf3](#). It provides detailed information on what each step does with the data, and how the data header keywords are used and modified as the data are processed. It also describes when the data should be manually recalibrated and how to do it.
- [Chapter 4 - WFC3 Images: Distortion Correction and AstroDrizzle](#) - provides the current knowledge about WFC3 detector distortion. A brief description of what the task AstroDrizzle does for WFC3 data is also presented.
- [Chapter 5 - WFC3 UVIS Sources of Error](#) - provides users with all the currently known error sources in the data for the UVIS detectors. Possible methods of mitigating the error sources are also discussed.
- [Chapter 6 - Charge Transfer Efficiency \(CTE\) in WFC3/UVIS](#)- describes various aspects of CTE losses and provides users with methods for mitigating CTE effects.
- [Chapter 7 - WFC3/IR Sources of Error](#) - provides users with all the currently known error sources in the data for the IR detector. Possible methods of mitigating the error sources are also discussed.
- [Chapter 8 - Persistence in WFC3/IR](#) - describes the effects of persistence on IR images, and strategies for correcting for persistence effects.
- [Chapter 9 - WFC3 Data Analysis](#) - includes discussions for improved analysis of WFC3 data. The discussions are primarily focused on photometry, astrometry and grism spectrophotometry.
- [Chapter 10 - WFC3 Spatial Scan Data](#) - provides some pointers on how to reduce and analyze spatial scanned images.

For the latest information regarding WFC3 performance and calibration, users are advised to consult the WFC3 web pages located at:

- <http://www.stsci.edu/hst/instrumentation/wfc3>

Typographic Conventions

[Visual Cues](#)
[Comments](#)
[List of Acronyms](#)

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


Visual Cues

The following typographic cues are used:


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

Comments

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

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

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

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

 *Information especially likely to be updated on the WFC3 Web site is indicated by this symbol.*

List of Acronyms

Acronym	
A-to-D	Analog to Digital conversion
ACQ	Acquisition
ACS	Advanced Camera for Surveys

ADC	Analog-to-Digital converter
APT	Astronomer's Proposal Tool
CASTLE	HST optical simulator
CSM	Channel Select Mechanism
CT	Crosstalk
CTE	Charge Transfer Efficiency
CVZ	Continuous Viewing Zone
DN	Data Number
DQ	Data quality
DS9	SAOimage astronomical image visualization application
ESA	European Space Agency
ETC	Exposure Time Calculator
FGS	Fine Guidance Sensor
FITS	Flexible Image Transport System
FOV	Field of View
FSM	Filter Select Mechanism in WFC3 IR channel
HOPR	Hubble Operations Problem Report
IDC	Image distortion correction
IR	Infrared channel on WFC3
IRAF	Image Reduction and Analysis Facility
ISR	Instrument Science Report
MAST	Mikulski Archive for Space Telescopes
MULTIACCUM	Accumulation mode (multiple non-destructive reads)
NICMOS	Near Infrared Camera and Multi-Object Spectrometer
OPUS	Operational Pipeline Unified System
OS	Optical Stimulus
OTA	Optical Telescope Assembly
PAM	Pixel Area Map
POM	Pick-off Mirror

PSF	Point Spread Function
QE	Quantum efficiency
SAA	South Atlantic Anomaly
SED	Spectral Energy Distribution
SIC&DH	Science Instrument Command and Data Handling unit
SIP	Simple Image Polynomial
SMOV	Servicing Mission / Orbital Verification
SOFA	Selectable Optical Filter Assembly in WFC3 UVIS channel
ST-ECF	Space Telescope - European Coordinating Facility
STScI	Space Telescope Science Institute
STSDAS	Space Telescope Science Data Analysis Software
TIR	Technical Instrument Report
TV3	Thermal Vacuum campaign 3
UVIS	Ultraviolet-Visible channel on WFC3
WCS	World Coordinate System
WF/PC-1	Wide-field Planetary Camera 1, on-board HST Apr 1990 to Nov 1993
WFPC2	Wide-Field Planetary Camera 2, on-board HST Dec 1993 to May 2009

Chapter 1: WFC3 Instruments

Chapter Contents


- [1.1 Instrument Overview](#)
- [1.2 The UVIS Channel](#)
- [1.3 The IR Channel](#)

1.1 Instrument Overview

Wide Field Camera 3 (WFC3) is a fourth-generation imaging instrument installed on the Hubble Space Telescope during Servicing Mission 4 in 2009. It replaced the extraordinarily successful Wide Field Planetary Camera 2 (WFPC2), thereby ensuring and enhancing the imaging capability of HST in the remaining years of its observing lifetime. WFC3 is the only HST instrument developed as a facility instrument by the HST Project.

WFC3 occupies WFPC2's spot in HST's radial scientific-instrument bay, where it obtains on-axis direct images. Light entering from the HST Optical Telescope Assembly (OTA) is intercepted by the flat 45 degree WFC3 pick-off mirror (POM) and directed into the instrument. A channel-select mechanism (CSM) inside WFC3 then diverts the light to the IR channel via a fold mirror, or when the CSM is moved out of the beam, the light enters the UVIS channel unimpeded. Because of this design, only a single channel, either UVIS or IR, can be used at any one time. Figure 1.1 shows a schematic diagram of the instrument's optical and mechanical layout. The main characteristics of each channel are summarized in the following sections. For a technical description of the instrument's properties, performance, operations, and calibration, please refer to the [WFC3 Instrument Handbook](#).

Figure 1.1: Schematic optical layout of the WFC3 instrument

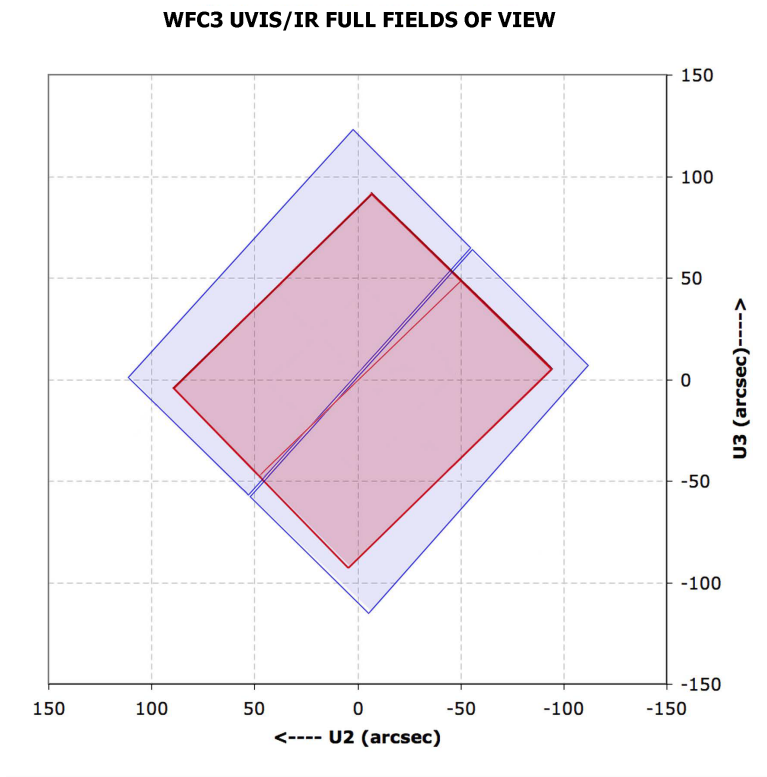


For schematic simplicity, the incoming OTA beam and POM have been rotated into the plane of the optical diagram. The actual incoming OTA beam direction is into the page and then reflected by the POM into the instrument (yellow path). The CSM is then either moved out of the beam to allow light to pass into the UVIS channel (blue path), or the CSM is moved into the beam to direct light into the IR channel (red path). Each channel contains mechanical and optical elements which allow for focusing and alignment, as well as anamorphic aspherical correctors to correct for the $\sim 1/2$ wave spherical aberration from the HST primary mirror. Filters and grisms are contained in the UVIS selectable optical filter assembly (SOFA) and the IR filter selection mechanism (FSM). The UVIS channel has a mechanical shutter, while the IR channel is shuttered electronically by the detector via a reset of the entire array. The IR channel's FSM also contains an aluminum blocker ('blank') in one of the filter slots that is in place when the channel is not in use. Light is detected by either the UVIS CCDs or the IR HgCdTe focal-plane array. A separate calibration subsystem with deuterium and tungsten bulbs provides flat-field illumination for both channels.

WFC3 provides HST with high-sensitivity, high-resolution, wide-field survey capability covering a broad wavelength range, from the near-UV at 200 nm to the near-IR at 1700 nm. WFC3 is comprised of two channels, each optimized for a specific wavelength regime:

- Ultraviolet-Visible channel (UVIS): sensitive to 200-1000 nm, with a pair of CCDs covering a $162 \times 162''$ total field of view (Figure 1.2), with a plate scale of $0.040''/\text{pixel}$ and a focal ratio of $f/31$.
- Infrared channel (IR): sensitive to 800-1700 nm, an IR focal plane array covering a $136 \times 123''$ field of view (Figure 1.2), with a plate scale of $0.13''/\text{pixel}$ and a focal ratio of $f/11$.

Figure 1.2: Schematic view of the WFC3 UVIS (blue) and IR (red) full fields of view.



WFC3 provides the user with:

- 62 wide-, medium-, and narrow-band filters in the UVIS channel
- 15 wide-, medium-, and narrow-band filters in the IR channel
- 3 grisms: 1 in the UVIS channel and 2 in the IR channel

1.2 The UVIS Channel

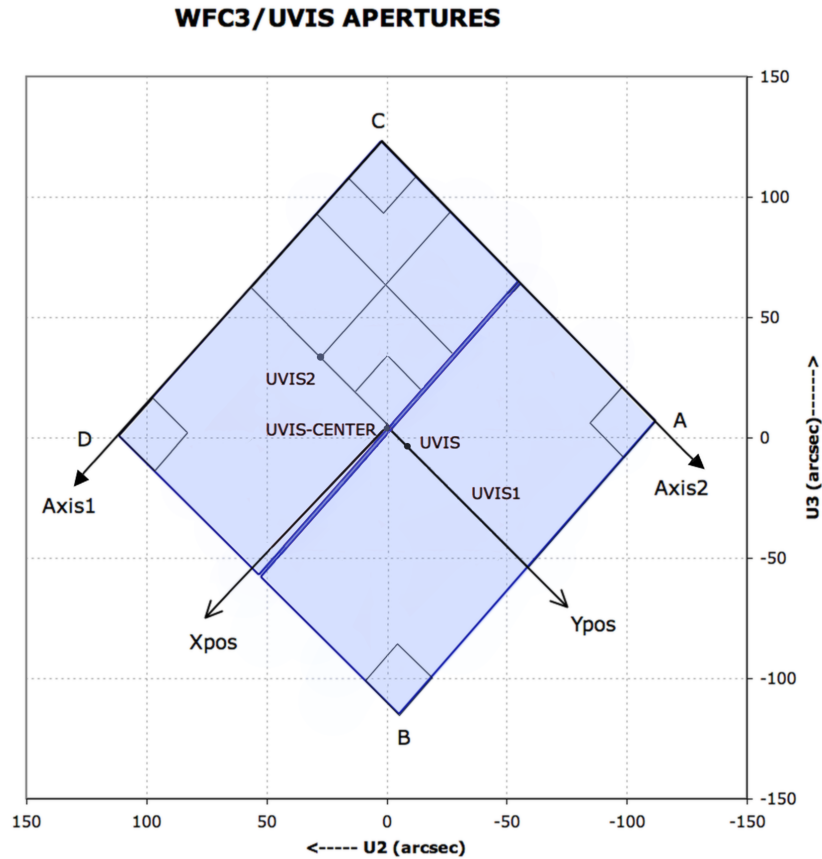
The UVIS channel employs a mosaic of two 4096×2051 Teledyne e2v (formerly Marconi Applied Technologies Ltd.) CCDs, with $\sim 0.040''/\text{pixel}$ plate scale, covering a nominal $162'' \times 162''$ field of view. These CCDs are thinned, back-illuminated devices cooled by thermo-electric cooler (TEC) stacks and housed in sealed, evacuated dewars with fused silica windows, nearly identical to the ones used for ACS. The spectral response of the UVIS CCDs is optimized for imaging from the near-UV at 200 nm to visible wavelengths at 1000 nm. The two CCDs are butted together with a 35-pixel gap between the two chips ($\sim 1.4''$ on the sky). The minimum UVIS exposure time is 0.5 sec. The dynamic range for a single exposure is ultimately limited by the depth of the CCD full well ($\sim 70,000$ electrons), which determines the total amount of charge that can accumulate in any one pixel during an exposure without saturation.

The UVIS detector operates only in ACCUM mode to produce time-integrated images. Cosmic rays affect all UVIS exposures, therefore observations should be broken into multiple exposures or dither patterns whenever possible, to allow removal of cosmic rays in post-observation data processing.

WFC3 recycles hardware used in WF/PC-1 to house the complement of filters for the UVIS channel. The Selectable Optical Filter Assembly (SOFA) contains a stack of 12 wheels housing 48 physical elements covering the UV/Visible range: 42 full-frame filters, 5 quad filters (2×2 mosaics providing 4 different bandpasses), and 1 grism, giving a total of 63 spectral elements. Each wheel has an open slot such that when an observation is taking place, the appropriate wheel is rotated to place the desired filter in the beam, while the other wheels place the open slot in the light path.

Figure 1.3 shows a schematic of the UVIS channel aperture projected onto the sky with respect to the U2/U3 reference frame. (For definitions of the coordinate systems in the figure, please refer to Section 7.4.3 of the *WFC3 Instrument Handbook*). The WFC3 optics cause the nominally square field of view of the UVIS detector to be projected onto the sky as a skewed rhombus, $162''$ on a side, with an angle of 86.1 degrees between the sides. This distortion affects both the photometric accuracy and astrometric precision of the UVIS images. For a thorough discussion of WFC3's geometric distortion, we refer the reader to Chapter 4.

Figure 1.3: This UVIS Aperture Diagram illustrates the regions imaged by the UVIS detector (represented by blue fill), the fiducial points of the full-detector apertures (UVIS, UVIS1, UVIS2, and UVIS-CENTER), and the outlines of the $2K \times 2K$, $1K \times 1K$, and 512×512 subarray apertures. Also indicated are the positions of the four readout amplifiers (A, B, C, and D). The POSition TARGET (POSTARG) coordinate system for the UVIS-CENTER aperture, with its origin at that aperture's fiducial point, is illustrated. Although the POSTARG coordinate systems for the other apertures are not illustrated, they are oriented the same, but have origins at each aperture's fiducial point. ($U2 = -V2$ and $U3 = -V3$).



1.3 The IR Channel

The IR detector employs a 1024×1024 Teledyne (formerly Rockwell Scientific) low noise, high quantum efficiency (QE) HgCdTe detector array with $\sim 0.13''$ pixels, covering a nominal $136 \times 123''$ field of view. Only the central 1014×1014 pixels are useful for imaging. The outer rind, 5 pixels in width, contains light-insensitive pixels that are used as reference. The HgCdTe array is actively cooled by a six-stage TEC that maintains the detector at a nominal operating temperature of 145 K. The spectral response of the IR detector is optimized for imaging at near-IR wavelengths from approximately 800 to 1700 nm.

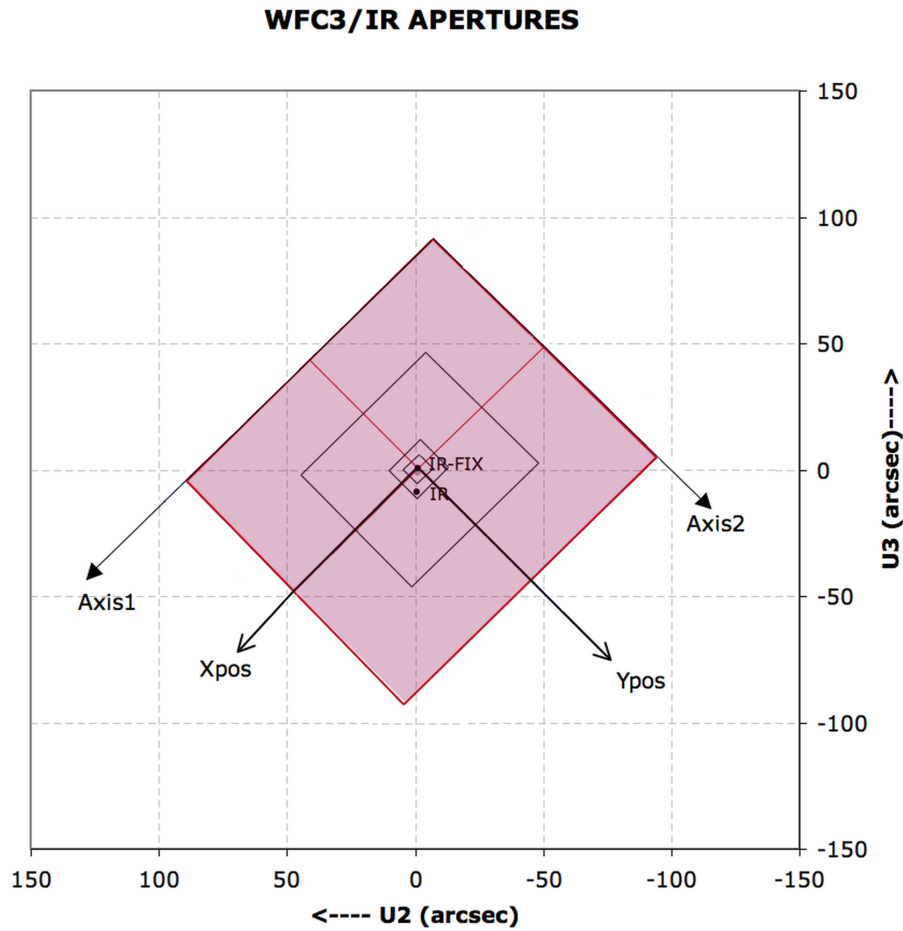
IR detectors allow accumulated signal to be read out non-destructively multiple times without affecting other pixels (a mode referred to as MULTIACCUM). This capability can be exploited to significantly reduce the effective read-out noise, enable recovery of sources that saturate during the integration time, and enable recovery of pixels affected by cosmic rays (CR), since CRs may be recognized and removed between adjacent reads.

The WFC3-IR detector is immune to the charge bleeding exhibited by CCDs at high signal levels; however, saturation can still be a problem because pixels subject to the highest signal levels show higher dark current rates ("image persistence") in subsequent exposures (see [Chapter 8](#)). IR detectors do not show long-term on-orbit charge transfer efficiency (CTE) degradation, because they do not employ the charge-transfer mechanism used in CCDs. However, they are intrinsically non-linear, although at low and intermediate count levels, the departure from linearity is quite modest and can be well calibrated.

The IR channel has a single filter wheel housing 17 spectral elements covering the near-IR wavelengths: 15 filters and 2 grisms. An 18th slot contains an opaque aluminum blocker (called a Blank). For IR observations, the requested element is simply rotated into the light beam. The IR channel operates only in MULTIACCUM mode. The WFC3 IR channel does not have a mechanical shutter, thus when the channel is not in use, the Blank is moved into the beam to block light from entering the detector, and the detector itself is continuously reset at the fastest possible rate (i.e. read out as fast as possible) for the full frame (about 2.93 seconds.)

[Figure 1.4](#) shows a schematic of the IR channel aperture projected onto the sky with respect to the U2/U3 reference frame. (For definitions of the coordinate systems in the figure, please refer to [Section 6.4.3](#) of the *WFC3 Instrument Handbook*) The IR focal plane is tilted 22 degrees with respect to the incoming beam, thus the field of view as projected onto the sky is rectangular, with an aspect ratio of ~ 0.90 . This distortion affects both the photometric accuracy and astrometric precision of the IR images. For a thorough discussion of WFC3 geometric distortion, we refer the reader to [Chapter 4](#).

Figure 1.4: The IR Aperture Diagram illustrates the fiducial points of the full-detector apertures (IR and IR-FIX), and the outlines of the concentric subarray apertures (512 × 512, 256 × 256, 128 × 128, and 64 × 64). The IR-FIX aperture has its fiducial point at the geometric center of the IR detector and the IR aperture at an optimum point near center of IR detector. Although the POSition TARGET (POSTARG) coordinate systems for the other apertures are not illustrated, they are oriented the same, but have origins at each aperture's fiducial point. ($U_2 = -V_2$ and $U_3 = -V_3$).



Chapter 2: WFC3 Data Structure

Chapter Contents

- [2.1 Types of WFC3 Files](#)
- [2.2 WFC3 File Structure](#)
- [2.3 Data Storage Requirements](#)
- [2.4 Headers and Keywords](#)

2.1 Types of WFC3 Files

- [2.1.1 Data Files and Suffixes](#)
- [2.1.2 Auxiliary Data Files](#)
- [2.1.3 Hubble Advanced Products](#)

Science data taken in orbit by WFC3 are received from the Space Telescope Data Capture Facility at Goddard Space Flight Center and sent to the STScI Operational Pipeline Unified System (OPUS) pipeline, where the data are unpacked, keyword values are extracted from the telemetry stream, and the science data are reformatted and repackaged into raw (uncalibrated) FITS files by the Generic Conversion process (see [Section 1.1.1](#) of the *Introduction to the HST Data Handbooks*). All WFC3 science data products are two-dimensional images stored in Multi-Extension FITS format files. For each exposure taken with WFC3, there is one FITS file with a unique 9-character rootname followed by a 3-character suffix: rootname_XXX.fits. The rootname identifies the observation and the suffix denotes what type of file it is (see [Chapter 5](#) of the *Introduction to the HST Data Handbooks* for more details on HST file names).

WFC3 data files are given the following definitions:

- An exposure is a single multi-extension FITS file, the atomic unit of HST data.
- A dataset is a collection of files having a common rootname.
- An association is a set of exposures that can be meaningfully combined into a single product (e.g. a set of dithered exposures, or a series of CR-SPLIT exposures)
- 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.

2.1.1 Data Files and Suffixes

The suffixes used for WFC3 raw and calibrated data products are described in [Table 2.1](#) and closely mimic those used by ACS and NICMOS.

Table 2.1: WFC3 data file suffixes.

File Suffix	Description	Units
Uncalibrated		
_raw	raw data	DN
Distortion Corrected		
_drz	UVIS and IR calibrated exposure, corrected for geometric distortion	e ⁻ /s
_drc	UVIS calibrated exposure, corrected for geometric distortion and CTE	e ⁻ /s
Intermediate		

_rac_tmp	UVIS CTE corrected raw data, no other calibration	DN
_blv_tmp	overscan-trimmed UVIS exposure	DN
_blc_tmp	overscan-trimmed UVIS exposure, CTE-corrected exposure	DN
_crj_tmp	uncalibrated, CR-rejected combined	DN
_crc_tmp	uncalibrated, CR-rejected, CTE-corrected	DN
_ima	calibrated intermediate IR multiaccum image	e ⁻ /s
Calibrated		
_flt	UVIS calibrated exposure, no CTE correction	e ⁻
_flc	UVIS calibrated exposure including CTE correction	e ⁻
_flt	IR calibrated exposure	e ⁻ /s
_flt	IR calibrated exposure (SCAN mode)	e ⁻
_crj	UVIS calibrated, CR-rejected image	e ⁻
_crj	IR calibrated, CR-rejected image	e ⁻ /s
_crc	UVIS calibrated, CR-rejected, CTE-corrected image	e ⁻
Auxiliary		
_asn	association file for observation set	
_spt	telescope and WFC3 telemetry and engineering data	
_trl	trailer file with processing history and messages	
_jit	3-second average jitter table with information on pointing, jitter, and other Pointing Control System (PCS) data; an Observation Log File	
_jif	two-dimensional histogram of jitter excursions the observation, accompanying the _jit file; an Observation Log File	

The initial input files to the calibration pipeline **calwf3** are the raw files (RAW) from Generic Conversion and the association (ASN) table, if applicable, for the complete observation set.

For UVIS images, temporary files are created once bias levels are subtracted and the overscan regions are trimmed. These temporary files have suffixes which also relate whether they have been corrected for Charge Transfer Efficiency (CTE) trailing. Files which have been corrected for CTE effects have a "c" as the third letter in their suffix. For example, "_blv_tmp" files have completed the `BIASCORR` stage, but were not corrected for CTE like the "_blc_tmp" files. If the save option is specified when `calwf3` is called on a "_raw" image, then an intermediate "_rac" file will also be saved. This contains the original data with just the CTE correction applied.

The final output files will be renamed with the "_flt" or "_flc" suffixes after the standard calibrations (flat fielding, dark subtraction, etc.) are complete. The "_blv_tmp" files serve as input for cosmic ray rejection, if required. For UVIS `CR-SPLIT` and `REPEAT-OBS` exposures, a temporary CR-combined image is created (cr*_tmp) and then renamed with either the "_crj" or "_crc" suffix once basic calibrations of that image are complete. With the addition of the CTE correction, which is the very first step, the pipeline (as of version 3.3) follows two distinct lines of processing, one with the CTE correction applied and one without. Users running raw data through `calwf3`, which has `PCTECORR` set to `PERFORM`, should expect to see both sets of calibrated files when the tasks are complete.

For the IR detector, an intermediate MultiAccum (ima) file is the result after all calibrations are applied (dark subtraction, linearity correction, flat fielding, etc.) to an individual readout of an IR exposure. A final step in `calwf3` processing of IR exposures produces a combined, CR-cleaned, image from the individual readouts, which is stored in an "_flt" product file.

The final products of the `calwf3` pipeline are not corrected for geometric distortion. `AstroDrizzle` is used to correct all WFC3 images for geometric distortion, whether they are taken as single exposures or as part of an association. `AstroDrizzle` is distributed as part of the `Drizzlepac` package, used for aligning and combining all HST images. For `CR-SPLIT` and `REPEAT-OBS`, `AstroDrizzle` supersedes the `calwf3` cosmic-ray rejection processing and uses the individual "_flt" or "_flc" files directly as input, performing cosmic-ray rejection in the process of producing the final drizzled image from multiple exposures (see [Table 2.2](#)). This has significant advantages in cases where small numbers of `CR-SPLIT` images were obtained at a small number of different dither positions, because `AstroDrizzle` will use all the information from all the input files to produce the best cosmic-ray rejection. The resulting drizzled images should generally be useful for science, although subsequent reprocessing off-line may be desirable in some cases to optimize the data for specific scientific applications.

Table 2.2: The `calwf3` and `AstroDrizzle` input and output products.

UVIS									
CALWF3						AstroDrizzle			
		Output	CTE Corrected?	CR Rejected?	Distortion Corrected?	Input	Output	CTE Corrected?	CR Corrected?
Single	raw	flt	No	No	No	flt	drz	No	No
		flc	Yes	No	No	flc	drc	Yes	No
CR-SPLIT	raw, asn	crj	No	Yes	No	flt, asn	drz	No	Yes
		crc	Yes	Yes	No	flc, asn	drc	Yes	Yes

RPT-OBS	raw, asn	crj	No	Yes	No	flt, asn	drz	No	Yes
		crc	Yes	Yes	No	flc, asn	drc	Yes	Yes
Dither Pattern	raw, asn	flt	No	No	No	flt, asn	drz	No	Yes
		flc	Yes	No	No	flc, asn	drc	Yes	Yes
IR									
CALWF3						AstroDrizzle			
Image Type	Input	Output	CTE Corrected?	CR Rejected?	Distortion Corrected?	Input	Output	CTE Corrected?	CR Corrected?
Single	raw	ima	N/A	No	No	flt	drz	N/A	Yes
		flt	N/A	Yes	No				
RPT-OBS	raw, asn	crj	N/A	Yes	No	flt, asn	drz	N/A	Yes
Dither Pattern	raw, asn	flt	N/A	Yes	No	flt, asn	drz	N/A	Yes

✔ For the UVIS channel, when there are multiple exposures, AstroDrizzle supersedes the calwf3 cosmic-ray rejection processing. It uses the "_flt" and "_flc" files produced by calwf3 as input, and performs cosmic-ray rejection in the process of producing the final distortion-corrected drizzled image.

For further information on drizzle, please refer to the DrizzlePac Handbook for http://www.stsci.edu/files/live/sites/www/files/home/scientific-community/software/drizzlepac/_documents/drizzlepac-handbook.pdf.

2.1.2 Auxiliary Data Files

Association Tables (asn)

Association tables are useful for keeping track of the complex set of relationships that can exist between exposures taken with WFC3, especially with REPEAT-OBS, CR-SPLIT, and dithered exposures. Images taken at a given dither position may be additionally CR-SPLIT into multiple exposures (e.g., UVIS observations). In these cases, associations are built to describe how each exposure relates to the desired final product. As a result, WFC3 association tables can be used to create one or more science products from the input exposures. The relationships defined in the association tables determine how far through the calibration pipeline the exposures are processed and when the calibrated exposures get combined into sub-products for further calibration.

An association file has a single extension that is a binary FITS table. That table has three columns:

- MEMNAME, the member name
- MEMTYPE, the role which that member plays in the association
- MEMPRSNT, a boolean value for whether the member is present or not.

The different MEMTYPE values are summarized in [Table 2.3](#).

Table 2.3: Exposure types in WFC3 associations. The suffix "n" is appended to the MEMTYPE to denote 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 2.4](#). 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, lxxxxxECQ and lxxxxxEGQ, are given the type EXP-CR1. The sub-product lxxxxx011 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 **calwf3**, the PROD-CRn files will have a MEMPRSNT value of "no". The sample association table in [Table 2.4](#) shows the values of MEMPRSNT prior to **calwf3** processing.

Table 2.4: Sample association table lxxxxx010_asn.

MEMNAME	MEMTYPE	MEMPRSNT
IxxxxxECQ	EXP-CR1	Yes
IxxxxxEGQ	EXP-CR1	Yes
Ixxxxx011	PROD-CR1	No
IxxxxxEMQ	EXP-CR2	Yes
IxxxxxEOQ	EXP-CR2	Yes
Ixxxxx012	PROD-CR2	No
Ixxxxx010	PROD-DTH	No

Table 2.5: Extensions in UVIS _flt or _flc files after standard processing with calwf3 and AstroDrizzle.

Extension Number	Extension Name	Extension Description	Extension Type	Extension Dimensions	Data Format	Imset /Chip /Axis
0	Primary	Global Header	PrimaryHDU	()		
1	[SCI,1]	Science Image	Image	4096 × 2051	Float32	Imset 1 = UVIS2 (CHIP2)
2	[ERR,1]	Error Array	Image	4096 × 2051	Float32	
3	[DQ,1]	Data Quality Array	Image	4096 × 2051	Int16	
4	[SCI,2]	Science Image	Image	4096 × 2051	Float32	Imset 2 = UVIS1 (CHIP1)
5	[ERR,2]	Error Array	Image	4096 × 2051	Float32	
6	[DQ,2]	Data Quality Array	Image	4096 × 2051	Int16	
7	[D2IMARR, 1]	Filter-independent distortion correction (lithography)	Image	64 × 32	Float32	CHIP2, X axis
8	[D2IMARR, 2]		Image	64 × 32	Float32	CHIP2, Y axis
9	[D2IMARR, 3]		Image	64 × 32	Float32	CHIP1, X axis
10	[D2IMARR, 4]		Image	64 × 32	Float32	CHIP1, Y axis
11	WCSDVARR	Filter-dependent distortion correction	Image	64 × 32	Float32	CHIP2, X axis

12	WCSDVARR		Image	64 × 32	Float32	CHIP2, Y axis
13	WCSDVARR		Image	64 × 32	Float32	CHIP1, X axis
14	WCSDVARR		Image	64 × 32	Float32	CHIP1, Y axis
15	WCSCORR	WCS change log	BINTABLE	14R × 24C		
16+	HDRLET	additional WCS solutions	HeaderletHDU	()		

Note 1: All UVIS filters except for the quad filters have a calibrated filter-dependent distortion correction and associated `NPOLFILE`. Images produced by the `calwf3` pipeline for these 52 filters have 15 extensions; images for the quad filters will have only 11 extensions (they have no `WCSDVARR`).

Note 2: With the release in Dec 2019 of [improved astrometric corrections for WFC3](#), additional "headerlet" extensions, which encapsulate WCS information, will be appended to the `flt/flc` files. The astrometric improvements primarily reduce the pointing errors present in WFC3 data (generally a few tenths of an arcsecond), but do NOT affect the distortion solution. Several astrometric solutions may be available for a given dataset, each stored in its own headerlet. Products retrieved from the MAST archive will have the best solution applied by default. For details on absolute astrometry in HST images, see [Section 4.5](#) of the [DrizzlePac Handbook](#).

In order to create a geometrically correct, drizzle-combined product, `PROD-DTH` exposures are combined only with **AstroDrizzle**, which executes after `calwf3` has finished processing all members. `PROD-RPT` and `PROD-CRJ` products are combined using `wf3rej` and all output files have the "cr" extension.

Support Files (spt)

The support files contain information about the observation and engineering data from the instrument and spacecraft that was recorded at the end of the observation. A support file can have multiple FITS image extensions within the same file. Each extension holds an integer (16-bit) image containing the data that populates the `*_spt.fits` header keyword values.

Trailer Files (trl)

Each task used by `calwf3` creates messages during processing that describe the progress of the calibration and are sent to `STDOUT`. In calibration pipelines written for other HST instruments, trailer files were created by simply redirecting the `STDOUT` to a file. Because multiple output files can be produced in a single run of `calwf3`, creating trailer files presents a unique challenge. Each task within `calwf3` must decide which trailer file should be appended with comments and automatically open, populate, and close each trailer file.

Note that **calwf3** will always overwrite information in trailer files from previous runs of **calwf3** while preserving any comments generated by Generic Conversion. This ensures that the trailer files accurately reflect the most recent processing performed. The string "CALWF3BEG" will mark the first comment added to the trailer file. If a trailer file already exists, **calwf3** will search for this string to determine where to append processing comments. If it is not found, the string will be written at 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.fits" suffix will be created. Further processing with **calwf3** will concatenate all trailer files associated with an output product into a single file. Additional messages will then be appended to this concatenated file. This duplicates some information across multiple trailer files but ensures that for any product processed within the pipeline, the trailer file will contain processing comments from all the input files.

Linking trailer files together can result in multiple occurrences of the "CALWF3BEG" string. Only the first, however, determines where **calwf3** will begin overwriting comments if an observation is reprocessed.

Jitter Files (jit/jif)

Jitter files are Observation Log Files that record pointing, jitter, and other Pointing Control System (PCS) data taken during an HST observation. The _jif.fits file is a FITS image of the jitter during the exposure; the associated FITS header contains information about file structure, observation details, modeled background light, point control system, jitter summary, orbital geometry, and problem flags and warnings. The _jit.fits file accompanies the _jif.fits file, and contains reconstructed pointing, guide star coordinates, derived jitter at the science instrument aperture, pertinent guiding-related flags, orbital data (e.g., latitude, longitude, limb angle, magnitude field values, etc.) and fine guidance sensor flags. Please refer to Section 5.1 and 5.2 of the *Hubble Space Telescope Data Handbooks* for more information.

2.1.3 Hubble Advanced Products

In late-2020, the HST data calibration and archive pipelines began producing new Hubble Advanced Products (HAP) to be distributed through MAST. These include two new types (or levels) of products for a given dataset, referred to as 'Single Visit Mosaics' (SVMs) and 'Multi Visit Mosaics' (MVMs), available for download from either the [MAST Discovery Portal](#) or the [MAST HST Search Interface](#). An example of the file naming convention for standard (HST) and advanced (HAP) products is shown in [Table 2.6](#).

SVMs comprise the data from a single HST visit which are aligned to a common astrometric reference frame and then drizzled onto the same north-up pixel grid. Compared to standard HST data products, these include improved relative alignment across different exposures (and/or filters) acquired within the same visit, enabling easy comparison of the drizzled products. When possible, sources in the HST images have been aligned directly to an external reference catalog (e.g. Gaia eDR3, GSC v2.4.2, or 2MASS) and include improved absolute astrometry in the image World Coordinate System (WCS), carried as HDRLET extensions in the FITS files (see [Table 2.5](#)). SVM data products with both relative alignment (by filter) and absolute alignment to Gaia will contain the string 'FIT_SVM_GAIAeDR3' in the 'WCSNAME' keyword in the science extension of the image header. The SVM drizzled images are used to generate point source and segment catalogs during pipeline processing. These catalogs supersede those produced by the [Hubble Legacy Archive](#) and will be the basis of the next version of the [Hubble Source Catalog](#). Availability of these SVM data products was announced in the December 2020 [MAST Newsletter article](#), and more details on the software used to compute these data products is available in the DrizzlePac documentation for [Single Visit Mosaic Processing](#).

In mid-2022, MAST began distributing a new type of Hubble Advanced Product (HAP) referred to as Multi-Visit Mosaics (MVM). These are cross-visit, cross-proposal mosaics, which combine public observations of fields observed multiple times by ACS/WFC, WFC3/UVIS, or WFC3/IR. MVM data products combine all exposures falling within a pre-defined 0.2° x 0.2° 'sky cell' for each detector+filter drizzled onto a common, pre-defined pixel grid. *Note that the MVMs are large files (up to 1TB, depending on the number of input exposures in each sky cell)*. Because they combine observations acquired over a range of dates, the MVMs may have photometric errors of several percent or systematic alignment errors when combining visits with different catalog solutions. These products are therefore recommend to be used as 'discovery images' for comparing observations in different detectors and passbands and not for precise photometry. Availability of the MVM products was announced in the May 2022 [MAST Newsletter article](#). More detail on the software used to compute these data products is available in the DrizzlePac documentation for [Multi-Visit Mosaic Processing](#).

For statistics on the alignment for each HST detector, uncertainties in the WCS solutions when aligning to different reference catalogs, and examples of SVM and MVM data products, see 'Improved Absolute Astrometry for ACS and WFC3 Data Products' ([WFC3 ISR 2022-06](#)).

Table 2.6: MAST standard data products (HST) and Hubble Advanced Products (HAP) available for a sample dataset 'idkl01020' taken in the F140W filter.

Data Level	Data Type	File Naming Convention	Example Dataset
HST	Drizzled Image (Exposure Level)	ippssoot_drz.fits	idkl01020_drz.fits
HAP-SVM	Drizzled Image (Visit Level)	hst_propid_obsetid_instr_detector_filter_ippss_drz.fits	hst_15132_0l_wfc3_ir_f140w_idkl0l_drz.fits
	Point Source Catalog	hst_propid_obsetid_instr_detector_filter_ippss_point-cat.ecsv	hst_15132_0l_wfc3_ir_f140w_idkl0l_point-cat.ecsv

	Segment Catalog	hst_propid_obsetid_instr_ detector_filter_ipppss_segment-cat. ecsv	hst_15132_0l_wfc3_ir_ f140w_idkl0l_segment-cat.ecsv
HAP- MVM	Drizzled Image (Sky cell)	hst_skycell_pPPPPxXXyYY_ instr_detector_filter_all_drz.fits	hst_skycell-p1857x08y11_ wfc3_ir_f140w_all_drz.fits

2.2 WFC3 File Structure


- [2.2.1 UVIS Channel File Structure](#)
- [2.2.2 IR Channel File Structure](#)
- [2.2.3 Contents of Individual Arrays](#)

All WFC3 science data products are two-dimensional images that are stored in Multi-Extension FITS files. All images taken during an exposure are bundled in a single FITS file, with each image stored in a separate FITS image extension (see [Section 2.2](#) of the *Introduction to the HST Data Handbooks*). The WFC3 file structure differs for UVIS and IR data, as explained in the following sections.

2.2.1 UVIS Channel File Structure

The WFC3 UVIS detector is similar in structure to the ACS WFC detector, with two chips butted together to form a complete detector array (there is a 1.2 arcsec gap between the two detectors). As shown in [Figure 2.1](#), each chip has 4096×2051 imaging pixels, with 19 rows and 30 columns of virtual overscan at the long and short inside edges respectively, and 25 columns of physical overscan on each side. As a result, full-frame raw images have a total of 4206×4140 pixels, and after overscan subtraction in the calibration process, calibrated images have a total of 4096×4102 pixels.

Figure 2.1: Schematic of a raw, full-frame WFC3 UVIS image.



The UVIS detector operates only in ACCUM mode to produce time-integrated images. The data read from the two chips are stored in separate image sets, or "imsets" (see [Section 2.2](#) of the *Introduction to the HST Data Handbooks*) within a single FITS file. Each imset contains three data arrays that are stored in three separate images extensions:

- the science image (SCI)
- the error array (ERR)
- the data quality array (DQ)

When the user retrieves their data from the archive, the system delivers calibrated data which are processed by both **calwf3** and **AstroDrizzle**. The **AstroDrizzle** step updates the header astrometric information, and adds 9 additional extensions to the final `_flt` or `_flc` files. In addition, with the release in Dec 2019 of [improved astrometric corrections for WFC3](#), additional "headerlet" extensions, which encapsulate WCS information, will be appended to the `_flt` or `_flc` files. Thus, the FITS file corresponding to a single full-frame calibrated UVIS exposure has a minimum of 16 extensions and potentially more depending on the available astrometric solutions (as summarized in [Table 2.5](#).)

The zeroth extension is the global or primary header unit, and the science, error, and data quality arrays are in extensions 1-6. As seen in [Figure 2.1](#) CHIP1 (UVIS1) is above CHIP2 (UVIS2) in y-pixel coordinates, but it is stored in imset 2 in the FITS file, shown graphically in [Figure 2.2](#). Thus, the chip-extension notation is counter-intuitive. To display the science image for UVIS1, the user must specify the second science extension "file.fits[sci,2]" or "file.fits[4]". Similarly, the data quality and error arrays for UVIS1 are specified as "file.fits[dq,2]" or "file.fits[5]" and "file.fits[err,2]" or "file.fits[6]", respectively.


The extensions 7-14 contain information about the geometric distortion. Extensions 7 to 10 are tabular data stored as imsets [`D2IMARR`] with one extension per each CCD chip's access. These imsets are filter-independent corrections for the CCD pixel-grid irregularities, an artifact of the manufacturing process imprinted on the detector itself. The lithographic-mask pattern correction is bi-linearly interpolated and used for pixel-by-pixel correction prior to correction for geometric distortion via the `IDCTAB`.

The extensions from 11 to 14 contain tabular data stored as imsets [`WCSDVARR`] with one extension per each CCD chip axis. These tabular data, which describe the fine-scale filter-dependent non-polynomial distortion corrections, are bi-linearly interpolated and used for pixel-by-pixel correction after correction for geometric distortion via the `IDCTAB`.

The 15th FITS extension, called [`WCSCORR`], contains a history of WCS changes, if the data were reprocessed with a new distortion correction reference file.

The 16th (and beyond) FITS extension, called [`HDRLET`], contains astrometric alignment information, one solution per [headerlet](#).

Figure 2.2: Format for the first 7 extensions of UVIS FITS files. These extensions are present in both uncalibrated and calibrated products. The final calibrated `_flc` or `_flt` files, that the user can retrieve from the archive, also have an additional 9 or more extensions with the astrometric information populated by AstroDrizzle.




2.2.2 IR Channel File Structure

The WFC3 IR channel uses a single 1024×1024 pixel detector. Reference (bias) pixels occupy the 5 rows and columns on each side of the detector, thus yielding bias-trimmed images with dimensions of 1014×1014 pixels, as shown in [Figure 2.3](#).

Like NICMOS, the IR channel operates only in MULTIACCUM mode, which starts an exposure by resetting all the detector pixels to their bias levels and recording those levels in an initial "zeroth" readout. This is then followed by n non-destructive readouts (n can be up to 15 and is set by the observer as parameter NSAMP in the Phase II proposal); the data associated with each readout are stored in a separate imset in the FITS file. The final FITS file will have $n+1$ imsets (one for each n plus one for the zeroth read).

Figure 2.3: Format of a raw full WFC3 IR image.



For IR data, each imset consists of five data arrays:

- the science image (SCI),
- the error array (ERR),
- the data quality array (DQ),
- the number of samples array (SAMP), and
- the integration time array (TIME).

An IR IMA (`_ima.fits`) FITS file will therefore contain: the primary header unit and N imsets, which all together form a single IR exposure. An IR FLT (`_flt.fits`) FITS file, on the other hand, contains only a single imset after the CRCORR ramp fitting step.

The primary header keyword `NSAMP` records the total number of readouts worth of data contained in the file. Note that the value of `NSAMP` keyword is increased by 1 relative to proposal parameter `NSAMP` because the keyword counts the zeroth read.

- ✔ The order of the IR imsets in the FITS file is in reverse time order. The first imset in the file contains the result of the longest integration time (the last readout of the MULTIACCUM series), the second imset contains the next-to-last readout and so on. The zeroth readout is stored last in the imset. This file organization has the advantage of placing the final readout first in the file, where it is easiest to access. This organization is shown graphically in [Figure 2.4](#)

Figure 2.4: Format for WFC3 IR data. Each read or image set (IMSET) of the (`_ima.fits`) FITS file consists of five data arrays: SCI, ERR, DQ, SAMP, and TIME. Consecutive MULTIACCUM readouts are stored in reverse chronological order, with `[SCI,1]` corresponding to the final, cumulative exposure. For more details on the FITS file structure, see [Table 2.6](#).

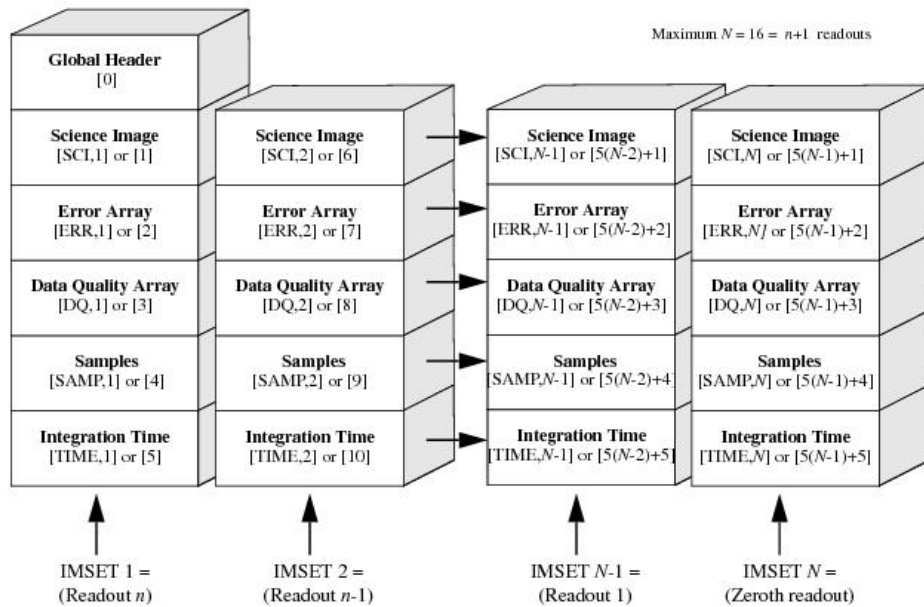


Table 2.6: File structure for a sample calibrated (*_ima*) data product, showing the IMSET, SAMPNUM, and SAMPTIME values for a full-frame IR SPARS100 exposure (*icqtbbsbxq_ima.fits*). Note that the image header keyword NSAMP reports a value of 16, but there are actually 15 science reads in the IMA file, following the 0th read (which has an exposure time of 0). While NSAMP keyword is reported in the primary header (extension 0), the SAMPNUM and SAMPTIME keywords may be found in the science header of each read, and these report the read (IMSET) number and the cumulative exposure time of each respective read. Note that SAMPNUM has the same value as NSAMP in Table 7.8 of the Instrument Handbook in the [Phase II Proposal Instructions](#).

IMSET	SAMPNUM	SAMPTIME
SCI, 16	0	0
SCI, 15	1	2.933
SCI, 14	2	102.933
SCI, 13	3	202.933
SCI, 12	4	302.933
SCI, 11	5	402.934
SCI, 10	6	502.934
SCI, 9	7	602.934
SCI, 8	8	702.935
SCI, 7	9	802.935
SCI, 6	10	902.935
SCI, 5	11	1002.936
SCI, 4	12	1102.936
SCI, 3	13	1202.936
SCI, 2	14	1302.936
SCI, 1	15	1402.937

2.2.3 Contents of Individual Arrays

The following sections explain the contents and origin of each of the individual arrays for WFC3 data products.

Science Image (SCI)

This image contains the data from the focal plane array (FPA) detectors. In raw data files, the science array is an integer (16-bit) image in units of data numbers, or DN. In calibrated data files, it is a floating-point value image in physical units of electrons (UVIS) or electrons per second (IR).

Error Array (ERR)

This is a floating-point image that contains an estimate of the statistical uncertainty associated with each corresponding science image pixel. It is expressed as a real number of signal units or signal rates (as appropriate for the units of the science image). The values for this array are calculated during calibration with the **calwf3** task, combining detector read noise, Poisson noise in the detected signal, and uncertainties from applied calibration reference data.

Data Quality Array (DQ)

This array contains 16 independent flags indicating various status and problem conditions associated with each corresponding pixel in the science image. Each flag has a true (set) or false (unset) state and is encoded as a bit in a 16-bit integer word. **Users are advised that this word should not be interpreted as a simple integer, but must be converted to base-2 and each bit interpreted as a flag.** [Table 2.7](#) lists the WFC3 data quality flags.


 In raw data files, the ERR and DQ arrays will usually have the value of zero for all pixels, unless, for the DQ array, errors are detected in the down linked data. In order to reduce data volume, if no errors exist, both ERR and DQ extensions will contain null data arrays with PIXVALUE equal to zero.

Table 2.7: WFC3 Data Quality flags

FLAG Value	Bit Setting	Data Quality Condition	
		UVIS	IR
0	0000 0000 0000 0000	OK	OK
1	0000 0000 0000 0001	Reed Solomon decoding error	Reed Solomon decoding error
2	0000 0000 0000 0010	Data replaced by fill value	Data replaced by fill value
4	0000 0000 0000 0100	Bad detector pixel	Bad detector pixel

8	0000 0000 0000 1000	(unused)	Deviant zero read (bias) value
16	0000 0000 0001 0000	Stable hot pixel	Stable hot pixel
32	0000 0000 0010 0000	Unstable pixel **	Unstable pixel
64	0000 0000 0100 0000	(Obsolete: Warm pixel)	(Obsolete: Warm pixel)
128	0000 0000 1000 0000	Bad pixel in bias	Bad reference pixel
256	0000 0001 0000 0000	Full well saturation	Full well saturation
512	0000 0010 0000 0000	Bad or uncertain flat value	Bad or uncertain flat value, including "blobs"
1024	0000 0100 0000 0000	Charge trap and sink pixels	(unused)
2048	0000 1000 0000 0000	A to D saturation	Signal in zero read
4096	0001 0000 0000 0000	Cosmic ray detected by AstroDrizzle	Cosmic ray detected by AstroDrizzle
8192	0010 0000 0000 0000	Cosmic ray detected by calwf3 during CR-SPLIT or RPT-OBS combination	Cosmic ray detected during calwf3 up the ramp fitting (flagged in _ima only)
16384	0100 0000 0000 0000	Pixel affected by ghost/crosstalk (not used)	Pixel affected by ghost/crosstalk (not used)

** Unstable hot pixel DQ flags (32) are available for observations after Nov 8, 2012, and are based on post-flashed dark reference files. For prior dates, both stable and unstable hot pixels are assigned a DQ flag value (16). See [WFC3 ISR 2018-15](#) for details.

Number of Samples Array (SAMP)

This array is present only for IR data. It is a 16-bit integer array and contains the number of samples used to derive the corresponding pixel values in the science image. For raw and intermediate data files, the sample values are set to the number of readouts that contributed to the science image. For calibrated files, the SAMP array contains the total number of valid samples used to compute the final science image pixel value, obtained by combining the data from all the readouts and rejecting cosmic ray hits and saturated pixels. Similarly, when multiple exposures (i.e., REPEAT-OBS) are combined to produce a single image, the SAMP array contains the total number of samples retained at each pixel for all the exposures.

Integration Time Array (TIME)

This array is present only for IR data. This is a floating-point array that contains the effective integration time associated with each corresponding science image pixel value. For raw and intermediate data files, the time value is the total integration time of data that contributed to the science image. For calibrated datasets, the TIME array contains the combined exposure time of the valid readouts or exposures that were used to compute the final science image pixel value, after rejection of cosmic rays and saturated pixels from the intermediate data.

✔ *In raw and intermediate data files, the SAMP and TIME arrays will each have the same value for all pixels. In order to reduce data volume, these image extensions contain null arrays, and the value of the number of samples and integration time is stored in the header keyword PIXVALUE in the SAMP and TIME extensions, respectively.*

2.3 Data Storage Requirements

The sizes of WFC3 data files (in MB) are given in [Table 2.7](#). The following assumptions were made when calculating these sizes:

- the ERR, DQ, SAMP, and TIME arrays are null in the raw files,
- the SAMP and TIME arrays are null in the ima files,
- IR images have the maximum of 16 MultiAccum readouts, and
- images are full-frame and unbinned.

Table 2.7: Sizes of WFC3 data files.

Channel	Size of FITS file			
	raw	ima	cal	drizzled
UVIS	35 MB		168 MB	217 MB
IR	34 MB	168 MB	17 MB	13 MB

The size of the drizzled image (`_drz`) is for a single image that is drizzled only for the purpose of performing geometric corrections and for combining the two UVIS chips. It assumes that the input and output image pixels are approximately the same size (drizzle scale parameter is 1.0).

2.4 Headers and Keywords

Both the primary and extension headers in a WFC3 science data file contain keywords. These keywords store a wide range of information about the observations themselves (e.g., observing mode, exposure time, detector, etc.), as well as any processing done or to be done on the data.

The primary header keywords apply to all extensions in the FITS file. The extension headers carry extension-specific keywords, i.e., information relevant only to the image in a particular extension. For example, observation parameters, calibration switches, and reference file names are stored in the primary header. The World Coordinate System and related parameters, on the other hand, are stored in the extension headers because this information can vary from one set of extensions to another.

In [Table 2.8](#) and [Table 2.9](#), we list the WFC3 keywords that appear in the primary header and SCI extension header, respectively. Columns 1 and 2 give the name of the keyword and a short description. The third column gives the format used for the value of the keyword: "L" for boolean values, "C" for characters, "R" for real values, and "I" for integer values, followed by the precision of the value. Columns 4 and 5 show to which detector the keyword applies.

[Table 2.10](#) lists the keywords that appear in the headers of the ERR, DQ, SAMP, and TIME extensions belonging to an imset in a WFC3 data file. Note that the SAMP and TIME extensions apply only to IR data.

AstroDrizzle adds additional keywords which are not listed in this handbook, and when it adds extensions does not update the `NEXTEND` keyword value. Consult the [DrizzlePac Handbook](#) for further information. Users can use the Python task `fits.info` to view all extensions in a file. The pattern keywords refer to the WFC3 dithering patterns which are discussed in greater detail in [Section 4.2](#).

Table 2.8: WFC3 Primary Header Keywords

Keyword	Description	Format	UVIS	IR
SIMPLE	data conform to FITS standard	L1	Y	Y
BITPIX	bits per data value	I2	Y	Y
NAXIS	number of data axes	I2	Y	Y
EXTEND	file may contain standard extensions	L1	Y	Y
NEXTEND	number of standard extensions	I2	Y	Y
GROUPS	image is in group format	L1	Y	Y
DATE	date this file was written (yyyy-mm-dd)	C10	Y	Y
FILENAME	name of file	C39	Y	Y

FILETYPE	type of data found in data file	C09	Y	Y
TELESCOP	telescope used to acquire data	C03	Y	Y
INSTRUME	identifier for instrument used to acquire data	C06	Y	Y
EQUINOX	equinox of celestial coordinate system	R4	Y	Y
DATA DESCRIPTION KEYWORDS				
ROOTNAME	rootname of the observation set	C34	Y	Y
IMAGETYP	type of exposure identifier	C18	Y	Y
PRIMESI	instrument designated as prime	C06	Y	Y
TARGET INFORMATION				
TARGNAME	proposer's target name	C30	Y	Y
RA_TARG	right ascension of the target (deg) (J2000)	R8	Y	Y
DEC_TARG	declination of the target (deg) (J2000)	R8	Y	Y
PROPOSAL INFORMATION				
PROPOSID	PEP proposal identifier	I4	Y	Y
LINENUM	proposal logsheet line number	C15	Y	Y
PR_INV_L	last name of principal investigator	C30	Y	Y
PR_INV_F	first name of principal investigator	C20	Y	Y
PR_INV_M	middle name initial of principal investigator	C20	Y	Y
EXPOSURE INFORMATION				
SUNANGLE	angle between sun and V1 axis	R4	Y	Y
MOONANGL	angle between moon and V1 axis	R4	Y	Y
SUN_ALT	altitude of the sun above Earth's limb	R4	Y	Y

FGSLOCK	commanded FGS lock (FINE, COARSE, GYROS, UNKNOWN)	C18	Y	Y
GYROMODE	number of gyros scheduled for observation	C1	Y	Y
REFFRAME	guide star catalog version	C8	Y	Y
MTFLAG	moving target flag; T if it is a moving target	L1	Y	Y
DATE-OBS	UT date of start of observation (yyyy-mm-dd)	C10	Y	Y
TIME-OBS	UT time of start of observation (hh:mm:ss)	C08	Y	Y
EXPSTART	exposure start time (Modified Julian Date)	R8	Y	Y
EXPEND	exposure end time (Modified Julian Date)	R8	Y	Y
EXPTIME	exposure duration (seconds)--calculated	R4	Y	Y
EXPFLAG	exposure interruption indicator	C13	Y	Y
QUALCOM1	data quality comment 1	C68	Y	Y
QUALCOM2	data quality comment 2	C68	Y	Y
QUALCOM3	data quality comment 3	C68	Y	Y
QUALITY	data quality summary	C68	Y	Y
DARKTIME	fiducial pixel dark time (secs)	R4	Y	N
POINTING INFORMATION				
PA_V3	position angle of V3-axis of HST (deg)	R4	Y	Y
TARGET OFFSETS (POSTARGS)				
POSTARG1	POSTARG in axis 1 direction	R4	Y	Y
POSTARG2	POSTARG in axis 2 direction	R4	Y	Y
DIAGNOSTIC KEYWORDS				
OPUS_VER	OPUS software system version number	C18	Y	Y

AWSYSVER	cloud infrastructure package version	C18	Y	Y
AWSDPVER	cloud docker image version	C18	Y	Y
CSYS_VER	calibration software system version id	C18	Y	Y
CAL_VER	CALWF3 code version	C24	Y	Y
PROCTIME	pipeline processing time (MJD)	R8	Y	Y
SCIENCE INSTRUMENT CONFIGURATION				
OBSTYPE	observation type - imaging or spectroscopic	C14	Y	Y
OBSMODE	operating mode	C10	Y	Y
SCLAMP	lamp status, NONE or name of lamp which is on	C14	Y	Y
NRPTEXP	number of repeat exposures in set: default 1	I2	Y	Y
SUBARRAY	data from a subarray (T) or full frame (F)	L1	Y	Y
DETECTOR	detector in use: UVIS or IR	C04	Y	Y
FILTER	element selected from filter wheel	C07	Y	Y
APERTURE	aperture name	C16	Y	Y
PROPAPER	proposed aperture name	C16	Y	Y
DIRIMAGE	direct image for grism or prism exposure	C09	Y	Y
NSAMP	number of MULTIACCUM samples	I2	N	Y
SAMP_SEQ	MULTIACCUM exposure time sequence name	C08	N	Y
SAMPZERO	sample time of the zeroth read (sec)	R4	N	Y
SUBTYPE	size/type of IR subarray	C08	N	Y
CTEIMAGE	type of Charge Transfer Image, if applicable	C04	Y	N

CTEDIR	serial or parallel	C08	Y	N
CRSPLIT	number of cosmic ray split exposures	I2	Y	N
CTE CORRECTION PARAMETERS				
CTE_NAME	name of the CTE algorithm	C23	Y	N
CTE_VER	Version number of the CTE algorithm	C23	Y	N
CTEDATE0	Date of WFC3/UVIS installation in HST (MJD)	R8	Y	N
CTEDATE1	Reference date of CTE model pinning (MJD)	R8	Y	N
PCTETLEN	Maximum length of CTE trail	R8	Y	N
PCTERNOI	Read noise amplitude	R8	Y	N
PCTENFOR	Number iterations used in CTE forward modeling	I4	Y	N
PCTENPAR	Number of iterations used in parallel transfer	I4	Y	N
PCTENSMD	read-noise mitigation algorithm	I4	Y	N
PCTETRSH	over-subtraction threshold	R8	Y	N
PCTEFAC	Scaling of CTE model (relative to CTEDATE1)	R8	Y	N
FIXROCR	Fix readout cosmic rays	I4	Y	N
SCAN KEYWORDS				
SCAN_TYP	Type of scanning path	C18	Y	Y
SCAN_WID	scan width (arcsec)	R8	Y	Y
ANG_SIDE	angle between sides of parallelogram (deg)	R8	Y	Y
DWELL_LN	dwel pts/line for scan pointing	I4	Y	Y
DWELL_TM	wait time at each dwell point (sec)	R8	Y	Y
SCAN_ANG	position angle of scan line (deg)	R8	Y	Y

SCAN_RAT	commanded rate of the line scan (arcsec/sec)	R8	Y	Y
NO_LINES	number of lines per scan	I4	Y	Y
SCAN_LEN	scan length (arcsec)	R8	Y	Y
SCAN_COR	scan coordinate frame of ref: (celestial,vehicle)	C18	Y	Y
CSMID	Channel Select Mechanism ID (UVIS /IR)	C7	Y	Y
CALIBRATION SWITCHES: PERFORM, OMIT, COMPLETE, SKIPPED				
DQICORR	data quality initialization	C08	Y	Y
BLEVCORR	subtract bias level measured from overscan, (UVIS) or reference pixels (IR)	C08	Y	Y
CRCORR	combine observations to reject cosmic rays/, identify cosmic ray hits	C08	Y	Y
DARKCORR	subtract dark image	C08	Y	Y
FLATCORR	flat-field data	C08	Y	Y
FLUXCORR	Scale the UVIS chips to get uniform flux correction	C08	Y	N
PCTECORR	Perform the pixel based CTE correction	C08	Y	N
PHOTCORR	populate photometric header keywords	C08	Y	Y
DRIZCORR	drizzle processing	C08	Y	Y
NLINCORR	correct for detector nonlinearities	C08	N	Y
RPTCORR	combine individual repeat observations	C08	N	Y
UNITCORR	convert to count rates (see also BUNIT)	C08	N	Y
ZSIGCORR	zero read signal correction	C08	N	Y
ZOFFCORR	subtract MULTIACCUM zero read	C08	N	Y

ATODCORR	correct for A to D conversion errors	C08	Y	N
BIASCORR	subtract bias image	C08	Y	N
EXPSCORR	process individual observations after cr-reject	C08	Y	N
FLSHCORR	post flash correction	C08	Y	N
SHADCORR	apply shutter shading correction	C08	Y	N
CALIBRATION REFERENCE FILES				
BPIXTAB	bad pixel table	C23	Y	Y
CCDTAB	detector calibration parameters	C23	Y	Y
OSCNTAB	table containing overscan (UVIS) or reference pixel (IR) locations	C23	Y	Y
CRREJTAB	cosmic ray rejection parameters	C23	Y	Y
DARKFILE	superdark image file name	C23	Y	Y
PFLTFILE	pixel to pixel flat-field file name	C23	Y	Y
DFLTFILE	delta flat-field file name	C23	Y	Y
LFLTFILE	low order flat	C23	Y	Y
GRAPHTAB	the HST graph table		Y	Y
COMPTAB	the HST components table		Y	Y
IDCTAB	distortion coefficients table	C23	Y	Y
DGEOFILE	distortion correction image (no longer used, kept in the header for compatibility)	C23	Y	Y
MDRIZTAB	AstroDrizzle parameter table	C18	Y	Y
NLINFIL	detector nonlinearities file	C23	N	Y
ATODTAB	analog to digital correction file	C23	Y	N
BIASFILE	superbias image file name	C23	Y	N

FLSHFILE	post flash correction file name	C23	Y	N
SHADFILE	shutter shading correction file	C23	Y	N
PCTETAB	The default parameters and reference data for the CTE correction	C18	Y	N
IMPHTTAB	The default values for the photometric keywords	C18	Y	Y
DRKCFILE	Superdark reference for CTE corrected files	C23	Y	N
BIACFILE	Superbias reference for CTE corrected files	C23	Y	N
SNKCFILE	map of sink-pixel locations	C23	Y	N
D2IMFILE	2-D non-polynomial distortion correction table	C23	Y	N
NOPOLFILE	2-D non-polynomial filter dependent distortion	C23	Y	N
COSMIC RAY REJECTION ALGORITHM PARAMETERS				
MEANEXP	reference exposure time for parameters	R4	Y	Y
SCALENSE	multiplicative scale factor applied to noise	R4	Y	Y
INITGUES	initial guess method (MIN or MED)	C03	Y	Y
SKYSUB	sky value subtracted (MODE or NONE)	C04	Y	Y
SKYSUM	sky level from the sum of all constituent images	R4	Y	Y
CRSIGMAS	statistical rejection criteria	C15	Y	Y
CRRADIUS	rejection propagation radius (pixels)	R4	Y	Y
CRTHRESH	rejection propagation threshold	R4	Y	Y
BADINPDQ	data quality flag bits to reject	I2	Y	Y
REJ_RATE	rate at which pixels are affected by cosmic rays	R4	Y	Y

CRMASK	flag CR-rejected pixels in input files (T/F)	L1	Y	Y
MDRIZSKY	sky value computed by AstroDrizzle	R4	Y	Y
PHOTOMETRY KEYWORDS (contained in primary extension for IR, in the SCI extensions for UVIS, see Table 2.9)				
PHOTMODE	observation configuration for photometric calibration	C50	Y	Y
PHOTFLAM	inverse sensitivity, ergs/cm2/Ang /electron	R8	Y	Y
PHTFLAM1	Ch1 Inv Sens, use PHOTFLAM for UV filters	R8	Y	N
PHTFLAM2	Ch2 Inv Sens, use PHOTFLAM if FLUXCORR=COMPLETE	R8	Y	N
PHTRATIO	PHTFLAM2/PHTFLAM1 ratio	R8	Y	N
PHOTFNU	inverse sensitivity, Jy* sec/electron	R8	Y	Y
PHOTZPT	ST magnitude zero point	R4	Y	Y
PHOTPLAM	pivot wavelength (Angstroms)	R4	Y	Y
PHOTBW	RMS bandwidth of filter plus detector	R4	Y	Y
POST FLASH PARAMETERS				
FLASHDUR	exposure time in seconds: 0.1 to 409.5	R4	Y	N
FLASHCUR	post flash current (zero, low, medium, high)	C07	Y	N
FLASHSTA	status: SUCCESSFUL, ABORTED, NOT PERFORMED	C16	Y	N
FLASHLVL	requested APT post flash, electrons /pixel	R4	Y	N
SHUTRPOS	shutter position: A or B	C05	Y	N
CHARGE INJECTION PARAMETERS				
CHINJECT	charge injection mode	C08	Y	N
OTFR KEYWORDS				

T_SGSTAR	OMS calculated guide star control	C18	Y	Y
PATTERN KEYWORDS				
PATTERN1	primary pattern type	C24	Y	Y
P1_SHAPE	primary pattern shape	C18	Y	Y
P1_PURPS	primary pattern purpose	C10	Y	Y
P1_NPTS	number of points in primary pattern	I2	Y	Y
P1_PSPAC	point spacing for primary pattern (arc-sec)	R4	Y	Y
P1_LSPAC	line spacing for primary pattern (arc-sec)	R4	Y	Y
P1_ANGLE	angle between sides of parallelogram patt (deg)	R4	Y	Y
P1_FRAME	coordinate frame of primary pattern	C09	Y	Y
P1_ORINT	orientation of pattern to coordinate frame (deg)	R4	Y	Y
P1_CENTR	center pattern relative to pointing (yes/no)	C03	Y	Y
PATTERN2	secondary pattern type	C24	Y	Y
P2_SHAPE	secondary pattern shape	C18	Y	Y
P2_PURPS	secondary pattern purpose	C10	Y	Y
P2_NPTS	number of points in secondary pattern	I2	Y	Y
P2_PSPAC	point spacing for secondary pattern (arc-sec)	R4	Y	Y
P2_LSPAC	line spacing for secondary pattern (arc-sec)	R4	Y	Y
P2_ANGLE	angle between sides of parallelogram patt (deg)	R4	Y	Y
P2_FRAME	coordinate frame of secondary pattern	C09	Y	Y
P2_ORINT	orientation of pattern to coordinate frame (deg)	R4	Y	Y

P2_CENTR	center pattern relative to pointing (yes/no)	C03	Y	Y
PATTSTEP	position number of this point in the pattern	I2	Y	Y
ENGINEERING PARAMETERS				
CCDAMP	CCD amplifier readout configuration	C04	Y	Y
CCDGAIN	commanded gain of CCD	R4	Y	Y
CCDOFSAB	commanded CCD bias offset for amps A\&B	I4	N	Y
CCDOFSCD	commanded CCD bias offset for amps C\&D	I4	N	Y
CCDOFSTA	commanded CCD bias offset for amplifier A	I4	Y	N
CCDOFSTB	commanded CCD bias offset for amplifier B	I4	Y	N
CCDOFSTC	commanded CCD bias offset for amplifier C	I4	Y	N
CCDOFSTD	commanded CCD bias offset for amplifier D	I4	Y	N
CALIBRATED ENGINEERING PARAMETERS				
ATODGNA	measured gain for amplifier A	R4	Y	Y
ATODGNB	measured gain for amplifier B	R4	Y	Y
ATODGNC	measured gain for amplifier C	R4	Y	Y
ATODGND	measured gain for amplifier D	R4	Y	Y
READNSEA	measured read noise for amplifier A	R4	Y	Y
READNSEB	measured read noise for amplifier B	R4	Y	Y
READNSEC	measured read noise for amplifier C	R4	Y	Y
READNSED	measured read noise for amplifier D	R4	Y	Y
BIASLEVA	bias level for amplifier A	R4	Y	Y

BIASLEVB	bias level for amplifier B	R4	Y	Y
BIASLEVC	bias level for amplifier C	R4	Y	Y
BIASLEVD	bias level for amplifier D	R4	Y	Y
ASSOCIATION KEYWORDS				
ASN_ID	unique identifier assigned to association	C10	Y	Y
ASN_TAB	name of the association table	C23	Y	Y
ASN_MTYP	role of the member in the association	C12	Y	Y

Table 2.9: WFC3 SCI extension header keywords

Keyword	Description	Format	UVIS	IR
XTENSION	extension type	C08	Y	Y
BITPIX	bits per data value	I2	Y	Y
NAXIS	number of data axes	I2	Y	Y
NAXIS1	length of first data axis	I4	Y	Y
NAXIS2	length of second data axis	I4	Y	Y
PCOUNT	number of group parameters	I2	Y	Y
GCOUNT	number of groups	I2	Y	Y
ORIGIN	FITS file originator	C30	Y	Y
DATE	date this file was written (yyyy-mm-dd)	C10	Y	Y
INHERIT	inherit the primary header	L1	Y	Y
EXTNAME	extension name	C08	Y	Y
EXTVER	extension version number	I2	Y	Y
ROOTNAME	rootname of the observation set	C34	Y	Y

EXPNAME	exposure identifier	C25	Y	Y
DATAMIN	the minimum value of the data	R8	Y	Y
DATAMAX	the maximum value of the data	R8	Y	Y
BUNIT	brightness units	C18	Y	Y
BSCALE	scale factor for array value to physical value	R8	Y	Y
BZERO	physical value for an array value of zero	R8	Y	Y
CCD CHIP IDENTIFICATION				
CCDCHIP	CCD chip (1 or 2)	I2	Y	N
WORLD COORDINATE SYSTEM AND RELATED PARAMETERS				
WCSAXES	number of World Coordinate System axes	I2	Y	Y
CRPIX1	xcoordinate of reference pixel	R8	Y	Y
CRPIX2	ycoordinate of reference pixel	R8	Y	Y
CRVAL1	first axis value at reference pixel	R8	Y	Y
CRVAL2	second axis value at reference pixel	R8	Y	Y
CTYPE1	the coordinate type for the first axis	C08	Y	Y
CTYPE2	the coordinate type for the second axis	C08	Y	Y
CD1_1	partial of first axis coordinate w.r.t. x	R8	Y	Y
CD1_2	partial of first axis coordinate w.r.t. y	R8	Y	Y
CD2_1	partial of second axis coordinate w.r.t. x	R8	Y	Y
CD2_2	partial of second axis coordinate w.r.t. y	R8	Y	Y

LTV1	offset in X to subsection start	R4	Y	Y
LTV2	offset in Y to subsection start	R4	Y	Y
LTM1_1	reciprocal of sampling rate in X	R4	Y	Y
LTM2_2	reciprocal of sampling rate in Y	R4	Y	Y
PA_APER	position Angle of reference aperture center (deg)	R8	Y	Y
VAFACOR	velocity aberration plate scale factor	R8	Y	Y
ORIENTAT	position angle of image y axis (degrees East of North)	R4	Y	Y
RA_APER	right ascension of aperture reference position	R8	Y	Y
DEC_APER	declination of aperture reference position	R8	Y	Y
REPEATED EXPOSURES INFORMATION				
NCOMBINE	number of image sets combined during CR rejection	I2	Y	Y
PHOTOMETRY KEYWORDS (contained in this extension for UVIS, in the primary extension for IR, see Table 2.8)				
PHOTMODE	observation configuration for photometric calibration	C50	Y	N
PHOTFLAM	inverse sensitivity, ergs/cm ² /Ang/electron	R8	Y	N
PHTFLAM1	PHOTFLAM on chip 1 when scales are not uniform	R8	Y	N
PHTFLAM2	PHOTFLAM on chip 2 when scales are not uniform	R8	Y	N
PHTRATIO	Ratio of PHTFLAM1/PHTFLAM2	R8	Y	N
PHOTFNU	inverse sensitivity, Jy*sec /electron	R8	Y	N
PHOTZPT	ST magnitude zero point	R4	Y	N
PHOTPLAM	pivot wavelength (Angstroms)	R4	Y	N

PHOTBW	RMS bandwidth of filter plus detector	R4	Y	N
READOUT DEFINITION PARAMETERS				
CENTERA1	subarray axis1 center point in unbinned detector pix	I4	Y	Y
CENTERA2	subarray axis2 center point in unbinned detector pix	I4	Y	Y
SIZAXIS1	subarray axis1 size in unbinned detector pixels	I4	Y	Y
SIZAXIS2	subarray axis2 size in unbinned detector pixels	I4	Y	Y
BINAXIS1	axis1 data bin size in unbinned detector pixels	I2	Y	Y
BINAXIS2	axis2 data bin size in unbinned detector pixels	I2	Y	Y
READOUT PARAMETERS				
SAMPNUM	MULTIACCUM sample number	I2	N	Y
SAMPTIME	total integration time (sec)	R4	N	Y
DELTATIM	integration time of this sample (sec)	R4	N	Y
ROUETIME	UT time of array readout (MJD)	R8	N	Y
TDFTRANS	number of TDF transitions during current sample	I4	N	Y
DATA PACKET INFORMATION				
FILLCNT	number of segments containing fill	I4	Y	Y
ERRCNT	number of segments containing errors	I4	Y	Y
PODPSFF	PODPS fill present (T/F)	L1	Y	Y
STDCFFF	science telemetry fill data present (T/F)	L1	Y	Y
STDCFFP	science telemetry fill pattern (hex)	C06	Y	Y

IMAGE STATISTICS AND DATA QUALITY FLAGS				
NGOODPIX	number of good pixels	I4	Y	Y
SDQFLAGS	serious data quality flags	I4	Y	Y
GOODMIN	minimum value of good pixels	R4	Y	Y
GOODMAX	maximum value of good pixels	R4	Y	Y
GOODMEAN	mean value of good pixels	R4	Y	Y
SNRMIN	minimum signal to noise of good pixels	R4	Y	Y
SNRMAX	maximum signal to noise of good pixels	R4	Y	Y
SNRMEAN	mean value of signal to noise of good pixels	R4	Y	Y
SOFTERRS	number of soft error pixels (DQF1)	I4	Y	Y
MEANDARK	average dark level subtracted	R4	Y	Y
MEANBLEV	average bias level subtracted	R4	Y	Y
MEANFLSH	mean number of counts in post flash exposure, DN/pixel	R4	Y	N
RADESYS	Equatorial coordinate system	C4	Y	Y
OCX10	original linear term from IDCTAB	R8	Y	Y
OCX11	original linear term from IDCTAB	R8	Y	Y
OCY10	original linear term from IDCTAB	R8	Y	Y
OCY11	original linear term from IDCTAB	R8	Y	Y
IDCSALE	pixel scale from the IDCTAB reference file	R4	Y	Y
IDCTHETA	orientation of detector's Y-axis w.r.t. V3 axis	R4	Y	Y

IDCXREF	reference pixel location in X	R4	Y	Y
IDCYREF	reference pixel location in Y	R4	Y	Y
IDCV2REF	reference pixel's V2 position	R8	Y	Y
IDCV3REF	reference pixel's V3 position	R8	Y	Y
D2IMERR1	Maximum error of D2IM correction for axis 1	R8	Y	N
D2IMDIS1	Detector to image correction type	C6	Y	N
D2IM1	Version number of WCSDVARR extension	C10	Y	N
D2IM1	Number of independent variables in D2IM function	C9	Y	N
D2IM1	Axis number of the 1st variable in a D2IM function	C10	Y	N
D2IM1	Axis number of the 2nd variable in a D2IM function	C10	Y	N
D2IMERR2	Maximum error of D2IM correction for axis 2	R8	Y	N
D2IMDIS2	Detector to image correction type	C6	Y	N
D2IM2	Version number of WCSDVARR extension	C10	Y	N
D2IM2	Number of independent variables in D2IM function	C9	Y	N
D2IM2	Axis number of the 1st variable in a D2IM function	C10	Y	N
D2IM2	Axis number of the 2nd variable in a D2IM function	C10	Y	N
D2IMEXT			Y	N
WCSNAMEO	Coordinate system title	C4	Y	Y
WCSAXESO	Number of coordinate axes	I1	Y	Y
CRPIX1O	Pixel coordinate of reference point	R4	Y	Y

CRPIX20	Pixel coordinate of reference point	R4	Y	Y
CUNIT10	Units of coordinate increment and value	C3	Y	Y
CUNIT20	Units of coordinate increment and value	C3	Y	Y
CTYPE10	Right ascension, gnomonic projection	C8	Y	Y
CTYPE20	Declination, gnomonic projection	C8	Y	Y
CRVAL10	[deg] Coordinate value at reference point	R8	Y	Y
CRVAL20	[deg] Coordinate value at reference point	R8	Y	Y
MJDREF	[d] MJD of fiducial time	R8	Y	Y
RADESYSO	Equatorial coordinate system	C4	Y	Y
CD1_10	partial of first axis coordinate w.r.t. x	R8	Y	Y
CD1_20	partial of first axis coordinate w.r.t. y	R8	Y	Y
CD2_10	partial of second axis coordinate w.r.t. x	R8	Y	Y
CD2_20	partial of second axis coordinate w.r.t. y	R8	Y	Y
A_ORDER	SIP polynomial order, axis 1, detector to sky	I1	Y	Y
A_0_2	SIP distortion coefficient	R8	Y	Y
A_0_3	SIP distortion coefficient	R8	Y	Y
A_0_4	SIP distortion coefficient	R8	Y	Y
A_1_1	SIP distortion coefficient	R8	Y	Y
A_1_2	SIP distortion coefficient	R8	Y	Y
A_1_3	SIP distortion coefficient	R8	Y	Y

A_2_0	SIP distortion coefficient	R8	Y	Y
A_2_1	SIP distortion coefficient	R8	Y	Y
A_2_2	SIP distortion coefficient	R8	Y	Y
A_3_0	SIP distortion coefficient	R8	Y	Y
A_3_1	SIP distortion coefficient	R8	Y	Y
A_4_0	SIP distortion coefficient	R8	Y	Y
B_ORDER	SIP polynomial order, axis 2, detector to sky	I1	Y	Y
B_0_2	SIP distortion coefficient	R8	Y	Y
B_0_3	SIP distortion coefficient	R8	Y	Y
B_0_4	SIP distortion coefficient	R8	Y	Y
B_1_1	SIP distortion coefficient	R8	Y	Y
B_1_2	SIP distortion coefficient	R8	Y	Y
B_1_3	SIP distortion coefficient	R8	Y	Y
B_2_0	SIP distortion coefficient	R8	Y	Y
B_2_1	SIP distortion coefficient	R8	Y	Y
B_2_2	SIP distortion coefficient	R8	Y	Y
B_3_0	SIP distortion coefficient	R8	Y	Y
B_3_1	SIP distortion coefficient	R8	Y	Y
B_4_0	SIP distortion coefficient	R8	Y	Y
WCSNAME	Coordinate system title	C30	Y	Y
MDRIZSKY	Sky value computed by AstroDrizzle	R4	Y	Y
CPERR1	Maximum error of NPOL correction for axis 1	R4	Y	N

CPDIS1	Prior distortion function type	C6	Y	N
DP1	Version number of WCSDVARR extension	C10	Y	N
DP1	Number of independent variables in CPDIS function	C9	Y	N
DP1	Axis number of the 1st variable in a CPDIS function	C10	Y	N
DP1	Axis number of the 2nd variable in a CPDIS function	C10	Y	N
CPERR2	Maximum error of NPOL correction for axis 2	R8	Y	N
CPDIS2	Prior distortion function type	C6	Y	N
DP2	Version number of WCSDVARR extension	C10	Y	N
DP2	Number of independent variables in CPDIS function	C9	Y	N
DP2	Axis number of the 1st variable in a CPDIS function	C10	Y	N
DP2	Axis number of the 2nd variable in a CPDIS function	C10	Y	N
CUNIT1	Units of coordinate increment and value	C3	Y	Y
CUNIT2	Units of coordinate increment and value	C3	Y	Y
WCSNAMEA	Coordinate system title	C20	N	Y
WCSAXESA	Number of coordinate axes	I1	N	Y
CRPIX1A	Pixel coordinate of reference point	R4	N	Y
CRPIX2A	Pixel coordinate of reference point	R4	N	Y
CUNIT1A	Units of coordinate increment and value	C3	N	Y
CUNIT2A	Units of coordinate increment and value	C3	N	Y
CTYPE1A	TAN (gnomonic) projection + SIP distortions	C12	N	Y

CTYPE2A	TAN (gnomonic) projection + SIP distortions	C12	N	Y
CRVAL1A	[deg] Coordinate value at reference point	R8	N	Y
CRVAL2A	[deg] Coordinate value at reference point	R8	N	Y
RADESYSA	Equatorial coordinate system	C4	N	Y
CD1_1A	partial of first axis coordinate w.r.t. x	R8	N	Y
CD1_2A	partial of first axis coordinate w.r.t. y	R8	N	Y
CD2_1A	partial of second axis coordinate w.r.t. x	R8	N	Y
CD2_2A	partial of second axis coordinate w.r.t. y	R8	N	Y

Table 2.10 : WFC3 extension header keywords (imset extensions 2-5)

Keyword	Description	Format	Extension Type			
			ERR	DQ	SAMP	TIME
			(UVIS and IR)		(IR Only)	
XTENSION	extension type	C08	Y	Y	Y	Y
BITPIX	number of bits per data pixel	I4	Y	Y	Y	Y
NAXIS	number of data axes	I1	Y	Y	Y	Y
NAXIS1	length of first data axis	I4	Y	Y	Y	Y
NAXIS2	length of second data axis	I4	Y	Y	Y	Y
PCOUNT	number of group parameters	I2	Y	Y	Y	Y
GCOUNT	number of groups	I2	Y	Y	Y	Y
ORIGIN	FITS file originator	C30	Y	Y	Y	Y

DATE	date this file was written (yyyy-mm-dd)	C10	Y	Y	Y	Y
INHERIT	inherit the primary header	L1	Y	Y	Y	Y
EXTNAME	extension name	C08	Y	Y	Y	Y
EXTVER	extension version number	I2	Y	Y	Y	Y
ROOTNAME	rootname of the observation set	C34	Y	Y	Y	Y
EXPNAME	exposure identifier	C25	Y	Y	Y	Y
BUNIT	brightness units	C18	Y	Y	Y	Y
WORLD COORDINATE SYSTEM AND RELATED PARAMETERS						
WCSAXES	number of World Coordinate System axes	I2	Y	Y	Y	Y
CRPIX1	x-coordinate of reference pixel	R8	Y	Y	Y	Y
CRPIX2	y-coordinate of reference pixel	R8	Y	Y	Y	Y
CRVAL1	first axis value at reference pixel	R8	Y	Y	Y	Y
CRVAL2	second axis value at reference pixel	R8	Y	Y	Y	Y
CTYPE1	the coordinate type for the first axis	C08	Y	Y	Y	Y
CTYPE2	the coordinate type for the second axis	C08	Y	Y	Y	Y
CD1_1	partial of first axis coordinate w.r.t. x	R8	Y	Y	Y	Y
CD1_2	partial of first axis coordinate w.r.t. y	R8	Y	Y	Y	Y
CD2_1	partial of second axis coordinate w.r.t. x	R8	Y	Y	Y	Y
CD2_2	partial of second axis coordinate w.r.t. y	R8	Y	Y	Y	Y
LTV1	offset in X to subsection start	R4	Y	Y	-	-
LTV2	offset in Y to subsection start	R4	Y	Y	-	-

LTM1_1	reciprocal of sampling rate in X	R4	Y	Y	-	-
LTM2_2	reciprocal of sampling rate in Y	R4	Y	Y	-	-
PA_APER	position angle of reference aperture center (deg)	R8	Y	Y	-	-
VAFACOR	velocity aberration plate scale factor	R8	Y	Y	-	-
IMAGE STATISTICS AND DATA QUALITY FLAGS						
NGOODPIX	number of good pixels	I4	Y	-	-	-
SDQFLAGS	serious data quality flags	I4	Y	-	-	-
GOODMIN	minimum value of good pixels	R4	Y	-	-	-
GOODMAX	maximum value of good pixels	R4	Y	-	-	-
GOODMEAN	mean value of good pixels	R4	Y	-	-	-

Chapter 3: WFC3 Data Calibration

Chapter Contents

- [3.1 The calwf3 Data Processing Pipeline](#)
- [3.2 UVIS Data Calibration Steps](#)
- [3.3 IR Data Calibration Steps](#)
- [3.4 Pipeline Tasks](#)
- [3.5 Manual Recalibration of WFC3 Data](#)

3.1 The calwf3 Data Processing Pipeline

- [3.1.2 Running calwf3](#)
- [3.1.3 Keyword Usage](#)
- [3.1.4 Using CRDS To Update Reference Files](#)

The main calibration pipeline which processes data taken with either the UVIS or IR detectors of the WFC3 instrument onboard the Hubble Space Telescope (HST) is called **calwf3**. It is a part of the software package (**HSTCAL**) that STScI uses for calibration of all HST science data. **calwf3** is an executable written in C that can be called via the command line interface, or run by the **wfc3tools** package in Python. The pipeline is organized into subroutines that are called by the **calwf3** executable, but can be invoked as standalone executables if users desire specialized processing of their dataset. The subroutines used for processing UVIS images are called **w3cte**, **wf3ccd**, and **wf32d**, while the main subroutine used for processing IR images is **wf3ir**. A subroutine, **wf3rej**, is shared between the UVIS and IR pipelines and is used for combining datasets obtained using **CR-SPLIT** or **REPEAT-OBS** modes. [Figure 3.1](#) shows the flow diagram for the UVIS pipeline as a whole, while [Figure 3.3](#) shows the flow diagram for the IR pipeline.

The **calwf3** pipeline processes every WFC3 image and the data are calibrated with reference files that are appropriate to the observing mode used. Currently, raw images and telemetry files are processed and calibrated, using the most up-to-date reference files, parameters, and software (see [Section 1.1.1](#) of the *Introduction to the HST Data Handbooks*). Then the data products, as well as the raw data are stored on MAST ([Mikulski Archive for Space Telescopes](#)). As of June 2024, the most recent version of **calwf3** is v.3.7.1, which was first released with **HSTCAL** v2.7.6. The **HSTCAL** software releases can be found on the "Releases" pages of the [STScI GIT repository](#).


During automatic pipeline processing on MAST, the calibrated output products generated by **calwf3** are run through **Astrodrizzle**, which is a part of the **DrizzlePac** software package. The main functions of **Astrodrizzle** are to correct calibrated images for geometric distortion, and combine dithered images into a single product.

3.1.1 Where to Find calwf3

As mentioned above, **calwf3** is part of the **HSTCAL** software, which is located in the [STScI GIT repository](#). There are three ways to install the **HSTCAL** software, two of which require using **Conda**. **HSTCAL** and its binaries are a part of the STScI software distribution **stenv**. It also independently resides on **conda-forge**. Alternatively, users can choose to install **HSTCAL** using the source code (requires compilation of C codes). The instructions for all three methods can be found under the [README](#) section of the **HSTCAL** repository.

3.1.2 Running calwf3

Before running the **calwf3** pipeline, the raw data need to be prepared for calibration by updating the calibration reference files section of the fits data header (see [Section 3.1.4](#) and [Section 3.5](#) for more details). **calwf3** can be run on a single input RAW file or an ASN table listing the members of an association.

 When processing an association of CR-SPLIT exposures, **calwf3** only retrieves the calibration switch and reference files from the first image listed in the ASN table, and applies them to the entire dataset.

calwf3 does not accept a user-defined list of input images on the command line (e.g. `*raw.fits` to process all raw files in the current directory). The **wf3ccd**, **wf32d**, **wf3cte**, **wf3ir** tasks on the other hand, will accept such user-defined input file lists, but they will not accept an association table (`_asn.fits`) as input.

Below are some basic examples illustrating how to run **calwf3** in Python and directly from a terminal. Although pure Python commands may still be used in a **PyRAF** session, please note that **IRAF/PyRAF** [have been deprecated](#) and support is no longer available. Users are strongly encouraged to switch to Python; more advanced usage examples are available in [Section 3.5](#).

Running calwf3 from a Python session using wfc3tools

```
from wfc3tools import calwf3
calwf3(filename)
```

Running multiple input files at the same time with Python

The recommended method for running **calwf3** in batch mode is to use Python and the **wfc3tools** package in the **conda-forge** distribution.

For example:

```
from wfc3tools import calwf3
from glob import glob

for fits in glob('i*_raw.fits'):
    calwf3(fits)
```

Displaying output from calwf3 in a Jupyter Notebook

When running the pipeline from a Jupyter notebook using the **calwf3** module from **wfc3tools**, informational text output from the underlying **calwf3.e** program is passed to the print function and is displayed in the user's cell output. This behavior can be customized by passing your own function as the `log_func` keyword argument to **calwf3**. As output is read from the underlying program, the **calwf3** Python wrapper will call `log_func` with the contents of each line (this provides a way to connect **calwf3** to the Python logging system by passing the `logging.debug` function or similar).

If `log_func=None` is passed, informational text output from the underlying program will be ignored, but the program's exit code will still be checked for successful completion.

Command line options for the calwf3 executable

calwf3 can also be called directly from a shell command line by supplying the executable **calwf3.e** with an input file and a list of options. This is the same executable that the **wfc3tools** package calls.

```
calwf3.e [options] input

calwf3.e -vts iaa012wdq_raw.fits

input : str
```

Name of input file

- single filename (`_raw.fits` or `_crj.fits`)

- filename of an ASN table (`_asn.fits`)

options

-d : print optional debugging statements

-q : print messages only to the trailer file

-r : print version number and date of software (e.g., Current version: 3.6.2 (May-27-2021))

-s : save temporary files

-t : print a detailed time stamp

-v : print verbose time stamps and information

-l : suppress the OpenMP parallel processing for the UVIS CTE correction

--help : print the syntax for executing this command

--version : print version number of software (e.g., 3.6.2)

--gitinfo : print git information (if it can be obtained)

Running many files at the same time from the command line

The command line executable only accepts one file at a time, but you can use operating system tools like `awk` to process many `_raw.fits` files in a directory:

```
ls *raw.fits | awk '{print "calwf3.e", $1}' | csh
```

3.1.3 Keyword Usage

`calwf3` processing is controlled by the values of keywords in the input image headers. Certain keywords, referred to as calibration switches, are used to control which calibration steps are performed. Reference file keywords indicate which reference files to use in the various calibration steps. Users who wish to perform custom reprocessing of their data may change the values of these keywords in the `_raw.fits` file primary headers and then rerun the modified file through `calwf3`.

Other keyword values record instrument and detector parameters that are used in the calibration and some record information that is computed or derived during calibration. [Table 3.1](#) provides a summary of these keywords used by `calwf3`, specifying whether they function as input or output to the task(s) listed in column 2. For a definition of each keyword see [Tables 2.8, 2.9, and 2.10](#).

Table 3.1: WFC3 keywords used with `calwf3`


Keyword	Task(s)	I/O	Header	Sample of Possible Values
APERTURE	wf3ccd, wf32d, wf3ir	Input	Primary	UVIS1, UVIS2, UVIS1-FIX, UVIS2-FIX, UVIS, IR, ...
ATODGNA, ATODGNB, ATODGNC, ATODGND	wf3ccd, wf32d, wf3ir, wf3rej	Output, Input	Primary	1.56, 2.26
BIASLEVA, BIASLEVB, BIASLEVC, BIASLEVD	wf3ccd	Output	Primary	2502.23, 2605.48
BINAXIS1, BINAXIS2	wf3ccd, wf32d, wf3ir	Input	SCI	1, 2, 3
CAL_VER	wf3ccd, wf32d, wf3rej, wf3ir	Output	Primary	2.1 (15-May-2010)
CCDAMP	wf3ccd, wf32d, wf3rej, wf3ir	Input	Primary	ABCD, AC, BD, A, B, C, D
CCDCHIP	wf3cte, wf3ccd, wf32d, wf3rej	Input	SCI	1, 2
CCDGAIN	wf3ccd, wf32d, wf3ir, wf3rej	Input	Primary	1.5, 2.5
CCDOFSTA, CCDOFSTB, CCDOFSTC, CCDOFSTD	wf3ccd, wf32d, wf3rej	Input	Primary	3
DETECTOR	wf3ccd, wf32d, wf3ir, wf3rej	Input	Primary	UVIS, IR
EXPSTART, EXPEND, EXPTIME	wf3ccd, wf32d, wf3ir, wf3rej	Input	Primary	51475.159
FILTER	wf3ccd, wf32d, wf3ir, wf3rej	Input	Primary	F606W, F160W, G102
FLASHDUR, FLASHSTA	wf3ccd, wf32d, wf3rej	Input	Primary	0.2, SUCCESSFUL
LTM1_1, LTM2_2	wf3ccd, wf3ir	Input	SCI, ERR, DQ	1.0, 0.5, 0.333
LTV1, LTV2	wf3ccd, wf32d, wf3ir	Input	SCI, ERR, DQ	0.0, 25.0, 19.0, 5.0
MEANBLEV	wf3ccd, wf3ir	Output	SCI	2554.763, 14201.36
MEANDARK	wf32d, wf3ir	Output	SCI	3.20642E-01
MEANFLSH	wf3ccd	Output	SCI	N/A
NEXTEND	wf3ccd, wf32d, wf3ir, wf3rej	Input	Primary	3, 6, 80

NSAMP	wf3ir	Input	Primary	2-16
OBSTYPE	wf3ccd, wf32d, wf3ir, wf3rej	Input	Primary	Imaging, Spectroscopic
PHOTMODE	wf32d, wf3ir	Output	SCI, Primary	“WFC3 UVIS1 F606W”
PHOTFLAM	wf32d, wf3ir	Output	SCI, Primary	Inverse sensitivity (erg cm ⁻² A ⁻¹ e ⁻¹)
PHOTFNU	wf32d, wf3ir	Output	Primary	Inverse sensitivity (Jy sec e ⁻¹)
PHOTZPT	wf32d, wf3ir	Output	SCI, Primary	ST magnitude zero point
PHOTPLAM	wf32d, wf3ir	Output	SCI, Primary	Pivot wavelength
PHOTBW	wf32d, wf3ir	Output	SCI, Primary	rms bandwidth of filter plus detector
PHTFLAM1	wf32d	Output	SCI, Primary	Chip1 Inverse Sensitivity for infinite aperture erg cm ⁻² A ⁻¹ e ⁻¹
PHTFLAM2	wf32d	Output	SCI, Primary	Chip1 Inverse Sensitivity for infinite aperture erg cm ⁻² A ⁻¹ e ⁻¹
PHTRATIO	wf32d	Output	SCI, Primary	PHTFLAM2/PHTFLAM1 ratio
READNSEA, READNSEB, READNSEC, READNSED	wf3ccd, wf32d, wf3ir, wf3rej	Output, Input	Primary	calibrated read noise for amplifier A, B, C, and D (electrons)
ROOTNAME	wf3ccd, wf32d, wf3ir, wf3rej	Input	Primary	rootname of the observation set
SAMP_SEQ	wf3ir	Input	Primary	RAPID, SPARS25,...
SAMPTIME	wf3ir	Input	SCI	Total integration time (sec)
SAMPZERO	wf3ir	Input	Primary	Sample time of MULTIACCUM zeroth read (sec)
SDQFLAGS	wf3ccd, wf32d, wf3ir	Input	SCI	Serious data quality flags considered “bad” by calwf3
SUBARRAY	wf3ccd, wf32d, wf3ir	Input	Primary	T, F
SUBTYPE	wf3ir	Input	Primary	FULLIMAG, SQ64SUB
TDFTRANS	wf3ir	Input	SCI	0, 1

IMAGE STATISTICS				
NGOODPIX	wf32d, wf3ir	Output	SCI, ERR	number of good pixels
GOODMIN , GOODMAX , GOODMEAN	wf32d, wf3ir	Output	SCI, ERR	min, max and mean values of good pixels (electrons)
SNRMIN , SNRMAX , SNRMEAN	wf32d, wf3ir	Output	SCI	min, max, and mean signal to noise of good pixels
CR-REJ PARAMETERS				
BADINPDQ	wf3rej	Output	Primary	data quality flag used for rejection
CRMASK	wf3rej	Output	Primary	T, F
CRRADIUS	wf3rej	Output	Primary	3.0
CRSIGMAS	wf3rej	Output	Primary	6.5, 5.5, 4.5
CRTHRESH	wf3rej	Output	Primary	rejection propagation threshold
EXPSTART , EXPEND , EXPTIME , TEXPTIME	wf3rej	Output	Primary	exposure start, and end times (modified Julian date)
EXPTIME , TEXPTIME	wf3rej	Output	Primary	total exposure duration (seconds) --calculated
INITGUES	wf3rej	Output	Primary	minimum, mode
MEANEXP	wf3rej	Output	Primary	Average exposure time (sec) for each image
NCOMBINE	wf3rej	Output	SCI	number of image sets combined during CR rejection
REJ_RATE	wf3rej	Output	Primary	rate at which pixels are affected by cosmic rays
SCALENSE	wf3rej	Output	Primary	Multiplicative term (in percent) for the noise model
SKYSUB	wf3rej	Output	Primary	Sky value subtracted (mode, none)
SKYSUM	wf3rej	Output	Primary	Sky level from the sum of all constituent images

3.1.4 Using CRDS To Update Reference Files

The HST Calibration Reference Data System (CRDS) is a Python library, set of command line tools and family of web servers used to manage and assign the HST and JWST calibration reference files. Users can query [CRDS](#) to identify and retrieve the best references files for their data. The following link explains how you can use this facility via a web interface, [Using CRDS to find the best reference files for your data](#), or from the [command line](#).

 *Users should ensure that the same version of all reference files and the pipeline software is used to analyze their entire dataset. This is particularly important for data taken at multiple epochs. Retrieving data from MAST will provide the most up-to-date calibration (software and reference files) available at the time of retrieval.*

For example, if the data are retrieved from MAST at different times (e.g., after the execution of each visit), it may be possible to observe systematic differences due to changes in the reference files, or more generally, to the whole data processing flow. Thus, it is recommended that the whole dataset pertaining to the given scientific investigation is retrieved simultaneously. Rarely there may be an instrument mode for which there is no calibration data. In these cases there will be a placeholder reference file in CRDS that is filled with zeros or ones, as appropriate. Such reference files are identified by having their header keyword `PEDIGREE` set to `DUMMY`. When `calwf3` encounters one of these reference files it will automatically skip the calibration step for which the file is used (e.g., `DARKCORR` will be set to the value `SKIPPED` if the `DRKFILE` is a dummy reference file).

3.2 UVIS Data Calibration Steps

- 3.2.1 Correction For Charge Transfer Efficiency
- 3.2.2 Error Array Initialization
- 3.2.3 Data Quality Array Initialization
- 3.2.4 Correct for Analog-to-Digital Conversion Errors
- 3.2.5 Overscan Bias Correction
- 3.2.6 Bias Structure Correction
- 3.2.7 Saturation Flagging and Sink Pixel Detection
- 3.2.8 Post-flash Correction
- 3.2.9 Dark Current Subtraction
- 3.2.10 Flat-field Correction
- 3.2.11 Shutter Shading Correction
- 3.2.12 Photometry Keywords Calculation
- 3.2.13 Flux normalization for UVIS1 and UVIS2
- 3.2.14 Cosmic-ray Rejection
- 3.2.15 UVIS Image Statistics Calculation

As of **calwf3** v3.7.1, the pipeline processes all UVIS data twice, once with the CTE correction applied as the first step, and a second time without the CTE correction. [Figure 3.1](#) shows a schematic representation of all the UVIS calibration steps, which are also briefly summarized below, in the order they are performed, with the corresponding calibration switch keyword in parentheses.

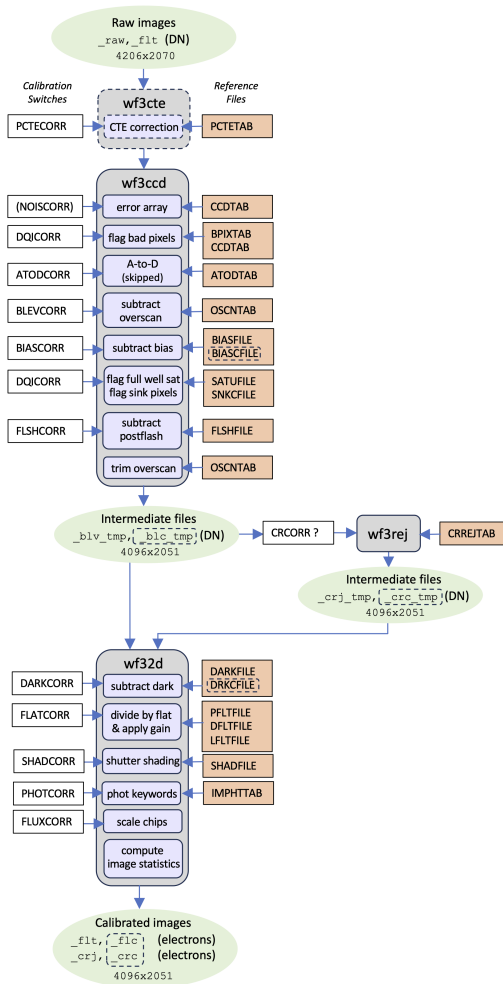
- Calculate and correct for effects due to declining charge transfer efficiency (CTE) in the image (`PCTECORR`), for UVIS only
- Calculate and record uncertainty (noise) for each pixel (`NOISCORR`) in the error, `ERR`, array and update its associated headers. Initialize the data quality, `DQ`, array of the image (`DQICORR`). Note the (`NOISCORR`) calibration switch is inaccessible to the user i.e. it is not listed explicitly in the image header and the correction is always performed by default
- Initialize the data quality (`DQ`) array of the image based on `BPIXTAB`, flag A-to-D saturation, and potentially flag full-well saturation (`DQICORR`)
- Correct for A-to-D conversion errors where necessary, currently skipped (`ATODCORR`)
- Subtract bias level determined from overscan regions (`BLEVCORR`)
- Subtract the residual bias structure from the image (`BIASCORR`)
- If applicable, employ a full-well saturation image, `SATUFILE`, to flag affected pixels.
- Detect and record `SINK` pixels in the `DQ` mask (performed if `DQICORR` is set to `PERFORM`)
- Subtract the post-flash image, if applicable (`FLSHCORR`)
- Scale and subtract the dark image (`DARKCORR`)
- Divide by flat-field image(s) and apply gain conversion (`FLATCORR`)
- Perform shutter-shading correction where necessary, currently skipped (`SHADCORR`), for UVIS only
- Populate photometric header keywords (`PHOTCORR`)
- Correct chips to use the same zero point (`FLUXCORR`)
- Calculate basic pixel statistics for the image and store the values in relevant header keywords (no switch)

The ATODCORR step is currently skipped as WFC3 ground tests did not show a bias toward the assignment of certain DN values, so this correction is not needed. The SHADCORR step, which corrects the science image for differential exposure time across the detector caused by the amount of time it takes for the shutter to open and close completely, is always set to OMIT as shutter shading was found to be an insignificant effect.

If BLEVCORR and BIASCORR are performed, a full-well 2-D saturation image (SATUFILE) is applied to flag affected pixels. This is the updated and preferred method of flagging such pixels. If the SATUFILE keyword is missing from the FITS header, then the flagging of full-well saturated pixels is done during the DQICORR step using a single constant value as the threshold.

Figure 3.1: Flow diagram of the calwf3 pipeline for UVIS data.

The name of the reference files used at each step, as well as the calibration switches that control each step are indicated. NOISCORR is a calibration switch inaccessible to the user i.e. it is not listed explicitly in the image header and the correction is always performed by calwf3.



3.2.1 Correction For Charge Transfer Efficiency

- Header switch: `PCTECORR`
- Reference file: `PCTETAB`

The charge transfer (CTE) of the UVIS detector has been declining over time as on-orbit radiation damage creates charge traps in the CCDs. During the charge transfer in the CCD readout process, these traps capture photo-electrons generated in other “upstream” pixels, thus lowering the detected charge. Faint sources, in particular, can suffer large flux losses or even be lost entirely if observations are not planned and analyzed carefully. The CTE depends on the morphology of the source, the distribution of electrons in the field of view, and the population of charge traps in the detector column between the source and the transfer register. Further details of our current understanding of the state of the WFC3/UVIS CTE are presented in Chapter 6, as well as on the [WFC3 CTE web page](#). The `PCTECORR` step aims to mitigate the flux loss incurred from CTE.

More information on this part of the pipeline can be found in Section 3.4.1, which details the functioning of the `wf3cte` routine.

3.2.2 Error Array Initialization

- Header switch: `NOISCORR` (not listed explicitly in image header, see text)
- Reference file: `CCDTAB`

In this step, the image error array is initialized and populated by the uncertainty signal measurement using the `NOISCORR` switch. The function examines the `ERR` extension of the input data to determine the state of the array. The input `_raw` image contains an empty `ERR` array. If the `ERR` array has already been expanded and contains values other than zero, then this function does nothing. Otherwise, it will initialize the `ERR` array by assigning pixel values based on a simple noise model for a CCD detector. The noise model uses the science (SCI) array and for each pixel calculates the uncertainty value σ in units of DN for each pixel. The `NOISCORR` calibration step keyword is not explicitly listed in the image header (i.e., it is not user-accessible), and it is always set to `PERFORM`.

$$\sigma_{UVIS(e-s)} = \frac{\sqrt{RN^2 + (flux * pflat + dark + PF) + \sigma_{dark}^2 + \sigma_{PF}^2 + (\sigma_{pflat} * flux)^2}}{pflat * exptime}$$

where `RN` (readnoise), `dark`, `PF` (postflash), and `flux` are in units of electrons, and `exptime` (exposure time) is in units of seconds.

The `CCD` refers to the various quadrants of the detector that are read by the four readout amplifiers (see [Chapter 5](#) of the WFC3 Instrument Handbook). The bias, gain and readnoise in the equation above for each quadrant is read from the `CCDTAB` reference file. The reference file is a FITS table that contains one row for each configuration that can be used during readout, and is uniquely identified by the list of amplifiers (replicated in the `CCDAMP` header keyword), the particular chip being read out (`CCDCHIP`), the commanded gain (`CCDGAIN`), the commanded bias offset level (`CCDOFST`) and the pixel bin size (`BINAXIS`). These commanded values are used to find the table row that matches the characteristics of the image that is being processed and reads each amplifiers characteristics, including readnoise (`READNSE`), A-to-D gain (`ATODGN`) and the mean bias level (`CCDBIAS`).

3.2.3 Data Quality Array Initialization

- Header switch: `DQICORR`

- Reference files: `BPIXTAB`

This step initializes the data quality array by reading a table of known bad pixels for the detector, as stored in the Bad Pixel reference table, `BPIXTAB`. In addition to the bad pixel types in the table, the types of bad pixels that can be flagged are:

Table 3.2: Data quality array flags for UVIS files

NAME	VALUE	DESCRIPTION
GOODPIXEL	0	OK
SOFTERR	1	Reed-Solomon decoding error
DATALOST	2	data replaced by fill value
DETECTORPROB	4	bad detector pixel or beyond aperture
DATAMASKED	8	unused
HOTPIX	16	hot pixel
CTETAILED	32	UVIS unstable pixel (post-November 2012)
WARMPPIX	64	warm pixel
BADBIAS	128	bad pixel in bias
SATPIXEL	256	full-well saturated pixel
BADFLAT	512	bad or uncertain flatfield value
TRAP	1024	UVIS charge trap and sink pixels
ATODSAT	2048	a-to-d saturated pixel
CR hit	4096	reserved for Astrodrizzle CR rejection
DATAREJECT	8192	rejected during image combination UVIS, IR CR rejection
CROSSTALK	16384	not used
RESERVED2	32768	not used

In some cases, the DQ array may already have been populated with some values to flag pixels affected by telemetry problems during downlink. Other DQ values will only be marked during further processing (such as cosmic-ray rejection). This function potentially performs two checks to determine saturation. If the newer (mid-2023) `SATUFILE` FITS keyword is missing or invalid in the input image header, the fallback is to flag full-well saturated pixels during this step using a single scalar as the threshold. This value is located in the `SATURATE` column in the CCD parameters table (`CCDTAB`). Any SCI array pixel value that is greater than the `SATURATE` value will be assigned a flag value of 256 in the DQ array. However, the newer technique is to flag the full-well saturated pixels in a sub-step after `BLEVCCORR`/`BIASCORR` using a full two-dimensional image as the threshold (see [Section 3.2.7](#)). Second, this function checks SCI array pixel values that have reached the limit of the detector's 16-bit A-to-D converters, flagging any pixel with a value > 65534 DN with the 'A-to-D saturation' DQ value of 2048. If a pixel is flagged as A-to-D saturated pixels, then they are also flagged as full-well saturated, so their DQ is actually updated (bitwise) to $256+2048$.

Full-well saturated pixels (but not A-to-D saturated) could still be used for aperture photometry, since charge keeps being created in saturated pixels though it's no longer retained in the original pixel and spills over to the neighboring ones (bleeding, see also [WFC3 ISR 2010-10](#)). Users should be aware that the photometric apertures in this case should be defined as to include all the bled charge. Users can check [Section 5.4.5](#) and [Section 5.4.6](#) of the *WFC3 Instrument Handbook* for further details on full-well and A-to-D saturation.

`DQICORR` combines the DQ flags from preprocessing, `BPIXTAB`, `SNKCFILE`, and saturation tests into a single result for the particular observation. These values are combined using a bit-wise logical OR operation 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 becomes a mask of all pixels that had some problem coming into the calibrations, so that the calibration processing steps can ignore masked pixels during processing. The `BPIXTAB` reference file maintains a record of the pixel position (x,y) and the DQ values for all known bad pixels in each CCD chip for a given time period.

Descriptions of the bad pixels and their associated values that can be flagged are indicated in [Table 3.2](#).

3.2.4 Correct for Analog-to-Digital Conversion Errors

- Header Switch: `ATODCORR`
- Reference File: `ATODTAB`

An analog-to-digital conversion correction is applied if the CCD electronic circuitry, which performs the analog-to-digital conversion, is biased toward the assignment of certain DN values. WFC3 ground test results showed that this correction is not currently needed, so the `ATODCORR` switch is currently always set to 'OMIT' so that this function is not performed.

3.2.5 Overscan Bias Correction


- Header Switch: `BLEVCCORR`
- Reference File: `OSCNTAB`

The overscan regions are used to monitor the instrument as well as provide a measure of the bias level at the time the detector was exposed (see [Figure 3.2](#)). `BLEV CORR` fits the bias level in the CCD overscan regions and subtracts it from the image data. The boundaries of the overscan regions are taken from the `OSCNTAB` reference file (Overscan Region Table). The `OSCNTAB` reference file describes the overscan regions for each chip along with the regions to be used for determining the actual bias level of the observation. Each row corresponds to a specific configuration, given by the gain, readout amplifier, charge injection, and binning.

The `OSCNTAB` columns `BIASSECTAn` and `BIASSECTBn` give the range of image columns to be used for determining the bias level in the leading and trailing regions, respectively, of the serial physical overscan regions, while columns `BIASSECTCn` and `BIASSECTDn` give the range of columns to be used for determining the bias level from the serial virtual overscan regions. The parallel virtual overscan regions are defined in the `OSCNTAB` in the `VXn` and `VYn` columns. To determine which overscan regions were actually used for measuring the bias level, check the `OSCNTAB` reference file. Users may modify the overscan region definitions in the reference table for manual calibration, but the `TRIMXn` and `TRIMYn` values must not be changed.

With these regions defined, the serial and parallel virtual overscans are analyzed to produce a two-dimensional linear fit to the bias level. The overscan level for each row of the input image is measured within the serial virtual overscan region, sigma-clipping is used to reject anomalous values (e.g., cosmic-ray hits that occur in the overscan) and a straight line is fit as a function of image line number. The same procedure is followed for the parallel overscan, resulting in a linear fit as a function of image column number. The parallel fit is computed in the form of a correction to be added to the serial fit result, in order to remove any gradient that may exist along the x-axis direction of the image. The serial fit and the parallel correction to it are then evaluated at the coordinates of each pixel and the computed bias value is subtracted from that pixel. This is done independently for each region of the image that was read out by one of the four CCD amplifiers. The mean bias value determined for each of the amplifier quadrants is recorded in the primary header keywords `BIASLEV[ABCD]` and the overall mean bias value is computed and written to the output SCI extension header as `MEANBLEV`.

Figure 3.2: Schematic view of the UVIS “raw” full frame, with overscan regions highlighted. See Section 6.7.2 of the *WFC3 Instrument Handbook* for further details.



Note: UVIS subarray images do not include virtual overscan, therefore the serial physical overscan will be used - if present - to perform the bias subtraction. If a subarray image does not include the physical overscan region of the detector, then the bias level cannot be determined. In this case a default value (CCDBIAS from the CCD parameters table) will be subtracted instead and a warning message is written to the processing trailer file.

The full bias level-subtracted image is retained in memory until the completion of all the processing steps in **wf3ccd**. The overscan regions will not be trimmed until the image is written to disk at the completion of **wf3ccd**.

3.2.6 Bias Structure Correction

- Header Switch: `BIASCORR`
- Reference File: `BIASFILE`

This step subtracts the two dimensional bias structure from the image using the superbias reference file listed in the header keyword `BIASFILE`. The dimensions of the input image are used to distinguish between full and subarray images. The `BIASFILE` has the same dimensions as a full-size science image, complete with overscan regions. The same superbias is used for full-frame and subarray input images; **calwf3** will extract the matching region from the full-size superbias file and apply it to the subarray image.

3.2.7 Saturation Flagging and Sink Pixel Detection

- Header Switch: `DQICORR`
- Reference File: `SATUFILE`, `SNKCFILE`

Full-well saturation occurs in a charge-coupled device (CCD) when accumulating charge from a central pixel begins to spill into neighboring pixels. The depths at which saturation occurs varies spatially over both UVIS CCDs. ([WFC3 ISR 2010-10](#)). Saturated pixels in UVIS images are identified by the `calwf3` calibration pipeline in the `DQICORR` step and subsequently flagged in the image data quality (DQ) array of the calibrated flat-fielded (FLT) file. If both the `BLEVCORR` and `BIASCORR` steps are performed, and the input image contains a valid FITS `SATUFILE` keyword in the primary header, then the full-well saturation image identified by the `SATUFILE` keyword will be used to define a per-pixel saturation threshold for flagging at this stage. For details on the new reference file, see [WFC3 ISR 2023-08](#) and [Section 5.6](#).

Sink pixels are a detector defect. These pixels contain a number of charge traps and under-report the number of electrons that were generated in them during an exposure. These pixels can have an impact on nearby upstream or downstream pixels, though they often only impact one or two pixels when the background is high, they can impact up to 10 pixels if the background is low (see [Section 6.7](#) for more information about sink pixels detection and flagging).

Flagging of sink pixels in the DQ extension of calibrated images is controlled with the `DQICORR` header keyword and happens after the bias correction has been performed. With `DQICORR` set to `PERFORM`, the sink pixels marked in the `SNKCFILE` reference image are flagged in the science image's DQ array. Given the `SNKCFILE` reference image, the procedure for flagging the sink pixel in the science data involves:

- Extract the MJD of the science exposure
- Examine the reference image, pixel by pixel, to find values greater than 999, which indicates that a given pixel is a sink pixel. The value of the reference file pixel corresponds to the date (in MJD format) at which this pixel started exhibiting sink pixel behavior.
- If the turn-on date of the sink pixel is after the exposure date of the science image, then we ignore the sink pixel in this exposure.
- If the turn-on date of the sink pixel is before the exposure date of the science image, then this science pixel was compromised at the time of the exposure. The corresponding pixel in the DQ extension of the science image is flagged with the "charge trap" value of 1024.
- If the pixel "downstream" of the sink pixel, along the readout direction, has a value of -1 in the `SNKCFILE` reference image, then it is also flagged with the "charge trap" value in the DQ extension. We then proceed vertically "upstream" from the sink pixel and compare each pixel in the reference file to the value of the sink pixel in the science exposure at hand. The reference file contains -at the positions of the upstream pixels- estimates of the minimum charge a sink pixel has to contain, in order not to affect the nearby upstream pixel. If the value of the sink pixel in the exposure is below the value of the upstream pixel in the reference image, we flag that upstream pixel with the "charge trap" value in the DQ extension. We continue to flag upstream pixels until the value of the pixel in the reference image is zero or until the value of the sink pixel in the exposure is greater than the value of the upstream pixel in the reference image.

[WFC3 ISR 2014-19](#) has a detailed analysis on detection of the sink pixels, while the strategy for flagging them is discussed in [WFC3 ISR 2014-22](#).

Sink pixels were originally only flagged in full frame science images, but since **calwf3** v3.4 sink pixel flagging has also been done in subarray images. **The pipeline flags pixels affected by sink pixels in the data quality extension; the pipeline does not change the science pixel values.**

3.2.8 Post-flash Correction

- Header Switch: `FLSHCORR`
- Reference File: `FLSHFILE`

WFC3 has post-flash capability to provide a means of mitigating the effects of Charge Transfer Efficiency (CTE) degradation (see [WFC3 ISR 2013-12](#)). When `FLSHCORR=PERFORM`, this routine subtracts the post-flash reference image, `FLSHFILE`, from the science image. This file has the same dimensions as a full-size science image complete with overscan regions. `FLSHFILE` is in units of electrons per second, scaled to an exposure time of 1 sec. The appropriate `FLSHFILE` is selected based on the values of the image header keywords: `USEAFTER`, `DETECTOR`, `CCDAMP`, `CCDGAIN`, `FLASHCUR`, `BINAXISi`, and `SHUTRPOS`.

The success of the post-flash operation during the exposure is first verified by checking the keyword `FLASHSTA` (`ABORTED`, `SUCCESSFUL`, `NOT PERFORMED`). **calwf3** selects the reference file which matches the science image's binning (`BINAXISi`), shutter position (`SHUTRPOS`, either A or B, i.e. the blade used reflect the postflash LED's light into the optical path to illuminate the detector), and flash current level settings (`LOW`, `MED`, `HIGH`, recorded in the `FLASHCUR` keyword). In addition, observation date is used for reference file selection as the `FLSHFILES` are now time-dependent ([WFC3 ISR 2023-01](#)); observers with post-flashed data retrieved prior to 2023 can re-retrieve their observations from MAST to obtain the updated calibration. The **calwf3** pipeline scales the reference file by the flash duration (stored in the `FLASHDUR` keyword) then subtracts it from the science image. The mean value of the scaled post-flash image is recorded in the `MEANFLSH` header keyword in units of DN. Since the blade used for post-flash typically alternates between exposures, different members of an association can have different values of `SHUTRPOS`. This does not constitute a problem for calibration as the reference files are populated separately for each exposure.

3.2.9 Dark Current Subtraction

- Header Switch: `DARKCORR`
- Reference File: `DARKFILE`, `DRKCFILE`

The dark current correction step subtracts the estimate of the dark current from the science image. The reference file listed under the `DARKFILE` header keyword is used in the non-CTE corrected UVIS pipeline branch, `DRKCFILE` is instead used in the CTE-corrected branch. The dark image (in units of electrons/sec) is multiplied by the exposure time, and subtracted from the input image.

The mean dark value is computed from the scaled dark image and used to update the `MEANDARK` keyword in the SCI image header. The dark reference file is updated frequently to allow the tracking of hot pixels over time. The time by which the dark reference file gets multiplied is simply the exposure time in the science image; it does not include the idle time since the last flushing of the chip or the readout time. Any dark accumulation during readout time is included automatically in the `BIASFILE`.


The reference file for dark subtraction, `DARKFILE` (`DRKCFILE`), is selected based on the values of the keywords `DETECTOR`, `CCDAMP`, `CHINJECT`, and `BINAXISi` in the image header. The dark correction is applied after the overscan regions are trimmed from the input science image. In a similar fashion to the `BIASCORR` step, `calwf3` requires the binning factors of the `DARKFILE` (`DRKCFILE`) and science image to match. Sub-array UVIS science images use the same reference file as a full-sized `DARKFILE` (`DRKCFILE`), `calwf3` simply extracts the appropriate region from the reference file and applies it to the sub-array input image.

3.2.10 Flat-field Correction

- Header Switch: `FLATCORR`
- Reference Files: `PFLTFILE`, `LFLTFILE`, `DFLTFILE`

This routine corrects for pixel-to-pixel and large-scale sensitivity variations across the detector by dividing the overscan-trimmed and dark-subtracted science image by a flat-field image.

Because of geometric distortion effects, the area of the sky seen by different pixels is not constant and therefore observations of a constant surface brightness object will have counts per pixel that vary over the detector, even if every pixel were to have the same intrinsic sensitivity. In order to produce images that appear uniform for uniform illumination, the same counts per pixel variation across the field is left in place in the flat-field images, so that when a science image is divided by the flat it makes an implicit correction for the distortion effects on photometry. A consequence of this procedure is that two point-source objects of equal brightness will not have the same total counts after the flat-fielding step, thus point source photometry requires the application of a [pixel area map \(PAM\) correction](#).


 ***The flat-field correction generates images that appear uniform for uniform illumination. As a result, aperture photometry extracted from a flat-fielded image (`flt`, `flc`) must be multiplied by the [effective pixel area map](#).***

The PAM correction is automatically included in pipeline processing by Astrodrizzle, which uses the geometric distortion solution to correct all pixels to equal areas. Photometry is therefore correct for both point and extended sources in drizzled (`drz`, `drc`) images.

Up to three separate flat-field reference files can be used: the pixel-to-pixel flat-field file (`PFLTFIELD`), the low-order flat-field file (`LFLTFIELD`), and the delta flat-field file (`DFLTFIELD`). The most recent information about the various UVIS flats currently in the pipeline can be found in [ISR 2016-05](#) and [ISR 2016-04](#). The `PFLTFIELD` is a pixel-to-pixel flat-field correction file containing the small-scale flat-field variations. Unlike the other flat fields, the `PFLTFIELD` is always used in the calibration pipeline. The `LFLTFIELD` is a low-order flat that corrects for any large-scale sensitivity variations across each detector. This file can be stored as a binned image, which is then expanded when being applied by `calwf3`. Finally, the `DFLTFIELD` is a delta-flat containing any needed changes to the small-scale `PFLTFIELD`. If the `LFLTFIELD` and `DFLTFIELD` are not specified in the SCI header, only the `PFLTFIELD` is used for the flat-field correction. If two or more reference files are specified, they are multiplied together to form a combined flat-field correction image. Currently the standard UVIS `calwf3` image processing uses a single flat field image, indicated by the `PFLTFIELD` keyword. This image, however, does not only contain pixel-to-pixel variations but includes low order polynomial corrections as well. All flat-field reference images must have detector, amplifier, filter, and binning modes that match the observation. A UVIS sub-array science image uses the same reference file as a full-size image; `calwf3` extracts the appropriate region from the reference file and applies it to the subarray input image.

Gain Correction

After the flat fielding step, the image is divided by the gain to convert into units of electrons.

 The units of some calibration reference files may be different than the units of the calibrated image at the stage when they are applied by `calwf3`. For example, `SATUFILE` and `SINKCFIELD` are in units of electrons, `FLSHFIELD` and `DARKFIELD` are in units of electrons/sec, while the image prior to `FLATCORR` is still in units of DN. Before applying these reference files, `calwf3` will use the gain to convert them to units of DN to ensure they are consistent with the input data.

3.2.11 Shutter Shading Correction

- Header Switch: `SHADCORR`
- Reference Files: `SHADFIELD`

This step corrects the science image for differential exposure time across the detector caused by the amount of time it takes for the shutter to completely open and close, which is a potentially significant effect only for images with very short exposure times (less than ~5 seconds). Pixels are corrected based on the exposure time using the relation:

$$corrected = uncorrected \times \frac{EXPTIME}{EXPTIME + SHADFIELD}$$

WFC3 tests have shown that the shutter shading effect is insignificant (< 1%) for even for the shortest allowed UVIS exposure time of 0.5 seconds ([WFC3 ISR 2007-17](#)). As a consequence, this step is set to `OMIT` in `calwf3` by default.

3.2.12 Photometry Keywords Calculation

- Header Switch: PHOTCORR
- Reference Files: IMPHTTAB

In order to extract calibrated magnitudes (or equivalently fluxes) directly from wfc3 images, a transformation to absolute flux units is necessary. Users who do not wish to use this feature should set the header keyword PHOTCORR to OMIT. However, users that intend to use the FLUXCORR step (see [Section 3.2.13](#)), must also set PHOTCORR to PERFORM as well.

The PHOTCORR step is performed using tables of precomputed values instead of calling **synphot** to compute keyword values on-the-fly as was done prior to November 2013. **calwf3** version 3.1.6 and greater uses an Image Photometry reference table (IMPHTTAB) specific to each WFC3 detector (IR or UVIS). The PHOTMODE keyword string reflects the observing configuration for the exposure (e.g., 'WFC3, UVIS1, F814W'). **calwf3** uses that PHOTMODE to retrieve the photometric values from the appropriate row in the reference table and update the science data header keywords listed below.

- PHOTFNU: the inverse sensitivity in units of: Jansky sec electron⁻¹
- PHOTFLAM: the inverse sensitivity in units of ergs cm⁻² A⁻¹ electron⁻¹
- PHOTPLAM: the bandpass pivot wavelength in Å
- PHOTBW: the bandpass RMS width in Å
- PHTFLAM1: the inverse sensitivity in units of ergs cm⁻² A⁻¹ electron⁻¹ for infinite aperture (chip1)
- PHTFLAM2: the inverse sensitivity in units of ergs cm⁻² A⁻¹ electron⁻¹ for infinite aperture (chip2)

Mathematical definitions of the above keywords, are defined on the [synphot ReadTheDocs page](#). For **calwf3** version 3.3 and beyond, the value PHOTFNU is calculated for each UVIS chip separately (see also [Section 3.2.13](#)). The SCI headers for each chip contain the PHOTFNU keyword, which is valid for its respective chip, and is calculated as:

For UVIS 1:

$$photF_{nu} = 3.33564 \times 10^4 \times PHTFLAM1 \times PHOTPLAM^2$$

For UVIS 2:

$$photF_{nu} = 3.33564 \times 10^4 \times PHTFLAM2 \times PHOTPLAM^2$$

The IMPHTTAB file format for WFC3 UVIS has the following structure.

EXT#	FITSNAME	FILENAME	EXTVE	DIMENS	BITPI
0	z7n21066i_imp	z7n21066i_imp.fits			16
1	BINTABLE	PHOTFLAM	1	5F \times 256R	
2	BINTABLE	PHOTPLAM	1	5F \times 256R	
3	BINTABLE	PHOTBW	1	5F \times 256R	
4	BINTABLE	PHTFLAM1	1	5F \times 256R	
5	BINTABLE	PHTFLAM2	1	5F \times 256R	


Each extension contains the photometry keyword information for that specific header keyword; rows are observation mode-dependent.

3.2.13 Flux normalization for UVIS1 and UVIS2

- Header Switch: `FLUXCORR`
- Reference Files: None

The `FLUXCORR` step was implemented in **calwf3** v3.3 on February 23, 2016 as part of the new UVIS chip-dependent photometric calibration ([WFC3 ISR 2016-01](#)). `FLUXCORR` multiplies the UVIS2 image (SCI, 1 in the data file) by the ratio of inverse sensitivities $PHTFLAM2/PHTFLAM1$ so that the same flux correction can be used for both chips. The ratio is stored for reference in the image header keyword `PHTRATIO`. By default, `FLUXCORR` is set to `PERFORM` during **calwf3** processing, and in this case the keyword `PHOTFLAM` will be valid for both chips after the correction is applied. If users do not wish to perform the correction, the `FLUXCORR` keyword may be set to `OMIT` and the raw data reprocessed through **calwf3**. In this case, the keywords `PHTFLAM1` and `PHTFLAM2` may be used to convert to flux units in the respective chips.

Chip-dependent flat fields (from 2016 or later) must be used with **calwf3** v3.3 onwards and should not be used with older versions of the pipeline, and vice versa, otherwise the data will be scaled incorrectly. For more details, see [Section 9.1](#).

 In order for `FLUXCORR` to work properly, the value of `PHOTCORR` must also be set to `PERFORM` since this populates the header of the data with the keywords `FLUXCORR` requires to compute the `PHTRATIO`.

3.2.14 Cosmic-ray Rejection

- Header Switch: `CRCORR`
- Reference Files: `CCREJTAB`

Associations with more than one member taken via `CR-SPLIT` or `REPEAT-OBS` will be combined using **wf3rej** (see [Section 3.4.5](#) for more details). The task uses the same statistical detection algorithm developed for ACS (**acsrej**), STIS (**ocrrj**), and WFPC2 (**crrej**), providing a well-tested and robust procedure. For all associations (including dithered observations), the DRZ and DRC products are created by **Astrodrizzle** which performs both cosmic ray detection (in addition to **wf3rej** for `CR-SPLIT` or `REPEAT-OBS` observations) and corrects for geometric distortion.

3.2.15 UVIS Image Statistics Calculation

- Header Switch: None
- Reference File: None

This routine computes several statistics for data values that are flagged as 'good' in the data quality array. These quantities are updated in the SCI image header: the minimum (`GOODMIN`), mean (`GOODMEAN`), and maximum (`GOODMAX`) values, as well as the minimum (`SNRMIN`), mean (`SNRMEAN`), and maximum (`SNRMAX`) signal-to-noise ratio (the ratio of the SCI and ERR pixel values). The number of good pixel is recorded in `NGOODPIX`. The minimum, mean, and maximum statistics are also computed for the ERR array.

3.3 IR Data Calibration Steps

- 3.3.1 Data Quality Initialization
- 3.3.2 IR Zero-Read Signal Correction
- 3.3.3 Bias Correction
- 3.3.4 IR Zero-read Image Subtraction
- 3.3.5 Error Array Initialization
- 3.3.6 Detector Non-linearity Correction
- 3.3.7 Dark Current Subtraction
- 3.3.8 Photometry Keywords
- 3.3.9 Conversion To Signal Rate
- 3.3.10 Up-the-ramp Fitting and Cosmic Ray Identification
- 3.3.11 Flat-field Correction
- 3.3.12 Image Statistics Calculation
- 3.3.13 Cosmic-ray rejection

The IR data reduction process begins with the raw IR image file. This data file contains all the non-destructive readouts from an exposure, stored in reverse time order i.e. the first extension corresponds to the last readout. Most of the calibration steps are applied independently to each readout. For example, the `DQICORR`, `NLINCORR`, and `DARKCORR` steps apply the bad pixel flags, non-linearity correction coefficients, and dark image, respectively, to each readout. On the other hand, the `CRCORR` step, which performs the up-the-ramp fit and removes the effects of cosmic ray hits, utilizes the values from all readouts of individual pixels simultaneously. Detailed descriptions of each step are provided in the following sections.

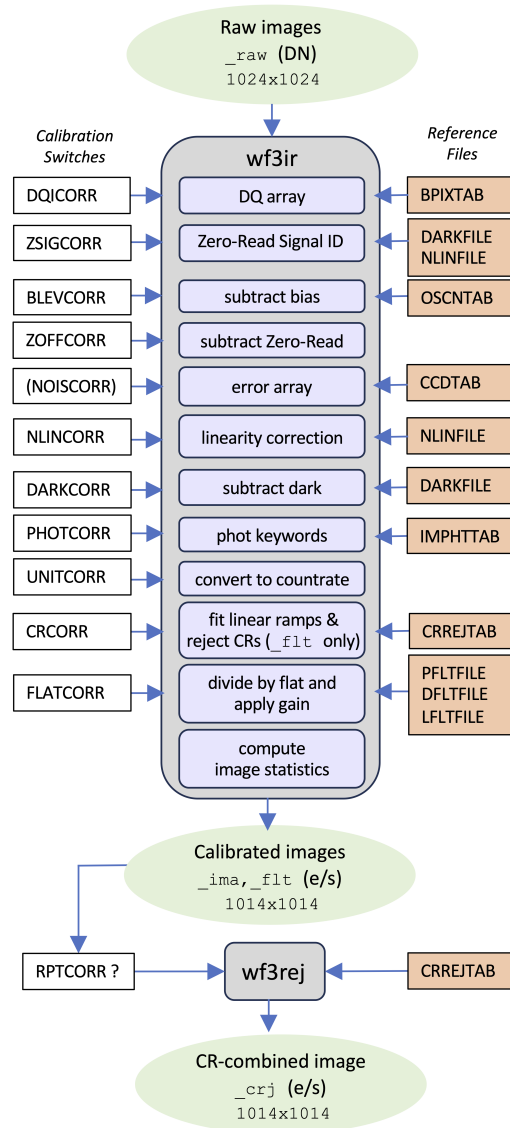
All steps up through `UNITCORR` are applied to an in-memory image stack that contains all the readouts. The `CRCORR` step produces an additional single image that gives the best-fit count rate for each pixel. The remaining steps in the process - `FLATCORR` and image statistics - are then applied to the full stack of readouts and to the single image produced by `CRCORR`.

Upon completion of the IR data calibration process, two output files are produced. The Intermediate MultiAccum (`ima`) file contains the full stack of calibrated readouts while the final calibrated image (`fit`) file is the single image produced by `CRCORR` with subsequent flat fielding applied.

Figure 3.3 shows a schematic representation of all the IR calibration steps, which are also briefly summarized below, in the order they are executed, with the corresponding calibration switch keyword in parenthesis:

- Initialize data quality, DQ, array (`DQICORR`)
- Estimate amount of signal in zeroth-read (`ZSIGCOR`)
- Subtract bias level from reference pixels (`BLEVCORR`)
- Subtract zeroth read image (`ZOFFCORR`)
- Initialize error array, ERR (`NOISCORR`)
- Correct for detector non-linear response (`NLINCORR`)
- Subtract dark current image (`DARKCORR`)
- Compute photometric keyword values for header (`PHOTCORR`)
- Convert to units of count rate (`UNITCORR`)
- Fit readout signal and identify CR hits (`CRCORR`)
- Divide by flat-field image(s) and apply gain conversion (`FLATCORR`)
- Compute image statistics (No switch)

Figure 3.3: Flow diagram for IR processing; each step, as well as the calibration switches that control each step are indicated. Similar to the UVIS, the gain is applied *after* the flat-fielding step. Note that NOISCORR is not a user-accessible keyword (see Section 3.3.5).



3.3.1 Data Quality Initialization

- Header Switch: DQICORR
- Reference File: BPIXTAB

This step populates the data quality (DQ) array in all IR readouts by reading a table of known bad pixels for the detector, stored in the 'Bad Pixel' reference table BPIXTAB. The appropriate BPIXTAB is selected based on the value of the DETECTOR and USEAFTER keywords.

The types of bad pixels that can be flagged are listed in [Table 3.3](#).

The DQ array is no longer updated to reflect any Take Data Flag (TDF) transition during the sample (see Section 5 of [STAN 12](#), Jan 2013). Other DQ values will only be marked during further processing (such as cosmic-ray rejection).

If the user wishes to update the DQ array themselves before running further processing, they should first complete the DQ initialization step, and remember that the data in the DQ extension is always in units of `UNSIGNED_INTEGER`.

Table 3.3: Data quality array flags for IR files

NAME	VALUE	DESCRIPTION
GOODPIXEL	0	OK
SOFTERR	1	Reed-Solomon decoding error
DATALOST	2	data replaced by fill value
DETECTORPROB	4	bad detector pixel
BADZERO	8	unstable IR zero-read pixel
HOTPIX	16	hot pixel
UNSTABLE	32	IR unstable pixel
WARMPIX	64	unused
BADBIAS	128	bad reference pixel value
SATPIXEL	256	full-well or a-to-d saturated pixel
BADFLAT	512	bad flat-field value
SPIKE	1024	CR spike detected during ramp fitting
ZEROSIG	2048	signal in addition to the bias in the zeroth-read
CR hit	4096	cosmic ray detected by Astrodrizzle
DATAREJECT	8192	rejected during up-the-ramp fitting
HIGH_CURVATURE	16384	not used
RESERVED2	32768	can't use

3.3.2 IR Zero-Read Signal Correction

- Header Switch: `ZSIGCORR`
- Reference Files: `DARKFILE`, `NLINFILE`

At the beginning of an IR observation, the detector pixels are reset and read out to record the bias level. An interval of approximately 2.9 seconds elapses between the time a given pixel is reset and then read. Because the IR channel does not have a shutter, signal from the field of view under observation, as well as persistent signal from previous observations, accumulates during that 2.9 second interval. When the initial (or 'zeroth') read is later subtracted from subsequent readouts, any signal in addition to the bias in the zeroth read will be subtracted.

The linearity correction and saturation checking, executed in the `NLINCORR` step (see [Section 3.3.13](#)), both depend on the absolute signal level in a pixel at the time it was read, and thus, any pixels that were exposed to bright sources during zeroth read will have a large signal and will over-subtract that signal from subsequent reads. This can lead to inaccurate linearity corrections, as well as the failure to detect saturation pixels. The `ZSIGCORR` step is used to estimate the amount of source signal in the field of view (plus any persistence signal) in the zeroth read and to supply this estimate to the `NLINCORR` step.

Such an estimate is given by the difference between the zeroth read in the linearity reference file (`NLINFILE`) and the science zero read exposure. In brief the `ZSIGCORR` module contains the following steps:

- Copy the zeroth signal image from the linearity reference file (super zero).
- Compute any subarray offsets.
- Subtract the super zero read reference image from the zero read science image (bias-subtracted zeroth read).
- Compute the noise in the subtracted zero image.
- A sigma clipping ($z_{\text{THRESH}} \times \text{noise}$) is done for the bias subtracted zeroth read and any pixel having signal above the threshold noise is flagged with a DQ value of 2048; pixels below the threshold noise are ignored. This sigma clipped bias subtracted zeroth read is provided as one of the inputs to the `NLINCORR` step.
- The linearity correction file (`NLINFILE`) also has an extension with saturation values. Pixels that are saturated in the zeroth or first reads are flagged in the DQ array and the number of saturated pixels is reported.

The `ZSIGCORR` step estimates the source signal in the science zeroth read, as outlined above, by subtracting the super zeroth read from the science zeroth read. Prior to March 2011, this module calculated an estimated source signal in the zeroth read based on the signals in the first read and zeroth read + the exposure time between them. The current method not only avoids the issue due to difference in readout times for subarrays, but also, the dark current subtraction is no longer necessary for the source signal estimate in the zeroth read (the `DARKFILE` is no longer used by this step).

Note that this technique will not work well for pixels that are exposed to very bright targets where the signal is already beginning to saturate either the zeroth or first reads. The zero read correction, `ZSIGCORR`, thus checks for saturation in the zeroth and first read arrays and flags the saturated pixels with a value of 256 in the DQ extension of all the corresponding readouts.

Pixels that are determined to have detectable signal in the zeroth read are flagged in the DQ arrays of the first readout of the output ima file with a data quality value of 2048 and carried over into the rest of the processing up to the final fit file. The `NLINFILE` is chosen based on the values of the `DETECTOR` and `USEAFTER` keywords, while the `DARKFILES` are selected based on the values of the `DETECTOR`, `CCDAMP`, `CCDGAIN`, `SAMP_SEQ`, `SUBTYPE` and `USEAFTER` keywords.

3.3.3 Bias Correction

- Header Switch: BLEVCORR
- Reference Files: OSCNTAB

The BLEVCORR step uses the reference pixels located around the perimeter of the IR detector to track and subtract the bias level during an exposure. For each raw readout array, the average signal level of the reference pixels is computed (via a sigma-clipped mean), subtracted from the array, and recorded in the MEANBLEV keyword in the SCI header of each readout.

The reference pixels located at the ends of each image row are used in this computation. Reference pixels are also located along the bottom and top of the detector, at the ends of each column, but those have been found to be less reliable and thus are not used. There are 5 reference pixels around the perimeter but the outermost pixel is ignored on each side, for a total of 8 reference pixels per row used in the BLEVCORR step.

As with the UVIS overscan correction, the boundaries of the reference pixel regions that are used in the computation are defined in the OSCNTAB reference table, in the BIASSECT* columns. The BIASSECTA[1,2] values indicate the starting and ending column numbers for the reference pixels on the left edge of the image, and the BIASSECTB[1,2] give the values for the right side of the image.

The reference pixel regions are retained in the fits file throughout the remainder of processing but are usually ignored or skipped over in the actual application of calibration algorithms. They are left in place in the calibrated data stored in the ima file at the end of processing but are trimmed from the flt image file.

The reference file for bias level correction, OSCNTAB, is selected based on the value of the DETECTOR keyword only.

3.3.4 IR Zero-read Image Subtraction

- Header Switch: ZOFFCORR
- Reference Files: None

The ZOFFCORR step subtracts the zeroth read from all readouts in the exposure, including the zeroth read itself, resulting in a zero-read image that is exactly zero in the remainder of processing. The zeroth-read image is propagated through the remaining processing steps and included in the output products, so that a complete history of error estimates and data quality (DQ) flags is preserved.

✔ **Note:** When interpreting the IR intermediate MultiAccum (ima) file, it is important to remember the file does NOT represent differences in adjacent reads, but always the difference between a given readout and the zero read. The signal rate recorded in each SCI extension of the ima file represents the average flux between that particular readout and the zero read.

3.3.5 Error Array Initialization

- Header Switch: NOISCORR (not listed explicitly in image header, see text)
- Reference Files: CCTAB

This step computes an estimate of the uncertainties associated with the raw science data based on a noise model for the detector. The `NOISCORR` keyword is not user-accessible; this step is always performed by `calwf3`. Currently, the noise model (in DN) is a simple combination of detector read noise (RN) and Poisson noise in the signal, such that:

$$\sigma_{IR(e-/s)} = \frac{\sqrt{RN^2 + (flux * pflat * dflat) + dark + \sigma_{dark}^2 + (\sigma_{pflat} * flux * dflat)^2 + (\sigma_{dflat} * flux * pflat)^2}}{pflat * dflat * exptime}$$

where RN (readnoise), dark, and flux are in units of electrons, and exptime (exposure time) is in units of seconds. The detector read noise and gain are read from the `CCDTAB` reference file and are different for each of the quadrants that are read out by the four amplifiers.

The uncertainties are then propagated, in quadrature, through each step of the calibration process, as appropriate. The `ERR` array for the final calibrated fit image is populated by the `CRCORR` step, based on the calculated uncertainty of the count rate fit to the MultiAccum samples (see [Section 3.3.13](#) for details).

The `CCDTAB` reference file used in this step is selected based on the value of the `DETECTOR` keyword only.

3.3.6 Detector Non-linearity Correction

- Header Switch: `NLINCORR`
- Reference Files: `NLINFILE`

In this step, the integrated counts in the IR science images are corrected for the non-linear response of the detector, flagging each pixel that is saturated (as defined in the saturation extension of the `NLINFILE` reference image.) The observed response of the detector can be represented by two regimes:

- At low and intermediate signal levels the detector response deviates from the incident flux in a way that is correctable using the following expression:

$$F_c = (1 + c_1 + c_2 \times F + c_3 \times F^2 + c_4 \times F^3) \times F$$

where `c1`, `c2`, `c3`, and `c4` are the correction coefficients (calculated from a large number of flat images that are median combined and then fit with a polynomial), `F` is the uncorrected flux in DN and `Fc` is the corrected flux. The current form of the correction uses a third-order polynomial, but the algorithm can handle an arbitrary number of coefficients. The number of coefficients and their uncertainties are given by the values of the `NCOEFF` and `NERR` keywords in the header of the `NLINFILE`.

- At high signal levels, as saturation sets in, the response becomes highly non-linear and is not correctable to a scientifically useful degree.

The signal in the zero read is temporarily added back to the zeroth read image of the science data before the linearity correction is applied and before the saturation is evaluated. Once the correction has been applied the zero read signal is once again removed. As with other calibration steps, this step is executed if `ZSIGCORR` is set to `PERFORM`. Saturation values for each pixel are stored in the `NODE` extension of the `NLINFIL`. Each pixel up the ramp (readout) is checked for saturation and if is saturated then the DQ array, for that readout, is updated and all subsequent readouts up the ramp are set as saturated. This is necessary since, in most cases, the readouts after a very saturated pixel can have a signal values that are lower than the saturation limit, and if not flagged will give incorrect flux values for these pixel. The `SAMP` and `TIME` arrays are not modified during this step. The `NLINFIL` reference files are selected based on the value of the `DETECTOR` keyword only.

3.3.7 Dark Current Subtraction

- Header Switch: `DARKCORR`
- Reference Files: `DARKFILE`

The `DARKCORR` step subtracts the detector dark current from the science data. The reference file listed under the `DARKFILE` header keyword is used to subtract the dark current from each sample in the ramps. Appropriate `DARKFILE` reference file are chosen based on the `DETECTOR`, `CCDAMP`, `CCDGAIN`, `SAMP_SEQ`, and `SUBTYPE` keywords of the science image. Due to potential non-linearities in some of the signal components, such as reset-related effects in the first one or two reads of an exposure, the dark current subtraction is not applied by simply scaling a generic reference dark image by the exposure time and then subtracting it. Instead, a library of dark current images are maintained, and includes darks taken in each of the available predefined `MULTIACCUM` sample sequences for both full-frame, as well as sub-array readout modes. The dark reference file is subtracted read-by-read from the stack of science image readouts so that there is an exact match in the timings and other characteristics of the dark image and the science image. The subtraction does not include the reference pixels. The `ERR` and `DQ` arrays from the reference dark file are combined with the `SCI` and `DQ` arrays from the science image, but the `SAMP` and `TIME` arrays are unchanged. The mean of the dark image is saved to the `MEANDARK` keyword in the output science image header.

3.3.8 Photometry Keywords

- Header Switch: `PHOTCORR`
- Reference Files: `IMPHTTAB`

The `PHOTCORR` step updates the image header with keywords that allow the user to convert their data from counts rates to absolute fluxes, and perform calibrated photometry. The step is performed using tables of precomputed values stored in the `IMPHTTAB` file. The appropriate entry in `IMPHTTAB` is selected according to the observing mode, whose value is stored in the image header keyword, `PHOTMODE`. The updated keywords are:

- `PHOTFLAM`: the inverse sensitivity in units of $\text{erg cm}^{-2} \text{\AA}^{-1} \text{electron}^{-1}$
- `PHOTFNU`: the inverse sensitivity in units of $\text{Jy sec electron}^{-1}$
- `PHOTZPT`: the `STMAG` zero point
- `PHOTPLAM`: the bandpass pivot wavelength in \AA
- `PHOTBW`: the band width in \AA

3.3.9 Conversion To Signal Rate

- Header Switch: UNITCORR
- Reference Files: None

This step converts the science data from a time-integrated signal to a signal rate by dividing the SCI and ERR arrays for reach readout by the TIME array. No reference file is needed when executing this step. The BUNIT keyword in the output data header is updated and reflects the appropriate data units. This step is skipped if the BUNIT value is already COUNTS/S. The flat fielding process (performed if FLATCORR is set to PERFORM), further changes the BUNIT by multiplying the image by the gain. Therefore, the final BUNIT value depends on the value of both UNITCORR and FLATCORR, as illustrated in [Table 3.4](#).

Table 3.4: Possible values of BUNIT after calibrating IR data with calwf3.

UNITCORR	FLATCORR	BUNIT AFTER CALIBRATION
OMIT	OMIT	COUNTS
OMIT	COMPLETE	ELECTRONS
COMPLETE	OMIT	COUNTS/S
COMPLETE	COMPLETE	ELECTRONS/S

3.3.10 Up-the-ramp Fitting and Cosmic Ray Identification

- Header Switch: CRCORR
- Reference Files: CCREJTAB

CRCORR combines the data from all readouts into a single image and identifies and flags pixels suspected of containing cosmic-ray (CR) hits. The method is described in detail in [Fixsen et al. \(2000\)](#).


The data from all readouts up the ramp are analyzed pixel-by-pixel, iteratively computing a linear fit to the accumulating counts-versus-exposure time relation. Pixels flagged as bad in the DQ arrays, such as when saturation occurs midway through the exposure, are rejected from the fitting process. CR hits are identified by searching for outliers from the fit results. The rejection threshold is set by the value in the CRSIGMAS column of the CR rejection parameters reference table CCREJTAB, and has a default value of 4. When a CR hit is detected, a linear fit is then performed independently for the sets of readouts before and after the hit; if a CR hit is identified to have occurred during a sample, the value measured for that sample is included in the 'after' ramp segment. Those fitting results are then again checked for outliers. This process is iterated until no new CRs are detected.

Pixel samples identified as containing a CR hit are flagged in the DQ arrays of the intermediate MultiAccum (ima) file, with DATAREJECT DQ value of 8192.

The `DATA REJECT` DQ flag is also set for all samples following the CR in order to indicate that the absolute value of the pixel is wrong after the first hit. However, this flagging smears the location of any hits which might occur after the first CR.

In addition to CRs, "negative CR hits" in the accumulated counts vs. time relation have occasionally been observed in WFC3 data. They appear as sudden "drops" in the accumulated counts vs. time plots for individual pixels (e.g. Figure 1 in [WFC3 ISR 2010-13](#)). These "negative CR hits" are also identified in the `CRCORR` step, and flagged with the SPIKE DQ flag, value = 1024. Appendix B of [WFC3 ISR 2009-40](#) gives a possible explanation for a sub-class of such events: normal cosmic rays that traverse the detector but instead of hitting the photo-sensitive HgCdTe pixel bulk, they strike other parts of the pixel (e.g. the electronics) that are sensitive to charged particles. These events are sometimes clearly associated with CRs by the physical trail visible in raw images e.g. as shown in [Figure 3.4](#). The trail has a lower signal than in the neighboring pixels (when the CR goes through the electronics) to undetectable (when the CR travels through layers that are unaffected by its passage) to higher signal than in the neighboring pixels (when charge is released in the active HgCdTe region of the pixel), see Appendix B of [WFC3 ISR 2009-40](#). Other negative spikes are sometimes observed in isolated pixels and are attributed to "burst noise", also known as popcorn noise or random telegraph signal.

Figure 3.4: Shown in the green ellipse, a negative cosmic ray trail associated with a positive one in a single read of a raw image. From bottom left we have the a trail of negative pixels, then pixels at the nominal noise level, and lastly higher signal in the top right in the ellipse.



The CR and SPIKES DQ flags are only present in the ima file and are not carried over into flt products which combine data from all readouts since the affected pixels are not included in the up-the-ramp fit.


Once all outliers have been identified, the slopes of each segment of non-flagged samples are computed via a linear fit to the counts vs. time data. This fit includes optimal weighting with individual data point uncertainties as well as contributions from the read noise and the Poisson noise for the source and the dark current. The linear fit reports the best fitting slope and its uncertainty. The final count rate value for a pixel (and its uncertainty) is determined by computing the weighted mean (and its uncertainty) of the individual sample's slopes. The result of this operation is stored in the output flt file, where the SCI array contains the final slope computed for each pixel, the ERR array contains the estimated uncertainty in the slope, the SAMP array contains the total number of non-flagged samples used to compute the slope, and the TIME array contains the total exposure time of those samples.

Pixels for which there are no unflagged samples, e.g., permanently hot or cold pixels, still have a slope computed (recorded in the SCI array of the output flt file) but they are flagged in the DQ array of the flt file. Users should therefore be always check the flt file DQ arrays to help determine whether a given SCI image pixel value is trustworthy for subsequent analysis.

The basic rule of thumb is that in order for a DQ value to propagate into the flt, it must be present in all the reads of the ima. The 8192 flag is not propagated into the flt because **calwf3** has already accounted for the effects of the CRs when performing the up-the-ramp fit. The DQ arrays in the ima files contain the complete record of when and where exactly each cosmic ray hit the detector.

A similar propagation scheme occurs for the saturation flag (DQ = 256). If, for example, a pixel is saturated in the last two reads of a ramp, then those two reads are flagged with 256 in the ima file, and **calwf3** ignores them during line-fitting. The resulting DQ value in the flt file is 0 because **calwf3** has dealt with the saturation and the effects are not present in the flt. If saturation occurs in the first read of a ramp, the SCI extension of the flt file for that pixel contains an estimate of the flux equal to the value in the input zeroth read image, but the DQ extension of the flt does not have a 256 value added to it. If the zeroth read is also saturated, the flt file still contains the same flux estimate as in the first-read saturation case, but in this case, the DQ flag 256 is propagated into the output flt DQ extension.

Pixels where **calwf3** finds 4 or more CRs up the ramp are flagged as UNSTABLE (32). These pixels are masked in the flt since that many signal jumps for a given pixel in a single ramp are an indication that likely the pixel should not be trusted. DQ values from any sample are carried through to the flt file if a pixel has no good samples.

 **Note for SCAN data:** With the release of **calwf3** v3.3, IR scan data is processed with **CRCORR** set to **OMIT** as up-the-ramp fits to scan data do not produce meaningful results. By setting **CRCORR=OMIT**, the ramp fit is not performed and instead an flt output image is produced, containing the first-minus-last read result. Note that since the SCAN fit output image is not a fit up-the-ramp, the output image units will not be a rate but instead be in counts (if **UNITCORR=OMIT**) or electrons (if **FLATCORR=COMPLETE**).


3.3.11 Flat-field Correction

- Header Switch: **FLATCORR**
- Reference Files: **PFLTFILE**, **LFLTFILE**, **DFLTFILE**

The **FLATCORR** step corrects for pixel-to-pixel and large-scale sensitivity variations across the detector by dividing the science images by one or more flat-field images. A combined flat is created within **calwf3** using up to three flat-field reference files: the pixel-to-pixel flat (**c**), the low-order flat (**LFLTFILE**), and the delta flat (**DFLTFILE**). **FLATCORR** also multiplies the science data by the detector gain, using the mean gain from all the amplifiers. Therefore the calibrated data will be in units of electrons per second (or electrons if **UNITCORR = OMIT**).

The **PFLTFILE** is a pixel-to-pixel flat-field correction file containing the small-scale flat-field variations. The **PFLTFILE** is always used in the calibration pipeline, while the other two flats are optional. The **LFLTFILE** is a low-order flat that corrects for any large-scale sensitivity variations across the detector. This file can be stored as a binned image, which is then expanded when being applied by **calwf3**. Finally, 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** is used for the flat-field correction. If two or more reference files are specified, they are read in and multiplied together to form a combined flat-field correction image.

The flat-field correction is applied to all readouts of the calibrated IR MultiAccum stack, as well as the single image produced by the `CRCORR` function. Due to geometric distortion effects, the area of the sky seen by different pixels is not constant and therefore observations of a constant surface brightness object will have counts per pixel that vary over the detector. To produce images that appear uniform with uniform illumination, the same counts per pixel variation across the field is left in place in the flat-field images, so that when a science image is divided by the flat it makes an implicit correction for the distortion effects on photometry. A consequence of this procedure is that two point-source objects of equal brightness will not have the same total counts after the flat-fielding step, thus point source photometry requires the application of a [pixel area map \(PAM\) correction](#).

 **Note:** All WFC3 observations, not just dithered images, processed with Astrodrizzle (drz files) will be corrected for geometric distortion and pixel area effects. However when using flt files to extract point-source photometry, the [pixel area map file](#) must be applied manually (calwf3 does not perform a PAM correction).

All flat-field reference images are selected based on the `DETECTOR`, `CCDAMP`, and `FILTER` used for the observation. A sub-array science image uses the same reference file(s) as a full-size image; `calwf3` extracts the appropriate region from the reference file(s) and applies it to the sub-array input image.

3.3.12 Image Statistics Calculation

- Header Switch: None
- Reference Files: None

This step computes image statistics using the “good pixels”, i.e. with DQ value equal to 0, and records them in image header keywords. The operation is performed for every readout in the calibrated MultiAccum stack (`ima`), as well as the final (`CRCORR`-produced) calibrated image (`flt`). The values computed and captured in keywords are: the minimum, mean, and maximum values (`GOODMIN`, `GOODMEAN`, `GOODMAX`, respectively), the number of good pixels (`NGOODPIX`), as well as the minimum, mean, and maximum signal-to-noise ratio (i.e. the ratio of the SCI and ERR pixel values) which are `SNRMIN`, `SNRMEAN`, `SNRMAX`, respectively. The minimum, mean, and maximum statistics are computed for the ERR arrays as well.

3.3.13 Cosmic-ray rejection

- Header Switch: `RPTCORR`
- Reference Files: `CCREJTAB`

Associations with more than one member, i.e. have been associated using `REPEAT-OBS`, are combined using `wf3rej` (see [Section 3.4.5](#) for more details). `CR-SPLIT` is not available for the IR channel. The task uses the same statistical detection algorithm developed for ACS (`acsrej`), STIS (`ocrrj`) and WFC2 (`crrej`), providing a well-tested and robust procedure. For all associations (including dithered observations), the DRZ products will be created by **Astrodrizzle**, which performs both cosmic ray detection (in addition to `wf3rej`, for `REPEAT-OBS` observations) and corrects for geometric distortion.

3.4 Pipeline Tasks

[3.4.1 wf3cte](#)
[3.4.2 wf3ccd](#)
[3.4.3 wf32d](#)
[3.4.4 wf3ir](#)
[3.4.5 wf3rej](#)

The following section contains information regarding the 5 high-level tasks that are called by the main program `calwf3`, namely `wf3cte`, `wf3ccd`, `wf32d`, `wf3ir`, `wf3rej`. These tasks can be called individually by the users from either the command line or within a Python interpreter. They generally are “umbrella” tasks that can perform multiple data processing steps. However, the user can control the processing flow of `calwf3` (and limit it to individual steps) by appropriate setting of the header keyword switches, as detailed in [Sections 3.2 \(UVIS\)](#) and [3.3 \(IR\)](#).

3.4.1 wf3cte

This routine is used to correct for Charge Transfer Efficiency (CTE) losses due to detector defects trapping electrons during readout. The `wf3cte` step corrects for Y-CTE losses which result in charge trailing in the vertical direction; as of the publication of this handbook (June 2024), there is no automated correction for the substantially smaller effect of X-CTE losses. The effect of the latter is primarily astrometric, most images will not be significantly impacted. Observers performing high-precision astrometry should consider applying the stand-alone X-CTE correction available on the [CTE tools](#) page (see [ISR 2024-07](#)).

More details on the CTE phenomenon are given in [Chapter 6](#), see also the [WFC3 CTE webpage](#). Note that while the CTE model has recently been updated ([WFC3 ISR 2021-09](#)), the interface to the code as described below remains the same.

The `PCTETAB` reference file contains extensions of calibration images and tables of parameters used during the CTE correction stage. The header of this file also contains parameters for the CTE correction algorithm, which are used by `wf3cte` to correct the data. Although not recommended, users who wish to use other settings for the CTE correction algorithm can adjust the pertinent keywords in their dataset header (see [WFC3 ISR 2016-02](#)).

`wf3cte` performs the CTE correction on raw data files. If the calibration step keyword `PCTECORR` is set to `PERFORM` then the CTE correction will be applied to the dataset. Some caveats for its use:

- CTE corrections can ONLY be performed on raw data which has not been calibrated in any way.
- Data which have `BLEVCORR`, `BIASCORR` or `DARKCORR` set to `COMPLETE` will be rejected.
- As of `calwf3` v3.4, both full frame and most subarray modes ([Table 3.5](#)) are corrected for CTE. The `calwf3` version used to process any data is stored in the header keyword `CAL_VER`. The CTE correction is unavailable for subarray images acquired with aperture `UVIS2-M1K1C-SUB` or `UVIS2-M512C-SUB` as these images, near the middle of the chip, lack physical overscan pixels. The CTE correction code requires overscan to calculate a secondary bias subtraction for the image before the CTE is measured.

Table 3.5: Sub-array modes for which CTE correction is enabled as of calwf3 version 3.4. For sub-arrays not listed here, a [workaround](#) is available.

```
UVIS1-2K2A-SUB
UVIS1-2K2B-SUB
UVIS2-2K2C-SUB
UVIS2-2K2D-SUB
UVIS2-C1K1C-SUB
UVIS2-C512C-SUB
UVIS2-C512D-SUB
UVIS1-C512A-SUB
UVIS1-C512B-SUB
UVIS1-2K4-SUB
UVIS-QUAD-SUB
UVIS2-2K4-SUB
```

The standalone call of **wf3cte** will produce a RAC fits (*_rac_tmp.fits) file by default. This contains only the CTE-corrected raw data, no other calibrations have been performed. The WFC3 CTE algorithm was updated in **calwf3/wf3cte** v3.6.0 (December 2020) where the pixel-based model is better constrained for both low and high pixel values. See the thorough discussion regarding the update in [WFC3 ISR 2021-09](#).

Displaying output from wf3cte in a Jupyter Notebook

When calling **wf3cte** from a Jupyter notebook, informational text output from the underlying **wf3cte.e** program will be passed through print as the calibration runs by default, and show up in the user's cell output. This behavior can be customized by passing your own function as the `log_func` keyword argument to **wf3cte**. As output is read from the underlying program, the **wf3cte** Python wrapper will call `log_func` with the contents of each line (this provides a way to connect **wf3cte** to the Python logging system by passing the `logging.debug` function or similar).

If `log_func=None` is passed, informational text output from the underlying program will be ignored, but the program's exit code will still be checked for successful completion.

Input Parameters for the Python interface

```
Parameters
  input : str or list
          Name of input files, such as
          - a single filename (iaa012wdq_raw.fits)
          - a Python list of filenames, a partial filename with wildcards (*raw.fits)
          - an at-file (@input)

  parallel : bool, default=True
            If True, run the code with OpenMP parallel processing turned on for the UVIS CTE
correction.

  verbose: bool, optional, default=False
          If True, print verbose time stamps.

  log_func : func(), default=print()
            If not specified, the print function is used for logging to facilitate use in the
Jupyter notebook.

Returns

  None

Usage

from wfc3tools import wf3cte
wf3cte('ib3805v0q_raw.fits', verbose=True)
```

Command Line Options for the wf3cte executable

```
wf3cte.e input [-options]
wf3cte.e ib3805v0q_raw.fits -vt1
wf3cte.e ib3805v0q_raw.fits,ib3805vlq_raw.fits -vt1

input : str or comma-separated list of string(s)
        Name of input filename or list of comma-separated input filenames
        - single filename: ippssoot_raw.fits
        - multiple filenames: ippssoot_raw.fits,ippssoot_raw.fits (Note: Do not include any blank
spaces.)

options
-r : print version number and date of software (e.g., Current version: 3.6.2 (May-27-2021)) and
exit
-t : print a detailed time stamp
-v : print verbose time stamps and information
-l : suppress the OpenMP parallel processing for the UVIS CTE correction

--help : print the syntax for executing this command and exit

--version : print version number of software (e.g., 3.6.2) and exit
--gitinfo : print git information (if it can be obtained) and exit

The output CTE corrected file is the ippssoot root of the input file with the suffix of _rac_tmp.fits (e.g.,
ib3805v0q_raw.fits -> ib3805v0q_rac_tmp.fits).
```

Basic Steps In The CTE Correction

- The reference bias image named in the `BIACFILE` header keyword is subtracted from the data.
- Parameters from the CTE parameter table, referenced in the `PCTETAB` header keyword, are read and stored.
- The data are reformatted: each quadrant is rotated such that the readout amp is located at the lower left of the array. The reoriented four quadrants are then arranged side-by-side into a

single 8412 × 2070 image (including the overscan pixels) with amps CDAB in that order. In this format, the pixels are all parallel-shifted down, then serial-shifted to the left.

- An additional bias correction is performed using the residual bias level measured for each amplifier from the steadiest pixels in the horizontal overscan, this value is then subtracted from all the pixels in each respective amp.
- The image is multiplied by the gain to convert to electrons.
- CTE correction v1.0 (**calwf3** v3.5.2 and earlier). The bias-subtracted image contained contribution from read-noise. Those read-noise “electrons”, however, are generated at the readout and are not affected by CTE. The smoothest possible image is constructed to be consistent with the original image plus pure readnoise (a Gaussian distribution with a mean of zero and a sigma of 3.25 e⁻). The CTE reconstruction algorithm then operates on the "smooth" version of the scene, determining which pre-read-out charge distribution could be CTE-trailed to result in the "smooth" scene. The difference between these two images is an estimate of how CTE likely redistributed charge in the exposure. This difference was added to the original raw, uncorrected, uncalibrated image.
- CTE correction v2.0 (**calwf3** v3.6.0 and later). While v1.0 worked well for a number of years, once the CTE correction became larger than a perturbation, read-noise was significantly amplified and required a new approach. In v2.0, pixel-to-pixel variations smaller than the background fluctuations (or 10 electrons, whichever is smaller) are essentially untouched by the reconstruction algorithm. The penalty is that the faintest sources will not be CTE-corrected and users will need to apply other strategies to correct the faintest measurements (see [WFC3 ISR 2021-09](#) for more details).
- The corrected image is now ready to continue through the rest of the pipeline. When the `DARKCORR` header keyword is set to `PERFORM`, the CTE corrected image will be dark-subtracted using the dark reference file assigned in the `DRKCFILE` header keyword.
- In the case of a subarray image, the same steps are performed as above after the image has been placed into the correct full-frame reference position since the correction is dependent on the distance of the pixels away from the read-out amplifier.

The `PCTETAB` and Algorithm Parameters

The following are new primary header keywords that will be updated in the data headers during the `wf3cte` step. They are also specified in the `PCTETAB` reference file.

KEYWORD	DESCRIPTION
CTE_NAME	name of CTE algorithm [string]
CTE_VER	version number of CTE algorithm [string]
CTEDATE0	date of WFC3/UVIS installation in HST, in modified Julian days (MJD)
CTEDATE1	reference date of CTE model pinning, in modified Julian days (MJD)
PCTETLEN	max length of CTE trail
PCTERNOI	read-noise amplitude for clipping
PCTENFOR	number of iterations used in CTE forward modeling
PCTENPAR	number of iterations used in the parallel transfer
PCTENSMD	read-noise mitigation algorithm

PCTETRSH	over-subtraction threshold
PCTEFRAC	CTE scaling fraction calculated from expstart and used in the algorithm
PCTERNOI	read-noise clipping level to use*
FIXROCR	make allowance for readout cosmic rays

*Note: The value `PCTERNOI` is the read-noise clipping level to use during processing. This value is no longer used from the `PCTETAB` file since `calwf3` v3.6.0 (December 2020). If the `PCTERNOI` keyword value in the raw science image header is non-zero, it will be used for the CTE computations. Otherwise, the value is computed on-the-fly based upon the raw image data.

The `PCTETAB` reference file has 4 extensions, two tables, and two images:

```

Filename: 5411347ei_cte.fits
No. Name          Type          Cards  Dimensions  Format
0  PRIMARY        PrimaryHDU   80         ()
1  QPROF          BinTableHDU  18         999R x 4C   ['I', 'J', 'E', '20A']
2  SCLBYCOL       BinTableHDU  22         8412R x 6C  ['I', 'E', 'E', 'E', 'E',
'20A']
3  RPROF          ImageHDU     33         (999, 100)  float32
4  CPROF          ImageHDU     33         (999, 100)  float32

```

The first extension lists the charge-trap levels. The columns are respectively the trap number, the charge-packet size it applies to (in electrons), the size of the trap (in electrons), and a description.

The second extension contains the CTE scalings as a function of column number. There are 6 columns, each with 8412 elements. The first column contains the integer column number in the readout amp-aligned large array. The other columns contain the CTE scaling appropriate for that column at the 512th, 1024th, 1536th and 2048th rows respectively. The final column provides a description.

The third extension contains the differential CTE trail profile as a function of charge level in the form of an image.

The fourth extension contains the cumulative CTE trail profile as a function of charge level, also in the form of an image.

Output Files

If the user is running the separate `wf3cte.e` step a `_rac.fits` file will be produced. This is the same as a `_raw.fits` file except the CTE correction has been applied to the data.

If the `PCTECORR` step is set to `PERFORM`:


- When the `_raw.fits` file enters `calwf3`, no intermediate `rac_tmp.fits` file will be saved unless you specify the `-s` flag, which instructs `calwf3.e` to save all intermediate files.
- The `calwf3` pipeline will produce both CTE calibrated products and non-CTE calibrated products. The CTE products have a 'c' at the end of their extension name, such as `_blc_c`, `_rac_c`, `_crc_c`, `_flc_c`, and the non-CTE calibrated products contain the familiar: `_blv`, `_crj`, `_flt`, etc.

Note: The CTE correction step uses all available CPUs (while the rest of the calibration steps do not).

3.4.2 wf3ccd

This routine performs the initial processing steps for all the WFC3 UVIS channel data. These steps are:

- `DQICORR` - initialize the data quality array with values from `BPIXTAB`, flag for A-to-D saturation, and potentially flag for full-well saturation using scalar value as the threshold (fall-back algorithm)
- `ATODCORR` - perform the a-to-d conversion correction
- `BLEVCORR` - subtract the bias level determined from the overscan regions
- `BIASCORR` - subtract the superbias image
- Flag for full-well saturation using a two-dimensional image (new algorithm)
- Detect and record SINK pixels in the DQ mask (performed if `DQICORR` is set to `PERFORM`)
- `FLSHCORR` - subtract the post-flash image

 The **wf3ccd** routine processes everything in data numbers (DN). If a calibration reference file is in units of electrons, **wf3ccd** will divide it by the gain before use. The conversion to electrons is performed by the **wf32d** routine.

wf3ccd first subtracts the bias and trims the overscan regions from the image. If an associated set of UVIS `CR-SPLIT` or `REPEAT-OBS` images is being processed, all of the overscan-trimmed images are sent through **wf3rej** to be combined and receive cosmic-ray rejection. The resulting combined image then receives final calibration with **wf32d**, which includes dark subtraction and flat fielding. If there are multiple subsets of `CR-SPLIT` or `REPEAT-OBS` images in an association, each set goes through the cycle of **wf3ccd**, **wf3rej** and **wf32d** processing.

Only those calibration steps with a switch value of `PERFORM` in the input files will be executed, after which the switch will be set to `COMPLETE` in the corresponding output files.

Displaying output from wf3ccd in a Jupyter Notebook

When calling **wf3ccd** from a Jupyter notebook, informational text output from the underlying **wf3ccd.e** program will be passed through `print` as the calibration runs by default, and show up in the user's cell output. This behavior can be customized by passing your own function as the `log_func` keyword argument to **wf3ccd**. As output is read from the underlying program, the **wf3ccd** Python wrapper will call `log_func` with the contents of each line (this provides a way to connect **wf3ccd** to the Python logging system by passing the `logging.debug` function or similar).

If `log_func=None` is passed, informational text output from the underlying program will be ignored, but the program's exit code will still be checked for successful completion.

Input parameters for the Python interface

Parameters

```
input : str or list
    Name of input files, such as
        - a single filename (iaa012wdq_raw.fits)
        - a Python list of filenames
        - a partial filename with wildcards (*raw.fits)
        - filename of an ASN table (*asn.fits)
        - an at-file (@input)

output : str, default=None
    Name of the output FITS file.

dqicorr : str, optional, default="PERFORM"
    Update the dq array from bad pixel table, as well as flag the A-to-D saturation. If the comparatively
    new FITS keyword (mid-2023) SATUFILE is missing or not populated in the input file primary header, the full-
    well saturation will also be flagged using a single value as the threshold. Allowed values are "PERFORM" and
    "OMIT".

atodcorr : str, optional, default="PERFORM"
    Analog to digital correction. Allowed values are "PERFORM" and "OMIT".

blevcorr : str, optional, default="PERFORM"
    Subtract bias level determined from overscan regions. Allowed values are "PERFORM" and "OMIT".

biascorr : str, optional, default="PERFORM"
    Subtract superbias image. Allowed values are "PERFORM" and "OMIT".

    NOTE: Strictly speaking, the application of the full-well saturation image is not a calibration step (i.
    e., there is no SATCORR), but the application of a 2D image to flag pixels versus using a single scalar value
    to flag saturated pixels as previously done in DQICORR will be done in doFullWellSat() after BLEVCORR and
    BIASCORR. This correction should only be done if both BLEVCORR and BIASCORR have been performed. This flagging
    is only applicable for the UVIS.

flashcorr : str, optional, default="PERFORM"
    Subtract post-flash image. Allowed values are "PERFORM" and "OMIT".

verbose : bool, optional, default=False
    If True, print verbose time stamps.

quiet : bool, optional, default=True
    If True, print messages only to trailer file.

log_func : func(), default=print()
    If not specified, the print function is used for logging to facilitate use in the Jupyter notebook.
```

Returns

None

Usage

```
from wfc3tools import wf3ccd
wf3ccd(filename)
```

Command Line Options for the wf3ccd C executable

```
wf32ccd.e input output [-
options]
```

Input may be a single filename and the options include:

```
-v: verbose
-f: print time stamps
-dqi: update the DQ array
-atod: perform a-to-d gain
correction
-blev: subtract bias from overscan
-bias: perform bias correction
-flash: remove post-flash image
```

3.4.3 wf32d

This routine performs the remaining series of tasks in the UVIS pipeline. The **wf32d** primary functions include:

- DARKCORR: dark current subtraction
- FLATCORR: flat-fielding
- SHADCORR: apply shutter shading correction (currently skipped)
- PHOTCORR: photometric keyword calculations
- FLUXCORR: photometric normalization of the UVIS1 and UVIS2 chips

Only those calibration steps with a switch value of `PERFORM` in the input files primary headers will be executed, after which the switch will be set to `COMPLETE` in the primary headers of the corresponding output files.

Displaying output from wf32d in a Jupyter Notebook

When calling **wf32d** from a Jupyter notebook, informational text output from the underlying **wf32d.e** program will be passed through `print` as the calibration runs by default, and show up in the user's cell output. This behavior can be customized by passing your own function as the `log_func` keyword argument to **wf32d**. As output is read from the underlying program, the **wf32d** Python wrapper will call `log_func` with the contents of each line (this provides a way to connect **wf32d** to the Python logging system by passing the `logging.debug` function or similar).

If `log_func=None` is passed, informational text output from the underlying program will be ignored, but the program's exit code will still be checked for successful completion.

Parameters

Parameters

```
input : str or list
    Name of input files, such as
        - a single filename (iaa012wdq_raw.fits)
        - a Python list of filenames
        - a partial filename with wildcards (*raw.fits)
        - filename of an ASN table (*asn.fits)
        - an at-file (@input)

output : str, default=None
    Name of the output FITS file.

dqicorr : str, optional, default="PERFORM"
    Update the dq array from bad pixel table. Allowed values are "PERFORM" and "OMIT".

darkcorr : str, optional, default="PERFORM"
    Subtract the dark image. Allowed values are "PERFORM" and "OMIT".

flatcorr : str, optional, default="PERFORM"
    Multiply by the flatfield image. Allowed values are "PERFORM" and "OMIT".

shadcorr : str, optional, default="PERFORM"
    Correct for shutter shading (CCD). Allowed values are "PERFORM" and "OMIT".

photcorr : str, optional, default="PERFORM"
    Update photometry keywords in the header. Allowed values are "PERFORM" and "OMIT".

verbose : bool, optional, default=False
    If True, print verbose time stamps.

quiet : bool, optional, default=True
    If True, print messages only to trailer file.

debug : bool, optional, default=False
    If True, print debugging statements.

log_func : func(), default=print()
    If not specified, the print function is used for logging to facilitate use in the Jupyter notebook.
```

Returns

None

Usage

```
from wfc3tools import wf32d
wf32d(filename)
```

Command Line Options for the wf32d executable

```
wf32d.e input output [-
options]
```

Input may be a single filename, and the options include:

```
-v: verbose
-t: print time stamps
-d: debug
-dark: perform dark subtraction
-dqi: update the DQ array
-flat: perform flat correction
-shad: perform shading correction
-phot: perform phot correction
```

3.4.4 wf3ir

This routine contains all the instrumental calibration steps for WFC3 IR channel images. The steps are:

- DQICORR - initialize the data quality array
- ZSIGCORR - estimate the amount of signal in the zeroth-read
- BLEVCORR - subtract the bias level from the reference pixels
- ZOFFCORR - subtract the zeroth read image
- NLINCORR - correct for detector non-linear response
- DARKCORR - subtract the dark current image
- PHOTCORR - compute the photometric keyword values
- UNITCORR - convert to units of count rate
- CRCORR - fit accumulating signal and identify the cr hits
- FLATCORR - divide by the flat-field images and apply gain conversion

The output images include the calibrated image ramp (ima file) and the accumulated ramp image (flt file).

Only those calibration steps with a switch value of `PERFORM` in the input files primary headers will be executed, after which the switch will be set to `COMPLETE` in the primary headers of the corresponding output files.

Displaying output from wf3ir in a Jupyter Notebook

When calling `wf3ir` from a Jupyter notebook, informational text output from the underlying `wf3ir.e` program will be passed through `print` as the calibration runs by default, and show up in the user's cell output. This behavior can be customized by passing your own function as the `log_func` keyword argument to `wf3ir`. As output is read from the underlying program, the `wf3ir` Python wrapper will call `log_func` with the contents of each line (this provides a way to connect `wf3ir` to the Python logging system by passing the `logging.debug` function or similar).

If `log_func=None` is passed, informational text output from the underlying program will be ignored, but the program's exit code will still be checked for successful completion.

Input Parameters for the Python interface

```
Parameters

input : str
    Name of input files, such as
    - a single filename (iaa012wdq_raw.fits)
    - a Python list of filenames
    - a partial filename with wildcards (\*raw.fits)
    - filename of an ASN table (\*asn.fits)
    - an at-file (@input)

output : str, default=None
    Name of the output FITS file.

verbose : bool, optional, default=False
    If True, print verbose time stamps.

quiet : bool, optional, default=True
    If True, print messages only to trailer file.

log_func : func(), default=print()
    If not specified, the print function is used for logging to facilitate use in the Jupyter notebook.

Returns

None

Usage

from wfc3tools import wf3ir
wf3ir(filename)
```

Command Line Options for the wf3ir executable

```
wf32ir.e input output [-
options]
```

Input may be a single filename, and the options include:

```
-v: verbose
-f: print time
stamps
```

3.4.5 wf3rej

wf3rej, the cosmic-ray rejection and image combination task in **calwf3**, combines **CR-SPLIT** or **REPEAT-OBS** exposures into a single image, first detecting and then replacing flagged pixels. The task uses the same statistical detection algorithm developed for ACS (**acsrej**), STIS (**ocrrej**), and WFPC2 data (**crrej**), providing a well-tested and robust procedure.

First, **wf3rej** temporarily removes the sky background from each input image (specified by the `SKYSUB` parameter in the `CRREJTAB` reference file or by the `skysub` parameter directly passed to the Python script or C executable), usually computed using the mode of each image. Sky subtraction is performed before any statistical checks are made for cosmic rays. Next, **wf3rej** constructs an initial comparison image from each sky-subtracted exposure. This comparison image can either be a median- or minimum-value sky-subtracted image constructed from all the input images, and it represents the ‘initial guess’ of a cosmic-ray free image. The comparison image serves as the basis for determining the statistical deviation of each pixel within the input images.

A detection threshold is then calculated for each pixel based on the comparison image:

$$\tau_n = \sigma^2 \times \left(\text{noise} + \frac{\text{value}}{\text{gain}} + (\text{scale} \times \text{value})^2 \right) / T_n^2$$

where:

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

The actual detection criterion for a cosmic ray is determined as:

$$\Delta = \left(\frac{\text{pixel}_n - \text{sky}_n}{T_n} - \text{median} \right)^2$$

where:

- pixel_n is the pixel value from input image n ,
- sky_n is the sky background of image n , and
- *median* is the median or minimum pixel value 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 some expansion radius (`CRRADIUS`) are marked as ‘`SPILL`’ pixels and are given less stringent detection thresholds.

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

$$\text{pixout}(x, y) = T \cdot \frac{\sum_n (\text{pixel}_n(x, y) - \text{sky}_n) m_n(x, y)}{\sum_n T_n m_n(x, y)} + \sum_n \text{sky}_n$$

where:

- T_n is the exposure time for image n ,

- $m_n(x,y)$ is the mask value (0 for 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 value recorded in the header keywords `TEXPTIME` and `EXPTIME`.

The following keywords are also derived from the variables in this equation:

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

The remaining keywords `EXPSTART`, `EXPEND` are updated based on the values corresponding to the first and last input images, respectively. In summary, the cosmic ray rejection task sums all non-rejected 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 increasingly stringent detection thresholds, as specified by `CRSIGMAS`.

Cosmic Ray Rejection Table

`wf3rej` 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 and observation, with suitable defaults being applied during pipeline processing. Observers may fine-tune the cosmic-ray rejection parameters when manually reprocessing data with `wf3rej` 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 3.6](#). 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`). If an exact match is not found for the exposure time, the table row with the closest value is used. If the `CR-SPLIT` value of the input images exceeds the values in the table, the table row with the largest `CR-SPLIT` value will be used. The sky fitting algorithm is controlled by the parameter `SKYSUB`, which can have values of 'mode', 'mean' or 'none'. The 'initial guess' image is created using the median or minimum value of the input exposures, as specified by the value of `INITGUES`.

Cosmic-ray detection requires the specification of a threshold above which a pixel value is considered a cosmic ray. This threshold was defined above and uses sigma rejection thresholds. The sigma threshold for each iteration and number of iterations corresponds to the `CRSIGMAS` column values in the `CRREJTAB` file. `SCALENSE` is a multiplicative term (in percent) for the noise model and is given as scale in the threshold equation above. 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 detector will cause stars to be mistakenly 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 'do no harm'. However, for cases in which there have been no image-to-image 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 the 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 3.6](#).

Table 3.6: Columns in the cosmic-ray rejection parameter table.

Column Name	Default Value	Contents
CRSPLIT	-	Number of exposures into which observation was split
CCDCHIP	-	Chip to which this conversion applies
MEANEXP	-	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)
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?

Displaying output from `wf3rej` in a Jupyter Notebook

When calling `wf3rej` from a Jupyter notebook, informational text output from the underlying `wf3rej.e` program will be passed through `print` as the calibration runs by default, and show up in the user's cell output. This behavior can be customized by passing your own function as the `log_func` keyword argument to `wf3rej`. As output is read from the underlying program, the `wf3rej` Python wrapper will call `log_func` with the contents of each line (this provides a way to connect `wf3rej` to the Python logging system by passing the `logging.debug` function or similar).

If `log_func=None` is passed, informational text output from the underlying program will be ignored, but the program's exit code will still be checked for successful completion.

Input Parameters for the Python interface

Parameters

```
input : str or list
        Name of input files, such as:
        - comma-separated (no spaces) filenames (iaao01k8q_flc.fits,iaao01k9q_flc.fits)
        - a Python list of filenames
        - a partial filename with wildcards (*flt.fits)
        - an at-file (@input)

output : str
        Name of the output FITS file.
```



```

crrehtab : str, default=""
    Reference filename.

scalense : float, default=0.
    Scale factor applied to noise.

initgues : str, default=""
    Initial value estimate scheme (min|med).

skysub : str, default=""
    How to compute the sky (none|mode|mean).

crsigmas : str, default=""
    Rejection levels in each iteration.

crradius : float, default=0.
    Cosmic ray expansion radius in pixels.

crthresh : float, default=0.
    Rejection propagation threshold.

badinpdq : int, default=0
    Data quality flag bits to reject.

crmask : bool, default=Setting to be read from CRREHTAB.
    If argument is present, write the CR flag value to the input DQ images.

shadcorr : bool, default=Setting to be read from SHADCORR keyword value in primary header of first image to
process.
    If argument is present, perform shading shutter correction.

verbose : bool, optional, default=False
    If True, print timestamps and other output.

log_func : func(), default=print()
    If not specified, the print function is used for logging to facilitate use in the Jupyter notebook.

```

Returns

None

Usage

```

from wfc3tools import wf3rej
from glob import glob
infiles = glob("*flt.fits")
wf3rej(infiles, "output.fits", verbose=True)
wf3rej("*flt.fits", "output.fits", verbose=True)
wf3rej("@input.lst", "output.fits", verbose=True)

```

Please see the warning at the bottom of the page regarding the parameter settings for **wf3rej.e** for more details as to the action taken when the parameters use their default values.

Command Line Options for the wf3rej C executable

```

wf3rej.e input output [-r] [-v] [-t] [-shadcorr] [-crmask] [-table <filename>]
[-scale <float>] [-init <med|min>] [-sky <none|mode|mean>] [-sigmas <string>]
[-radius <float>] [-thresh <float>] [-pdq <short>]

```

Example - Process data with timestamps and a custom cosmic ray rejection table:

```

wf3rej.e iaao01k8q_flc.fits,iaao01k9q_flc.fits output.fits -t

```

Example - Print the code version and exit:

```

wf3rej.e -r

input : comma-separated list of strings
       - Input filenames as a list of comma-separated input names
         - ippssoot_raw.fits,ippssoot_raw.fits (Note: Do not include any blank spaces.)

output : str
        Name of output filename

options
  -r : print version number/date of software and exit (no other options selected)
  -v : verbose mode
  -t : print the timestamps
-shadcorr : perform shading shutter correction    (not used)
-crmask : set CR flags in input DQ images

-table <filename>: string, the crrejt看 filename
-scale <number>: float, scale factor for noise
-init <med|min>: string, initial value estimate scheme
-sky <none|mode|mean>: string, method to compute sky

-sigmas <string of numbers>: string, rejection levels for each iteration (e.g., "3.5,4.5,5.5")

-radius <number>: float, CR expansion radius

-thresh <number> : float, rejection propagation threshold

-pdq <number>: short, data quality flag bits to reject

```

! Observers should not invoke the `shadcorr` option as there is no `SHADFILE`.

Including the `crmask` option on the command line will have the CR flag values placed into the DQ extension of the input images. Omitting this option does not turn off the insertion, but rather the program will follow the default setting for the `crmask` option as indicated in the `CRREJTAB` calibration file.

If any of the CR-rejection options are omitted from the command line, default values from `CRREJTAB` will be used; those options are: `crmask`, `scale`, `init`, `sky`, `sigmas`, `radius`, `thresh`, and `pdq`. In verbose mode, all of the option values are printed to the output logfile.

3.5 Manual Recalibration of WFC3 Data

[3.5.1 Requirements for Manual Recalibration](#)

[3.5.2 calwf3 Examples](#)

3.5.1 Requirements for Manual Recalibration

Retrieving Software and Input Data Files

Data retrieved from MAST (Barbara A. Mikulski Archive for Space Telescopes), via the [Portal](#) or the [search form](#), will have been processed with the most up-to-date calibration reference files and software. Nevertheless, some users may wish to reprocess existing WFC3 observations with e.g. alternative reference files. To do so, the following must be available on their system:

- Software for data processing (distributed in the [stenv](#) package)
- Reference files may be obtained from [CRDS](#), as described in [Section 3.1.4](#).
- Uncalibrated (raw) data files from MAST
- Any association tables that describe observation sets.

Setting the Path to Reference Files

Before any recalibration can be done, the user must define a local directory containing the calibration reference files. The directory for WFC3 is given the environment variable name 'iref'. The raw image headers will already contain the appropriate keywords, 'iref' variable, and the reference file names that were assigned during STScI pipeline processing. The user must simply define the 'iref' directory as the local directory in their Unix environment:

```
csh shells: setenv iref /mydisk
/myiref/
sh shells: export iref /mydisk
/myiref/
```

If executed from the command line, this setup must be done in the same window in which Python will be started. (Setting 'iref' from within Python will not work, even though subsequently typing 'show iref' would suggest it might.) For convenience, this setup command can be added to the '.setenv' file for csh or '.bash_profile/.bashrc' file for sh, so that the 'iref' environment variable will always be defined when you start a new terminal. The 'iref' environment variable is required whether reprocessing through Python or from the OS command line.

An alternate method of accessing WFC3 reference files is shown in a new Python tutorial in [Section 3.5.2](#). In this example, the best WFC3 reference files for a given set of input images are retrieved directly from CRDS and placed in a local 'crds_cache/' directory. The notebook shows how to update the header keywords to point to these new reference files and sets the 'iref' environment variable for subsequent **calwf3** processing.

Selecting Calibration Switches

The MAST HST Data Processing uses the most up-to-date calibration reference files by default. In order to use non-default reference files, manual recalibration is required and, in this case, the calibration reference file keywords will need to be updated manually in the primary header of the raw data files with the desired file names before running **calwf3**. In addition, the user can choose to change which calibration steps are performed by **calwf3** by resetting the values of the calibration switch keywords. These keywords are listed in [Table 3.7](#) along with their default values as used in the STScI pipeline. To change the values of any of the keyword switches, a FITS keyword editor such as the Python package `astropy.io.fits` may be used.

```
from astropy.io import fits
fits.setval('path/to/myfile_raw.fits', keyword= 'DARKCORR', value= 'OMIT',
ext=0)
```

Table 3.7: Calibration switch and default settings.

SWITCH	DESCRIPTION	CRITERIA
PCTECORR	Perform CTE correction (UVIS)	DEFAULT = 'PERFORM' (cte-corrected) else 'OMIT' (non-cte-corrected)
DQICORR	Data Quality Array Initialization	DEFAULT = 'PERFORM'
ATODCORR	Analog-to-Digital Correction	DEFAULT = 'OMIT'
ZSIGCORR	Zero-read Signal Correction (IR)	DEFAULT = 'PERFORM'
BLEVCORR	CCD Overscan Region Subtraction (UVIS) Reference Pixel Bias Correction (IR)	DEFAULT = 'PERFORM'
BIASCORR	Superbias Image Subtraction (UVIS)	DEFAULT = 'PERFORM'
FLSHCORR	Post-flash Image Subtraction (UVIS)	If FLASHLVL > 0 then 'PERFORM' else 'OMIT'
CRCORR	UVIS Cosmic Ray Rejection:	If CRSPLIT or NRPTEXP > 1 then 'PERFORM', else 'OMIT'
	IR Cosmic Ray Rejection:	DEFAULT = 'PERFORM' (else 'OMIT' for scan observations)
ZOFFCORR	Zero-read Image Subtraction (IR)	DEFAULT = 'PERFORM'
NLINCORR	Non-linearity Correction (IR)	DEFAULT = 'PERFORM'
DARKCORR	Dark Image Subtraction	DEFAULT = 'PERFORM'
FLATCORR	Flat-Field Correction	DEFAULT = 'PERFORM'

SHADCORR	Shutter Shading Correction (UVIS)	DEFAULT = 'OMIT'
UNITCORR	Units Conversion (IR)	DEFAULT = 'PERFORM'
PHOTCORR	Photometric Processing	DEFAULT = 'PERFORM'
FLUXCORR	Photometric normalization of the two UVIS chips	DEFAULT = 'PERFORM'
RPTCORR	Repeat-Obs Processing (IR)	If NRPTEXP > 1 then 'PERFORM', else 'OMIT'
EXPSCORR	Full calibration of individual exposures in an association	DEFAULT = 'PERFORM'
DRIZCORR	Drizzle Processing	DEFAULT = 'PERFORM'

3.5.2 calwf3 Examples

Documentation for the **calwf3** software is provided at [readthedocs](#), with version history available on the [STScI github repository](#) or [readthedocs software history](#). Users may reprocess either a single RAW image, an image association, or a list of raw files in Python:

A single RAW exposure

```
from wfc3tools import calwf3
calwf3('ibohbfb9q_raw.fits')
```

An image association (ASN)

```
calwf3('ibohbf040_asn.fits')
```

A list of RAW files (not part of an ASN)

```
import glob
for raws in glob.glob('ibohbf*_raw.fits'):
    calwf3(raws)
```

A new [WFC3 notebook](#) tutorial provides **calwf3** reprocessing examples to improve calibrated WFC3/IR images affected by time-variable background. The notebook shows how to diagnose images with poor-quality ramp fits and rerun calwf3 with the 'CRCORR' step turned off.

✔ Step 3 of this notebook also provides an alternative way to query CRDS for the best WFC3 reference files.

Time-variable background, described in detail in [Section 7.10](#), can be caused by either scattered Earth shine and/or helium I emission in the Earth's atmosphere at 1.083 microns ([WFC3 ISR 2014-03](#)). The latter effect can appear in F105W and F110W images as a flat excess signal which is added to the total background for a subset of reads. When strong enough, this non-linear signal may compromise the ramp fitting performed by `calwf3`, which is designed to flag and remove cosmic rays and saturated reads. The affected calibrated FLT data products will have much larger noise and a distinctly non-gaussian background. In the tutorials described in [Section 7.10](#), the images are reprocessed using the 'Last-minus-first' technique described in [WFC3 ISR 2016-16](#). `calwf3`'s ramp fitting step (`CRCORR`) is turned off and treats the IR detector like a CCD that accumulates charge and is read out only at the end of the exposure. In this case, the observed count rate is determined by simply subtracting the first from the last read of the detector and dividing by the time elapsed between the two reads.

Note: while time-variable background also impacts the IR grisms, the methods used for imaging data should not be used to correct G102 and G141 observations, which are affected by a combination of zodiacal background, helium line emission, and scattered Earth light, each of which vary spatially across the detector. More detail on background corrections for grism data is provided in [WFC3 ISR 2020-04](#).

Chapter 4: WFC3 Images: Distortion Correction and AstroDrizzle

Chapter Contents

- [4.1 WFC3 Geometric Distortion](#)
- [4.2 Distortion Corrections and Image Combination](#)
- [4.3 References](#)

4.1 WFC3 Geometric Distortion

[4.1.1 The UVIS Channel](#)

[4.1.2 The IR Channel](#)

WFC3 images exhibit significant geometric distortion, similar to that seen in ACS images. The required folding of the light paths in both channels to fit them within the instrument's optical-bench envelope results in substantial tilts of the focal surfaces with respect to the chief rays. The WFC3 UVIS detector is tilted at ~ 21 degrees producing a rhomboidal elongation of $\sim 7\%$. The IR detector has a ~ 24 degree tilt about its x-axis, creating a rectangular elongation of $\sim 10\%$. The orientations of the WFC3 detector edges for both detectors are at approximately 45 degrees with respect to the V2 and V3 coordinate axes of the telescope. [Figure 2.2 of the WFC3 Instrument Handbook](#) shows the WFC3 apertures in the telescope's V2, V3 reference frame.


The first on-orbit measurements of the geometric distortion for the WFC3 detectors were made in 2009 during SMOV (Servicing Mission Observatory Verification). The standard astrometric fields in the vicinity of globular cluster 47 Tuc (NGC 104) and the Large Magellanic Cloud (LMC) were observed with multiple offsets (proposal 11444 for UVIS with F606W and 11445 for IR with F160W). The derived WFC3 geometric distortion is computed as a fourth-order polynomial model and is accurate to a precision level of better than 0.05 pixel in both UVIS and IR or about 2 and 5 mas, respectively ([WFC3 ISR 2009-33](#) and [WFC3 ISR 2009-34](#)).

In Cycles 17-21 (2009-2014) the globular cluster Omega Centauri was observed through all UVIS filters and five IR filters. These observations were used to derive the high-order geometric distortion coefficients as a function of filter ([Table B.1](#)) on both channels. The unique polynomial coefficients of the geometric distortion are stored in the Instrument Distortion Coefficient Tables (IDCTABs) that are used to correct for distortion in MAST drizzled data products. The name of the IDC table used during calibration is stored in the primary header keyword IDCTAB in the WFC3 science image header. These reference files are also necessary for users who wish to reprocess their data with the [DrizzlePac/AstroDrizzle](#) software.

For the UVIS CCDs, in addition to large-scale geometric distortion, there is a micro-distortion in the form of fine-scale systematics in the residuals from the best-fit polynomial solutions. These systematic residuals typically extend ~ 0.15 pixel and vary in amplitude depending on the location within a CCD chip. Such residuals cannot be removed by a polynomial model. These fine-scale and low-amplitude distortions are the result of the photo-lithographic mask pattern used during the chip manufacture which imprints a large 6x2 grid onto each chip. The correction for the lithographic pattern imprint is implemented in [DrizzlePac/AstroDrizzle](#) as a 2-D bi-linear interpolation (D2IMFILE, Detector to Image distortion correction reference file) which is applied prior to the large-scale distortion (polynomial model) correction (IDCTAB) ([WFC3 ISR 2013-14](#)). The look-up table is accessed by [AstroDrizzle](#) in the pipeline using the filename specified in the D2IMFILE keyword in the primary image header. [AstroDrizzle](#) attaches new fits extensions to the calibrated image of type 'D2IMARR' (four extensions total, two for each chip in each X and Y axis). See [Table 2.5](#) for details on these additional image header extensions.

In addition to the pixel-grid irregularities there are fine-scale non-polynomial filter-dependent distortions that have been calibrated for a subset of 36 UVIS filters ([WFC3 ISR 2014-12](#) and [WFC3 ISR 2018-10](#)). These filter-dependent distortion solution tables are specified in the `NPOLFILE` keyword in the primary image header and attached to the image as FITS extensions (four extensions, two for each CCD chip and each X and Y axis), [WFC3 ISR 2014-12](#). Each element of the 64×32 table is bi-linearly interpolated pixel-by-pixel in UVIS images by **AstroDrizzle**.

Figure 4.1: Non-linear components (vectors magnified by 10) of the geometric distortion on the WFC3-UVIS detector.




4.1.1 The UVIS Channel

[Figure 4.1](#) illustrates the shape of the UVIS channel field of view as projected onto the sky. As noted above, its rhomboidal shape is due primarily to the tilt of the CCD around its detector diagonal with respect to the chief ray. The angle between the x- and y-axes is ~ 86.1 degrees. The field diagonals are tilted slightly from the V2- and V3-axes (see [Figure 1.2](#)). By design, there is a ~ 1.2 " gap between the two CCD chips. The dots in the diagram indicate where points in the image would be located without non-linear distortion, and the vectors, scaled up by a factor of 10, indicate the actual locations of the points on the sky, including the non-linear distortion components. The corner displacements are about 140 pixels, corresponding to 5.5". The principal effect is the diagonal variation of scale. At the center of UVIS1 (CCD CHIP1, top), the scale in the x-direction is 0.0396"/pixel, and 0.0393"/pixel in the y-direction. For UVIS2 (CCD CHIP2, bottom), these scales are 0.0398"/pixel, and 0.0398"/pixel, respectively. UVIS1 forms a slightly distorted rectangle 162" \times 81" in size, while UVIS2 subtends 163" \times 82". In the presence of distortion, it is not possible to flat field the image in such a way as to preserve both surface brightness and flux. The convention that has been adopted in the pipeline preserves surface brightness, therefore a [pixel area map correction](#) is necessary if the goal is to measure accurate fluxes via point-source photometry (see [Section 9.1.11](#) and [WFC3 ISR 2010-08](#)).

4.1.2 The IR Channel

The IR detector field of view is nominally concentric with the UVIS field, but it subtends a smaller area on the sky, 138" \times 123". The detector tilt is about the x-axis, so the projected aperture shape is nearly a rectangle, with the angle between the x- and y-axes on the sky nearly 90 degrees as shown by the outline in [Figure 4.2](#). At field center, the x- and y-scales are 0.135 and 0.121"/pixel, respectively. A vector plot of the deviation from linearity is also shown in [Figure 4.2](#), where the deviations have been magnified by a factor of 10 for illustrative purposes. The largest deviation is ~ 10 pixels, corresponding to about 1.4".

Figure 4.2: Non-linear components (vectors magnified by 10) of the geometric distortion on the WFC3-IR detector.



4.2 Distortion Corrections and Image Combination

[4.2.2 Manual AstroDrizzle Reprocessing](#)

[4.2.3 AstroDrizzle Documentation](#)

Due to the tilt of the WFC3 focal plane with respect to the incoming light beam, the projected pixel area on the sky varies across the field of view in the raw (`raw.fits`) and calibrated (`flt/flc.fits`) images. As a result of the different projected pixel area on the sky, pixels in different regions of the detector collect a different amount of light, i.e. observations of a constant surface brightness object would have count rates per pixel that vary over the detector even if every pixel had the identical sensitivity. In WFC3, the pixel area on the sky varies by about 7% along the diagonal in UVIS images and about 8% from top to bottom in IR images.


In order to produce images that appear uniform for uniform illumination, the WFC3 flat fields include the effect of the variable pixel area across the field (see [Section 5.4](#)). However as a consequence of dividing by the flat field, two stars of equal brightness in an `flt/flc.fits` image falling on different portions of the detector would not have the same total counts. To perform point source photometry on calibrated '`flt/flc.fits`' images, they must be multiplied by the effective [pixel area map](#) (see [WFC3 ISR 2010-08](#)). Alternatively, the pixel area effect is accounted for in the pipeline by drizzling, where the geometric distortion solution is used to correct all pixels to equal areas on the sky in the '`drz/drc.fits`' data products. Because drizzling conserves flux, users should recover the aperture photometry from `drz/drc` images and from `flt/flc` images multiplied by the pixel area map.

4.2.1 AstroDrizzle in the Pipeline

WFC3 data obtained from MAST are corrected for geometric distortion with **AstroDrizzle**, which replaced MultiDrizzle in the OPUS pipeline for WFC3 data on June 7, 2012. During pipeline processing, calibrated data that belong to an association (e.g. as defined by the user in APT via a standard dither pattern, a `REPEAT-OBS`, or `CR-SPLIT` pair, see [Table 2.2](#)) are corrected for distortion and drizzle-combined with cosmic-ray rejection. If the associated images are dithered, they are aligned using the World Coordinate System (WCS) information in their headers before being combined. If there is no association table, each single-exposure WFC3 image is drizzled to correct for geometric distortion. Additional WFC3-specific rules define that images obtained with the same filter within a given visit are associated and thus combined.

AstroDrizzle uses the `MDRIZTAB` reference table to define a default set of parameter values that work well for most use cases. Each detector has its own `MDRIZTAB`, and each row provides the parameter settings specific to the number of input images per association. To correct for distortions, **AstroDrizzle** relies on the following reference files.

- Image Distortion Correction Table (`IDCTAB`, high-order polynomial coefficients)
- Detector to Image Distortion Correction (`D2IMFILE`, lithographic mask pattern correction for UVIS only)
- Non-polynomial Filter-Dependent Distortion (`NPOLFILE`, 2-D look-up table for each calibrated filter)

 The names of the reference files used by **AstroDrizzle** are stored in the WFC3 UVIS and IR primary header keywords `IDCTAB`, `D2IMFILE`, and `NPOLFILE`. Users interested in these files may obtain them from [CRDS](#).

As of December 2019, the MAST pipeline updates the astrometry (alignment) of WFC3 data with two approaches. The first, an “a priori” method, updates the positions of the guide stars used in a given image with the coordinates from Gaia, as these are typically more accurate than the positions from the Guide Star Catalog. The second approach, an “a posteriori” alignment of the data, when possible matches sources in the image to an external reference catalog (such as Gaia) and corrects for any offsets. These updates improve the absolute astrometry of WFC3 data products, especially when an a posteriori Gaia alignment solution is present. Further details regarding these improvements are discussed in [WFC3 ISR 2022-06](#) as well as the **Drizzlepac** documentation.

During **AstroDrizzle** processing, the geometric distortion is extracted from these reference files and stored as Simple Image Polynomial (SIP) header keywords and as additional FITS extensions in the *_flt.fits/flc.fits images. Please refer to [Table 2.4](#) or to Section 3.2.3 in the [DrizzlePac Handbook](#) for details.

The steps performed by **AstroDrizzle** include the following:

- Correct the geometric distortion.
- Project images onto a single pixel grid using the header World Coordinate System.
- Match the sky background levels across overlapping images.
- Perform cosmic-ray rejection using the *_flt/flc.fits files as input. Note: IR flt data are already corrected for cosmic rays, but the drizzling process does identify further detector artifacts and rejects them during the combination process.
- Convert the UVIS data from units of electrons to electrons per second; IR data are already in e⁻/s.
- Combine associated (e.g. dithered) observations into a single product.

In addition to the improved astrometry in the standard pipeline products, two further versions of products are created in the MAST pipeline, referred to as Hubble Advanced Products. The first of these are the Single Visit Mosaics, which improve the relative alignment of all images taken within a given visit (whilst the standard pipeline products mentioned above are aligned at the association level), and also attempt to align the images to the Gaia frame. The second type of advanced products are the Multi Visit Mosaics (MVM), which combine all images in a given filter within a given part of the sky. The MVM products do not perform any further realignment from the Single Visit Mosaics, so they may have alignment errors (e.g. when visits were taken at different epochs) and thus should be used as discovery products. For more information, see [WFC3 ISR 2022-06](#) as well as the [Drizzlepac documentation](#).

4.2.2 Manual AstroDrizzle Reprocessing

Drizzled images combined in the pipeline were produced using a default set of parameters that are suitable for the widest range of scientific applications. These defaults, however, may not produce the optimum science data quality for many programs, and those images will require post-pipeline processing.

Four main areas for improvement may include: (1) image alignment, (2) sky subtraction, (3) cosmic ray rejection, and (4) final image resolution. While single visit data with small dithers (like the 4-point dither box) are usually aligned to better than 0.1 pixel, the drizzled products are created using the native detector plate scale, and with a drop size, defined as the projected pixel shrinking factor (`PIXFRAC`), of 1.0. In these cases, the resolution of the drizzled products can be improved by fine-tuning the final sampling i.e. experimenting with the scale and `PIXFRAC` parameters.

While a significant portion of WFC3 data is aligned to Gaia by the MAST pipeline, the resulting relative astrometry for data taken in different visits of the same target may still have small errors which could degrade image combination. Thus, it may be necessary to manually refine alignment of data, and re-drizzle. Furthermore, there may be cases where the automated alignment does not obtain the best solution, in which case some manual reprocessing may be needed. In these cases, see the [Drizzlepac Notebooks](#) for guidance.

Poor alignment can lead to improper cosmic-ray rejection, and inadvertent flagging of some astronomical sources as cosmic rays, compromising the photometric accuracy of the final data products. Additionally, a poor estimate of the sky background, for example in images where a bright target fills the frame, may also affect the accuracy of cosmic-ray rejection, and in turn, the resulting photometry.

4.2.3 AstroDrizzle Documentation

AstroDrizzle is available as part of the **DrizzlePac** software, which contains all the tools for manually reprocessing (aligning and combining) ftc/flt HST images. This software may be obtained from the [DrizzlePac](#) web page. This page also provides useful resources such as the [DrizzlePac Handbook](#) and [documentation](#), a 'Quick Start Guide' to drizzling, a set of example Jupyter notebooks, and some basic video tutorials.

4.3 References

- M. Marinelli and L. Dressel, 2024, *Wide Field Camera 3 Instrument Handbook*, Version 16.0 (Baltimore: STScI).
- J. S. Kalirai et al., 2010, *WFC3 Pixel Area Maps*, [WFC3 ISR 2010-08](#).
- V. Kozhurina-Platais, et al., 2009, *WFC3 SMOV Proposal 11445 - IR Geometric Distortion Calibration*, [WFC3 ISR 2009-34](#).
- V. Kozhurina-Platais, et al., 2009, *WFC3 SMOV Proposal 11444 - UVIS Geometric Distortion Calibration*, [WFC3 ISR 2009-33](#).
- V. Kozhurina-Platais, M. Dulude, T. Dahlen, C. Cox, 2012, *WFC3/UVIS and IR Multi-Wavelength Geometric Distortion*, [WFC3 ISR 2012-07](#).
- V. Kozhurina-Platais, D. Hammer, N. Dencheva, & W. Hack, 2013, *Astrometric Correction for WFC3 /UVIS Lithographic-Mask Pattern*, [WFC3 ISR 2013-14](#).
- V. Kozhurina-Platais, 2014, *Astrometric Correction for WFC3/UVIS Filter-Dependent Component of Distortion*, [WFC3 ISR 2014-12](#).
- S. L. Hoffmann et al. 2021, *The DrizzlePac Handbook*.
- J. Mack et al., 2022, *Improved Absolute Astrometry for ACS and WFC3 Data Products*, [WFC3 ISR 2022-06](#).
- C. Martlin et al., 2018, *Updates to the WFC3/UVIS Filter-Dependent and Geometric Distortions*, [WFC3 ISR 2018-10](#).

Chapter 5: WFC3 UVIS Sources of Error

Chapter Contents

- [5.1 Gain and Read Noise](#)
- [5.2 Bias Subtraction](#)
- [5.3 Dark Current and Hot Pixels](#)
- [5.4 UVIS Flat Fields](#)
- [5.5 Image Anomalies](#)
- [5.6 Generic Detector and Camera Properties](#)
- [5.7 UVIS Photometry Errors](#)
- [5.8 References](#)

5.1 Gain and Read Noise

- 5.1.1 Gain
- 5.1.2 Read Noise

5.1.1 Gain

Photoelectrons that accumulate in a CCD pixel are read out as a voltage and converted to data numbers (DN), often called Analog-to-Digital Units (ADUs), by the analog-to-digital converter (ADC). The ADC output is a 16-bit number, capable of producing a maximum of 65,535 DN ($2^{16} - 1$) for each pixel. A straightforward scheme in which one DN corresponds to one electron would make it impossible to measure signals larger than 65,535 electrons. The conversion gain, defined as the number of electrons per DN, provides a way of changing the dynamic range of electrons that can be measured. Although it is possible to operate the WFC3 CCD detector at gains of 1, 1.5, 2, and 4 e^- /DN, only a gain of 1.5 e^- /DN is supported. This gain permits sampling of the entire dynamic range of the detectors with negligible impact on the readout noise. The absolute gain values for the WFC3 CCDs, routinely measured every year, have been stable to within 1-2% since launch in May 2009 ([WFC3 ISR 2013-02](#), [WFC3 ISR 2015-05](#), [WFC3 ISR 2016-13](#), [WFC3 ISR 2018-17](#), and [WFC3 ISR 2022-08](#)). The most recent absolute gain values, as a function of chip, amp, and binning (1x1, 2x2, 3x3) are summarized in [Table 5.1](#). The gains were measured in flat-field images via the standard mean-variance technique: the inverse slope of the mean signal level plotted versus the variance yields the gain.

Table 5.1: WFC3/UVIS absolute gain values (measurement errors are: ~0.01 e-/DN)

CCD Chip	Amp	Gain at 1 x 1 (e-/DN)	Gain at 2 x 2 (e-/DN)	Gain at 3 x 3 (e-/DN)
1	A	1.58	1.58	1.56
1	B	1.58	1.57	1.57
2	C	1.60	1.59	1.57
2	D	1.59	1.60	1.58

5.1.2 Read Noise

The average read noise level measured in the bias frame science area pixels at the default gain setting (1.5 e^- /DN) are summarized in [Table 5.2](#). Scatter in the measurements based on data acquired between June 2009 and July 2015 is less than 1%, depending on the amplifier (WFC3 ISRs [2017-17](#), [2015-13](#), [2009-26](#)).

Table 5.2: WFC3/UVIS readout noise, in electrons, and uncertainty for unbinned and binned modes.

Mode	Amplifier A			Amplifier B			Amplifier C			Amplifier D		
Binning	1 x 1	2 x 2	3 x 3	1 x 1	2 x 2	3 x 3	1 x 1	2 x 2	3 x 3	1 x 1	2 x 2	3 x 3

Mean	2.97	3.11	3.22	3.03	3.15	3.26	2.95	2.99	3.09	3.06	3.29	3.38
Uncertainty	0.01	0.02	0.04	0.01	0.01	<0.01	0.00	<0.01	0.01	0.01	0.02	<0.01

An analysis of the statistical behavior of the WFC3 ADCs shows some tendency for the least significant bit to be slightly biased at the readout speed adopted by the WFC3 electronics (see [WFC3-ISR 2005-27](#)). This minor effect should not degrade the photometric and noise characteristics of the WFC3/UVIS images.

5.2 Bias Subtraction

[5.2.1 Bias Calibration](#)

[5.2.2 Bias Correction for WFC3 Subarrays](#)

5.2.1 Bias Calibration

Bias reference frames are acquired frequently for scientific calibration purposes and for monitoring the detector performance. Multiple bias frames are stacked into a reference superbias image, removing the cosmic rays accumulated during the readout time enhancing the signal-to-noise ratio of the final results (e.g. [WFC3 ISR 2023-03](#)). `calwf3` performs the bias correction in two steps (see [Section 3.2.5](#) and [Section 3.2.6](#)). The first step consists of subtracting a global bias level, computed from the overscan regions which provide a measure of the instrument bias level at the time when the images were taken. The location of the overscan regions in a raw image varies, depending upon the type of readout that is performed (see [Figure 3.2](#)). This overscan-based bias level is subtracted from the raw image through the `BLEV CORR` step in the WFC3 calibration pipeline. In the second step, residual two-dimensional bias structure is removed via the superbias reference file correction applied via the `BIAS CORR` step.

5.2.2 Bias Correction for WFC3 Subarrays

When science data are obtained in a subarray format, the requisite dark and flat-field corrections are obtained from the full-frame calibration files, extracted from the appropriate subregion. The superbias correction is also extracted from full-frame 4-amp readout superbias files (if the subarray resides entirely within a single detector 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. For subarrays that span detector quadrants, special superbias files must be constructed from individual bias frames read out through the same amplifier as the subarray (a subarray is always read out through a single amplifier). For example, a single-chip readout (the largest possible subarray) read out through amp A must be calibrated with a superbias constructed from single-chip bias frames read out through amp A; similarly, science data read out through amp B must be calibrated with a superbias constructed from bias frames read out through amp B.

5.3 Dark Current and Hot Pixels

[5.3.1 Dark Current](#)
[5.3.2 Hot Pixels](#)
[5.3.3 Pixel Stability](#)

5.3.1 Dark Current

Superdark reference files are generated on a daily basis, with typically between 10 - 18 individual dark images (acquired over 4-5 days) contributing to each superdark. The individual darks are recalibrated with the appropriate superbias and post-flash files (see [Section 3.2.8](#) for more information on post-flash) using the most recent **calwf3** software version, stacked to remove cosmic rays, and converted from DN to e^-/sec . Any pixels with values $> 54 e^-/\text{hr}$ are considered hot; their values are left unchanged in the science extension and flagged with a value of 16 in the DQ extension, which is propagated into the final science image DQ extensions. In this way, observers can decide whether to ignore hot pixels or to allow the dark subtraction to stand (for instructions on how to control which bit masks are used during drizzling, please consult the [HST DrizzlePac Handbook](#)).

Because the mean dark current in the WFC3 CCDs is low ($\sim 11 e^-/\text{hr}/\text{pixel}$ in early-2024) it is very difficult to achieve a useful signal-to-noise for pixels that have normal levels of dark current with data from only 10-18 individual dark frames. Subtracting these uncertain values from science images during calwf3 processing would introduce noise into the calibrated images. Therefore, all cold pixels in the SCI extensions of superdark reference images are set to an anneal-cycle-averaged value. In other words, each good pixel in the superdark reference image is set to the average value of that pixel over all individual darks taken in an anneal cycle (about 100 frames). The generation of the dark calibration reference files is described in detail in [WFC3 ISR 2016-08](#). Users can verify whether the superdark file most appropriate to their observations has been installed for pipeline use in two different ways:

- Re-retrieve the images from the MAST, which automatically updates the headers and recalibrates the data
- Check the CRDS ([Calibration Reference Data System](#)).


Using an out-of-date superdark reference file can produce a poor dark correction by 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 warmed to anneal hot pixels in the interim).

5.3.2 Hot Pixels

Two types of bad pixels are routinely monitored using on-orbit WFC3 data: hot pixels and dead pixels. Hot pixels, i.e., those pixels with a higher than normal dark current, are identified in dark frames using a threshold of $0.015 e^-/\text{sec}$ ($54 e^-/\text{hr}$). The cutoff was chosen based on the tail of the dark histogram in early on-orbit data (see [Figure 5.1](#)) as well as visual examination of the dark frames. The number of hot pixels increases over time due to on-orbit radiation damage. Periodic anneal procedures, where the UVIS detector is warmed to $\sim 20^\circ\text{C}$, successfully fix a percentage of the hot pixels. Hot pixel locations and levels are provided by the UVIS superdark reference files, although dithering can mitigate their effect as well.

Dead pixels, specifically dead columns, are identified through visual inspection of both individual and stacks of internal frames. Bad pixel locations are propagated into the bad pixel mask (header keyword `BPIXTAB` and the file name `*_bpx.fits`) that is applied by `calwf3` in the standard data reduction pipeline. Currently, there are ~ 8000 dead pixels in chip 1 (amps A and B) and ~ 16000 in chip 2 (amps C and D); these comprise $\sim 0.1\%$ and 0.2% of each chip, respectively.

Figure 5.1: Histograms of three dark images (hot pixel threshold of $0.015 \text{ e}^-/\text{sec}$ marked with vertical dotted line).



Trending

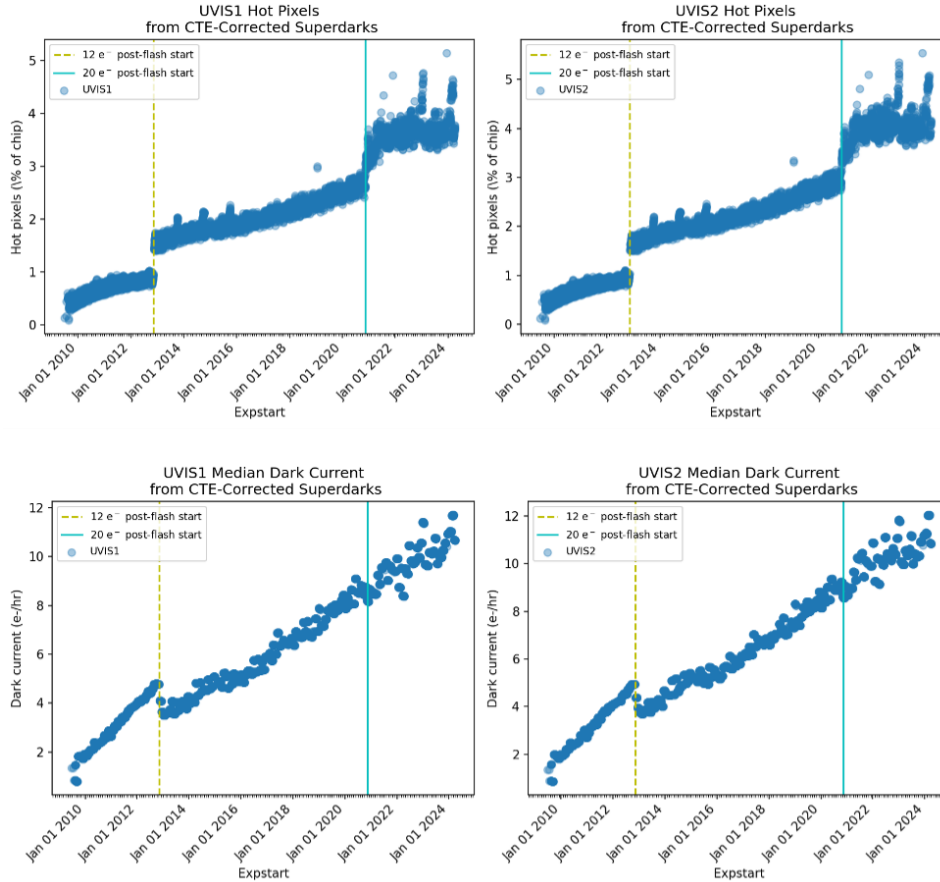
We have chosen a limit of $0.015 \text{ e}^-/\text{s}/\text{pixel}$ ($54 \text{ e}^-/\text{hr}/\text{pixel}$) to be the threshold above which we consider a pixel to be "hot" based on the tail of the histogram as well as a visual examination of 900-second dark frames taken during Cycle 17. Figure 5.1 shows a histogram of CR-free pixels from 900-second darks taken at three different times after the April 2010 anneal procedure: immediately following the procedure (red line), about 10 days later (green line), and about 18 days later (blue line). The increase in hot pixels due to on-orbit radiation damage is apparent; the anneal procedures have been found to fix a fraction of the hot pixels which accumulate over time. The hot pixel cutoff is shown with a vertical line at $54 \text{ e}^-/\text{hr}$; at this threshold, the growth rate for WFC3 hot pixels is ~ 1000 pixels/day (see [WFC3 ISR 2016-08](#) for more information on darks and hot pixels).

[Figure 5.2](#) shows the number of hot pixels as a function of time since the installation of WFC3 on HST. Between Nov 8, 2012 (dashed yellow vertical line) and Nov 12, 2020 (solid cyan vertical line) the UVIS darks were post-flashed, i.e., an LED applied a background of $12\text{ e}^-/\text{pixel}$ to each image. Post-flashing preserves faint signal in pixels far from the amps that otherwise would have been undetectable due to CTE losses. With the ongoing accumulation of damage due to high-energy radiation, as of Nov 12, 2020 the UVIS darks have had a 20 e^- post-flash to further mitigate diminishing CTE. The result is a jump in the number of *detected* hot pixels as the increased background improves CTE, allowing the hot pixels to be identified as hot. In addition, the post-flash's suppression of hot pixel and cosmic ray CTE trails also causes a commensurate reduction in the overall measured dark current. Dark current and hot pixel plots are updated, as resources allow, on the [WFC3 Performance Monitoring webpage](#) or the user may email the STScI helpdesk.

As discussed above, the WFC3 CCD detectors degrade over time due to exposure to the space environment. This damage manifests itself in the darks as an increase in the number of individual hot pixels as well as in an overall higher dark current (CTE loss, another manifestation of damage, is discussed in more detail in [Chapter 6](#)). Based on a fit to non-post-flashed dark frames taken since launch, the median dark current (excluding hot pixels) is increasing by $\sim 0.5\text{ e}^-/\text{hr}/\text{pixel}/\text{year}$ and is currently at $\sim 11\text{ e}^-/\text{hr}/\text{pixel}$ in 2024 ([Figure 5.2](#)). The number of permanent hot pixels, i.e., pixels that the anneals are unable to restore to normal levels, are growing by about 35/chip/day, or 0.05-0.1% per month.

Figure 5.2: UVIS hot pixels and dark current over time

Top row: Percent of UVIS hot pixels ($>0.015 \text{ e}^-/\text{s}$) as a function of time for Chip 1 (left) and Chip 2 (right). Bottom row: median UVIS dark current (e^-/hr) as a function of time. The changes to the post-flash levels for UVIS dark frames are marked by vertical lines. The plots include data up to March 31, 2024.



Several events led to a temporary increase in the number of hot pixels flagged in the UVIS superdarks. Recoveries from safe mode of either the instrument or observatory near the end, or at the beginning, of an anneal cycle can cause apparent elevated numbers of hot pixels in the last superdarks during which the safing occurred, or the first superdarks immediately following recovery from safe mode.

5.3.3 Pixel Stability

The dark current present in WFC3 UVIS can be used to monitor a pixel's behavior and determine whether the pixel varies excessively over time compared to its expected noise. In 2018 we began using a data structure known as a 'pixel history' image to efficiently evaluate each WFC3 UVIS pixel over time. The pixel stability is analyzed in steps, anneal cycle by anneal cycle, because the population of hot pixels changes after each anneal procedure. The goal of the analysis is to identify and flag pixels that vary significantly in dark current over an anneal cycle compared to their expected uncertainty; such pixels are deemed "unstable" and thus more unreliable compared to their stable counterparts. Conversely, we also aim to identify pixels that are stable despite being flagged as hot pixels; since their variance is low and their uncertainties are well calculated, such pixels need not be discarded during nominal calibration of hot pixels. The metric for "stability" is determined from a ratio defined as "F", hereafter referred to as the stability ratio (as described in [WFC3 ISR 2018-15](#)):

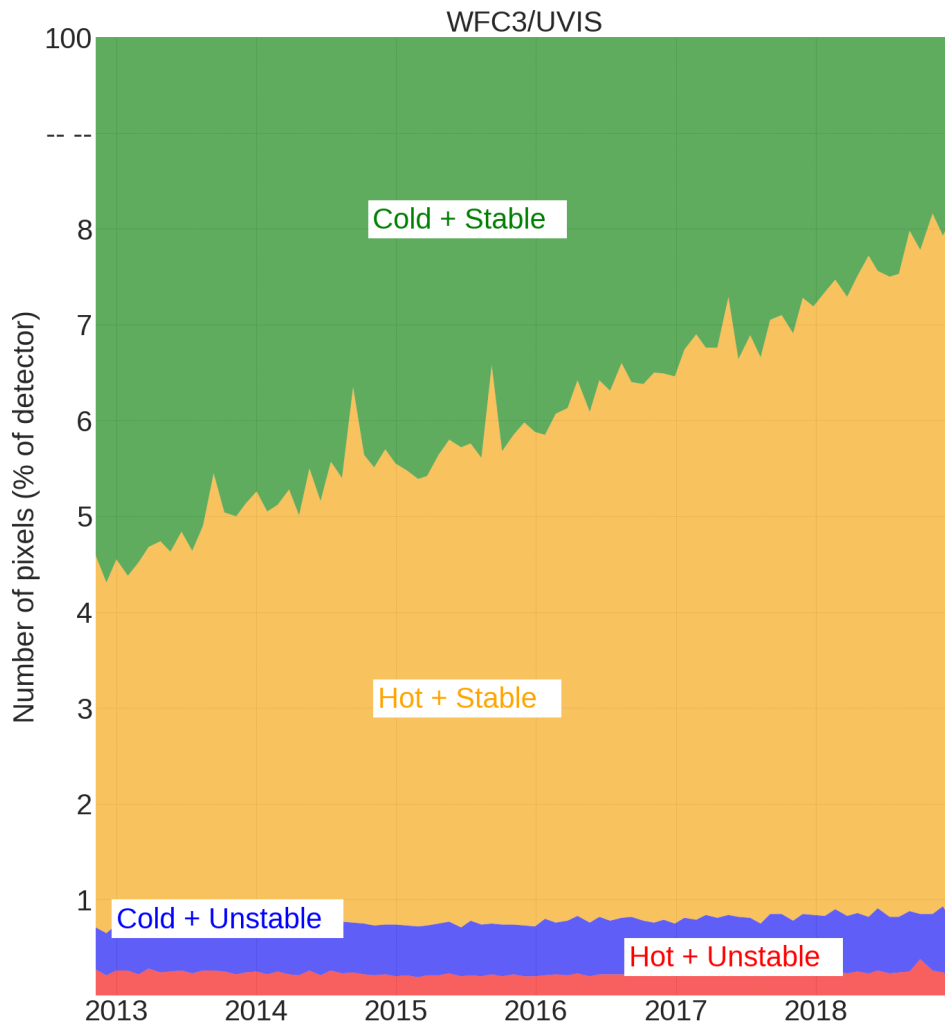
$$(1) \quad F = \frac{\text{Variance}(SCI) - \text{Mean}(ERR^2)}{\text{Mean}(SCI)} + 1$$

where SCI is the science value of the pixel and ERR is the uncertainty of the pixel value. Using the stability ratio, we define four classes of pixels that we monitor each anneal cycle:

Classification Name	Definition
Cold and Stable	Mean(SCI) < 54 e-/hr and F < 2
Cold and Unstable	Mean(SCI) < 54 e-/hr and F ≥ 2
Hot and Stable	Mean(SCI) ≥ 54 e-/hr and F < 2
Hot and Unstable	Mean(SCI) ≥ 54 e-/hr and F ≥ 2

The vast majority of hot pixels (~4% of the detector in 2021) are found to be stable and can be calibrated. A small portion of pixels (<1.5%) are genuinely unstable and should not be used. The data quality flag 32 is used to mark the unstable pixels in the bad pixel reference files, and hot pixels continue to be flagged with the usual value of 16. In [Figure 5.3](#) we show a breakdown of each of the four classes of pixels for each anneal cycle from the November 2012 anneal through 2019. The y-axis is a broken logarithmic axis to simplify the visualization. Approximately 90% of the distribution consists of cold stable pixels. While the hot stable pixels would normally be discarded in science observations via the hot pixel flag in UVIS dark reference files, such pixels need not be thrown away since they do not vary significantly throughout the anneal cycle relative to their expected uncertainty. As a result, ~4-8% (depending on epoch) of the detector pixels are sufficiently stable to be calibrated in science observations rather than discarded. The prior bad pixel methodology would have treated the cold and unstable pixels as nominally 'good'; such pixels are now flagged and provided in the bad pixel tables for users to choose whether or not to use them. More information about the UVIS pixel stability and bad pixel tables is outlined in [WFC3 ISR 2018-15](#) and [WFC3 ISR 2021-14](#).

Figure 5.3: The number of pixels, as a percentage of the detector, in each of the categories for each annual cycle from the November 2012 annual cycle through the December 2018 annual. Figure is from [WFC3 ISR 2018-15](#).



5.4 UVIS Flat Fields

[5.4.1 Ground Flats](#)


[5.4.2 On-orbit Corrections](#)

[5.4.3 Chip-Dependent Flats](#)

The UVIS flat fields provide a 2-D spatial map of the detector response and represent the combined throughput of multiple components along the optical path. The flat fields include both a high spatial frequency component that accounts for variations in the pixel-to-pixel response (P-flat) and a low spatial frequency component (L-flat) that corrects for large-scale modulations across the detector and as such are often referred to as LP-Flats.

The flat field reference files in use during WFC3's commissioning in 2009 (t*_pfl.fits) were derived from ground test data ([WFC3 ISR 2008-46](#)), and these flats were found to contain a large internal window reflection affecting ~45% of the detector field of view at a level of 1-2%. A simplified geometric model was used to remove this feature from the ground flats.

Because of geometric distortion, the area of the sky seen by a given pixel is not constant; therefore, observations of a uniform 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 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 falling on different portions of the detector will not have the same total counts in a calibrated (*_flt.fits, *_flc.fits) image. To correct for this effect, point source photometry extracted from flat-fielded images should be multiplied by the effective pixel area map (see [Section 9.1.11](#)). Alternatively, this correction is accounted for in the pipeline by AstroDrizzle (see [Section 4.2](#)), where the distortion reference files are used to correct all pixels to equal area on the sky for drizzled (*_drz.fits, *_drc.fits) data products.

 **Because of geometric distortion, the area of the sky seen by a given pixel is not constant; therefore, observations of a uniform 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 include the effect of the variable pixel area across the field.**

Corrections for additional low-frequency components, due to differences between the ground and in-flight optical paths, were computed from on-orbit photometry of bright stars in Omega Centauri dithered across many positions on the detector ([WFC3 ISR 2013-10](#)). In-flight solutions (v*_pfl.fits) were delivered on December 14, 2011 for the 42 full-frame filters (excluding the QUAD filters, which currently still use the ground flats). Flats for 2 × 2 and 3 × 3 binned modes (w*_pfl.fits) were delivered on August 29, 2012 (WFC3 TIR 2012-04).

Finally, to support the UVIS chip-dependent photometric calibration ([WFC3 ISR 2016-03](#)), an updated set of flat fields (z*_pfl.fits) were delivered on February 23, 2016. These include revised L-flat corrections based on CTE-corrected Omega Centauri photometry and a temperature-dependent correction for the UV filters. The photometric repeatability of white dwarf standards placed at different areas (or stepped) across the detector field of view is now better than 1% r.m.s. and 3% peak-to-peak over the full wavelength range of the detector ([WFC3 ISR 2016-04](#), [WFC3 ISR 2016-05](#)).

The following three sections describe in more detail the development and processing which has gone into generating the UVIS flatfields currently in use in the MAST automated calibration pipeline.

5.4.1 Ground Flats

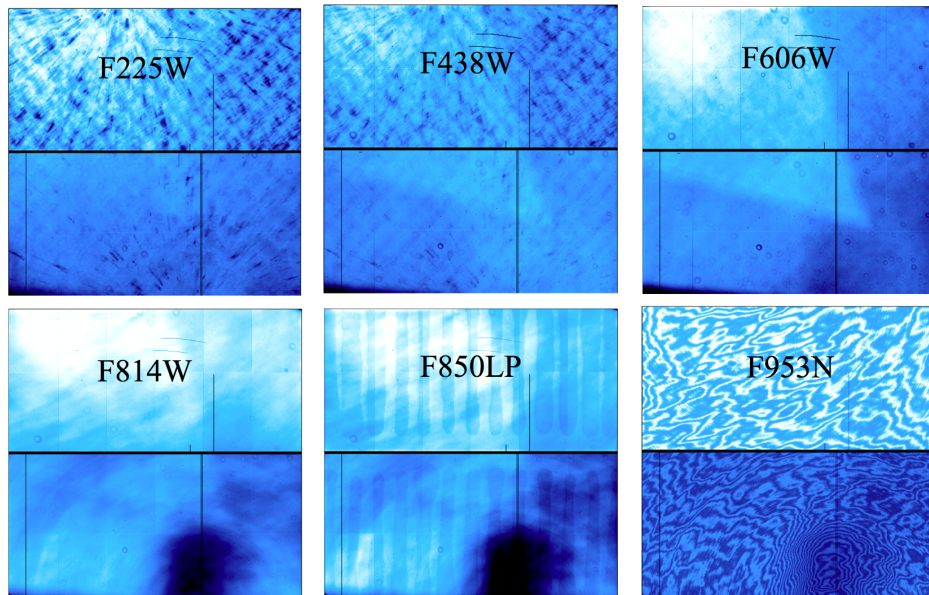
In the spring of 2008, flat fields for the UVIS channel were produced at the Goddard Space Flight Center during the third and final thermal vacuum campaign (TV3) using the CASTLE Optical Stimulus (OS). CASTLE (Calibration Stimulus from Leftover Equipment) is an HST simulator capable of delivering OTA-like, external, monochromatic point source and broadband full field illumination. Flats were acquired in only the standard CCD readout configuration of four amplifiers (ABCD), gain=1.5 and binning=1 × 1. Flat fields with the OS Xenon and Halogen lamps were taken with the detector at its nominal operating temperature of -82°C. Flats for the ultraviolet (UV) filters were acquired at a warmer detector temperature (-49°C) using the deuterium lamp to achieve higher count rates.

With a total signal of at least 75,000 electrons, the UVIS ground flats achieve an r.m.s. error better than 0.4% per pixel. The flats were normalized by the median value in a 100 × 100 pixel region in UVIS1 amplifier A, coordinates [1032:1133,328:429]. This region was selected to avoid the small dark rings ("droplets") that are spread across the UVIS field of view. The droplets are likely mineral residue on the outer window of the flight detector caused by a condensation event that occurred before TV3 ([WFC3 ISR 2008-10](#)). About one-third of these droplets moved in a coherent way during launch ([WFC3 ISR 2009-27](#)).

Unlike the ACS/WFC detector, the two UVIS chips were not cut from the same silicon wafer and should be considered as two independent detectors. For the 42 full-frame filters, the UVIS2 flats were divided by the UVIS1 normalization value in order to preserve the overall sensitivity difference between chips. In this way, calibrated flat-fielded (*_flt.fits, *flc.fits) or drizzled (*_drz.fits, *drc.fits) data products only required the use of a single zeropoint value for each filter. For the QUAD filters, on the other hand, each quadrant was normalized using a 100 × 100 pixel box in the corresponding amplifier, and therefore each carries its own unique zeropoint.

A subsample of ground test UVIS flat fields are shown in [Figure 5.4](#). A complete set may be found in the 'UVIS CASTLE Photometric Filter Flat-Field Atlas' ([WFC3 ISR 2008-46](#)). On-orbit corrections to the flat fields are described in the next two sections.

Figure 5.4: Flat fields obtained during ground testing for a subset of UVIS filters, where dark regions correspond to lower response pixels. Notable features include a crosshatch pattern at wavelengths less than 500 nm due to structure in the detection layer, a large diffuse spot in UVIS2 for wavelengths greater than ~ 600 nm due to variations in the detector thickness, vertical striping in the long pass filters which see through the detector to glue adhesive on the other side of the CCD, and fringing in narrow band filters at wavelengths greater than ~ 600 nm.



5.4.2 On-orbit Corrections

Due to differences in the optical path, flat fields acquired in flight were expected to differ from those obtained during ground testing, since neither the laboratory flat-field illumination nor the on-orbit internal lamp flats provides an accurate simulation of the OTA. Corrections to the spatial response were characterized by observing dense star clusters with multiple telescope orientations and large dithered steps across the detector ([WFC3 ISR 2009-06](#)). Placing the same stars over different portions of the detector and measuring relative changes in brightness, local variations in the response (L-flats) were computed.

Initial on-orbit data confirmed that the ground flats did not fully correct low- and mid-frequency structure, and SMOV calibration program [11452](#) (UVIS Flat-Field Uniformity) showed differences of ~ 1.5 to 4.5% in a subset of 6 filters based on observations of two galactic globular clusters (see [WFC3 ISR 2009-19](#)). Supplemental observations were obtained in Cycles 17 and 18 (programs [11911](#) & [12339](#)) in a key set of 10 broadband filters used most frequently by observers: F225W, F275W, F336W, F390W, F438W, F555W, F606W, F775W, F814W, and F850LP.

The same mathematical algorithm for deriving ACS inflight corrections ([ACS ISR 2003-10](#)) was used to characterize the accuracy of the UVIS ground flats. Aperture photometry was performed on the cluster data with a radius of 0.2" (5 pixels) to minimize uncertainties due to crowding. Then, a spatially variable aperture correction to 0.4" (10 pixels) was applied to each image to account for variations in the PSF with detector position and telescope focus, including temporal effects such as breathing and long term focus drifts ([WFC3 ISR 2013-11](#)). Because of spatial and temporal variability in the UVIS encircled energy, the aperture corrections for radii smaller than 0.4" (10 pixels) are most accurate when computed directly from bright stars in each image.

The first set of L-flat corrections showed evidence of an extended wedge-shaped artifact in the ground flats (see [Figure 5.5](#) where the wedge is most prominent in the F606W flat). This same feature was noticed in on-orbit flats obtained by observing the dark side of the Earth during periods of full moon illumination (programs [11914](#) and [12709](#)). This extended feature, dubbed the UVIS 'flare', is a result of the tilted UVIS focal plane, where light is reflected multiple times between the detector and the two chamber windows. When a point source is positioned in the lower right quadrant of the UVIS detector, out of focus reflections between the CCD and the two windows appear along a diagonal from the source towards the upper left (see [Section 5.5.1](#)). When the detector is illuminated by a uniform surface brightness source, the summation of defocused ellipses creates the wedge shaped 'flare' apparent in both the ground flats and the on-orbit Earth flats. The L-flat residuals showed a negative imprint of the flare, implying that it is not a true feature of the L-flat, but instead an internal reflection in the ground flat which is imprinted on the Omega Centauri images during the flat fielding process.

A geometric model ([WFC3 ISR 2011-16](#)) was used to predict the shape and extent of the flare and the relative strength of the four internal reflections with respect to the primary source (see [Figure 5.5](#)). The absolute strength of the flare as a function of wavelength was estimated from the preliminary L-flat solutions, and this was used to scale and divide the model from the ground flats. L-flats were then recomputed from cluster photometry recalibrated with flare-free ground flats. Sensitivity residuals for the 10 broadband filters are shown in [Figure 5.6](#), represented as a chessboard grid of basis functions of order $2n$, where the 4th and 5th order solutions result in a 16×16 and 32×32 pixelated version of the UVIS detector (with each grid pixel representing an independent solution). With the exception of the two bluest UV flats, which were obtained with the detector running warmer than usual during ground testing, the sensitivity residuals show a general wavelength dependence, where the required correction deviates from unity more at longer wavelengths compared to shorter ones.

Figure 5.5: Left: schematic diagram of Figure 8 ghosts, a co-added image of 37 individual frames obtained during ground testing prior to installation in HST. Each star in the lower right quadrant produces four Figure 8 ghosts to the upper left. Color-coding links each star to its corresponding ghosts (from [WFC3 ISR 2011-06](#)). Right: geometric model of the UVIS flare as a set of 4 internal reflections for a source that uniformly illuminates the detector.




Figure 5.6: Low-frequency corrections to the flare-corrected ground flats for 10 filters. Blue indicates regions of the detector where the flats over-corrected the data, making the observed photometry too faint. The solutions are presented as a 32×32 grid, such that one grid pixel corresponds to 128×128 detector pixels. The exceptions are F225W, F275W, and F850LP for which 16×16 grid solutions were used, and one grid pixel corresponds to 256×256 detector pixels.





Figure 5.7: Ratio of the 2011 pipeline flats to the 2009 ground flats. Blue indicates that the ground flats overcorrect the data, making sources too faint. The residuals include contributions from both the flare and the stellar L-flat.



Rather than carry a separate set of L-flat reference files, revised LP-flats were created by multiplying the ground P-flats (flare removed) with smoothed versions of the pixelated L-flat. Figure 5.7 shows the ratio of the improved 2011 pipeline flats with respect to their 2009 solutions for the 10 broadband filters with on-orbit calibration data. The 2011 solutions differ by 0.6 to 1.8% rms, with a maximum change of ~ 2.8 to 5.5%, depending on filter (WFC3 ISR 2013-10).

For the remaining 32 full-frame filters, flat-field corrections (including both the flare and the stellar L-flat) were computed via interpolation based on the filter pivot wavelength, using the weighted average of the residual flat for the two filters nearest in wavelength. The wavelength dependence of both the flare and the L-flat in Figures 5.6 and 5.7 suggests that interpolation is a valid approximation for filters with no available on-orbit cluster data.

To quantify the accuracy of the new flat fields, white dwarf standards were stepped across the detector field of view to measure their photometric repeatability. For filters with pivot wavelength greater than 300nm, photometry in a 10-pixel aperture was repeatable to better than 0.7% rms and 2.7% peak-to-peak. The flux residuals showed a weak correlation with position (e.g. the number of parallel transfers required for each source during read-out), suggesting that some of the spatial variation was due to CTE losses during readout at a level of 0.5-1.0%. For UV filters, the photometric residuals from stepped data calibrated with the 2011 flats were 1.8% rms and 6.7% peak-to-peak (WFC3 ISR 2015-18).

The 2011 in-flight flats corrected for sensitivity offsets between the two chips using dithered stellar photometry in Omega Centauri, but the average cluster population is significantly redder than the white dwarf used to derive the zeropoints. When dividing the cluster stars into separate populations by their intrinsic color, differential photometry of sources dithered between the chips showed offsets in the measured count rate that correlated with the source's intrinsic color. For the UV filters, this offset was as large as 5% in F225W for blue versus red populations, though red leaks ([Section 6.5.2](#) of the *WFC3 Instrument Handbook*) limit the accuracy of this prediction at the 1-2% level.


5.4.3 Chip-Dependent Flats

The quantum efficiencies of the two UVIS chips differ at wavelengths shorter than about 400 nm (see [Section 5.4.1](#) of the *WFC3 Instrument Handbook*). Motivated by this difference, a chip-dependent approach to the photometric calibration was adopted in February 2016 to better track any changes with time as the two detectors age. Prior to this date, sensitivity offsets between chips were derived from stars in Omega Centauri dithered between UVIS1 and UVIS2, as described in [Section 5.4.2](#), and these were accounted for by **calwf3** as part of the flat fielding step (`FLATCORR`). Photometric keywords such as `PHOTFLAM`, the inverse sensitivity, were then populated in the image header via the `PHOTCORR` photometry switch, where a single value of `PHOTFLAM` applied to both UVIS1 and UVIS2.

In the chip-dependent approach, the L-flats were computed from CTE-corrected star cluster data dithered over a single chip only. The LP-flats were then normalized by the median value of each chip. To ensure that count rate photometry in calibrated data remains continuous across the two chips, **calwf3** (starting with version 3.3+) made use of a new calibration switch (`FLUXCORR`), which multiplies the UVIS2 science extension by the sensitivity ratio (`PHTRATIO`) of the two chips (see [Section 9.1](#) for details). This ratio is empirically derived from photometry of white dwarf standards measured for UVIS1 and UVIS2 separately as part of the UVIS photometric calibration program ([WFC3 ISR 2016-03](#)).

For the majority of filters, the improvement in the chip-dependent flats is small and is primarily attributed to correcting for CTE losses in the star cluster data used to determine the L-flat. [Figure 5.8](#) shows the ratio of the CTE-corrected L-flat to the 2011 L-flat for six broadband filters, where the total correction from the center to the edge of each chip is on average 0.6%.


Figure 5.8: Ratio of the 2016 to 2011 L-flat (low spatial frequency) correction, showing the change after applying CTE-corrections to the Omega Centauri calibration data. The color bar spans a range of 1%.



Flats for the four bluest UVIS filters (F218W, F225W, F275W, F280N) include an additional correction to the spatial response with temperature. Due to ground equipment limitations, these flats were obtained at a warmer detector temperature (-49°C) during ground testing and have now been corrected with in-flight data (acquired at -82°C). White dwarf standard stars stepped across the two chips showed flux residuals which correlate with a crosshatch pattern in the UV flat fields at spatial scales of $\sim 50 - 100$ pixels. These data were used to derive a linear correction model which improves the spatial repeatability from $\sim 7\%$ to $\sim 3\%$ peak-to-peak ([WFC3 ISR 2016-05](#)). The ratio of the 2016 to 2011 flats for the F275W filter is shown in [Figure 5.9](#), where the crosshatch correction is more pronounced for UVIS1 (top chip).

The UVIS chip-dependent calibration is intended to produce uniform count rates across the two chips for sources similar in color to the white dwarf standards. As discussed in [Section 5.4.2](#), count rate offsets of several percent are found in the UV for red versus blue stars, but these appear to impact only the photometric zeropoints and not the flat fields themselves. Monochromatic flats obtained during ground testing with an optical stimulus at wavelengths 310, 530, 750, and 1250 nm showed no evidence of spatial residuals in the flat ratio with wavelength.

Figure 5.9: Flat field ratio 2016/2011 for the F275W filter, with UVIS1 on top and UVIS2 on the bottom. The color bar for each chip spanning a total range of 4%, but offset by 5% to account for the chip sensitivity offset in this filter, with UVIS2 being more sensitive. The 2016 UV flats correct for mid-frequency residuals in a crosshatch pattern that are a result of the colder in-flight operating temperature.



5.5 Image Anomalies

[5.5.1 Ghosts](#)

[5.5.2 Dragon's Breath and Scattered Light](#)

[5.5.3 Crosstalk](#)

Some UVIS images may contain features that are not direct images of astronomical sources. The causes of these features include multiple reflections between optical surfaces (detector, filters, and windows) of light from the astronomical scene, scattered light from bright sources outside the detector FOV, light from the bright Earth that is scattered in the OTA, and electronic cross talk between readout amplifiers. In general, these artifacts are not calibrated and cannot be removed by the WFC3 pipeline.

Examples of the different kinds of image anomalies can be found on the [WFC3 website](#) as well as in [WFC3 ISR 2017-22](#). Two new XML overlay files have been added for users to explore the potential for the Dragon's Breath and Figure 8 ghost artifacts on UVIS images (see [STAN 46](#) for more information). It can be loaded through the FOV tab in the Aladin file viewer, on top of the planned visit. Users can then tweak the overlay or their observation ORIENT to explore whether particular ORIENT restrictions might compromise their science results.

5.5.1 Ghosts


Ghosts appear as images of the pupil formed from the light of a bright target in, or near, the UVIS detector FOV. The target light is scattered twice (or more) by optical surfaces forming one (or more) out of focus image(s). The separation of the ghost from the source depends upon the separation of the scattering surfaces and the angle of scattering. Filter ghosts are formed by scattering of the near-normal source light at the surfaces of a filter and, as such, are found close to, or overlapping, the source image. Further details concerning these ghosts may be found in [WFC3 ISR 2007-09](#). Some examples of the ghosts can be seen in [Figure 5.10](#).

Scattering between the UVIS CCD and either the detector, or dewar window forms so-called 'optical' ghosts. These ghosts are separated by $\sim 80''$ from their source generally in pairs (a "figure eight"; see the right panel of [Figure 5.10](#)). Further details may be found in the following WFC3 ISRs; note that documents released before 2009 were based on ground test data. The reports include [WFC3 ISR 2001-17](#) (ghost modeling), [WFC3 ISR 2004-04](#) (initial filter ghosts, most fixed before flight), [WFC3 ISR 2007-09](#) (filter ghosts after filter replacements with overall performance of those replacement filters in [WFC3 ISR 2007-01](#)), [WFC3 ISR 2007-21](#) (channel glint in spare UVIS detector ground-testing images, **not** present in on-orbit images) and [WFC3 ISR 2011-16](#) ("figure eight" ghosts).

Figure 5.10: Ghosts in the WFC3/UVIS detector

Examples of filter-detector ghosts and optical ghosts in the WFC3/UVIS detector:

- Left: a filter-detector ghost (from dataset `ibsq56igq`) due to the bright star Fomalhaut in the center.
- Middle: A filter ghost (indicated by the arrow) formed near the source (`ic9t32gbq`).
- Right: An optical ghost known commonly as a figure-eight. Multiple figure-eights are seen in the image caused by two sources in the field (`ibvk02ccq`).




5.5.2 Dragon's Breath and Scattered Light

Bright sources immediately outside the image frame on WFC3/UVIS can create (mostly minor) light anomalies in the form of spurious signal redirected into the field of view and affecting the image (see [Figure 5.11](#)). Dragon's Breath (when the light originates from the edge) or Scattered Light (when the light is somewhat displaced from the edge) affects about 12% of a subset of images which include those taken in full-frame wide and long pass filters with exposure times ≥ 300 seconds ([WFC3 ISR 2017-02](#)). It occurs when a bright source lies within a $30''$ ring around the field of view, and peaks when the neighboring source has a V band magnitude around 11.9 or 14.6 ([interactive visualization of known dragon's breath sources](#)). The cause for the scattering along the optical path is not completely understood. Most cases of Dragon's Breath and Scattered Light are low-level features and will not adversely affect a science image. Note that dithering, done by the vast majority of observers, results in either a shift or sometimes complete disappearance of the feature.

Figure 5.11: Dragon’s Breath and Scattered Light on WFC3/UVIS

Examples of Dragon’s Breath and Scattered Light on WFC3/UVIS. Left: Dragon’s Breath spills into the image frame from the bottom of the image. Right: Scattered light projects into the center of the image from a source just off the bottom of the FOV. These are worst-case examples shown with an extreme image stretch; most instances of Dragon’s Breath and scattered light are much less dramatic both in number of affected pixels and flux from the effect.



5.5.3 Crosstalk

Generally whenever two or more quadrants of an electronic device are read out simultaneously, there is a chance of crosstalk (Janesick 2001) and in fact, both the UVIS and IR channels in WFC3 exhibit some very low-level crosstalk (CT). In the UVIS, amps A+B and amps C+D are coupled, thus bright sources in A (or B, C, D) will give rise to CT in B (or A, D, C) respectively. In the IR channel, the CT is also a low-level mirror image although in this case, the coupled amps are 1+2 (upper left and lower left, when image is displayed with $x=1,y=1$ at lower left) and 3+4 (upper right and lower right). In both channels, the CT appears as a negative image; thus, these electronically induced features are unlikely to be confused with e.g. optical ghosts.

Figure 5.12 shows a sample UVIS image with CT in both amplifiers due to multiple bright sources. The UVIS CT is linear, negative, and appears at the level of $\sim 10^{-4}$ to 10^{-5} of the source. Specifically, in full-frame, unbinned UVIS readouts, the CT level is $\sim 2 \times 10^{-4}$ of the source when the target is in quadrants A or C and about 8×10^{-5} when the target is in quadrants B or D (WFC3 ISR 2009-03).

Crosstalk in the UVIS channel only occurs in the chip containing the target, it does not cross between chips. To within the errors, the CT due to hot pixels and cosmic rays is the same as for point or extended sources. Dithering of observations can help mitigate the low-level effects of CT: the mirror image nature of the CT moves the features in a direction opposite to the target motion, i.e., they will appear to be transients and thus be removed during the drizzling procedure.

A standalone procedure is available for correcting UVIS data for crosstalk (the automated pipeline does not remove the effect). Pixels are effectively corrected to a mean which is well within 1-sigma of the mean of surrounding pixels. The code is described in [WFC3 ISR 2012-02](#) and the software ([IDL](#) and [python version](#)) are available for use.

[Figure 5.13](#) shows a subsection of amplifier B from the image in [Figure 5.12](#) and the corresponding crosstalk corrected image.

Figure 5.12: UVIS Crosstalk

UVIS crosstalk for chip 1 (dataset iaby011eq). Circles show bright stars in amplifier A causing crosstalk in amplifier B and squares show bright stars in amplifier B causing crosstalk in amplifier A.





Figure 5.13: UVIS Crosstalk Correction

Subsection of amplifier B (dataset iaby01leq) showing the original (left) and the CT-corrected (right) image sections.



5.6 Generic Detector and Camera Properties

- [5.6.1 Full Well Depth](#)
- [5.6.2 Linearity at Low to Moderate Intensity](#)
- [5.6.3 Linearity Beyond Saturation](#)
- [5.6.4 Shutter Stability](#)
- [5.6.5 Fringing](#)

5.6.1 Full Well Depth

Conceptually, full well depths can be derived by analyzing images of a rich starfield 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, it is also necessary to correct for a ~10% "piling up" effect of higher values being reached at significant levels of over-saturation relative to the value at which saturation and bleeding to neighboring pixels in the column begins (see [WFC3 ISR 2010-10](#)).

Since the full well depth varies over the CCDs, it is desirable to observe a rich star field with a gain that samples the full well depth (e.g. the default WFC3 UVIS gain), and for which a large number of stars saturate. Calibration and GO programs have serendipitously supplied the requisite data of rich fields observed at two different exposure times.

There is a real and significant large-scale variation of the full well depth on the UVIS CCDs. The variation over the UVIS CCDs is from about 63,000 e⁻ to 72,000 e⁻ with a typical value of about 68,000 e⁻. There is a significant offset between the two CCDs, as visible in Figure 5 of [WFC3 ISR 2010-10](#).

Prior to version 3.7.1, the **calwf3** calibration pipeline used a single scalar threshold based on the average of these values (`SATURATE`, in the `CCDTAB` file) for full well to flag saturated pixels in the science image data quality file. Installed in December 2023, **calwf3** version 3.7.1 adopted a new method of UVIS full well saturation flagging that uses a two-dimensional array of pixel threshold values instead of a single scalar threshold applied to the entire image. This is now a deliverable saturation map reference file (`SATUFILE`) in CRDS (see [Section 3.2](#) for where it is implemented in the pipeline. The first delivery of the `SATUFILE` still effectively applies the same scalar threshold ([WFC3 ISR 2023-08](#)). However, the change lays the foundation for a future delivery of a spatially-dependent saturation map to account for the variation in full well depth across the detectors.

5.6.2 Linearity at Low to Moderate Intensity

Linearity at low and moderate exposure levels was measured by comparing back-to-back exposures of NGC 1850 made in October 2009. [Figure 5.14](#) shows the response of UVIS2, where aperture sums for stars with flux greater than about 2,000 e⁻ in a short exposure (central pixel greater than about 350 e⁻) show apparently perfect linear response when compared to the counts in the same aperture in an exposure 50 times as long. However, below a level of 2,000 e⁻ the ratio of long to short exposure counts deviates from a linear response. At a total aperture flux of about 200 e⁻ in the short exposure, the total flux values are ~ 5% lower than expected based on scaling from the corresponding long exposure. This is due to the nature of CTE loss which has a greater effect on fainter sources.

[Figure 5.15](#) shows data from both UVIS CCDs for stars with short exposure aperture sums of 500 to 2000 e⁻. Stars near the readout amplifiers ($y_{\text{pix}} = 0,4095$) behave as expected i.e. the short exposure sums scaled by the difference in exposure time are equivalent to the long exposure but losses grow with distance from the amplifiers. This effect is consistent with losses induced by finite charge transfer efficiency in successive parallel shifts in clocking the charge packets off the CCDs.

The WFC3 team monitors the extent to which CTE losses influence faint object photometry. CTE degradation in flight has increasingly affected the measurement of unsaturated images over the years, requiring increasingly greater flux corrections with larger uncertainties. *The most recent CTE loss results are summarized in [Chapter 6](#).*

Figure 5.14: Loss of linearity at low and moderate intensity due to finite Charge Transfer Efficiency.

The upper panel shows the counts in $r = 3$ pixel apertures for several thousand stars in back-to-back 500s and 10s exposures (from 2009) plotted against each other. To better illustrate small differences, the lower panel shows the same data after binning and plotting as the ratio of long to short exposure time; 50.0 denotes perfect linearity.





Figure 5.15: UVIS1 and UVIS2 CTE Losses For a Short Exposure

Ratio of fluxes in short and long exposures as a function of Y position. Plot is based on data from the previous figure for UVIS2, plus similar data for UVIS1 plotted against y-position in the concatenated detector space (images from 2009). Y-values near 2050 correspond to the maximum distance from the readout amplifiers, and hence the most parallel shifts inducing CTE losses.



5.6.3 Linearity Beyond Saturation

The response of the WFC3 UVIS CCDs is linear well beyond the point of saturation. [WFC3 ISR 2010-10](#) shows the well-behaved response of WFC3: electrons are clearly conserved after saturation -- in some locations with the need for a minor calibration, as provided in the ISR, while in other regions no correction is needed. This result is similar to those of the STIS CCD ([Gilliland et al. 1999](#)), the WFPC2 camera ([Gilliland, 1994](#)), and the ACS ([ACS ISR 2004-01](#)). It is possible to easily perform photometry on point sources that remain isolated simply by summing over all of the pixels into which the charge has bled.

To characterize the accuracy of point source photometry for sources in which one or more pixels have exceeded the physical full well depth, we used a dataset consisting of alternating short and long exposures of a moderate-to-rich star field, resulting in both unsaturated and saturated images for many stars. Counts were summed over an area designed to include all pixels in the core of the stellar image and all pixels that were bled into in the long exposures.

Departures from linearity were greater for UVIS1 than for UVIS2. Results for Amp A and UVIS1 are summarized in [Figure 5.16](#) and [Figure 5.17](#). Over a range of nearly 7 magnitudes beyond saturation, photometry remains linear to $\sim 1\%$ after a simple correction ([WFC3 ISR 2010-10](#)). For Amp C on UVIS2, the response is sufficiently linear beyond saturation that no correction is required. These results were obtained when CTE was near pre-flight levels in October 2009. See [Section 5.6.2](#) for the level of CTE loss that was already measurable at that time. CTE degradation in flight has increasingly affected the measurement of unsaturated images over the years, requiring increasingly greater flux corrections with larger uncertainties.

While these results are based on the use of *_flt.fits images, the flux conservation ensured by **AstroDrizzle** leads to equally good results for linearity beyond saturation when comparing long and short *_drz.fits images.

Figure 5.16: Linearity Analysis for Amp A

Linearity beyond saturation for Amp A on UVIS1 in data from 2009. The upper panel shows the ratio of counts in the long exposure divided by the counts in the short exposure multiplied by the relative exposure time; a value of unity indicates perfect linearity. The x-axis shows the multiplicative degree by which a star is over-saturated in the long exposure. The middle panel shows the long exposure aperture sums versus short exposure aperture sums. The lower panel shows the peak data value in the long exposure relative to the short exposure value. The response is linear up to 68,000 e⁻ (y-axis) where the long exposure encounters saturation. Amp A data show significant deviations from a linear response for over-saturations near and beyond a factor of 10.





Figure 5.17: Long vs. Short Exposure Ratios of Linearity

Top panel shows the count rate ratio between long and short exposures for both amplifier quadrants of UVIS1 (Amps A and B) vs. the over-saturation level in the long exposure in data from 2009. Bottom panel shows the same data but after applying the corrections as described in [WFC3 ISR 2010-10](#). At high levels of saturation, the Long to Short count ratios are restored to unity and the scatter is significantly reduced.



5.6.4 Shutter Stability

The WFC3-UVIS shutter is a circular, rotating blade divided into two open and two closed quadrants (See [Section 2.3.3](#) of the [WFC3 Instrument Handbook](#) for details). Operationally, the shutter mechanism has two distinct modes, based on the commanded exposure time. At the shortest commanded exposure time of 0.5 seconds, the shutter's motion is continuous during the exposure, rotating from the closed position through the open position and on to the next closed position. For commanded exposure times of 0.7 seconds and longer, the shutter rotates into the open position, stops and waits for an appropriate amount of time, and then rotates to the closed position.

Using series of short exposures during SMOV, the “shutter shading,” i.e. position-dependent exposure time, shutter blade dependence (A blade or B blade), shutter timing repeatability, and shutter timing accuracy were all assessed (see [WFC3 ISR 2009-25](#) and [WFC3 ISR 2015-12](#)). Shutter shading does not exceed $\sim 0.2\%$ across the detector, so no correction for shutter shading is performed in `calwf3`.


Stellar PSFs in short exposures taken with shutter B show slightly larger FWHM and lower sharpness compared to those taken with shutter A, especially in the shortest, 0.5-sec, exposures ([WFC3 ISR 2014-09](#)). In order to facilitate the science where short exposures are necessary and a good and stable PSF is crucial, use the available-mode option in in APT to select the shutter blade A. No other systematic difference in shutter behavior (exposure time, repeatability, etc.) has been found when comparing the A and B blades of the shutter.

The stability of shutter timing is a bit more problematic. Very early on-orbit results were based on eleven white dwarf standard star images; for exposure times of 0.5, 0.7, 0.8, and 1.0 seconds, the rms variation was measured at 1.9%, 1.6%, and 1.5%, and 1.0%, respectively ([WFC3 ISR 2009-25](#)). These results for direct images must include photometric errors from variations in pointing, PSF variations, and flat field errors in addition to shutter timing variation thus the results are upper bounds to the variation in shutter timing alone.

In a more recent study, a spectrophotometric standard star is observed using the G280 grism in order to disperse its light across hundreds of pixels, allowing for high photometric precision in very short exposures to more precisely measure the shutter timing jitter ([WFC3 ISR 2023-04](#)). The actual exposure time was found to vary by 2.43 ± 0.32 milliseconds for exposure times of 1.0, 2.0, and 4.0 seconds, which is nearly four times better than the system requirement and a much smaller rms variation measured in prior analyses. The results also show that the repeatability for 0.5 second exposures is ten times smaller than in longer exposures, which is most likely due to the shutter mechanism not stopping during its open state. Conversely, the timing jitter in 0.7 second exposures is approximately twice that measured in longer exposures, though still within the system requirements. The 0.7 second exposures are most likely affected by the vibration of the shutter system, leading to the somewhat higher variations in the shutter timing.

- ✔ In short exposures, the shutter jitter dominates Poisson noise as shown in [Figure 5.18](#). Observers should consider avoiding 0.7 second exposures due to its greater shutter timing variation and use the 0.5 or 1.0 second instead when possible. For precision photometry with short exposures, the shutter jitter error should be included in the noise model ([WFC3 ISR 2023-04](#)).

Figure 5.18: Photometric noise sources for summed G280 spectra acquired with a variety of exposure times. Two theoretical components, shutter jitter and Poisson noise, are plotted as straight lines with power law indices of -1 and -0.5 respectively. Their root-sum-of-squares combination is also shown. Data from the positive orders in Table 1 of [WFC3 ISR 2023-04](#) are plotted (blue points) with associated 1- σ uncertainties. For comparison, the Poisson noise for a nearly saturated star in a direct image is $\sim 2 \times 10^{-3}$. Figure from WFC3 ISR 2023-04.



Photometric measurements indicate that for commanded exposure times of 0.5 s and 0.7 s, the shutter actually is open for 0.48 s and 0.695 s, respectively (cf. citations earlier in this subsection). The `EXPTIME` header keyword value in the science image headers reflects these actual (as opposed to commanded) exposure times.

5.6.5 Fringing

At wavelengths longer than about 650 nm, silicon becomes transparent enough that multiple internal reflections in the UVIS detector can create patterns of constructive and destructive interference, or fringing. Fringing produces wood-grain type patterns in images as a result of narrow-band illumination at long wavelengths.

The amplitude and phase of the fringes is a strong function of the silicon detector layer thickness ([Figure 5.19](#)) and the spectral energy distribution of the illumination. Detector layer thickness maps are used to model/predict fringing behavior, whose shape generally follows variations in thickness (see [Figure 5.19](#)). Fringe amplitude i.e. the contrast between constructive and destructive interference is greatest a) at the longest wavelength (high transparency allows for more internal reflections) and b) for the narrowest spectral energy distributions (SEDs). For broad SEDs, interference is averaged over phase, so that the amplitude of the fringing is reduced. Thus fringing is significant for UVIS imaging data only if red narrow-band filters (e.g. [Figure 5.20](#)) are used or sources with red line emission are observed.

Figure 5.19: Left: UVIS CCD detection layer thickness map based on TV3 data. For Chip 1, the median and standard deviation of thickness is $16.04 \mu\text{m} \pm 0.23 \mu\text{m}$, with 1% of pixels thinner than $15.51 \mu\text{m}$ and 1% thicker than $16.46 \mu\text{m}$. For Chip 2, the thickness is $15.42 \mu\text{m} \pm 0.58 \mu\text{m}$, with 1% thinner than $13.87 \mu\text{m}$ and 1% thicker than $16.19 \mu\text{m}$ (from [2010 HST Calibration Workshop](#), pg 194). Right: a quad filter (FQ889N in Amp A, FQ937N in Amp B, FQ906N in Amp C, and FQ924N in Amp D) flat field which shows the effects of fringing.





Figure 5.20: Fringing

Quadrant B of three flat fields: the broadband F606W control (left), one narrowband unaffected by fringing (middle), and one affected by fringing (right). Dark blue region of the FQ937N flat is masked to avoid areas impacted by the quad filter edge effects; the same region of F606W and F673N is masked for consistency.



Flat fields from ground tests ([WFC3 ISR 2008-46](#)) have been used to estimate the magnitude of fringing effects, for a continuum light source, in the narrow-band red filters (see [Table 5.3](#) and [WFC3 ISR 2010-04](#)). Each column lists a different metric of fringe amplitude for filters in which fringing effects could be detected in the flat-field data along with a control filter (F606W) without fringing. These metrics can best be understood by examining the histograms ([Figure 5.21](#)) of the flat fields shown in [Figure 5.20](#).

Table 5.3: Metrics of fringe amplitude based on ground flat fields


Filter	Distance Between Histogram Peaks percent	Manual Peak-to Trough percent	Quadrant	RMS Deviation percent	Full Width At 20% Maximum percent
F606W			A	0.9	2.9
F606W			B	1.0	3.0
F606W			C	1.2	3.3
F606W			D	1.2	3.3
F656N		2.2 +/- 1.2	A	1.4	4.9
F656N		1.7 +/- 1.1	B	1.3	4.5
F656N		1.7 +/- 1.2	C	1.6	5.1
F656N		3.2 +/- 1.3	D	1.5	5.1
F658N			A	1.3	4.4
F658N			B	1.2	3.9

F658N			C	1.4	4.4
F658N		0.9 +/- 1.1	D	1.3	3.7
F673N			A	0.9	3.1
F673N			B	1.0	3.2
F673N			C	1.3	3.6
F673N		0.5 +/- 1.1	D	1.4	4.2
F953N	17.8	19.8 +/- 1.7	A	7.6	23.7
F953N	17.3	20.8 +/- 1.8	B	7.8	24.0
F953N	14.9	17.8 +/- 1.7	C	7.1	23.5
F953N	15.0	11.5 +/- 1.6	D	6.7	22.9
FQ672N		4.6 +/- 1.3	D	2.0	7.0
FQ674N		2.4 +/- 1.3	B	2.1	7.4
FQ727N		2.3 +/- 1.1	D	1.6	5.3
FQ750N		1.2 +/- 1.1	B	1.3	4.5
FQ889N	7.8	10.0 +/- 1.3	A	3.8	13.5
FQ937N	12.5	12.2 +/- 1.4	B	5.4	18.2
FQ906N	6.9	10.1 +/- 1.4	C	4.4	16.5
FQ924N	9.4	14.2 +/- 1.4	D	5.0	18.4

Values are given in units of percentage of the normalized flat-field signal level. Each metric is described in the text and graphically represented in [Figure 5.20](#).

The fifth data column in the table is simply the root mean square deviation from the mean of the sample and is indicated by triangles with horizontal error bars in the histograms. Filters/quadrants with rms deviations greater than corresponding values for the control filter (F606W) may be influenced by fringing. The last column is full width at 20% maximum, rather than full width at 50% maximum because this metric is more effective for bimodal pixel brightness distributions in filters with strong fringing, such as FQ937N (pictured). The second data column gives the separation between histogram peaks, which can be detected in flat-field data for only the five reddest of the twelve filters affected by fringing. Squares in [Figure 5.21](#) mark the histogram peaks. Adjacent fringes were also manually sampled, and the results reported in the second data column.

Figure 5.20: Histograms of the three flat-field samples shown in Figure 5.17. Symbols correspond to fringe amplitude metrics listed in Table 5.3: rms deviation (triangles with error bars), full width at 20% maximum (circles with error bars), and bimodal histogram peaks (squares).



For estimating photometric uncertainties in data taken with these filters, it can be useful to consider how much the rms deviation in the flat fields departs from that of the control filter F606W. For exposure time calculations of targets affected by fringing, the full width at 20% maximum or the distance between histogram peaks may be the most useful metric. To understand how small dithers might affect photometry of targets that happen to fall on adjacent positive and negative fringes, the manually determined amplitude might be most appropriate. For sources with SEDs very similar to the calibration lamp, application of the pipeline flat fields should considerably reduce the effect of fringing on the data.

5.7 UVIS Photometry Errors

The photometric internal precision for the new (2020) time-dependent WFC3 UVIS inverse sensitivities, computed from observations of the four white dwarf standards GD153, GD71, G191B2B, GRW+70 5824, and the G-type star P330E are: $\sim 0.5\%$ for the wide-, medium-, and narrow-band filters and ~ 10 to 15% for the quad filters. However, the absolute error due to systematics and uncertainties in the white dwarf models used as absolute standards can be up to 1.5% . The signal to noise ratio, per observation, for the standard stars is such that the Poisson error is $\sim 0.2\%$ or less for most filters.

For those users who need precise photometry, a complete error analysis would include Poisson noise of the science target as well the additional sources of error of flat fielding (see [Section 5.4](#)), spatial repeatability ($\sim 0.5\%$, but see [Section 5.4.2](#)), readout noise ([Section 5.1.2](#)), uncertainty in the encircled energy ([Section 9.1](#)), and charge transfer efficiency corrections (CTE, see [Chapter 6](#)).

5.8 References

- S. Baggett, 2009, *WFC3 TV3 Testing: UVIS-1 Crosstalk*, [WFC3 ISR 2009-03](#).
- S. Baggett, 2009, *Performance of WFC3 Replacement UVIS filters*, [WFC3 ISR 2007-01](#).
- S. Baggett, E. Sabbi, P. McCullough, 2009, *WFC3 SMOV Proposal 11422/ 11529: UVIS SOFA and Lamp Checks*, [WFC3 ISR 2009-27](#).
- T. Borders, S. Baggett, 2009, *WFC3 SMOV Proposals 11419, 11426, 11431, and 11448: On-Orbit Biases*, [WFC3 ISR 2009-26](#).
- M. Bourque, S. Baggett, 2016, *WFC3/UVIS Dark Calibration: Monitoring Results and Improvements to Dark Reference Files*, [WFC3 ISR 2016-08](#).
- M. Bourque, D. Borncamp, S. Baggett, T. Desjardins, & N. Grogin, 2018, *Using Dark Images to Characterize Pixel Stability in the WFC3/UVIS Detector*, [WFC3 ISR 2018-15](#).
- T. M. Brown & O. Lupie, 2004, *Filter Ghosts in the WFC3 UVIS Channel*, [WFC3 ISR 2004-04](#).
- T. Brown, 2007, *UVIS Channel Filter Ghosts after Filter Replacement*, [WFC3 ISR 2007-09](#).
- T. Brown, 2007, *WFC3 TV2 Testing: UVIS Channel Glint*, [WFC3 ISR 2007-21](#).
- T. Brown, G. Hartig, S. Baggett, 2008, *WFC3 TV3 Testing: UVIS Window Contamination*, [WFC3 ISR 2008-10](#).
- S. E. Deustua et al., 2016, *UVIS 2.0 Chip-dependent Inverse Sensitivity Values*, [WFC3 ISR 2016-03](#).
- J. Fowler, L. Markwardt, M. Bourque, J. Anderson, 2017, *Analysis of Dragon's Breath and Scattered Light Detector Anomalies on WFC3/UVIS*, [WFC3 ISR 2017-02](#).
- J. Fowler, 2018, *WFC3/UVIS Gain Stability Results for Cycles 24 and 25*, [WFC3 ISR 2018-17](#).
- R. L. Gilliland, 1994, *Stellar Photometry Including Saturated Images: Results on M67 with WFPC2*, [Apj](#), **435**, L63-66.
- R. L. Gilliland, P. Goudfrooij, and R. A. Kimble, 1999, *Linearity and High Signal-to-Noise Performance of the STIS CCD*, [PASP](#), **111**, 1009-1020.
- R. L. Gilliland, 2004, *ACS CCD Gains, Full Well Depths, and Linearity up to and Beyond Saturation*, [ACS ISR 2004-01](#).
- R. L. Gilliland, A. Rajan and S. E. Deustua, 2010, *WFC3 UVIS Full Well Depths, and Linearity Near and Beyond Saturation*, [WFC3 ISR 2010-10](#).
- C.M. Gosmeyer & the Quicklook Team, *WFC3 Anomalies Flagged by the Quicklook Team*, [WFC3 ISR 2017-22](#).
- H. Gunning, C. Pavlovsky, S. Baggett, 2013, *WFC3 Cycle 19 Proposal 12690: UVIS Gain*, [WFC3 ISR 2013-02](#).
- H. Gunning, 2015, *WFC3 Cycle 21 Proposal 13561: UVIS Gain*, [WFC3 ISR 2015-05](#).
- B. Hilbert, 2009, *UVIS Channel Shutter Shading*, [WFC3 ISR 2009-05](#).

- B. Hilbert, 2010, *WFC3 SMOV Program 11427: UVIS Channel Shutter Shading*, [WFC3 ISR 2009-25](#).
- K. Huynh and B. Kuhn, *WFC3/UVIS Gain Stability Results for Cycles 26-29*, [WFC3 ISR 2022-08](#).
- K. Huynh and P. R. McCullough, *UVIS Shutter Timing Jitter*, [WFC3 ISR 2023-04](#).
- J. Janesick, 2001, *Scientific charge-coupled devices*, SPIE Optical Engineering Press, (Bellingham, WA).
- H. Khandrika and S. Baggett, 2015, *WFC3 UVIS Read Noise*, [WFC3 ISR 2015-13](#).
- H. Khandrika, 2017, *WFC3/UVIS Read Noise Aug 2009 - Apr 2017*, [WFC3 ISR 2017-17](#).
- B. Kuhn and H. Khandrika, *2018 Superbias Reference File*, [WFC3 ISR 2019-11](#).
- J. Mack, E. Sabbi, & T. Dahlen, 2013, *In-flight Corrections to the WFC3 UVIS Flat Fields*, [WFC3 ISR 2013-10](#).
- J. Mack, A. Rajan, & A.S. Bowers, 2015, *Spatial Accuracy of the UVIS Flat Fields*, [WFC3 ISR 2015-18](#).
- J. Mack, T. Dahlen, E. Sabbi, and A. S. Bowers, 2016, *UVIS 2.0: Chip-Dependent Flats*, [WFC3 ISR 2016-04](#).
- J. Mack, 2016, *UVIS 2.0: Ultraviolet Flats*, [WFC3 ISR 2016-05](#).
- R. van der Marel, 2003, *Determination of Low-Frequency Flat-Field Structure from Photometry of Stellar Fields*, [ACS ISR 2003-10](#).
- C. Martlin, 2016, *WFC3 Cycle 23 Proposal 14373: UVIS Gain*, [WFC3 ISR 2016-13](#).
- P. McCullough, 2011, *Geometric model of UVIS window ghosts in WFC3*, [WFC3 ISR 2011-16](#).
- J. V. Medina, 2021, *UVIS Pixel Stability: Updates to The UVIS Bad Pixel Table Pipeline*, [WFC3 ISR 2021-14](#).
- M. V. Newberry, 1991, *Signal-to-Noise Consideration For Sky-subtracted CCD Data*, [PASP, 103, 122](#).
- M. V. Newberry, 1999, *Increasing Precision and Accuracy in Photometric Measurements in Precision CCD Photometry*, [ASP Conference Series, Vol. 189](#), Edited by E. R. Craine, D. L. Crawford, and R. A. Tucker. ISBN: 1-58381-015-3 (1999), p.74.
- I. Rivera and B., Kuhn, *WFC3/UVIS: 2021 and 2022 Superbias Reference File Procedural Updates*, [WFC3 ISR 2023-03](#).
- I. Rivera, M. Marinelli, J. Mack, *WFC3/UVIS: New Full-Well Saturation Map Reference File*, [WFC3 ISR 2023-08](#).
- M. Robberto & B. Hilbert, 2005, *The Behaviour of the WFC3 UVIS and IR Analog-to-Digital Converters*, [WFC3 ISR 2005-27](#).
- E. Sabbi, M. Dulude, A.R. Martel, S. Baggett, H. Bushouse, 2008, *WFC3 Ground P-flats*, [WFC3 ISR 2008-46](#).
- E. Sabbi et al., 2009, *WFC3 Calibration Using Galactic Clusters*, [WFC3 ISR 2009-06](#).
- E. Sabbi, 2009, *WFC3 SMOV PROGRAM 11452: UVIS Flat Field Uniformity*, [WFC3 ISR 2009-19](#).

- Sabbi, E. and Bellini, A., 2013, *UVIS PSF Spatial & Temporal Variations*, [WFC3 ISR 2013-11](#).
- K. Sahu, S. Baggett, J. MacKenty, 2014, *Use of the Shutter Blade Side "A" for UVIS Short Exposures*, [WFC3 ISR 2014-09](#).
- K. Sahu, C.M. Gosmeyer, S. Baggett, 2015, *WFC3/UVIS Shutter Characterization*, [WFC3 ISR 2015-12](#).
- M. Stiavelli, J. Sullivan, and C. Hanley, 2001, *Modeling the Impact of Ghosts in WFC3/UVIS*, [WFC3 ISR 2001-17](#).
- A. Suchkov & S. Baggett, 2012, *WFC3/UVIS Crosstalk and Crosstalk Correction*, [WFC3 ISR 2012-02](#).
- M. H. Wong, 2010, *WFC3/UVIS Fringing: Ground test data and analysis codes*, WFC3 TIR 2010-01, Available by request from <http://hsthelpp.stsci.edu>.
- M. H. Wong, 2010, *Amplitude of fringing in WFC3/UVIS narrowband red filters*, [WFC3 ISR 2010-04](#).

Chapter 6: WFC3 UVIS Charge Transfer Efficiency - CTE

Chapter Contents

- [6.1 Overview](#)
- [6.2 CTE Losses And Background](#)
- [6.3 The Nature Of CTE Losses](#)
- [6.4 The Pixel-Based Model](#)
- [6.5 Empirical Corrections](#)
- [6.6 Dealing With CTE Losses in WFC3 UVIS Images](#)
- [6.7 Sink Pixels](#)
- [6.8 References](#)

6.1 Overview

WFC3/UVIS was installed on board HST in May of 2009. As of the summer of 2024, the instrument has completed its fifteenth year of operation in the harsh radiation environment of low-Earth orbit.

There are three main consequences of radiation damage: (1) an increase in the number of hot/warm pixels; (2) an increase in the overall dark-current level; and (3) an increase in the number of charge traps, which cause CTE (charge-transfer-efficiency) losses and [sink pixels](#). Some of the hot pixels can be annealed away by warming the detector to $\sim 20^\circ$ C during the monthly anneal procedures (e.g. Figure 6 in [WFC3 ISR 2016-08](#)), but there is a gradual build-up of permanent hot pixels over time. There is no indication that annealing reduces CTE losses.

This chapter describes the various aspects of CTE losses in the WFC3/UVIS detector. The first section provides a general overview and compares CTE losses in WFC3/UVIS to those in ACS/WFC. The next section discusses the nature of WFC3/UVIS losses in more detail. [Section 6.4](#) describes the pixel-based model in terms of how it was constructed and how the pixel-based-reconstruction procedure can be run. [Section 6.5](#) gives a brief overview of the empirical formula-based corrections available for point sources. The final section discusses how best to deal with CTE losses in WFC3/UVIS images in terms of prevention, mitigation, and correction.

The CTE losses described and characterized in this chapter are in the parallel transfer of charge, aka y-CTE. WFC3/UVIS does experience some charge-transfer losses in the serial direction (x-CTE), but these losses are much smaller than the parallel-transfer losses. [WFC3 ISR 2024-07](#) provides a discussion of the current impact of serial CTE on the WFC3/UVIS detector. In brief, while y-CTE affects both the astrometry and photometry of bright sources at the 5% level and can have a pathological effect on faint sources, the impact of x-CTE is almost entirely astrometric and is only 1% of a pixel for bright stars and 2.5% for fainter stars. The ISR provides more information and describes the stand-alone pixel-based corrections for those few projects that may require them.

6.2 CTE Losses And Background

ACS/WFC was installed on board HST in March 2002 and has been in space about 7 years longer than WFC3. While it took several years for CTE to become a concern for ACS observations, CTE losses became apparent for WFC3/UVIS even before it had been in orbit for a year. In fact, CTE losses for even moderately bright targets were 2-3 times greater than the expected CTE losses based on ACS's performance in its first years on-orbit ([ACS ISR 2009-01](#)). It is worth noting that the faster degradation seen by UVIS 2009/2010 was also occurring on ACS as well in those years (see [Massey 2010, MNRAS 409L, 109](#)). As discussed in a [CTE White paper](#) in 2011, one contributing factor to WFC3 /UVIS's rapid CTE losses was likely the solar minimum:

1. The strength and extent of the South Atlantic Anomaly, the region in HST's orbit where most of the radiation damage is thought to occur, is known to be negatively correlated with solar activity, and
2. The installation of WFC3 on HST coincided with solar minimum while the installation of ACS occurred during a period of higher solar activity.

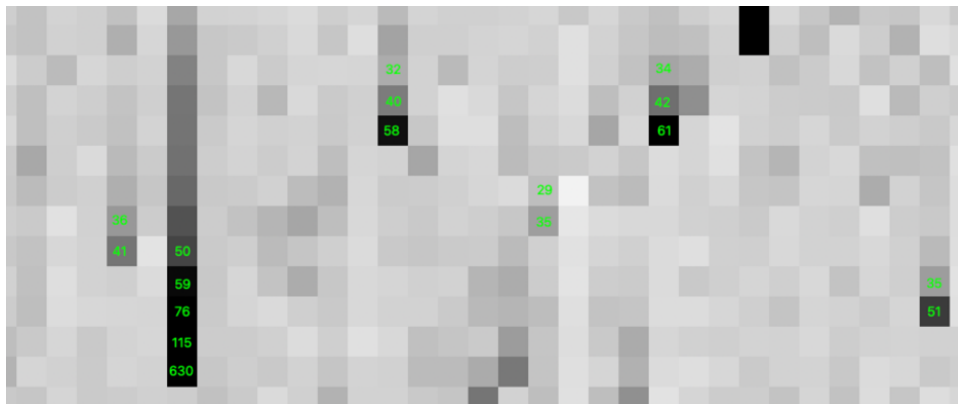
An additional factor in the WFC3/UVIS CTE behavior is its typical image background. It has long been known that CTE losses are greater in images with low backgrounds. For several reasons, WFC3/UVIS images tend to have much lower backgrounds than ACS images. First, WFC3/UVIS was designed to be sensitive to low S/N objects, with low readnoise (~ 3 electrons/pixel) and low dark current (~ 12 electrons/pixel/hr in 2024 and increasing by ~ 0.5 electron/hr with every additional year on orbit). The dark current of ACS is about 50-55 electrons/hour, several times higher. Second, the WFC3/UVIS pixels are 40% smaller than the ACS/WFC pixels, and as such they intercept correspondingly fewer sky photons, again keeping the WFC3/UVIS background low. Finally, images obtained with WFC3 /UVIS's many UV and narrow-band filters often have extremely low sky backgrounds. For all these reasons, even long UVIS science exposures frequently have backgrounds of less than 5 electrons, whereas similar ACS/WFC science exposures rarely have backgrounds less than 25 electrons. (For a summary of WFC3/UVIS and ACS backgrounds, see [WFC3 ISR 2012-12](#), and [ACS ISR 2012-04](#), respectively.)

A recent report ([WFC3 ISR 2021-09](#)) has shown that, in addition to having lower backgrounds than ACS/WFC, the WFC3/UVIS detector is also more sensitive to CTE losses when the background is low. Figure 9 in that document shows that even though WFC3/UVIS has been in orbit for only 65% of the time ACS/WFC has, it already suffers more CTE losses than ACS/WFC when the background is below 20 electrons. The loss spectrum for ACS follows a simple power law, whereas that for WFC3/UVIS shows an upturn from the power law at low backgrounds (e.g. Figure 9 in [WFC3 ISR 2021-09](#)). This finding was unexpected, as WFC3 was designed to have a mini channel that would preserve charge from losses at the low-signal end.

6.3 The Nature Of CTE Losses

CTE losses arise during the readout process, as the charge packet for each pixel is transferred pixel-by-pixel down the detector, in parallel, to the serial register. As the packet moves through the silicon of the detector, it often encounters imperfections (traps) in the lattice that have been generated by radiation damage. These traps can temporarily detain individual electrons, and once an electron is trapped, it becomes separated from its original charge packet. This trapped electron is often released some time later during the readout and finds itself in an upstream pixel. For this reason, CTE takes charge away from downstream pixels and deposits it into upstream pixels. Visually, the effect results in “trails” of charge that extend out from sources in the direction opposite the readout amplifier (See [Figure 6.1](#)).

Figure 6.1: A close up of a 31x12-pixel region of a recent stack of 900s darks far from the amplifier. Parallel shifting is downward toward the readout amplifier.



Electrons are confined within pixels by the electrostatic potential wells created by the electrodes. Charge packets that have more electrons tend to occupy a physically larger volume within the pixel than smaller charge packets, on account of their self-repulsion pushing back against the pixel's potential well. The larger volume occupied by larger charge packets provides a larger cross-section to traps as the packets are shuffled through the silicon lattice. For the WFC3/UVIS model we discuss below in [Section 6.4](#), a cloud with 10 electrons, for example, will see about 3 times more traps than a cloud with only 1 electron. Thus, even though losses increase in an absolute sense when a cloud has more electrons, the per-electron losses decrease with increasing packet size.

Figure 6.2: Effect of Electron Traps

Left: The dark line shows the cumulative number of traps in each column as a function of the size of the electron packet. The dotted line shows the extrapolation of the power law from the bright end; Right: The derivative of the cumulative curve on the left, showing the marginal number of traps seen by the Nth electron in a packet. The dashed line denotes where there is one trap per marginal electron.





Figure 6.2 shows the current model for charge traps (WFC3 ISR 2021-09, see also Section 6.4). The model was based on the analysis of the blurring experienced by hot and warm pixels in dark exposures taken in late 2020. On the left, we show the cumulative number of traps as a function of packet size. When there is no background, a charge packet that contains just one electron and is located far from the readout amplifier will encounter 40 traps on its ~2000-pixel journey to the serial register. Clearly, such an electron is unlikely to survive the journey. A charge packet with ten electrons will encounter 100 traps and is also not likely to be detected.

A packet with 100 electrons will encounter 150 traps. Of course, an isolated packet that starts with 100 electrons will not maintain that size all the way down, so in practice it will see fewer than 150 traps during the latter part of its journey. Packets with 1000 electrons on zero background will lose about 300 electrons, so over 70% will survive to the register. Packets with 10,000 electrons on an image with no background will lose only about 6% of their electrons.

In the right panel of Figure 6.2 we show the marginal number of traps, i.e. the number of traps that would be seen by the Nth electron in a charge packet. The first electron will encounter almost 40 traps on its path to the readout register, and thus, have a very small chance of making the journey with its original pixel. The second electron will see 20 traps, the third 15, and the tenth about 2, and so on.

Note that these curves cannot directly predict the original number of electrons from the observed number, since as a charge packet gets shuffled down the detector, it loses electrons and therefore its electron-loss-rate changes in a non-linear way. In addition, downstream packets can “prefill” traps, such that upstream packets can be shuffled down the chip with fewer losses. CTE losses can also cause asymmetries in source shapes which is why a forward-modeling approach is required to reconstruct the pixel distribution.

 The fact that the marginal losses drop sharply as a function of charge packet size means that if we have a small level of background in an image, many of the traps will be kept filled by the background and will not therefore affect the science photons above the background.

The curve shows that in an image with a background of zero, a marginal electron will see 20 traps, but if the background is 10, it will likely see less than two traps. If the background is around 20, in fact, a marginal science electron has about a 50% chance of making it to the readout register (as of late 2020).

Finally, we note that the delta-function charge packets that we have been considering in our study of warm pixels in dark exposures are of course not representative science sources. Even unresolved objects on UVIS have a point-spread-function shape (PSF) where the central pixel receives about 20% of the star's light and the immediately adjacent pixels each receive about 10% of the light. As a consequence, the outer pixels of sources will fill some traps that the sources' inner pixels then will not have to encounter. This can increase the overall transfer efficiency of real-world sources over the warm pixel delta functions in the darks.

6.4 The Pixel-Based Model

The previous section made reference to the pixel-based model of CTE losses. The WFC3/UVIS model is based on the empirical model for ACS/WFC that was constructed by [Anderson & Bedin \(2010\)](#), which itself was inspired by the model in [Massey et al. \(2010\)](#).

The pixel-based model has two sets of parameters. The first set describes the number and distribution of electron-grabbing traps, and the second set describes the release of trapped charge.

The basic readout model assumes that each column of pixels contains a number of charge traps distributed evenly among its pixels from row $j=1$ to $j=2051$. Charge packets for pixels at the top of the chip (far from the serial register) will encounter all of the traps in the column. Those in the middle pixels, at row $j\sim 1024$, will see half of the traps, and those at $j\sim 200$ will experience a tenth of them. According to the model, each trap will grab and hold a particular electron in a charge packet (i. e., the first electron, the second, the hundredth, the ten-thousand-third, etc.).

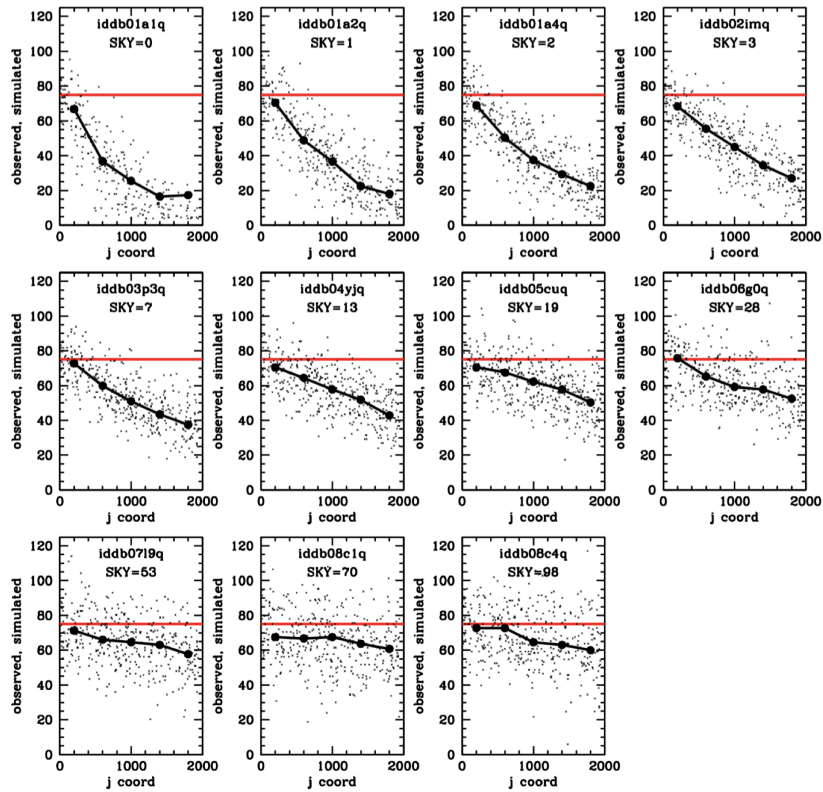
The model also has a release profile for each trap. With most traps, there is a 20% chance the electron will be released into the first upstream pixel, a 10% chance it will be released into the second upstream pixel, etc. For WFC3/UVIS, essentially all trapped electrons are released after 60 transfers. In addition, the traps that grab the first few electrons in a pixel appear to have a much steeper release profile wherein all the charge is released after only 6 transfers (see Figure 5 in [WFC3 ISR 2021-09](#)).

The model simulates the readout process by going up through each column, pixel by pixel, from row $j=1$ (close to the serial register) to $j=2051$ (at the top of the detector) and determining (1) which electrons will be trapped during the journey down the detector and (2) when these trapped electrons will be released. Even though in reality this is a stochastic process, the model must treat it deterministically. The end result of the readout process is that the observed distribution of pixel values gets blurred in the upstream direction relative to the original distribution, leading to the characteristic CTE trails seen behind sources in CTE-affected UVIS images.

In order to determine optimal parameters for the pixel-based model, we examined the blurring experienced by warm pixels (WPs) in a variety of dark exposures taken in December 2020 with both short (30s) and long (900s) exposure times, and with various post-flash background levels (from zero up to 100 electrons). With this strategy, the WP intensity from the long exposures (where they did not suffer much loss) can be scaled down, in order to estimate the true intensity of the WPs in the short exposures where their fractional losses will be greater and more dependent on the background. This procedure enabled us to directly examine the losses as a function of background.

[Figure 6.3](#) tracks a WP that started with 75 electrons and shows how many of these electrons remain in the original pixel packet as a function of how many parallel shifts it experiences. The different panels correspond to different post-flash background levels showing losses are more than 50% for any background below about 19 electrons for a source transferred down the entire detector.

Figure 6.3: The observed number of electrons that survive to readout for a 75 electron warm pixel as a function of the number of parallel shifts to the amplifier. Each panel shows a different background. The red line shows the starting value from the warm pixel. Taken from Figure 3 of [WFC3/ISR 2021-09](#).



CTE Losses as a Function of Transfers and Background.

Plots such as [Figure 6.3](#) allow us to pin down the CTE losses as a function of background level, and thus as a function of packet size. The right panel in [Figure 6.2](#) shows the number of traps experienced by the marginal electron. A faint source will experience pathological losses on any background where this is greater than 1.

The curves in [Figure 6.2](#) can be used in a forward model to simulate how an initial pixel distribution would be blurred by the readout process. In order to use the model on science images, this procedure must be inverted, i.e. we must determine what original image - when pushed through the readout algorithm - generates the observed image. This is formally a non-linear deconvolution process.


There are several challenges to a pixel-based reconstruction. One challenge is that the observed images represent not just the charge that arrived at the readout, but they represent this charge plus a contribution from the readnoise (3.1-3.2 electrons for UVIS) along with any noise associated with the charge-trapping and -releasing, which is a random process (in other words, a noisy process). The other challenge is that the model assumes an even distribution of traps throughout the detector, even though in actuality this distribution is stochastic: each column will naturally have a slightly different number of traps and even the distribution of traps within each column will surely not be perfectly uniform.

The issue of readnoise introduces a serious complication. An empty image will be read out to have a variance of about 3.2 electrons in each pixel due to readnoise. If the original image on the detector had a pixel-to-pixel variation of 3.2 electrons on a low background, then the charge-transfer process would blur the image out so much that the image that reached the amplifier would have a variance of less than 0.5 electron (see Figure 12 in [WFC3/ISR 2021-09](#)). The CTE reconstruction algorithm determines what the original image would have to be in order to be read out as the observed image. If we include readnoise in this observed target image, then the original image would have to start with an extremely large amount of pixel-to-pixel variation (a sigma of more than 10 electrons) for it to end up with 3.2 electrons of pixel-to-pixel variation after the blurring readout process. In other words, we would have to add a factor of five more noise to arrive at the image that was read out which would not produce a useful reconstruction.

To mitigate this "readnoise amplification" issue, we take the observed image (with readnoise) and determine the smoothest possible image that is consistent with the image, modulo readnoise and noise from the charge-transfer process. We choose the smoothest possible image, since the smoother the image, the less the readout algorithm will end up redistributing charge. While this may not represent the true counts that arrived at the readout register, it should provide the most conservative possible CTE correction. The pixel-based reconstruction algorithm then operates on this "smooth" image, and the redistribution of flux is applied to the observed image (with readnoise).

[Figure 6.4](#) shows an example of a star field that has considerable CTE losses. The images show the uncorrected `_flt` image on the left and the pipeline CTE-corrected `_flc` image on the right, a 4-s exposure (iehr02lkq) taken at the center of globular cluster Omega Centauri. The portion of the image shown is near the top of the chip, furthest from the amplifier where CTE losses are greatest. The image background here is 15 electrons/pixel (5 electrons natural and 10 electrons post-flash). The CTE trails are obvious in the image on the left. The pixel-based algorithm does a reasonable job removing the trails and restoring flux into the sources. The few instances of over-subtracted trails (which appear as white pixels in this inverted stretch) come from cosmic rays that struck the detector during readout. The electrons from these CRs did not undergo the 2000 parallel transfers that the other electrons in the pixels did.

Figure 6.4: A 200×250-pixel portion of a 4-second exposure (iehr02lkq) taken of the center of globular cluster Omega Centauri. The scene shown is near the top of UVIS2, about 2000 parallel shifts from the readout amplifier. The panel on the left shows the uncorrected `_flt` image. The panel on the right shows the `_flc` image, which has had the pixel-based CTE correction performed in the pipeline.



Before 2016, the pixel-based CTE correction was a standalone FORTRAN program available for download. Since then, however, the pixel-based CTE correction has been incorporated into the automated calibration pipeline for all UVIS full-frame images (`calwf3` version 3.3 and later) and most¹ UVIS subarray images (`calwf3` version 3.4 and later). Invoking the pixel-based CTE correction within the `calwf3` pipeline is controlled via an image header calibration switch (`PCTECORR`), and associated calibration table (`PCTETAB`) and CTE-corrected reference files (e.g. `DRKCFILE`). `calwf3` will, by default, produce two sets of products: the standard non-CTE-corrected files (e.g. `*_flt.fits` and `*_drz.fits`), as well as the CTE-corrected results (`*_flc.fits`, `*_drc.fits`). Observers can use the `*_flc.fits` and `*_drc.fits` data products in the same way as `*_flt.fits` and `*_drz.fits` files.

In general, the pixel-based correction is accurate to about 25%. However, in low-signal/low-background situations where CTE losses can become greater than 25%, the reconstruction procedure can have difficulties as it uses the surviving charge distribution as a "scaffold" upon which to reconstruct the lost charge. When the surviving charge no longer resembles the original charge distribution, reconstruction becomes impossible. Until 2015, images with backgrounds of 12 electrons/pixel or greater experienced CTE losses less than 25%, but now in 2024, image backgrounds of 20 electrons/pixel keep CTE losses for faint source to less than ~23% ([WFC3 ISR 2024-04](#)). In the future, higher image background are likely to be required to continue to keep the CTE losses under 25%, where the pixel-based correction works best.

One other feature of pixel-based reconstruction worth mentioning is that when the background is low and CTE losses are large, there are two undesired consequences: the correction algorithm cannot reconstruct the scene and it amplifies noise. To avoid this, the pipeline algorithm is designed to suppress noise amplification, and as a result, little reconstruction is attempted for sources with fluxes near the background level (see Section 5 of [WFC3 ISR 2021-09](#) for more details). Given the limitations of reconstruction at low signal and background levels, we have made available some empirical corrections that do not suffer from noise amplification in [WFC3 ISR 2021-13](#).

Another way to extract CTE corrected fluxes post-observation is to use the software routine, **hst1pass**. **hst1pass** performs PSF-fitting photometry to measure stellar (or point-source) fluxes and positions and the user can specify output parameters to include CTE-corrected fluxes. For more information on **hst1pass**, including guidance on how to download and run this routine, please see [WFC3 ISR 2022-05](#).

The latest information on the status of CTE and how to mitigate its effect is always available on the [WFC3 CTE webpage](#).

 The WFC3 CTE web page is at: <http://www.stsci.edu/hst/instrumentation/wfc3/performance/cte>

¹For observers with subarrays that do not contain overscan columns (UVIS2-M1K1C-SUB and UVIS2-M512C-SUB) - these are not supported by the calwf3 CTE correction - there is a workaround available that uses the standalone FORTRAN CTE-correction code ([STAN 18](#)).

6.5 Empirical Corrections

Empirical corrections provide an alternate way to correct point-source observations for CTE losses or estimate future CTE loss levels. Noeske et al. ([WFC3 ISR 2012-09](#)) studied star clusters NGC6791 and 47 Tuc and examined how the measured fluxes of stars varied when particular stars were placed close to and far from the readout register. They documented the losses experienced by point sources as a function of (1) the total flux (in electrons) of the point source, (2) the detector Y position of the source, (3) the image background level, and (4) the date of the observation.

These corrections are provided in the form of polynomials and can be used to estimate how many electrons were originally present in the aperture before the readout process. The conclusions in [WFC3 ISR 2012-09](#) and in more recent follow-up studies ([WFC3 ISR 2016-17](#), [WFC3 ISR 2017-09](#), [WFC3 ISR 2021-06](#), and [WFC3 ISR 2024-04](#)) mirror the conclusions from the pixel-based-model CTE study: namely, that as of 2024, a background of ~ 20 electrons or more can have a significantly ameliorative effect on charge-transfer efficiency. The most recent empirical corrections are described in [WFC3 ISR 2024-04](#), which also provides an analysis of the long term behavior of the UVIS CTE using observations of star clusters. This information is also kept updated on the WFC3 CTE page at: <http://www.stsci.edu/hst/instrumentation/wfc3/performance/cte>.


6.6 Dealing With CTE Losses in WFC3 UVIS Images

There are several ways of mitigating the effects of imperfect charge-transfer efficiency in WFC3/UVIS images. Sections 7 and 8 of [WFC3 ISR 2021-09](#) provide a detailed discussion of the many issues and trade-offs; this section provides an overview.

The best way to deal with CTE losses is by minimizing them in the first place. Neither the reconstruction algorithm nor the empirical correction is perfect, so minimizing the need for these corrections is the first line of defense in reducing CTE losses. The easiest way to do so is to place the source close to the readout register, thereby minimizing the number of readout transfers during which CTE losses can occur. Of course, this is only possible for fields-of-interest that are relatively small.

Another way to minimize CTE losses is to ensure that the total image background (sky + dark + postflash if needed) has a minimum level of 20 electrons per pixel (as of 2024; see [WFC3 CTE page](#) for any updates). A total image background at the recommended value will enable at least half of the electrons of a marginal source to survive to the readout (see Figure 18 of [WFC3 ISR 2021-09](#)).

In order to predict the number of background electrons in a planned image through a particular filter with a given exposure time, users should use the [WFC3/UVIS imaging exposure time calculator](#) (ETC). The ETC provides an estimate of the predicted total background based on an input observing mode and will provide a warning if the background falls below the current recommendation. In images where the predicted natural background is less than the recommended level, observers should use post-flash to make up the difference. The ETC does not include an estimate of the actual astronomical scene in its calculations, and knowledge of what background to expect from your source can help you add just the right amount of post-flash. In general, though, it is better to flash too much rather than not enough.

 **NOTE:** The post-flash (PF) levels in APT and the ETC represent the average across the FOV. The [PF illumination pattern](#) varies by $\sim \pm 20\%$ across the FOV, and is brightest in B/D quadrants, falling off towards A&C corners. As a consequence, observers desiring to ensure a specific electron/pixel level across the entire FOV will want to increase the requested levels by $\sim 20\%$.

Backgrounds above the recommended level provide increased protection against CTE losses, but adding background also adds noise, and the net signal to noise ratio for marginal sources ends up being essentially flat between backgrounds of 20 and 50 electrons. Splitting up long exposures into multiple dithered exposures is commonly done to improve PSF sampling and remove detector artifacts and cosmic rays. However, from a CTE-minimization perspective, it can be better to take fewer, longer exposures to reduce the amount of postflash required (and the resulting noise), thereby improving the signal to noise for faint sources.

Once CTE losses have been minimized in a data set, it will still likely be important to correct the observations for whatever losses have occurred, or at the very least obtain an estimate of how the readout process may have degraded the original image. One can use either the pixel-based CTE correction in `calwf3` for general scenes of point and extended sources ([Section 6.4](#)) or the empirical formula correction for point sources (see [Section 6.5](#)).

6.7 Sink Pixels

With the advent of post-flashing in 2012, a new type of image defect called sink pixels was identified (see [WFC3 ISR 2014-04](#) for the initial discovery and [WFC3 ISR 2014-19](#) for additional analysis). Sink defects are caused by pixels that contain a modest number of charge traps (typically 20–100 electrons). During exposure integration, the pixels fill with charge, and some of that charge gets stuck in traps. When the parallel shifting begins, the trapped charge does not shuffle down with the other electrons associated with the pixel, resulting in less charge recorded in that pixel than expected.

Sink pixels stand out in post-flashed bias images, as when they are read out, they contain fewer counts than adjacent “normal” pixels (see Figure 4 of [WFC3 ISR 2014-22](#)). This phenomenon is distinct from normal pixel-to-pixel sensitivity variations, in that photons that interact with sink pixels do generate electrons, but some of those electrons do not shuffle out of the pixel during the readout process, and are thus not recorded with that pixel. Sensitivity variations apply to all impacting photons irrespective of how bright the pixel is, whereas sink-pixel trapping generally has a smaller fractional effect on brighter pixels.

Only a very small number of sink pixels were found in data taken before launch. Investigations suggest that most of the sink pixels are likely a consequence of on-orbit radiation damage ([WFC3 ISR 2014-19](#), [WFC3 ISR 2014-22](#)), in fact that sink pixel traps are likely the same traps that cause CTE losses. The distinction is that CTE losses correspond to charge getting shuffled through traps. Sink pixels represent trapping before any shuffling starts. As such, sink pixel effects are localized losses, while CTE losses are more generalized. At present, no sink pixel has been found to heal or recover, consistent with the monotonically increasing CTE losses.

The impact of sink pixels depends both on their locations in an image and the pre-readout pixel value. For images with high backgrounds (~ 85 electrons/pixel) and for sink pixels near the readout register, the sinks are single low pixels and have little effect on upstream or downstream pixels in the same column. However, for images with low backgrounds or for pixels far from the readout register, the interplay between CTE losses and sink pixels can extend the sink pixel profile to more than 10 pixels (see Figure 3 of [WFC3 ISR 2014-19](#)). So although sink pixels are rare (at last count, $\sim 0.25\%$ of the detector), in low-background imaging they can corrupt as much as $\sim 1\%$ of the detector. The impact of the sink pixels is also scene-dependent. For example, if a source lands on a sink pixel (or the streak of a sink pixel), then the electrons in the source will affect the trail behind the sink in the same way that a higher background would. Thus, the WFC3 team has adopted a conservative approach to flag all pixels that are likely to be affected in a given image; observers can choose whether to use or ignore the sink pixel flags (DQF value of 1024, see [Table 3.2](#)) in the science image DQF files.

The calibration pipeline **calwf3** (version 3.3+) uses the sink pixel reference file (`SNKCFILE`) to populate the data quality array of a science image with 1024 for all flagged sink pixels. The sinks are identified in the `SNKCFILE` with the modified-Julian date (MJD) of the appearance of the sink pixel; any upstream/downstream adjacent pixels affected by the sink pixel are flagged as well (see Section 6 in [WFC3 ISR 2014-22](#)). More details on the flagging process are given in [Section 3.2.7](#), as well as in [WFC3 ISR 2014-22](#). Note that the sink pixels flagging is performed regardless of whether the CTE-correction is performed, i.e. both `flt` and `flc` will have sink pixels flagged in the DQ extensions; the science data pixels are not changed in any way.

6.8 References

- J. Anderson, *One-Pass HST Photometry with `hst1pass`*, [WFC3 ISR 2022-05](#).
- J. Anderson, *Revisiting x-CTE in WFC3/UVIS*, [WFC3 ISR 2024-07](#).
- J. Anderson, *Table-Based CTE Corrections for flt-Format WFC3/UVIS*, [WFC3 ISR 2021-13](#).
- J. Anderson & S. Baggett, *Flagging the Sink Pixels in WFC3/UVIS*, [WFC3 ISR 2014-22](#).
- J. Anderson & S. Baggett, *Sink Pixels and CTE in the WFC3/UVIS Detector*, [WFC3 ISR 2014-19](#).
- J. Anderson, & L. Bedin, 2010, *An Empirical Pixel-Based Correction for Imperfect CTE. I. HST's Advanced Camera for Survey*, [PASP 122 1035](#).
- J. Anderson, et al. 2012, *The Efficacy of Post-Flashing for Mitigating CTE-Losses in WFC3/UVIS Images*.
- J. Anderson, et al., 2021, *Updating the UVIS CTE Model and Mitigation Strategies*, [WFC3/ISR 2021-09](#).
- S. Baggett, & J. Anderson, 2012, *WFC3/UVIS Sky Backgrounds*, [WFC3 ISR 2012-12](#).
- S. Baggett, et al., 2011, [CTE White Paper](#).
- S. Baggett, C. Gosmeyer, and K. Noeske, 2015, *WFC3/UVIS Charge Transfer Efficiency 2009 -2015*, [WFC3 ISR 2015-03](#).
- J. Biretta & M. Bourque, *WFC3 Cycle 19 and 20 Dark Calibration: Part I*, [WFC3 ISR 2014-04](#).
- M. Bourque & S. Baggett, *WFC3/UVIS Dark Calibration: Monitoring Results and Improvements to Dark Reference Files*, [WFC3 ISR 2016-08](#).
- M. Chiaberge et al., *Updated CTE Photometric Correction for WFC and HRC*, [ACS ISR 2009-01](#).
- C. M. Gosmeyer & S. Baggett, *WFC3/UVIS External CTE Monitor: 2016 Updates on Coefficients and Analysis Pipeline*, [WFC3/ISR 2017-09](#).
- C. M. Gosmeyer & S. Baggett, *WFC3/UVIS External CTE Monitor: Single-Chip CTE Measurements*, [WFC3 ISR 2016-17](#).
- B. Kuhn, *WFC3/UVIS External CTE Monitoring 2009-2024*, [WFC3 ISR 2024-04](#).
- B. Kuhn & J. Anderson, *WFC3/UVIS New FLC External CTE Monitoring 2009-2020*, [WFC3 ISR 2021-06](#).
- J. Mackenty, & L. Smith, 2012, [CTE White Paper](#).
- R. Massey, 2010, *Charge Transfer Inefficiency in the Hubble Space Telescope since Servicing Mission 4*, [MNRAS 409L, 109](#).
- R. Massey et al., 2010, *Pixel-based correction for Charge Transfer Inefficiency in the Hubble Space Telescope Advanced Camera for Surveys*, [MNRAS 401 371](#).
- K. Noeske, et al., 2012, *WFC3 UVIS Charge Transfer Efficiency October 2009 to October 2011*, [WFC3 ISR 2012-09](#).
- J. Sokol, et al., 2012, *Assessing ACS/WFC Sky Backgrounds*, [ACS ISR 2012-04](#).

Chapter 7: WFC3 IR Sources of Error

Chapter Contents

- [7.1 WFC3 IR Error Source Overview](#)
- [7.2 Gain](#)
- [7.3 WFC3 IR Bias Correction](#)
- [7.4 WFC3 Dark Current and Banding](#)
- [7.5 Blobs](#)
- [7.6 Detector Nonlinearity Issues](#)
- [7.7 Count Rate Non-Linearity](#)
- [7.8 IR Flat Fields](#)
- [7.9 Pixel Defects and Bad Imaging Regions](#)
- [7.10 Time-Variable Background](#)
- [7.11 IR Photometry Errors](#)
- [7.12 References](#)

7.1 WFC3 IR Error Source Overview

In comparison to data from the earlier generation of IR detector flown on the Hubble Space Telescope ([NICMOS](#)), the WFC3/IR data's improved quality and performance simplifies the tasks of data reduction and analysis. We present and discuss the known error sources to allow a careful data processing, which will yield improved results. Each is identified, its impact on the data is illustrated, and, if available, possible processing solutions are presented.

7.2 Gain

Gain, a fundamental scaling parameter of the IR detector, relates the registered number of analog-to-digital units (ADUs) or equivalently, data numbers (DNs), to the number of accumulated photoelectrons. Like the UVIS channel, the Analog to Digital Converter (ADC) outputs a 16-bit number, allowing output signal values ranging from 0 to 65535 (i.e. $2^{16} - 1$) for each pixel.

While it is possible to operate the WFC3/IR detector at gain settings of 2.0, 2.5, 3.0, and 4.0 e-/DN, only a gain setting of 2.5 e-/DN is supported. This gain provides a complete sampling of the full-well of the IR detector ($\sim 80,000$ electrons, which is $\sim 32,000$ DN at the 2.5 e-/DN setting). Due to slight hardware differences, effective gain values in each quadrant of the detector deviate slightly from the nominal gain setting. Twice yearly, the effective gain is measured via the mean-variance technique on pairs of internal flat fields: the mean signal level is plotted against the variance in signal, and the gain is the inverse of the slope. See [WFC3 ISR 2015-14](#) for details on this method and the status of the IR gain from 2010-2015; results covering 2016-2023 are in TIR 2024-01. [Table 7.1](#) summarizes the average values of the effective gain measured in each quadrant of the IR detector. The WFC3 group monitors the effective gain of the detector twice yearly to evaluate its stability over time.

Table 7.1: Gain measurements for each quadrant, in e-/DN, averaged over all epochs of observations. Quadrants 1, 2, 3, and 4 are the 512×512 pixel detector quadrants that are read out independently of each other (see [Section 2.2.2](#)).

Quadrant.	Gain	Error
1	2.28	0.026
2	2.221	0.027
3	2.24	0.027
4	2.265	0.032

7.3 WFC3 IR Bias Correction

At the beginning of each MULTIACCUM observation, when the detector is reset (i.e. read out), a net DC bias with a value of order 11,000 DN is introduced. [Section 3.3](#) provides details on how this bias level is corrected (see also [Section 7.7.2](#) of the *WFC3 Instrument Handbook (IHB)*). The overall bias level is different for each readout quadrant, and is automatically removed during **calwf3** processing. In addition, there is a fixed pattern signal representing the inherent pixel-to-pixel reset levels. In standard MULTIACCUM processing, a two-step process is used to remove these signals.

First, reference pixels are used to remove any changes in the overall bias level that may occur during the exposure. The 5 columns of reference pixels on the left and right sides of the IR detector are insensitive to illumination; the four innermost of the five are used to measure the bias level of the detector at the time of each readout (see [Section 5.6](#) of the *IHB*). During **calwf3** processing, the mean measured signal of the reference pixels in each read is calculated and subtracted from all of the science pixels on the detector for that read. This method has the advantage of removing any bias drifts with time. [WFC3 ISR 2002-06](#) details the method of reference pixel bias subtraction, while [WFC3 ISR 2012-05](#) shows reference pixel behavior, and compares various methods of subtracting the reference pixel signal. A long-term drift in the raw pixel values of the reference pixels and the zeroth read signals has been observed over the course of on-orbit observations (see Figures 5 and 6 of [WFC3 ISR 2017-04](#)) but has had no adverse effect on the final calibration. The reference pixel level subtraction is performed in the `BLEVCORR` step of **calwf3**.

Next, pixel-to-pixel and quadrant-to-quadrant differences in reset levels are removed by subtracting the zeroth readout of the science exposure from all subsequent readouts. This is performed in the `ZOFFCORR` step of **calwf3** ([Section 3.3](#)) The IR bias signal is therefore not a component of any of the calibration reference files, because it is removed using data contained in the science exposure itself.

7.4 WFC3 Dark Current and Banding

[7.4.1 Dark Current Subtraction](#)

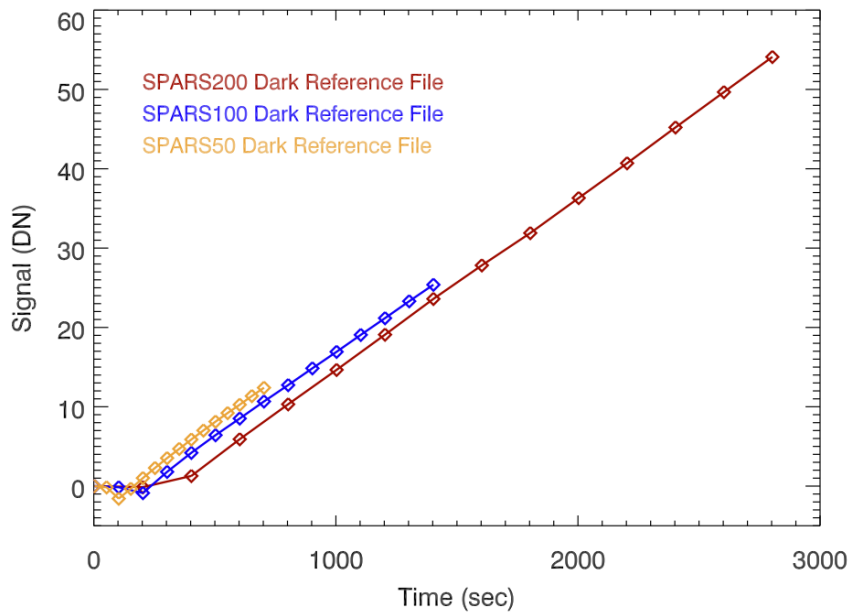
[7.4.2 Banding](#)

7.4.1 Dark Current Subtraction

The dark current in the IR detector is the signal measured when no illumination source is present. In an ideal detector, this signal would grow linearly with integration time. In practice, the dark current behavior of the IR detector is dependent upon the timing pattern used to collect each observation and is not constant for the duration of a given MULTIACCUM ramp (for more information, see [Section 7.7](#) in the *WFC3 Instrument Handbook*). The measured dark current can even be negative in certain situations (e.g. detector reset effects, see [WFC3 ISR 2010-16](#) and references therein). [Figure 7.1](#) shows a plot of the mean measured dark current signal versus time for three different timing patterns. Note that the three curves are not superimposed one another, nor do they show a straight line for the entire duration of the ramps. Details are presented in [WFC3 ISR 2009-21](#). For these reasons, there is a separate MULTIACCUM dark current reference file for each sample sequence and mode (full-frame, subarray) combination. During pipeline processing, `calwf3` subtracts the appropriate dark current ramp read-by-read from the science observation (see [WFC3 ISR 2012-11](#)).


The behavior of the IR darks has been analyzed in on-orbit data spanning Sep 2009 - June 2016. Although the darks possess a similar signal pattern across the detector, the median dark rates vary by as much as ~ 0.03 e⁻/s ([WFC3 ISR 2017-04](#)). The distribution of these median values has a triangular shape with a mean and standard deviation of 0.049 ± 0.007 e⁻/s. There is no apparent systematic long-term trend in the dark signal (see Figure 3 from [WFC3 ISR 2017-04](#)).

Figure 7.1: Dark Current Signal vs Time.



The dark current calibration files are created from a stack of dark up-the-ramp datasets which are taken on a regular basis throughout each observing cycle. For each sample sequence, the dark current calibration file is created by calculating the robust (outlier rejected) mean signal for each pixel in each read. Calculated uncertainties in the dark current calibration signals (in the error arrays of these files) are propagated into the error arrays of the calibrated science observations at the time of the dark current subtraction by **calwf3**. Figure 7.2 illustrates the large-scale dark current structure, showing the measured signal rate in a high signal-to-noise dark current calibration ramp. In general, the upper left quadrant of the detector has the highest dark current, while the upper right has the lowest.

Figure 7.2: Image showing the large-scale dark current structure of the IR detector. The image is shown with a histogram equalization stretch between 0 and 0.07 e⁻/sec.



The time-variable behavior of the WFC3/IR darks has been analyzed in on-orbit data covering May 2009 - Oct 2017 and new WFC3/IR dark calibration files were created for all allowed observing modes ([WFC3 ISR 2019-04](#)). Unlike previous superdark releases, the values of any pixels exhibiting time-dependent behavior are updated on a cycle-by-cycle basis, which allows for better calibration of those ~3.5% of detector pixels that were found to have fluctuating dark levels ([WFC3 ISR 2019-03](#)). The new superdarks also incorporate an additional four years of dark calibration data since the previous release, which reduces the errors associated with the superdark values by 8.9-67%, with an average decrease of 30% (see [WFC3 ISR 2019-04](#)), depending on the observing mode. The superdarks were installed in 2019 in the **calwf3** pipeline along with the new accompanying time-dependent WFC3/IR bad pixel tables ([WFC3 ISR 2019-03](#); see also [STAN 29](#) from June 2019); all prior data were reprocessed with the new reference files. For drizzled data products, a new **AstroDrizzle** parameter table should be used together with these new reference files in order to recover the stable hot pixels ([WFC3 ISR 2019-05](#)).

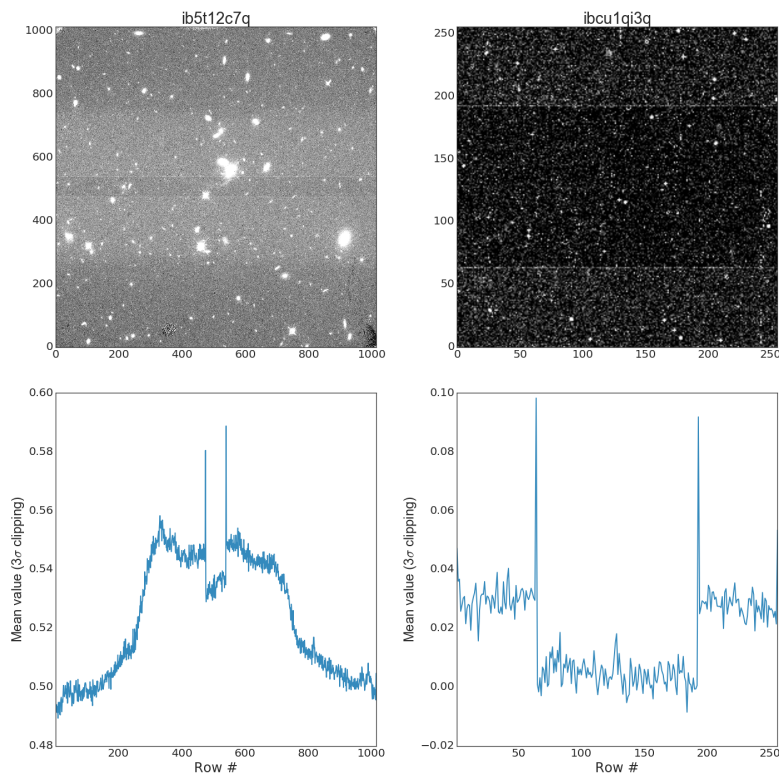
Observers with archival data prior to May 2019 may re-request their data from the MAST portal to obtain data reprocessed and redrizzled with the new WFC3/IR superdark, bad pixel, and **AstroDrizzle** reference files.

7.4.2 Banding

Banding occurs when an IR observation (or observations) in a smaller readout format immediately precedes an observation in a larger subarray or a full-frame. The banded images exhibit a rectangular region containing pixels whose brightness levels are offset by typically ± 3 -5 DN from values in the rest of the image. The band is centered vertically in the larger (second) image, extending all the way across the image horizontally, and has a height equal to the height of the smaller (first) subarray observation, as illustrated in [Figure 7.3](#) (top panels). The banded region is outlined on top and bottom by single rows of pixels with even brighter levels. The bottom panels of [Figure 7.3](#) show the respective vertical brightness profiles. For more information, please refer to [WFC3 ISR 2011-04](#) and the last paragraph of [Section 7.4.4](#) of the [WFC3 Instrument Handbook](#).

Figure 7.3: Examples of banded images.

Top Left: 64-pixel-high band in a SPARS50 full-frame science image. Top Right: 128-pixel-high band in a SPARS10 256 × 256 subarray dark calibration image. Bottom panels: 3-sigma clipped robust mean along the x-axis of the two images in the top panels. Note the central banded regions and the two higher spikes from the rows that bound them.



Mitigation

- ✔ The best mitigation is during observation planning: banding can be prevented by avoiding mixed aperture sizes within the same orbit. If a variety of aperture sizes are necessary, the observations should be sequenced from largest to smallest aperture sizes.

Some early dark reference files exhibited banding which imprinted the feature onto science images during the calibration process. Observations retrieved from the archive are automatically reprocessed with the latest calibration files (which are now band-free), so simply re-retrieving one's observations from MAST may solve the problem. If recalibration does not resolve the problem, it may be possible to subtract the features as an additional 'sky' component.

7.5 Blobs

[7.5.1 Blob catalogs](#)

[7.5.2 Blob Map](#)

[7.5.3 Blob Correction Strategies](#)

Small features called “blobs” are present in all WFC3/IR images and may be noticeable in observations with high background or in observations of large, extended objects. These artifacts affect ~2% of IR detector pixels and are caused by de-focused particulates on the surface of the mirror mounted on the Channel Select Mechanism (CSM), which is used to switch between UVIS and IR channels. Blobs are typically 10-20 pixels in diameter and can reduce the incoming light by up to ~30% in their central pixel, although the effect is typically much less. Since the CSM mirror is moved out of the beam in order for light to enter the UVIS channel, blobs do not appear in UVIS images (although UVIS data contain a different set of artifacts referred to as “droplets” (cf. [Section 5.4.1](#))).

New delta-flat reference files were delivered in October 2020 to correct for blobs during **calwf3** processing for six IR filters: F098M, F105W, F110W, F125W, F140W, F160W (cf. [Section 7.5.3](#) and [Section 7.8.4](#)). These flats correct for the majority of blobs appearing in IR images through July 2018 (see [Figure 7.10](#) and [WFC3 ISR 2021-10](#) and include 148 of 153 total known blobs cataloged in [WFC3 ISR 2018-06](#)).

7.5.1 Blob catalogs

Regions of low-sensitivity were first noticed in IR flat fields acquired during ground testing ([WFC3 ISR 2008-28](#)) and were described as 'small particulate features with reduced QE' (quantum efficiency). The authors recommended that on-orbit images be used to properly flat field these features, which were expected to shift position following the in-flight IR channel optical alignment.

Blobs were observed in inflight images shortly after WFC3 was installed on HST, and [WFC3 ISR 2010-06](#) provided an initial catalog of position and radius for 19 blobs appearing through March 2010. This report also describes a set of calibration test data which confirmed that blobs are located on the CSM and not on the IR detector itself. [WFC3 ISR 2012-15](#) expanded the catalog to 40 total blobs appearing through July 2012 and included a new estimated 'date of appearance' and blob 'color' which reported the strength of the blob in different filters. The blobs were found to be more opaque to bluer light and their effect is therefore slightly stronger in the F105W and F125W filters compared to the F160W filter.

Starting in mid-2010, calibration observations of the dark Earth were used to identify many smaller and weaker blobs which were not apparent in science images. By using airglow as a uniform sky background illumination in the F153M filter, the catalog was expanded to 127 blobs through March 2014. [WFC3 ISR 2014-21](#) describes the procedure for detecting blobs and for measuring their 'strength' based on the integrated sensitivity loss over their total area. Identification numbers were assigned to each blob in the three strength categories. Of these, only the 46 strongest blobs (losses exceeding 2.6%) were included in a new set of time-dependent bad pixel tables, with a separate table for each unique appearance date. The remaining 81 weaker blobs were cataloged but were not assigned any bad pixel flags.

The blobs have been increasing in number monotonically in time; no blob, once it appears, has disappeared. The rate of appearance was higher in the first year after launch (2009) than it was later. To monitor the appearance of new blobs, the schedulers insert dark Earth flats into appropriate gaps in the observing schedule, providing typically ~2-3 observations per week (e.g. WFC3 program [CAL-17371](#)). The catalog was updated again in 2018 ([WFC3 ISR 2018-06](#)) to include four additional years of calibration data ([WFC3 ISR 2018-06](#)). For convenience, blob IDs were renumbered according to their estimated 'Appearance Window' rather than by their strength, so that new blobs may simply be assigned the next consecutive number. **The catalog in this report ([WFC3 ISR 2018-06](#)) is continually updated when new blobs appear.** As of May 2024, there are 153 known blobs, with 58 classified as strong ([Figure 7.4](#)).

Since all new blobs are manually flagged by a WFC3 team member, we aimed to automate this process using machine learning. A convolutional neural network (CNN) was built, trained, and evaluated to classify if a new blob was present in 256x256 cutouts of IR images ([WFC3 ISR 2021-08](#)). This first step in automatic anomaly detection explored how machine learning can be used in astronomical operations, especially in the dawn of astronomy's era of big data. The model's accuracy was 91%, and is available for use on the [DeepWFC3 GitHub repository](#).


7.5.2 Blob Map

The CSM mirror is near to and slightly tilted with respect to the telescope's focal plane. As a result, blobs' radii systematically increase from the detector's upper right corner (where they are nearly in focus, with radii less than ~4 pixels) toward the lower left corner (where they have radii of ~13 pixels). [Figure 7.4](#) provides a map of blobs on the IR detector, color coded by their strength, where green and red circles highlight the strongest blobs flagged in the bad pixel tables.

Possible overlaps between known blobs and the IR grism aperture reference positions and the IR dither patterns have been investigated ([WFC3 ISR 2017-16](#)). Apertures are designed to place the science target on a cosmetically clean area of the IR detector. Similarly, dither patterns are designed to mitigate cosmetic defects by rarely (ideally never) placing such targets on known defects. As reported in ([WFC3 ISR 2017-16](#)), only two potential overlaps were found but no changes were made to the defined apertures or dither patterns because 1) one of the overlaps occurs with a dither /aperture combination that is most commonly used for wide-field surveys/mosaics and 2) the other overlap is ~9 pixels from a blob that has a radius of ~10 pixels i.e. an already conservative distance. Since blobs accumulate over time, the originally defined apertures and dither patterns are periodically checked to verify that they still accomplish their goals for science observation.

Figure 7.4: Blobs on the WFC3 IR detector

Detector map (1024x1024 pixels) showing 153 blobs appearing through May 2024. Blobs are detected using observations of the dark Earth in F153M, normalized by a stack of internal flats which do not contain blobs. The ratio image has been subtracted from 1.0 so that blobs appear as positive deviations. The position and radius of each blob are indicated with color-coded circles according their strength from weak to strong: cyan, blue, green, red, respectively. Starting in July 2014, time-dependent IR bad pixel tables were delivered to flag the strongest blobs (green and red) in the DQ array of calibrated FLT images. In October 2020, a new set of delta-flats were delivered to correct for the blobs in six filters and populate DQ flags in calibrated images for all known blobs appearing through July 2018, regardless of their strength. This increased the fraction of flagged pixels (DQ=512) from 1.1% to 2.3%.



7.5.3 Blob Correction Strategies

Appropriate dithering during Phase II planning (e.g. PATTERN='WFC3-IR-DITHER-BLOB' or a custom ~5" dither) can permit the cleaning of blobs from combined images during advanced reprocessing, as described in [WFC3 ISR 2010-09](#). Nominal drizzled images (*drz.fits) produced by MAST ignore the blobs as flagged in the DQ array of calibrated (*flt.fits) images, assuming that not every program utilizes a 'blob dither'. Even with effective dithers (+/- 15 pixels), the images cannot maintain pixel-phase coherence. Multiple sets of four-point dithers are necessary to compensate for both blobs and pixel-phase coverage ([WFC3 2023-05](#)).

An alternative method of correcting for blobs is to apply a delta-flat correction image, the first of which was derived from dark Earth observations ([WFC3 ISR 2014-21](#)). This correction was initially recommended for cosmetic purposes only; the authors suggested that while it may improve the photometry of extended sources, it was not expected to work well for stellar photometry due to the small optical cross section of the blobs. However a follow-up study that applied a delta-flat to dithered exposures in a crowded region of 47 Tucanae determined that applying a blob flat field is effective for improving stellar photometry in blob-affected areas ([WFC3 ISR 2015-06](#)).


In October 2020, a new set of delta-flat (D-flat) reference files were delivered to CRDS to correct for blobs during pipeline processing for six IR imaging filters: F098M, F105W, F110W, F125W, F140W, F160W ([WFC3 ISR 2021-10](#)). These flats were computed by stacking observations of sparse astronomical targets after the appearance of each blob to produce a deep image of the sky background (see [Section 7.8.4](#)). While the bad pixel tables continue to carry DQ flags for strong blobs (green and red in [Figure 7.4](#)), the delta-flats now provide DQ flags for all cataloged blobs, regardless of their strength. This increases the total number of flagged pixels (DQ= 512, 'bad in flat') from 1.1% to 2.3% in calibrated IR (non-grism) images. The D-flats correct for 148 of the 153 known blobs at 49 unique "appearance dates." Additional D-flats will be delivered for the most recent blobs once more inflight data is available.

The new D-flats (*_dfl.fits) are intended to be used together with an updated set of P-flats, which improve the 'pixel-to-pixel' sensitivity calibration with wavelength for all imaging filters (see [Section 7.8.3](#)). As part of the flat field calibration step, **calwf3** uses the product of the P-flat and D-flat reference files to correct for spatial variations in the detector sensitivity with time. The D-flats represent an average of the CSM angle positions sampled with in-flight images, and any given image may contain blobs at an off-nominal position, producing the characteristic dark and light residuals in [Figure 7.5](#). In this case, users may elect to ignore both the delta-flat correction and the additional DQ flags for weaker blobs by setting the `DFLTFILE` keyword to 'N/A' in the header of the raw image and then reprocessing with **calwf3**. Details on manual reprocessing of WFC3 data are provided in [Section 3.5](#).

After comparing a large set of dark Earth flats acquired over time, slight offsets in the blob positions were discovered and attributed to changes in the CSM position when switching from UVIS to IR. To explore this effect, a set of seven delta-flats were computed by stacking images at slightly different positions, in increments of 0.004 degrees and spanning a total range of 0.024 degrees of rotation ([WFC3 ISR 2021-10](#)). This total range corresponds to a diagonal shift of ~1 pixel for blobs near the center of the detector. The left panel of [Figure 7.5](#) shows a small region of the 2014 delta-flat, highlighting two blobs with radii ~10 pixels and a total separation of ~100 pixels. To illustrate the effect of flat-fielding an image acquired at an off-nominal CSM position with a delta-flat from the nominal position, the center and right panels of [Figure 7.5](#) show the ratio of two delta-flats with CSM angles differing by 0.012 and 0.024 degrees, respectively. Due to the unpredictability of the <1 pixel blob shifts, it is not possible apply shift-dependent blob delta-flats in the automated calibration pipeline; these would need to be applied manually by the user during recalibration. Alternatively, observers concerned about the blobs can consider expanding the outer rim of the existing flagged blobs in the DQF; while this will unavoidably flag some good pixels, it will cover the extent of shifts seen to date.

Figure 7.5: Blob flat residuals due to CSM offsets

(Left) F153M delta-flat displayed with a scale from 0.90 to 1.05. This zoomed region shows the position of blob #116 (lower) and #33 (upper) observed at the nominal CSM position and separated by ~ 100 pixels. (Center, Right) The ratio of two delta-flats with CSM angles differing by 0.012 and 0.024 degrees, displayed with a scale from 0.97 to 1.03. Figure taken from [WFC3 ISR 2021-10](#).



7.6 Detector Nonlinearity Issues

7.6.1 Nonlinearity Calibrations


7.6.2 Non-Zero Zeroth Read Correction for Bright Sources

7.6.1 Nonlinearity Calibrations

HgCdTe detectors, such as that in the WFC3/IR channel, are inherently non-linear in their response to incident photons. [Figure 7.6](#) shows the measured signal up the ramp for a single pixel during a flat-field ramp. The black diamonds show that the measured signal is clearly non-linear as the signal increases. The red line is a linear fit to reads at low signal levels, extended out to the end of the ramp. The difference between these two lines (normalized by the measured signal values) represents the amount of non-linearity in the measured signal. By measuring this difference for each pixel using a set of flat-field ramps, the non-linearity behavior of each pixel is characterized, and a correction produced. This correction takes the form of a 3rd order polynomial, as seen in the equation below, that relates the measured and idealized signal and is applied by **calwf3** in the `NLINCORR` step of the processing. Further details can be found in [WFC3 ISR 2008-39](#).

The non-linearity correction in **calwf3** currently uses quadrant-averaged polynomials derived from ground testing data. Only small improvements in the nonlinearity calibration are anticipated from using on-orbit flat fields, and noticeable differences have been measured between nonlinearity corrections applied to flat fields compared to stars. For those reasons, the methods and results described in [WFC3 ISR 2014-17](#) have not been implemented.

Figure 7.6: Measured non-linear response of an IR pixel.




7.6.2 Non-Zero Zeroth Read Correction for Bright Sources

As described in [Section 7.3](#), the first non-destructive readout after the detector reset provides the reference bias level for each pixel in a WFC3/IR science image. This image is referred to as the 'zeroth readout' or 'zeroth read'. Due to the timing patterns in use, this read is collected at a finite time after the detector is reset. Because the IR channel has no shutter, when a bright source is present in the field, a non-negligible amount of charge may accumulate on the detector between the time of the detector reset and the zeroth read. This implies that after the subtraction of the zeroth read from all subsequent reads in a ramp, the pixels at the location of the bright source will be under-reporting the amount of signal they were subjected to during the ramp. While this has no effect on the **calwf3** calculation of the signal rate, it can lead to errors in the non-linearity correction step, since this correction is based on the absolute signal level.

Prior to March 2011, in order to correct for this situation, the **ZSIGCORR** step in **calwf3** (very similar to that in the **calnica** software for NICMOS) computes an estimate of the number of counts for each pixel in the zeroth readout of a **MULTIACCUM** ramp, based on the count rate measured between the first two reads. After March 2011, the bias is read from a linearity reference file and subtracted from the zeroth read of the file being processed. (see [Section 3.3.2](#) for details). This information is then used in the non-linearity correction (**NLINCORR**) step to estimate the absolute signal in each pixel and apply the appropriate linearity correction and saturation checks for that signal level.

7.7 Count Rate Non-Linearity

Previous HgCdTe detectors on HST have suffered from a count-rate dependent non-linearity, motivating a search for a similar effect on the WFC3/IR detector. Utilizing a color term to account for differences in the observed bandpasses, statistically-significant non-linearity was detected in WFC3 /IR photometry, a few times smaller than that observed in NICMOS.

 The direction (or sign convention) of the count-rate non-linearity is such that photometry of faint (i.e., sky dominated) sources calibrated with the published WFC3/IR zeropoints (derived from bright standard stars) will appear 0.04 +/-0.01 mag too faint. The calwf3 pipeline does not correct WFC3 IR data for count-rate non-linearity; an appropriate correction to stellar photometry will need to be made manually by the observer, if desired.

Combining all available measurements from direct images and those from the WFC3 grism, [WFC3 ISR 2019-01](#) provides a consistent and improved characterization of WFC3 IR's count-rate dependent non-linearity of 0.0075 +/- 0.0006 per dex¹, with no apparent wavelength dependence, measured across a dynamic range of more than a million in stellar apparent brightness. It compares the latest and best estimates to prior ones which also are summarized in the following paragraph. Prior results are consistent with but approximately four times less precise than the newer results, and hence we recommend the values in [WFC3 ISR 2019-01](#).

The work prior to 2019 consisted of two studies. An initial measurement of the effect was made by comparing the photometry of star clusters observed over a wide dynamic range and at overlapping wavelengths in WFC3/IR and NICMOS and/or ACS-WFC (see [WFC3 ISR 2010-07](#)). Based on 235 stars in 47 Tuc observed with WFC3/IR in F110W and F160W and in similar bandpasses in NICMOS Camera 2, the non-linearity of WFC3/IR was measured at 0.011 +/- 0.0023 and 0.010 +/- 0.0025 mag per dex, respectively, over a range of 10 magnitudes (4 dex). An independent measurement utilizing 1390 stars in NGC 3603 observed with ACS/WFC F850LP and WFC3/IR F098M yielded a very similar result, 0.010 +/- 0.0033 mag/dex. The consistency of the results from two different comparison detectors of different technology suggests the result is robust. Furthermore, an additional analysis using 24 bright stars observed by WFC3 and by 2MASS found that the count rate non linearity of the WFC3 IR detector is less than or equal to 0.01 mag per dex ([WFC3 ISR 2011-15](#)).

¹ One dex is a factor of 10. Thus 1 mag per dex means a difference of 1 magnitude per factor of 10 in brightness.

7.8 IR Flat Fields

- [7.8.1 Ground Flats](#)
- [7.8.2 On-orbit Corrections](#)
- [7.8.3 Filter-Dependent Sky Flats](#)
- [7.8.4 Blob Flats](#)

Standard **calwf3** processing allows for up to three types of flat field reference files for calibration: 1) a P-flat, which corrects for high-frequency 'pixel-to-pixel' variations in the spatial sensitivity as a function of bandpass, 2) a delta-correction or D-flat which accounts for any 'pixel-to-pixel' changes in the P-flat with time, and 3) a low-frequency L-flat which provides a smooth correction to the P-flat. When two or more types of flats are provided, they are multiplied together by **calwf3** to form a combined flat field image.

Prior to 2020, only P-flats were used in the **calwf3** pipeline for IR images. These flats were first acquired in 2008 during ground testing (see [Section 7.8.1](#)), and then updated in 2010 using sky flats to derive a smooth, in-flight correction to the ground flats (see [Section 7.8.2](#)). In 2020, new sky flats were computed in using over 10 years of on-orbit imaging. For six filters (F098M, F105W, F110W, F125W, F140W, and F160W), the sky flats have sufficient signal-to-noise to fully replace the ground-based solutions. For the remaining nine IR imaging filters, new P-flats were computed by multiplying each ground flat by an interpolated in-flight correction based on sensitivity residuals from sky flats at similar wavelengths (see [Section 7.8.3](#); [WFC3 ISR 2021-01](#)). At the same time, new delta-flats were computed for six filters to correct for low-sensitivity artifacts known as 'blobs' appearing in the IR observations over time ([WFC3 ISR 2021-10](#)). A detailed description of IR blobs is provided in [Section 7.5](#) and a discussion of the new delta-flats in [Section 7.8.4](#).


Because of the effect of geometric distortion, the area of the sky seen by a given detector 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 a uniform illumination, the flat fields for both UVIS and IR include the effect of the variable pixel area across the field. A consequence of dividing by the flat fields is that two stars of equal brightness positioned on different regions of the detector will not have the same count rate after the flat fielding step. Thus, point source photometry extracted from calibrated FLT (*flt.fits) images must be multiplied by an effective [pixel area map](#), as described in [Section 9.1.1](#). Alternatively, this correction is accounted for by **AstroDrizzle**, which uses the geometric distortion solution to correct all pixels to equal areas, i.e. for drizzled DRZ (*drz.fits) images, photometry is correct for both point and extended sources at any position on the detector.

7.8.1 Ground Flats

During the spring of 2008, flat field images for the IR channel were produced in the laboratory during the third and final thermal vacuum campaign (TV3) using the CASTLE Optical Stimulus (OS) system ([WFC3 ISR 2008-28](#)). The CASTLE is an HST simulator designed to deliver an OTA-like external beam to WFC3. It can provide either point-source or flat-field illumination in either monochromatic or broadband mode. Flat fields with the OS tungsten lamp were taken with the detector at its nominal operating temperature. These were acquired using the SPARS10 sample sequence, with varying numbers of readouts (samples) per exposure, chosen to obtain a signal of about 60,000 electrons per pixel in the final read. Four exposures were obtained per filter, providing a mean signal-to-noise ratio of ~500, yielding flat field corrections accurate to ~0.2%.

Figure 7.7 shows the science and error array of the F125W flat field acquired in TV3. A faint, extended cross-pattern in the flats is an artifact of the optical stimulus support structure and would therefore need to be corrected using inflight data. Additionally, the blob positions were expected to shift once an on-orbit optical alignment was performed. While blobs are not visually apparent in the flat field science arrays, they are easily identifiable in the error arrays due to a realignment which was performed midway through the IR ground flat calibration. As a result, a single blob appears in two slightly offset positions with $\sqrt{2}$ larger errors than the surrounding regions of the P-flat. Similarly, the cross-pattern appears in the flat field error arrays at two unique positions.

Figure 7.7: Left: Science array of the F125W ground flat reference file (sca20262i_pfl.fits) displayed with a linear stretch from 0.85 (black) to 1.05 (white). Right: Error array of the F125W ground flat shown with a stretch from 0 to 0.5%. Due to an optical realignment that was performed midway through the flat calibration program, blobs (particulates on the channel select mechanism (CSM)) appear as pairs of white circles with $\sqrt{2}$ larger errors. A broad cross-shaped residual caused by the optical stimulus support structure also appears in two distinct positions.



To test the large-scale uniformity of the detector response as provided by the ground flat calibration, dithered observations of the globular cluster Omega Centauri were acquired during Servicing Mission Orbital Verification, just after WFC3 was installed on HST (program [11453](#)). By placing the same group of stars on different regions of the detector and measuring relative changes in brightness, low-frequency spatial variations in the response of the detector were measured, with average photometric errors of $\pm 1.5\%$ attributed to differences between the CASTLE OS and the inflight illumination ([WFC3 ISR 2009-39](#)).

Omega Centauri was revisited in Cycles 17 and 18 (2009-2011, programs [11928](#) and [12340](#)) at additional dither positions and orientation angles. Local variations in the detector response were computed using the same software used to derive the ACS inflight correction ([ACS ISR 2003-10](#)). For application to WFC3, the IR detector was divided into a 16×16 grid and a unique solution was computed for each grid point, representing the deviations from unity. L-flat corrections were derived for five filters (F098M, F110W, F125W, F139M, and F160W) and showed rms variations of $\sim 1.1\%$ and peak-to-peak residuals of $\pm 2.5\%$. The inflight sensitivity residual for the F160W filter is shown in the lower-right panel of [Figure 7.8](#) and highlights the cross-pattern seen in the ground flats. Due to the relatively small number of stars in the calibration field, these low-order corrections were insufficient to correct for the complex spatial residuals. Instead, deep images of sparse astronomical targets were used to produce sky flats in order to correct for differences in the inflight sensitivity, as described in the next section.


7.8.2 On-orbit Corrections

Combining IR external imaging data acquired between September 2009 and December 2010, sky flats were used to compute the required correction to the ground flats ([WFC3 ISR 2011-11](#)). In short, all science observations taken with the F098M, F105W, F110W, F125W, F140W, and F160W filters with exposure times exceeding 300s were calibrated using the ground flats and then combined in each filter, after masking out objects. This produced a high-quality sky image for the F160W filter, and somewhat noisier sky images for F098M and F125W. For the remaining three filters, an insufficient number of input science frames were available to derive sky images with adequate signal-to-noise. Comparing the resulting sky image between filters, no clear indication of any color dependence of the low-frequency structure was found thus, a 'gray' sky image was constructed using all available data. The top left panel of [Figure 7.8](#) shows the gray correction, which is similar in structure to the L-flat derived from stellar observations (bottom right).

Similar differences between the CASTLE optical stimulus and inflight illumination patterns were also apparent from calibration observations of the moonlit Earth limb (program [11917](#)). One of those Earth flats, acquired with the F105W filter and flat fielded with the ground flat reference file, is shown in the top right panel of [Figure 7.8](#). Both the shape and the amplitude of the low-frequency residuals are very similar to the image derived from sky observations. To compare the sky image with the photometric residuals from Omega Centauri, the sky image was binned to the same 16×16 grid. The bottom left panel of [Figure 7.8](#) shows the 16×16 sky image grid for comparison with the bottom right panel which shows the F160W stellar L-flat. Similar low frequency structures are present in both images, i.e. the results are consistent, albeit with lower signal-to-noise in the stellar L-flat.

In December 2010, the ground flat reference files ('s*pfl.fits') were replaced with an updated set of P-flats ('u*pfl.fits'). These were constructed by multiplying the ground flats with the gray sky image, rather than generating separate L-flat reference files. The gray correction was smoothed to reduce noise, causing the blobs to be slightly larger and reduced in strength. As done for the ground flats, each new P-flat was normalized to 1.0 over the image section [101:900,101:900], which excludes areas of the detector known to contain anomalies such as the "Death Star" and the "Wagon Wheel" (see [Section 7.9.2](#)). After recalibrating the Omega Centauri observations with the new P-flats, the photometric accuracy was found to be $\sim 0.5\%$ (with peak-to-peak variations of -1.5 to $+1.6\%$) when avoiding the outermost 128 pixels of the edge of the detector. Within the "Wagon-Wheel" region and the edge of the detector, the photometric uncertainties increased to $\sim 0.8\%$ (with peak to peak variations of -2.0 to $+1.9\%$).

Figure 7.8: Top left: Filter-independent 'gray' correction derived from on-orbit sky flats and applied to the ground-based P-flats in 2010. Top right: residuals for F105W after application of the ground flat to an image of the moonlit Earth limb. Bottom left: 16 x 16 binned image of the 'gray' sky correction for comparison with the stellar L-flat. Bottom right: L-flat correction for the F160W filter based on aperture photometry of Omega Centauri acquired at multiple dithered positions (from [WFC3 ISR 2011-11](#)). The flat field residuals in each panel are displayed with a linear stretch from 0.98 to 1.02.



7.8.3 Filter-Dependent Sky Flats

Sky flats were computed from a large set of in-flight science observations acquired over more than 10 years of on-orbit data. Delivered to the Calibration Database Reference System ([CRDS](#)) in October 2020, these new P-flat reference files ('4*pfl.fits') correct for wavelength-dependent sensitivity residuals of $\pm 0.5\%$ in the central 800x800 pixel region of the detector and up to 2% near the detector edges, depending on filter ([WFC3 ISR 2021-01](#)).

For the six most commonly used filters (F098M, F105W, F110W, F125W, F140W, and F160W), the combined sky image is now deep enough, with errors less than or equal to 0.2%, to fully replace the ground flat reference files. The new P-flats improve the 'pixel-to-pixel' sensitivity calibration and correct for differences in the IR gain across the four amplifiers. **Unlike the 2010 solutions, the P-flats no longer include the blobs, which are now corrected for by a separate set of D-flat reference files (see next Section).**

For the remaining nine IR imaging filters, wavelength-dependent sensitivity residuals from the six most frequently used filters were used to correct the P-flats acquired on the ground. An L-flat correction image for each filter was derived by interpolating the pixel values of the flat field residual images (sky flat divided by ground flat) from the two filters closest in wavelength. These residuals were smoothed to avoid degrading the filter-dependent pixel-to-pixel sensitivity and were then multiplied by the original ground flat.

To check the stability of the IR detector sensitivity, spatial-scanning photometric observations of a sparse region of M35 were acquired over a ~6 year baseline ([WFC3 ISR 2021-05](#)). To avoid the effects of image persistence, the target was dithered by large offsets to place the stars at multiple, 'clean' (no prior illumination) positions across the central 512x512 subarray. The photometry of each star shows an intrinsic scatter of order 0.5%, suggesting either residual errors in the flat field, subtle effects due to the scanning mode, and/or systematic effects related to the detector illumination history. For more details on IR photometry errors, see [Section 7.11](#).

7.8.4 Blob Flats

A new set of delta-flats (D-flats), delivered together with the new P-flats in October 2020, now correct for 'blobs' in the six most frequently used filters ([WFC3 ISR 2021-10](#)). The blob delta-flats were computed by stacking archival sky flat images after the appearance of each blob. Conversely, P-flats in blob-affected regions are computed by stacking images prior to each blob's appearance. D-flats with statistical errors < 1% are available for 148 (of 152 total) blobs, and flats for the most recent and newer blobs will be delivered as and when more inflight data are available. For a current status of blob monitoring and catalog, see [WFC3 ISR 2018-06](#).


The new D-flat reference files ('4*df1.fits') are intended to be used in combination with the most recent P-flats described in [Section 7.8.3](#). While there is only one set of P-flats (one for each of the 15 imaging filters), D-flats are available only for six filters at 49 unique blob appearance dates. The appropriate D-flat is selected according to the USEAFTER keyword in the flat field header corresponding to the observation date (DATE-OBS) keyword in the header of the raw image. During pipeline processing, **calwf3** multiplies the P-flat and the D-flat reference files together to form a combined flat field image. D-flat values are set to 1.0 everywhere except in the blob regions which contain a high-frequency correction to the P-flat. Since blobs do not change position (aside from the occasional <1 pixel shifts discussed at the end of [Section 7.5](#)), the D-flat for an individual blob does not change over time, and new blobs are simply added to the previous D-flat and assigned an updated USEAFTER date. Previously, only the strongest set of blobs were assigned a DQ flag (value=512, 'bad in flat') via the IR bad pixel tables. Now, the D-flats provide DQ flags for all cataloged blobs, regardless of their strength, increasing the total number of flagged flat field pixels from 1.1% to 2.3% in the DQ array of calibrated WFC3/IR (*flt.fits) images.

For blobs appearing in ground test data or in early inflight science images, sky flats will have little to no signal in the affected pixels. We therefore replaced their pixel values in the P-flats with the median P-flat value in a 20-pixel wide circular annulus surrounding each blob. For the early set of blobs (those present by September 2009), there are insufficient number of external images from which to compute a pixel-to-pixel flat, so the delta-flat includes a combination of the blob and the P-flat. While these blobs affect all inflight data, they were not included in the P-flats in order to make those reference files more useful for grism calibrations. A model of the multi-component dispersed IR sky background in G102 and G141 is described [WFC3 ISR 2020-04](#). Blobs appear in the Zodiacal light component as elongated shapes, similar to those of normal dispersed spectra but with lower signal due to their reduced sensitivity. Blobs are also apparent in the Helium I airglow component although offset from their nominal positions in imaging data due to the grism's dispersion.

The left panel of [Figure 7.9](#) shows the product of the F125W P-flat and D-flat for 148 blobs appearing through July 2018. The right panel shows the F125W filter-dependent residual with respect to the 2010 'gray' inflight correction, which was heavily weighted toward the F160W filter as that had the most inflight images from which to form a sky flat. The residuals are similar in both shape and amplitude to those in Figure 2 of [WFC3 ISR 2011-11](#) which shows the ratio of the F125W and F160W sky flat computed from early inflight data.

Figure 7.9: Left: Combined F125W flat field, computed as the product of the P-flat and D-flat (4ac1921ii_pfl.fits and 4ac1825sj_dfl.fits, respectively) which corrects for blobs appearing through July 2018, shown with a stretch from 0.8 to 1.1. The large cross-pattern visible in the left panel of Figure 7.7 (an artifact of the optical stimulus used in ground testing) is now removed.



Right: Ratio of the combined F125W flat field (P-flat times D-flat) and the 2010 P-flat (uc72113qi_pfl.fits), shown with a stretch from 0.98 to 1.02. Filter-dependent residuals are present at ~0.5% in the center of the detector and ~1% in the "Wagon wheel" feature in the lower right corner. Blobs arising inflight after 2010 appear as dark residuals, while blobs included in the 2010 P-flat appear as dark residuals encircled with a white ring (the result of smoothing in the early in-flight correction).



A sample drizzled product from MAST is shown in [Figure 7.10](#), comprised of four FLT images in the F140W filter. The drizzled image has a combined exposure time of ~ 5200 seconds and a total background of $\sim 10,000$ electrons. The top panel shows the drizzle-combined image (ibkdf2010_drz.fits) from FLT images calibrated with the 2010 flats. Residuals in the sky background are similar to the wavelength-dependent structure seen in the right panel of [Figure 7.9](#). The drizzle-combined image from FLT images calibrated with the new (2020) P-flat and D-flat reference files is shown in the bottom panel in [Figure 7.10](#). The large-scale uniformity of the background is significantly improved and regions of the detector impacted by blobs appear visually clean.

For precise photometry in regions affected by blobs, observers planning WFC3 IR programs should consider utilizing a 'blob' dither in the Phase II proposal implementation (e.g. PATTERN= 'WFC3-IR-DITHER-BLOB', see also WFC Instrument Handbook) in order to permit the cleaning of blobs from drizzle-combined images. Calibrated FLT images can be reprocessed with **AstroDrizzle**, rejecting any pixels flagged with a value of 512 in the DQ array, and 'filling in' with clean regions of the sky from the associated dithered exposures. For more detail on blob correction strategies, see [Section 7.5.3](#).

Figure 7.10: MAST drizzled dataset (ibkdf2010_drz.fits), in electrons/sec, observed in the F140W filter with a total exposure time of 5,200 seconds and a cumulative background of 10,400 electrons. The image in the top panel was calibrated with the 2010 flats while the image in the bottom panel was calibrated with the 2020 P-flats (which correct for filter-dependent residuals) and D-flats (which correct for IR blobs).



7.9 Pixel Defects and Bad Imaging Regions

- [7.9.1 Bad Pixels](#)
- [7.9.2 Off-nominal Detector Regions](#)
- [7.9.3 Dead Pixels](#)
- [7.9.4 Bad Zeroth Read Pixels](#)
- [7.9.5 Unstable Pixels](#)
- [7.9.6 Snowballs](#)

7.9.1 Bad Pixels


Data from a variety of on-orbit calibration programs have been used to identify bad pixels and regions on the IR detector. All pixels found anomalous enough to potentially impact data analysis results are flagged and listed in the current bad pixel table reference file. While these pixels are listed in the bad pixel table and propagated into the data quality (DQ) arrays of the observer's ima and flt files produced by **calwf3**, the flagged science pixels will still have **calwf3**-calculated signals and signal rates in the science (SCI) arrays. It is therefore important for observers to examine the DQ arrays in their data, and use them to identify which pixels to ignore in any post-**calwf3** data reduction and analysis (such as **AstroDrizzle**).

! *It is the responsibility of the observer to determine which types of bad pixels are acceptable and which are to be avoided during data analysis. Use the DQ arrays for this purpose.*

During **calwf3** processing, the bad pixel table is imprinted onto the data quality (DQ) array associated with each MULTIACCUM ramp. Using the pixel values in the DQ array (see [Table 3.3](#)) observers can tailor the types of bad pixels used in their analysis. The current default for IR channel data is for **Astrodrizzle** to ignore (treat as bad) all pixels with any flag except 512 and 64. The 512 flag is used to identify pixels affected by blobs, while the 64 flag is currently not attached to any type of bad pixel and not used. Observers can change the `driz_sep_bits` parameter within **Astrodrizzle** to modify this default and adjust which types of bad pixels to use or ignore.

A collection of time-dependent WFC3/IR bad pixel tables are generated annually in order to reflect the observed behavior changes of ~3.5% of the detector pixels (Figure 7.9., [WFC3 ISR 2019-03](#); [WFC3 ISR 2024-02](#)). These new tables also incorporate enhanced bad pixel detection methods that identify and flag bad-in-zeroth-read, dead, hot and unstable pixels on a cycle-by-cycle basis as well as the IR blobs as they appear on the instrument. No evidence has been found for a long-term trend in the number or locations of the bad-in-zeroth-read and dead pixels. However, evidence for a linear increase of ~200 new stable hot pixels per cycle have been found and the number of unstable pixels can change by up to ~4000 pixels per cycle. The bad pixel tables used by the **calwf3** pipeline along with the time-dependent WFC3/IR superdarks ([WFC3 ISR 2019-04](#)). For drizzled data products, a new **AstroDrizzle** parameter table can be used together with these new reference files in order to recover the stable hot pixels ([WFC3 ISR 2019-05](#)). Observers with archival data prior to May 2019 may request their data from the MAST portal to have their data automatically reprocessed and redrizzled with the most up-to-date reference files.

Figure 7.9: An IR Bad Pixel Mask from 2023. Black pixels are flagged in the bad pixel table.



7.9.2 Off-nominal Detector Regions

As noted in [Section 5.7.7](#) of the *WFC3 Instrument Handbook*, there are several coherent features on the IR detector composed of poorly performing pixels. The “Death Star”, the large circular feature at the left bottom edge of the detector, is a collection of pixels which have extremely low QE and exhibit unstable behavior. Unbonded pixels in the upper left and right corners and along the top edge of the detector are unresponsive to illumination. These pixels have all been flagged as dead in the bad pixel mask and should be avoided during analysis. In the lower right corner of the detector is the feature known as the “Wagon Wheel”, a collection of pixels with quantum efficiencies 25% to 50% below normal. This does not mean that these pixels cannot be used during data analysis, but sources in this region will have a lower signal-to-noise ratio than they would elsewhere on the detector, an aspect reflected in the error arrays of `calwf3` calibrated data. A more detailed description of these detector regions is given in [WFC3 ISR 2008-28](#).

7.9.3 Dead Pixels

These are pixels with a very low quantum efficiency i.e. they measure little or no signal when illuminated. In addition to the dead pixels found through the analysis of on-orbit data, the pixels comprising the “Death Star” are considered dead. In total, about 3300 pixels are flagged as dead (0.3% of the detector’s light-sensitive pixels), marked with a 4 in the bad pixel table (see [WFC3 ISR 2010-13](#) and [WFC3 ISR 2019-03](#) for details). Other than those pixels within the “Death Star” and the small areas of unbonded pixels along the FOV edge, dead pixels are scattered randomly across the detector. No long-term trend is evident in the dead pixel population. We recommend that observers ignore any pixel marked as dead.

7.9.4 Bad Zeroth Read Pixels

These are pixels which exhibit anomalous signals in the zeroth read of a data ramp, usually due to being shorted or unbonded ([WFC3 ISR 2003-06](#), [WFC3 ISR 2019-03](#)). As a result, many of the bad zeroth read pixels are also flagged as dead. By flagging bad zeroth read pixels in the bad pixel table, the pipeline adopts a conservative approach to bad pixel behavior. Historically, pixels with an off-nominal signal in the zeroth read displayed other non-nominal behaviors thus these pixels are clearly of questionable quality. As with all flavors of bad pixels, observers should determine whether or not using these pixels will have a significant impact on their analysis.

In a typical cycle, there are about 3100 pixels in total (~0.30% of the detector pixels) flagged as bad in the zeroth read (flag value of 8); no long-term trend is evident in the bad in zeroth read population. These pixels are largely concentrated in the areas of the “Death Star”, the upper corners, and along the quadrant boundaries of the detector ([WFC3 ISR 2019-03](#)).

7.9.5 Unstable Pixels

There are a total of about 15,000 unstable pixels (1.5% of all science pixels) on the IR detector, flagged with a DQ value of 32 ([WFC3 ISR 2019-03](#)). Due to the unpredictable behavior of these pixels, we recommend against including them in data analysis.

Unstable pixels display abnormal variations in signal in a set of nominally identical ramps. Initial studies early in the WFC3 mission found a wide range of unstable pixel behaviors: some pixels appear stable and repeatable in almost all ramps, but will measure appreciably different signal values in only one or two ramps. Other pixels display signal values that vary wildly from ramp to ramp in all observations of a data set ([WFC3 ISR 2010-13](#), [WFC3 ISR 2012-11](#)). A recent analysis of more than 10 years of data ([WFC3 ISR 2019-03](#)) identified and flagged unstable pixels (as well as bad-in-zeroth, dead pixels) on a cycle-by-cycle basis and delivered updated bad pixel tables to the calibration pipeline. About 3.5% of the IR science pixels have exhibited unstable behavior for some period of time since launch. The number of unstable pixels can change by up to ~4000 pixels per cycle.


For any epoch of WFC3/IR data retrieved from MAST after May 2019, the automated calibration pipeline images (flt, ima) will have unstable pixels flagged with 32 in the DQF. Observers working with such data can choose to reclaim calibrated stable hot pixels for their analysis by discarding only those hot pixels flagged as both hot (16) AND unstable (32); hot pixels which are stable are calibrated via the superdark calibration reference file. Observers working with drizzled data (drz) will find that the stable hot pixels (16) are by default reclaimed and treated as good; only pixels flagged as unstable (32) continue to be treated as bad ([WFC3 ISR 2019-05](#)).

7.9.6 Snowballs

Curious but relatively innocuous anomalies which occasionally appear in IR observations are "snowballs". These sources have an extended, fuzzy appearance in the data. Snowballs are transient, extended sources that appear in IR channel data at rates of roughly 1.2 snowballs per hour of exposure time. The snowballs are suspected to be caused by natural radioactivity within the detector itself. Specifically, alpha particles emitted by thorium and/or uranium at ~ 1 ppm concentration in the detector can explain the observed characteristics of the snowballs ([WFC3 ISR 2009-44](#)). Similar to the manner in which cosmic rays appear, the entire flux of a snowball is deposited into the detector's pixels instantaneously. A typical snowball affects about 10 pixels, depositing between 200,000 and 500,000 electrons, and saturating 2-5 pixels. [Figure 7.10](#) shows a 7×7 mosaic of snowballs gathered from ground testing and on orbit data. Since the snowball behavior mimics that of cosmic ray impacts, **calwf3** is able to remove snowballs from WFC3/IR data during standard pipeline processing. That, combined with low rate of occurrence, means snowballs should have a minimal impact on science observations. Further details on snowballs can be found in [WFC3 ISR 2009-43](#) (initial characterization) and in [WFC3 ISR 2015-01](#) (long-term characterization).

The snowball occurrence rate, intensity, shape, and variation over time provides important constraints on candidate species by their decay rates, and HST's WFC3/IR has by far the longest time baseline for analysis ([WFC3 ISR 2020-03](#)). The IR 'Snowball' occurrences in WFC3/IR have been analyzed from in-flight data since launch finding more than 10,000 snowballs over an 11 year baseline of observations, distributed roughly equally across that time period.

Figure 7.10: A mosaic of snowballs based on ground and on-orbit data.



7.10 Time-Variable Background

7.10.1 Scattered Earthlight

7.10.2 Metastable Helium 1.083 μm Emission Line

The background in the WFC3/IR channel is a combination of zodiacal light, scattered light from the bright Earth limb, and line emission at 1.083 μm from helium atoms excited by sunlight in the Earth's upper atmosphere. The strength of the zodiacal light depends on the orientation of the target with respect to the sun, which varies throughout the year but is effectively constant within a given exposure/orbit/visit. The scattered light and line emission components can vary within an orbit and even within a single exposure, and the line emission is filter-specific.

Scattered light can often be present for observations made when the limb angle, which is the angle between HST's pointing direction and the nearest limb of the bright Earth, is less than ~ 30 degrees ([WFC3 ISR 2002-12](#), [WFC3 ISR 2009-21](#)). The total amount of scattered light increases as the target-to-limb angle decreases throughout an orbit, producing a background which is both time- and spatially-variable. This primarily impacts the left side of the detector ([Figure 7.11](#)) and can affect all filters and both grisms.

The helium emission line is seen when the spacecraft leaves the Earth shadow and enters the illuminated atmosphere. The strength of the increased background depends on the observed path-length through the atmosphere, and as a diffuse source the helium emission shows no spatial structure in the F105W and F110W filters ([WFC3 ISR 2016-16](#)). For the G102 grism, the helium emission has a unique spectral signature and must therefore be corrected using a background model. A discussion of variable background subtraction methods for the IR grisms and an updated set of dispersed background models are provided in [WFC3 ISR 2020-04](#); Jupyter notebooks available for correcting the background are discussed later in this section.

Strong time variation in the background during a MULTIACCUM ramp can corrupt the `calwf3 (wfc3ir)` cosmic-ray identification algorithm (`CRCORR`, [Section 3.3.10](#)), which assumes that a given pixel sees a constant count rate from a combination of sources and diffuse background (i.e., the "ramps" are linear). Strong time variation in the background can trip the CR thresholds, with most or all of the image identified as a CR at a given read. Furthermore, since the background variation is fairly smooth from read to read any algorithm that tries to iteratively identify clean reads before and after a CR hit will likely fail.

The primary impact of the strong background variations is to increase noise as it reduces the available exposure time in the final fit products (e.g. only one or two reads out of 15 are used to form the fit). Furthermore, the distribution of background pixel values frequently shows multi-modal non-Gaussian shapes as different parts of the image trip ([Fig 7.12](#)) and confuse the CR algorithm in different ways. **As a consequence, default fit products generated from ramps with time-variable background are not recommended for scientific analysis and should be manually reprocessed using one of the strategies described below.**

- ✔ Four Jupyter notebooks were created to aid users in identifying images affected by variable background. These notebooks are briefly described in [Table 7.2](#), and in more detail below and in [Section 9.5.3](#). These visualization tools may be used to identify issues with the data, for example, a [guidestar \(GS\) failure](#), a [satellite trail](#) in a specific read, or time variable background.

Table 7.2: Jupyter notebooks that can be used to help identify images affected by variable background (see also Section 3 of [STAN 42](#))

Notebook Title	Method(s)	Usage
IR IMA Visualization Tools with an Example of Time Variable Background	Inspect the IMA <code>MULTIACCUM</code> images and produce signal ramp plots to determine whether a dataset needs reprocessing.	Any IR imaging filter. Compares cumulative vs instantaneous signal in ramp.
Manual Recalibration with <code>calwf3</code>: Turning off the IR Linear Ramp Fit	Use the final read of the IMA <code>MULTIACCUM</code> image, skip the ramp fit, and create an FLT image.	Any IR imaging filter. Uses Astrodrizzle to combine FLTs and reject cosmic rays.
Correcting for Helium Line Emission Background in WFC3/IR Exposures using the "Flatten-Ramp" Technique	Subtract the median background per IMA read and re-run the ramp fit with <code>calwf3</code> .	All filters affected by the helium line emission background. Corrects uncharacteristically 'noisy' FLT images with poor ramp fits.
Correcting for Scattered Light in WFC3/IR Exposures: Using <code>calwf3</code> to Mask Bad Reads	Mask specific reads in the RAW image and reprocess with <code>calwf3</code> , including running the ramp fit.	Removes scattered light present in reads at the beginning or end of an exposure. Uses AstroDrizzle to combine FLTs.
Correcting for Scattered Light in WFC3/IR Exposures: Manually Subtracting Bad Reads	Subtract specific reads from the IMA <code>MULTIACCUM</code> image, skip the ramp fit, and create an FLT image.	Removes scattered light or anomalies present in any read. Uses AstroDrizzle to combine FLTs and reject cosmic rays.

7.10.1 Scattered Earthlight

Observations made when HST is pointing near the bright Earth limb can result in the leftmost ~200-400 columns of the detector being subjected to background levels up to twice as bright as that on the rest of the chip. The overall shape of this high background region is similar from one affected image to another, but the overall brightness and the number of affected pixels varies as the HST pointing approaches or recedes from the bright Earth limb. Details on the nature of this effect in IR darks can be found in [WFC3 ISR 2009-21](#).

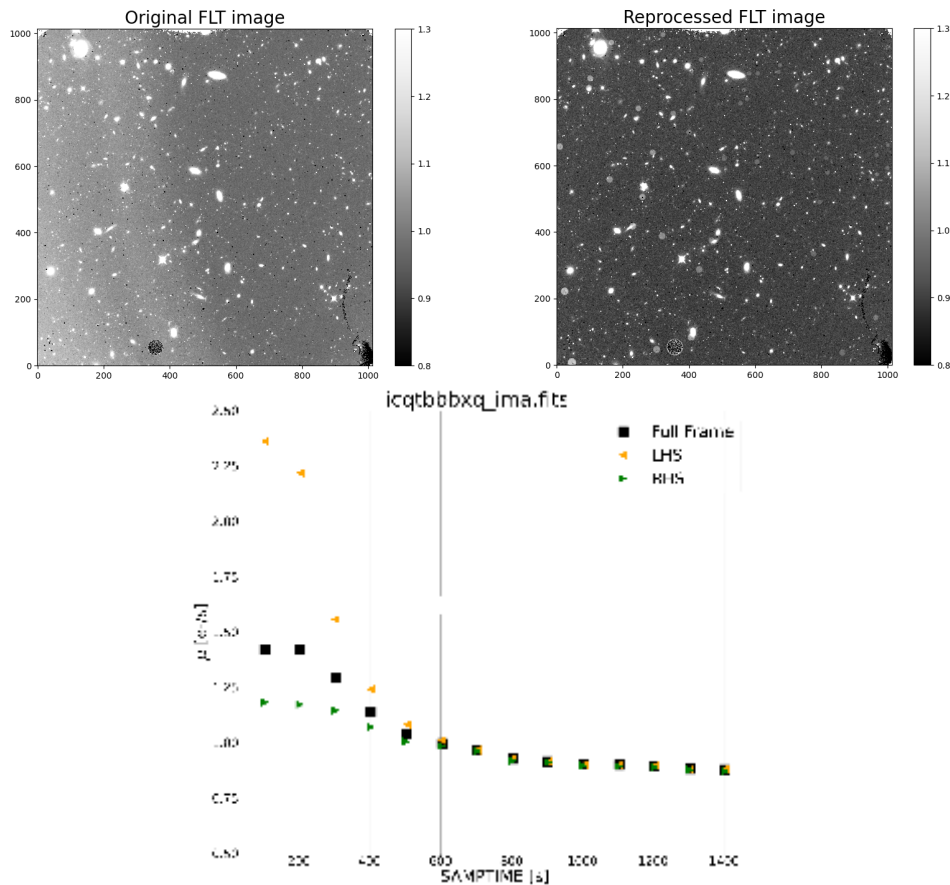
Two new Jupyter notebook tutorials show how to reprocess datasets affected by scattered light using different methods. The first notebook, [Correcting for Scattered Light in WFC3/IR Exposures: Using calwf3 to Mask Bad Reads](#), is based on the F140W exposure (icqtbbsbxqflt.fits) shown in the top-left panel of [Figure 7.11](#). This image was acquired during the first half of the orbit and is contaminated by scattered light in multiple reads. To identify the impacted reads, the average flux between consecutive IMA reads is computed for two different regions of the detector and plotted in bottom panel of [Figure 7.11](#).

Because scattered light always occurs on the left side of the detector (due to the geometry of the optical bench, see [Section 1.1](#)), it is easy to determine which reads are impacted. When the ratio of the differential count rate for two regions of the detector exceeds a user-defined threshold, those reads may be flagged as 'bad' in the RAW frame and then excluded during **calwf3** reprocessing. For this sample dataset, the first five reads show a strong excess at the left hand side (LHS) of the detector.

The recalibrated FLT image is displayed in the top-right panel of [Figure 7.11](#) with the same color stretch as the original image. While the total exposure is reduced from 1403 seconds to 1000 seconds, the background in the reprocessed image is now uniform over the entire field of view. Bright circular residuals show a higher background for regions affected by blobs (see [Section 7.5](#)). Because **calwf3** assumes that any pixel flagged with a DQ value of 512 (bad in flat) is bad in every read, the software fills in these regions with the pixel value from the total exposure. For associated images with a large dither to step over the blobs, the calibrated FLT data may be re-drizzled while rejecting any pixels with 512 flags. When combining observations, **AstroDrizzle** will then replace those pixels with unflagged regions from the second (dithered) exposure in the association.

Figure 7.11: Top left: Calibrated FLT image (*icqtbbsxq_flt.fits*) impacted by scattered light from the Earth limb, seen as a bright excess at the left side of the detector. Top right: The reprocessed FLT image, displayed with the same scaling as the original. This new version was produced by *calwf3* after masking the first 5 reads (not including the zero read) in the RAW file. The reprocessed image is free of Earth limb scattered light and the background is uniform over the entire field of view though with an unavoidable reduction in effective exposure time, from 1403 to 1000 seconds (see notebook [Correcting for Scattered Light in WFC3/IR Exposures: Using calwf3 to Mask Bad Reads](#)). Blob are filled by *calwf3* with their original, higher signal value as discussed in the text however these regions are flagged in the DQ array and may be rejected when combining dithered exposures.

Bottom: The difference in count rate between individual reads in a WFC3/IR *_ima.fits* file (see notebook [WFC3/IR IMA Visualization Tools with An Example of Time Variable Background](#) particular scattered light exposure, excess signal affects the first several reads, manifested as a higher median count rate on the left side of the image (orange triangles) than on the right side of the image (green triangles). The left and right side count rates level out in later reads to a Zodiacal light level of +0.9 e-/s.



A second Jupyter notebook shows an alternative strategy for removing scattered light: [Correcting for Scattered Light in WFC3/IR Exposures: Manually Subtracting Bad Reads](#). This method does not leave residuals in blob regions in the recalibrated image though it is more computationally intensive.

Alternatively, exposures impacted by scattered light may be reprocessed using the 'Last-minus-first' technique described in [WFC3 ISR 2016-16](#) and in the Jupyter notebook described in Section 7.10.2. While this method is the most straightforward, it will result in a spatially varying sky background in the recalibrated image. Photometry on such a calibrated FLT image will therefore require a local background correction to account for the sky however the exposure will have the expected noise distribution. While this strategy is insensitive to any background variations during the exposure, the lack of an up-the-ramp fit requires the user to identify CRs by other means, for example, comparing multiple exposures with **AstroDrizzle**.

7.10.2 Metastable Helium 1.083 μm Emission Line

Line emission at 1.083 microns due to metastable helium in the Earth's upper atmosphere affects the F105W and F110W filters and the G102 grism. When present, this emission can increase the IR background by up to a factor of 6 above the nominal Zodiacal background. This spatially diffuse source affects portions of HST orbits where both the telescope and the atmosphere are illuminated by sunlight. [WFC3 ISR 2014-03](#) describes this effect in more detail while [WFC3 ISR 2016-16](#) demonstrates how to identify exposures affected by time-variable background and suggestions for reprocessing the affected exposures (see also the Jupyter notebook [WFC3/IR IMA Visualization Tools with An Example of Time Variable Background](#)). The correction techniques can be summarized as follows.

The first method is the 'Last-minus-first' technique described in the previous Section and in the Jupyter notebook: [Manual Recalibration of Images using calwf3: Turning off the WFC3/IR Linear Ramp Fit](#) which is discussed in [Section 9.5.3](#). In this workflow, the **calwf3** ramp fitting step is turned off (`CRCORR=OMIT`) so that the calibrated (`_flt`) image is the result of subtracting the first read from the last read, then dividing by the time elapsed between the two reads. The top panel of [Figure 7.12](#) compares two back-to-back F105W exposures acquired in a single orbit, where the first had a stable zodiacal background and the second was impacted by variable helium line emission. The central panel shows both images reprocessed with the ramp fitting (`CRCORR`) turned off. While the noise properties of the images are now improved, omitting the up-the-ramp fit leaves CRs in each image which must be identified by other means, such as **AstroDrizzle**.

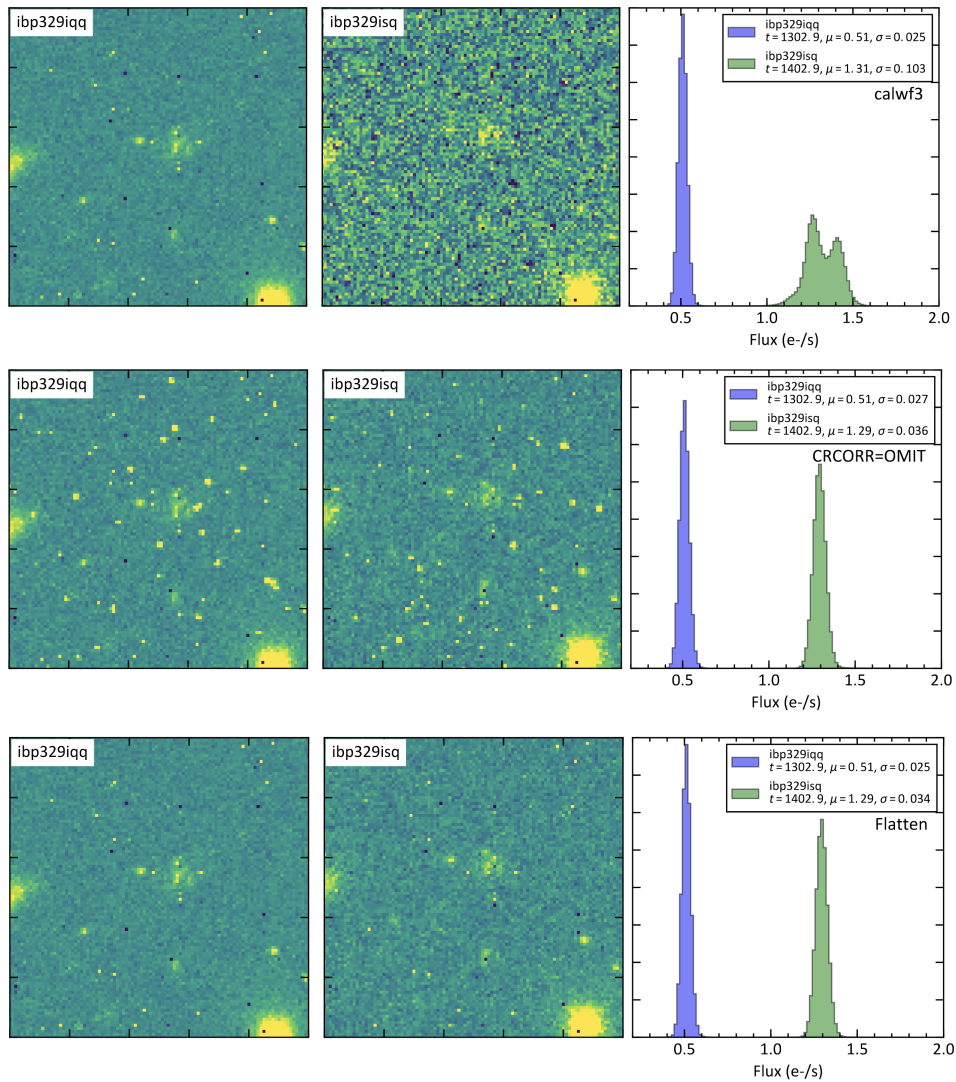
The bottom panel of [Figure 7.12](#) shows the second image as reprocessed using the 'Flatten-ramp' technique which equalizes the background signal rate for each read prior to performing the ramp fit (see [Section 9.5.3](#) and the Jupyter notebook [Correcting for Helium Line Emission Background in WFC3 /IR Exposures using the "Flatten-Ramp" Technique](#)). First, **calwf3** is run up to the point of generating the calibrated IMA files but omitting the cosmic-ray identification step (`CRCORR`). The median background is then subtracted from each read, assuming a constant excess signal for every pixel. Next, a constant value, representing the average count rate of the full exposure, is added back to preserve pixel statistics. Processing with **calwf3** is then resumed to perform the ramp fit (`CRCORR`) step. This approach works well for relatively sparse fields where sky background is easily determined and has been used successfully e.g. for reprocessing the large sample of archival images used to generate IR sky flats (see [WFC3 ISR 2021-01](#)). The statistics of the reprocessed image are indistinguishable from the "Last-minus-first" technique with the added benefit of having the cosmic rays identified and removed by **calwf3**.

Figure 7.12: Top: Comparison of back-to-back F105W exposures in an orbit, where the first exposure (left column) shows a stable zodiacal light background and the second exposure (center column) is impacted by time-variable helium line emission. The pixel distribution of the second exposure is poorly behaved (right column): the noise is not only higher due to the elevated background level but also exhibits a bimodal non-Gaussian shape. Exposures such as these require manual reprocessing (see text).

Center: The same two images and pixel distributions are shown after reprocessing with the 'Last-minus-first' technique ($CRCORR=OMIT$). The noise properties of the products have improved although the cosmic rays in each exposure must now be identified and removed.

Bottom: Results from reprocessing with the "Flatten-ramp" technique. The statistics of the re processed image are indistinguishable from the "Last-minus-first" image with the added benefit of cosmic ray removal by *calwf3*.


Figure from [WFC3 ISR 2016-16](#).



7.11 IR Photometry Errors

The estimated errors for the inverse sensitivity (zeropoints) of the WFC3 IR channel computed from observations of the white dwarf standards plus a G-type star are $\sim 1\%$ (statistical) and 1-2% absolute (systematic, uncertainties in the models used as absolute standards, [WFC3 ISR 2020-10](#)). The signal to noise ratio (SNR), per observation, for the standard stars is such that the formal Poisson error is $\sim 0.2\%$ for the wide band filters. However, Poisson noise is not the limiting factor at $\text{SNR} > 100$ due to poorer than expected repeatability ([WFC3 ISR 2020-10](#)). Data taken using different observation configurations (apertures, sample sequences and numbers, staring versus scan mode) can show differences in photometry. For example, 1 sigma relative photometric repeatability values have been measured at $\pm 1.5\%$ when using the staring mode ([WFC3 ISR 2019-07](#)) and 0.65% when using the scanning mode ([WFC3 ISR 2021-05](#)).


Figure 7.13: Measured normalized standard deviation of measurements of spectrophotometric standard star GD-153 in F160W (points) as a function of estimated signal to noise ratio. The blue line is the theoretical normalized standard deviation ($1/\text{SNR}$) for a given SNR.



A new calibration for the WFC3/IR inverse sensitivities (zeropoints) was released in 2020 ([WFC3 ISR 2020-10](#)). The zeropoint changes (0.5-2%, depending on filter) were due to updated HST CALSPEC models and an increase in the Vega reference flux ([Bohlin et al. 2020](#)) as well as new IR flatfields ([Section 7.8](#)). For optimum results as well as consistency with WFC3/UVIS and ACS, the new 2020 WFC3/IR zeropoints should be used; data retrieved from MAST after Oct 2020 will have the most current photometric calibration incorporated into the science image header keywords.

The IR photometric repeatability is occasionally worse than the expected nominal calibration/Poisson error. Some scatter in measurements within a visit can be due to self-persistence of bright sources releasing charge into pixels used in subsequent exposures (WFC3 ISR 2019-07, WFC3 ISR 2016-11). Self-persistence causes sources to appear brighter by about 2% on average when many saturating stars have been exposed onto the array (individual stars may be even brighter). Sufficient dithering (~10-pixel steps) can mitigate persistence effects, bringing average repeatability for many stars down to ~0.5% (WFC3 ISR 2019-07). On the other hand, high-precision photometry from spatial scans suggests that dithering over a larger area on the detector may cause observations of the same source to have a higher level of scatter (WFC3 ISR 2021-05). Some of the dependence of repeatability on the extent of dithers may be due to limitations in the flat field accuracy although visit-to-visit comparisons of scan photometry of the same sources at the same positions also show offsets (Figure 7.14), implying other causes as yet not understood. As evident from Figure 7.14, the overall photometric repeatability, a combined effect of the intra-visit scatter and the inter-visit offsets, typically surpasses the Poisson error in the high Signal-to-Noise regime.

Figure 7.14: (left) Photometric repeatability in the IR measured from spatially-scanned F140W observations over 10 epochs between 2015 and 2024. Measurements from each star, plotted relative to their mean over all ten epochs, are plotted against time. Data for the different stars from a given epoch are plotted with small artificial time offsets between them for display purposes only. The mean photometry is represented by the black dotted line at 100%. The solid purple line represents a linear fit to the data. The corresponding slope, quantifying the rate of sensitivity evolution is also shown. Figure reproduced from WFC3 ISR 2024-01; (right) Same as left side, but with F098M data distributed over six epochs



Monitoring of staring mode images of individual standard stars (2009-2024) does not show a discernible time-dependence although data are noisy with short-term 1-sigma repeatability of +/- 1.5% (WFC3 ISR 2019-07). However, analyses of grism observations of spectrophotometric standard stars and scan and staring mode observations of cluster stars show low levels of sensitivity declines ranging from ~ 0.06%/yr to ~ 0.12%/yr (WFC3 ISR 2022-07, WFC3 ISR 2024-01, WFC3 ISR 2024-06). For more discussion on the IR detector sensitivity changes please see Section 9.1.6.

For users who need precise photometry, a complete error analysis should include Poisson noise of the science target as well as the additional sources of photometric error of flat fields, spatial repeatability (~1% peak to peak), temporal repeatability (2% peak to peak), readout noise, uncertainty in the encircled energy (approximately 0.5% at the infinite aperture), uncertainty in the gain ([Section 7.2](#)), and processing errors (e.g. sky subtraction, dark subtraction, correction for persistence).

7.12 References

- J. Anderson and N. Grogin, *Dithering for ACS and WFC3 Primes and Parallels*, [WFC3 ISR 2023-05](#).
- V. Bajaj, A. Calamida, J. Mack, 2020, *Updated WFC3/IR Photometric Calibration*, [WFC3 ISR 2020-10](#).
- V. Bajaj, 2019, *WFC3/IR Photometric Repeatability*, [WFC3 ISR 2019-07](#).
- R. C. Bohlin, *New Grids of Pure-hydrogen White Dwarf NLTE Model Atmospheres and the HST/STIS Flux Calibration*, *AJ* 160, 1, 2020.
- G. Brammer, 2016, *Reprocessing WFC3/IR Exposures Affected by Time-Variable Backgrounds*, [WFC3 ISR 2016-16](#).
- G. Brammer, N. Pirzkal, P. McCullough, J. MacKenty, 2014, *Time-varying Excess Earth-glow Backgrounds in the WFC3/IR Channel*, [WFC3 ISR 2014-03](#).
- H. Bushouse, 2008, *WFC3 IR Ground P-Flats*, [WFC3 ISR 2008-28](#).
- T. Dahlen, L. Dressel, and J. Kalirai, 2010, *Dithering strategies for WFC3*, [WFC3 ISR 2010-09](#).
- F. Dauphin, J. V. Medina, P. R. McCullough, *WFC3 IR Blob Classification with Machine Learning*, [WFC3 ISR 2021-08](#).
- M. J. Dulude, S. Baggett, H. Bushouse, & B. Hilbert, 2011, *WFC3/IR Banding*, [WFC3 ISR 2011-04](#).
- M. J. Durbin, M. Bourque, S. Baggett, 2015, *IR “Snowballs”: Long-Term Characterization*, [WFC3 ISR 2015-01](#).
- M. J. Durbin, and P. R. McCullough, 2015, *The Impact of Blobs on WFC3/IR Stellar Photometry*, [WFC3 ISR 2015-06](#).
- M. Giavalisco, K.Sahu, R. Bohlin, 2002, *New Estimates of the Sky Background for the HST Exposure Time Calculator*, [WFC3 ISR 2002-12](#).
- C. M. Gosmeyer & S. Baggett, 2015, *WFC3 IR Gain from 2010 to 2015*, [WFC3 ISR 2015-14](#).
- J. D. Green & H. Olszewski 2019, *IR ‘Snowball’ Occurrences in WFC3/IR: 2009-2019*, [WFC3 ISR 2020-03](#).
- M. Gennaro et al., 2018, *A characterization of persistence at short times in the WFC3/IR detector*, [WFC3 ISR 2018-05](#).
- B. Hilbert, *IR Channel Subarray Dark Current Behavior*, [WFC3 ISR 2010-16](#).
- B. Hilbert, S. Baggett, M. Robberto, 2003, *Masking Technique on WFC3-IR Images*, [WFC3 ISR 2003-06](#).
- B. Hilbert, 2008, *WFC3 TV3 Testing: IR Channel Nonlinearity Correction*, [WFC3 ISR 2008-39](#).
- B. Hilbert, & P. McCullough, 2009, *WFC3 SMOV Results: IR Channel Dark Current, Readnoise, and Background Signal*, [WFC3 ISR 2009-21](#).
- B. Hilbert, V. Kozhurina-Platais, and E. Sabbi, 2009, *WFC3 SMOV Program 11453: IR Flat Field Uniformity*, [WFC3 ISR 2009-39](#).

- B. Hilbert, 2009, *Snowballs in the WFC3-IR Channel: Characterization*, [WFC3 ISR 2009-43](#).
- B. Hilbert & H. Bushouse, 2010, *WFC3/IR Bad Pixel Table: Update Using Cycle 17 Data*, [WFC3 ISR 2010-13](#).
- B. Hilbert, 2012, *WFC3/IR Reference Pixel Characterization #1: Comparison of Bias Subtraction Methods*, [WFC3 ISR 2012-05](#).
- B. Hilbert, 2014, *Updated non-linearity calibration method for WFC3/IR*, [WFC3 ISR 2014-17](#).
- B. Hilbert, and L. Petro, 2012, *WFC3/IR Dark Current Stability*, [WFC3 ISR 2012-11](#).
- K. Huynh and H. Khandrika, *Improvements and Updates to the WFC3/IR Bad Pixel Tables: Cycle 28-30*, [WFC3 ISR 2024-02](#).
- V. Kozhurina-Platais & S. Baggett, 2020, *WFC3/IR Sensitivity over Time*, [WFC3 ISR 2020-05](#).
- K. Long, et al., 2016, *The Effect of Repeated Exposures on Measured Fluxes in the WFC3/IR Detector*, [WFC3 ISR 2016-11](#).
- J. Mack, H. Olszewski, and N. Pirzkal, 2021, *WFC3/IR Filter-Dependent Sky Flats*, [WFC3 ISR 2021-01](#).
- J. Mack & V. Bajaj, 2019, *Improved Drizzled Data Products for the WFC3/IR Detector*, [WFC3 ISR 2019-05](#).
- R. van der Marel, 2003, *Determination of Low-Frequency Flat-Field Structure from Photometry of Stellar Fields*, [ACS ISR 2003-10](#).
- P. R. McCullough, 2009, *Radioactivity in HgCdTe devices: potential source of “snowballs”*, [WFC3 ISR 2009-44](#).
- P. R. McCullough, J. Mack, M. Dulude, and B. Hilbert, 2014, *Infrared Blobs: Time-dependent Flags*, [WFC3 ISR 2014-21](#).
- H. Olszewski & J. Mack, 2021, *WFC3/IR Blob Flats*, [WFC3 ISR 2021-10](#).
- N. Pirzkal, A. Viana, A. Rajan, 2010, *The WFC3 IR “Blobs”*, [WFC3 ISR 2010-06](#).
- N. Pirzkal, J. Mack, T. Dahlen, and E. Sabbi, 2011, *Sky Flats: Generating Improved WFC3 IR Flat-fields*, [WFC3 ISR 2011-11](#).
- N. Pirzkal, B. Hilbert, 2012, *The WFC3 IR “Blobs” Monitoring*, [WFC3 ISR 2012-15](#).
- N. Pirzkal, R. Ryan, 2020, *The dispersed infrared background in WFC3 G102 and G141 observations*, [WFC3 ISR 2020-04](#).
- A. G. Riess, 2010, *First On-orbit Measurements of the WFC3-IR Count-rate Non-Linearity*, [WFC3 ISR 2010-07](#).
- A. G. Riess, 2011, *An Independent Determination of WFC3-IR Zeropoints and Count Rate Non-Linearity from 2MASS Asterisms*, [WFC3 ISR 2011-15](#).
- A. G. Riess, G. Narayan, A. Calamida, 2019, *Calibration of the WFC3-IR Count-rate Non-linearity, Sub-percent Accuracy for a Factor of a Million in Flux*, [WFC3 ISR 2019-01](#).

- M. Robberto, C. Hanley, I. Dashevsky, 2002, *The reference pixels on the WFC3 IR detectors*, [WFC3 ISR 2002-06](#).
- R. Ryan and P.R. McCullough, 2017, *Possible Overlaps Between Blobs, Grism Apertures, and Dithers*, [WFC3 ISR 2017-16](#).
- D. Som, R. Bohlin, J. Mack, V. Bajaj, A. Calamida, *Sensitivity Evolution of WFC3/IR Using Spatial Scanning Photometry and Grism Spectrophotometry*, [WFC3 ISR 2024-01](#).
- D. Som, V. Bajaj, J. Mack, and A. Calamida, 2021, *Photometric Repeatability and Sensitivity Evolution of WFC3/IR*, [WFC3 ISR 2021-05](#).
- B. Sunnquist, S. Baggett, and K. Long, 2017, *An Exploration of WFC3/IR Dark Current Variation*, [WFC3 ISR 2017-04](#).
- B. Sunnquist, 2018, *WFC3/IR Blob Monitoring*, [WFC3 ISR 2018-06](#).
- B. Sunnquist, G. Brammer & S. Baggett, 2019, *Time-dependent WFC3/IR Bad Pixel Tables*, [WFC3 ISR 2019-03](#).
- B. Sunnquist, M. McKay & S. Baggett, 2019, *Time-dependent WFC3/IR Superdarks*, [WFC3 ISR 2019-04](#).

Chapter 8: Persistence in WFC3 IR

Chapter Contents

- [8.1 Persistence in WFC3 IR](#)
- [8.2 Evaluating the Amount of Persistence in Science Images](#)
- [8.3 Mitigating the Effects of Persistence](#)
- [8.4 References](#)


8.1 Persistence in WFC3 IR

Image persistence, a phenomenon commonly observed in HgCdTe IR detectors, is an afterglow of earlier images, that in the case of the WFC3 IR detector, is present when pixels are exposed to fluence¹ levels greater than about 40,000 electrons. In cases where portions of the detector are heavily saturated in the initial image, the afterglow can be detectable at levels comparable to the background for several hours. We note that on rare occasions an afterglow can appear from bright sources taken more than several hours earlier. The effect, dubbed 'burping', is somewhat different than standard persistence in that it occurs after a sequence of persistence-free images followed by a well-exposed internal flat field (intflat). That is, exposures prior to the intflat contain no discernible persistence yet an exposure after the intflat can show persistence which can be traced back to a bright source observed hours up to several days earlier ([WFC3 ISR 2015-11](#)). The effect is not well-understood but speculation is that a portion of persistence from bright sources can sometimes remain trapped and hidden longterm in the IR array, to be released later in response to high signal from the intflat.

An obvious example of persistence is shown in [Figure 8.1](#). An image of a high galactic latitude field, the observation was taken to search for the optical counterpart to a γ -ray burst. However, two visits from separate IR programs had preceded the observation of this field. The afterglow of the bright sources in these dithered observations (taken in the preceding 12 hours) is clearly visible as 5-point line patterns in [Figure 8.1](#).

Figure 8.1: Persistence in an IR image.

Example of an IR image (ia21h2e9q), in which persistence due to earlier visits is very obvious. The amplitude of the persistence for the brightest pixels in the diagonal trails of stars is about 0.1 electrons/sec. The image is shown using a linear grey scale from 0.7 to 0.9 electrons/sec.





Pixels in the WFC3 IR array and other HgCdTe IR detectors are operated as reversed biased diodes. IR detector pixels are effectively p-n junctions; when in a steady state, the diffusion current is equal to the drift current. A very small current can be induced by applying a reverse-bias voltage (the reset switch is briefly closed); charge is removed and the depletion region is widened. Subsequently, when light falls on the detector, the depletion region and diode capacitance shrinks, until the next reset read (Dereniak, 1996). Persistence is understood to arise from imperfections, or traps, within the detector pixels that are exposed to free charge as the bias is reduced. The number of traps exposed is determined by the amount of charge accumulated in the diode. A small percentage, of order 1%, of the free charge (WFC3 ISR 2013-07) is captured by the traps and released later, creating the afterimages known as persistence. Resets prevent more charge from being trapped because they remove free charge from the location of the imperfections, however they do not affect the charge that has already been trapped. A detailed theory of persistence has been presented by Smith et al. (2008a, 2008b) and Long et al. (WFC3 ISR 2013-07).

Different HgCdTe IR detectors have different persistence characteristics. In WFC3, a pixel exposed to an effective fluence level of 105 electrons produces a signal of about 0.3 e⁻/sec 1000 s after the exposure. This signal decays with time as a power law with a slope of about -1. Thus after 10,000 s, the persistence signal will be about 0.03 e⁻/sec, just under the dark current of 0.048 e⁻/sec (median). As shown in the left panel of Figure 8.2, the amount of persistence in the WFC3 IR detector depends strongly on the fluence of the earlier exposure. The shape of the curve reflects 1) the density of traps in different regions of the pixels and 2) the fact that once the detector is saturated the voltage levels within the diodes do not increase sharply with increasing fluence. The right panel of Figure 8.2 shows the power law decay of the persistence at different fluence levels.

Figure 8.2: Persistence as a function of stimulus and time.

Top: Persistence as a function of stimulus in a series of dark frames after an observation of the globular cluster Omega Cen. Each color represents one dark image, thus the color sequence represents a time sequence in the darks series (blue is the first dark, green the second, red the third, and so on). Note that the persistence decreases with each subsequent dark exposure.

Bottom: The persistence decay as a function of time. Each color represents the persistence decay as a function of time for pixels grouped by different fluence values in the stimulus exposure (Long et al. 2012).



A further complexity to modeling the persistence is that persistence actually depends not just on the total fluence, but on the complete exposure history since the traps have finite trapping times ([WFC3 ISR 2013-06](#), [WFC3 ISR 2013-07](#)). A short exposure of a source that results in a fluence of 105 electrons produces less persistence than a longer exposure of a fainter source that reaches the same fluence level because in the second case the traps have more time to capture free electrons before the diode gets reset. While tracking the complete history of each WFC3/IR pixel is a huge task, to a good level of approximation it is possible to model persistence as a function of only total fluence and total exposure time in the stimulus image. This model, which is currently used by the WFC3 team to predict persistence and derive the persistence products (see [Section 8.2](#)) is parameterized as:

$$P(t) = A \left(\frac{t}{1000} \right)^\gamma$$

where t is the time, in seconds, since the end of the stimulus exposure and A and γ are functions of both exposure time and fluence level in the stimulus exposure. We refer to this model as the “A- γ ” model ([WFC3 ISR 2015-15](#)).

Additionally, clear evidence of spatial variation in persistence across the IR detector has been measured ([WFC3 ISR 2015-16](#)). One quadrant (upper left) has a higher persistence amplitude than the other three. The shapes of the power law exponents also appear to differ between quadrants. Using a correction flat provides a factor of two reduction in the peak to peak uncertainties; this flat is incorporated into the persistence prediction software and is available from MAST (Version 3.0.1 of the persistence software).

Persistence of the magnitude (and impact) seen in [Figure 8.1](#) is rare. In part, this is because the STScI contact scientists check all Phase II submissions to identify any programs that are likely to cause large amounts of persistence so that the mission schedulers can insert 2 WFC3/IR-free orbits before executing another WFC3/IR program. However, this process is only intended to identify the worst cases. A large fraction of WFC3/IR exposures have some saturated pixels and all of these pixels have the potential to generate persistence in a following IR observation. Inhibiting IR observations of all exposures that could generate persistence would make it impossible to schedule the large numbers of IR observations that are carried out with HST, and in most cases, small amounts of persistence do not affect the science quality of the data, as long as observers and data analyzers take time to examine their IR images for persistence.

¹Fluence here is expressed in electrons. We use this nomenclature regardless of whether or not the pixel's full-well capacity is reached. Therefore when values larger than the typical $\sim 80,000$ e^- full-well capacity for the WFC3/IR channel are reported throughout this Chapter, their meaning is not that of “detected” electrons, but rather that of “electrons that would have been detected for an infinite full-well capacity”. This number is proportional to the impinging photon flux multiplied by the exposure time.

8.2 Evaluating the Amount of Persistence in Science Images

As described above, persistence in a WFC3 IR image depends on images that have been obtained previously. However, when an observer retrieves data, the observations that came before will most likely be proprietary and inaccessible to the observer. To address this difficulty, the WFC3 team has developed software that estimates the persistence in an observer's image ([WFC3 ISR 2015-15](#) and <https://github.com/kslong/Persistence>). Although not currently part of the standard calibration pipeline, the WFC3 team runs these persistence estimation tools on all WFC3 IR images. The tools produce an HTML file summarizing the amount of persistence in each image as well as several FITS files that can be used by observers to analyze the impact of persistence in their own images. Observers can retrieve the output products via MAST using a special screen shown in [Figure 8.3](#) and located at <http://archive.stsci.edu/prepds/persist/search.php>. As an exercise to illustrate the tool's functionality, the reader can type "ia21h2e9q" into the box under "Dataset" then click "Search": a new browser tab will appear with a list; clicking on the hyperlink "ia21h2e9q" in the left column will open another HTML file with the available persistence information, including a graphic for each instance of persistence predicted in the image.

Figure 8.3:MAST search screen used to retrieve information about persistence.

The screenshot shows the MAST PERSIST Search Form interface. At the top, there is a header for "Barbara A. MIKULSKI ARCHIVE OF SPACE TELESCOPES" with navigation links for MAST, STScI, Tools, Mission Search, Search Website, Follow Us, Register, and Forum. Below this is a secondary navigation bar with links for HST Home, About HST, Getting Started, Registration, Archive Status, and HST Search. The main content area is titled "PERSIST Search Form" and includes a "Help Field Descriptions" link. There are two tabs: "Standard Form" (selected) and "File Upload Form". The form contains several sections: 1. Search controls: "Search", "Reset", and "Clear Form" buttons. 2. Search criteria: "Target Name", "Right Ascension", "Declination", "Equinox" (J2000), "Resolver" (Resolve), and "Radius (arcmin)" (3.0). 3. Dataset and Proposal ID: "Dataset" and "Proposal ID" fields, each with a "User-specified field 1" and "Field Descriptions" dropdown. 4. Search Output: A "Search Output Columns" list with "up", "down", "remove", and "reset" buttons. "Sort By:" options (null, Reverse) and "Display Coords:" options (Sexagesimal, Degrees, Hours). "Search Output Format" (HTML_Table) and checkboxes for "Remove Null Columns", "Make Rows Distinct", and "Skip formatting". 5. Record limits: "Maximum Records" (5001) and "Records per Page" (500) dropdowns. A second "Search", "Reset", and "Clear Form" row is at the bottom.

Normally, these output data products are available within few weeks after an IR observation has taken place.

There are two types of persistence that generally affect data analysis. We define external persistence as persistence that is generated by an earlier visit, and internal persistence as persistence generated within the same visit as the image in question. External persistence is not within the control of the observer, and can appear anywhere in an IR image. Internal persistence is usually less important, since at least in the case of small dithers, it affects regions of the detector which are in the wings of the point spread function of bright sources. The results of a query from the persistence query screen for Visit H of program ID 11189 are shown in [Figure 8.4](#). The columns EXT1, EXT2 and EXT3 indicate that 0.3%, 1.89% and 6.83% of the pixels were predicted to have external persistence exceeding 0.1, 0.03, and 0.01 e⁻/sec. The columns labeled TOT1, TOT2, and TOT3 provide the percentages for the sum of external and internal persistence. If all of the estimates of persistence for a dataset in the results screen, particularly those having to do with external persistence, are zero then one should not need to investigate further.

Observers should check for persistence as part of their evaluation of the science data quality of their observations and do so early enough that they are able to submit a [Hubble Observation Problem Report](#) (HOPR) if their data are so compromised by external persistence that the proposed science cannot be achieved. Filing a HOPR for persistence is rare; more often, investigators must determine whether they need to take persistence into account during their analysis.

If the summary estimates in the results screen indicate there is persistence, then more information about the amount of persistence and its location in the images can be obtained by clicking on the links associated with data set column in the results screen (if the data are still proprietary, you will be asked for your MAST login name and password at this point). This will bring up a page that contains additional text as well as a number of images showing where persistence is predicted. There are also various thumbnail images, showing how well the persistence is removed by the model in selected areas of an image. The utility of these thumbnails varies, depending upon the crowding of the image with real sources. An algorithm is used to try to select regions of the image which have the largest amounts of persistence but which are not very crowded, so that it is easy to establish the sky level and automatically generate thumbnails. But this algorithm is not very sophisticated, and in many cases, all of the regions with significant persistence lie in regions of the image that contain brighter *real* objects. Nevertheless, inspecting the images is recommended, and should help one to determine, qualitatively, whether persistence is likely to affect the science goals.

Figure 8.4: Results of MAST persistence query.

Dataset	Asn ID	Proposal ID	Visit	Aper	Target	RA (J2000)	Dec (J2000)	Exp Time	Filter	PI Last Name	Release Date	Status	EXT1	EXT2	EXT3	TOT1	TOT2	TOT3
ia21h2e9q	IA21H2010	11189	H2	IR	GRB090423	09 55 32.674	+18 09 01.32	1302.900	F160W	Tanvir	2011-01-26 07:18:02	Complete_3.0.2	0.030	1.890	6.830	0.040	1.930	6.780
ia21h2e9a	IA21H2010	11189	H2	IR	GRB090423	09 55 32.632	+18 09 01.53	1302.900	F160W	Tanvir	2011-01-26 07:18:02	Complete_3.0.2	0.010	1.580	6.660	0.020	1.640	6.630
ia21h2e9c	IA21H2010	11189	H2	IR	GRB090423	09 55 32.634	+18 09 01.28	1302.900	F160W	Tanvir	2011-01-26 07:18:02	Complete_3.0.2	0.000	0.750	6.170	0.010	0.820	6.150
ia21h2e9d	IA21H2010	11189	H2	IR	GRB090423	09 55 32.592	+18 09 01.49	1302.900	F160W	Tanvir	2011-01-26 07:18:02	Complete_3.0.2	0.000	0.520	6.040	0.010	0.600	6.030
ia21h2e9f	IA21H2010	11189	H2	IR	GRB090423	09 55 32.639	+18 09 01.63	1302.900	F160W	Tanvir	2011-01-26 07:18:02	Complete_3.0.2	0.000	0.130	5.540	0.000	0.180	5.560
ia21h2e9g	IA21H2010	11189	H2	IR	GRB090423	09 55 32.597	+18 09 01.85	1302.900	F160W	Tanvir	2011-01-26 07:18:02	Complete_3.0.2	0.000	0.090	5.400	0.010	0.130	5.430
ia21h2e9h	IA21H2010	11189	H2	IR	GRB090423	09 55 32.679	+18 09 01.67	1302.900	F160W	Tanvir	2011-01-26 07:18:02	Complete_3.0.2	0.000	0.020	4.910	0.010	0.050	4.940
ia21h2e9j	IA21H2010	11189	H2	IR	GRB090423	09 55 32.637	+18 09 01.89	1302.900	F160W	Tanvir	2011-01-26 07:18:02	Complete_3.0.2	0.000	0.010	4.750	0.010	0.050	4.790
ia21h2e9k	IA21H2020	11189	H2	IR	GRB090423	09 55 32.674	+18 09 01.32	1302.900	F160W	Tanvir	2011-01-26 07:18:02	Complete_3.0.2	0.000	0.000	4.160	0.000	0.030	4.210
ia21h2e9m	IA21H2020	11189	H2	IR	GRB090423	09 55 32.632	+18 09 01.53	1302.900	F160W	Tanvir	2011-01-26 07:18:02	Complete_3.0.2	0.000	0.000	2.820	0.000	0.020	2.890

The persistence products software also produces several FITS files (summarized in [Table 8.1](#)) for each flt file in a dataset. The *_extper.fits and *_persist.fits files contain predictions of the external and total persistence in a dataset. The *_flt_cor.fits file is an flt file with the total persistence subtracted. These files can be retrieved from MAST as a tar.gz file by clicking one of the Visit links (see Visit column in [Figure 8.4](#)); each Visit must be retrieved separately. In situations where there is some persistence predicted in a science image, comparing these persistence product FITS files to the

original data in DS9 or some other image display is ultimately the best way to determine whether persistence needs to be taken into account in the analysis of the data.

Table 8.1: Files Associated with Persistence Correction

Generic Filename	Contents
_extper.fits	Prediction of external persistence
_persist.fits	Prediction of total persistence
_flt_cor.fits	flt file with persistence model subtracted

8.3 Mitigating the Effects of Persistence

Two possible ways to mitigate persistence are to 1) exclude the affected pixels from the analysis, or 2) subtract an estimate of the persistence signal from the data and use the modified data for the analysis.

Simple procedures can be used to mark compromised pixels and modify the data quality extensions of the `flt` files. Typically one would use the `persist.fits` products (described in the previous section) to define an acceptable level of persistence, e.g. $0.01 \text{ e}^-/\text{sec}$, based upon a visual comparison of the data and the corresponding persistence images. The choice of level should be based on the science objective and the fraction of pixels impacted. The selected level can be used to flag all pixels in the `persist.fits` image that exceed predicted persistence values.

This method is typically very effective at flagging persistence from previous observations in staring mode. However, the “A-gamma” model predictions can differ somewhat from the observed amount of persistence for images taken shortly before the image of interest (short-term persistence) causing a mismatch between the estimated internal persistence and what is observed. In a similar fashion, when persistence is from prior spatially scanned observations, the model significantly underestimates the level of persistence observed in the data. The “A-gamma” model relies on fluence estimates from the `flt` images of prior observations, i.e. implicitly assumes that the estimated flux in the `flt` accurately reflects the fluence absorbed by specific pixels, which is not the case for scanned observations. The `CRCORR` step of `calwf3` works by fitting a line to the accumulating charge between reads, but for scan-mode observations, the flux is only accumulating in a very short interval, thus the pipeline is unable to report the true flux ([Section 10.2](#)). This is true not only for the scans but also for other very rapidly moving targets such as like asteroid trails ([WFC3 ISR 2017-18](#)) or bright light sources (e.g. cities) that streak across some of the dark-Earth IR flats used to monitor the IR blobs ([Section 7.5](#)). Persistence in those cases is very poorly estimated and can be inaccurate by orders of magnitude.

Note, however, that for both internal persistence and for persistence from moving targets, the “A-gamma” model does predict the location of persistence on the detector very well. Therefore, the user is advised to tune the tolerance level to account for possible mismatches if moving targets are the suspected causes of persistence.

Once the pixels are flagged, down-stream analysis proceeds as it normally would, assuming the tools that are used take the data quality flags into account. Persistence is a property of the pixel: as long as the observer who planned the original observations created a set of appropriately-dithered images, using this procedure to flag bad-persistence pixels should not significantly impact the science.

For example, using `AstroDrizzle` to exclude those flagged pixels would result in a combined image that is largely cleansed of the effects of persistence. In some cases, the threshold defined from inspection of the `flt.fits` may not be low enough to detect lower-level persistence, so the deeper, combined `drz.fits` images may be used to verify whether persistence was fully removed.


An alternative method is to use the persistence-subtracted files in down-stream analysis. If one adopts this approach, one should carefully inspect the `_flt_cor.fits` corrected data products (produced by simply subtracting the `persist.fits` model directly from the `flt.fits` image) to evaluate how well persistence has been subtracted from the image. Observers should inspect all of the images in a particular science observation as the persistence signal will decline with time. Typically, about 90% of the persistence signature is removed by the model. Because there are variations in persistence which are still being studied, the corrected data products can be either under- or over-subtracted, particularly in the case of internal persistence where the power law model is not optimal for estimating the short term decay ([WFC3 ISR 2019-02](#)).

The persistence-subtracted image of our example is shown in [Figure 8.5](#). The resulting image is much cleaner, especially at the center of the image where the γ -ray burst had occurred. However, a careful inspection shows some residual signatures of the persistence, particularly due to the set of early images that produced the horizontally-stepped faint features in the original image.

✔ Sometimes one can do improve the persistence correction by scaling the persistence model and producing one's own persistence-subtracted image. It is also, of course, possible to use a blended approach: flagging the worst pixels in the persistence-corrected images then using `AstroDrizzle` to produce a combined image.

Figure 8.5: Impact of persistence subtraction

Left: Same as the image shown in [Figure 8.1](#); **Middle:** Same as the image shown in [Figure 8.1](#) but with persistence subtracted. Much, although not all, of the persistence has been removed; **Right:** Difference image showing just the pixels with persistence that were removed.



If persistence remains a significant problem in the image analysis in spite of the approaches suggested above, please contact the help desk (<http://hsthelpp.stsci.edu>). It is possible that by changing some of the parameters in the model from their nominal values, better estimates of the persistence in an `flt` file can be obtained, particularly in situations where several earlier visits have affected a science image.

In all cases when dealing with persistence, as is true of most problems associated with data analysis, first assess the severity of the problem and then choose a method of handling the issue that is consistent with the science being carried out. Users concerned about the effects of persistence on their data should check for updates on the [WFC3 Persistence pages](#) as well as the main [WFC3 web page](#).

8.4 References

- L.E. Bergeron and M.E. Dickinson, 2003, *Removal of Cosmic Ray Persistence from Science Data using the Post-SAA Darks*, [NICMOS ISR 2003-010](#).
- T. Dahlen, 2013, *WFC3/IR Internal Flat Fields*, [WFC3 ISR 2013-04](#).
- Dereniak, E., and Boreman, G.D., 1996, *Infrared Detectors and Systems*, *Wiley Series in Pure and Applied Optics*, 1996.
- M. Gennaro, S. Baggett and V. Bajaj, 2019, *A characterization of persistence at short times in the WFC3/IR detector. II*, [WFC3 ISR 2019-02](#).
- B. Hilbert and L. Petro, 2012, *WFC3/IR Dark Current Stability*, [WFC3 ISR 2012-11](#).
- B. Hilbert, 2012, *WFC3/IR Reference Pixel Characterization #1: Comparison of Bias Subtraction Methods*, [WFC3 ISR 2012-05](#).
- B. Hilbert and P. McCullough, 2010, *WFC3 SMOV Results: IR Channel Dark Current, Readnoise, and Background Signal*, [WFC3 ISR 2009-21](#).
- K. S. Long, S. M. Baggett, J. W. MacKenty, and A. G. Riess, 2012, *Characterizing Persistence in the IR detector within the Wide Field Camera 3 instrument on the Hubble Space Telescope*, [Proceedings of the SPIE, 8442, 84421W-9](#).
- K. S. Long, S. M. Baggett, and J. W. MacKenty, 2013a, *Characterizing Persistence in the WFC3 IR Channel: Finite Trapping Times*, [WFC3 ISR 2013-06](#).
- K. S. Long, S. M. Baggett, and J. W. MacKenty, 2013b, *Characterizing Persistence in the WFC3 IR Channel: Observations of Omega Cen*, [WFC3 ISR 2013-07](#).
- K. S. Long, S. M. Baggett, and J. W. MacKenty, 2015, *Persistence in the WFC3 IR Detector: An Improved Model Incorporating the Effects of Exposure Time*, [WFC3 ISR 2015-15](#).
- K. S. Long, S. M. Baggett, and J. W. MacKenty, 2015, *Persistence in the WFC3 IR Detector: Spatial Variations*, [WFC3 ISR 2015-16](#).
- M. Robberto, C. Hanley, and I. Dashevsky, 2002, *The Reference Pixels on the WFC3 IR Detectors*, [WFC3 ISR 2002-06](#).
- R. E. Ryan Jr. & S. M. Baggett, *The Internal Flat Fields for WFC3/IR*, [WFC3 ISR 2015-11](#).
- R. M. Smith, M. Zavodny, G. Rahmer, and M. Bonati, 2008a, *A Theory for Image Persistence in HgCdTe Photodiodes*, [Proceedings of the SPIE, 7021, 70210J-1](#).
- R. M. Smith, M. Zavodny, G. Rahmer, and M. Bonati, 2008b, *Calibration of Image Persistence in HgCdTe Photodiodes*, [Proceedings of the SPIE, 7021, 70210K-1](#).
- B. Sunnquist, J. Mack, and H. Khandrika, *Asteroids in the WFC3/IR Frontier Fields Images*, [WFC3 ISR 2017-18](#).
- D. Thatte, T. Dahlen, et al., 2009, *NICMOS Data Handbook*, version 8.0, pg 1.

Chapter 9: WFC3 Data Analysis

Chapter Contents

- [9.1 Photometry](#)
- [9.2 Astrometry](#)
- [9.3 Spectroscopy](#)
- [9.4 STSDAS, STSCI_PYTHON and Astropy](#)
- [9.5 Specific Tools for the Analysis of WFC3](#)
- [9.6 References](#)

9.1 Photometry

We present below a checklist of steps for measuring photometry of sources on images collected with the WFC3 UVIS and IR detectors. A detailed description of each step is provided in the following sections.

- Download data products appropriate to your science case from [MAST](#):
 - Single exposures (FLT) or drizzled (DRZ) images
 - [CTE-corrected¹](#) (FLC/DRC) or not-corrected (FLT/DRZ) UVIS images
- Apply the Pixel Area Map ([PAM](#)) correction (only when performing aperture photometry on FLT/FLC images; not needed for surface brightness photometry)
- Perform Aperture or Point-Spread Function ([PSF](#)) photometry according to the science case
- Apply an aperture correction to the measured count rates if performing PSF photometry or using a small (≤ 5 pixels) aperture radius
- Apply the WFC3 tabulated encircled energy ([EE](#)) corrections to infinity to the aperture corrected count rates
- Convert the count rates to the appropriate photometric system (VEGAMAG , STMAG, or ABMAG) by using the WFC3 provided [inverse sensitivities or zero points](#)
- Apply [count-rate non-linearity correction](#) (IR only, if needed)
- Check [blue leak levels](#) for possible impact (rare, IR only; [WFC3 ISR 2008-05](#))
- Apply a [color term correction](#) (UV filters, if needed)
- Correct for [red leak](#) (UV filters, if needed; [WFC3 ISR 2008-49](#))

Summarizing, after downloading the selected image type, applying the PAM correction if needed, and performing small-radii aperture or PSF photometry on the data, to obtain calibrated magnitudes, M , from the instrumental ones, m , the following equation can be used:

$$M = m + AP_{10} + EE_{inf} + 2.5\log(exptime) + ZP_{VEGA,ST,AB} + CRNL + CT + RL$$

where AP_{10} is the correction to a 10-pixel aperture radius (recommended), EE_{inf} is the encircled energy correction to infinity, $exptime$ is the exposure time of the image (for UVIS data only),

$ZP_{VEGA,ST,AB}$ is the magnitude zero point in the VEGAMAG, STMAG or ABMAG photometric system, $CRNL$ is the count-rate non-linearity correction for the IR detector, and CT and RL are the color term and red leak corrections for UV filters.

- 9.1.1 Photometric Systems
 - Photometric Zero Points
 - Photometric Transformations
- 9.1.2 Pixel Area Maps
 - PAM Concept Illustration
 - Example #1 Constant Surface Brightness Object
 - Example #2 Photometry of a point source
- 9.1.3 Aperture and Encircled Energy Corrections
- 9.1.4 UVIS Photometric Calibration
 - Time-Dependent Solutions
 - Prior UVIS Photometric calibrations
 - In-flight Calibration in 2009 and 2012
 - Chip-Dependent Solutions
- Quad Filters
- 9.1.6 UV Filters
 - Photometric Calibration
 - Color terms
 - Contamination
- 9.1.6 IR Photometric Calibration
 - Time-dependent solutions
 - Count Rate Non-Linearity
- 9.1.7 Image Photometry Reference Table
- 9.1.8 Narrow-Band Filter Calibration for Emission-Line Objects

9.1.1 Photometric Systems

The WFC3 filters naturally define their own photometric system and users are encouraged to refer their photometric results to this native system. The natural instrumental magnitude of an object observed in a WFC3 filter is instrumental mag = $-2.5 \log(\text{count rate})$, where the count rate is in units of electrons per second (e^-/sec). WFC3 supports the STMAG, ABMAG and VEGAMAG photometric systems and provides zero points for these, in addition to the filter-dependent instrument sensitivity, which converts the measured count rate (e^-/sec) to a mean flux density F_λ in units of $\text{erg cm}^{-2} \text{s}^{-1} \text{\AA}^{-1}$ and F_ν in units of $\text{erg cm}^{-2} \text{s}^{-1} \text{Hz}^{-1}$.

The STMAG and ABMAG systems define an equivalent flux density for a source, corresponding to the flux density of a source of predefined spectral shape that would produce the observed count rate, and convert this equivalent flux to a magnitude (see e.g. [stsynphot](#)). 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, and the reference spectrum is flat in F_λ . An object with $F_\lambda = 3.63 \times 10^{-9} \text{ erg cm}^{-2} \text{ s}^{-1} \text{ \AA}^{-1}$ will have STMAG=0 in every filter, where the STMAG zero point is -21.10 (PHOTZPT header keyword).

- $\text{STMAG} = -2.5 \log F_\lambda - 21.10$

In the ABMAG system, the flux density is expressed per unit frequency, and the reference spectrum is flat in F_ν . An object with $F_\nu = 3.63 \times 10^{-20} \text{ erg cm}^{-2} \text{ s}^{-1} \text{ Hz}^{-1}$ will have magnitude ABMAG=0 in every filter, where the ABMAG zero point is -48.6.

- $\text{ABMAG} = -2.5 \log F_\nu - 48.6$

The relationship between ABMAG and STMAG is:

- $\text{ABMAG} = \text{STMAG} - 5 \log (\text{PHOTPLAM}) + 18.692$

where F_ν is expressed in $\text{erg cm}^{-2} \text{ s}^{-1} \text{ Hz}^{-1}$, F_λ in $\text{erg cm}^{-2} \text{ s}^{-1} \text{ \AA}^{-1}$, and PHOTPLAM is the bandpass pivot wavelength in angstroms. PHOTPLAM values for all filters are listed in Column 3 of Table 6.2 of the [WFC3 Instrument Handbook](#).

Formally, the HST VEGAMAG system is defined by the absolute spectral energy distribution of Vega, such that Vega has VEGAMAG=0 at all wavelengths. Thus, the VEGAMAG magnitude of an object with flux F is:

- $\text{VEGAMAG} = -2.5 \log (F_{\text{object}} / F_{\text{vega}})$

where F_{vega} is the CALSPEC observed flux density of Vega. For the equations that define the average flux, see [Bohlin 2014](#) and [Bohlin et al. 2020](#). In the Johnson-Cousins magnitude system, the average value of six AOV stars sets the zero point values so that U-B=0 and B-V=0 ([Johnson & Morgan, 1953](#)) and by extension V-R=0 and V-I =0 ([Cousins 1974](#)). In this system Vega has the following magnitudes: U=0.03, B=0.03, V=0.03, R=0.07, I=0.10, J=-0.18, H=-0.03, K=0.13. The VEGAMAG system is convenient for many observers because of its long heritage; however, the ABMAG system is popular with large imaging surveys.

A detailed discussion of these three photometric systems within the context of HST observations is provided in [Sirianni et al. 2005](#) as well as [WFC3 ISR 2009-31](#) and in [stsynphot](#). Further information on the VEGAMAG system is also provided in [Bohlin & Gilliland \(2004\)](#), the ABMAG system in [Oke \(1964\)](#) and the STMAG system in [Koorneef et al. \(1986\)](#). Although convenient, transformation to these (as well as other) photometric systems always has limited precision and is dependent on the color range, surface gravity, and metallicity of the source stars considered ([Sirianni et al. 2005](#)). A Jupyter notebook illustrating filter transformations between the WFC3/UVIS photometric system and others is available [here](#) (see also [Section 9.5.2](#)).

Photometric Zero Points

The photometric zero point of a telescope/instrument/filter combination is a convenient way to characterize the overall sensitivity of the system. By most definitions, the zero point represents the magnitude of a star-like object that produces one count per second within a given aperture (see [Maiz Apellaniz 2007](#)). For WFC3, this throughput measures the performance within a given bandpass taking into account the HST Optical Telescope Assembly (OTA), pick-off mirror and relay mirror reflectivity, filter throughput, transmission of the outer and inner detector package windows, and the quantum efficiency (QE) of the detector. For HST instruments such as WFC3, the zero points depend on the absolute flux calibration of HST white dwarf model atmosphere spectra, and therefore they will change whenever that calibration is improved.

The photometric zero point can be determined using several techniques. In [stynphot](#), a photometric and spectral simulator, a user can renormalize a spectrum to 1 count/sec in the appropriate WFC3 bandpass and output the zero point in the selected magnitude system (assuming that updated throughput tables are included in the local [synphot](#) installation). A [Jupyter notebook](#) illustrating how to calculate WFC3 zeropoints with [synphot](#) is also available.

Similarly, the STMAG and ABMAG zero points for WFC3 data can be computed using photometric keywords in the SCI extension(s) of the image header as follows:

$$ZP_STMAG = -2.5 \log(\text{PHOTFLAM}) + \text{PHOTZPT} = -2.5 \log(\text{PHOTFLAM}) - 21.10$$

$$ZP_ABMAG = -2.5 \log(\text{PHOTFLAM}) - 21.10 - 5 \log(\text{PHOTPLAM}) + 18.692$$

where, fits header keyword PHOTFLAM is the inverse sensitivity and represents the flux density ($\text{erg cm}^{-2} \text{s}^{-1} \text{\AA}^{-1}$) of a star that produces a response of one electron per second in this bandpass and PHOTPLAM is the pivot wavelength in Angstroms, where pivot wavelength is a measure of the effective wavelength of a filter ([Tokunaga & Vacca 2006](#)).

For more information, please see [WFC3 ISR 2021-04](#), [Calamida et al. \(2022\)](#) and the [WFC3 Photometric Calibration webpage](#).

Photometric Transformations

In some cases it may be desirable to compare WFC3 photometric results with datasets in different photometric systems (e.g., WFPC2, ACS, SDSS, 2MASS, Johnson-Cousins). Since the WFC3 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 can be introduced.

Transformation coefficients for different spectral types and astronomical sources have been published ([WFC3 ISR 2014-16](#)). The photometric transformation coefficients between Johnson-Cousins UBV filters and WFC3 UVIS wide-band filters for a given object spectrum can also be computed. A WFC3 Jupyter notebook tutorial has been developed to assist users with these computations and is described in [Section 9.5.2](#).

9.1.2 Pixel Area Maps

If one decided to perform photometry (aperture or PSF) on flt/flc images then the data needs to be corrected for pixel area variations in the data array. The WFC3/UVIS CCDs and WFC3/IR detector contain pixels that vary in their area on the sky as a result of the geometric distortion. As a consequence of this, a larger pixel collects more photons relative to a smaller pixel, leading to an overall gradient in an image of an intrinsically uniform background. However, the flat-fielding process in the HST [calwf3](#) pipeline is designed to correct for that gradient and produce images that have a flat background. As a result, while surface photometry measurements on flat-fielded science data (flt) will be correct, the measured total brightness of sources will vary depending on the position of the object i.e. the areas of the pixels underlying the source.

To achieve uniform aperture photometry of point sources over the detector, observers can either use flt/flc images after applying a pixel area map (PAM) correction or use distortion-free images (drz/drc). The **AstroDrizzle** software corrects for distortion, yielding images (drz/drc) which have a flat sky and contain pixels that are uniform in area (i.e., corrected for distortion and related pixel area variations). Therefore, photometry of any source in an **AstroDrizzle** (drz/drc) image will yield the same count rate irrespective of the position of the source on the image. Photometry measured on a calibration pipeline image (flt/flc) requires a field-dependent correction factor (PAM) to achieve uniformity in the measured count rate of an object across the field.

This PAM correction, in the form of an image, comes from the derivatives of the geometric distortion polynomial. The size of the PAM image is the same as the calibrated (flt/flc) image and each pixel value is set to the normalized area of that pixel. By multiplying the calibrated (flt/flc) images by the PAM image, users will recover the same count rate on both types of images (flt/flc and drz/drc) and the same zero point will apply to both data products, e.g. $drz = flt * PAM$, where the (flt/flc) image has been converted to counts per second.

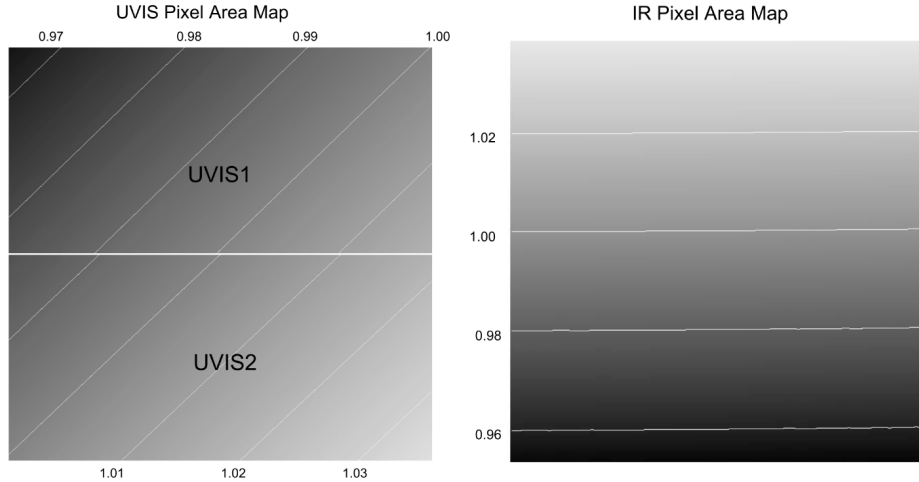
A contour plot of the relative pixel size across the UVIS image, normalized to the central pixel, is shown in the left panel of [Figure 9.1](#); the ratio of maximum to minimum pixel area is 1.074. The variation of pixel area across the IR channel is shown in the right panel of [Figure 9.5](#); the maximum deviation from the central value is 4.1%.

A detailed description of the WFC3 UVIS and IR PAMs is provided in [WFC3 ISR 2010-08](#), which also discusses a unique choice for normalizing the WFC3 PAMs that differs from previous instruments. This approach ensures that the PAM images do not artificially scale the flux in flt/flc images by large amounts. Instead, the PAM images simply serve to provide a relative correction of the counts based on the size of pixels as compared to the size of a reference pixel near the center of the detectors ([WFC3 ISR 2010-08](#)).

i The PAMs are available at <http://www.stsci.edu/hst/instrumentation/wfc3/data-analysis/pixel-area-maps>. See also the Jupyter notebook [WFC3/UVIS Pixel Area Map Corrections for Subarrays](#).

Figure 9.1: Left: Variation of the effective pixel area with position on the UVIS detector. Darker shading indicates pixels with smaller area. Contours are drawn at 1% increments.

Right: Variation of the effective pixel area on the IR detector. Darker shading indicates pixels with smaller area. Contours are drawn at 2% increments.



PAM Concept Illustration

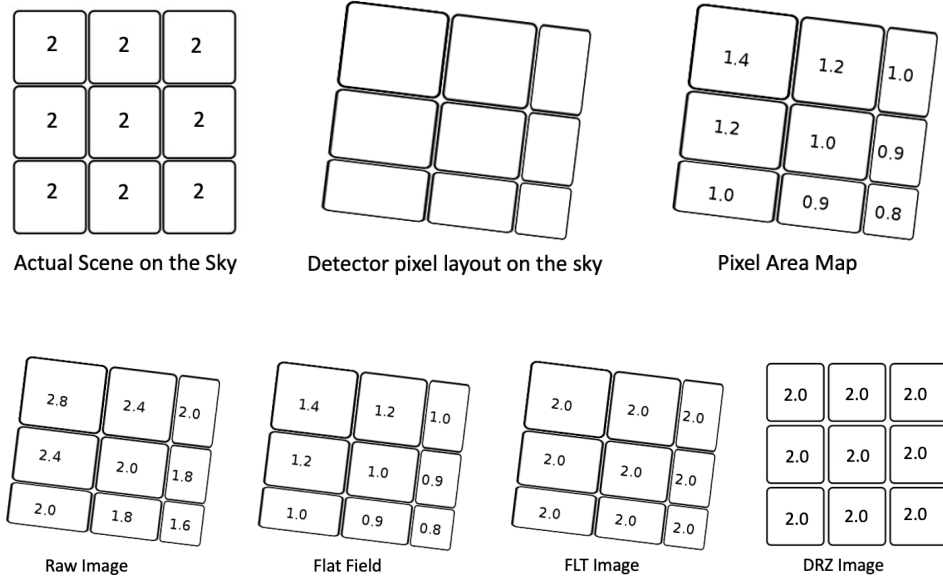
To illustrate the concepts of extended source and point source photometry on images (flt, drz) we consider a simple idealized example of a 3×3 pixel section of the detector. We assume that the bias and dark corrections are zero and that the quantum efficiency is unity everywhere.

Example #1 Constant Surface Brightness Object

Assume an extended object has a surface brightness of $2 \text{ e}^-/\text{pixel}$ in the undistorted case. Without any geometric distortion the image is shown in [Figure 9.2](#) as the 'Actual Scene on the sky'. With geometric distortion, the pixels are not square and the pixel area varies across the detector, as shown in the 'Detector pixel layout on the sky', with a corresponding PAM in the top-right panel.

As a result of the distortion, there will be an apparent variation in surface brightness in the raw image (no flat applied). The geometric 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. WFC3 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 science image (flt/flc) has the correct surface brightness and can be used to perform surface photometry. However, the image morphology is distorted. **AstroDrizzle** can be run on the flat fielded image (flt/flc); the resulting image (drz/drc) will have pixels free of geometric distortion.

Figure 9.2: PAM Illustration for a Constant Surface Brightness, Extended Object



Example #2 Photometry of a point source

Assume we are observing a point source and that all the flux is included in the 3×3 grid with the count distribution as shown in the left panel of Figure 9.3. The total counts are 100. Due to geometric distortion, the PSF as seen in the raw image is distorted, i.e. the total counts are conserved but they are redistributed on the CCD. The flat-field correction, however, does not conserve the total counts since the counts now add up to 99.08, instead of 100. **For accurate integrated photometry on flat fielded WFC3 images, the pixel area variation must be taken into account.** This can be done by multiplying the science image (flt/flc) by the PAM image or by running AstroDrizzle and performing the photometry on the drz/drc image.

Only by running AstroDrizzle can the geometric distortion be removed, but both approaches (using either drz/drc or flt/flc \times PAM) will correctly recover integrated 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 flt \times PAM image, they should calculate a new aperture correction to the total flux of the star as the aperture corrections discussed in Section 9.1.8 are only for drizzled output images (drz/drc).

⚠ In most cases, the aperture correction for distorted flt/flc images will be quite different from the same star measured in a drz/drc image, particularly for small radius apertures.

Figure 9.3: PAM Illustration for a Point Source

Figure 9.3: PAM Illustration for a Point Source



9.1.3 Aperture and Encircled Energy Corrections

In order to reduce errors due to background variations and source crowding, aperture photometry and PSF-fitting photometry are often performed by measuring the flux within a small radius (≤ 5 pixels) around the center of the source. However, small aperture photometric measurements need to be adjusted to a 'total' count rate of the source, i.e. to an 'infinite' aperture radius, by applying an aperture correction (in magnitude units) or an Encircled Energy (EE) fraction (in flux units). For example, the measured UVIS flux in the F606W filter within a 10-pixel (0.4") aperture radius is $\sim 91\%$ of the total flux and within a 50-pixel (2.0") radius aperture is $\sim 98\%$ of the total flux. For the IR detector, the flux in F140W within an aperture radius of 3 pixels is $\sim 84\%$ of the total flux and $\sim 97\%$ within an aperture radius of 15 pixels (2.0").

To compute aperture corrections, a two-step process is recommended. First, users should determine the offset between their own small aperture photometry and photometry within a 'standard' aperture radius, beyond which the EE does not change due to orbital breathing or detector position. This radius was measured to be approximately 10 pixels for the UVIS and 3 pixels for the IR detector, respectively. Therefore, the recommendation is to correct small aperture and PSF photometry to these radii before applying the EE corrections to infinity. For more details on this procedure please see the WFC3 [UVIS](#) and [IR](#) EE webpages.

For the first step, the aperture correction can be calculated by measuring magnitudes or fluxes of a few bright stars in an uncrowded region of the image at the selected small aperture radius and at 10 (UVIS) or 3 (IR) pixel radius and calculating the magnitude difference or flux ratio. These offsets should be applied to the photometric measurements of all sources in the image. If such stars are not available, tabulated EE values for different aperture radii might be applied instead. Alternately, the MAST PSF search tool can be used to download PSFs extracted from WFC3 archival data for a similar detector position and focus as the observation of interest, and if appropriate aperture corrections can be calculated using these. The WFC3 observed PSF database can be accessed by choosing the 'Select a collection' to 'WFC3 PSF' on the [MAST Portal interface](#) (for details, see [WFC3 ISR 2021-12](#)). There are a total of about 32 million UVIS PSFs (29 unsaturated and 3 saturated) and 25.2 million IR PSFs (25 unsaturated and 0.2 saturated) in all the WFC3 images taken prior to January 2024, and all of these observed PSFs are available through the [MAST Portal interface](#). In order to download appropriate PSFs, the user can first use the HST [focus model](#) (the "Annual Summary" provides the model focus values at 5-minute intervals) to estimate the focus value at the time of the observation. Then, the WFC3 PSF database can be searched for PSFs observed under similar focus, detector position, etc. as described in [WFC3 ISR 2021-12](#).


In the second step, the photometry should be corrected to an 'infinite' aperture which encloses all of the light of the source. This filter-dependent correction may be taken from the EE tables, which are derived from high signal-to-noise observations of isolated stars out to large radii. Model EEs have been tabulated for the IR ([WFC3 ISR 2009-37](#)) and the UVIS ([WFC3 ISR 2009-38](#)) detectors. The IR EEs were derived from deep exposures in the F098M and F160W filters and extrapolated to other wavelengths using an optical model. Chip-dependent, filter-based EE fractions were computed for the UVIS detector and spliced to the 2009 in-flight models at $r = 35$ pixels ($\sim 1.4''$). These may be found in [WFC3 ISR 2017-14](#) for radii between 3 and 10 pixels and on the [WFC3 photometry web pages](#) for radii up to 50 pixels ($\sim 2.0''$). The 2017 filter-dependent EE values agree with the 2009 models to $\sim 1\%$ for most filters. However, users should keep in mind that accurate aperture corrections are a function of time and position on the detector.

! Blind application of tabulated EEs should be avoided for small apertures (i.e., $r < 7$ pixels for UVIS, $r < 3$ pixels for IR) where the measured photometry (and the EE fraction) is strongly dependent on the telescope focus and orbital breathing ([WFC3 ISR 2013-11](#)).

New EE corrections for the UVIS detector were calculated in 2020 for two UVIS filters, F275W and F814W, by correcting for time-dependent sensitivity changes of the detector before stacking the images. Values derived for UVIS1 and UVIS2 showed improved agreement with each other and with the 2009 EE values. Following these results, EE values for the other UV filters were also changed by $\sim 1\%$. The EE at wavelengths greater than $7,500 \text{ \AA}$ were changed by $\sim 0.5\%$ to be in close agreement with the F814W values. [Figure 9.4](#) shows the updated EE values at $r = 10$ pixels ($0.3962''$) as a function of wavelength and compares to prior calibration. For more details, see [WFC3 ISR 2022-02](#).

Figure 9.4: UVIS Encircled Energy within a 10-pixel radius.

EE measurements for the UVIS1 and UVIS2 detectors (red, solid, and dashed lines) as a function of wavelength ([WFC3 ISR 2022-02](#)). Prior EE values from 2016 are shown in cyan while the EE model values ([WFC3 ISR 2009-38](#)) are shown in blue and the 2006 pre-launch values ([WFC3 ISR 2005-10](#)) in black.



9.1.4 UVIS Photometric Calibration

The UVIS imaging channel consists of two separate e2v (now Teledyne e2v) CCDs mounted side by side. The two detectors (chips), UVIS1 and UVIS2, have different quantum efficiencies at wavelengths $< 3500 \text{ \AA}$ where UVIS2 is up to $\sim 30\%$ more sensitive than UVIS1 ([WFC3 Instrument Handbook](#) Figure 5.2 and [WFC3 ISR 2016-03](#) Figure 1). At longer wavelengths, the response of the two detectors is similar to within 0.5%. Monitoring observations of multiple CALSPEC standards acquired over time show the sensitivity of the two detectors is slowly changing ([WFC3 ISRs 2016-17](#), [2017-15](#), [2018-16](#), [2021-04](#)). The count rate ratio of the two detectors also changes with time and shows offsets which vary with the color of the source (see Section 9.1.6).

The absolute photometric calibration of the two UVIS detectors has evolved over the years, including: a single detector solution computed in 2009 and 2012, a chip-dependent calibration in 2016 and 2017, and more recently, a time- and chip-dependent calibration in October 2020 (see [WFC3 ISR 2021-04](#)). Any data retrieved from MAST prior to October 2020 can be re-retrieved to obtain the time- and chip-dependent calibration.

When the chip-dependent calibration was implemented, two new header keywords, `PHTFLAM1` and `PHTFLAM2`, were added to the FITS image headers and populated with the inverse sensitivity values for UVIS1 and UVIS2, respectively. For backward compatibility with existing user software, the original header keyword `PHOTFLAM` was populated with the value of `PHTFLAM1`. A new keyword switch, `FLUXCORR`, was also added; this was needed to enable the calibration pipeline, `calwf3`, to scale fluxes measured on UVIS2 to match those measured on UVIS1 by multiplying the UVIS2 science array by the inverse sensitivity ratio, defined as $\text{PHTRATIO} = \text{PHTFLAM2} / \text{PHTFLAM1}$ (see Sections 3.2.13 and 3.4.3). After applying `PHTRATIO`, a source should produce approximately the same number of electrons on UVIS1 and UVIS2 in calibrated (flt/flc) images, corrected for distortion using the pixel area map (see Section 9.1.11). UVIS2 subarray data obtained are also scaled by the `PHTRATIO`, ensuring that sources have the same count rate in calibrated images, regardless of the chip on which they are observed.

As part of the chip-dependent calibration, updated flat fields were computed for all full-frame filters (excluding the QUAD filters and the grism), with low-frequency corrections for the in-flight sensitivity computed separately for each chip ([WFC3 ISR 2016-04](#), [WFC3 ISR 2016-05](#)). The primary change is that the flat fields are normalized separately for each chip and no longer corrected for sensitivity offsets between UVIS1 and UVIS2. Instead, the chip-dependent calibration relies on the `FLUXCORR` step in `calwf3`, which uses the inverse sensitivity values in the `IMPHTTAB` to correct for the chip count rate ratio.

⚠ *While the chip-dependent solutions represent a significant change in the calibration software and reference files, this change was designed to be transparent to the majority of users who will still only need to keep track of a single set of inverse sensitivity values (`PHOTFLAM`) for both chips.*

For photometry in the UV filters, where bandpass differences between the two chips are significant, a flowchart for determining which photometric keywords to use is provided in [Figure 9.2](#).

Time-Dependent Solutions

The chip-dependent calibration assumed a constant detector sensitivity with time. As more monitoring data were acquired, small sensitivity changes up to about 0.2% per year were found, depending on the filter and the detector ([WFC3 ISR 2018-16](#), [WFC3 ISR 2021-04](#)). Over WFC3's on-orbit timespan to date (2009-2024), the change can amount to about 2%. Differences in the time-dependent sensitivity for UVIS1 and UVIS2, as well as small errors in the flat field between different amplifiers, resulted in count rate ratios across the two UVIS detectors differing by as much as 2% over about 10 years ([WFC3 ISR 2018-08](#), [WFC3 ISR 2021-04](#)).

The models for the HST primary spectrophotometric standard white dwarfs (GD153, GD71 and G191B2B) provided by the [CALSPEC calibration database](#) were updated in March 2020 (Bohlin et al. 2020). Also, the Vega reference grey flux at 0.5556 μm , as reconciled with mid-IR absolute flux measures from the SPIRIT III instrument on MSX (Midcourse Space Experiment), increased by $\sim 0.9\%$. The spectral energy distribution for Vega used in the new 2020 calibration is [alpha_lyr_stis_010.fits](#). The standard white dwarf absolute fluxes are determined by the normalization of their modeled spectral energy distributions to their respective relative responses to Vega, using STIS precision spectrophotometry of all three stars along with the flux of Vega at 0.5556 μm ($3.47 \times 10^{-9} \text{ erg cm}^{-2} \text{ s}^{-1} \text{ Hz}^{-1}$). This method provides the basis for HST's entire calibration system. With the adoption of the new models, the HST primary standard white dwarf absolute fluxes increased overall by $\sim 2\%$ for wavelengths in the range 0.15 - 0.4 μm , and $\sim 1.5\%$ in the range 0.4 - 1 μm , necessitating an update to the WFC3 inverse sensitivities.


Major changes in the 2020 time-dependent photometric calibration compared to the previous 2017 calibration can be summarized as follows.

1. The latest photometric values are based on updated SEDs for the standard stars and a new reference flux for Vega (Bohlin et al. 2020);
2. Encircled energy (EE) values for a few filters were updated by correcting for the slow changes in sensitivity over time, then re-drizzling the standard star images to compute new EE curves for each chip;
3. Ten additional years of data acquired from 2009-2019 were used;
4. Data from an additional white dwarf spectrophotometric standard (GRW+70d5824) was included in the analysis;
5. Time-dependent corrections were applied to standard star photometry before deriving the inverse sensitivities. Furthermore, the standard star photometric measurements were weighted according to their photometric errors and the number of collected measurements.

Figure 9.5 shows an example of the time-dependent behavior of UVIS photometry over about 12 years, in this case for UVIS2 and the F814W filter.


Figure 9.5: UVIS Photometry in F814W

Aperture photometry in a 10-pixel radius for five CALSPEC standards in the F814W filter expressed as the count rate percentage versus observation date (in MJD). Monitoring observations of four white dwarfs GRW+70d5824 (filled circle), GD153(diamond), GD71 (square), G191B2B (triangle), and the G-type star P330E (horizontal triangle) are shown for two UVIS2 subarrays, where black points indicate drizzled (drc) photometry for Amplifier C (Amp C) and red points for Amplifier D (Amp D). For comparison, photometry in the corresponding single calibrated (flc) images is shown in grey. Photometry on scanned images is overplotted for GRW+70d5824 (blue) and GD153 (green). The solid line shows the fit to the photometry for all five standards, for both staring and scanned images, and indicates a loss in sensitivity of 0.11% per year +/- 0.07% in this filter.



The new time-dependent inverse sensitivities provide a photometric internal precision of $\leq 0.5\%$ for wide-, medium-, and narrow-band filters, with a significant improvement compared to the prior 2016-era values ($\leq 1\%$ for wide-, $\leq 2\%$ for medium-, and $\leq 5 - 10\%$ for narrow-band filters). The 2020 photometric recalibration is based on an updated correction for the inflight detector response, new encircled energy corrections, and new filter throughput tables for all filters delivered to CRDS. A new image photometry table (`IMPHTTAB`) was also delivered ('51c1638pi_imp.fits', see [Table 9.3](#)) and all WFC3/UVIS archival data were reprocessed through an updated version of the pipeline (`calwf3` v3.5.2). The latest models of the spectrophotometric standard stars can be downloaded from the [CALSPEC webpage](#) and used in `synphot` simulations.

The 2020 calibration now ties the inverse sensitivity values to a common reference epoch, i.e. MJD = 55008 (June 26, 2009) for all filters. Values of the inverse sensitivities for both detectors at each observing epoch can be found in the image header as the `PHTFLAM1` and `PHTFLAM2` keywords. We also provide a tutorial in [Section 9.5.2](#) for running `synphot` with the updated filter curves in order to derive the inverse sensitivities, zeropoint values and throughput curves for any detector, observation epoch, filter or aperture. A second tutorial in [Section 9.5.2](#) shows how to use the new time-dependent solutions to work with UVIS data obtained at different observation dates.

 For data retrieved after October 15, 2020, new time-dependent photometry keyword values (*PHOTFLAM*, *PHTFLAM1*, *PHTFLAM2* and *PHTRATIO*) are populated in the image headers. For observers with data retrieved prior to this date, we recommend re-retrieving to obtain the most up-to-date calibration of their science data.

Before combining **flt.fits* or **flc.fits* products which span multiple epochs (and possibly orientations) with *AstroDrizzle*, the change of the *PHOTFLAM* keyword value in the header must be taken into account. More details on working with time-dependent UVIS zero points and on combining **flt.fits* or **flc.fits* images collected at different epochs is provided in [Section 9.5.2](#).

Prior UVIS Photometric calibrations

In-flight Calibration in 2009 and 2012

The first in-flight photometric calibration (in 2009) was based on the average measurements of two white dwarf spectrophotometric standard stars (GD153 and GRW+70d5824; [WFC3 ISR 2009-31](#)). The total system throughput of both CCDs was found to be significantly better than expected from the Thermal Vacuum 3 (TV3) testing campaign, with efficiency gains of ~10% at the blue and red ends of the UVIS wavelength range and ~20% near the central wavelength 5500 Å ([Figure 6 in WFC3 ISR 2009-31](#)). A smooth polynomial fit across wavelength was used to correct for the observed increase in on-orbit sensitivity.

By 2012, a larger cumulative set of calibration observations made it possible to replace the polynomial fits with more accurate filter-dependent corrections. These revised solutions were based on the average of three white dwarfs (GD153, GD71, G191B2B) plus the G-type standard star P330E. The revised solutions were not documented in a formal ISR but were posted to the WFC3 photometry web page and populated in the image header keywords *PHOTFLAM*, *PHOTPLAM*, and *PHOTBW*. The updated calibration made use of improved flat fields delivered in 2011 which removed a large internal reflection, known as the UVIS 'flare', and corrected for low-frequency spatial variations in the in-flight sensitivity measured from dithered star cluster observations ([WFC3 ISR 2013-10](#)).

Chip-Dependent Solutions

On February 23, 2016, a chip-dependent photometric calibration was implemented for the UVIS detector, and updated inverse sensitivity values were computed ([WFC3 ISR 2016-03](#)) and later updated with improved CALSPEC models ([WFC3 ISR 2017-14](#)). The inverse sensitivity values were recomputed using calibration observations of only the three spectroscopic standard white dwarfs (GD153, GD71, G191B2B), obtained over about six years and measured at multiple positions on the detector. The 2016 inverse sensitivity values were systematically ~3% smaller than the 2012 set of solutions across the full wavelength range of the UVIS detector and were a result of improvements in the photometric reduction procedure ([WFC3 ISR 2016-03](#)). This systematic change, applied retroactively to all UVIS archival data, brought the UVIS photometric system closer to ACS/WFC system ([WFC3 ISR 2018-02](#)).

For prior versions of the inverse sensitivities and EE tables, see [Previous UVIS Calibration.](#), and for a summary of prior photometry reference files and **synphot** tables, see [UVIS Calibration History](#).

Quad Filters

Updated inverse sensitivities for the 20 UVIS quad filters have been computed using the updated 2020 CALSPEC models for the HST flux standard stars as well as the updated reference flux of Vega ([WFC3 ISR 2021-04](#)). The new values do not include any time-dependent correction since insufficient observations were available to reliably measure any changes over time. Aside from the new models, the quad filter photometric calibration is unchanged from 2012 and still makes use of pre-flight (ground) flats that contain the UVIS flare (see [Figure 5.4](#)). As a result, the photometric calibration for QUADs filters on amplifier A will have systematic errors of 1 - 2% due this effect.

The WFC3 UVIS channel contains five quad filters: each is a 2×2 mosaic of filter elements occupying a single filter slot, with each amplifier (quadrant) providing a different bandpass. A subarray observation with a quad filter only covers a single detector quadrant and hence, the science data will have the correct filter keyword reported in their headers. **As there is only a single filter keyword in the science image headers, full-frame quad observations always have the filter keyword (FILTER) populated with the quad element that corresponds to quadrant A, regardless of which quad was requested in the Phase II submission.** [Table 9.1](#) lists the four spectral elements associated with a single quad element, where the value of the filter keyword reported in the image header corresponds to amplifier A. Users can instead query the value of the “ASN_ID” keyword in the image header and search for that association in MAST, where the archive database is populated with the correct ‘FILTER’ keyword. This discrepancy is due to different software systems creating the MAST database and populating the science image header keywords.

To avoid problems with using `PHOTFLAM` values from the wrong quadrant, **all observations using QUAD filters have the `PHOTCORR` and `FLUXCORR` switches set to OMIT.** As a result, the photometric keywords are not populated in the image header during `calwf3` processing. The `PHOTFLAM` values thus have to be retrieved elsewhere, for example from [Table 7](#) in [WFC3 ISR 2021-04](#) or from the [quad filter](#) tables on the photometry website.

Table 9.1: Detector quadrant (amp) imaging locations for the WFC3 quad filters.

Filter wheel	Slot	UVIS1	UVIS1	UVIS2	UVIS2
		Amp A	Amp B	Amp C	Amp D
11	4	FQ508N	FQ674N	FQ575N	FQ672N
12	1	FQ437N	FQ378N	FQ232N	FQ243N
12	2	FQ387N	FQ492N	FQ422M	FQ436N
12	3	FQ889N	FQ937N	FQ906N	FQ924N
12	4	FQ619N	FQ750N	FQ634N	FQ727N

9.1.6 UV Filters

Photometric Calibration

One motivation of the chip-dependent calibration described above was to quantify and correct for bandpass differences, i.e. the response functions, between the two detectors in the UV. Even when two systems (telescope + UVIS1 or UVIS2) use the same filter, due to chip QE differences, the effective bandpass is different in the UV ([WFC3 ISR 2017-07](#), [WFC3 ISR 2018-08](#), [WFC3 ISR 2021-04](#)).

Calibrated WFC3 data products (*flt.fits, *flc.fits) are combined in the pipeline using the **AstroDrizzle** software to create distortion-free data products (*drz.fits, *drc.fits) with the two CCD chips drizzled to the same output frame. AstroDrizzle assumes the images are in units of counts or count rate and not in the flux units used for photometry. This means that the same source must have the same count rate on both chips for AstroDrizzle to properly combine dithered data in which the same sources fall on both chips. Because of the different effective bandpasses in the UV, the measured count rate ratio between chips for the white dwarf standards does not equal the PHTRATIO value provided in the 2016 IMPHTTAB ('z*imp.fits') for the UV filters.

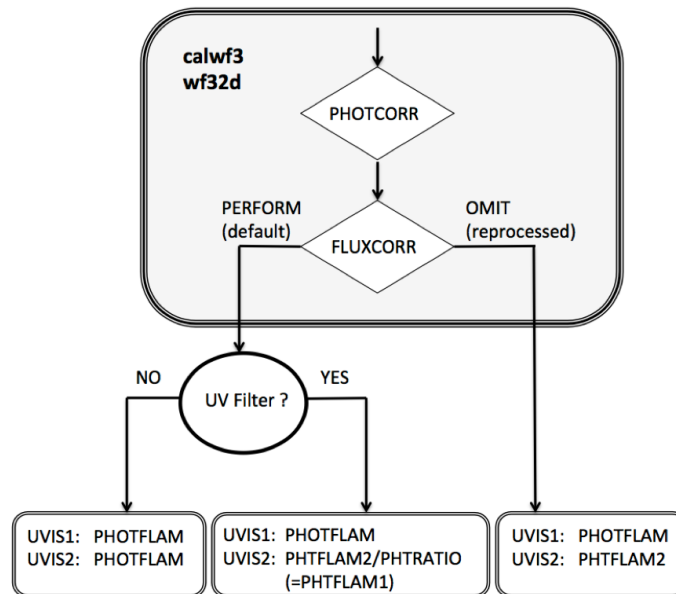
Therefore, on November 21, 2016, the photometric calibration of the UV filters was improved by the delivery of a revised IMPHTTAB ('0*imp.fits'), which scales the PHTFLAM1 values by a factor reflecting the empirical count rate ratio of the white dwarf standards. **These modified values of PHTFLAM1 are provided for the four UV filters (F218W, F225W, F275W and F200LP), such that the PHTRATIO values in the image header match the observed count rate ratios for the white dwarf standards. In this case, the PHOTFLAM value should be used to perform flux calibration, and not PHTFLAM1, which has been modified to normalize the count rate ratio.**

For programs which require very precise (< 1%) photometry in the UV filters, users are advised to treat the two chips as separate detectors and use the chip-dependent keyword values to perform the flux calibration, PHTFLAM1 and PHTFLAM2. For UVIS1, this means simply multiplying the measured count rates by PHOTFLAM (which is equal to PHTFLAM1). For UVIS2, the science array has already been scaled by PHTRATIO, so flux calibration may be achieved by dividing the measured count rates by PHTRATIO and then multiplying by PHTFLAM2, as reported in the image header. In this case, photometry can be computed directly from the calibrated data products (*flt.fits, *flc.fits) multiplied by the pixel area map to account for geometric distortion. Alternately, users can correct for variations in the pixel area by drizzling each chip separately prior to performing photometry. This makes it easier to keep track of which inverse sensitivity value to use with which detector pixels, of particular importance for observations obtained at different orientations or with large dithers.

A flowchart for determining which photometric keywords to use for photometry with UV filters is provided in [Figure 9.6](#). The light gray box shows the final **calwf3** processing steps in the **wf32d** module. The data products (*flt.fits, *flc.fits) obtained from MAST are created with **FLUXCORR** set to **PERFORM** in **calwf3**. Subsequently, the header keyword **PHOTFLAM**, i.e. a unique inverse sensitivity value, can be used to perform the flux calibration of both chips. For the UV filters F218W, F225W, F275W, F200LP, bandpass differences between the two chips results in the count ratio not being equal across them. In order to combine the two chips to create the drizzled (*drz.fits, *drc.fits) data products, **calwf3** multiplies the **UVIS2** science array by **PHTRATIO** (the inverse sensitivity ratio) to bring both chips on the same flux level. Hence, using a single **PHOTFLAM** value for both chips may lead to errors of up to 2% in these filters ([WFC3 ISR 2017-07](#)) since the measured count rate ratio between chips for the white dwarf standards does not equal the **PHTRATIO** value provided in the **IMPHTTAB**. *To achieve high photometric precision (< 1%) in these UV filters, users are advised to treat the two chips as independent detectors.* This is accomplished either by using an alternate set of keywords (bottom center) or by setting **FLUXCORR** to **OMIT**, reprocessing the RAW (*raw.fits) data with **calwf3**, and using a third set of keywords (bottom right).

Figure 9.6: Photometry keywords for UV filters

For observations using the chip-dependent calibration, this flowchart is useful to decide which set of keywords to use to calibrate the photometry across the detector.




Color terms

The **UVIS1** and **UVIS2** detectors have different quantum efficiencies in the ultra-violet (UV) regime ($\lambda < 4,000 \text{ \AA}$), where count rate ratios change as a function of spectral type. When calibrating photometry of stars cooler than $T_{\text{eff}} \sim 30,000 \text{ K}$ in the UV filters (e.g. when observing open and globular clusters, resolved local group galaxies, Galactic stellar populations), color term transformations must be applied to **UVIS2** magnitudes. Sources of any spectral type observed only on a single **UVIS** detector will not require any magnitude offset.

The upper panel of [Figure 9.7](#) shows the synthetic ST magnitude difference, UVIS1 - UVIS2, for a sample of CALSPEC standard stars of varying spectral type (the white dwarfs (WDs) include G191B2B, GD71 and GD153), as computed for the three UV filters F218W, F225W, and F275W. Note that cool red sources such as the G-type star P330E measured on UVIS2 have a magnitude offset relative to UVIS1 up to ~0.08 mag when observed with the F225W filter while F-type stars such as HD160617 have a magnitude offset of ~0.04 mag.

The lower panel of [Figure 9.7](#) shows the ST magnitude difference for the F225W and F275W filters for ω Cen Extreme Horizontal Branch (EHB), Horizontal Branch (HB), Main Sequence (MS), and Red-Giant Branch (RGB) stars measured on different detectors and amplifiers, A (UVIS1) and C (UVIS2). WFC3 observations of ω Cen validate the results of synthetic photometry: red stars (RGBs) observed on UVIS2 have a magnitude offset up to 0.08 mag relative to UVIS1.

Lookup tables with color term transformations are provided in [WFC3 ISR 2018-08](#). The corrections may be applied to UVIS2 magnitudes when observing with the three UV filters F218W, F225W, and F275W.

 The inverse sensitivity ratio, PHTRATIO, is derived using photometry of the CALSPEC spectrophotometric standard white dwarfs and is valid for hot stars, $T_{\text{eff}} > 30,000$ K. For cooler stars, when observing with UV filters, PHTRATIO is not equal to the ratio of the two detectors' count rates but changes with the stellar spectral type. Thus photometry for cooler stars measured on UVIS2 will require a magnitude offset depending on their color, temperature or spectral type.

Before applying the offset to magnitudes measured on UVIS2, photometry must first be calibrated using the UVIS1 inverse sensitivities:


$$\text{ZP_STMAG (UVIS1)} = -2.5 \log (\text{PHOTFLAM}) - 21.1$$

$$\text{ZP_STMAG (UVIS2)} = -2.5 \log (\text{PHOTFLAM}) - 21.1 + \text{Delta (Mag)}$$

where the Delta (Mag) = Mag(UVIS1 - UVIS2) correction is listed in lookup tables in [WFC3 ISR 2018-08](#).

Figure 9.7: Color terms for UV filters

Left: Synthetic ST magnitude difference between the UVIS1 (UV1) and UVIS2 (UV2) detectors for a set of standard stars of varying spectral type (which include the white dwarfs G191B2B, GD71 and GD153, and the solar-analog P330E), as computed for three UV filters F218W, F225W, and F275W. Right: ST magnitude difference for the F225W and F275W filters for EHB, HB, MS, and RGB stars in Omega-Cen measured on different detectors and amplifiers: UVIS1, Amp A and UVIS2, Amp C (see text).



Red Leak

The design and manufacture of the UV filters was based on a careful balance of the achievable in- and out-of-band transmissions. In general, higher in-band transmission results in poorer suppression of out-of-band transmission, and vice versa. The WFC3 filters represent an attempt to achieve an optimum result, maximizing the in-band transmission while keeping the out-of-band transmission as low as possible in order to minimize red leaks.

Table 9.1 below summarizes the red-leak levels for the WFC3 UV filters based on the results in WFC3 ISR 2008-49. The table lists the fraction of the total signal that is due to flux longward of 400 nm, as a function of effective temperature, calculated by convolving a blackbody of the given effective temperature (T_{eff}) with the system throughput in the listed filter. As can be seen from the table, red leaks should not be an issue for observations of any objects taken with F275W or F336W. The other UV filters have some red leaks, whose importance depends on stellar temperature. The red leaks in F218W and F300X, for example, exceed $\sim 1\%$ for objects cooler than ~ 6000 K, while in F225W the red leak reaches $\sim 1\%$ for objects with even cooler temperatures. The most extreme red leaks arise from F218W and F225W observations of objects with effective temperature (T_{eff}) of ~ 4000 K or cooler, necessitating appropriate corrections.

Table 9.1: Fraction of flux longward of 400 nm as a function of effective temperature.

T_{eff} (K)	F218W	F225W	F275W	F300X	F336W
1000	1	1	1	1	1
2000	9.9E-01	9.9E-01	8.4E-01	5.5E-01	3.0E-02
3000	6.0E-01	2.7E-01	3.0E-02	8.9E-02	8.4E-04
4000	1.1E-01	1.8E-02	3.1E-03	3.3E-02	1.4E-04

5000	2.7E-02	3.2E-03	8.6E-04	1.7E-02	4.5E-05
6000	9.9E-03	1.0E-03	3.8E-04	1.0E-02	2.2E-05
7000	4.9E-03	4.6E-04	2.2E-04	7.3E-03	1.3E-05
8000	2.8E-03	2.5E-04	1.5E-04	5.5E-03	9.0E-06
9000	1.9E-03	1.6E-04	1.1E-04	4.4E-03	6.8E-06
10000	1.3E-03	1.1E-04	8.6E-05	3.7E-03	5.4E-06
11000	1.0E-03	8.6E-05	7.1E-05	3.2E-03	4.5E-06
12000	8.3E-04	6.9E-05	6.0E-05	2.8E-03	3.9E-06
13000	6.9E-04	5.7E-05	5.3E-05	2.6E-03	3.5E-06
14000	5.9E-04	4.8E-05	4.8E-05	2.3E-03	3.1E-06
15000	5.1E-04	4.2E-05	4.3E-05	2.2E-03	2.9E-06
20000	3.3E-04	2.6E-05	3.2E-05	1.7E-03	2.2E-06
30000	2.1E-04	1.7E-05	2.4E-05	1.3E-03	1.7E-06
40000	1.8E-04	1.4E-05	2.1E-05	1.2E-03	1.5E-06
50000	1.6E-04	1.3E-05	2.0E-05	1.1E-03	1.4E-06

Contamination

The UVIS detector is regularly monitored for contamination effects i.e. a decline of sensitivity which could be due to volatile molecules accumulating on either the detector itself or on other optical surfaces. When present, contamination is expected to manifest as a wavelength-dependent decline in the photometric throughput, i.e. with the strongest effect present in the bluest filters. Historically, contamination monitoring has been done via observations of the spectrophotometric white dwarf standard GRW+70d5824 in several key filters from 200 nm to 600 nm, with red filters acting as a control ([WFC3 ISR 2011-18](#), [WFC3 ISR 2014-20](#)). In late 2015, a second white dwarf standard star, GD153, was added to the monitoring program ([WFC3 ISR 2017-15](#)). In addition, scanning-mode observations of two standard stars were added to improve the precision of relative photometry over time ([WFC3 ISR 2017-21](#)). [WFC3 ISR 2021-04](#) shows that spatial scan measurements are more precise than staring mode ones: scans have ~ 2.5 less residual noise than their staring mode counterparts (see [WFC3 ISR 2022-04](#)).

For scans, sensitivity losses are relatively flat independent of wavelength on both UVIS detectors, with no evidence of contamination. UVIS2 appears to have slightly higher losses (-0.17 ± 0.01 %/yr) compared to UVIS1 (-0.12 ± 0.01 %/yr). When measured over the same time period, spatial scans and staring mode observations yield filter-dependent loss rates that agree well with each other in most filters within computed uncertainties ([WFC3 ISR 2021-04](#)).

9.1.6 IR Photometric Calibration

For the IR detector, the initial on-orbit photometric calibration in 2009 was based on the average of two CALSPEC spectrophotometric standards (GD153 and P330E, a hot white dwarf and G-type star, respectively). Following the same procedure used for UVIS, a smooth polynomial fit was used to correct for the increase in on-orbit sensitivity with wavelength relative to pre-launch ground tests ([WFC3 ISR 2009-30](#)). By 2012, a larger cumulative set of calibration data allowed for more accurate filter-dependent corrections to the sensitivity and the revised solutions were based on the average of three white dwarf standards (GD153, GD71, G191B2B) plus the G-type star (P330E). While these 2012-era solutions were not documented in a formal ISR, they are available from the [WFC3 photometry webpage](#). Independent calibrations from the four standards agree to within $\sim 1\%$ for most wide- and medium-band filters, and the inverse sensitivity per filter ($PHOTFLAM$) is set to the average of the measurements.

In October 2020, an updated set of inverse sensitivities was calculated and delivered, using the same data from the 2012 calibration plus 8 additional years of monitoring for the same four standard stars plus the white dwarf GRW+70d5824 ([WFC3 ISR 2020-10](#)). The new photometric calibration also included updates to the flat field reference files described in [Section 7.8.3](#) as well as updates to the CALSPEC models. These new model spectra of the standard stars changed their flux, and thus the resulting inverse sensitivities, by approximately 1.5% in the IR ([Bohlin et al. \(2020\)](#); [Section 9.1.3](#)). Observers with data retrieved prior to October 2020 can re-retrieve their data to obtain the most up-to-date calibration.

Current estimates of the photometric internal precision of the inverse uncertainties are $\sim 1\%$ for wide-, medium-, and narrow-band filters. A discussion on the photometric errors of the IR detector is provided in [Section 7.11](#). The IR photometry header keywords are populated by `calwf3` using the `IMPHTTAB` reference file, which is described in [Section 9.1.7](#).


Time-dependent solutions

Photometry of fields centered on external (uncrowded) regions of Galactic globular clusters indicated time-dependent sensitivity losses for WFC3/IR. From F160W stare mode imaging of 47 Tuc, the rate of sensitivity loss was measured to be -0.07 ± 0.01 %/yr ; for F110W, changing sensitivity was measured to be -0.12 ± 0.02 % /yr from M4, and 0.13 ± 0.005 %/yr from Omega Centauri ([WFC3 ISR 2024-06](#)). These results agree well with those obtained by analyzing spatial scan observations of stars in the open cluster M35 in two filters ([WFC3 ISR 2024-01](#)): a sensitivity loss at the rate of -0.065 ± 0.006 %/yr in F140W and -0.157 ± 0.021 %/yr in F098M. Similar loss rates are also seen in integrated grism observations of spectrophotometric standard white dwarfs, with an average sensitivity loss of 0.115 ± 0.008 %/yr in G102 and 0.061 ± 0.007 %/yr in G141. [Figure 9.8](#) presents the stare and scan mode results; for more detail see [WFC3 ISR 2024-01](#).

Summarizing, the sensitivity loss of the IR detector appears to be wavelength-dependent, with the greatest losses in the bluest filters. The rates of sensitivity loss as a function of pivot wavelength measured by using the different methods discussed are shown in [Figure 9.8](#), with the recommended correction factors plotted as colorful dashed lines that span the pivot wavelengths of the filters for which they are recommended.

Currently, these slopes must be manually applied by users correct their photometry as described in [WFC3 ISR 2024-06](#). A new image photometry reference table (IMPHTTAB) with these time-dependent inverse sensitivities will be incorporated into the automated calibration pipeline later in 2024.

Figure 9.8: IR filter pivot wavelength versus measured rates of sensitivity loss, from both directly-derived (solid symbols) and indirectly-derived (G102 and G141, open symbols) photometry. For filters with loss rates calculated from both imaging and spectral data, the points are slightly offset in the x-direction for readability. Recommended correction factors are plotted in their negated form as adopted loss rates and are represented by colorful dashed lines that span the pivot wavelengths of the filters for which they are recommended (see [WFC3 ISR 2024-06](#) for more details on applying these corrections).



Count Rate Non-Linearity

The IR detector exhibits a low-level count rate non-linearity (CRNL) at $\sim 1\%$ per dex over a range of 12 magnitudes (see [Section 7.7](#)). Since bright standard stars (11th magnitude) are used to calibrate the detector, this means that faint source photometry (about 23rd magnitude) using the computed set of IR zero points will be systematically $\sim 4.5\%$ too faint and require manual correction ([WFC3 ISR 2010-07](#) and [WFC3 ISR 2011-15](#)).

Additional data to further quantify and correct for this effect as a function of wavelength were obtained in 2016-2017 (calibration programs 14868 and 14870). [WFC3 ISR 2019-01](#) shows results based on these data and on photometry for a set of faint ($V \geq 17$ mag) spectrophotometric standard DA white dwarfs (the White Dwarf Flux Standards, WDFsS, [Axelrod et al. 2023](#)). The measured CRNL is 0.0077 ± 0.0008 mag/dex, characterized over 16 magnitudes with no apparent wavelength dependence.

9.1.7 Image Photometry Reference Table

After December 2012 (`calwf3` version 3.1.6 and higher), the HST calibration pipeline software (HSTCAL) no longer invoked `synphot` to calculate the photometric keyword values on-the-fly. Instead, it uses an image photometry lookup table (`IMPHTTAB`) for each WFC3 channel which contains the values for `PHOTFLAM` (the inverse sensitivity), `PHOTPLAM` (the filter pivot wavelength), and `PHOTBW` (the filter bandwidth). The active `IMPHTTAB` reference files are listed in [Table 9.2](#) for the UVIS and IR channels and include the latest photometric calibration.

The UVIS chip-dependent calibration implemented in **calwf3** v3.3 required a new 5-extension **IMPHTTAB** reference file to carry additional keywords reflecting the inverse sensitivity for each CCD chip, **PHTFLAM1** and **PHTFLAM2**. When UVIS time-dependent calibration was implemented in **calwf3** v3.5.2, the **IMPHTTAB** was updated to contain additional columns to specify the inverse sensitivity as a function of date (MJD) for each CCD chip.

⚠ NOTE: For observations spanning multiple epochs, users are advised to verify that the same **IMPHTTAB** reference file was used to process all data.

If data from multiple-visit programs are retrieved from MAST at different times (e.g. after the execution of each visit), there may be systematic differences due to changes in the reference files and/or software. In this case, users have two options: 1) retrieve all pertinent datasets for the given science program from **MAST** again; 2) reprocess the RAW files offline with a self-consistent version of **calwf3** and reference files. For manual reprocessing instructions, see the examples in [Section 3.5](#).

Table 9.2: The IMPHTTAB reference file used to populate the photometry keywords in the image headers of UVIS and IR data. For a history of previous calibration reference files, see the WFC3 Photometry Webpage.

Activation Date	Reference File	Description	Documentation
2020 Oct 15	<i>51c1638pi_imp.fits</i>	Provides UVIS time-dependent photometry keywords in the image header corresponding to the observation date; used with <i>calwf3_v3.5.2±</i> . Zero points are now tied to a common reference epoch (June 26, 2009MJD 55008). Adds GRW+70d5824 to the prior set of calibration standards (GD153, GD71, G191B2B, P330E) and uses updated 2020 CALSPEC models.	WFC3 2021-04 ISR
2020 Oct 15	<i>4af1533ai_imp.fits</i>	IR Solutions based on 11 years of monitoring data for five CALSPEC standard stars: GRW+70d5824, GD153, GD71, G191B2B, P330E. Uses updated 2020 CALSPEC models and IR flat fields.	WFC3 2020-10 ISR

9.1.8 Narrow-Band Filter Calibration for Emission-Line Objects

WFC3 has several [narrow-band filters](#) which are often used to study emission-line objects like planetary nebulae and HII regions. To derive the emission-line fluxes in such cases, one must correct the observed flux for signal arising from contaminating lines and the underlying continuum. A detailed method to carry out this correction is described in Appendix A of [O’Dell et al. \(2013\)](#).

¹ The WFC3 UVIS detector suffers degradation of their charge transfer efficiency (CTE), the effect of which is reducing the apparent brightness of sources. There is evidence for CTE degradation over time (see [Section 6.3](#)), and this can be mitigated by using either a pixel-based model ([Section 6.4](#)), or an empirical correction ([Section 6.5](#)).

9.2 Astrometry

[9.2.1 Coordinate Transformations](#)
[9.2.2 Absolute and Relative Astrometry](#)
[9.2.3 Impact of Guide Star Failure](#)

9.2.1 Coordinate Transformations

There are three coordinate systems applicable to WFC3 Images. First, there is the position of a pixel on the geometrically distorted flat-fielded images (flt) after pipeline processing through **calwf3**. Second, there is the pixel position on the drizzled images (drz) created by **AstroDrizzle**, which corresponds to an undistorted pixel position on a tangent plane projection of the sky. Third, there is the corresponding world coordinate (RA, Dec) position on the sky.

Utilities to transform between these coordinate systems are built on the **WCSLIB**, a C library that implements the FITS standards for World Coordinate System (WCS) information. For transformations involving (flt) images, the distortion information in the FITS header is used to correct for the very large effects of geometric distortion. **Drizzlepac** and **astropy** WCS are two commonly-used Python interfaces to this library that implement these transformations.

For example, consider an object found at x,y pixel position (152,156) on UVIS chip 1 on an image (testflt.fits). The position on the celestial sphere and the corresponding pixel position on a distortion-corrected file (testdrz.fits) can be found as follows.

Example #1 Transform pixel position from flat fielded frame (flt) to sky frame using astropy wcs.all_pix2world

```
>>> from astropy.io import fits
>>> from astropy.wcs import wcs
>>> import numpy as np
```

Create wcs object containing WCS keywords from header.

```
>>> flt_wcs = wcs.WCS(fits.getheader('testflt.fits', 1), fits.open('testflt.fits'))
```

Find RA and Dec corresponding to pixel position (152,156) on the image (flt).

```
>>> pixcrd_flt = np.array([[152,156]])
>>> flt_skycoords = flt_wcs.all_pix2world(pixcrd_flt,1)
```

The task `astropy.wcs.all_pix2world` does all three transformations in series (table lookup distortions, SIP, and core WCS) from pixel to world coordinates. SIP (Simple Imaging Polynomial) is a convention for storing distortion information in FITS file headers (e.g. [Shupe et al., 2005](#)).

Example #2 Transform pixel position from flat fielded frame (flt) to drizzled frame (drz) using `astropy.wcs`

Continuing from the code in Example 1, also create wcs object from drizzled image (drz):

```
>>> drz_wcs = wcs.WCS(fits.getheader('test_drz.fits', 1))
```

Use `all_world2pix` to convert the sky position (found in Ex 1.) into pixel position in the drizzled image (drz).

```
>>> drz_pix = drz_wcs.all_world2pix(flt_skycoords, 1)
```

9.2.2 Absolute and Relative Astrometry

Beginning in 2020, the World Coordinate System (WCS) in the header of WFC3 images is corrected by a variety of methods to improve absolute astrometry. All the solutions are stored in an astrometry database and placed in extra [headerlet](#) extensions of WFC3 images (flt, flc). The first method is the update of guide star positions (RA/Dec) in the newest version of the guide star catalog. Due to the motion of the Fine Guidance Sensors in the focal plane, the absolute uncertainty of these astrometric solutions may still be as large as 200 mas. Next, an automatic alignment is attempted, where sources in the image are detected and matched to the [Gaia EDR3](#) catalog. If successful, the absolute astrometry of the images is substantially improved, reducing astrometric uncertainty to less than 20 mas. If sources are unable to be matched to Gaia EDR3, matching to GSC v2.4.2 followed by 2MASS is tried. However, in the cases where this fails, the realigned guide star solution remains active. Please refer to [WFC3 ISR 2022-06](#) for more details on improved absolute astrometry for WFC3 data products.

While aligning to Gaia may provide better absolute astrometry, the alignment uncertainties may result in individual exposures with slightly poorer relative alignment to each other. Restoring the original, not-aligned-to-Gaia solution, where the WCS is based on guide star pointing information and image distortion correction table (`IDCTAB`), may improve relative astrometry, at the cost of absolute astrometric precision. In general, the default `IDCTAB` solution used before aligning to Gaia has very good relative astrometry (typically a few mas), except in the cases of guide star loss or very large dithers. The various alignment solutions can be activated for individual exposures through the `stsci.wcsutil.headerlet` interface. Examples of this interface are presented in a Jupyter notebook [Using Updated Astrometry Solutions](#). Users requiring either precise relative or precise absolute astrometry will want to assess the available solutions, database and interface described in detail in [section 4.5 of the DrizzlePac Handbook](#).


In cases where the astrometric solutions included with the data are insufficient for science purposes, WFC3 images can be realigned using the HST software **DrizzlePac**. Specifically, images can be aligned to each other or to external catalogs using the **DrizzlePac** task **TweakReg**. This is often necessary for comparing images of the same target from different epochs (or visits), or producing mosaicked images with **AstroDrizzle**. In some cases, the aligned-to-Gaia astrometric solutions may also provide sufficient accuracy to create mosaic images, particularly if there are many Gaia stars covering the extent of the mosaic. Examples of various alignment procedures are presented in the [DrizzlePac Notebooks](#).

9.2.3 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. Such observations will experience small drifts around the guide star. To determine the quality of tracking during these observations, please review [Observation Logs](#) in the [Introduction to the HST Data Handbooks](#). In recent years, as gyroscopes have failed and been replaced, the typical drift rate under single-star guiding has increased from 1.5 mas/sec to ≤ 17 mas/sec. This drift causes a rotation of the observed field around the single guide star, which in turn introduces a small translational drift of the targets 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 WFC3, with the current drift rate, the roll about the guide star produces a translation of up to 60 mas (1.5 UVIS pixel, 0.45 IR pixel) in 1000 sec.

The **TweakReg** task may be used to measure and correct for any shifts between successive exposures. The drift over an orbital visibility period can be calculated from this number; the typical visibility period in an orbit (outside the Continuous Viewing Zone [CVZ]) is in the range 52-60 minutes, depending on target declination ([HST Primer](#)). The drifts inherent to single-star guiding are not represented in the image header astrometric information, with two important consequences:

- There will be a slight drift of the target on the detector within a given exposure. Originally, for the majority of observations and scientific programs, this did not degrade the results, especially for short exposures. Normally the drift was smaller than the FWHM of the point spread function (PSF), comparable to the typical jitter during an HST observation (0.003-0.005 arcsec) even when two guide stars were used. Currently however, the drift on the detector could measurably affect the size and shape of the PSF.
- There will be small shifts between consecutive exposures which can build up between orbits in the same visit. This will affect the automated **AstroDrizzle** products from the pipeline since the pipeline relies on the header astrometry, e.g. the structure of sources (i.e. PSFs and extended targets) in the image will be degraded during the cosmic ray rejection routine. However, the user can address these issues by first measuring the shifts and then running **AstroDrizzle** off-line using those shifts.

 Note that even when two guide stars are used, there is often a slow drift of the telescope, up to 0.01" per orbit, due to thermal effects. So, it is generally advisable to check the image shifts, and if necessary measure them to improve the alignment of exposures before running AstroDrizzle off-line to perform the cosmic ray rejection and image combination.

In summary, for most scientific programs, single-star guiding will not necessarily degrade the usefulness of WFC3 data, provided the shifts are measured post-facto and **AstroDrizzle** is re-run offline using those shifts. However, we do not recommend single-star guiding for programs requiring very accurate knowledge of the PSF (e.g. astrometric programs) or for programs that rely critically on achieving a dithering pattern that is accurate on the sub-pixel scale (which often are not achievable even with two-star guiding).

9.3 Spectroscopy

- [9.3.1 Using the WFC3 Grisms](#)
- [9.3.2 Pipeline Calibration](#)
- [9.3.3 Slitless Spectroscopy Data and Dithering](#)
- [9.3.4 Spectroscopy with the WFC3 G280 Grism](#)
- [9.3.5 Spectroscopy with the WFC3 IR Grisms](#)
- [9.3.6 Extracting and Calibrating Slitless Spectra](#)
- [9.3.7 Accuracy of Slitless Spectra Wavelength and Flux Calibration](#)

9.3.1 Using the WFC3 Grisms

WFC3 contains three grism elements: the G280 in the UVIS channel and the G102 and G141 in the IR channel. The grisms provide slitless spectra over the whole field of view of each channel. These spectroscopic modes have the well-known advantages and disadvantages of slitless spectroscopy. The chief advantage is large-area coverage, which enables 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.

In the UVIS channel, the G280 grism provides spectroscopy over a useful wavelength range of 200-800nm, at a dispersion of ~ 14 Å per pixel in the first order. The two grisms for the IR channel cover the wavelength ranges 800-1150 nm (G102) and 1075-1700 nm (G141). The dispersions are 24.5 and 46.5 Å per pixel, respectively. The primary aim of the reduction of WFC3 slitless spectra is to provide one-dimensional wavelength- and flux-calibrated spectra of all objects with detectable spectra. The reduction presents special problems because of the dependence of the wavelength zero point on the position of the object in the field, the blending of spectra in crowded fields, and the need for flat-field information over the whole available wavelength range. **HSTaXe** (Section [9.3.6](#) and [9.5.5](#)), a STScI-supported Python/C based package, is designed for the automated extraction, calibration, and visualization of spectra from slitless spectroscopic instruments. **Slitlessutils**, a new grism analysis package written entirely in Python, is under development at STScI, with an initial version released in 2024 ([STAN 44](#)).

The normal method for taking WFC3 slitless spectra is to take a pair of images, one direct and one dispersed, of each target field, without a shift in position between the two. 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 WFC3 UVIS and IR grisms have some unique properties that result in different types of issues associated with data for the two different channels. These are discussed in more detail below. Some common issues associated with all WFC3 grism observations are highlighted here.

Bright Stars

The brightest objects produce spectra that can extend far across the detector. This is especially problematic for the UVIS G280 grism, where the relative throughput of the higher spectral orders is significant. These spectra provide a strong source of contamination for fainter sources. In addition, the higher order spectra are increasingly out of focus and thus signal spreads in the cross-dispersion direction. Bright stars also produce spatially extended spectra formed by the wings of the PSF.

Resolution and Object Size

In slitless spectroscopy the object itself provides the ‘slit’. The WFC3 PSF has a high Strehl ratio ([Burrows et al. 1991](#)) over most of the accessible wavelength range of the gratings and therefore the degradation of point sources beyond the theoretical resolution is minimal. The spectral resolution for an extended object, however, will be degraded depending on the size and light distribution in the object and spectral features will be diluted.

Zeroth Order

The grism 0th order is only detectable for brighter objects observed with the IR gratings because that order contains only a small fraction of the total flux. This faint feature is therefore easily mistaken for an emission line. The direct image should be used to determine the position of the 0th order for each source, which allows the 0th order feature in the dispersed images to be distinguished from emission lines. For the UVIS G280 grism, the 0th order has high throughput and is therefore more readily distinguished from emission features. The high throughput of the G280 0th order also means that it will often be saturated in long exposures, which leads to CCD charge bleeding and potential contamination of adjacent spectra.

Background

The background in a single grism image pixel is the result of the transmission across the whole spectral range of the disperser and can therefore be high, depending on the spectrum of the background signal. The IR grism background includes not only signal from the sky, but also thermal emission from the telescope and WFC3 optics. The detected background in the IR gratings exhibits variability on timescales shorter than a HST orbit and shows a distinct two-dimensional structure that is due to the overlapping of the background spectral orders. This background needs to be carefully removed before extracting the spectra of targets (see [WFC3 ISR 2020-04](#) for details). The G280 grism, on the other hand, produces relatively low background compared to the IR gratings, because of the faintness of the sky in the near-UV and optical; however, due to the geometry of the optical bench, there is structure in the background. G280 sky frames have been provided on the [Grism Resources](#) page and are described in more detail in [WFC3 ISR 2023-06](#).

Crowding

Because of the high sensitivities of the WFC3 gratings, observations of moderately crowded fields can produce many instances where spectra overlap. It is important to know if a given spectrum is contaminated by that of a neighbor and to choose a telescope roll angle which eliminates or minimizes contamination for specific sources of interest. Contamination can also be mitigated by obtaining grism observations of the same field at different telescope roll angles, which improves the chances of cleanly extracting the spectrum for a given target. The [slitlessutils](#) package is optimized for analyzing spectra obtained at multiple orientations.

Extra-field Objects

There will inevitably be cases where objects outside the field of view result in spectra getting dispersed into the field, contaminating sources within the field. This is more serious for the G280 where the spectra are long relative to the size of the detector ([WFC3 ISR 2020-09](#)). In these cases, reliable wavelengths can not be determined for the extra-field object unless the 0th order is also present. Even then, the wavelength zero point will be relatively uncertain because the 0th order is somewhat dispersed and therefore difficult to localize. The extent of the out-of-field area has been determined for the G102 and G141 IR grisms: extending to pixel -189 on the left of the detector and +85 on the right ([WFC3 ISR 2016-15](#)).

9.3.2 Pipeline Calibration

The direct image of a direct-plus-grism image pair can be fully reduced by **calwf3**, including bias subtraction, dark subtraction, flat fielding, and computation of the photometric zero point in the header. In contrast to direct images, no single flat-field image can be correctly applied to grism images, because each pixel contains signal arising from different wavelengths. Flat fielding must therefore be applied during the extraction of spectra once the wavelength corresponding to each pixel is known from post-pipeline processing with **HSTaXe** or the new [slitlessutils](#). With **HSTaXe**, the user can apply flat-field corrections which are dependent on the wavelength falling on each pixel, as specified by the position of the direct image and the dispersion solution. So during **calwf3** processing the `FLATCORR` step is still performed, but the division is done using a special flat-field reference file that only contains information on the relative gain offsets between the different detector amplifier quadrants. This allows the `FLATCORR` step to still apply the gain correction (converting the data to units of electrons for UVIS or electrons per second for IR) and thus also corrects for offsets in gain between the various quadrants of the detectors. Details of the flat-fielding done by [slitlessutils](#), released in 2024, can be found in the [ReadTheDocs](#).

The **calwf3** fit products should then be the starting point for all subsequent reduction of slitless data with **HSTaXe**, or other software such as [slitlessutils](#). The units of the data in the SCI and ERR extensions of these files are electrons for UVIS and electrons per second for IR. The primary output of **HSTaXe** is a file of extracted, flux calibrated spectra, provided as a FITS binary table with as many table extensions as there are extracted spectra (see [Section 9.3.6](#) for more details). For [slitlessutils](#) software, see the [ReadTheDocs](#) for details of the output products.

9.3.3 Slitless Spectroscopy Data and Dithering

The common approach to dithering WFC3 imaging data, in order to improve the sampling of the PSF and to allow for the removal of bad pixels, applies equally well to slitless spectroscopy data. For long grism observations the data taking is typically broken into several sub-orbit, dithered exposures.

The **AstroDrizzle** task, which is normally used to correct for the geometrical distortion of WFC3 and combine dithered exposures, is not generally applicable to grism observations. This is due to the fact that the spatial distortion correction would only be applicable to the cross-dispersion direction of grism images. For similar reasons, the combining of dithered grism images before extracting spectra is not generally recommended. Every detector pixel has a different spectral response, which has not yet been corrected in the calibrated two-dimensional images (see [Section 9.3.2](#) on flat fielding). Combining dithered grism images before extraction will combine data from different pixels, making it difficult or impossible to reliably flat field and flux-calibrate the extracted spectra. Extracted spectra from dithered images can be properly combined into a final spectrum using the `aXedrizzle` task in the **HSTaXe** package.

AstroDrizzle processing of dithered grism exposures can, however, be useful for simple visual assessment of spectra in a combined image and for the purpose of flagging cosmic-ray (CR) hits in the input fit images. When **AstroDrizzle** detects CR's in the input images, it inserts flags to mark the affected pixels in the DQ arrays of the input fit files. The **HSTaXe** spectral extraction can then be run on these updated fit images utilizing the DQ flags to reject bad pixels; please refer to the `slitlessutils` [ReadTheDocs](#) for its treatment of bad pixels. CR removal is particularly helpful for long UVIS G280 exposures; it's not typically necessary for IR grism images since the IR fit files have already had CR's rejected by the `calwf3` up-the-ramp fitting process.

9.3.4 Spectroscopy with the WFC3 G280 Grism

The filters most often used for obtaining a direct image in tandem with the G280 grism are the F300X and F200LP. 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 G280 wavelength zero point is generally calibrated to an accuracy of about 1 pixel. It is not possible to use the 0th order image of a source in a G280 exposure to establish the source position, because the 0th order is weakly dispersed and prone to saturation effects.

Spectra produced by the G280 grism are oriented in WFC3 images with the positive spectral orders to the left (lower x-axis pixel index) of the 0th order spot, with wavelength increasing to the left. Negative orders are located to the right, with wavelength increasing to the right. The +1st order extends to the left of the 0th order a distance of about 1/4 of the image width. The throughput of the +1st order of the G280 is ~20% larger than that of the negative orders. This leads to heavy overlap of the orders at wavelengths redder than ~400nm. In addition, there is curvature of the spectral traces at the blue ends of the orders. The amplitude of the curvature is about 30 pixels in the detector y-axis. Due to the significant throughput at higher orders, the spectra of very bright objects may extend across nearly the entire field of view of the detector. A number of reports are available with details on the characteristics and calibration of the G280 grism (WFC3 ISRs [2020-09](#), [2017-20](#), [2011-18](#), and [2009-01](#)).

As an example, [Figure 9.8](#) shows a G280 image of the Wolf-Rayet star WR-14, which is used as a wavelength calibrator. Superimposed on the dispersed image is a F300X image, which illustrates the relative location of the direct image of the source (circled in [Figure 9.8](#)). The full 4096-pixel x-axis extent of the detector is shown, which is completely filled by the positive and negative orders of this bright source.

Figure 9.8: The G280 spectral orders on the detector. The circled source is the superimposed position of the F300X direct image. The 0th order is the bright source in the center, with the positive and negative orders extending to the left and right, respectively. The image shows the full 4096-pixel extent of the detector in the x-axis.





Figure 9.9 shows a zoomed view of the first several positive spectral orders of this source, where wavelength increases to the left. The blue end of each order curves upwards, and at longer wavelengths (greater than ~400nm) there is significant overlap of adjacent orders. Very bright sources produce spectra in which orders up to 6-8 can be detected. These spectra, which in principle can be analyzed (although dispersion solutions have been determined for only the first few orders), provide a strong source of contamination for the spectra from fainter objects. In addition, the higher order spectra are increasingly out of focus and thus spread in the cross-dispersion direction.

Figure 9.9: Zoomed view of the G280 positive spectral orders. Overlap between the +1st and +2nd orders occurs for wavelengths greater than about 400nm.



9.3.5 Spectroscopy with the WFC3 IR Grisms

The dispersion of the G102 grism is high enough that only the +1st and +2nd order spectra generally lie within the field of the detector. For the lower-dispersion G141 grism, the 0th, +1st, +2nd, and +3rd order spectra lie within the field for a source that has the +1st order roughly centered. The IR grisms have the majority (~80%) of their throughput in the +1st order, resulting in only faint signals from the other orders. The trace of the observed spectra are well described by a first-order polynomial, however the direct-to-dispersed image offset is a function of the source position in the field. The tilt of the spectra relative to the image axes is small, only 0.5-0.7 degrees. Typical filters used for obtaining companion direct images are F098M and F105W for the G102 grism, and F140W and F160W for the G141 grism. Other medium- and narrow-band filters can be used when necessary to prevent saturation of very bright targets. The image centroids of sources at a given telescope pointing will depend on the filter used. These filter-dependent systematic variations, documented in [WFC3 ISR 2012-01](#), are generally at the sub-pixel level, but must be taken into account during the spectral extraction and reduction process if the filter used to take the direct images is different from the filter in which the trace is calibrated. Version 4.32 of the G102 and G141 grism configuration files now include filter-specific configuration files which include the direct filter wedge corrections (the version number is part of the tar file name, e.g. G102.F098M.V4.32.conf).

The dispersion direction of the IR grisms is opposite to that of the G280, with the positive spectral orders appearing to the right of the 0th order and wavelength also increasing to the right. Examples of G102 and G141 observations of the flux calibration standard star GD153 are shown in [Figure 9.10](#) and [Figure 9.11](#), respectively.

Figure 9.10: A G102 grism observation of the star GD-153. The location of the source in the accompanying direct image is shown superimposed. The entire x-axis extent of the detector is shown.




Figure 9.11: A G141 grism observation of the star GD-153. The location of the source in the accompanying direct image is shown superimposed. The entire x-axis extent of the detector is shown.



9.3.6 Extracting and Calibrating Slitless Spectra

The software package **aXe** provides a streamlined method for extracting spectra from WFC3 slitless spectroscopy data. Due to the deprecation of IRAF/PyRAF, an updated version of **aXe**, called **HSTaXe**, was developed in Python/C that is completely independent of **IRAF/PyRAF**. The **HSTaXe** software requires an environment that is independent of the standard Space Telescope Environment (**stenv**). It is recommended that **HSTaXe** be installed with the [environment file](#) hosted on the [HSTaXe GitHub repository](#). A new STScI package, **slitlessutils**, is currently under development to support wide-field slitless spectroscopy. An early release with basic functionality is available in [STAN 44](#). The [Grism Data Analysis page](#) provides the latest status on the available analysis software.

There is a detailed [aXe manual](#) and [Jupyter Notebook tutorials](#) specific to WFC3 grism data reduction, both of which are available on the [HSTaXe GitHub repository](#), so only a brief outline of its use is presented here. In addition to the detailed cookbooks, we also have a complementary Instrument Science Report ([WFC3 ISR 2023-07](#)) that describes the **HSTaXe** installation, recommended preprocessing steps, file outputs, and an advanced extraction method. Other software packages are also available and a current list can be obtained on the [Grism Resources](#) webpage.

The basic steps involved in extracting spectra from grism images are:

- **Make a direct image source catalog.** This step consists of identifying and cataloging sources from a direct image in the field of the grism image. The source positions and sizes are used later to define extraction boxes and calculate wavelength solutions in the extraction step. The source information is often derived from an **AstroDrizzle** combination of direct images. If there is only a single direct image, it is still recommended to process the file through **AstroDrizzle**.
- **Prepare the grism images and remove sky background.** In this step a scaled master sky background image is subtracted from the grism images.
- Project source catalog positions to coordinate system of direct images. If the source catalog was derived from dithered or drizzled direct images, then the catalog positions need to be transformed back to the coordinate system of each direct image.
- **Extract sets of pixels for each object spectrum.** The spectra of all objects in the transformed catalog are extracted from each grism image.
- **Combine all spectra of each object using aXedrizzle.** All 2-dimensional spectra for each object are combined and CR-rejected using drizzle techniques. The results are 2-d spectral images and 1-d tables.

The starting point is always a set of dispersed slitless images and the derived catalog of objects in the images. 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 per electron per second) table are stored in a configuration file, which enables the full calibration of extracted spectra. For each instrumental configuration, the configuration files and all necessary calibration files for flat fielding and flux calibration can be downloaded from the WFC3 instrument website (<http://www.stsci.edu/hst/instrumentation/wfc3/documentation/grism-resources>).

Background Subtraction

aXe has two different strategies for removal of the sky background from the spectra. The first strategy is to perform a global subtraction of a scaled "sky" frame from each input grism image at the beginning of the reduction process. This removes the background signature from the images, so that the remaining signal can be assumed to originate from the sources only and is extracted without further background correction in the **aXe** reduction. Sky frames are available for download from the WFC3 grism resource webpages ([UVIS](#), [IR](#)). While a "sky" image is available for both UVIS and IR, we recommend an alternate background subtraction method for IR grism observations, which is independent of **aXe**.

When reducing G280 spectra, application of a sky image is recommended for use with your slitless spectroscopy reduction code (e.g., **HSTaXe**) of choice. We provide the G280 sky images for both calibrated, flat-fielded individual FLT exposures, as well as their corresponding CTE-corrected FLC frames on the [Grism Resources](#) webpage. This sky image was generated using an inverse-variance weighted stack of all low-background G280 exposures taken to-date with sources masked, to characterize the stray light present in the G280 science frames ([WFC3 ISR 2023-06](#)).

For both G102 and G141, the background sky signal is variable and made of multiple components. Currently, the WFC3 calibration pipeline, **calwf3**, does not have the capability to model and remove the dispersed 2D background. Therefore, we highly advise that IR grism observers use the WFC3 Backsub Python script to process uncalibrated RAW images into calibrated, flat-fielded FLT files. A description of the three background components and the methods used in WFC3 **Backsub** can be found in [WFC3 ISR 2020-04](#). The software and supporting reference files for WFC3 **Backsub** are available for download on [Box](#). An example workflow incorporating the use of WFC3 **Backsub** is given in the [HSTaXe IR cookbook](#).

The second strategy for background subtraction with **aXe** is to make a local estimate of the sky background for each beam by interpolating between the adjacent pixels on either side of the beam. In this case, an individual sky estimate is made for every beam in each science image. This individual sky estimate is processed (flat fielded, wavelength calibrated) parallel to the original beam. Subtracting the 1D spectrum extracted from the sky estimate from the 1D spectrum derived from the original beam results in the pure object spectrum. The second approach needs to account for the fact that the background of an observation can vary during the course of an observation. As is the case for direct imaging, there can be a significant amount of dispersed HeI light as part of the background. Steps to mitigate this problem are described in [WFC3 ISR 2017-05](#). An updated discussion of variable background subtraction methods for the IR grisms and an improved set of dispersed background models are provided in [WFC3 ISR 2020-04](#).

Output Products from aXe

The primary output of **aXe** is a file of extracted and calibrated spectra, which is provided as a multi-extension FITS binary table with as many table extensions as there are extracted spectra. The table contains 15 columns, including wavelength, total and extracted and background counts and their errors, the calibrated flux and error, the weight and a contamination flag. The primary header of this "SPC" table is a copy of the header of the frame from which the spectrum was extracted

aXe also creates a 2-d "stamp" image for each beam. The stamp images of all spectra extracted from a grism image are stored as a multi-extension FITS (STP) file with each extension containing the image of a single extracted spectra. It is of course also possible to create stamp images for 2-d drizzled grism images.

Table 9.5 below, from [WFC3 ISR 2023-07](#), lists the file categories and short descriptions for the various files outputted after a basic extraction with **HSTaXe**. The output products from **aXe** consist of ASCII files, FITS images and FITS binary tables. The [WFC3 grism cookbooks](#) show end-to-end examples of using HSTaXe and handling the output files to display the 2-d stamps and make plots of the extracted spectra.

Table 9.5: HSTaXe Output Files from Basic Extraction

File Suffix	Description
cat	Input object catalog
OAF	Original aperture file
BAF*	Background aperture file
BCK.fits*	Background image
BCK.PET.fits*	Pixel extraction tables for the background
CONT.fits	Contamination estimate for spectra
PET.fits	Pixel extraction tables
SPC.fits	1D extracted spectra
STP.fits	2D stamp image of extracted traces
opt.SPC.fits†	Optimally extracted 1D spectra
opt.WHT.fits†	Optimal weight image

* Files produced if the argument **back=True** in the **axecore** task

† Files produced if the argument **weights=True** in the **axecore** task

9.3.7 Accuracy of Slitless Spectra Wavelength and Flux Calibration

Wavelength Calibration

The WFC3 grism dispersion solutions were established by observing both astronomical sources with known emission lines (e.g., the Wolf-Rayet star WR-14 and the planetary nebula Vy2-2 (WFC3 ISRs [2009-17](#) and [2009-18](#)) and ground-based monochromator sources (see WFC3 ISRs [2009-01](#) and [2008-16](#)). The field variation of the dispersion solution was mapped by observing the same source at different positions over the field. The internal accuracy of these dispersion solutions is good to ~ 0.25 pixels for the IR grisms ($\sim 6\text{\AA}$ and $\sim 9\text{\AA}$ for the G102 and G141, respectively), and to ~ 1 pixel ($\sim 14\text{\AA}$) for the UVIS G280.

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 (recorded in spt file header keywords `DGESTAR` and `SGESTAR`), 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. The wavelength calibration of the G102, G141 and G280 grisms were recently updated (WFC3 ISR [2016-15](#) and WFC3 ISR [2020-09](#)). The new aXe format calibration files also include the filter-dependent wedge offsets that should be applied when pairing G102 and G141 observations with imaging obtained using a specific direct filter.

Flux Calibration

The sensitivity of the dispersers was established by observing a spectrophotometric standard star at several positions across the field. The sensitivity (HSTaXe uses a sensitivity tabulated in $\text{erg cm}^{-2} \text{s}^{-1} \text{\AA}^{-1}$ per detected electron) was derived using data flat fielded by the flat-field cube. Results for the IR grisms show 4-5% differences in the absolute flux of spectra located near the center of the field as compared to those near the field corners. This is clear evidence for a large-scale variation in the overall illumination pattern in the grism flat-field data cubes. Additional field-dependent flux calibration observations have been completed for G280 in 2020 (WFC3 ISR [2020-09](#)). The last full-scale flux calibration for G102 and G141 were completed in 2011 (WFC3 ISR [2011-05](#)).

UVIS Flux Calibration

The sensitivity of the G280 grism was established by observing a spectrophotometric standard star at several positions across the field. The sensitivity (HSTaXe uses a sensitivity tabulated in $\text{erg cm}^{-2} \text{s}^{-1} \text{\AA}^{-1}$ per detected electron) was derived using data flat fielded by the flat-field cube. The results of a full calibration of the WFC3 UVIS G280 slitless spectroscopic mode from the entire G280 calibration data (nearly 600 datasets) obtained by 2020 are detailed in WFC3 ISR [2020-09](#).

IR Flux Calibration


The sensitivities of the G102 and G141 grisms were established by observing the white dwarf spectrophotometric standard stars GD153 and GD71 at several positions across the field (WFC3 ISRs [2009-17](#), [2009-18](#), and [2011-05](#)).

The flux calibration is primarily based on observations of GD153 at the center of the detector. Variation of the sensitivity over the field-of-view was measured by observing GD71 over a grid of nine points over the detector, which shows that the throughput falls off by about 2% in the y direction and by about 4% in the x direction to the edges of the FoV. The large-scale variation is included in the grism flat-field cube (see [WFC3 ISR 2011-05](#)). Regular monitoring of the flux calibration is performed, and has shown that the calibration has been stable to within 1% since the installation of the camera on HST (WFC3 ISRs [2012-06](#) and [2014-01](#)).

Flux Monitor

The IR grism flux monitor checks the sensitivity of the grisms using observations of CALSPEC white dwarf standards. Recent analysis shows that the relative sensitivity of both grisms is decreasing at a rate of 0.115+/-0.008% per year for G102 and 0.061+/-0.007% per year for G141 ([WFC3 ISR 2024-01](#) ; [Bohlin & Deustua \(2019\)](#)), shown in [Figure 9.12](#). The vertical scatter in the data for a given date is due to dithered observations in the calibration data, designed to sample many different regions of the detector. This scatter will be much smaller for observations with the supported [IR grism apertures](#) G102-REF and G141-REF near the center of the detector.

Figure 9.12: Relative photometry of four CALSPEC white dwarf flux standards for the two IR grisms. A slight loss in sensitivity of ~0.115 percent per year is observed for G102 and 0.06 percent per year for G141.



9.4 STSDAS, STSCI_PYTHON and Astropy

STSCI/IRAF has been deprecated as described in the [2018 STScI Newsletter](#). To ease this transition for long-time IRAF users, a number of IRAF routines have been ported to [python](#) and documented via a set of Jupyter Notebook Tutorials, located in a [Central Repository](#). The WFC3 team continues to develop [tools](#) aimed at the user community. New tools will continue to be added. We note that many of the software tools that are useful for analyzing HST data can now be found in the [Astropy](#) (the current community-supported Python Astronomy package) library of python modules.

9.5 Specific Tools for the Analysis of WFC3

- [9.5.1 wfc3tools](#)
- [9.5.2 Point Spread Function Modeling](#)
- [9.5.3 WFC3 Photometry Tools](#)
- [9.5.4 Code for mitigation of variable IR background](#)
- [9.5.5 DASH \(Drift-And-SHift\) Reduction](#)
- [9.5.6 Grism Reduction Tools](#)
- [9.5.7 PandExo: A community tool for transiting exoplanet science with JWST & HST](#)
- [9.5.8 Crosstalk correction software](#)
- [9.5.9 Satellite trail detection](#)
- [9.5.10 IDL procedures for simulating trajectories of multi-lined spatial scans](#)

This section describes existing tools and packages that can be used for the analysis of WFC3 data. Some of these tools are distributed as STScI affiliated packages, while others have been developed by STScI scientists for their own scientific projects, but have also been made available to the community. The latter type of software is not directly supported by the WFC3 team; thus users are directed to the software developers for assistance. For our latest software releases, refer to our [Software Tools](#) page.

New tools for WFC3 analysis will be stored in [WFC3 Notebooks](#), which is a child GitHub repository to [HST Notebooks](#). The goal of the repository is to organize and standardize our Jupyter notebooks by providing internal team reviews on documentation and code testing. This repository, aimed at the user community, is maintained and supported through any future software updates by employing continuous integration and continuous delivery (CI/CD).

9.5.1 wfc3tools

wfc3tools is a Python package containing several WFC3-specific tools. Online documentation for **wfc3tools** can be found on read the docs at: <http://wfc3tools.readthedocs.io/en/latest/index.html>

The package is available on [Github](#) and is also distributed on the STScI-maintained Space Telescope Environment (**stenv**) channel. **wfc3tools** contains the python wrapper modules that call the **calwf3** pipeline executables (whose source code is written in C), as well as other auxiliary functions. The pipeline modules (**calwf3**, **wf3cte**, **wf3ccd**, **wf32d**, **wf3rej**, **wf3ir**) are described in detail in [Section 3.4](#) of this book, along with an example of **calwf3** manual reprocessing in [Section 3.5.2](#). Here we briefly describe the other tools. The boldface paragraph titles correspond to the module name and link to the read-the-docs resources which contain a more detailed documentation.

[embedsub](#)

Given an image specified by the user which contains a subarray readout, return a full-frame image with the subarray implanted at the appropriate location.

[pstat](#)

Plot statistics for a specified image section up the ramp of an IR MultiAccum image. Sections from any of the SCI, ERR, DQ, image extensions can be plotted. A choice of mean, median, mode, standard deviation, minimum and maximum statistics is available.

[pstack](#)

Plot the stack of MultiAccum sample values for a specified pixel in an IR multiaccum image. Pixels from any of the SCI, ERR, DQ, or TIME image extensions can be plotted.

[sampinfo](#)

Prints information about a WFC3/IR MultiAccum image, including exposure time information for the individual samples (readouts). The global information listed (and the names of the header keywords from which it is retrieved) includes:

- the total number of image extensions in the file (`NEXTEND`)
- the name of the MultiAccum exposure sample sequence (`SAMP_SEQ`)
- the total number of samples, including the “zeroth” read (`NSAMP`)
- the total exposure time of the observation (`EXPTIME`).

[sub2full](#)

Given an image specified by the user which contains a subarray readout, return the location of the corner of the subarray in a full frame reference image (including the full physical extent of the chip), in 1-indexed pixels. If the user supplies an X and Y coordinate, then the translated location of that point will be returned.

9.5.2 Point Spread Function Modeling

A new WFC3 notebook, located on the [HST notebooks webpage](#), demonstrates how to generate PSF models for WFC3 observations. The new tool provides users with several workflows depending on their science goals and available data, including 1) downloading empirical library PSF models for high-precision stellar photometry and astrometry; 2) extracting and median stacking stars for characterizing extended PSF wing emission and diffraction spikes; and 3) querying the MAST PSF image library to retrieve and stack stars from archival observations when analyzing sparse fields with very few stars. This tool has the ability to model PSFs for both individual exposures (FLT/FLCs), as well as drizzled data products (DRCs/DRZs). While the majority of the examples are based on WFC3, the code can also model ACS and WFPC2 observations.

9.5.3 WFC3 Photometry Tools

WFC3 zeropoints

A [Jupyter Notebook](#) (also linked in the WFC3 section of the [HST Notebooks](#) repository) shows how to use `stsynphot` to compute photometric keyword values such as the inverse sensitivity (`PHOTFLAM`), pivot wavelength (`PHOTPLAM`) and filter bandwidth (`PHOTBW`) for any WFC3 'obsmode', which is a combination of 'instrument, detector, filter, date, and aperture'. The tool also computes zeropoint values (`STMAG`, `ABMAG`, `VEGAMAG`) and is especially useful for Vegamag zeropoints which require an input spectrum. The notebook may also be used for determining time-dependent WFC3/UVIS zeropoints for any observation date, as the values given in [WFC3 ISR 2021-04](#) are defined for the June 2009 reference epoch (as of mid-2021, the WFC3/IR zeropoints are not time-dependent). The python code works for both the UVIS and IR detectors, has the capability to loop over multiple filters, and optionally creates and plots the 'total system throughput' tables for each obsmode.

UVIS time-dependent photometry

A [Jupyter Notebook](#) (also linked in the WFC3 section of the [HST Notebooks](#) repository) shows how to work the new time-dependent UVIS calibration for observations of which span a range of dates and therefore have different photometric keyword values populated in the image headers. The notebook uses sample images of the standard star GD153 acquired at three epochs and shows how to compute aperture photometry, apply the new time-dependent `PHOTFLAM` keywords, and plot the corresponding countrates and magnitudes.

The unique zeropoint values must be accounted for prior to combining UVIS observations over multiple epochs with **AstroDrizzle**, and the notebook shows how to equalize the countrate values in the science array of each input FLC image prior to drizzling.

WFC3 synthetic photometry examples

A [Jupyter Notebook](#) (also linked in the WFC3 section of the [HST Notebooks](#) repository) replaces pysynphot examples from the 2018 version of the Data Handbook and demonstrates how to use stsynphot for several use cases:

- Compute the inverse sensitivity, zeropoint, and encircled energy correction for any WFC3 'obsmode'
- Renormalize a spectrum to 1 count/sec in a given bandpass and output the predicted magnitude or flux for a different bandpass
- Determine the color transformation between two bandpasses for a given spectrum
- Compute color terms for UV filters for a blue versus a red standard star observed on UVIS2.

WFC3 photometric conversion tool

A [Jupyter notebook](#) (also linked in the WFC3 section of the [HST Notebooks](#) repository) demonstrates how to calculate photometric transformation coefficients between WFC3/UVIS wide-band filters and any other non-HST filter system for a given object spectrum. The new tool uses the latest WFC3 synthetic throughput tables and replaces functionality provided in the WFC3 Photometric Conversion Tool, which is no longer supported. For more detail on photometric transformations to other systems, see [WFC3 ISR 2014-16](#).

Flux converter tool

A [Jupyter notebook](#) (also linked in the [HST Notebooks](#) repository) provides a framework for users to convert between multiple magnitude and flux unit systems based on a user-defined input spectrum. This tool is based on the NICMOS unit conversion form and replaces the HST Unit conversion tool, which is no longer supported. The notebook incorporates the latest WFC3/UVIS ([WFC3 ISR 2021-04](#)) and WFC3/IR ([WFC3 ISR 2020-10](#)) photometric calibration as well as recent changes in the Vega spectrum of up to ~1.5% ([Bohlin et al. 2020](#)).

9.5.4 Code for mitigation of variable IR background

Strategies for reprocessing images with variable background are described in [WFC3 ISR 2016-16](#). Examples include 1) correcting for Helium I atmospheric emission at 1.083 microns and 2) excising reads impacted by scattered light from the bright Earth limb. This ISR provides a full description of the methods and a link to the original [python software](#). This software has since been implemented in WFC3 [Jupyter notebooks](#) (linked in the [HST Notebooks](#) repository) for reprocessing IR images with variable background, as described below. Additional examples are provided in the Appendix of [ISR 2021-01](#), which describes the reprocessing of archival observations impacted by variable background in order to compute sky flats. For more details, see [Section 7.10](#).

To aid the user in identifying WFC3/IR images affected by variable background, the Jupyter notebook “[WFC3/IR IMA Visualization Tools with an Example of Time Variable Background](#)” found on the [WFC3 Notebooks GitHub](#) page (also linked in the [HST Notebooks](#) repository) provides visualization tools to inspect individual reads and plot the accumulated signal through the exposure.

Example commands in [Section 3.5.2](#) demonstrate how to diagnose calibrated WFC3/IR images with poor quality ramp fitting due to time-variable background during the exposure. Images are reprocessed using the 'Last-minus-first' technique described in [WFC3 ISR 2016-16](#). This turns off **calwf3**'s ramp fitting step (`CRCORR`) and treats the IR detector like a CCD that accumulates charge and is read out only at the end of the exposure. A step-by-step walkthrough of this method of reprocessing WFC3/IR images can be found in the Jupyter notebook “[Manual Recalibration with calwf3: Turning off the IR Linear Ramp Fit](#)” on the [WFC3 Notebooks GitHub](#) page.

Three other methods for correcting WFC3/IR images affected by variable background can be found in Jupyter notebooks on the [WFC3 Notebooks GitHub](#) page. Each notebook illustrates a unique method for correcting the IR images, including: 1) 'flattening' the ramp by subtracting the excess background per read in the notebook “[Correcting for Helium Line Emission Background in IR Exposures using the “Flatten-Ramp” Technique](#)”, 2) manually excluding reads impacted by scattered light in the notebook “[Correcting for Scattered Light in WFC3/IR Exposures: Manually Subtracting Bad Reads](#)” or 3) masking specific reads in the RAW image and reprocessing with **calwf3** in the notebook “[Correcting for Scattered Light in WFC3/IR Exposures: Using calwf3 to Mask Bad Reads](#)”.

While time-variable background also impacts the IR grisms, the methods used for imaging data should not be used to correct G102 and G141 observations which are affected by a combination of Helium I, Zodiacal background, and "Scatter" Earth light, as each of these varies spatially across the detector. More detail on correcting grism data for time-variable background is provided in [WFC3 ISR 2017-05](#) and [WFC3 ISR 2020-04](#).

9.5.5 DASH (Drift-And-SHift) Reduction

Software is available to aid users in properly reducing IR DASH observations. While there are portions of the pipeline that will need to be specialized depending on the specific observation strategy, the general outline of steps should be useful to all users in their reduction. [WFC3 ISR 2021-02](#) walks users through this package and provides an accompanying [Jupyter notebook](#) which outlines a strategy based on best practices ([Momcheva et. al \(2017\)](#)).

9.5.6 Grism Reduction Tools

HSTaXe

HSTaXe is a Python/C package for the calibration, extraction, and visualization of spectra from HST slitless spectroscopic observations. **HSTaXe** replaced **aXe** (now retired) as the official STScI supported tool for the reduction of slitless grism/prism data from HST. **HSTaXe** has similar functionalities as **aXe** but does not require IRAF/PyRAF (no longer supported by STScI) for its functioning. **HSTaXe** can be obtained from [this github repository](#) which also hosts a collection of [Jupyter Notebook tutorials](#) showcasing cookbook-style WFC3 grism data reduction workflows. A description of the **HSTaXe** data reduction workflow including various recommended data preprocessing procedures is presented in [WFC3 ISR 2023-07](#). Please see the [WFC3 grism data analysis page](#) for additional and most current information on grism calibration and data analysis.

Slitlessutils

Slitlessutils is a Python-only package for simulating and extracting slitless spectroscopy data from all active HST grism and prism modes (ACS/WFC; ACS/SBC; WFC3/IR; and WFC3/UVIS). It is planned to succeed **HSTaXe** as an STScI-supported slitless spectroscopy analysis tool. **Slitlessutils** implements the LINEAR algorithm ([Ryan, Casertano, & Pirzkal 2018](#)) for multi-oriented fields, and the **HSTaXe** extraction methodology for single-orients. It also includes several utilities to help prepare spectroscopic data for analysis. An early release version of **Slitlessutils** can be downloaded from [GitHub](#) or through [pypi](#), with active development and support for the package ongoing. Please see the [WFC3 grism data analysis page](#) for the most current information on grism calibration and data analysis.

Grizli

Grism redshift & line analysis software for space-based slitless spectroscopy (**Grizli**) is not an official STScI software, but was developed by WFC3 scientists. It is intended to offer general techniques for manipulating HST slitless spectroscopic observations. **Grizli** provides software kernels for the end-to-end processing, quantitative and comprehensive modeling, and fitting of WFC3/IR grism data and is available, along with examples and documentation, on [github](#).

9.5.7 PandExo: A community tool for transiting exoplanet science with JWST & HST

PandExo is not official WFC3 software, but WFC3 scientists have contributed significantly to its development. Similar to an exposure time calculator (ETC), **PandExo** is a transiting exoplanet noise simulator which can be used to create simulated observations and estimate realistic transit depth precisions. It is based on **Pandeia**, the ETC for JWST, and has been expanded to include HST's WFC3 instrument. **PandExo_HST** can be called and used locally for detailed calculations, but it is also available as a web-based tool in the [Exoplanet Characterization Tool Kit \(ExoCTK\)](#) for one-off calculations. A description of PandExo and how it works can be found in [Batalha et al. \(2017\)](#) and the accompanying [github Readme](#).

9.5.8 Crosstalk correction software

Electronic crosstalk between the UVIS amplifiers during readout induces faint, negative, mirror-symmetric ghost images in the other quadrant of the same CCD chip at $\sim 10^{-4}$ levels (see [Section 5.5.3](#) for more details). Standalone software for crosstalk removal ([WFC3 ISR 2012-02](#)) can be downloaded from: <http://www.stsci.edu/hst/instrumentation/wfc3/software-tools/crosstalk>.

9.5.9 Satellite trail detection

The WFC3 team performs a visual inspection of all images acquired and flags those containing satellite trails (as well as other artifacts). The resulting database is available for download from the [WFC3 Performance page](#); please see the accompanying report ([WFC3 ISR 2020-02](#)) for details.

The ACS instrument team has developed a code for finding satellite trails in their data, and flagging the interested pixels' DQ extension accordingly. The software is part of the **acstools** package with detailed documentation at: <http://acstools.readthedocs.io/en/latest/satdet.html>.

The ACS instrument team is also developing upgrades to this package (see, e.g., [ACS ISR 2022-08](#)), which will be included in future versions of **acstools**. Flagging such trails can be useful e.g., in combining multiple images with **Astrodrizzle**.



Disclaimer: the software is developed and tested only for ACS data, but should work on WFC3 data as well.

9.5.10 IDL procedures for simulating trajectories of multi-lined spatial scans

The [WFC3 ISR 2017-06](#) describes simulations of spatial scans using a simple physical model of HST motions during rarely-used multi-lined spatial scans (single-line scans are much more common, e.g. for time-series spectrophotometry of exoplanet transits). The document contains an IDL code in the appendix that can be used both in designing multi-lined spatial scans, as well as for analyzing existing ones (e.g. PID 16983).

9.6 References

- T. Axelrod, 2023, *All-sky Faint DA White Dwarf Spectrophotometric Standards for Astrophysical Observatories: The Complete Sample*, [ApJ](#), **951**, 1.
- V. Bajaj, A. Calamida, and J. Mack, 2020, *Updated WFC3/IR Photometric Calibration*, [WFC3 ISR 2020-10](#).
- N. Batalha, et al., 2017, *PandExo: A Community Tool for Transiting Exoplanet Science with JWST & HST*, [PASP](#) **129**, 064501.
- R. C. Bohlin and S. E. Deustua, 2019, *CALSPEC: Wide Field Camera 3 Infrared Grism Spectrophotometry*, [AJ](#) **157**, 229.
- R. C. Bohlin and R. Gilliland, 2004, *Hubble Space Telescope Absolute Spectrophotometry of Vega from the Far-Ultraviolet to the Infrared*, [AJ](#), **127**, 3508.
- R. C. Bohlin, 2014, *Hubble Space Telescope CALSPEC Flux Standards: Sirius (and Vega)*, [AJ](#), **147**, 127.
- R. C. Bohlin, I. Hubeny, and T. Rauch, 2020, *New Grids of Pure-hydrogen White Dwarf NLTE Model Atmospheres and the HST/STIS Flux Calibration*, [AJ](#), **160**, 21.
- Burrows, Christopher J., Holtzman, Jon A., Faber, S. M., Bely, Pierre Y., Hasan, Hashima, Lynds, C. R., Schroeder, Daniel, *The Imaging Performance of the Hubble Space Telescope*, [ApJL](#), **369**, L21.
- G. Brammer, 2016, *Reprocessing WFC3/IR Exposures Affected by Time-Variable Backgrounds*, [WFC3 ISR 2016-16](#).
- T. Brown, *WFC3 TV3 Testing: IR Channel Blue Leaks*, [WFC3 ISR 2008-05](#).
- T. Brown, *WFC3 TV3 Testing: Red Leak Checks for the UV Filters*, [WFC3 ISR 2008-49](#).
- A. Calamida, J. Mack, S. Deustua, and E. Sabbi, 2018, *WFC3 color term transformations for UVfilters*, [WFC3 ISR 2018-08](#).
- A. Calamida et al., 2022, *New Photometric Calibration of the Wide Field Camera 3 Detectors*, [AJ](#), **164**, 32.
- A. Calamida, J. Mack, J. Medina, C. Shahanan, V. Bajaj, and S. Deustua, 2021, *New time-dependent WFC3 UVIS inverse sensitivities*, [WFC3 ISR 2021-04](#).
- A. Calamida, V. Bajaj, J. Mack, M. Marinelli, J. Medina, A. Pidgeon, V. Kozhurina-Platais, C. Shanahan, and D. Som, 2022, *New Photometric Calibration of the Wide Field Camera 3 Detectors*, [AJ](#), **164**, 32.
- A. W. J. Cousins, 1974, *Standard Stars for VRI Photometry with S25 Response Photocathodes*, [MNRAS](#), **33**, 149.
- F. Dauphin, J. Anderson, V. Bajaj, L. Dressel, K. Sahu, M. Bourque, C. Shanahan, 2021, *The WFPC2 and WFC3 PSF Database*, [WFC3 ISR 2021-12](#).
- S. E. Deustua et al., 2016, *UVIS 2.0 Chip-dependent Inverse Sensitivity Values*, [WFC3 ISR 2016-03](#).
- S. E. Deustua, R. C. Bohlin, J. Mack, V. Bajaj, H. Khandrika, and E. Sabbi, 2017, *WFC3 Chip Dependent Photometry with the UV filters*, [WFC3 ISR 2017-07](#).

- S. E. Deustua, J. Mack, V. Bajaj, and H. Khandrika, 2017, *WFC3/UVIS Updated 2017 Chip-Dependent Inverse Sensitivity Values*, [WFC3 ISR 2017-14](#).
- S. E. Deustua and J. Mack, 2018, *Comparing the ACS/WFC and WFC3/UVIS Calibration and Photometry*, [WFC3 ISR 2018-02](#).
- C. M. Gosmeyer, S. Baggett, S. E. Deustua, and D. M. Hammer, 2014, *Update on the WFC3/UVIS Stability and Contamination Monitor*, [WFC3 ISR 2014-20](#).
- C. M. Gosmeyer and S. Baggett, 2017, *WFC3/UVIS External CTE Monitor: Single-Chip CTE Measurements*, [WFC3 ISR 2016-17](#).
- G. F. Hartig, *WFC3 UVIS PSF Evaluation in Thermal-Vacuum Test #1*, [WFC3 ISR 2005-10](#).
- G. F. Hartig, 2009, *WFC3 SMOV Programs 11437/9: IR On-orbit PSF Evaluation*, [WFC3 ISR 2009-37](#).
- G. F. Hartig, 2009, *WFC3 SMOV Programs 11436/8: UVIS On-orbit PSF Evaluation*, [WFC3 ISR 2009-38](#).
- J. Holtzman et al. 1995, *The Photometric Performance and Calibration of WFPC2*, [PASP, 107, 1065](#).
- H. L. Johnson and W. W. Morgan, 1953, *Fundamental stellar photometry for standards of spectral type on the revised system of the Yerkes spectral atlas*, [ApJ, 117, 313](#).
- J. S. Kalirai et al, 2009, *WFC3 SMOV Proposal 11450: The Photometric Performance and Calibration of WFC3/UVIS*, [WFC3 ISR 2009-31](#).
- J. S. Kalirai et al, 2009, *WFC3 SMOV Proposal 11451: The Photometric Performance and Calibration of WFC3/IR*, [WFC3 ISR 2009-30](#).
- J. S. Kalirai et al, 2010, *WFC3 Pixel Area Maps*, [WFC3 ISR 2010-08](#).
- H. Khandrika, S. E. Deustua, and J. Mack, 2018, *WFC3/UVIS - Temporal and Spatial Variations in Photometry*, [WFC3 ISR 2018-16](#).
- J. Koorneef et al., 1986, in [Highlights of Astronomy IAU, Vol.7, p833](#), ed. J.-P. Swings.
- V. Kozhurina-Platais and S. Baggett, 2020, *WFC3 IR sensitivity over Time*, [WFC3 ISR 2020-05](#).
- B. Kuhn, D. Som, A. Pidgeon, N. Hathi, R. Ryan, R. Avila, N. Pirzkal, *HSTaXe - ACS & WFC3 Cookbook Tutorials*, [WFC3 ISR 2023-07](#).
- M. Kümmel et al., 2009, *The Slitless Spectroscopy Data Extraction Software aXe*, [PASP, 121, 87, 59](#).
- H. Kuntschner, M. Kummel, J. R. Walsh, and H. Bushouse, 2011, *Revised Flux Calibration of the WFC3 G102 and G141 grisms*, [WFC3 ISR 2011-05](#).
- H. Kuntschner, H. Bushouse, M. Kummel, and J. R. Walsh, 2009, *WFC3 SMOV proposal 11552: Calibration of the G141 grism*, [WFC3 ISR 2009-17](#).
- H. Kuntschner, H. Bushouse, M. Kummel, and J. R. Walsh, 2009, *WFC3 SMOV proposal 11552: Calibration of the G102 grism*, [WFC3 ISR 2009-18](#).
- H. Kuntschner, H. Bushouse, J. R. Walsh, and M. Kummel, 2008, *The TV3 ground calibrations of the WFC3 NIR grisms*, [WFC3 ISR 2008-16](#).

- H. Kuntschner, H. Bushouse, M. Kummel, and J. R. Walsh, 2009, *The ground calibrations of the WFC3 /UVIS G280 grism*, [WFC3 ISR 2009-01](#).
- J. C. Lee, N. Pirzkal, B. Hilbert, *Flux Calibration Monitoring: WFC3/IR G102 and G141 Grisms*, [WFC3 ISR 2012-06](#).
- J. C. Lee, N. Pirzkal, B. Hilbert, *Flux Calibration Monitoring: WFC3/IR G102 and G141 Grisms*, [WFC3 ISR 2014-01](#).
- J. Mack, E. Sabbi, and T. Dahlen, 2013, *In-flight Corrections to the WFC3 UVIS Flat Fields*, [WFC3 ISR 2013-10](#).
- J. Mack, T. Dahlen, E. Sabbi, and A. S. Bowers, 2016, *UVIS 2.0: Chip-Dependent Flats*, [WFC3 ISR 2016-04](#).
- J. Mack, 2016, *UVIS 2.0: Ultraviolet Flats*, [WFC3 ISR 2016-05](#).
- J. Mack, H. Olszewski, and N. Pirzkal, 2021, *WFC3/IR Filter-Dependent Sky Flats*, [WFC3 ISR 2021-01](#).
- J. Mack, W. Hack, M. Burger, R. L. White, V. Bajaj, R. J. Avila, G. S. Anand, and M. de la Pena, 2022, *Improved Absolute Astrometry for ACS and WFC3 Data Products*, [WFC3 ISR 2022-06](#).
- J. Maiz Apellaniz, 2007, *A Uniform Set of Optical/NIR Photometric Zero Points to be Used with CHORIZOS*, [ASPC, 364, 227](#).
- M. Marinelli, V. Bajaj, A. Calaminda, H. Khandrika, J. Mack, A. Pidgeon, C. Shanahan, and D. Som, 2022, *Monitoring WFC3/UVIS Photometric Sensitivity with Spatial Scans*, [WFC3 ISR 2022-04](#).
- M. Marinelli, V. Bajaj, A. Calamida, and J. Mack, 2024, *Time-Dependent Sensitivity of the WFC3/IR Detector (2009-2023)*, 2024, *Monitoring WFC3/UVIS Photometric Sensitivity with Spatial Scans*, [WFC3 ISR 2024-06](#).
- C. Martlin, R. O'Brien, I. Momcheva, and M. Gennaro, 2021, *Reducing Drift and Shift (DASH) Data Using wfc3dash and Accompanying Notebook Workflow*, [WFC3 ISR 2021-02](#) .
- J. Medina, S. Baggett, and The Quicklook Team, 2020, *WFC3 Quicklook Anomalies Database*, [WFC3 ISR 2020-02](#).
- J. Medina, J. Mack, and A. Calamida, 2022, *WFC3/UVIS Encircled Energy*, [WFC3 ISR 2022-02](#).
- P. R. McCullough, 2017, *Trajectories of Multi-lined Spatial Scans*, [WFC3 ISR 2017-06](#).
- I. Momcheva, P. G. van Dokkum, A. van der Wel, G. B. Brammer, J. MacKenty, E. J. Nelson, J. Leja, A. Muzzin, and M. Franx, 2017, *A New Method for Wide-Field Near-IR Imaging with the Hubble Space Telescope*, [PASP, 129, 5004](#).
- C. R., O'Dell, G. J. Ferland, W. J. Henney, and M. Peimbert, 2013, *Studies of NGC 6720 with Calibrated HST/WFC3 Emission-Line Filter Images. I. Structure and Evolution*, [Apj, 145, 19](#).
- B. Oke, 1964, *Photoelectric Spectrophotometry of Stars Suitable for Standards*, [Apj, 140, 689](#).
- A. Pagul et al., 2023, *The WFC3/UVIS G280 Grism Sky*, [WFC3 ISR 2023-06](#).
- N. Pirzkal, 2020, *Updated Calibration of the UVIS G280 Grism*, [WFC3 ISR 2020-09](#).

- N. Pirzkal and R. Ryan, 2020, *The dispersed infrared background in WFC3 G102 and G141 observations*, [WFC3 ISR 2020-04](#).
- N. Pirzkal et al., 2017, *Trace and Wavelength Calibrations of the UVIS G280 +1/-1 Grism Orders*, [WFC3 ISR 2017-20](#).
- N. Pirzkal, R. Ryan, *Variable He I emission in grism data*, [WFC3 ISR 2017-05](#).
- N. Pirzkal, R. Ryan, G. Brammer, *Trace and Wavelength Calibrations of the WFC3 G102 and G141 IR Grisms*, [WFC3 ISR 2016-15](#).
- A. G. Riess, 2010, *First On-orbit Measurements of the WFC3-IR Count-rate Non-Linearity*, [WFC3 ISR 2010-07](#).
- A. G. Riess, 2011, *An Independent Determination of WFC3-IR Zeropoints and Count Rate Non-Linearity from 2MASS Asterisms*, [WFC3 ISR 2011-15](#).
- A. G. Riess, G. Narayan, and A. Calamida, 2019, *Calibration of the WFC3-IR Count-rate Non-linearity, Sub-percent Accuracy for a Factor of a Million in Flux*, [WFC3 ISR 2019-01](#).
- B. Rothberg et al., 2011, *First Results from Contamination Monitoring with the WFC3 UVIS G280 Grism*, [WFC3 ISR 2011-18](#).
- R. Ryan, S. Casertano, and N. Pirzkal, 2018, *LINEAR: A Novel Algorithm for Reconstructing Slitless Spectroscopy from HST/WFC3*, [Publ. Astron. Soc. Pac. 130 034501](#).
- E. Sabbi, 2012, *Proposal 11923-UVIS Filter Wedge Check*, [WFC3 ISR 2012-01](#).
- E. Sabbi and A. Bellini, 2013, *UVIS PSF Spatial & Temporal Variations*, [WFC3 ISR 2013-11](#).
- K. Sahu, S. E. Deustua, and E. Sabbi, 2017, *WFC3/UVIS Photometric Transformations*, [WFC3 ISR 2014-16](#).
- C. E. Shanahan, C. M. Gosmeyer, and S. Baggett, 2017, *Update on the WFC3/UVIS Stability and Contamination Monitor*, [WFC3 ISR 2017-15](#).
- C. E. Shanahan, P. McCullough, and S. Baggett, 2017, *Photometric Repeatability of Scanned Imagery: UVIS*, [WFC3 ISR 2017-21](#).
- D. L. Shupe, et al., 2005, *The SIP Convention for Representing Distortion in FITS Image Headers*, ASP Conference Series, Vol 347, pg 491.
- D. V. Stark, N. Grogin, J. Ryon, R. Lucas, *Improved Identification of Satellite Trails in ACS/WFC Imaging Using a Modified Radon Transform*, [ACS ISR 2022-08](#)
- M. Sirianni et al., 2005, *The Photometric Performance and Calibration of the Hubble Space Telescope Advanced Camera for Surveys*, [PASP, 117, 1049](#).
- D. Som, V. Bajaj, J. Mack, and A. Calamida, 2021, *Photometric Repeatability and Sensitivity Evolution of WFC3/IR*, [WFC3 ISR 2021-05](#).
- D. Som, R. Bohlin, J. Mack, V. Bajaj, and A. Calamida, 2024, *Sensitivity Evolution of WFC3/IR Using Spatial Scanning Photometry and Grism Spectrophotometry*, [WFC3 ISR 2024-01](#).
- A. Suchkov and S. Baggett, *WFC3/UVIS Crosstalk and Crosstalk Correction*, [WFC3 ISR 2012-03](#).

A. Tokunaga & W. Vacca, 2006, *The Mauna Kea Observatories Near-Infrared Filter Set. III. Isophotal Wavelengths and Absolute Calibration*, [PASP](#), **117**, 421.

Chapter 10: WFC3 Spatial Scan Data

Chapter Contents

- [10.1 Analysis of Scanned Data](#)
- [10.2 IR Scanned Data](#)
- [10.3 UVIS Scanned Data](#)
- [10.4 References](#)

10.1 Analysis of Scanned Data

Spatial scans enable high-fidelity observations of bright sources such as host stars of exoplanets. In the case of photometry of bright, isolated sources, spatial scans have two key advantages which can lead to higher photometric precision over staring mode observations. First, the ability to collect millions of source photons per exposure without saturating by spreading the light over hundreds of pixels, yields much higher signal to noise ratios than conventional staring-mode observations. Second, illuminating hundreds of pixels averages out spatially-dependent sources of noise (such as flat-field errors) and enables sampling of different pixel phases along the direction of the scan. Spatial scans also enable high precision astrometry of bright sources.

The analysis of spatially scanned data obtained with HST WFC3, however, is different from that of the nominal, standard pipeline (**calwf3**, see [Chapter 3](#)) calibrated, staring-mode observations. Observers are sometimes required to customize, depending on the usage, a large fraction of their analysis of spatially scanned data.

In this chapter, we give a general overview of the custom analysis flows for scanned A) IR spectroscopy, especially of exoplanet transits or eclipses, and B) UVIS and IR imaging.

We recommend the observers use the following pipeline generated files as the starting point of their own custom analysis:


- IR: *_ima.fits
- UVIS: *_flt.fits or *_flc.fits

10.2 IR Scanned Data

- 10.2.1 IR Scanned Spectra
- 10.2.2 IR Scanned Photometry

The nominal calibration and cosmic-ray cleaning of IR observations by **calwf3** produces `_flt.fits` files containing average count rate information for each pixel. The count rate is derived by fitting a straight line to the up-the-ramp slope of accumulating charge within a pixel. This method, however, is inappropriate for spatially scanned data where the charge accumulation in any given pixel crossed by a source occurs within a very short period of time and thus appears as a step rather than a ramp (see [Figure 10.1](#)).

Figure 10.1: Schematic diagram of charge accumulation in a pixel sampled with the RAPID sequence for a source observed in staring mode (blue diamonds, in a ramp) or in spatially scanned mode (orange squares, in a step function).



As a consequence, analysis of IR scanned data has typically been done using the MULTIACCUM files: `_ima.fits`, which contain the individual readouts of the detector, adjusted for the appropriate reference pixel, dark and non-linearity corrections. As mentioned in [Section 3.3.9](#) and in the note in [Section 3.3.10](#), depending on the values of the `UNITCORR` and `FLATCORR` switches, the `ima` file will report the average count (or electron) rate up to a given read or the total accumulated counts (or electrons) up to that read. Given that WFC3/IR sampling sequences may have unevenly spaced readouts, converting from one to the other requires using the values of the readout timings (recorded in the `ima` image as either the `PIXVALUE` keyword of the `TIME` extension of each `imset`, or as the `SAMPTIME` keyword of the `SCI` extension of the same `imset`). This conversion also requires manually calculating differences of the N th readout and the $(N-1)$ th readout to form a set of difference images which can be subsequently analyzed with custom procedures. For the SPARS sample sequences specifically, the time interval between the 0th readout and the 1st readout is much shorter than the intervals between subsequent readouts; hence for SPARS data, the first interval can be discarded in subsequent analysis if desired for consistency. For example, with a GRISM512 subarray and SPARS10 readout, the first interval is 0.85 seconds, whereas subsequent intervals are 7.92 seconds (see [WFC3](#)

ISR 2006-06 or Phase II Proposal Instructions). Please note that, depending on the usage, specialized treatment of data and/or customization of analysis techniques may be required (e.g., see [Riess et al 2018a](#), [Riess et al 2018b](#)).


Starting with **calwf3** version 3.3 in early 2016, the `_flt.fits` files for scanned data are constructed without the up-the-ramp fit by setting `CRCORR` to `OMIT`. The result is a calibrated science data file (`flt.fits`) consisting of the first-minus-last science extension from the `ima` file (where the individual reads are stored in reverse time order), a more reasonable representation of the image than the up-the-ramp fit. Observers with calibrated scan data predating 2016 may re-retrieve their files from the MAST archive to obtain the improved `flt.fits` files. Since `CRCORR` is set to `OMIT`, these calibrated files may contain cosmic rays requiring removal via traditional routines e.g., stacking of dithered images, Laplacian cosmic ray identification (e.g., [L.A.Cosmic](#); [van Dokkum 2001](#)), and [Astrodrizzle](#).

We note that meta data related to spatial scans can be found in the `_spt.fits` files; please see [Table 10.1](#) for an example block of keywords associated with spatial scans. Additionally, with **calwf3** version 3.3 and later, all the UVIS and IR scan-related keywords formerly accessible only via the engineering file headers (`_spt.fits`), are now also present in the calibrated science data headers (e.g., `_flt.fits`).

Table 10.1: Spatial Scan Keywords in the `_flt.fits` files

Keywords	Description
<code>SCAN_TYP= 'C'</code>	C:bostrophidon; D:C with dwell; N:N/A
<code>SCAN_WID=0.000000000000E+00</code>	scan width (arcsec)
<code>ANG_SIDE=9.000000000000E+01</code>	angle between sides of parallelogram (deg)
<code>DWELL_LN=0</code>	dwell pts/line for scan pointing (1-99,0 if NA)
<code>DWELL_TM=0.000000000000E+00</code>	wait time (duration) at each dwell point (sec)
<code>SCAN_ANG=2.248278045654E+02</code>	position angle of scan line (deg)
<code>SCAN_RAT=5.000000074506E-02</code>	commanded rate of the line scan (arcsec/sec)
<code>NO_LINES= 1</code>	number of lines per scan (1-99,0 if NA)
<code>SCAN_LEN=2.969249963760E+00</code>	scan length (arcsec)
<code>SCAN_COR='V'</code>	scan coordinate frame of ref: celestial,vehicle

Figure 10.2:A spatially-scanned spectrum labeled with its 0th and +1st order light, and compared to a nominal staring-mode slitless spectrum of the same field (red outlined inset); reproduced from [WFC3 ISR 2012-08](#).



10.2.1 IR Scanned Spectra

MAST provides the user with calibrated images taken in scanned mode. However the analysis of spatially scanned WFC3/IR spectroscopy (see [Figure 10.2](#) for an example) involves ad-hoc post-processing steps that are left to the users. Some guidelines for the analysis of spatially-scanned IR spectroscopy can be found in [WFC3 ISR 2012-08](#), [WFC3 ISR 2019-12](#) & [WFC3 ISR 2019-13](#). Additionally, papers exist in the literature that illustrate custom data reduction procedures for the analysis of scanned data, e.g. [Zieba & Kreidberg \(2022\)](#) and references therein for exoplanet transits observed with spatial scanning. The WFC3 team does not endorse any specific paper; the interested reader should search the literature for such examples and use their own judgement in adopting the reduction strategy best suited for their purposes.

As a final note, observers should be aware that spatially scanned spectra have been found to exhibit sensitivity losses of $\sim 0.06\text{-}0.12\%$ /year ([WFC3 ISR 2024-01](#)).

10.2.2 IR Scanned Photometry

An example workflow to do photometry using IR scanned observations can be found in [WFC3 ISR 2021-05](#). This work uses total counts from the calibrated MULTIACCUM ima files as input and performs aperture photometry of stellar scans to study time-dependent sensitivity of the IR channel. However, as mentioned before, the details of the workflow depend on the usage of the data and several examples of customized analysis techniques can be found in the literature. Furthermore, the IR scanned photometry exhibits similar repeatability issues as seen in the staring mode described in [Section 7.11](#).

As is the case for scanned spectra, observers should be aware that sensitivity losses for spatially scanned IR photometry have been shown to be $\sim 0.06\text{-}0.16\%$ /year ([WFC3 ISR 2024-01](#)).

10.3 UVIS Scanned Data

[10.3.1 Astrometry using UVIS Scanned Observations](#)

[10.3.2 Photometry of Bright Targets using UVIS Scanned Observations](#)

For UVIS data, the calibrated `flt/flc` files of scanned data can be used directly for analysis, just as for staring-mode observations.

10.3.1 Astrometry using UVIS Scanned Observations

There are several existing papers in the literature that illustrate the use of UVIS scanning-mode data for astrometric measurements (e.g., [Riess et al. 2021](#), [Riess et al. 2014](#), [Casertano et al. 2016](#)). In these investigations, the authors have been able to achieve an astrometric accuracy of ~ 30 micro arcseconds in the measurements of trigonometric parallaxes, more than 10 times the precision achievable with staring mode observations. The interested reader is directed to such references for more details about the data analysis process.

10.3.2 Photometry of Bright Targets using UVIS Scanned Observations

HST program 14878 was a Cycle 24 calibration program intended to demonstrate the photometric repeatability of spatial scans of bright, isolated stars with WFC3/UVIS. Analysis of two identical visits showed that the photometric repeatability was $\sim 0.1\%$ r.m.s. ([WFC3 ISR 2017-21](#)), an improvement of more than a factor of 5 over the traditional results using staring mode. Sensitivity losses for UVIS spatial scans are relatively flat ($\sim 0.1\text{-}0.2\%$ /year), independent of wavelength on both UVIS CCDs, and show no evidence of contamination ([WFC3 ISR 2022-04](#)).

As an example of the analysis procedure, the steps performed on program 14878 are summarized below (the aforementioned report contains additional details).

- Calibrated (`*_flt.fits`) products, processed with the `calwf3` calibration pipeline, were retrieved from the MAST archive. Vertical scans and corner subarrays were used in this program to mitigate CTE losses.
- Cosmic rays (CRs) were removed using an algorithm originally developed for CR rejection in STIS CCD images. This algorithm identifies cosmic ray hits in the scanned images and replaces them with an interpolated value from the surrounding ‘good’ pixels.
- Images were sky subtracted. The sky region is defined as all pixels excluding a 10-pixel border around the perimeter of the subarray, and a conservatively large 400×75 pixel rectangular aperture around the source center. The pixel values in the sky region are sigma clipped (iteratively) and the mean of the remaining sky pixels is subtracted from the science array to remove the sky.
- The pixel area map was applied. This step can be important if the position of the scan on the subarray drifts significantly between visits.
- A rectangular aperture was centered on the scan using `photutils.RectangularAperture`. For higher precision, the overlap of the apertures with the science array was set to the ‘subpixel’ mode to allow for subpixel centering of the scans. Finally, `photutils.aperture_photometry` was used to sum up source counts.

10.4 References

- S. Casertano et al., 2016, *Parallax of Galactic Cepheids from Spatially Scanning the Wide Field Camera 3 on the Hubble Space Telescope: the Case of SS Canis Majoris*, [Apj](#), 825, 11.
- D. Deming et al., 2013, *Infrared Transmission Spectroscopy of the Exoplanets HD209458b and XO-1b Using the Wide Field Camera-3 on the Hubble Space Telescope*, [Apj](#), 774, 95.
- M. Marinelli et al., 2022, *Monitoring WFC3/UVIS photometric Sensitivity with Spatial Scans*, [WFC3 ISR 2022-04](#).
- P. McCullough, and J. MacKenty, 2012, *Considerations for using Spatial Scans with WFC3*, [WFC3 ISR 2012-08](#).
- L. Petro and T. Wheeler, 2006, *New IR Detector Sample Times*, [WFC3 ISR 2006-06](#).
- A. Riess, S. Casertano et al., 2014, *Parallax Beyond a Kiloparsec from Spatially Scanning the Wide Field Camera 3 on the Hubble Space Telescope*, [Apj](#), 785, 161.
- A. Riess, S. Casertano et al., 2018, *New Parallaxes of Galactic Cepheids from Spatially Scanning the Hubble Space Telescope: Implications for the Hubble Constant*, [Apj](#), 855, 136.
- A. Riess, S. Casertano et al., 2018, *Milky Way Cepheid Standards for Measuring Cosmic Distances and Application to Gaia DR2: Implications for the Hubble Constant*, [Apj](#), 861, 126.
- A. Riess, S. Casertano, et al., 2021, *Cosmic Distances Calibrated to 1% Precision with Gaia EDR3 Parallaxes and Hubble Space Telescope Photometry of 75 Milky Way Cepheids Confirm Tension with Λ CDM*, [Apj](#) 208, L6.
- C. E. Shanahan, P. McCullough, S. Baggett, *Photometric Repeatability of Scanned Imagery: UVIS*, [WFC3 ISR 2017-21](#).
- D. Som et al., 2024, *Sensitivity Evolution of WFC3/IR Using Spatial Scanning Photometry and Grism Spectrophotometry*, [WFC3 ISR 2024-01](#).
- D. Som, V. Bajaj, J. Mack, A. Calamida, *Photometric Repeatability and Sensitivity Evolution of WFC3/IR*, [WFC3 ISR 2021-05](#).
- K. Stevenson, and J. Fowler, 2019, *Analyzing Eight Years of Transiting Exoplanet Observations Using WFC3's Spatial Scan Monitor*, [WFC3 ISR 2019-12](#).
- K. Stevenson, and W. Eck, 2019, *Pre-Flashing WFC3/IR Time-Series, Spatial Scan Observations*, [WFC3 ISR 2019-13](#).
- P. van Dokkum, 2001, *L.A. Cosmic., Laplacian Cosmic Ray Identification*.
- S. Zieba, and L. Kreidberg, 2022, *PACMAN: A pipeline to reduce and analyze Hubble Wide Field Camera 3 IR Grism data*, [Journal of Open Source Software](#), 7(80), 4838.