

Pipeline Calibrations of ACS Data

Max Mutchler, Marco Sirianni¹, & Ray A. Lucas

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218
E-mail: mutchler@stsci.edu, sirianni@stsci.edu, lucas@stsci.edu

Abstract. Our strategies for collecting ACS calibration data, and converting it into reference files for use in the calibration pipeline, have been continually evolving since the instrument was installed during Hubble Servicing Mission 3B in March of 2002. We provide an historical overview of basic ACS pipeline calibrations, including the impact that some of the more recent changes have on downstream data processing (e.g. drizzling). This paper emphasizes bias and dark calibrations for the ACS CCDs, and it represents a significant update to previous documentation of this subject (Mutchler et al. 2004). We describe the expected detector degradations that these calibrations are designed to track and correct, and also some unexpected anomalies we've encountered and the steps we've taken to ameliorate them.

1. Introduction: Philosophy and Practices

We have been continually improving the quality and degree of automation in producing and delivering calibration reference files for use in the ACS pipeline (**CALACS**, see Pavlovsky et al. 2005). The detector characteristics we are striving to calibrate are described in great detail in Sirianni et al. (2004). The emphasis of this paper is on bias and dark calibrations for the Wide Field Channel (WFC) and High Resolution Channel (HRC) of the Advanced Camera for Surveys (ACS). In October 2004, we implemented several improvements which we document here.

Our goal has always been to strike the best balance between quality and timeliness. We strive to produce the best calibrations possible within 2-3 weeks following any given ACS observation. For this reason, ACS users should retrieve (or re-retrieve) their data via on-the-fly-reprocessing (OTFR) several weeks after they occur, mainly to ensure that the best superbias and superdark reference files have been applied.

In August 2004, we migrated our reference file production system to the native pipeline environment (SunFire), with more streamlined production and overall quality control (Lucas et al. 2006, these Proceedings). The delivery of reference files to the Calibration Database System (CDBS) for use in the pipeline has also become more consolidated and efficient (Diaz-Miller 2005). We also deliver our reference files to the Multimission Archive at STScI (MAST) and make them available for downloading from our 'jref' directory (see Section 8).

2. Bias features and calibration

For the WFC and HRC, the bias *level* is measured from each frame's overscan region, and subtracted from each amplifier quadrant independently. Bias *features* are subtracted by the 'superbias' reference file, which is identified in the BIASFILE keyword in image headers.

¹Affiliated to the Space Telescope Division, European Space Agency

Before October 2004, we were obtaining one bias frame every day (for both WFC and HRC), and producing a WFC superbias every week, and an HRC superbias every two weeks. Bias (mainly bad CCD columns) structure was identified (with data quality flag 128) only in a bad pixel table (BPIXTAB) produced in July 2002.

Since October 2004, we have been obtaining one bias frame every *other* day, and producing superbias nominally every two weeks, for both WFC and HRC. The superbias are simple cleaned combinations of the 8 bias frames obtained during each bi-week. Bias structure is now identified with flag 128 in each superbias data quality [DQ] array, which propagates to the [DQ] array of science data. Therefore, the growth of bad CCD columns is now being tracked and flagged much better, typically at two week intervals. Figure 2 illustrates the WFC bias structure.

3. Dark features and calibration

We conduct a ‘monthly’ (roughly every 26 days) CCD annealing which restores many hot pixels to the nominal dark current, although the population of non-annealing hot pixels grows continually (Sirianni, these Proceedings). The annealing cycle also sets the cadence for our reference file production and delivery: we make a batch of superbias and superdark reference files for every half of one anneal period (roughly every 2 weeks).

Before October 2004, we were obtaining four 1000-second dark frames every day, and producing a superdark reference file for every day. We were flagging only hot pixels in superdark data quality or [DQ] arrays, using flag 16.

Since October 2004, we began collecting 4 dark frames every *other* day, and therefore we now produce a superdark reference file for every other day. We reduced the number of bias and dark frames we obtain to lower the profile of this (the largest) ACS calibration program, with what we felt was minimal impact on the quality of the resulting calibrations. In addition to hot pixels, we also began flagging several other dark features: warm pixels, CTE tails of hot pixels, and saturated pixels, which we describe in detail in the next section.

Our dark reference files are actually a hybrid combination of two types of superdark. First, we make a 2-week ‘basedark’, which is a high signal-to-noise combination of 32 dark frames. Then we also make a 4-frame ‘daydark’ for each day in the bi-week, which gives the best snapshot of the warm and hot pixel populations on that date. The reference superdark is a copy of the basedark, with the warm and hot pixels inserted from the relevant daydark. Figure 3 illustrates the more subtle WFC dark structures. The histogram in Figure 4 illustrate how the reference files get the majority of their pixels (normal dark current pixels, which exhibit only Poisson noise) from the ‘basedark’, and get their warm and hot pixels from the ‘daydark’.

Due to a small bias level variation (described in Section 5), we also began equalizing any residual bias level variations seen in the superdark, so that at least the superdark will not propagate this problem to science data.

4. Data quality flagging and monitoring

In Table 1, we define all the flags used in the data quality arrays of ACS data (i.e. the [DQ] image extension) to identify good, bad, and corrected pixels (see Figures 5 and 6). Some of these flags emanate from the ‘permanent’ bad pixel tables (identified in the BPIXTAB image header keyword), although as of October 2004 we began using the bad pixel tables much less. Currently, most of these flags emanate from the generic conversion of the science data, or they propagate from the data quality arrays of the various reference files used to calibrate and combine data in the pipeline (CALACS and MultiDrizzle).

We have also re-defined some flags for new uses – warm pixels (flag 64) and the CTE tails of hot pixel (flag 32) – although in the pipeline we set parameters (bits=96) to ignore

these new flags. So for now, they are provided only as optional data processing leverage for stand-alone processing. Note that in the stand-alone environment, the MultiDrizzle default is `bits=0`, so the user would need to manually set `bits=96` to similarly ignore these new flags (i.e., to include these pixels) during image combination.

5. Unexpected anomalies

Most of the detector calibrations were begun during pre-flight thermal vacuum testing, so their characteristics were reasonably well understood before launch. But inflight experience has revealed some unexpected anomalies. We briefly mention some notable anomalies here, and the steps we have taken to ameliorate them.

1. Random bias level variations can occur between the ‘real’ pixels in an image and their own overscans (Sirianni et al. 2003), where the bias level is measured, to be subtracted from the rest of the image. There is nothing we can do to correct this when it occurs in science exposures. But as of October 2004, we have been correcting this effect when seen in our superdarks (i.e. when it occurs in any of our input dark frames), so that at least our superdarks do not impose this effect on science data.
2. We have sometimes seen faint diffuse scattered light in dark frames, which would survive cosmic ray rejection and appear in the corresponding superdarks, and propagate to science data. After this problem was discovered, reference files which excluded the affected dark frames were delivered. We noticed that this problem seemed to occur following CCD annealings, when we had been leaving the filter wheels open (with CLEAR,CLEAR filters). We have not seen this problem after modifying our annealing program by rotating crossed filters into the optical path following each anneal.
3. In the first science data taken with ACS in April 2002, very faint negative imprints of bright objects were evident as mirror-image ‘ghosts’ replicated in each amplifier quadrant (see Figure 1). This is an electronic effect, caused by amplifier crosstalk, which shouldn’t have much impact on science data. Nonetheless, we recently changed the default WFC gain (from `gain=1` to `gain=2`) which minimizes this effect.

6. Pipeline drizzling

Although the most significant recent change to the ACS pipeline was the addition of **MultiDrizzle** for combining and cleaning associated datasets, this topic is being covered more completely elsewhere (Koekemoer 2006, these Proceedings), so we only briefly mention here some details of its implementation in the pipeline environment.

1. Pre-defined drizzling parameter sets (determined mainly by the number of images being combined) are applied to associated datasets via a FITS table, which is identified in image headers by the MDRIZTAB keyword. While this table define reasonably good default parameters for combining images in the pipeline, the user can experiment considerably with these parameters in the stand-alone environment, where trial-and-error iterations will often lead to more optimal parameters for a specific dataset.
2. As of mid-2005, pointing patterns defined with POS TARGs are now recognized as associated datasets, along with patterns defined with ‘convenience’ pattern forms (Mutchler & Cox 2001). Only associated datasets are automatically combined by MultiDrizzle in the pipeline. Large datasets involving multiple visits/epochs or large mosaics are not automatically or fully associated, and therefore must be combined by the user in the stand-alone environment.

3. In early 2005, filter-dependent distortion solutions (4th-order polynomials) were introduced, along with residual distortion correction images. These are identified in image headers in the IDCTAB and DGEOFILE keywords, respectively.

7. Max's wish list

The following items are some remaining issues, and thoughts on how they could be addressed. They do not necessarily represent issues for which the ACS group at STScI has given high priority. Rather, they are included here for consideration and discussion.

1. We now have two eras of data quality flagging. It would be desirable to somehow make the new flagging scheme retroactive to ACS launch (March 2002). Creating all new superbias and superdarks for 2002-2004 and delivering them is impractical at this point. Perhaps monthly bad pixel tables could be created for 2002-2004, which reflect the new flagging scheme (contain warm pixels, CTE tails, saturated pixels). With such bad pixel tables in place in the pipeline, subsequent data retrievals (via OTFR) would have flagging much more like we have had since October 2004.
2. AS CTE worsens, it will become increasingly desirable to reject the growing CTE tails of bright artifacts such as cosmic rays and hot pixels. More complete and unique flagging of hot pixel CTE tails in superdark reference files would help with the latter. As of October 2004, the first trailing pixel is flagged uniquely (data quality flag 32), while any subsequent pixels in the CTE tail are flagged along with (i.e. not distinguished from) the warm pixels (flag 64). Note that MultiDrizzle allows the user to 'grow' the rejection of artifacts, preferentially along the trailing detector y-axis, so that their CTE tails are also rejected. However, this feature is currently not utilized in the pipeline.
3. At present, only a small fraction of ACS data is associated, so most ACS datasets are unable to take advantage of MultiDrizzle being in the pipeline. For existing archival data, logical and complete association tables could be generated offline, and ingested into the archive, such that subsequent data retrievals (via OTFR) would produce clean drizzle-combined products.
4. Add image registration (e.g. tweakshifts) to the pipeline, to allow the combination of data taken with different guide stars (e.g. data from different visits, epochs, and/or observing programs). Another way to approach this problem is to measure shifts offline, and in conjunction with the above item, ingest association tabs with the shifts and rotations embedded in them.
5. Improved handling of moving target (planetary) observations in the pipeline. Associated moving target datasets are improperly combined by MultiDrizzle, which uses the World Coordinate System (WCS) to register images. Hubble's moving target tracking is generally accurate to within a fraction of a pixel, so a combination which simply ignores the WCS would produce good results. Also, since moving targets often exhibit additional complex motions (e.g. planets rotate between exposures), single-image drizzled output is always desirable, for associated or unassociated datasets (i.e. single exposures which have only been distortion corrected).

8. More information

More information about the basic ACS pipeline calibrations described in this paper can be found on the web at: http://www.stsci.edu/hst/acs/analysis/reference_files.

ACS reference files in FITS format can be directly downloaded from:
<ftp://ftp.stsci.edu/cdb/cdb7/jref/>.




Figure 1: Amplifier crosstalk evident in the first ACS science data. The signal from the Tadpole Galaxy (in the amplifier A quadrant) generates faint negative ‘ghosts’ in the other amplifier quadrants.

The ACS Data Handbook and Instrument Science Reports referenced in this document are available online at: <http://www.stsci.edu/hst/acs/documents>.

Acknowledgments. We thank the following people for essential contributions to the overall success of ACS pipeline calibrations and drizzling: Warren Hack, Anton Koekemoer, Chris Hanley, Mike Swam, and Rossy Diaz-Miller.

References

- Diaz-Miller, R. I., 2005, *Technical Instrument Report* CDBS-2005-01 (Baltimore: STScI), available through <http://www.stsci.edu/hst/observatory/cdb>s
- Lucas, R. A., Swam, M., Mutchler, M., & Sirianni, M., 2006, *The 2005 HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 61
- Koekemoer A., et al., 2006, *The 2005 HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 423
- Mutchler, M., & Cox C., 2001, *Instrument Science Report* ACS 2001-07 (Baltimore: STScI), available through <http://www.stsci.edu/hst/acs>
- Mutchler, M., Sirianni, M., van Orsow, D., & Riess, A., 2004, *Instrument Science Report* ACS 2004-07, (Baltimore: STScI)
- Pavlovsky, C., et al., 2005, *ACS Data Handbook*, Version 4.0 (Baltimore: STScI)
- Sirianni, M., et al., 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 82
- Sirianni, M., Mutchler, M., Clampin, M., Ford, H., Illingworth, G., Hartig, G., van Orsow, D., & Wheeler, T., 2004, *Optical and Infrared Detectors for Astronomy*, Proc. SPIE, Vol. 5499
- Sirianni, M., et al., 2006, *The 2005 HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 45

Table 1: ACS data quality flag definitions

Flag	Definition
0	Good pixels
1	Reed-Solomon decoding error; e.g. data lost during compression .
2	Data replaced by fill value; e.g. neighboring cosmic ray contaminated pixels.
4	Bad detector pixel , or beyond aperture. In the HRC BPIXTAB, this identifies a small detector defect in the upper right corner.
8	Pixels masked by aperture feature, e.g. the HRC occulting finger.
16	Hot pixels with dark current greater than $0.08 \text{ e}^-/\text{sec}$; flagged in superdark data quality [DQ] arrays.
32 ^a	CTE tails of hot pixels; flagged in superdark [DQ] arrays. For now, we only flag the first pixel trailing each hot pixel, but this flagging may become more sophisticated and complete as CTE worsens. Note that these CTE tails are flagged more completely as warm pixels (flag 64).
64 ^b	Warm pixels with dark current between 0.02 and $0.08 \text{ e}^-/\text{sec}$; flagged in superdark [DQ] arrays.
128	Bias structure (mostly bad columns). Since 8 Oct 2004, bias structure has been flagged in the bi-weekly superbias [DQ] arrays. Before 8 Oct 2004, bias structure was only flagged in an increasingly outdated bad pixel table (BPIXTAB), which was created in July 2002.
256	Both full-well (useable at higher gain setting) and A-to-D (never useable) saturated pixels are flagged by ATODCORR, based on the CCDTAB. But A-to-D saturation is also flagged 2048, so it can be distinguished from full-well saturation. Also, since 8 Oct 2004, full-well saturated pixels in superdark are flagged in their corresponding [DQ] arrays (note that they are also flagged as hot pixels with flag 16). Before 8 Oct 2004, saturated pixels in the superdark were flagged only in a BPIXTAB created in July 2002.
512	Bad pixel in reference file. Used in the flatfield [DQ] arrays to indicate a portion of the flat which is not defined or not calibrated with the same accuracy as the other regions, often around the detector edges. Used for F892N and WFC polarizer observations, where the filter only subtends a portion of the chip. Used to identify dust mote replacement patches.
1024	Charge traps ; flagged in the bad pixel table (BPIXTAB).
2048	A-to-D saturated pixels which are never useable, even at higher gain settings; flagged by ATODCORR, using thresholds in the CCD table (CCDTAB).
4096	Cosmic rays and detector artifacts rejected during MultiDrizzle image combination. These flags are not present in the drizzled output images (*drz.fits). Rather, they are propagated back to the [DQ] arrays of the input images (*ft.fits). Only data from pointing patterns (using pattern forms or POS TARGs) and CR-SPLITS are automatically associated and combined by MultiDrizzle in the pipeline.
8192	Cosmic rays rejected during the combination of CR-SPLIT images in the pipeline (ACSREJ); flagged in the [DQ] arrays of the output (*crj.fits) images.

^aRe-defined as of 8 October 2004. Was previously used to flag ‘blobs’ of bleeding around saturated pixels (at the end of bad columns). These blobs are also flagged as hot or saturated pixels, so this was redundant.

^bRe-defined as of 8 October 2004. Was previously used to flag ‘permanent’ (non-annealing) hot pixels which had persisted through the first four CCD annealing cycles after launch in 2002. These flags existed only in an increasingly outdated bad pixel table (BPIXTAB), and they were redundant with flag 16.




Figure 2: Bias structure in a typical WFC superbias. This image has been binned and smoothed to enhance subtle features. The two WFC chips are mosaicked: chip 1 or [sci,2] is on top, and chip 2 or [sci,1] is on bottom, with the interchip gap between them. Many bad columns are evident (their numbers are increasing with time, due to radiation damage), and each amplifier quadrant exhibits its own distinct structure. The HRC bias structure is relatively featureless, and is not displayed in this paper.




Figure 3: Dark structure in a typical WFC superdark. This image has been binned and smoothed to enhance the more subtle dark features (unobscured by the many hot pixels), and the two WFC chips are mosaicked, with the interchip gap between them. The growth of hot pixels has been well-documented (Sirianni, these Proceedings), but the more subtle structures seen here have not changed significantly since launch, and are mostly the result of CCD manufacturing processes. The HRC dark structure is relatively featureless, and is not displayed in this paper.




Figure 4: Overlaid HRC superdark histograms (number of pixels vs dark current in e^-/sec): a daydark, a basedark, and the resulting hybrid superdark. Our superdarks get most of their pixels from a high signal-to-noise 2-week basedark. A superdark for a given observation date is a copy of the basedark, with the warm pixels (0.2 to $0.8 e^-/\text{sec}$, identified with data quality flag 64) and hot pixels (above $0.8 e^-/\text{sec}$, identified with data quality flag 16) added from the corresponding 4-frame daydark.




Figure 5: A closeup of some typical WFC dark structures: Scattered warm and hot pixels (some with prominent CTE tails), and columns of saturated pixels. The corresponding data quality flagging is displayed in Figure 6.




Figure 6: Data quality [DQ] flagging corresponding to Figure 5. A colorful lookup table was chosen here, to give every flag value (or combination of flag values) a unique color (which can be seen in the on-line version). Hot pixels always have one trailing pixel flagged for CTE, but more of the CTE tails are also flagged as warm pixels. Saturated columns of pixels are also flagged uniquely.