Instrument Science Report WFC3 2014-22

Flagging the Sink Pixels in WFC3/UVIS

Jay Anderson and Sylvia Baggett
Nov 17, 2014

ABSTRACT

Fig13.Post-flashed calibration products have highlighted a previously undocumented type of image defect called "sink pixels" (SPs). These pixels apparently contain a number of charge traps and, as a result, they under-report the number of electrons that were generated in them during the exposure. Sink pixels can also have an impact on nearby upstream or downstream pixels. We explored this phenomenon in WFC3 ISR 2014-19. Here we take the next step and develop a strategy to flag them in science exposures.

1. INTRODUCTION

In 2012, WFC3/UVIS began making use of its post-flash capability in order to mitigate the extreme charge-transfer efficiency (CTE) losses that can result when the detector background is low. When we began to construct reference files of these ~12-e^- post-flashed exposures, it became clear that a small fraction of the detector pixels register systematically low for reasons unrelated to quantum-efficiency (QE) variations. Science images frequently have similar backgrounds to this, but since the dithering procedure ends up diluting the contributions of these low pixels to the drizzled stacks, their ubiquity has not been appreciated until recently. Even though dithering does lessen the impact of these pixels on stacks, they still represent a set of defects that will introduce a bias to an image if they are not corrected or flagged. It is particularly worth identifying the impact of sink pixels (SPs) for those datasets that do not have enough dithered exposures to allow adequate outlier rejection.

In CAL-13638 (PI-Anderson), we took a set of 10 post-flashed bias-type[1] exposures with three different levels of post-flash. These exposures were used to identify the low "sink" pixels and gauge their impact on images with various backgrounds.

2. THE OBSERVATIONS

Program CAL-13638 was allocated 10 internal orbits and was executed between Jun 30 and Jul 1 in 2014. The program took ten "bias" exposures with each of three post-flash (PF) levels: 25, 50, and 100 e^-. In addition to these data, we also analyzed 24 "bias" exposures taken with 12 e^- post-flash taken between May 21 and Jun 26 of the same year as a part of the standard dark-calibration pipeline program (CAL-13554, CAL-13557, and CAL-13558).

In total, we had 10 images for each of the three higher PF levels and 24 images for the 12-e^- level. We work here with only the `_raw.fits` images, since we wanted to include in our analysis the post-flash electrons (which are by default removed from the `_flt.fits` images). For each raw image, we first converted the pixels into a wide "`_raz.fits`" format. This orients all the amplifiers such that the readout is down and to the left (see **Figure 1**). We next subtracted a long-term average bias reference image (which will be described in a forthcoming ISR) and measured a bias level from the central 2070×27[2] pixels in the horizontal post-scan in each amplifier and subtracted this value from each pixel in the amplifier.

We then took the 10 (or 24) exposures taken at each post-flash level and determined a robust sigma-clipped average value for each pixel. This left us with a single average image for each of the four PF levels. **Figure 2** on the next page shows the stacked raz-format image for the 25 e^- background exposure ("PF25"). **Figure 3** shows a close up of **Figure 2** for the central region of Amplifier D.

---

[1] Technically, the images were post-flashed 0.5s darks, since flight-software constraints prevent post-flashing in bias mode. This is not a major problem, since even the official biases have effective dark times of about 3s.

[2] There are 30 columns of serial virtual overscan with each full-amp readout, but we discard the first two columns and the last column from the bias measurement since these columns often are not settled.

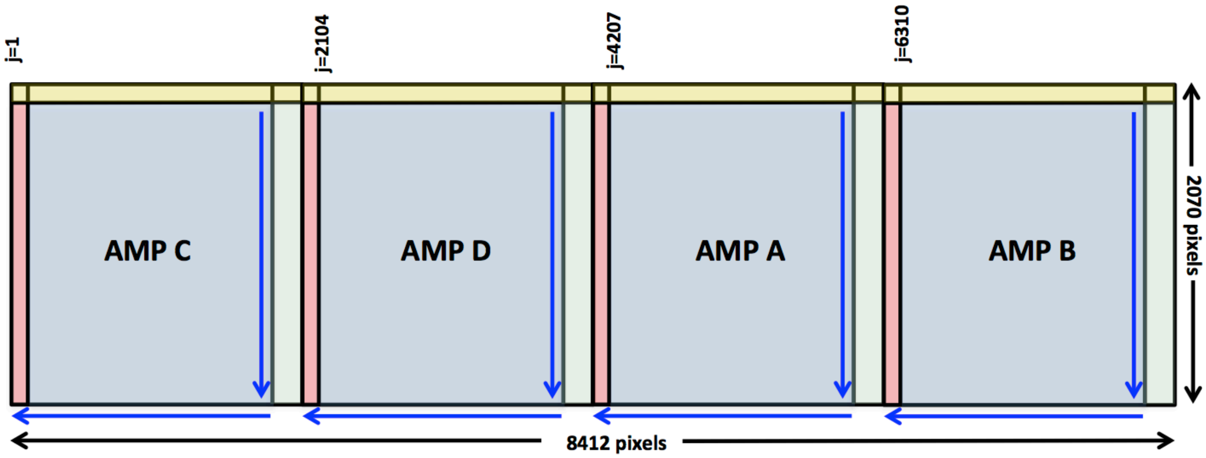

Figure 1: This figure shows the setup of the wide "raz" image format. The light blue regions are the 2048×2051 science pixels in each amp, each red region corresponds to the 25×2051 pre-scan pixels, each green region is the 30×2051-pixel horizontal post-scan, and the 2070×19-pixel yellow region is the vertical overscan. Readout directions are marked with blue arrows.

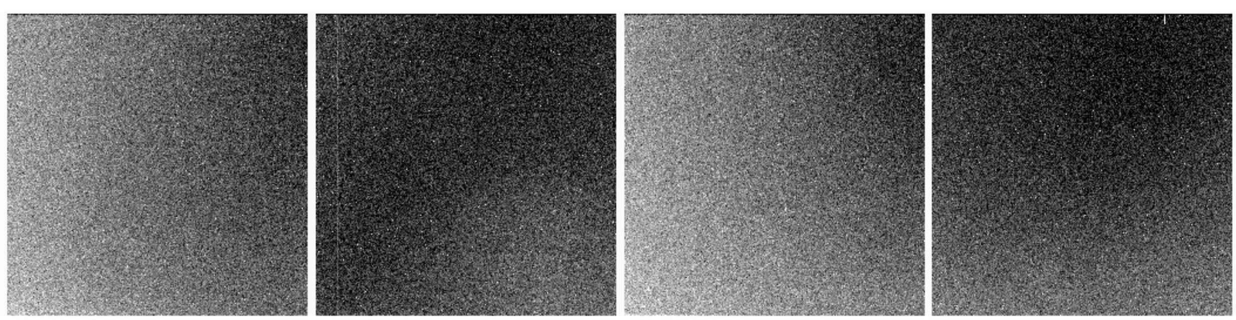

Figure 2: Average post-flashed _raz.fits file for the 25 e- post-flash. All figures here are shown in inverse grayscale, such that darker means more electrons. The flash level here goes from 28 e^- at the darkest points to 18 e^- at the lightest points.

Figure 3: This is a close up of a 350×325-pixel region in the upper part of AMP D, [3500:3850,1700:2025] in the raz-format image from Figure 2. The white pixels are the low (sink) pixels.

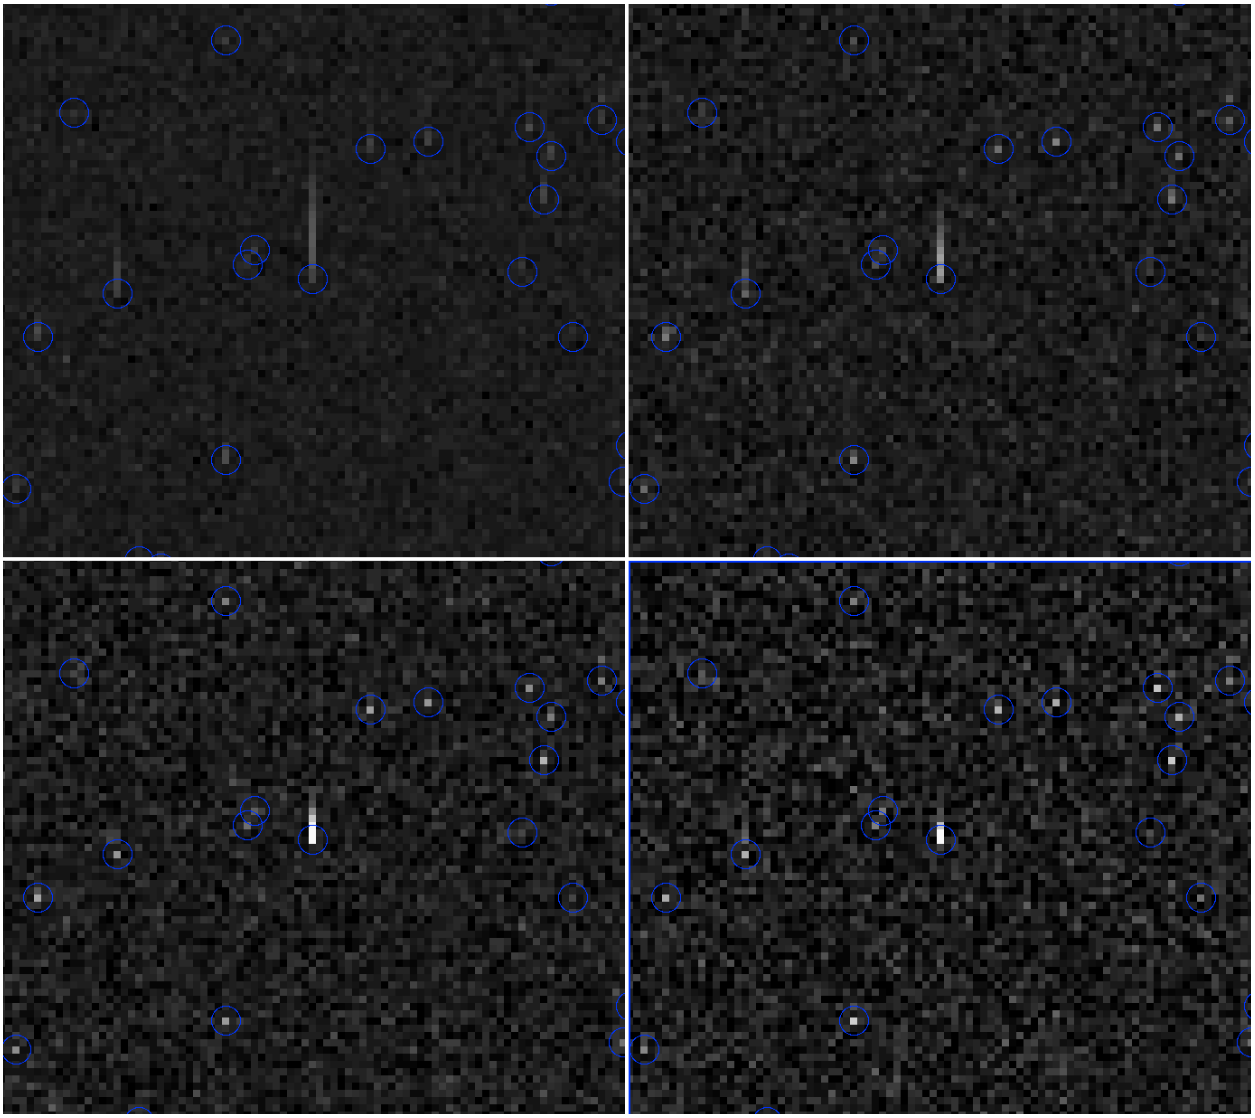

Figure 4: The roughly 100×100-pixel region in AMP D centered on (3680,3884) in the raz file; the white pixels are the low (sink) pixels and the upstream pixels that are affected by them. The upper left shows the background-subtracted stack of 24 exposures that were post-flashed at 12 e$^-$, the upper right shows the same, but for PF = 25 e$^-$, the lower left shows PF = 50 e$^-$, and the lower right shows PF = 100 e$^-$. The open blue circles show the SPs identified in our SP-finding algorithm, which required a deficit of greater than 7 electrons in both of the bottom two images.

The next thing we did was to subtract the smooth background of post-flash electrons from each post-flashed bias stack. Figure 4 shows the central portion of Figure 3 for all four smooth-background-subtracted post-flashed stacks. It is clear that the sink pixels occur in the same place in all exposures, but that they affect more pixels above them in the images with lower background (shown in the upper panels). This phenomenology is discussed in Anderson & Baggett (2014).

In order to find SPs to study and flag, we went through the PF=50 and PF=100 stacks (the bottom two panels in Figure 4) pixel by pixel and identified every pixel that was more than 7 electrons low in both stacks. We then threw out all low pixels that were directly above other low pixels, so as to identify only the bottom-most pixel of a contiguous streak as the SP. There are some scratches on AMP B, and we flagged those regions ahead of time so as not to identify any SPs among the ~13,000 pixels within ~5 pixels of the scratches. This procedure ended up identifying 41,762 candidate SPs (about 1 pixel in 400). These are shown with blue circles in **Figure 4**.

Figure 5 below shows the distribution of sink-pixel depths in terms of the residual in the 100-e^- PF stack. It is clear that we are not seeing *all* of the SPs. There are likely many pixels with fewer than 7 e^- worth of traps in them, but these are hard to reliably identify because the Poisson noise in the stack of 10 exposures with the 100 e^- of background is $\sqrt{BKGD} / \sqrt{N_{EXP}} = \sqrt{100}/\sqrt{10} = 3.1$ e^-. Consequently, our 7-e^- limit is at about the 2-σ level in the PF=100 e^- stack. If we were to identify all such pixels, then we would have a great many negative noise spikes in our list. The fact that we also require a similar detection in the 50 e^-, has a noise of 2.3 e^-, implying that only 3-σ excursions would qualify. Requiring a 7-e^- in both exposures means that we should be impacted only by ~5-σ noise excursions. If we were to lower our threshold below 7 e^-, then we would start identifying a significant number of noise excursions as SPs.

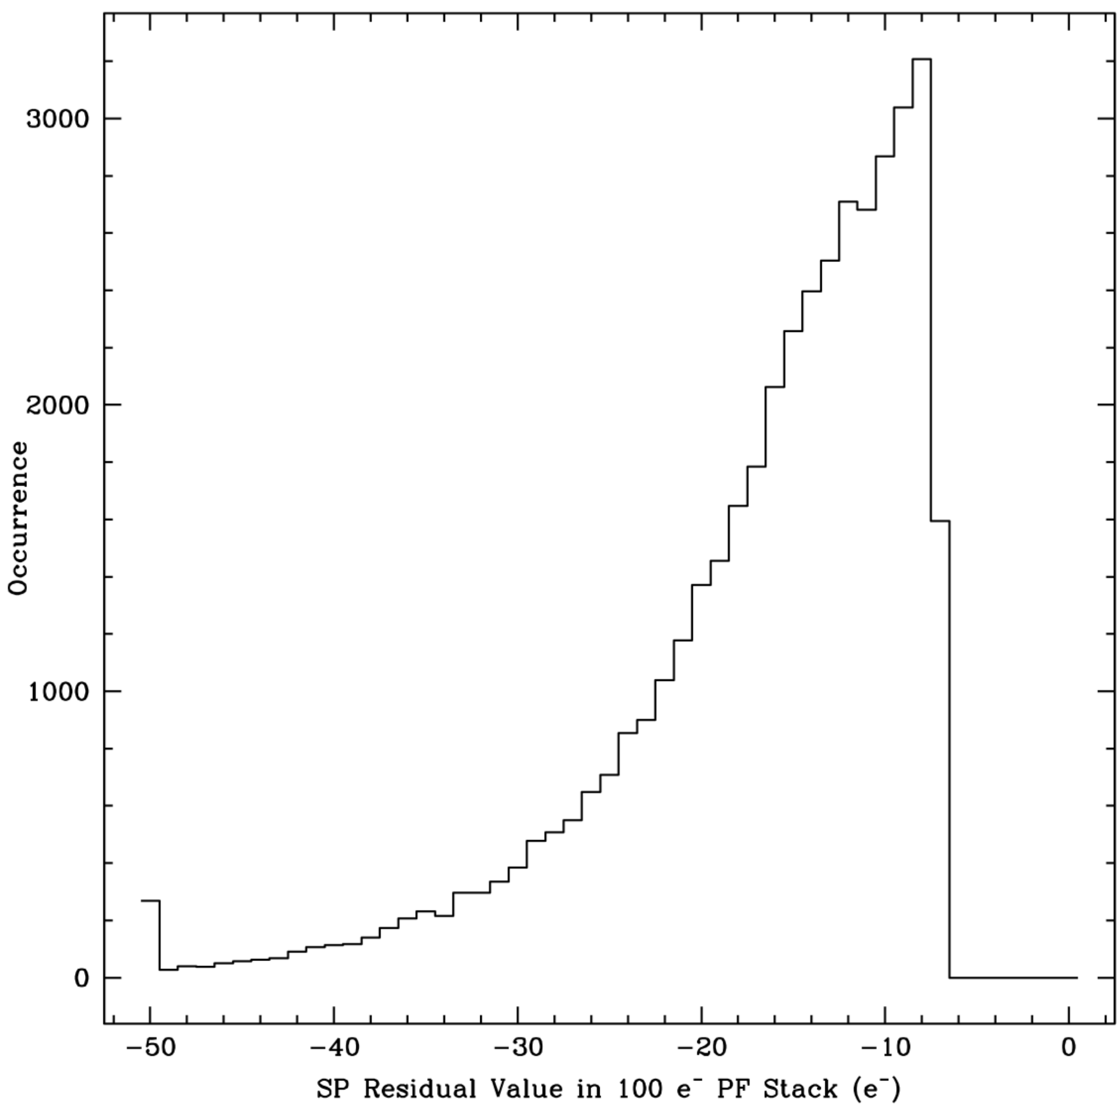

Figure 5: Distribution of SP depths. The occurrence corresponds to the number of SPs over the 16 million pixels in the four amplifiers. The bin at −7 goes from −7.5 to −6.5 and since the cutoff it at −7, it is only half filled.

3. The Disparate Impact of Sink Pixels

We noted in ISR 2014-14 that sink pixels typically appear more like delta functions in the high-background exposures, but they often extend upwards in the low-background exposures. This can be seen in **Figure 4** and is shown graphically in **Figure 6** below for the current data set. The different panels show SPs with varying intensity, and the colors show images with low and high backgrounds, from red to black, respectively. The dotted curves correspond to SPs at the top of the detector (where the PF background is higher) and the solid curves correspond to SPs at the bottom of the detector.

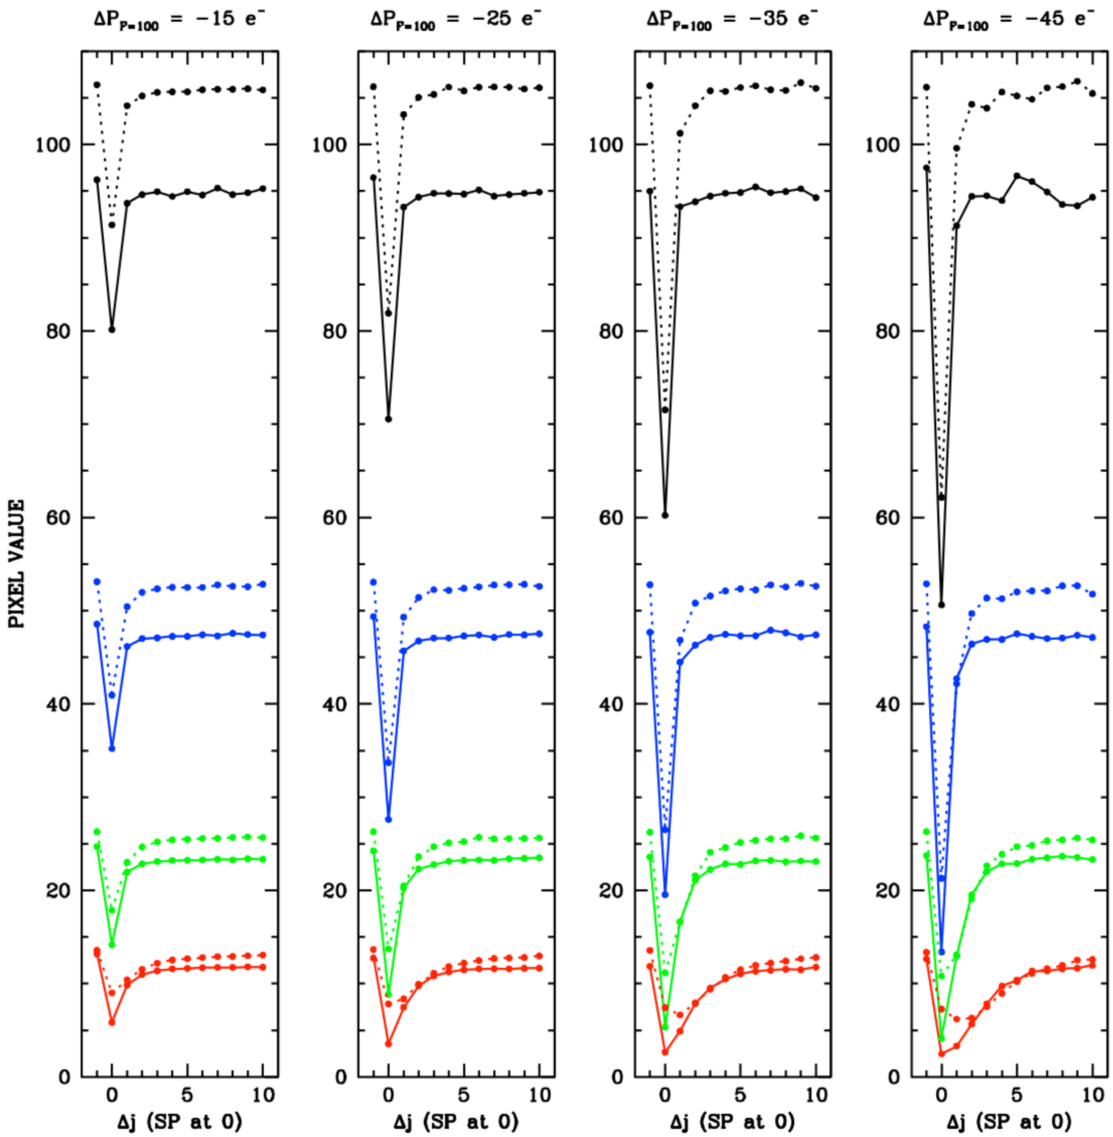

Figure 6: The average profiles of sink pixels with different depths in exposures with different backgrounds. Panels from left to right show, respectively, a composite average of sink pixels with 15, 25, 35, and 45 electrons of traps in each (according to the PF100 exposure). The different colors show the profiles for exposures with backgrounds of 12 (red), 25 (green), 50 (blue) and 100 (black) electrons. The solid curves show the profiles at the bottom of the raz image (close to the amplifier), and the dotted curves show the profiles at the top of the raz image (far from the amplifier and thus subject to more blurring and back-filling due to imperfect CTE).

It is clear that SPs often impact only one or two pixels when the background is high, but they can impact up to 10 pixels if the background is low. Because the impact of a given SP is scene-dependent, they will not be as easy to remove by calibration as, for example, hot pixels.

At first, we considered coming up with a correction strategy for the SPs, since it does seem that their impact can be characterized in terms of a predictable streak of electron deficits as a function of background. But then several questions arose: Which background should we use? And what would result if there happens to be a source somewhere within the streak? The images we have taken here are essentially bias frames, so they contain nothing but a smoothly varying background — science images will not generally be so simple.

It is not clear how to correct these pixels for an unknown scene, so we recommend simply flagging the pixels that might possibly be impacted by a given SP. Fortunately, if there is a source present, then the impact of the SP will be *less* (in both amplitude and extent) than in the case of no source, since a source will only fill up the traps faster. Thus, we will conservatively over-estimate the number of pixels to flag.

A remaining question, then, is how to characterize the background so that we can estimate which pixels to flag as questionable. From **Figure 6**, we see that the value within the SP varies monotonically with the background, so we can characterize the background implicitly from the value of the sink pixel itself. For each sink pixel in each of the four post-flashed stacks, we tabulate (1) the value of the sink pixel in that stack and (2) the range of pixels impacted by the sink pixel ($j_{1ow}\ldots j_{high}$).

To determine this impacted range we need to consider two things. The first has to do with the fact that sink pixels sometimes have some impact on the pixel below them, and the second has to do with the fact that the influence of SPs can extend upwards by several pixels, depending on the background.

To determine j_{low} (the lowest pixel impacted by the SP), we note that, as was first observed in Biretta & Bourque (2013), SPs often have the curious feature that the pixel downstream from the low SP (i.e., at $[i_{SP},j_{SP}-1]$) is sometimes *high* relative to the background. For some unknown reason, this tends to be more common closer to the amplifier than farther away.

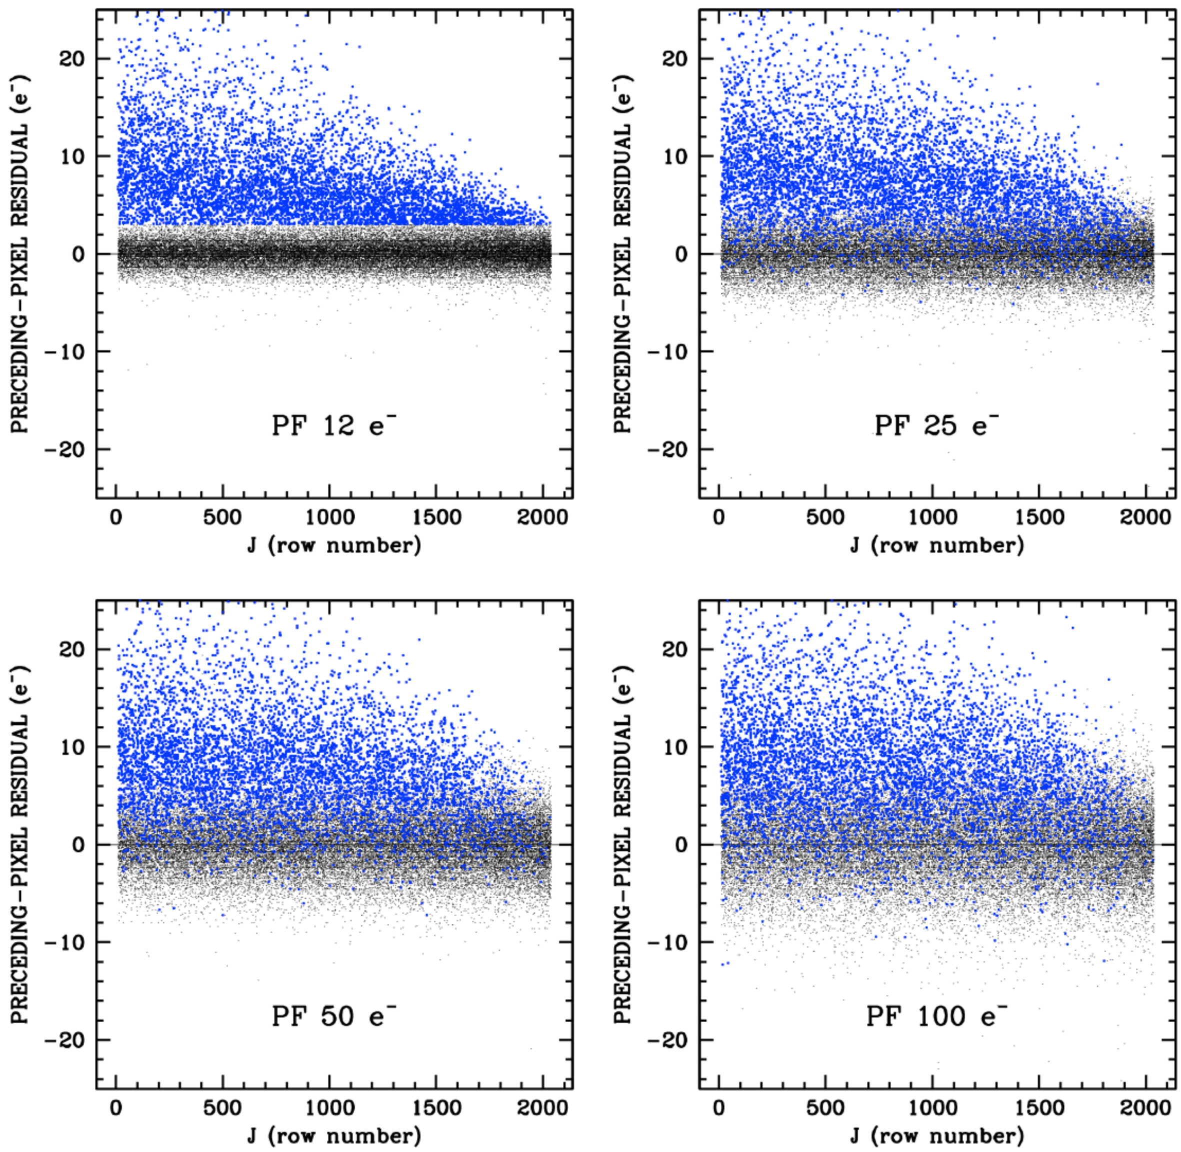

Figure 7: This shows the value of the pixel immediately downstream of the sink pixel relative to the local background in each of the four stacks, as a function of row number for each of the 41,379 SPs identified. The pixels that are more than 3 electrons bright in the 12-e⁻ PF stack are color-coded blue in all panels. They are shown with increased weight in order to help us distinguish them from the larger population of black points.

Figure 7 above shows the value of the downstream-of-sink pixel with respect to the local background for the four stacks with different levels of post-flash. We identified in blue the pixels that were at least 3 e⁻ above the local background in the 12-e⁻ PF stack (shown in the upper-left panel), and then we follow these same pixels in the stacks with different backgrounds. There appears to be a clear distinction between the normal SPs in black and the SPs that have a high downstream pixel, so we will adopt this discriminant as an indicator that the SP's impact starts one pixel below the actual SP itself.

There is a clear change in the distribution of enhanced downstream-of-sink pixels, in that the enhancement is greater closer to the amplifier. One might think that this could be due to CTE losses that whittle down the bright pixel on its way to the amplifier. But when the background is 12 electrons, we would expect at most a 25% reduction in the pixel from CTE losses, not the near 100% reduction that we observe here. The reduction takes place even when the post-flash background is 100 electrons, where the CTE losses should be considerably lower. For this reason, it does not seem reasonable to attribute the downward trends in the blue points to predictable CTE losses.

Our next task is to determine j_{high}, the uppermost pixel impacted by the SP. To do this we looked at how far up from the SP we had to go before the pixel value returns to the background level. We did this independently for the stacks from the four different post-flash levels. We started with the first pixel above the sink pixel (*n*=1) at $[i_{SP},j_{SP}+n]$ and flagged it as bad, since this pixel is almost always lower than the background (see **Figure 6**). We next asked whether the pixel value was more than 0.5-sigma below the local-background level (where sigma was defined by the RMS of the local background); if this pixel is more than a half-sigma low, we continued the procedure to the pixel above it. We iterated this procedure, increasing *n* by one until the 0.5-sigma-low condition was no longer met. In this way, we always flag one pixel beyond the last pixel that is more than 0.5-sigma low, just to be sure.

For the 41,762 SPs identified, the PF100 stack identified 136,000 impacted pixels to flag, the PF 50 stack identified 142,000, the PF25 stack identified 156,000, and the PF12 stack had 202,000 pixels to flag (about 1 in 100 of the detector's 16 million pixels). **Figure 8** shows the distribution of the number of pixels flagged at the four different post-flash levels as a function of the SP depth in the 100-e^- post-flashed stack. It is not surprising that the more traps we find in a pixel, the more pixels get flagged when the background is low (red curves).

By the above procedure, we generate a preliminary reference file that contains all the information we need in order to start flagging the pixels impacted by the SP traps. For each of the four PF exposures we record: (1) P_{SP}, the value of the sink-pixel itself, (2) j_{low}, the bottom pixel of the SP's influence, and (3) j_{high}, the top pixel of the SP's influence. In general, the lower the value of the SP itself, the more pixels upstream from it are.

In order to know how many pixels have been impacted by the traps in a particular sink pixel, we need to know the number of electrons that were recorded in the sink pixel itself. If the pixel value is high, then chances are most of the traps will already be filled and perhaps only one other pixel will be affected. If the pixel value is low, then chances are that the charge in the sink pixel will not have filled all of the traps, and several more pixels are likely to be impacted. The `_flt.fits` images have had their post-flash and dark-current electrons removed, so we cannot determine from them the true number of electrons recorded in the sink pixel. We therefore must use the _raw.fits image to determine how many pixels may be impacted. The P_{SP} value in the above preliminary database comes directly from bias-subtracted `_raw.fits` images.

For each sink pixel, the reference file tells us which pixels are impacted for four different observed values of P_{SP}: $P_{SP}^{[100e]}$, $P_{SP}^{[50e]}$, $P_{SP}^{[25e]}$, and $P_{SP}^{[12e]}$. For a particular science exposure, if the value of P_{SP} is greater than $P_{SP}^{[100e]}$, then we will (conservatively) flag all the pixels recommended by the PF=100 e^- stacks. If the value of P_{SP} is between $P_{SP}^{[100\,e]}$ and $P_{SP}^{[50\,e]}$, then we will flag the pixels recommended by the PF=50 e^- stacks, etc. If the value of P_{SP} is less than that in the 12-e^- stacks ($P_{SP}^{[12e]}$), then we will simply flag the number of pixels recommended by the 12-e^- stacks. Most WFC3/UVIS exposures now have at least 12 e^-, either naturally or via post-flash, particularly those taken after mid 2013, so the fact

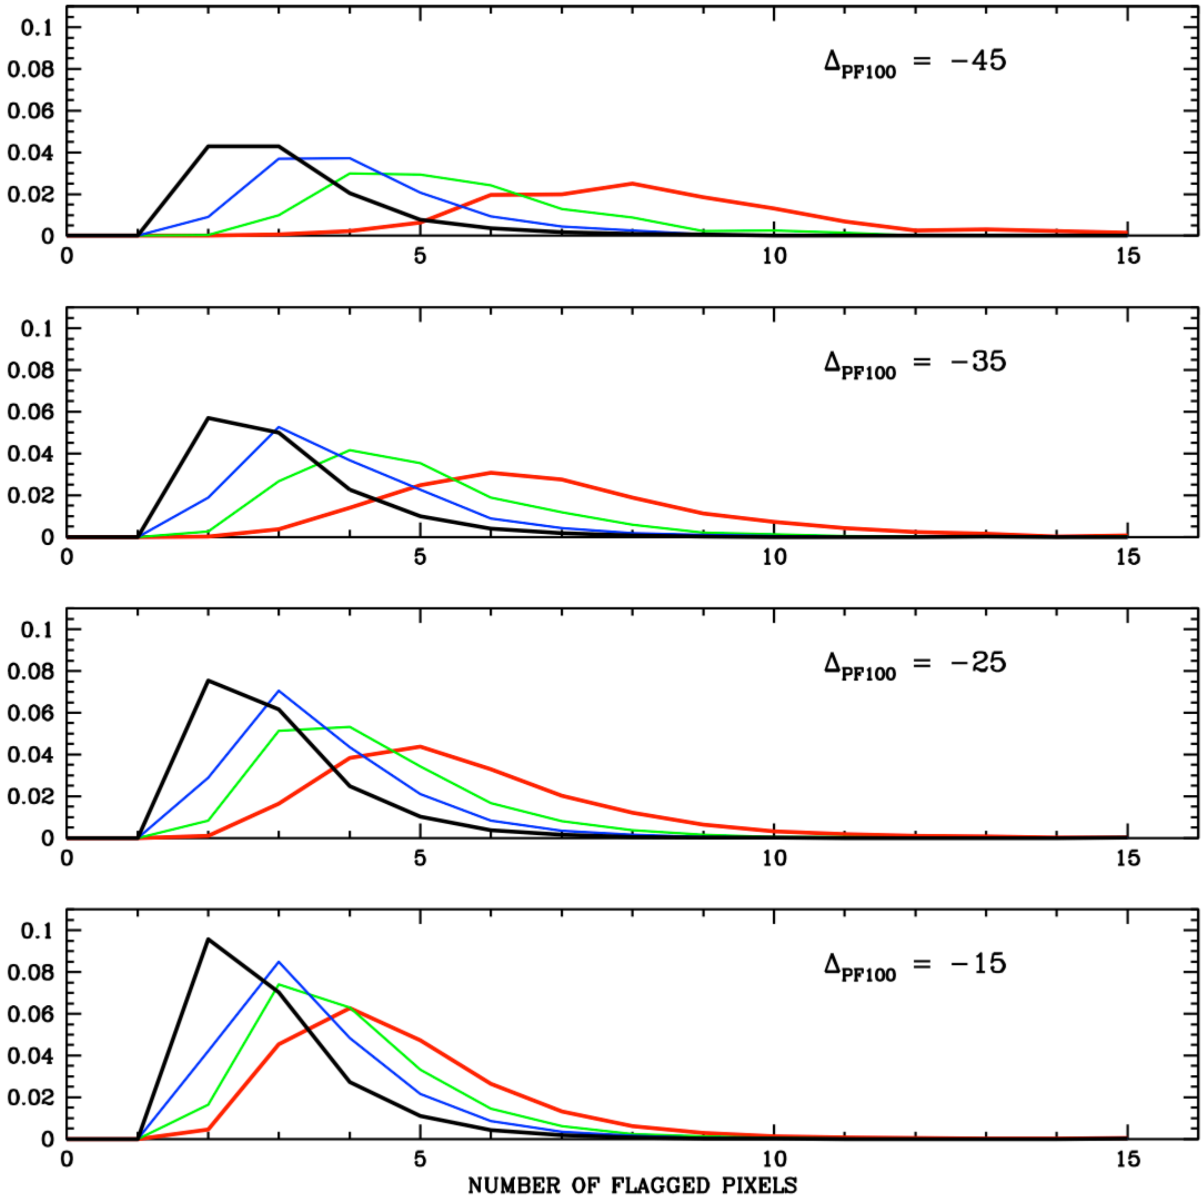

Figure 8: Distribution of the number of pixels flagged as affected by each sink pixel. The different panels correspond to SPs with different depths in the 100-e$^-$ PF stack, denoted by Δ_{PF} label in each plot. The black, blue, green, and red curves correspond respectively to the 100-e$^-$, 50-e$^-$, 25-e$^-$, and 12-e$^-$ post-flashed stacks.

that we cannot flag even more pixels when the background is even lower than 12 e$^-$ should not impact very many exposures.

4. WHEN SHOULD A SP BE FLAGGED?

The above procedure tells us which pixels to flag as a function of P_{SP}, the observed value of the SP in the bias-subtracted `_raw.fits` exposure. The only question remaining is *when* to flag the pixels.

Anderson & Baggett (2014) showed that it is often possible to examine the time-history of WFC3/UVIS pixels to determine when a particular pixel became a sink pixel. The vast majority of the SPs appear to have arisen during the lifetime of WFC3/UVIS on orbit, so for many of them there is a time before which they were perfectly normal and a time after which they are compromised. We would like our new SP correction to be valid for images taken before 1 Jul 2014, such that when the pipeline is run on older images, it will flag only the SPs that were active at the time.

To do this, we must have a "use after" date for each SP in the reference file. To determine the creation time for each SP, we extracted the entire time-history of WFC3/UVIS from installation to the beginning of 2014 for each of the 41,762 pixels from the "Master" images that have been described in Anderson & Baggett (2014). Along with the pixel value, we also record the average of the 5×5 surrounding pixels, the exposure time, and the gain.

To determine when a SP acquires its traps, we work only with the observations that have backgrounds between 8 and 15 electrons. **Figure 9** shows all the background-qualifying time-series data for the first thirty sink pixels in the list. To determine the time at which the SP turned on, we go through the exposures one by one, under the presumption that the "break point" could be just after any of them. For each putative break point, we compute a robust average residual between the SP and the local background for the preceding data and a robust average for the data taken after the break point. We then determine an average absolute residual for all the data, after removing the adopted pre/post averages. We identify the optimal break point as the location where the total absolute residual is minimized. The break point and the before/after averages are shown in each panel.

Most of the SPs in the list have easy-to-identify break points and before/after averages. **Figure 10** below plots the pre-break average against the post-break average. About 72% of the SPs (shown in cyan) have obvious break-points with a pre-break average level of about zero and a post-break average level that is negative. These are clearly SPs that arose on orbit before 2014. About 5% of the identified SPs (shown in blue) have negative residuals both before and after; these are likely SPs that were present before launch. About 14% (shown in red) exhibit no clear sink before or after. These are likely the SPs that have arisen since the beginning of 2014 (the last date that was included in the time-series database). About 6% (shown in green) have a peculiar "high" late average. About 2.5% (shown in black) do not follow any particular trend.

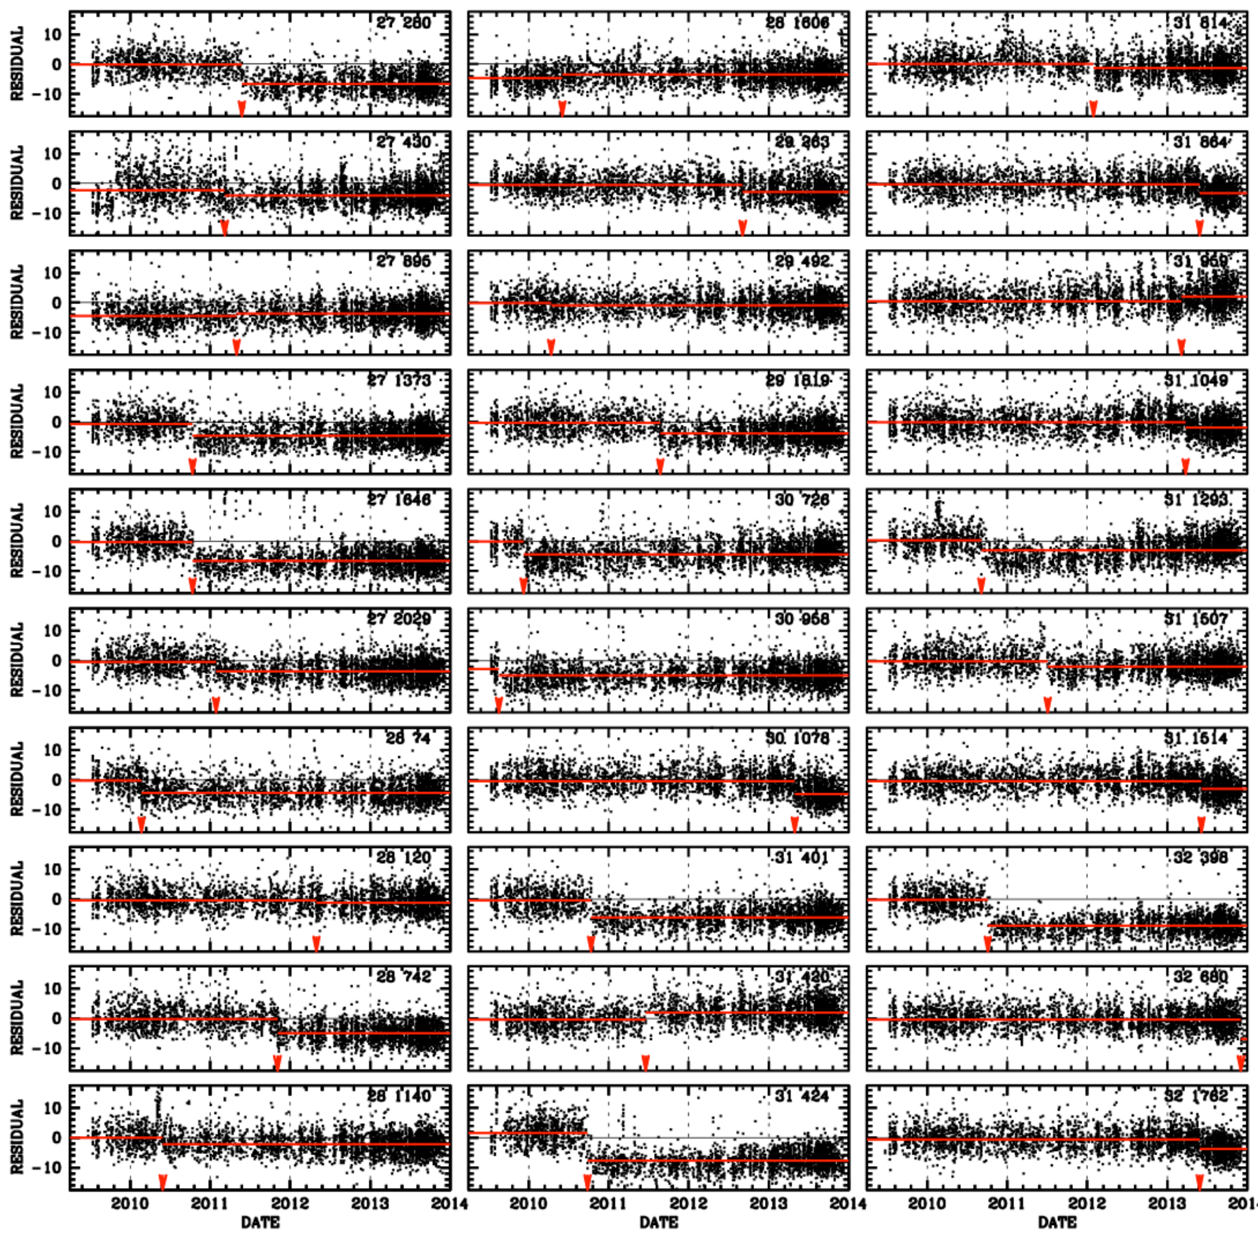

Figure 9: Time-series of residuals (in electrons) for the first 30 sink pixels. The residuals shown represent the difference between the value of the sink pixel and the local background in each exposure. The break-point is shown as an arrow at the bottom of each plot, and the before/after averages are indicated as solid horizontal lines. The majority of these SPs show clear evidence of a higher average before the break than after the break (e.g., pixels [27,280], [27,1373], [27,1646]). Some pixels became sinks early in the mission (e.g. [30,958], [30,726]) while others have only recently changed (e.g. [32,680],[31,864],[31,1049]. Others show extremely small differences in the average before and after the break (e.g. [27,895], [28,120] and a small number show higher average after the break than before (e.g.[31,959],[31,420]). The increase in the density of points after 2013 stems from the large-scale adoption of post-flash.

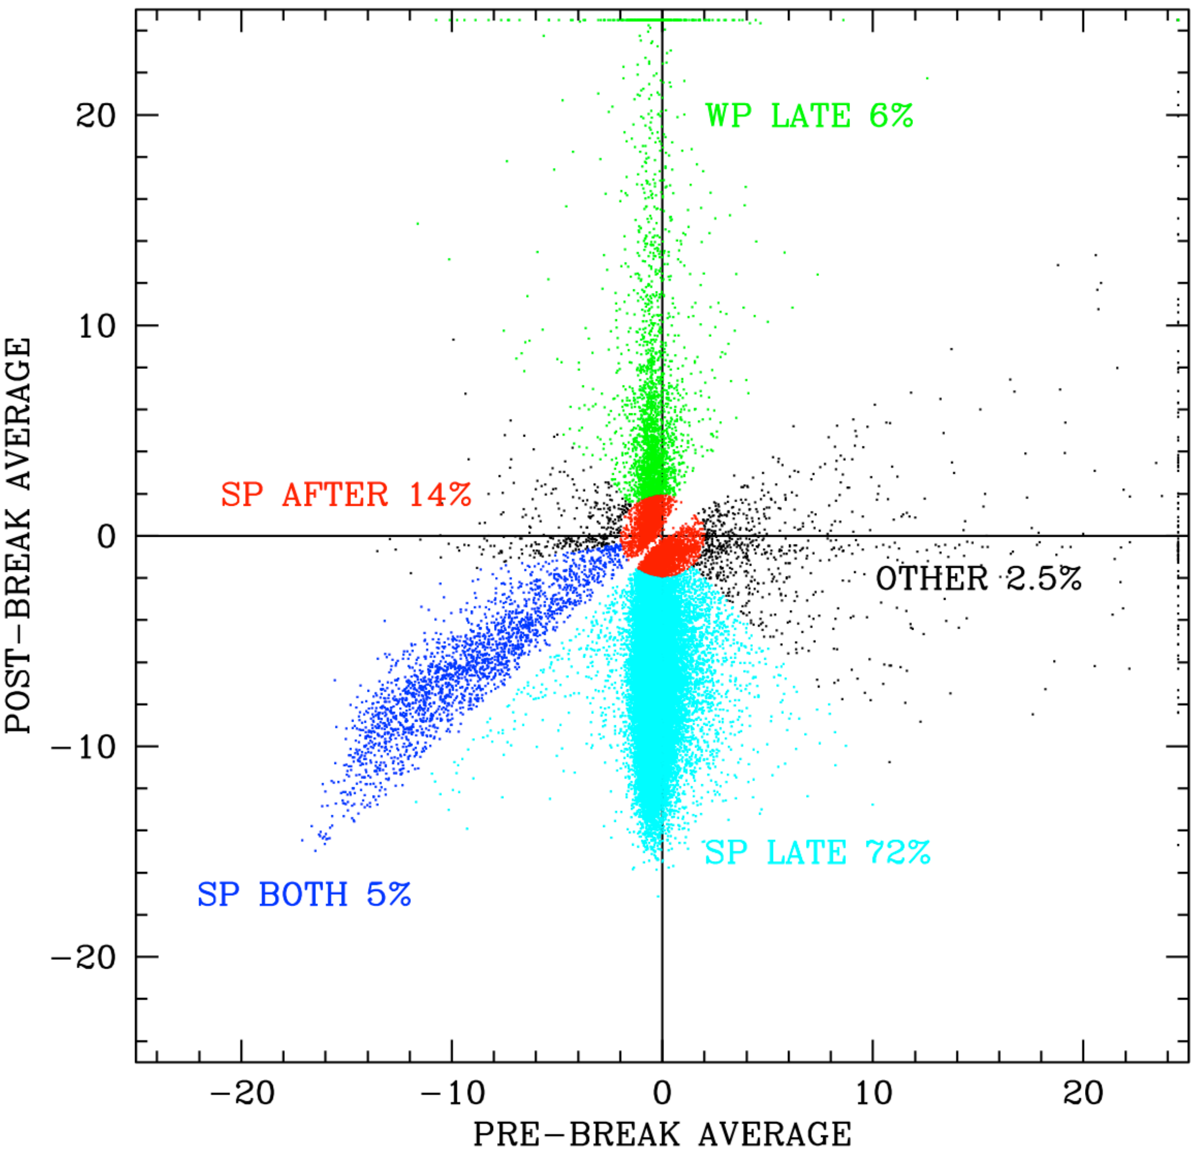

Figure 10: The pre-break average plotted against the post-break average (in electrons) for pixels identified as sinks (i.e., 7 e^- or lower relative to the background in PF=50 e^- and PF=100 e^- bias image stacks). The populations are color-coded as follows: The cyan pixels marked as 'SP LATE' became sinks sometime on-orbit. The blue pixels marked 'SP BOTH' were already sinks before launch. The green ones marked 'WP LATE' may be warm and not simple sinks. The red 'SP AFTER' pixels show no clear sink before or after. Finally, the black pixels marked 'OTHER' show no particular trend. Perhaps they were healed. The empty wedges at 45° and 225° arise from the fact that the algorithm tends to find *some* shift between the before/after level. The preponderance of blue points with higher pre-break averages to post-break averages can likely be attributed to CTE filling of the post-break SPs.

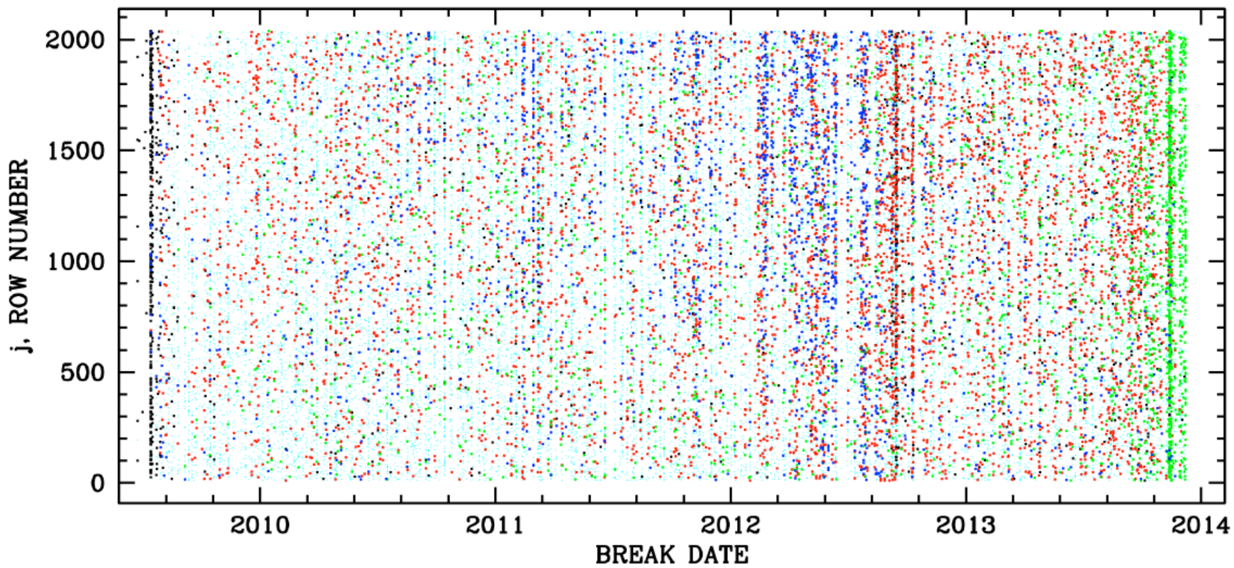

Figure 11: This shows the row number of the SP as a function of the break date for the variously classified SPs.

Figure 11 above shows the distribution of break-dates for all the SPs. The “easy” cyan cases have been shown with a small point-size so that it will be easier to study the harder cases. The black and green points appear to be clustered near the ends, indicating they may be spurious measurements without much data on the two sides of the break. The blue points appear to be relatively evenly distributed, with perhaps a concentration in mid-2012, which happens to be when PF began operation; this may have changed the distribution of achieved background in science images, so it is not surprising the algorithm might want to find a break-point there, all other things being equal.

We will identify the break-date for each SP as follows: For the cyan points, we will identify the observed break-date as the SP’s official break-date. For the red points, which show no clear SP in the data, we will conservatively assume the SP arose after the beginning of 2014 and will assign them a break-date of 2014.0. For the blue points, we will assign 2009.0, since they appear to have traps both before and after the solved-for break date. This leaves the black and green points. They also likely arose after 2014. In total, we have 22% arising after 2014, even though we would expect only 10% based on the fraction of the time post-2014 represents relative to WFC3/UVIS’s total time on-orbit. Perhaps these are not all true sink pixels, but it seems safest to assume that they are and flag them. There are simply too many of them to inspect by eye one at a time.

5. WERE THERE ANY SPS BEFORE LAUNCH?

The fact the blue-coded SPs were low even from the very first WFC3/UVIS exposure on orbit likely means that perhaps these pixels had traps even before launch. There was not much thermal-vacuum ground test data taken with the moderately low backgrounds that we require in order to see sink pixels, but we were able to identify a few images and examine the SP distribution in them. We found ~59 images from thermal-vacuum tests that had backgrounds between 20 and 200 electrons.

We have taken the 2115 “early” SPs indicated by the blue points in **Figure 10** and computed for each pixel in each exposure the residual between the value of that pixel at [i,j] and the average of the pixels just below it in that row, leaving out the pixel directly below it, in case that pixel happened to be elevated (i.e., we took [i, j−10 : j−2]). **Figure 12** shows the average of each of these low-both-early-and-late SPs in the data from the thermal-vacuum ground tests. It appears that perhaps 100 of the 2115 were low before even launch, accounting for about 5%.

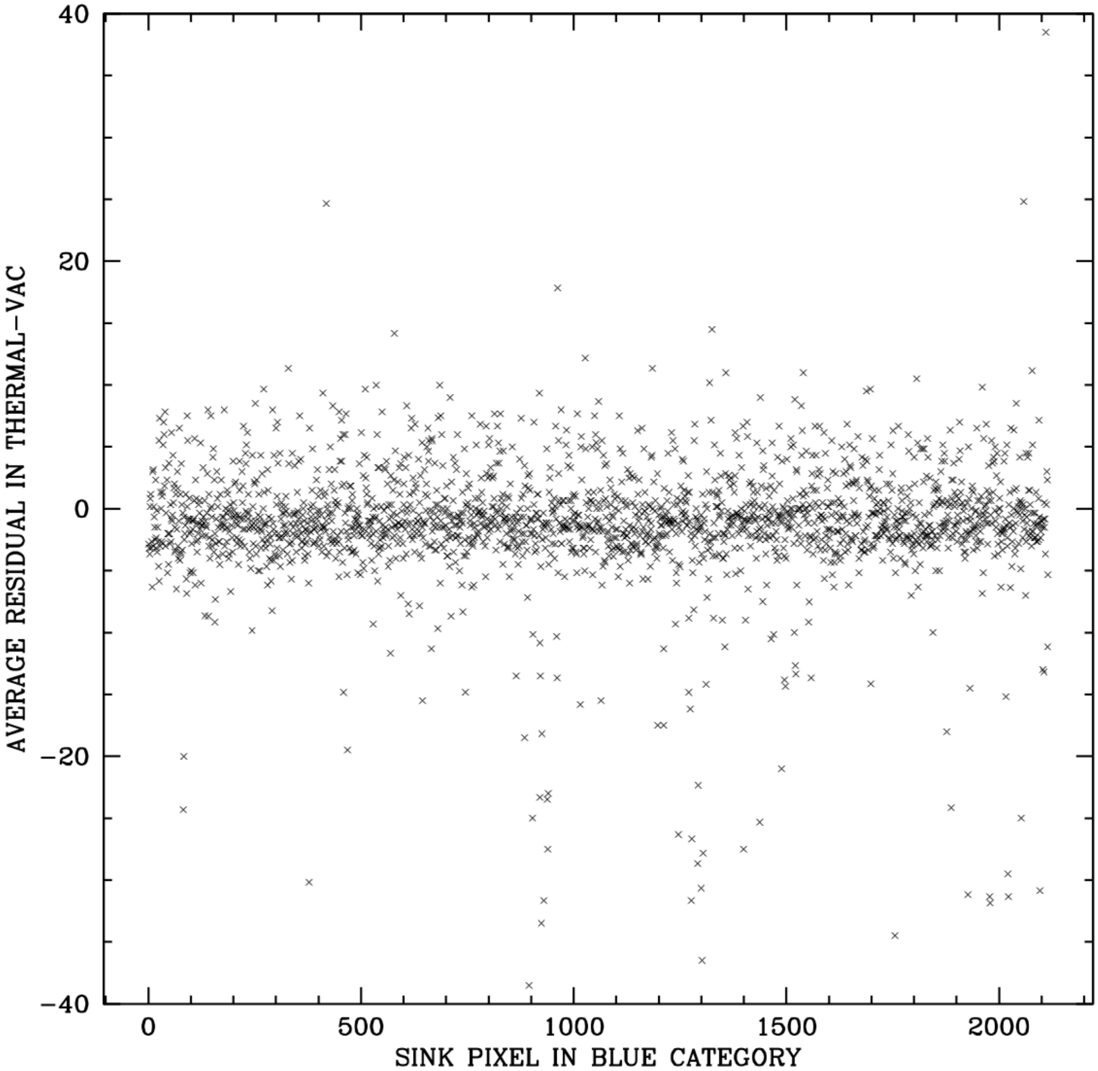

Figure 12: A plot of the residuals for the blue 'SP-BOTH' color-colored points from Figure 10 in the thermal vacuum test images.

To make sure that the pixels that we had identified as low both early and late (the dark blue points in **Figure 10**) were in fact properly identified, we examined the on-orbit time-history for the first 25 of these by eye to verify that they were indeed low the entire lifetime of WFC3/UVIS. We found that 20 out of 25 were indeed low from the first available exposure, and 5 became low within the first few exposures (which would have been hard to catch with the algorithm we used). We conclude that it is safe to flag these as SPs from the beginning of WFC3/UVIS's time on orbit.

6. THE FORMAT OF THE REFERENCE IMAGE

We considered generating a simple reference table that indicated which pixels to flag as SPs based on the time of the observation and the value of the SP. However, we decided to encapsulate all this information into a single reference fits image in order to provide an efficient, compact way to control the flagging as well as a convenient visualization of the sink-pixel locations.

The reference image has in each of the SPs the "flag-after" date in fractional years (2009 Jul 1 is 2009.50, etc.)[3]. If the pixel below the SP is one of the ~20% that have an excess, we assign that pixel the value of −1 in the sink-pixel reference file. Any pixel above the pixel that was part of the SP trough in the PF = 100 e^- image gets a value of 999 in the sink-pixel reference file, meaning that it will be flagged for all but the very brightest of backgrounds. Any pixel that was good in the PF = 100 e^- stack, but bad in the PF=50 e^- stack is assigned the value of the SP in the 100 e^- image. If the SP is below this value in a particular image, then this pixel will be flagged. Similarly, any pixel in the sink-pixel reference file that is good in the PF = 50 e^- stack, but not the PF = 25 e^- stack, gets assigned the value of the SP in the PF = 50 e^- stack. Finally, pixels in the sink-pixel reference file that are bad only in the PF = 12 e^- stack get assigned the value of the SP in the PF = 25e^- stack. **Figure 13** shows a portion of the reference image.

Given the reference image, the procedure for flagging SPs in science data is straightforward. We first read in the raw science exposure and do the bias-subtraction. We extract the MJD of the exposure from the header and convert it into fractional years (assuming that the reference file remains in fractional years rather than MJD, which is TBD). Next, we go through the reference image pixel by pixel. If we encounter a pixel value greater than 999, then that indicates that this is a sink pixel with its value representing the turn-on date. If the turn-on date of the SP is *after* the exposure date, then we ignore the SP in this exposure and move on to the next SP.

If the SP's turn-on-date is *before* the exposure date, then the pixel was compromised at the time of the observation. We thus flag the SP with the "charge trap" flag, 1024, in the science-image data-quality file (DQ extension). If the pixel "below" the SP in the raz-format reference image has a value of −1, then we activate the charge-trap flag for that pixel as well. We then proceed vertically up from the SP and compare each pixel in the reference file to the value of the SP in the science exposure at hand. If the value of the SP in the exposure is below the value of the upstream pixel in the reference image, we flag that pixel with '1024' in the data quality file of the exposure. We continue to flag until there are no more flag-able pixels for this SP (i.e., until the value of the pixel in the reference image is zero), or until the value of SP in the exposure is greater than the value of the upstream pixel in the reference image.

---

[3] Note that we may ultimately use MJD in this reference file rather than fractional years. A date in fractional years is more intuitive for humans to inspect, but the pipeline deals always deals with time-related corrections in MJD. To go from MJD to rYR, we use: rYR = 2009.0 + (MJD-54832.5)/365.24, with the inverse transformation naturally being: MJD = 54832.5 + (rYR-2009.0)*365.24.

The reference image that we constructed here has 47,762 SPs (0.25% of the 2048×2051×4 total science pixels), with 6,296 pixel-below-SPs flagged with "−1". There are 167,260 pixels (about 1% of the total) that are to be flagged according to the observed value in the SP.

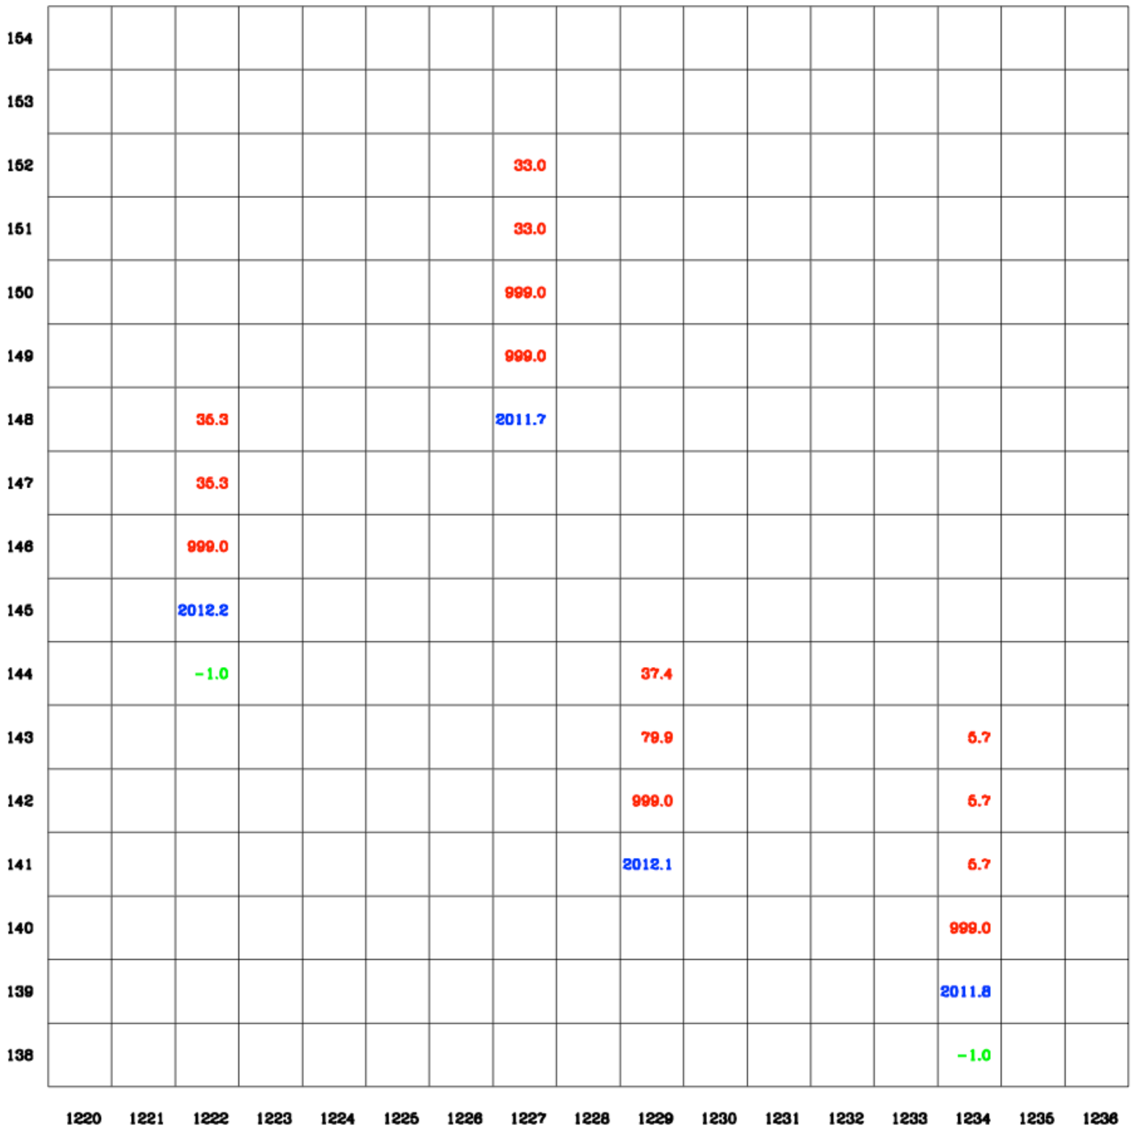

Figure 13: A portion of the sink-pixel reference file image, containing four sink pixels. The blue values indicate the SPs themselves, the green values indicate the enhanced downstream pixels, and the red values indicate low upstream pixels. Blank pixels here have values of zero and are not impacted by SPs.

7. DEMONSTRATING THE IMPROVEMENT FROM SP FLAGGING

The FORTRAN code provided in the APPENDIX shows how to generate a list of pixels to flag from the SP reference image and a given raw science image. The routine reads in a `_raw.fits` image in the "raz" format and the reference image. It goes through the reference image and identifies the SPs and flags them in the science-image data-quality file. For each SP, it also determines whether the downstream pixel is likely to be enhanced by the SP and flags it accordingly. Finally, it proceeds upstream from the sink pixel and flags as many pixels as are expected to be impacted.

We ran this routine on two post-flash bias images, one for a relatively low background and another for a high background and ended up flagging about 1% and 0.6% of the pixels, respectively. **Figure 14** shows histograms of the values of the pixels flagged (where the pixel value reported is in terms of the local background). We have color-coded the histograms in terms of whether the pixel was flagged because it was a SP itself (blue), an enhanced pixel below a SP (green), or whether it was in the trough associated with the SP (red). The green and blue histograms show pixels that are clearly impacted by the SP; the typical pixel in the trough is affected at about the 3-electron level, which is the essentially the readnoise.

The routine that performs this flagging is being prepared for inclusion in the WFC3/UVIS pipeline. Stay tuned.

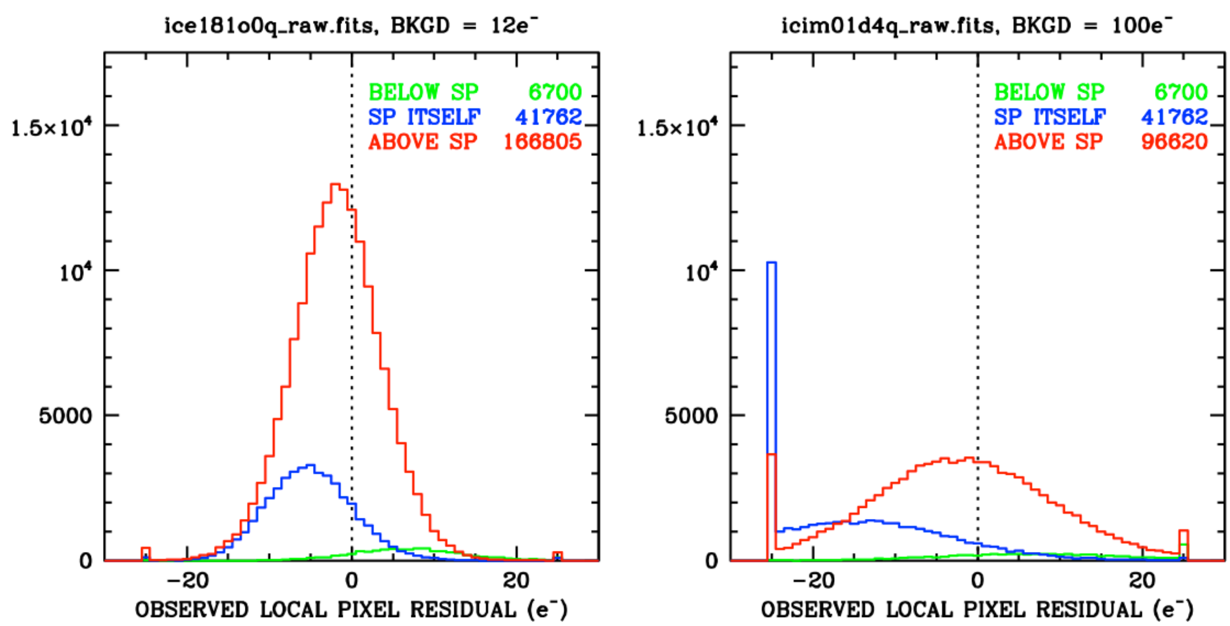

Figure 14: Histogram of the flagged pixel values relative to the local background for post-flashed bias images with low background (12 e^-, left) and moderate background (100 e^-, right). The green histograms show the pixels downstream of the SPs that are often high. The blue histograms show the SP itself, and the red histograms show the pixels that are upstream from the SP and likely affected.

References

Anderson, J. & Baggett S. “Sink Pixels and CTE in the WFC3/UVIS Detector” WFC3/ISR 2014-14

Biretta, J. & Bourque, M. “WFC3 Cycle 19 & 20 Dark Calibration Plan: Part 1” WFC3/ISR 2013-12

Note: sample code in Appendix has been removed.