



WFC3 Instrument Science Report 2009-13

CALWF3 Testing

E. Sabbi, H. Bushouse, S. Deustua, M. Dulude, B. Hilbert, J. Kalirai,
V. Kozhirina-Platais, A. Martel, A. Rajan
August 06, 2009

ABSTRACT

The IRAF/STSDAS program calwf3, which is used to calibrate WFC3 images, has been subjected to a rigorous testing campaign. The operation of each individual calibration step for WFC3 UVIS and IR images has been tested, as well as the ability of the overall calwf3 infrastructure to handle certain error conditions. Several identified problem areas have been corrected and included in the public release of calwf3 version 1.4.

Introduction

The standard IRAF/STSDAS program that is used to calibrate WFC3 images is calwf3. Detailed descriptions of each calibration and data reduction step performed by calwf3 are contained in the WFC3 Data Handbook (2009, available at http://www.stsci.edu/hst/wfc3/documents/handbooks/currentDHB/WFC3_longdhbcover.html). As with other HST instrument calibration programs, keyword values in the headers of raw images to be processed are used to identify the steps to be performed and the reference images and tables to be used. The calwf3 calibration steps are segregated into two main branches: one for processing WFC3 UVIS channel images and the other for IR images. The UVIS calibration steps are performed by the sub-tasks wf3ccd and wf32d, while the IR processing is performed by wf3ir. A third major portion of calwf3 is the wf3rej task, which can be used to combine multiple images taken with either the UVIS or IR channel. This procedure will stack images taken as part of a CR-SPLIT or REPEAT-OBS grouping, but will not combine images taken at different pointings. The

`calwf3` task itself can be thought of as a higher-level procedure that is used to execute the various calibration steps contained within the lower-level tasks `wf3ccd`, `wf32d`, `wf3ir`, and `wf3rej`.

Between November 2008 and February 2009 we performed extensive testing of each calibration step within `calwf3`, in order to verify that each step was producing scientifically correct results. As input to the tests, we used data that had been acquired during the third WFC3 thermal-vacuum (TV3) testing campaign, so that the properties of the input images are representative of flight-like operations.

The tests were organized into four broad categories.

1. Infrastructure: These tests were designed to check that `calwf3` handles input files properly and to verify that, when used in an incorrect way, error messages are correct and can guide the user in finding the problem.
2. UVIS calibration: These tests were designed to check that the basic image reductions, such as bias and dark subtraction and flatfielding are correctly performed on UVIS CCD data. We performed tests on full-frame unbinned, full-frame binned, and sub-array images in order to verify that all of these modes are handled correctly. We selected sub-array images that included some portion of the serial physical overscan within the sub-array, as well as those that don't contain any overscan, and also come from at least two different and opposing amplifiers (e.g. A and D, or B and C). Sub-array testing included the verification that any calibration step involving the application of a reference image extracted and applied the correct sub-array portion of the full-frame reference image.
3. IR calibration: These tests checked all of the IR image calibration steps, such as dark subtraction, flatfielding, and up-the-ramp slope fitting. We again performed tests using all possible imaging modes, including full-frame and sub-array images.
4. Image combination: The `wf3rej` task operations were checked in detail both on UVIS and IR images to verify that CR-SPLIT and REPEAT-OBS exposures are properly combined and that Cosmic Ray (CR) hits are identified and rejected.

Test Results

The following sections describe each test that was performed and the results, including any bug fixes, changes, or upgrades that were made to `calwf3` in response to the tests. All of the tests were performed using `calwf3` version 1.1, which was released in October 2008. Most of the changes that resulted from these tests were included in `calwf3` v1.2 and a few were incorporated in subsequent versions up through v1.4, which was released in April 2009.

In addition to the formal test plan activities, `calwf3` operation was subjected to detailed scrutiny during the same time period as it was being used extensively to process and create calibration reference files from the TV3 data. This use also led to the

discovery of some software bugs and desired upgrades, which have now been incorporated into the latest version of `calwf3`.

Infrastructure tests

No command-line arguments: When `calwf3` is run without any command-line arguments it correctly prompts for the name of an input file.

Single raw FITS file as input: `Calwf3` properly accepts as input a single `_raw.fits` file name and processes it, producing appropriately named output calibrated image file(s) and a trailer (`.tra`) file containing the processing log.

Single association table as input: When running `calwf3` with the name of an `_asn.fits` file as input, the raw images listed in the `asn` table are correctly processed. When the `asn` table contains the names of two UVIS CR-SPLIT images, for example, the individual images are calibrated and combined, producing an output `_crj.fits` file. Furthermore, when the keyword switch `EXPSCORR=PERFORM`, fully calibrated `_flt.fits` files are produced for each of the input images.

Non-existent input file: Running `calwf3` with the name of a non-existent file as input resulted in an error message, saying that the file was not found, but the process then eventually crashed before finishing the shutdown process. The bug that caused the crash was associated with the trailer (`.tra`) file management and has been corrected in `calwf3` versions 1.2 and later.

Existing flt file given as input: `Calwf3` is designed to operate on raw images only. When an `_flt.fits` file was used as input, `calwf3` correctly produced an error message saying that the file could not be opened, but then crashed before completing a shutdown, as in the previous test. The same bug was responsible, which has been fixed in `calwf3` v1.2.

Multiple raw files given as input: `Calwf3` is designed to process multiple images in one execution, but only when those images are members of an association listed in an association (`_asn`) table. It is not designed to accept multiple images listed on the command line. When given a list of input images on the command line, `calwf3` responded with “cannot open image” error messages, which is appropriate, but then crashed before finishing the shutdown process. This is due to the same bug discovered in the previous two tests, which has been fixed in `calwf3` v1.2 and later.

Wildcard used in input file argument: Running "calwf3 *raw.fits" resulted in a list of "cannot open image" error messages, ending with a program crash. The long list of error messages is due to calwf3 trying to open different possible incarnations of the input file, e.g. _raw, _asn, etc. The crash was caused by the same bug responsible for the previous three results, which has been fixed.

Use of task parameter "savetmp": When task parameter "savetmp=yes" calwf3 correctly produces and saves the intermediate results file _blv_tmp.fits file. This was tested using single _raw.fits files as input, as well as association tables (_asn.fits). When "savetmp=no", the intermediate files are deleted at the end of the processing run.

Non-existent reference file: If a designated reference file does not exist, calwf3 correctly shuts down with a "file not found" error message. The only case involving inaccessible reference files that could use enhancement is when the "mtab" environment variable is not defined, which gives the directory path for the Synphot tables used in the PHOTCORR step. In this case, calwf3 only issues a generic message saying that the data could not be processed, but without any specifics as to the real problem. The error messaging should be improved in a future version.

Non-existent reference file for an association: When processing an association, calwf3 uses only the reference files designated in the header of the _raw.fits image listed first in the asn table. Therefore if any raw file other than the first one points to a non-existent ref file, calwf3 proceeds without error. If the first raw file points to a missing reference file, calwf3 shuts down with a "file not found" error.

Incorrect reference files: We tested how calwf3 handles a single raw file that has inappropriate reference files listed in the header. For UVIS data we found that when a bias or flatfield image is used in place of a dark image calwf3 ran with no errors or warnings. When a flat or dark is used instead of a bias, calwf3 produced a "biascorr size mismatch" error and stops. When a dark or bias is used instead of a flat, calwf3 crashed with a floating-point overflow. For IR data we found that if reference files are used inappropriately, calwf3 produced an error message and stopped. As a result of these tests, version 1.2 of calwf3 was enhanced to perform checking on the value of the FILETYPE keyword in all reference files to make sure they are the appropriate type for a given operation, and to check that other properties of the reference file, such as the CCD gain setting or the filter value for a flatfield, also match the science data being processed.

Mismatch in reference image binning: When calwf3 was run with one or more UVIS reference images that have a different binning than the science image it correctly produced an error message saying that there is a mismatch and stopped.

Mismatch in flatfield FILTER value: When processing an IR image with a flat file having a FILTER value that is different from that of the science image `calwf3` correctly reported the error and stopped. Processing a UVIS image with a flat file having a different FILTER value, however, did not result in any warnings or errors and processing proceeded to completion. This bug was fixed in `calwf3` version 1.2, as part of the larger set of reference file checks mentioned above.

Mismatch in IR dark SAMP_SEQ value: When processing an IR image with a dark file that has a different SAMP_SEQ value, `calwf3` correctly reported the error and stopped.

Vital keywords missing: When the DETECTOR or CCDAMP keywords were missing from the input image header, `calwf3` issued an error and stopped. On the UVIS side, however, if either CCDGAIN or FILTER were missing, it proceeded with processing with no messages, because `calwf3` is setup to use a default value of 1.5 (UVIS) or 2.5 (IR) for the gain if the CCDGAIN keyword is missing and, at least on the UVIS side, is similarly setup to use a default value of "" (null string) if the FILTER keyword is missing. Version 1.2 of `calwf3` was upgraded to check for the value of the FILTER keyword for both UVIS and IR images.

Corrupted input file: When `calwf3` was run with a single `_raw.fits` file as input that was missing one or more image extensions the error was correctly reported and processing stopped.

UVIS calibration tests

NOISCORR: We confirmed that the ERR values computed by `calwf3` are equal to $\sqrt{(\text{DN}-\text{bias})/\text{gain} + (\text{readnoise}/\text{gain})^2}$

where bias, gain, and readnoise are values from the CCDTAB (and as reported in the image headers) and DN is the image value (in units of DN). This was verified for all 4 amps, confirming that the proper unique values are used for each amp. This was checked for full-frame unbinned, subarray, and binned UVIS images.

DQICORR: We created a custom bad pixel table (BPIXTAB) with some flags set in all 4 UVIS image quadrants. We found that all pixels are correctly flagged in the output DQ arrays. When there is no flag within the boundary of the image the output DQ image array is null. Testing included full-frame unbinned and binned, as well as sub-array, UVIS images.

BLEVCORR: We processed full-frame unbinned and binned UVIS images that contained an artificial uniform bias level in the SCI extensions, along with an artificial border of pixels along the edges of the active image area. `calwf3` correctly measured the bias level, leaving a residual of exactly zero in the subtracted images. The artificial border of pixels appeared in the first rows and columns of the overscan-trimmed images, indicating proper trimming parameters.

BIASCORR: We compared the values of individual pixels before and after the bias image subtraction, using full-frame unbinned and sub-array UVIS images, which showed that the correct pixel value from the BIASFILE is subtracted from the science image in all cases, except for sub-arrays contained within the B and D amplifier quadrants. For the quadrant B and D sub-arrays, the presence of the columns of serial virtual overscan in the middle of the full-frame bias image were not taken into account when computing the location of the sub-array to be extracted from the BIASFILE. This was fixed in `calwf3` version 1.3. We also tested the propagation of DQ flags from the BIASFILE into the calibrated output image and found that all flag values are propagated. The only exception to this is when multiple bias-subtracted images are subsequently combined using the `wf3rej` routine, in which case any DQ=8192 (CR hit) flags contained in the BIASFILE are not propagated through to the output of `wf3rej`, because all input CR flags are intentionally cleared upon entry into the `wf3rej` routine.

CRCORR: When processing a UVIS CR-SPLIT association, `calwf3` calls the `wf3rej` routine to combine the bias-subtracted science images and correctly populates the EXPTIME and TEXPTIME values in the combined output image. The correct identification and removal of CRs was verified using combined bias, dark, and flat field images. The identification thresholds may need some optimization once on-orbit images are available. Another open issue is the treatment of negative outliers; they are rejected from the combined image, but flagged in the input image DQ arrays as a CR hit even though a negative outlier is not actually a CR. We may consider using a different flag value for negative outliers.

The correct summation of unrejected input values into the output image SCI and ERR arrays was also checked by comparing the `calwf3` results with by-hand summation of the known input values. Initially this seemed to give mixed results, with some exact matches and others that were slightly different. Subsequent investigation revealed that this was due to our incomplete understanding of the exact way in which `wf3rej` computes the summed output value for a pixel that has one or more input values rejected. The output pixel value is computed as the sum of the unrejected sky-subtracted input values, plus the sum of the sky values from *all* input images (both rejected and

unrejected). Using this approach for our by-hand calculations then yielded an exact match in all cases.

Finally, during testing it was noticed that running `calwf3` with an association (asn) table containing only one input UVIS image resulted in `calwf3` correctly noticing and reporting that there was only one image and therefore the CRCORR step would be omitted, but it then subsequently stopped with an error message saying that "CRCORR and/or RPTCORR values not consistent with ASN table". Other tests showed that CRCORR always gets performed if the input to `calwf3` is an asn table, even when CRCORR is set to OMIT in all the input raw image headers, suggesting that the presence of the asn table is overriding the CRCORR settings. No changes have yet been made to `calwf3` to address either of these issues, but may be in the future.

EXPSCORR: We verified that, as expected, running `calwf3` on an association with EXPSCORR=OMIT produces only a single output crj file, while when EXPSCORR=PERFORM it also produces a calibrated flt file for each input UVIS image.

DARKCORR: We verified that the DARKFILE is scaled properly by the gain and exposure time of the science image and then subtracted correctly from the science image. We also verified that DARKFILE DQ values are propagated into the DQ array of the output science image. These tests were carried out on unbinned and binned full-frame images and on subarray UVIS images.

FLATCORR: We verified that flat field values are correctly applied to the science image and DQ values in the FLATFILE are included in the DQ array of the output science image (whether flt or crj). The use of both a PFLT (pixel-to-pixel flat) and DFLT (delta flat) results in the correct combination of the two flats and subsequent application to the science image. This was verified for both full-frame and subarray UVIS images. Finally, for both full frames and sub-arrays, the output SCI and ERR array values are correctly scaled by the gain for each amplifier quadrant, which converts the science data from units of DN's to electrons. The BUNIT header keyword is also correctly updated to reflect this change.

We discovered, however, that the application of an LFLT (low-order flat) caused `calwf3` to crash in the FLATCORR step, due to bugs in the methods used to interpolate a binned LFLT to match the pixel scale of the science image. These bugs were fixed in `calwf3` version 1.4, but the interpolation routines still need further work in order to produce the best results. The use of a binned LFLT is therefore not currently allowed, with an appropriate error message being produced. It is possible, however, to apply an LFLT that has the same binning as a science image, in which case no interpolation of the LFLT is needed.

PHOTCORR: After running `calwf3`, UVIS output files have all photometry-related keywords populated and the values agree with results computed directly with the Synphot bandpar task.

DRIZCORR: DRIZCORR currently produces an empty, placeholder drz file for both UVIS and IR processing. The UVIS drz file has the correct number of imsets (1 for each CCD chip), as does the IR drz file (always just one imset). The imsets consist of SCI, ERR, and DQ extensions for both UVIS and IR images. All extensions, including SCI, have null data arrays. This step will be superseded by MultiDrizzle processing in the OPUS science pipeline.

IR calibration tests

DQICORR: Full-frame and sub-array IR images were processed through `calwf3` with only the DQICORR step turned on. The resulting products had all of the pixels listed in the BPIXTAB flagged with the proper values in all of the DQ extensions of both the ima and flt files. The flagged pixels in the sub-array images were an exact match to the equivalent sub-array portion of the full-frame image.

BLEVCORR: Full-frame and sub-array IR images were processed through `calwf3` with only the BLEVCORR turned on. We compared the mean reference pixel values that were computed by `calwf3` – as stored in the MEANBLEV keyword associated with each readout in the resulting ima file – with our own computations and found an exact match for all reads of both types of images. Furthermore, manual subtraction of the computed reference pixel values from each readout in the raw files resulted in images that were an exact match to the `calwf3` results in the ima files.

ZOFFCORR: Full-frame and sub-array IR images were processed through `calwf3` with only the ZOFFCORR step turned on. The resulting science extensions in the ima files were compared with manual subtractions of the zero-read image from each readout in the raw file, which produced an exact match.

NOISCORR: An IR point-source exposure was processed through `calwf3` with only the DQICORR, ZSIGCORR, BLEVCORR, and ZOFFCORR steps turned on. This yielded an output ima file with the ERR array values initialized to include contributions from only the readout noise and Poisson source noise in each pixel. The ERR values, in units of DN, in the ima file were an exact match to values computed manually, using the formula

$$\text{ERR} = \sqrt{\text{readnoise}^2 + \text{gain} * \text{SCI}} / \text{gain}$$

where the readnoise and gain are from the CCDTAB and are unique to each amplifier quadrant. We verified that the appropriate quadrant-specific values were applied to the science image pixels.

We also verified that when DARKCORR is turned on, the values in the ERR arrays of the output ima file are the combined error values of the readout noise and Poisson noise from the earlier run (ERR_run1) combined with the ERR values from the DARKFILE (ERR_drk), according to the equation

$$\text{ERR} = \text{sqrt}(\text{ERR_run1}^{**2} + \text{ERR_drk}^{**2}).$$

DARKCORR: Comparison of `calwf3` results with manual dark subtraction showed a perfect match, indicating that the correct dark file readouts were subtracted from each of the science IR image readouts. DQ values from the dark file were also correctly propagated into the science image DQ arrays in the ima files. These tests used science IR images with varying values of NSAMP, as well as full-frame and sub-arrays.

NLINCORR: NLINCORR was verified to be working properly for a full-frame NSAMP=14 IR image and a 512x512 sub-array IR image with NSAMP=16. This was done by comparing corrected pixel values and DQ flags with independent calculations using the same formulae.

PHOTCORR: We verified that IR output files have all photometry-related keywords populated, with values that match those computed independently using the Synphot bandpar task.

UNITCORR: We checked the output ima file SCI and ERR extensions for a full-frame, NSAMP=9, SAMP_SEQ=STEP25 IR image and for a Sq512Sub, NSAMP=16, SAMP_SEQ=RAPID IR image. For each sample, SCI and ERR values had been correctly divided by the appropriate sample times, converting the SCI and ERR values from units of counts to countrate. The BUNIT keywords in the output file headers were also updated correctly to reflect this change in units.

CRCORR: We compared the “up the ramp” line fitting results of `calwf3` with independent computations. The independent computations did not include the detailed weighting schemes used within the `calwf3` CRCORR methods and therefore do not provide an exact match in results. The differences were quite small, however, which gives us confidence that the `calwf3` routines are working properly. For pixels with high signal levels the `calwf3` line fitting agrees with the manual methods to within 0.2%, with somewhat larger differences for the low-signal regime.

The identification and rejection of CR hits was verified to be working properly, expect for the following behavior. When a CR was identified in read N of a ramp, a CR

flag was placed in the DQ extension of the ima file for read N, but not for any of the subsequent reads that are also affected by the CR. This was fixed in `calwf3` version 1.2 so that all reads following a hit are also flagged. Testing of the CR rejection threshold was also performed, which showed that when using the default 4-sigma threshold any CR that would increase the slope of the best-fit line by more than 2.25-2.5% above the nominal value is identified and rejected. This test was carried out by inserting CRs of different strengths into the middle of the IR ramps.

Finally, we also tested the handling of the SAMP and TIME extension values in the output flt image when CRs are rejected from a pixel. We found some errors that were leading to incorrect SAMP and/or TIME values when one or more samples were rejected. In some cases the SAMP and/or TIME values were too low by 1 sample amount, due to the zero-read not being accounted for, and in other cases they were too high due to double counting of samples that contained a CR. These bugs were fixed in `calwf3` version 1.3.

FLATCORR: Manually applying the flat field and gain values to an input IR full-frame ramp showed an exact match to the results from `calwf3` when FLATCORR is performed. Testing FLATCORR on IR sub-array exposures revealed a bug in the procedure that applies the gain correction to sub-arrays. This bug was fixed in `calwf3` version 1.2.

RPTCORR: We verified that when `calwf3` is executed on an association of IR REPEAT-OBS exposures, the `wf3rej` routine is called to combine the individually-calibrated flt files, with the proper list of files to be combined, and a combined crj file is produced as a result.

Summary

Our investigation allowed us to identify several previously unknown bugs in `calwf3`. We prioritized the necessary software upgrades, giving highest weight to those bugs that would have affected the on-going construction of reference files that will be used during SMOV and Cycle 17. All known issues have been resolved as of the release of `calwf3` version 1.4, with the exception of further work on the interpolation of binned UVIS low-order flat field (LFLT) images.

Acknowledgment

We thank Sylvia Baggett for useful suggestions and for reporting odd `calwf3` behavior that she discovered independently of the formal test plan, and Nor Pirzkal for reviewing this ISR.