Updating the WFC3/UVIS CTE Model and Mitigation Strategies

Jay Anderson, Sylvia Baggett, and Ben Kuhn
June 29, 2021

ABSTRACT

The pixel-based charge transfer efficiency (CTE) correction was last updated for WFC3/UVIS in 2016. Since the strength of CTE generally increases linearly with time, as does the population of warm pixels, the effect is almost twice as strong after five additional years of HST operations, and there are twice as many warm pixels available to use in deriving the correction. Unfortunately, the new model confirms that charge-transfer losses have continued to increase steadily since WFC3/UVIS was installed. Now, even with a background of 20 electrons, CTE losses are almost 50% for sources near the chip-gap (i.e., sources that undergo the maximum number of parallel transfers), compared to about 30% in 2016. The pixel-based algorithm works well only when the correction is small enough to be considered a perturbation on the signal received. If the correction is too large, then the algorithm tends to amplify noise, which causes more harm to the image than benefit. For this reason, the default pipeline setting is now designed to suppress the correction for faint sources in order to avoid noise amplification. In general, for observers with relatively bright targets (S/N > ~30) on image backgrounds of at least 20 e⁻ per pixel, the new pixel-based CTE correction works well, correcting targets to within 5%. Observers with fainter targets, which the new algorithm treats minimally in order to avoid noise amplification, will need to make additional adjustments. We provide advice to users about how to plan observations and how to carry out reductions in order to get the most out of WFC3/UVIS observations. A future ISR will provide more detailed prescriptions on how to correct sources bright and faint.

1. Introduction

WFC3/UVIS was installed on board HST in May of 2009 during Servicing Mission 4, and the moment it was carried above the earth's atmosphere by the space shuttle Atlantis, the detector began to suffer unrelenting radiation damage. Since the existing wide-field imager, ACS/WFC,

didn't start showing significant CTE (charge-transfer efficiency) degradation until many years after its installation in March of 2002, it was expected that WFC3/UVIS would similarly enjoy several years of nearly perfect charge-transfer. Unfortunately, WFC3/UVIS began showing CTE issues very soon after its own installation.

The reason that CTE became an issue for WFC3/UVIS so quickly is that UVIS was designed to operate with extremely low background, thanks to its low dark current, low readnoise, UV sensitivity, and large complement of narrow-band filters. As it turns out, the first few electrons in a pixel are particularly vulnerable to charge-transfer losses. As such, UVIS started to show imperfect-CTE issues very early in its lifetime. The WFC3/UVIS chips were designed with a mini-channel that was supposed to help with CTE losses for faint sources, but it is not clear how much it helps in practice (more on this below).

In 2009, when ACS was exhibiting continually degrading CTE after 7 years on-orbit, Massey et al. (2010) examined the hot/warm pixels in the GOODS dataset and constructed a model that corrected CTE in that specific dataset. Seeing the success of this approach, Anderson & Bedin (2010, AB10) studied the trails behind hot and warm pixels in dark exposures and developed a generalized model that worked on *all* datasets. This correction quickly made it into the HST pipeline, which now produces two kinds of level-2 products for each exposure: uncorrected `_flt` images and CTE-corrected `_flc` images, and the associated `_drz` and `_drc` images.

After just a few years on-orbit, it became clear that a pixel-based CTE model for UVIS would be useful, so we repeated the ACS analysis on UVIS darks. One difference between the ACS model and the UVIS model was that the ACS images rarely had backgrounds below 20 e$^-$, so the ACS model did not need to be particularly accurate below this level. In contrast, the UVIS model would have to be able to deal with very low backgrounds. To address this, we took a series of shorter darks, so that we could explore the CTE of WPs at lower charge levels. Analysis of these exposures made it clear that the charge-transfer losses increased significantly when the background was low, with losses being particularly bad for backgrounds below about 12 electrons. This realization led to the practice of post-flashing WFC3/UVIS images to ensure this minimum background level.

This "post-flash to 12" guidance worked well for several years, but by 2017 it was clear that losses were becoming problematic, even with a background of 12 electrons. Clearly, we needed to revisit the WFC3/UVIS model. For this reason, we took some detailed calibration data designed to address two deficits in the correction models. First, we needed to improve the parameters of the pixel-based correction, and second, we needed to make a detailed determination of the CTE-loss profile with background: was 12 e$^-$ truly a "sweet spot" in the curve? Or could it be that higher levels of background might provide more mitigation?

The WFC3 web pages on CTE[1] provide a broad repository of information about how our understanding of CTE has evolved. ISR-2020-08 by Anderson introduces the new correction and describes the losses as of 2019. The current report describes the new model in more detail, and also analyzes data taken in early 2021 to provide the most updated possible picture of CTE. It is clear that 10+ years of unrelenting radiation damage has brought the detector into a new phase of its life. It is no longer possible to simply post-flash and assume that the pixel-based correction can make the images as good as new.

---

[1] https://www.stsci.edu/hst/instrumentation/wfc3/performance/cte

Section 1 summarizes the results of the first pixel-based model from 2012 and the analysis that led us early on to recommend post-flashing at 12 e⁻. **Section 2** describes the theoretical basis for the pixel-based model, how it is parametrized, and how the readout process is simulated. **Section 3** describes the original correction of WFC3/UVIS that started as a standalone supplemental routine in 2013 and was implemented in the pipeline in 2016. **Section 4** describes the new model, which is better constrained for both low and high pixel values. **Section 5** discusses the issue of noise amplification, and the fundamental limitations it places on any pixel-based model. Stars and other astronomical sources are different from warm pixels in that they have some self-shielding that should lessen their losses. **Section 6** describes a recent calibration program that shows empirical CTE losses for actual faint stars near the detection limit. **Section 7** puts together all the above wisdom and provides advice for GOs planning future observations, in order to help them preserve as much of their signal as possible. Finally, **Section 8** provides advice on how one might measure point sources in exposures, either from the uncorrected `_flt` images, the new corrected `_flc` images, or other custom versions of the `flc` images.

2.   The Pixel-Based CTE Model

The pixel-based CTE algorithm has at its core a model for the trapping and release of charge during parallel transfer. The basic "forward" model takes an image as the electrons would be initially registered on the detector and simulates what the image would look like after it gets pushed through the CTE-blurring parallel-transfer process and read out at the serial register.

The "reverse" model does the opposite: it takes an observed image that has been pushed through the readout process and tries to determine the most likely input image that could have resulted in the image that was read out. We use an iterative procedure on the forward model to arrive at the reverse model. This is described in detail in AB10.

2.1 Some Phenomenology

The pixel-based model simulates the bucket-brigade transfer of flux down a column and it deals with one trap at a time. The details of charge trapping are not perfectly well understood. A useful way of thinking about this is to think of a pixel as a 3-D box-like region in the silicon lattice. Its boundaries are defined by the electric-field imposed by the electrodes. This electric potential well keeps the electrons that are associated with a pixel confined within the pixel boundaries.

When a pixel contains a small number of electrons, they do not have much self-repulsion, so they are confined by the field to a small region at the center of the pixel. When a pixel collects more and more electrons, their increased self-repulsion causes the electron cloud to take up more volume within the pixel, pushing back against the potential walls. At some point, a pixel collects so many electrons that the cloud cannot be contained within the potential well, and electrons dribble out into neighboring pixels. It turns out that the potential walls are "higher" between the columns than they are between the rows, so charge tends to bleed along the columns in the ±y direction.

Radiation damage results when a particularly energetic particle — such as a cosmic ray (CR) — travels through the silicon detector and displaces atoms in the lattice. Depending on which atoms are displaced and where these displacements lie within the pixel lattice, this can result in either a hot or warm pixel or a cold/sink pixel (see Anderson & Baggett 2014).

In the case of a hot or warm pixel, the damage is at the edge of the pixel and electrons are able to trickle into the pixel from outside, resulting in dark current. More electrons trickle in when the exposure is longer.

In the case of a cold/sink pixel, the radiation damage results in a "trap" in the silicon that prevents electrons in the pixel from moving freely out of the pixel. All traps will not affect all electron clouds. If the trap is near the center of the pixel's potential well, then it can affect even small clouds, but if it is closer to the edge of the well, then only very full electron clouds will be impacted by the trap.

The impact of a charge trap is two-fold. First, when the detector is exposed to the sky or to post-flash, the pixels get filled with charge before they are read out. If there is a trap in a pixel and the pixel's charge cloud is large enough to reach the location of the trap, then the trap grabs an electron. And when the charge cloud is shifted out of the pixel, that electron stays with the trap rather than following along with the original pixel's electron cloud. Later, perhaps, when a very small cloud is shuffled through, the trap will release the electron and it will be shuffled down the detector. However, the electron will now be associated with the cloud of a different pixel. The initial pixel will be read out as having one fewer electron than it started with. If we could somehow flash the detector with 5 noiseless electrons per pixel and record the image without read-noise, then the pixel with the trap will reach the readout register with 4 electrons (we call such a pixel a 'sink' pixel), and another pixel upstream will reach it with 6 electrons.

Given the presence of readnoise and Poisson noise, it would be hard to see the impact of a single trap in a pixel. Nevertheless, we *do* see the systematically low sink pixels in our images (see Anderson & Baggett WFC3/ISR 2014-22). This leads us to conclude that some pixels have many more traps than other pixels. A new analysis by Montes-Quiles (in prep) shows that the amount of charge that a sink pixel can grab depends on the charge level in a pixel. Pixels with ~20 electrons might find all of their electrons trapped, but pixels with over 100 electrons might find 30 electrons trapped, and pixels with over 1000 might find 40 trapped. One can imagine that the radiation damage from a particularly energetic CR would lead to a 1-D line of defects in the 3-D silicon lattice of the pixel, and the overlap between that line and the pixel's initial charge cloud will tell us how many electrons in the charge cloud can be delayed in the sink.

The second impact of a trap is a loss in charge-transfer efficiency (CTE). The electron cloud from a particular pixel is "parallel" shuffled down the detector to the readout register and then "serial" shuffled to the readout amplifier. Some of the pixels through which the cloud gets shuffled on its way to the serial register have trapping defects in them. If the charge cloud is large enough, some of its electrons may get trapped in the pixel and will not be included with the original cloud when it is read out at the amplifier. This results in CTE loss.

2.2 Parametrization of the Model

The typical electron cloud gets parallel-shuffled through many hundreds of different pixels on its way to the serial register. From the upcoming ISR by Montes-Quiles, over 0.5% of the

WFC3/UVIS detector pixels have traps in them that will grab 5 or more electrons out of a 100-electron cloud. The fact that there are so many traps in so many different pixels suggests a statistical approach wherein the ensemble of lines of defects through pixels result in a "typical" distribution of traps within the "typical" pixel. In our model, we treat all pixels as having the same distribution of (fractional) traps.

The first parameter of our model is simply a one-dimensional cumulative function $\phi(q)$ that gives the number of traps in the average 2048-pixel column that affect pixel-packet sizes q or greater (q is the number of electrons in the packet/cloud). This function is parameterized by its value at 17 points: q = (1, 2, 3, 5, 7, 12, 20, 30, 50, 70, 100, 300, 1000, 3000, 10000, 30000, 99999). The model turns this cumulative function into a discrete list of traps, each of which affects a particular ordinal electron in a cloud (for example, the first, fifth, seventy-first or the one thousand-third, etc.). If a particular electron cloud does not have at least q electrons, then the cloud will not be affected by a trap that affects the q^{th} electron.

The second aspect of the pixel-based readout model is the release of charge. The charge that is trapped in one pixel is eventually released in an upstream pixel. A second set of parameters regulates the probability that an electron that is trapped will be released after 1, 2, 3, etc., shifts. This model specifies the fractional release probability as a function of pixel shifts, $\tau(\Delta j)$. The function τ can also be a function of q (which ordinal electron is impacted). This could make sense if the probability of release is not just a simple time constant but rather is some function of the intensity of the oscillating electric current involved in shifting the charge down the detector. It could be that the traps closer to the bottom of the pixel's potential well could be more (or less) sensitive to this forcing.

In summary, the parameters of the model are ϕ and τ. The first parameter $\phi(q)$ corresponds to the number of traps affecting packets with a charge of q, and $\tau(\Delta j; q)$ describes the release profile of the charge affecting traps that grab the q^{th} charge. To simulate the readout, we turn $\phi(q)$ into a discrete list of traps, q_1 through q_N, where N is the total number of traps that affect a full charge packet.

2.3 Simulating Readout with the Model

The readout algorithm deals with one trap (q_N) at a time. As we mentioned above, each trap affects a particular ordinal electron in the packet (say the 44th). The algorithm goes up the column and when it encounters a packet that has at least 44 electrons, it removes an electron from the packet and resets the trap to full. It then shuffles the charge down. If the next pixel packet has fewer than 44 electrons, the trap releases a fraction of the grabbed electron into that pixel— $\tau(1; 44)$ — then moves on to the next pixel. If that pixel also has fewer than 44, it releases $\tau(2; 44)$, and so on. If it encounters a pixel that has 44 electrons or more, then it determines what fraction of the trap is still empty, fills it with electrons from that pixel, then resets the trap to $\Delta j = 0$.

In practice, pixels closest to the readout register experience fewer traps than those that are far from the readout register, so we scale the size of traps from 0 to 1 according to $j/2048$, where j is the row number relative to the serial register. Also, the distribution of charge in a pixel cloud can actually change as the charge is shuffled down the detector: a packet with 100 electrons at the top of the detector may have only 85 halfway down the detector, and as such it would be

affected by fewer traps. Therefore, we perform the shuffle in 7 stages, shuffling a seventh of the way down in each stage.   Similarly, the new model describes CTE losses as of November 12, 2016, the date the pinning data were taken. For dates before or after this time, the model should be scaled down or up, and the easiest way to do this is to just change the "size" of each trap. So, in summary, every pixel experiences the same number of traps, but the size of each trap varies linearly with j and with the observation date.

The code that simulates this transfer was provided in FORTRAN in an appendix of AB10. This algorithm has changed very little since then, but some of the loops have been reorganized to speed-up performance.   In what follows, we used the FORTRAN version of the code to optimize the model parameters $\phi(q)$ and $\tau(\Delta j; q)$, then translated it into C for its use in the pipeline. The list of traps and release profiles is provided in a reference file.

3.   The Initial Pixel-Based Correction

In 2012, the initial WFC3/UVIS CTE model was constructed in a similar manner to the procedure used in AB10 for ACS/WFC data. As had been done with ACS, we studied the trails behind warm pixels in stacks of 900s dark exposures to infer how much flux had been lost from them as a function of the WP intensity.   We found that the ACS model did a similarly good job restoring the electrons in UVIS images, even though the detectors are quite different. In particular, the WFC3/UVIS detector was constructed with a "mini-channel", which was designed to provide some shielding from CTE loss for small electron clouds.

Thanks to the lower dark current and lower readnoise in the WFC3 detector, we were able to follow WPs down to fainter levels than had been possible with ACS. This turned out to be important for UVIS, since many exposures are taken with very low backgrounds.

The model was able to confirm what we had suspected: when the background is low, the CTE losses can be severe for faint sources (50% or more). But when the background is above ~12 e⁻, the losses are diminished considerably (less than 15% as of 2012). This realization led the WFC3 team to work with the engineers at Goddard to quickly enable a procedure that could post-flash the images with an LED lamp in order to ensure a well-calibrated low level of illumination (Biretta & Baggett, 2013). At the same time, Baggett & Anderson (2012) performed a study of the natural backgrounds in UVIS images to advise users how much additional post-flash they needed to add to their images in order to ensure a minimum background of 12 electrons so as to minimize CTE losses.

The WFC3/UVIS detector was designed with low readnoise and low dark current in order to take advantage of its sensitivity in the ultraviolet, where objects are often faint on very low backgrounds. Adding 12 electrons of background noise (which came with 3.5 e⁻ of Poisson noise) was not ideal for such observations, but this added noise was not much greater than the readnoise, so UVIS continued to be able to easily observe faint objects. Also, users could reconfigure their observations to optimize their signal-to-noise. Rather than take many short exposures with low background and added post-flash, they could take fewer exposures with more natural background and less post-flash.

The correction worked well for several years. But by 2017, losses even in images with 12-electron backgrounds were now greater than 25%. The pixel-based reconstruction procedure works best when the losses are a perturbation on the initial charge distribution. Indeed, it uses the remaining charge as a "scaffold" to gently redistribute flux. When that scaffold ceases to resemble the original distribution, we lose valuable and unrecoverable information about the original distribution. In other words, when losses approach 25%, the algorithm either does nothing, or worse: it amplifies noise.

So, in 2017 we began a new study of CTE in WFC3/UVIS. The initial part of this study involved determining the actual CTE loss profile with pixel-packet size $\phi(q)$, so that we can recommend optimal levels for post-flashing. This has led to the more recent recommendation of at least 20 e⁻ post-flash; that is, images should have at least 20 e- total background (dark+sky+postflash). We also used the calibration data taken in 2017 and later to re-pin the model so that the pixel-based reconstruction can be as accurate as possible.

4.   A New Direct Way to Constrain the CTE Moodel

The original CTE correction was based completely on a study of warm/hot pixels in the standard calibration darks. It was therefore indirect: we had to infer from the trail behind a WP how many electrons it had lost so that we could map the losses as a function of WP intensity. We could not see the losses directly. Iteration allowed us to develop a consistent model for the trapping and release.

This indirect approach worked well for the brighter WPs, where losses were a small fraction of the remaining WP. As a consequence, the model worked well for larger charge packets, but it did not work so well for fainter WPs, where losses were larger than 25%, and where it was hard to detect and measure the entire trail. This led us to devise a more direct approach.

In CAL-14880 (PI-Anderson), we took a series of 50s "short" darks with a variety of post-flash levels. These can be combined with 900s "long" darks from CAL-14534 (unflashed) and similar 900s "long" darks from CAL-14531 that were post-flashed to have 12 e⁻ background. The goal is for the long darks to provide the "true" intensity of each warm pixel (the number of electrons per second), and the short darks would then allow us to examine directly how much of the deposited charge was successfully charge-transferred to the readout register, as a function of row number and sky background.

Our plan in re-pinning the CTE model parameters will be to start at the faint end (the traps that affect the smallest packets) then move on to the bright end. This is because the "low" traps that grab the first few electrons affect all packets, big and small, whereas the "higher" traps that grab only the 1000th or 10,000th electrons in a pixel will have no effect on smaller packets. This allows us to focus on the low traps in isolation and know that they are properly accounted for when we get to the high traps.
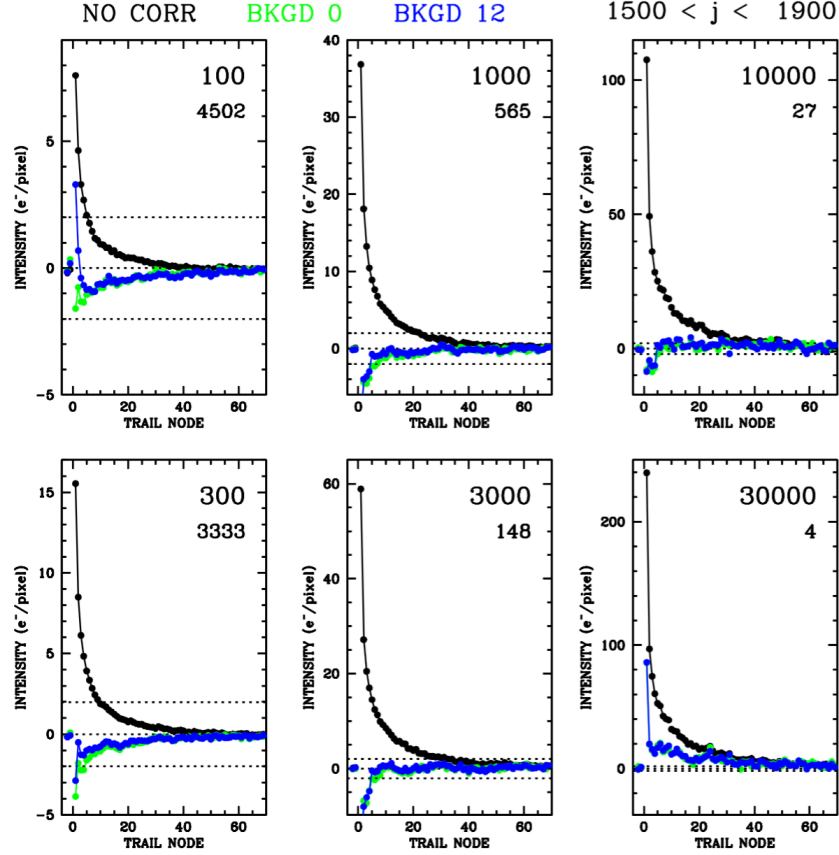
Figure 1: Trails behind six fiducial WPs with 900s intensities of 100, 300, 1000, 3000, 10000, and 30000 electrons. Each trail represents a stack of between 4502 and 4 individual WPs extracted from the upper quarter of the detector. The black curve corresponds to the trail in the uncorrected stack that had been 12-e⁻ post-flashed. The green curve shows the trail behind the v1.0 CTE-corrected unflashed stack, and the blue curve shows the trail behind the corrected 12-e⁻ flashed stack.

4.1 Identifying the Warm Pixels for Analysis

We start by making a master list of the warm pixels from the long darks. We have two sets of long darks — flashed and unflashed. At the time of this initial analysis (2017), the long darks from the official calibration program had a post-flash level of 12 e⁻; we had thirty-seven such 900s exposures and stacked them. We identified a WP as any pixel in this stack that was brighter than its eight adjacent and semi-adjacent neighbors in 30 out of 37 exposures. There were 56,938 qualifying warm pixels.

We also stacked ten 900s darks that had no post-flash, for comparison. Finally, we ran the v1.0 CTE correction on the flashed and unflashed stacks. For each of the identified WPs in each of the stacks, we extracted the region from 2 pixels downstream ($-y$) to 100 pixels upstream ($+y$) in the same column.

Figure 1 shows the tails behind six fiducial warm pixel levels (100, 300, 1000, 3000, 10000, and 30000 electrons) for the uncorrected and the corrected unflashed images and the corrected 12e⁻-

flashed images. All the WPs shown come from the upper quarter of the detector, far from the serial register, meaning that they suffer maximal CTE losses. The smaller number underneath the WP level in each panel shows the number of WPs identified in each bin.

The v1.0 CTE-reconstruction model does a decent job correcting the middle-intensity WPs, but it appears to over-correct the faintest WPs and under-correct the brightest ones. It is easy to understand the model's limitations at the bright end: even in late 2016, after more than seven years of sustained radiation damage in orbit, there were only *four* extremely bright WPs in the top quarter of the detectors. There were even fewer in 2012 when the model was first developed.

4.2 Calibrating the Dark Time

The goal in the new v2.0 CTE correction approach is to take the WP intensity from the corrected long darks as the "truth" so that we can examine directly the losses in the short darks as a function of background level. However, before we can simply "scale down" the WPs observed in the long darks to compare against the WPs observed in the short darks, we must pay some attention to the darktime. The darktime includes the commanded dark time as well as image-setup-overhead and time spent post-flashing.

We compared the intensity of the WPs observed in the CTE-corrected long darks and the uncorrected short darks. To get a proper darktime ratio, we compare the WPs that are close to the serial register ($j<100$), where CTE losses should be minimal. We compared WPs that had between 1000 and 3000 electrons in the long exposures so that the short exposures would have between 50 and 150 electrons and thus would not suffer large fractional losses. **Figure 2** shows this comparison in the various panels.

The commanded darktime was 900s for the long darks and 50s for the short darks. This would suggest a nominal ratio of 0.055555. The lower-right panel in **Figure 2** shows that this is close to the WP intensity ratio for the first exposure (which has no flash time), but the other exposures exhibit progressively longer effective darktimes, on account of the longer postflash (PF; both lamp-time and setup time). The down-shift at exposure number 8 reflects the transition to the medium-current postflash (which produces a higher count rate than the low-current postflash).
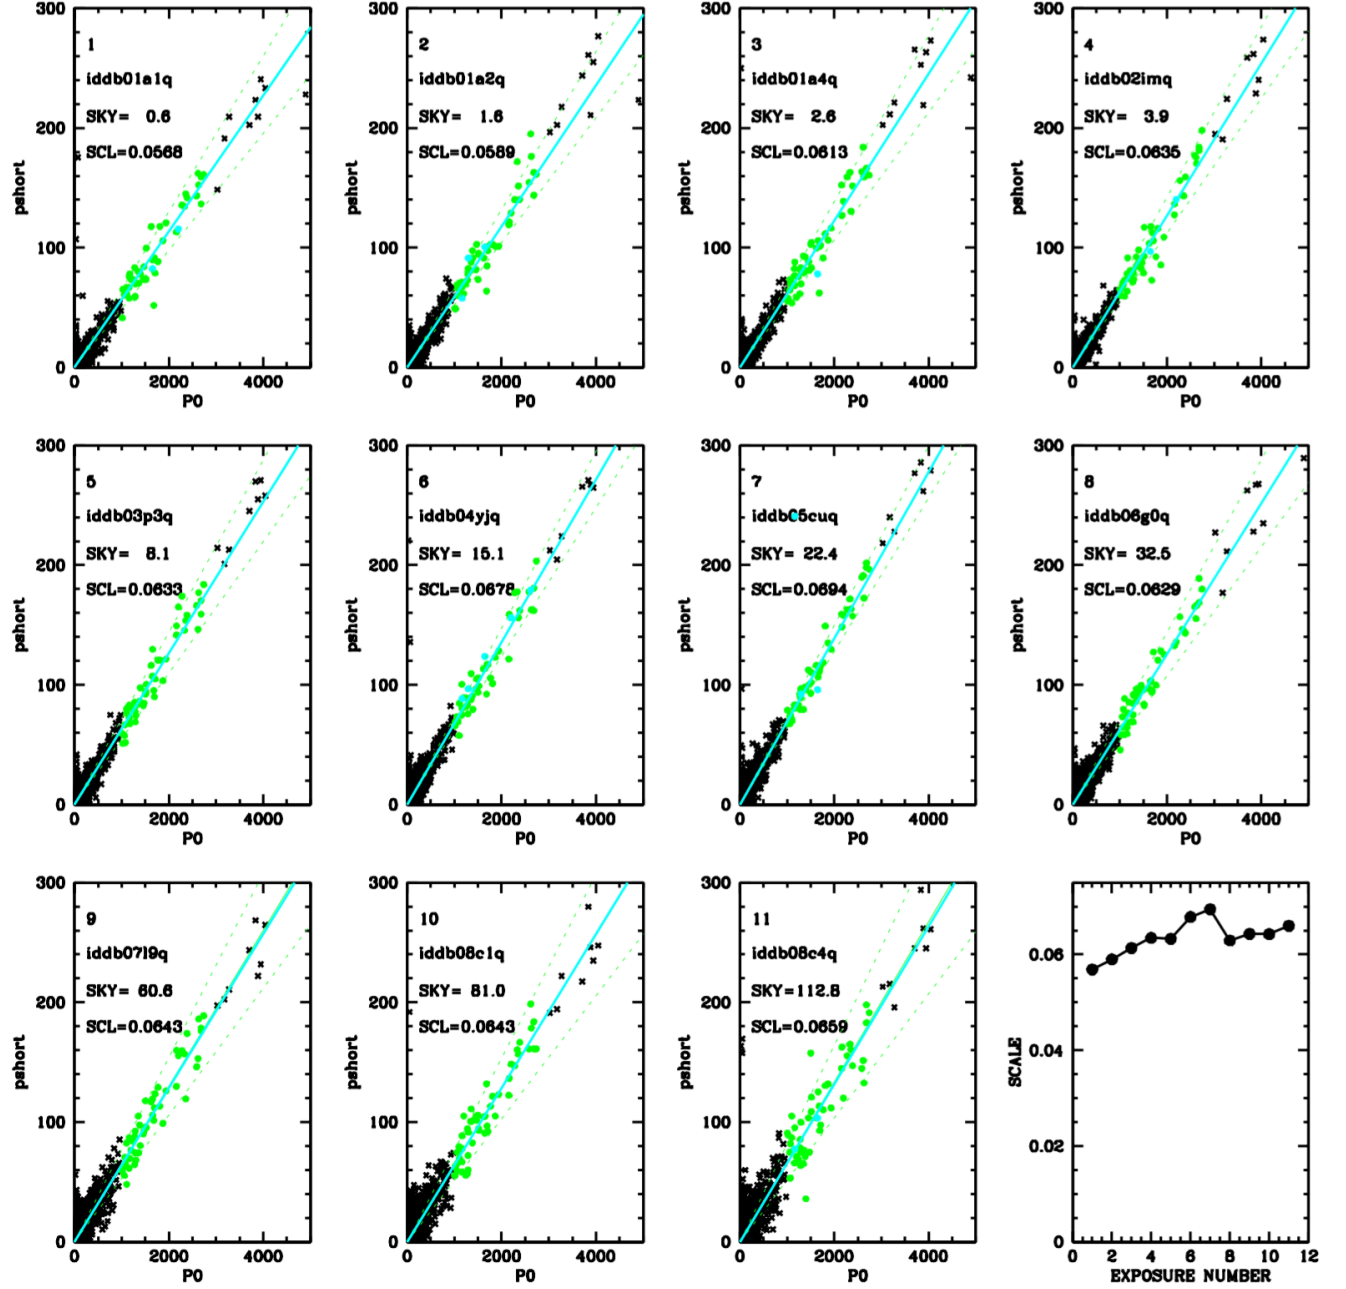
Figure 2: These panels show the determination of the scaling ratio between the 900-s long-dark WPs and the various 50-s short-dark WPs, taken with different post-flash levels. The green points show the WPs used to fit the slope with the 900-s flux along the horizontal axis and the 50-s flux on the vertical axis (both in electrons). The slope between these two provides the scaling ratio between the long dark and each individual short dark. The lower right panel shows the derived scale factor as a function of the exposure number.
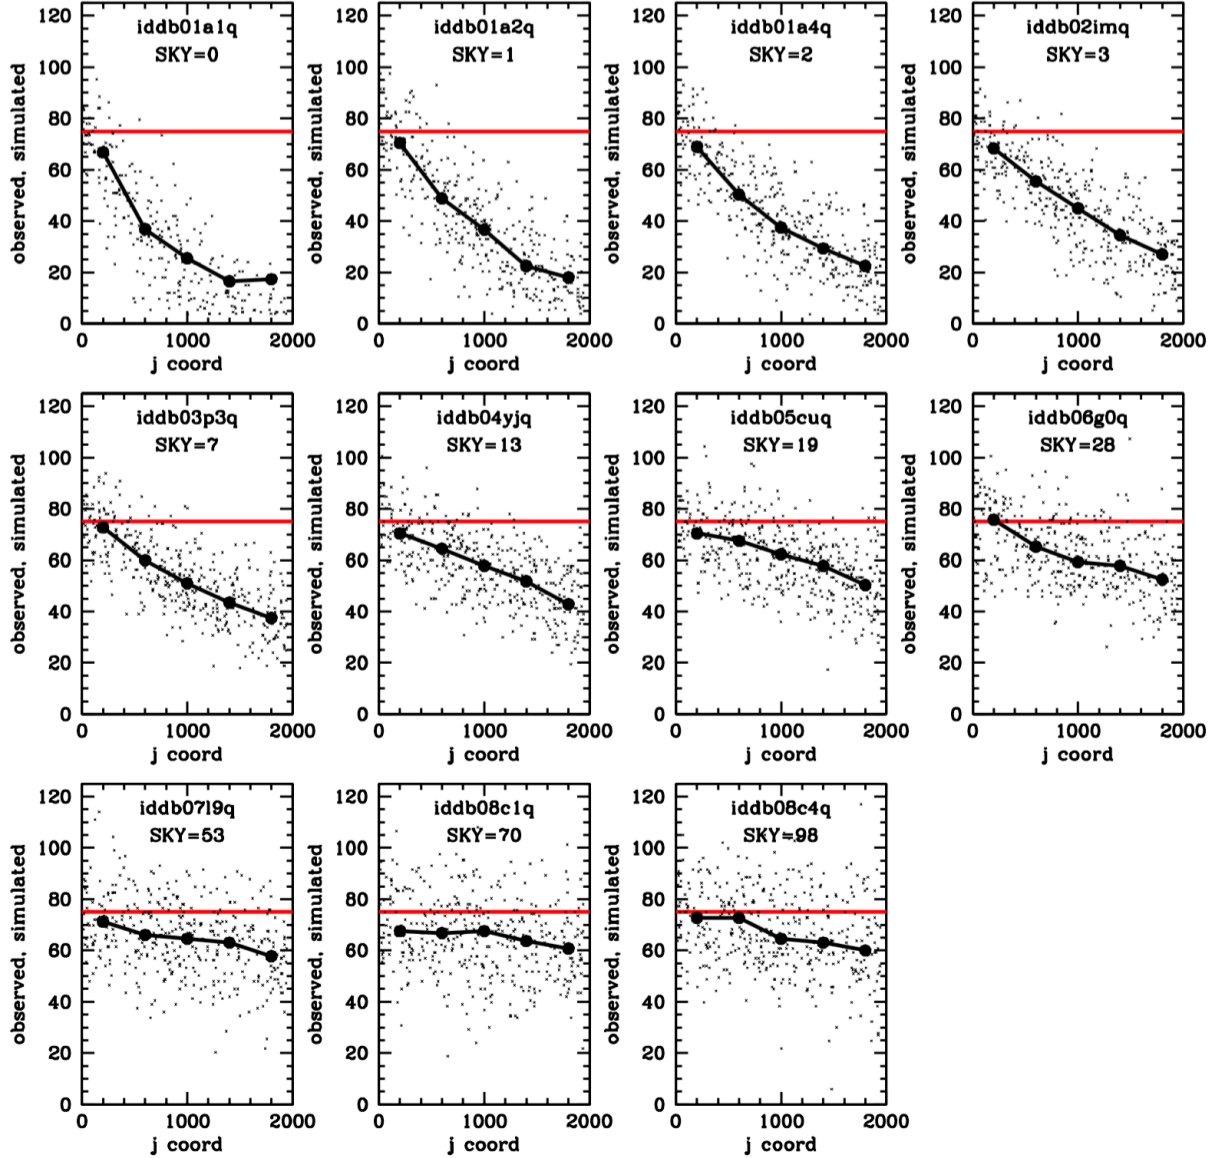
Figure 3: CTE-related losses as a function of number of parallel shifts for a WP that started with 75 electrons above the background. Each panel shows the losses for a different 'sky' (i.e., post-flash) background. The connected filled circles are an empirical binned average of the individual dots. The data in this plot were taken in late 2016.

4.3 Examining the Direct Losses

Now that we have an estimate of the true, intrinsic flux per second in each warm pixel and the effective darktime of each exposure, we can examine the observed counts in each warm pixel as a function of both the number of transfers and background to directly examine CTE losses. **Figure 3** shows such a plot.

The various panels show how many counts survive for a WP that starts with 75 electrons and is transferred by various numbers of parallel shifts (*j*). Each panel shows the results for different exposures at various post-flashed sky backgrounds. We selected those exposures that probed sky

11

(i.e. post-flash) backgrounds at the fiducial levels where our cumulative $\phi(q)$ model could be specified. The beneficial aspect of this analysis is that it is *completely* independent of the trap-release profile, $\tau(\Delta j; q)$. The small black points are the individual raw observations of surviving WPs. The black curve averages the observations in 400-pixel-wide bins. It is clear that until the sky background is greater than 13 e⁻, even a WP with 75 electrons will lose *most* of its charge in its trip down the detector.

The goal is to specify the CTE model such that it agrees with these observations. Doing a model-observation comparison for this direct pinning is different from doing as we did in **Figure 1**, where we examined subtracted trails. Here, we want to examine the *absolute* losses, so we need to simulate the observations.

To do this, we took the warm-pixels list generated in **Section 4.1** and scaled them down according to the dark-time ratios in **Section 4.2**. We then added them to background-only versions of the post-flashed images. This allowed us to add the artificial WPs into the same environment that their real-world analogs were generated in. We then ran these model exposures through the forward-model readout simulation and produced simulated observations.

It is worth noting that when we add a WP of 75 e⁻ to a background of s electrons, then we are not just exploring the losses at 75 electrons, but rather we are exploring the losses between 75+s electrons and s electrons. Furthermore, as the WP loses counts down the detector, we then start exploring the losses between some intermediate number (less than 75) and the background s. So, we cannot simply tweak the model at a single charge level to fit the observations. This is one of the many complications involved in solving for the CTE model.
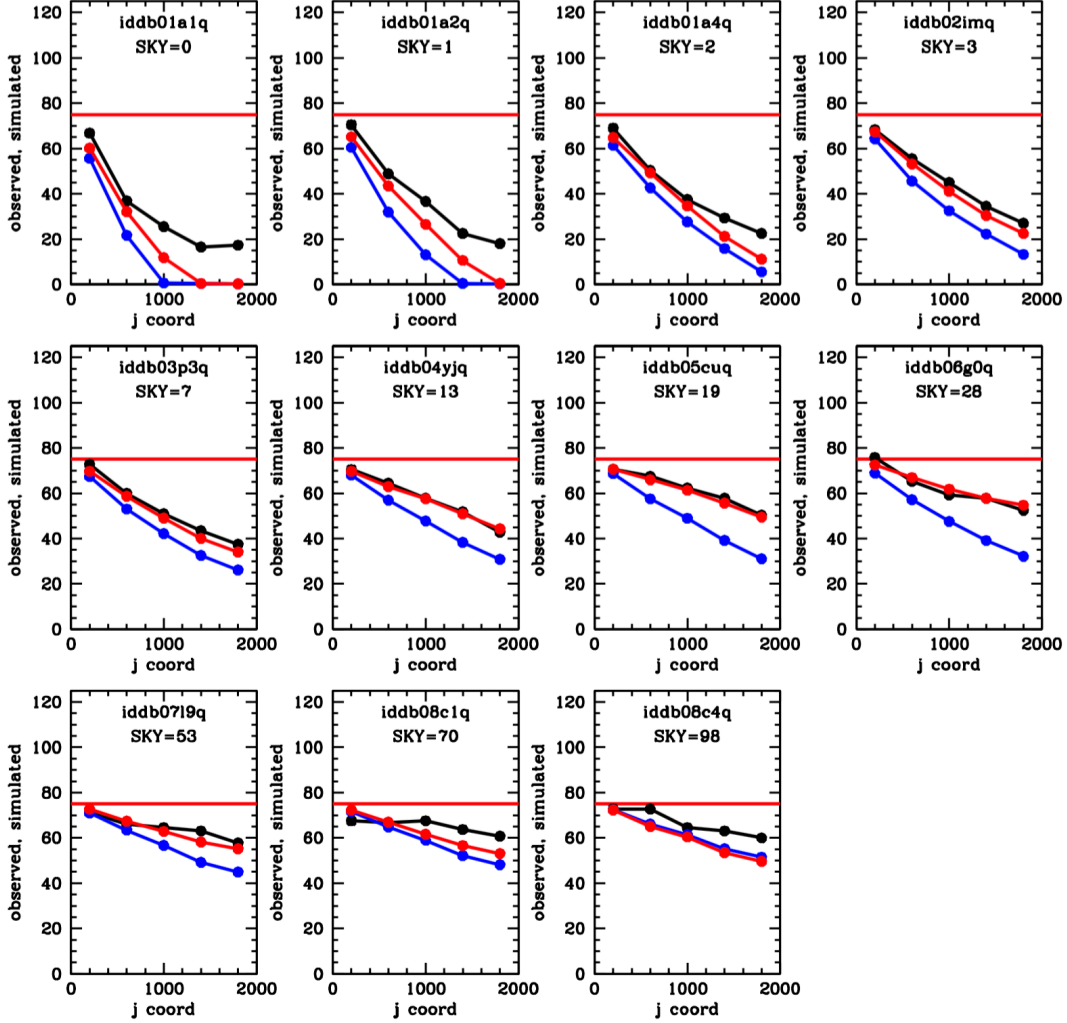
Figure 4: The average empirical losses seen in Figure 3 (black) compared with the v1.0 model (blue) and with the improved v2.0 model (red).

4.4 Re-pinning the Model

We started with the original v1.0 model. The blue curves in **Figure 4** compare the losses predicted from the model with the short-dark observations. The original v1.0 model clearly over-predicts losses at the faint end.

To remedy this, we tweaked the v1.0 fiducial values of $\phi(q)$ at the fiducial locations from $q \sim 1$ up to $q \sim 100$ electrons in an effort to better match the model to the observations. The red curve shows the best fit. The new model does an extremely good job for backgrounds between 3 and 70 electrons. Below 3 electrons, the model over-predicts losses. It is hard to model things perfectly at this low background level, since it is very hard to know the true background in a given pixel, given the readnoise of 3.25 e⁻. Above 70 e⁻, the model also is not perfect, but the total packet size is 70 e⁻+75 e⁻ = 145 e⁻, which is beyond the faint-end regime we are exploring here.

13

4.5 Revisiting the Shape of the Faint WP Trails

The original model did not include any variation in the trail shape. In our initial 2012 study based on the long darks, we did not notice any change in trail shape. Such a change would have been hard to detect, since there is a degeneracy when the trails themselves are used to measure the losses. The "direct" approach employed here for v2.0, enabled by the taking of post-flashed short darks, allows us to break this degeneracy. The variety of backgrounds we have available also helps us to examine the trails from faint WPs.

One complication of examining trails from faint WPs is that they are even fainter than the WPs themselves. In addition, the trails get fainter the farther downstream from the WP we go (Δj), so it is particularly difficult to assess their shape and to know how much is in the entire trail. Finally, the trails themselves tend to have few electrons in them and, as such, they are susceptible to getting trailed themselves. All of these things make it hard to accurately measure the trailing.

In an effort to examine the "pure" trails from faint WPs, we extracted the trail profiles for medium-brightness WPs on zero sky that were very close to the readout register ($250 < j < 750$), so that the self-trailing of the trails could be minimized. Since the same investigation was done for ACS the last time its CTE model was re-pinned, we include the ACS results in **Figure 5** on the left, with the results for WFC3/UVIS on the right.

The trail shape for ACS is very similar for WPs with an intensity of 20 electrons and for WPs with an intensity of 70 electrons. By contrast, WFC3/UVIS appears to have a profile that drops to zero at $\Delta j \sim 4$ for WPs smaller than 40 electrons. For WPs brighter than this, there appears to be an increasing amount of flux beyond $\Delta j \sim 4$. Whereas the trail for bright WPs has 23% of the released charge in the first pixel, 9% in the second, 7% in the third pixel, 6% in the fourth, and 50% beyond the fifth pixel, we modeled the trail for the traps that grabbed the smallest charge packets as releasing over 50% of their trapped flux in the first pixel, 17% in the second pixel, 11% in the third and *all* of the flux released by 6 pixels. We thus had a two-tier trail model for WFC3, with a sharp trail for charge traps that affect the first ~70 electrons and a broader trail for traps that affect the larger clouds.
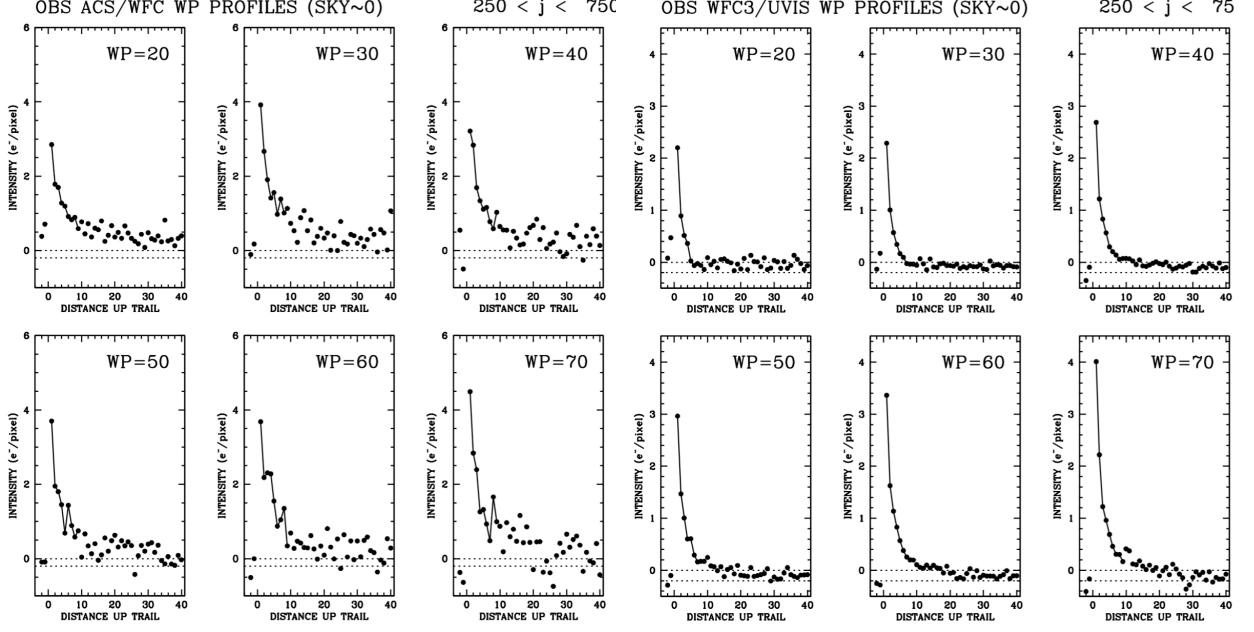
Figure 5: Trail profiles for faint WPs near the readout register, where trailing of the trails could be minimized, for ACS (left) and WFC3/UVIS (right).

4.6 Adjusting the Brighter End of the Model

Now that we have re-determined the parameters for the new CTE model for the first 90 electrons in a packet, we can focus on the model for larger packets. For these, it is possible to use the familiar "indirect" strategy of examining the trail and iteratively converging on a trail shape.

Figure 1 showed what the trails behind the WPs looked like for the v1.0 CTE parameters for WPs that are far from the readout register. The trails behind WPs that had between 300 and 10000 electrons looked good, but at the faint and bright ends ($q \sim 100$ e$^-$ and $q \sim 30000$ e$^-$) there were clear problems.

To converge on the new model, we adjusted $\phi(100)$ until the blue subtracted trail behind the upper-left panel was flat. Then we turned our attention to $\phi(300)$ and the lower-left panel. We continued until all the panels were as flat as flat as possible. We then made small adjustments to the trail profile when it was clear that all the subtracted trails showed similar residuals. It took about eight iterations between adjusting $\phi(q)$ at various points, and then running the reverse model to see how the new model improved the WP trail subtraction.
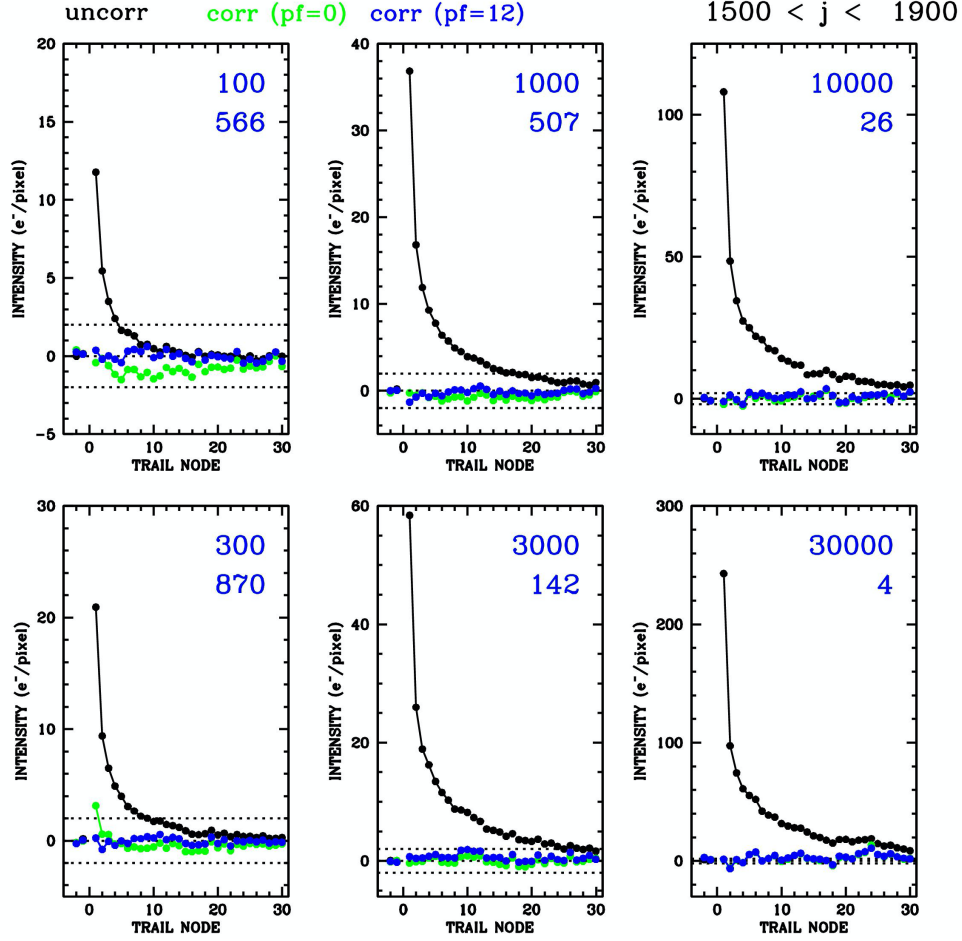
Figure 6: Efficacy of the new v2.0 UVIS model. The black curve shows the uncorrected (PF=12 e⁻) trails behind six fiducial WPs. As with Figure 1, the green curve shows the trails after running the pixel-based reconstruction on the unflashed darks. The blue curve shows the same for the PF=12 e⁻ darks. Note the improvement at the faint end and at the bright end.

Figure 6 shows the results for the new model. As before, the black curve shows the uncorrected trails behind six fiducial WPs: 100, 300, 1000, 3000, 10000, and 30000 e⁻. The blue curve shows the trails behind the corrected PF = 12 e⁻ exposures, and the green curve shows the trails behind the corrected PF = 0 e⁻ exposures. The PF = 0 e⁻ curves are shown just for reference; a minimum image background of 12e- was recommended until June 2020 (when it was increased to 20e-).

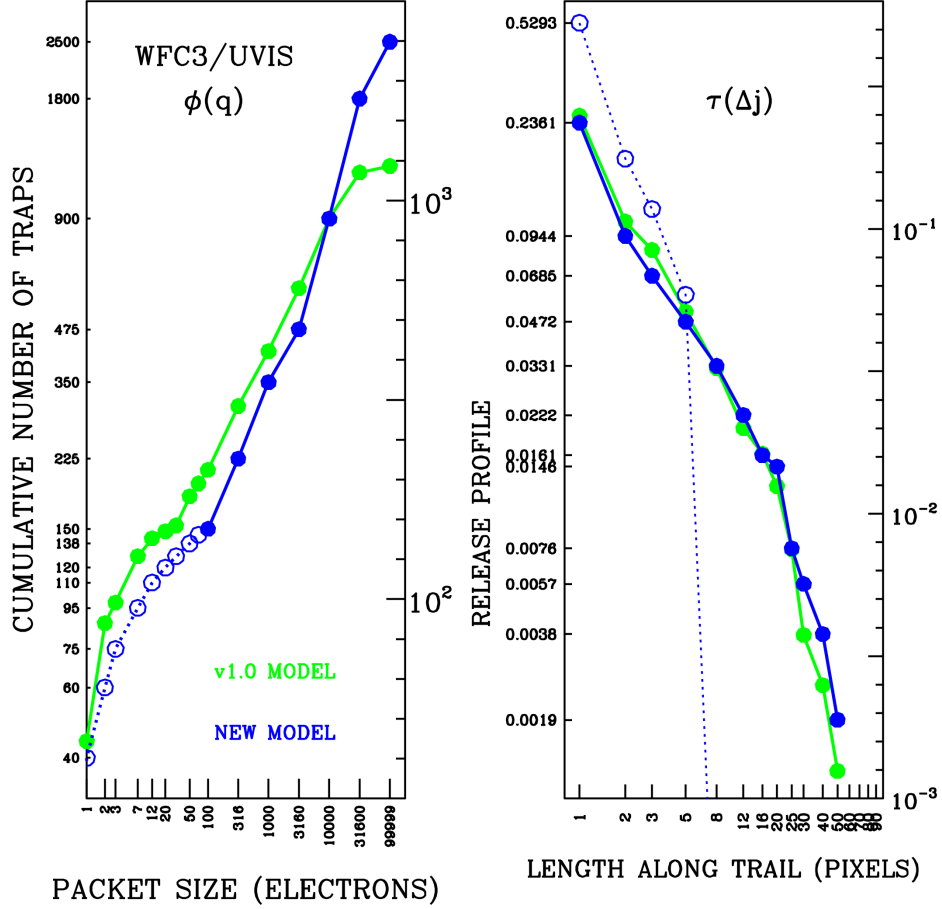**Figure 7:** **The parameters of the v2.0 WFC3/UVIS model. On the left, we show the cumulative number of traps ϕ(*q*) against packet size *q* and, on the right, the trap profile. The filled circles show the bright end and the open circles the faint end (*q* < 70 electrons). The green and blue curves trace the parameters for the v1.0 and v2.0 models, respectively.**

4.7 The New CTE Model

In **Figure 7**, we show the converged-upon parameters ϕ(*q*) and τ(*q*;Δ*j*) for the new v2.0 WFC3/UVIS CTE model, compared to those for the v1.0 model. The new model has more traps at the bright end and fewer traps—and a different shape—at the faint end. The old model inferred twice the actual losses at the faint end because it (1) treated the trail as the same for all packets and (2) relied on the first few pixels of the trail to infer the losses. Thus, the good news is that the losses are not as severe as predicted by the old model.

17
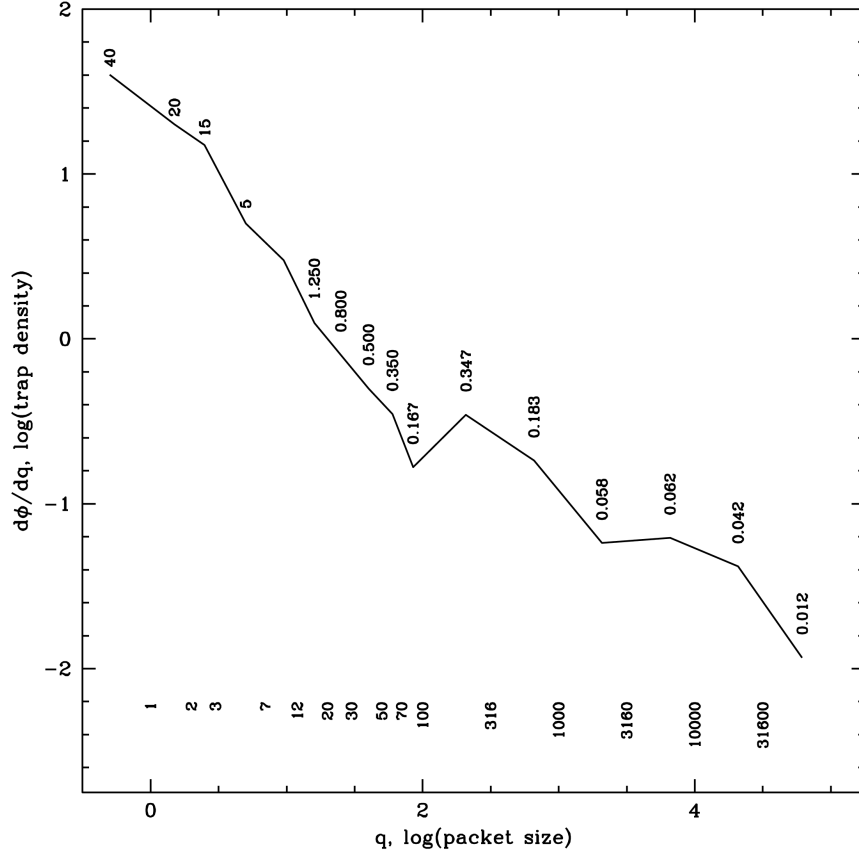
Figure 8: The number of traps per column per marginal electron for various charge-packet sizes for the new v2.0 model.

Figure 8 shows the model parameters in a somewhat more intuitive way. The curve shows the number of traps per column that will impact a marginal electron as a function of charge-packet size, based on the warm pixel analysis. One cannot use this curve directly to correct observed CTE losses for stars, since there is some self-shielding involved for stars. Nevertheless, it is clear that there are just as many traps in a 2048-pixel column as there are source electrons for packets of size $q \sim 20$ e$^-$.

The dip we see at around 70 e$^-$ may or may not be real. The previous v1.0 model had an apparent dip between 12 e$^-$ and 30 e$^-$ that turned out not to be real when we took more data that better probed those background levels. We need to take more data to sample the loss curve better between 30 e$^-$ and 100 e$^-$. In CAL-13567, we were focused on the losses around 12 electrons, since that was the recommended total background level at the time. Calibration program 16440 (PI-Anderson) should provide improved insight into the behavior of the loss curve beyond 30e-.
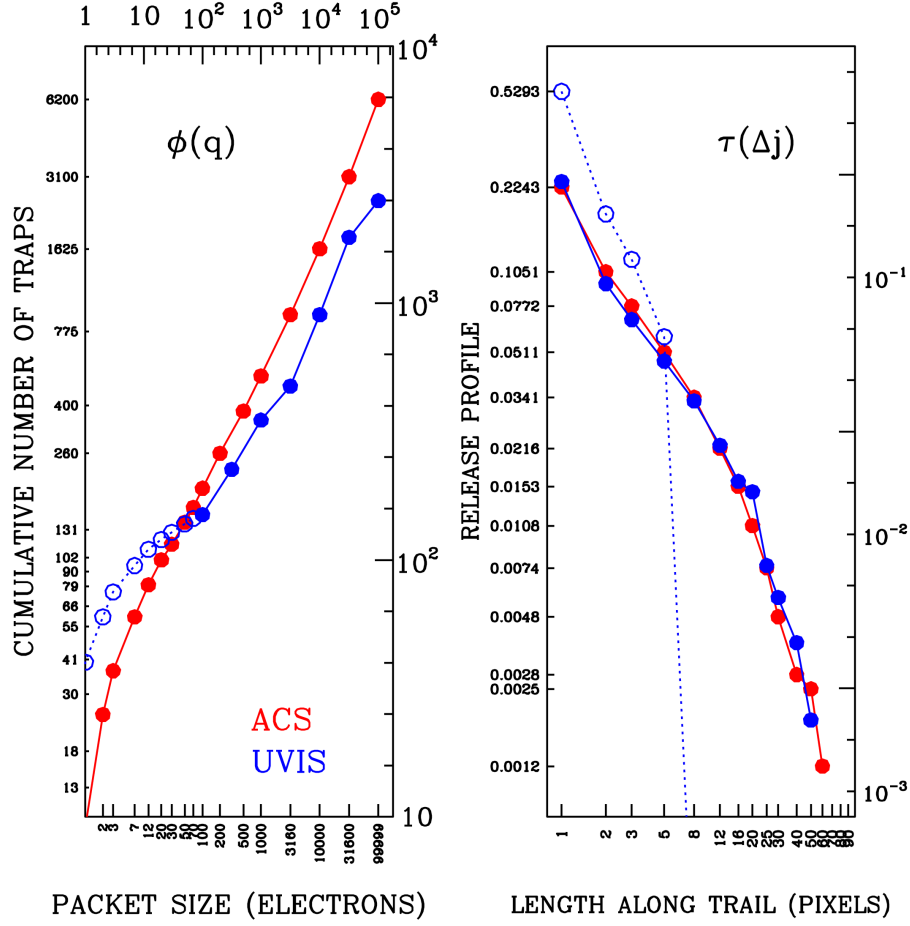
Figure 9: Comparison between the WFC3/UVIS and ACS models, in blue and red, respectively. The left plot shows φ(*q*) as a function of the packet size (*q*) and the right plot shows the trail profile. The dotted portion of the UVIS profile corresponds to *q* < 70 e⁻.

ACS had its model re-pinned at about the same time as the WFC3 data were taken (see Anderson & Ryon 2019). The two models are essentially the same in terms of their inner workings, so it is easy to compare their parameters. In **Figure 9,** the left panel shows the cumulative number of traps as a function of packet size and the right panel the trap profile, with ACS parameters in red and the WFC3 parameters in blue.

The detectors of the two instruments show similar cumulative-trap slopes at the bright end, with ACS exhibiting about 75% higher losses — not unreasonable given the 15 years ACS has spent in orbit compared to the 8 years UVIS has been in orbit (as of 2017, when the model-pinning data were taken). However, the two detectors show quite different behaviors at the faint end ($q <$ 70 e-), where UVIS *already* has more losses than ACS/WFC. It is possible this difference is related to the "mini channel" that was manufactured into WFC3 to help improve CTE for faint sources, although it runs contrary to what one might have expected.

5.   The New v2.0 Model: Avoiding Noise Amplification

In 2012 when we formulated the v 1.0 CTE model for WFC3/UVIS and recommended a minimum image background of 12 e⁻ (dark + sky + post-flash, if needed), the CTE losses at this background level were about 15%. The fact that the losses were a small perturbation on top of the surviving scene meant that the pixel-based correction could be done accurately. The only complication was the readnoise.

5.1 Initial Strategy

Readnoise is added after the charge-transfer process has already taken place. As such, it is added to the already blurred image. If we blindly try to unblur the image (which is a deconvolution-type process) with the readnoise present, then the algorithm will amplify the noise.

In order to avoid the amplification of readnoise, in version 1.0 we took the original image **O** and attempted to determine the smoothest possible image **S** that was consistent with being **O** + **R**, where image **R** is consistent with being pure readnoise (a Gaussian distribution with a mean of zero and a sigma of 3.25 e⁻). We then performed the correction on this smoothest-possible image **S** to get **S′** (the pixel-corrected smooth image) and associated the corrected observed image **O′** with **O** + (**S′** − **S**). The thinking was that if the correction is performed on the smoothest possible image, then we will get the minimum possible correction that is consistent with the data — a good conservative goal.

5.2 How the Initial Strategy Is Working Now

The v1.0 approach worked well for several years. However, now that at typical sky background levels the CTE correction is much larger than merely a perturbation, we find that the v1.0 approach amplifies readnoise significantly.

Figure 10 shows a subsection from an (uncorrected) `flt` image on the left and a v1.0 CTE-corrected `flc` image on the right. The exposure is a short 30s dark, so there should be only background (with Poisson and readnoise), along with a few warm pixels, sink pixels, and cosmic rays. The `flt` image on the left does indeed appear to be largely noise with a couple obvious warm pixels and a few cosmic rays. The image on the right, however, has a large number of high and low pixels, arising from seemingly unremarkable locations in the original uncorrected source image (the `flt`).

The histogram below the plots shows the distribution of pixels relative to the local background (which is 20 electrons, mostly from post-flash). The blue histogram shows the distribution for the `flt` image, which has a sigma of 5.66 electrons, largely consistent with being readnoise (3.25 electrons) plus Poisson noise ($\sqrt{20}$ electrons) added in quadrature. About 0.5% of the points are beyond four sigma.

The green histogram shows the local pixel-to-pixel variations in the `flc` image from the v1.0 correction. The pixel-to-pixel distribution now has a sigma of 6.31 e⁻, which is characteristic of an image with a background of 30 electrons. There are two times as many 2-sigma pixels and *nine* times as many 3-sigma pixels as one would expect in a normal distribution.
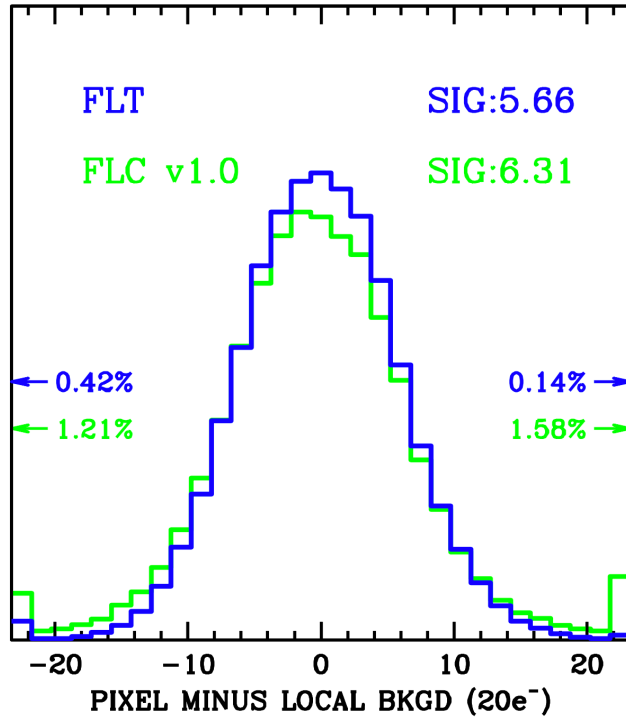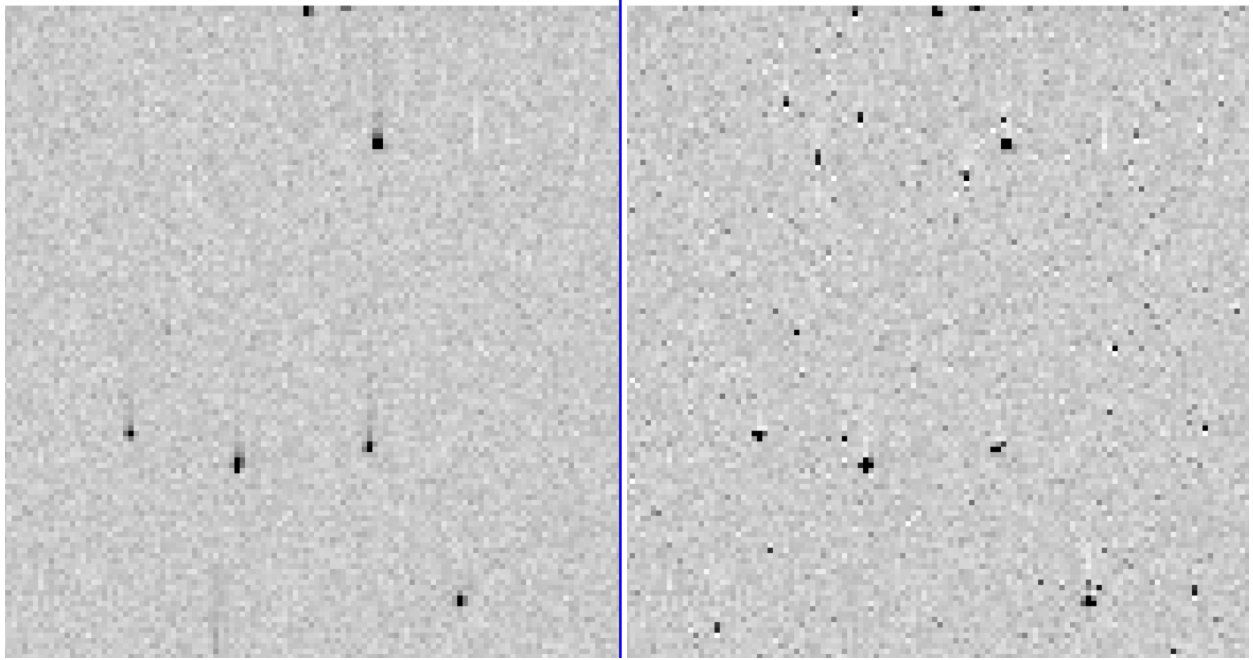
Figure 10: A short dark frame (`iehq19kyq`) taken in December 2020 with a postflash of 20 electrons. Upper left: the `flt`; upper right: the `flc` generated by the v1.0 pixel-based correction, showing the same image subsection near the top of the detector, ~2000 parallel shifts from the readout register where CTE losses are largest. Bottom: histogram of the distribution of pixel values about the local ~20 e⁻ background for the two images.
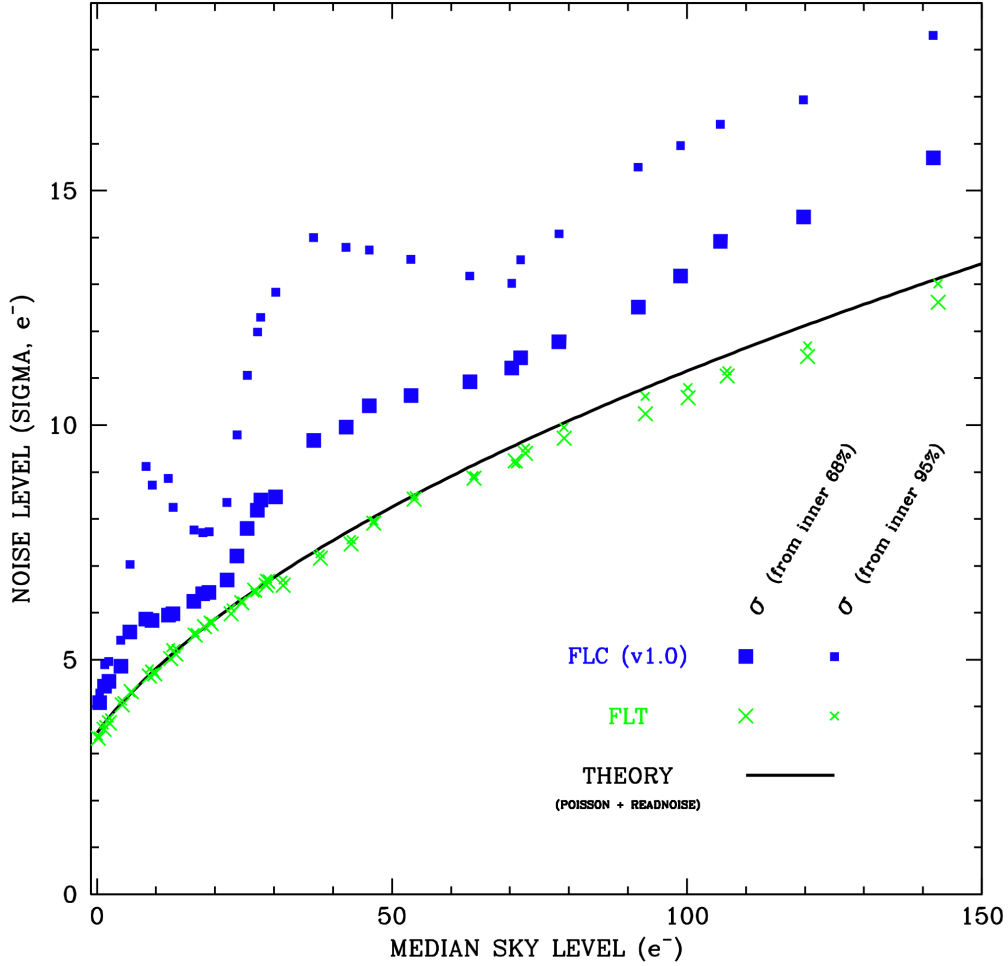
Figure 11: **The noise measured in `flt` images (green points) and v1.0 `flc` images (blue points); the theoretical level (Poisson+read noise) is the black line. Symbol size denotes the percentile region within which sigma is measured: from the inner 68% (large) and by the inner 95% (small).**

What is even more troubling than the increased noise distribution is that the change to the background is not just a matter of increasing the noise, but it dramatically changes the *properties* of the noise. The noise distribution becomes asymmetric: in the deconvolved image, the positive noise excursions are fewer but greater than the negative excursions[2] (see also the histogram in **Figure 10**).

This asymmetric noise can have deleterious effects on many analysis strategies. For instance, if we use sigma-clipping to arrive at a sky value, it will tend to reject the few high-sigma pixels, but keep the more low-sigma low ones, thus biasing the sky downwards, which could bias measured-flux results upwards.

---

[2] This makes sense when we consider that a positive noise excursion will be modeled similar to a warm pixel: if the pixel is 5 electrons higher than its neighbors in the observed image that is read out, then the corrected image will perhaps have the pixel ~10 electrons high, and the model will pair the positive excursion pixel with a negative-excursion trail: it will take out 3 electrons from the first upstream pixel and one from the following two, etc. This results in a few large positive excursions and more small negative ones, thus resulting in asymmetric noise.

The increased noise in the v1.0 `flc` images is not just an artifact of one particular sky level. It is present at all sky levels. **Figure 11** shows the noise present in images taken with various backgrounds for the `flt` images and for the v1.0 `flc` images taken as a part of CAL-16440 (PI-Anderson) in December 2020. We examine the pixels between j=1975 and j=2025 (far from the readout register) and characterize the noise by first subtracting from each pixel the local background then quantifying the noise in two ways. The large points show sigma as determined from the inner 68% of the distribution (which for a Gaussian distribution should be $2 \times \sigma$) and the small points show sigma as determined from the inner 95% of the distribution (which for a Gaussian distribution should be $4 \times \sigma$).

The `flt` image shows that the observed noise is consistent with theoretical estimates from Poisson noise from the background plus a readnoise of 3.25 e⁻. The noise in the v1.0-corrected flc images is considerably higher than this. The increase in noise due to the pixel-based correction is much worse for the highest and lowest 5% of the pixels, but it is still quite bad for the pixels in the bulk of the distribution. Note that the dip in the errors at a background of around 20 e⁻ is an artifact of the v1.0 model, which was based on limited data that seemed to show a paucity of traps between q=12 electrons and q=36 electrons (see the green curve in **Figure 7**). As it turns out, there is no such plateau (see the blue curve in **Figure 7**). Therefore, the seeming dip of noise amplification at sky ~ 20 e⁻ in **Figure 11** is illusory: with an accurate model, there will be considerable noise amplification everywhere.

When we devised the v2.0 correction, we examined the same plot and found that the amplification was only very slightly improved. The reconstructed `flc` images were still much noisier than the `flt` images.

5.3 Searching for a Better Strategy

One of the WFC3/UVIS detector's main advantages is its low noise (low dark current, low readnoise, etc.), so it was unacceptable to allow such noise amplification. Therefore, we explored several other algorithms to prevent the amplification. Unfortunately, the only algorithms we could find that prevented the amplification involved allowing for much more noise than readnoise alone could account for. This was frustrating because the more allowance we must make for noise, the less restoration we can provide for faint levels of charge.

The need to allow for more than 3.25 electrons of noise added during readout ran counter to our long-standing basic assumptions. These assumptions were that once the shutter is closed and post-flash is added, the pixels contain a realization of the scene that freezes in whatever Poisson-type processes generated it (source noise, background noise, dark-current noise, post-flash noise, etc.). Under these assumptions, the parallel-shuffle readout process ends up blurring this scene according to the trapping and releasing that takes place as electrons are shuffled down the detector. The image background uniformly fills traps below a certain level, and pixel levels above this level get preferentially trapped and released into downstream pixels. The pixel-based CTE model we have created treats this as a smooth, deterministic process.
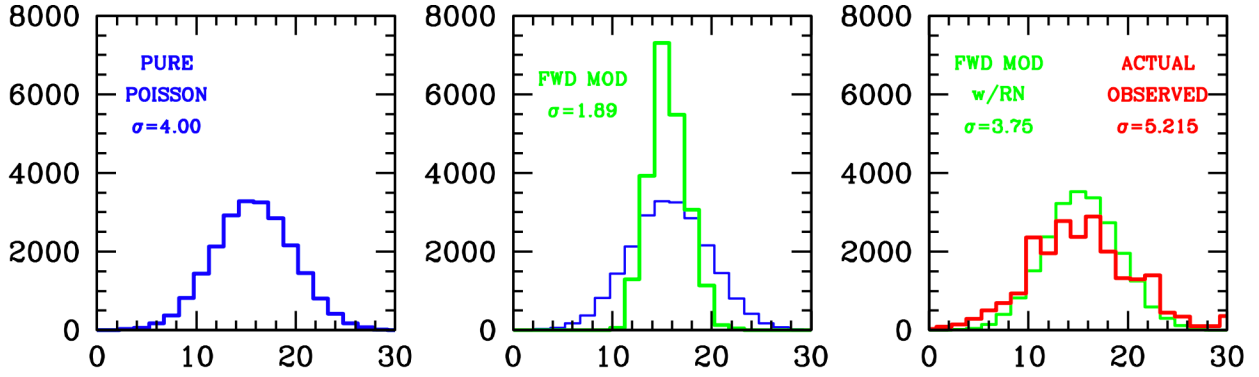
Figure 12 (Left panel) The distribution of pixel values in a simulated image with 16 e- background starting with pure Poisson noise at the top of the detector. (Middle panel) The same simulation after parallel-shifting all the way down the detector. (Right panel) The forward-model distribution with read-noise added. The red curve shows an actual observed image with the same 16 e⁻ background, which has considerably more noise than our simplistic model would predict.

Figure 12 shows a simulation we performed to explore the behavior of our model on an image full of background noise. The left panel shows an image of pure 16-e⁻ background, which has an RMS of exactly 4 e⁻. This image is then pushed through the forward model to yield the green histogram in the middle panel (the original blue histogram is shown for reference). In line with our expectations, the noise peaks are cut down by the traps, the noise troughs are filled in by the trails, and the RMS goes down by more than half, to 1.89 e⁻. Once read out, this would then have readnoise added to it, which would increase the apparent sigma to 3.75 e⁻ (shown as green in the third panel). We can compare this to what we see in an actual WFC3/UVIS image that has a background of 16 e⁻ (shown in red).

Somehow, the parallel shuffle process did not result in reduced pixel-to-pixel variation, rather it is similar to what we would get if we just took the original Poisson noise and added readnoise: $4.00^2 + 3.25^2 = 5.15^2$. Our model appears to be missing something fundamental about the readout process.

We know from examination of the impact of CTE on warm pixels in **Figure 3** and **Figure 5** that the model does indeed clip electrons off of high pixels and deposit them in trails. But it seems that the process is not as smooth and deterministic as our model assumes.

Surely, at some level, we know that it cannot be smooth, since nature cannot deal with fractional electrons, but only whole ones. Also, each pixel does not in reality contain an identical distribution of fractional traps, as our model assumes. In fact, the sink-pixel studies tell us that traps are localized in a relatively small fraction of the pixels. There is no way that we could map the location of every trap in every detector pixel — and even if we could do that today, tomorrow there would be new traps in different pixels. And, finally, even if we could have a map of all the traps and which particular electron they grab in a cloud, then we still could not deterministically predict the readout process: the probabilistic trap and release will always be a random process.
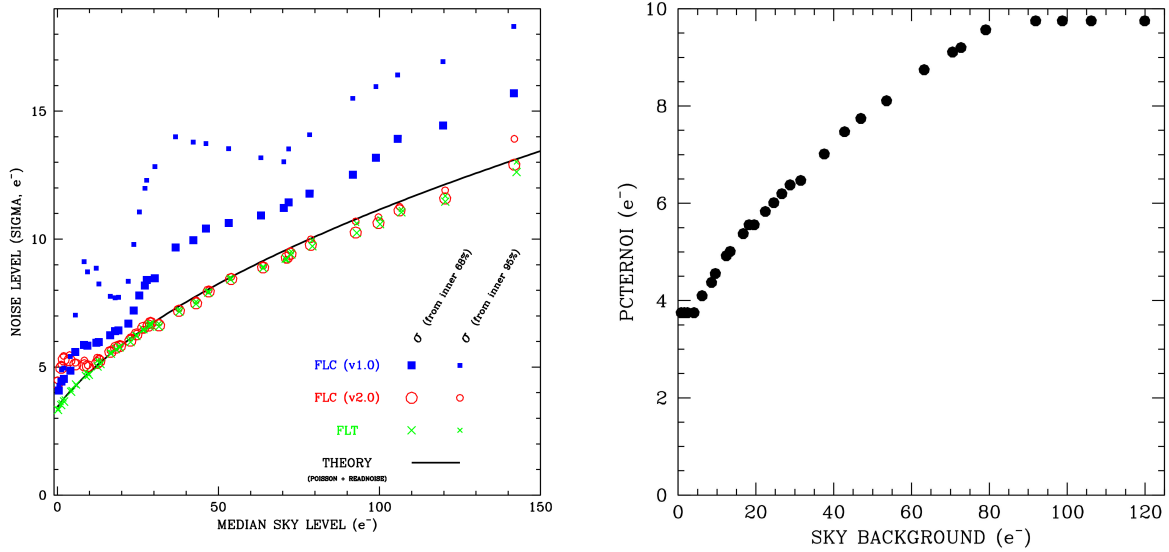
Figure 13: **(Left) Same as Figure 11, but with the `flc` from the v2.0 pixel-based correction added in red circles. (Right) The resulting pixel-to-pixel noise allowance.**

5.4 Living with Fundamental Limitations

The smooth, deterministic model is the best way we have to approximate what is going on in the detector, but it is clearly limited. And we are starting to bump up against these fundamental limits. There is no way to directly correct the image for the charge-transfer process. The charge-transfer process itself adds noise along the way, with a result that is similar to the noise that gets added at the readout amplifier. When the level of added noise is comparable to the CTE correction itself, the process is clearly not going to be reversible.

Given this, we decided to do the best we could. Our fundamental goal was to "do no harm". We should do nothing that will increase the noise significantly, since that will not help us find and measure faint sources. This requires us to increase the noise-amplification-mitigation levels beyond the 3.25 electrons of readnoise to a noise level that corresponds to a background where losses are indeed at the perturbation level and hence can be deconvolved without adding much noise.

The new procedure is to evaluate the pixel-to-pixel noise in the background, and if it is less than 10 electrons, to set the noise-mitigation keyword (PCTERNOI) at that level. When the reconstruction algorithm recommends an adjustment in a pixel value lower than this level, then this adjustment will be damped. The net result is that the pixel-to-pixel variations smaller than the background fluctuations (or 10 electrons, whichever is smaller), will be largely untouched by the reconstruction algorithm.

Figure 13 shows the run of observed noise as a function of background with the new v2.0 `flc` images added in (red). It is clear that there is essentially no amplification of noise up beyond a background of 100 electrons. This goes for the bulk of the pixels, and even for the 5% most extreme pixels.

The plot on the right shows the solved-for value of PCTERNOI, the effective "readnoise" allowed for in the pixel-based CTE reconstruction. The minimum value used for PCTERNOI is 3.75 electrons, and the maximum value is 9.75 electrons.

Users can override this decision to mitigate the background noise. By default, the raw images are populated with the PCTERNOI keyword set to 0.00 electrons. If calwf3 encounters this null value for the keyword, it will solve for it and populate it with the measured value from the background, and will use that value in the restoration. If, however, a user pre-populates the PCTERNOI keyword in the image header with a number between 3.75 and 9.75 electrons, then CAL/WF3 will use that value. So, if users want to minimize the mitigation and suffer the noise-amplification consequences, they are able to do that.

It is clear that we have entered a different regime for the pixel-based CTE correction. With the v2.0 correction, we have tuned the MAST pipeline default correction such that it will almost never amplify the noise. The penalty is that the faintest sources will not be corrected and users will have to use other strategies to correct the faintest measurements. We will discuss these strategies in the following section. At least, though, the signal from faint sources will not be further hidden in added noise.

## 6. Recent Data Showing Current Losses

When we began the work of re-pinning the CTE model, we realized that the long-recommended 12-e$^-$ level for post-flash was no longer providing the protection that it used to provide. We took some additional data in 2019 (CAL-16029) and 2020 (CAL-16440 and CAL-16441) to investigate the optimal level for post-flash. In the meantime, the noise noise-amplification issue had worsened.

6.1 Losses in 2021, According to the WPs

CAL-16029 and CAL-16440 were internal calibration programs that provided data for the direct short-dark/long-dark strategy to determine absolute losses, similar to the CAL-14880 data that was used to re-pin the model (see **Figures 2**, **3**, and **4**). The original focus of the CAL-14880 data was to have the background levels around 12 e$^-$, since that was presumed to be the sweet spot based on our understanding of the mini-channel. Later, when CTE losses for sources on 12 e$^-$ of background became significant, we began to explore whether higher post-flash levels might provide more protection from losses. As a consequence, we explored a broader range of more finely spaced background levels.

CAL-16440 was taken in December 2020. The program contains a set of eight 900s "long" darks with a background of ~100 electrons. The high background in the long darks allows us to measure the intensity of the WPs without significant CTE-correction uncertainty.
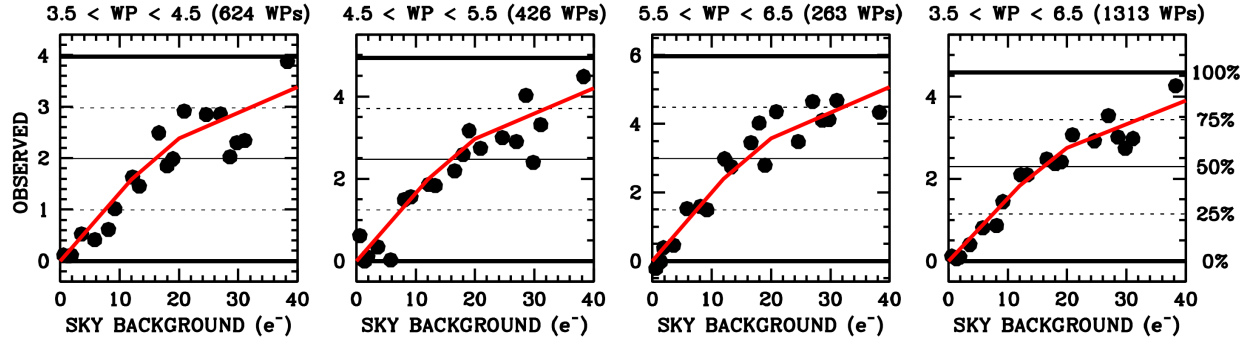
Figure 14: The marginal losses as a function of sky background from faint WPs (levels are noted in the title bar of each panel).

As we did in **Section 4**, we made a list of 36,325 warm pixels in the long darks and divided their intensity by a factor of ~30 to estimate each WP's starting intensity in the nineteen 30s "short" darks that were taken with a post-flash background between 0 and 40 electrons. A WP with 150 electrons in the long dark should have ~5 electrons in the short dark. These low-intensity ~5-e⁻ warm pixels allow us to probe marginal losses as a function of background level.

The first three panels of **Figure 14** show what fraction of a 4, 5, and 6 e⁻ warm pixel in the upper quarter of the WFC3/UVIS detectors survives the 1500+ parallel shifts to the readout. The final panel combines the results from the other three panels. The red line shows the average trend. Basically, if the background is 12 e⁻, the WP now loses roughly 60% of its original electrons. If the background is 20 e⁻, it loses 45%, and if the background is 30 e⁻, it loses about 25%. With losses of ~50% or greater for backgrounds between 12 and 20 e⁻, it is not surprising that the algorithm ends up amplifying noise. Even marginal losses of 25% will not allow a noiseless restoration.

These loss levels are indeed harrowing for users that want to use UVIS to find faint objects. It is worth noting that even if the pixel-based CTE reconstruction algorithm did not have any noise-amplification issues, these losses would still have a significant impact on signal-to-noise. When the restoration algorithm is able to correct the shape and amplitude of sources, it must do so by taking the flux that survives to be read out and then extrapolate that flux to what must have been there in the first place. *Therefore, the resulting S/N is always controlled by the surviving flux, not by the initial flux (or the reconstructed flux), irrespective of any noise that may be added.*

6.2 Losses in 2021, According to the Stars

The analysis up to this point has dealt with warm pixels. Stars are different from warm pixels in that they are not delta functions — they have a width to them on account of the PSF. We would expect that the downstream pixels of a star might protect the upstream pixels from loss such that the net losses for point sources may be considerably lower than **Figure 14** shows for warm pixels.

CAL-16441 is a companion program to CAL-16440. It imaged the center of globular cluster Omega Centauri with a series of short and long exposures, all through the F606W filter. There were four un-dithered long exposures, each 800s in length. The background in these deep
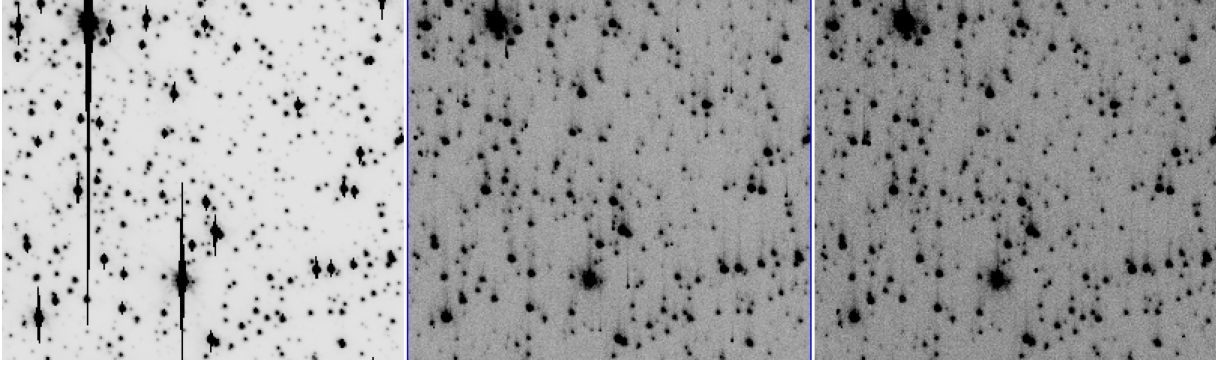
Figure 15: Left: image section of the 800s deep stack of Omega Cen, from the top of the detector where CTE losses are greatest. Middle: the corresponding 4s exposure with a background of 12 electrons. Right: the corresponding 4s exposure with a background of 20 electrons.

images is around 500 electrons, so CTE losses are minimal and the pixels can be easily corrected for CTE blurring via the pixel-based algorithm. The deep `_flc` exposures can be stacked to provide us the "true" astronomical scene.

CAL-16441 also took eight "short" 4s exposures with various amounts of post-flash to achieve average background levels of 12, 14, 16, 18, 20, 24, 27, and 32 electrons. **Figure 15** shows a portion of the deep and short images at the top of the detector, far from the readout register, where CTE losses are greatest.

Stars were identified and measured in the stacked deep exposures. Their fluxes and profiles can be scaled down by a factor of 200 to examine how various faint stars appear in images with various levels of background. For example, a star with 1,500 total counts in the deep exposures will have about 300 counts in its central pixel in the deep exposures. Such a star should have about 15 counts in in its central pixel in the short exposures. **Figure 16** shows the observed profiles for such a faint star on backgrounds from 10 to 30 electrons.

Even with no CTE losses, an individual faint star will have a considerable amount of noise in each exposure. So, in an effort to focus on the average trends, we summed together the images of many similar stars (namely, stars of the same brightness at the same location with respect to the readout).
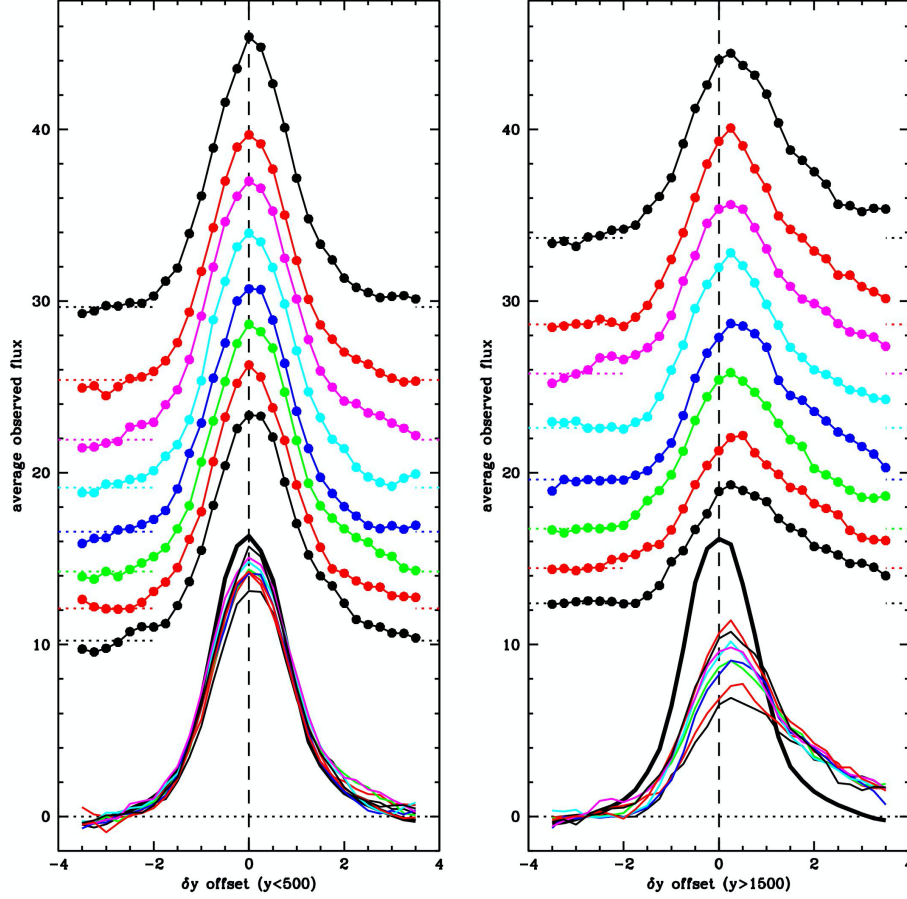
Figure 16 Vertical profile through the central pixel for stars identified in the long exposures to have about 15 electrons in their central pixels in the short exposures. The series of eight dotted profiles show how the star registers in the eight short exposures with different background levels. The left panel shows stars in the 500 pixels closest to the readout, and the right panel shows the star in the 500 pixels farthest from the readout. All the profiles are shown together at the bottom of each panel with sky subtracted. The heavy black profile shows the expectations based on the deep exposure.

The left panel of **Figure 16** shows the stars that are close to the readout (j < 500). These stars are marginally affected by CTE. When the background is 10 e⁻, the star loses only a couple electrons out of the 15 electrons in its central pixel. Yet, even with these small losses, it is clear that the profile is already becoming asymmetric. There is an excess of flux on the right of the profile relative to the expectations based on the deep exposure.

The right panel of **Figure 16** shows the empirical profiles for stars that are far from the readout. These star profiles have gone through over 1500 parallel shifts and their profiles are profoundly impacted. Even with a background of 30 e⁻, stars lose a third of the electrons in their central pixels. When the background is 12 e⁻, central-pixel losses are almost two thirds. The profile also becomes increasingly asymmetric, with flux being lost from downstream pixels (on the left) and deposited in upstream pixels (on the right).   Note that the post-flash background, although stable over time, is not constant across the detector (Biretta & Baggett, 2013). The post-flash tends to be about 15% higher at the top of the detector (near the chip-gap) than at the bottom (close to the readout register).
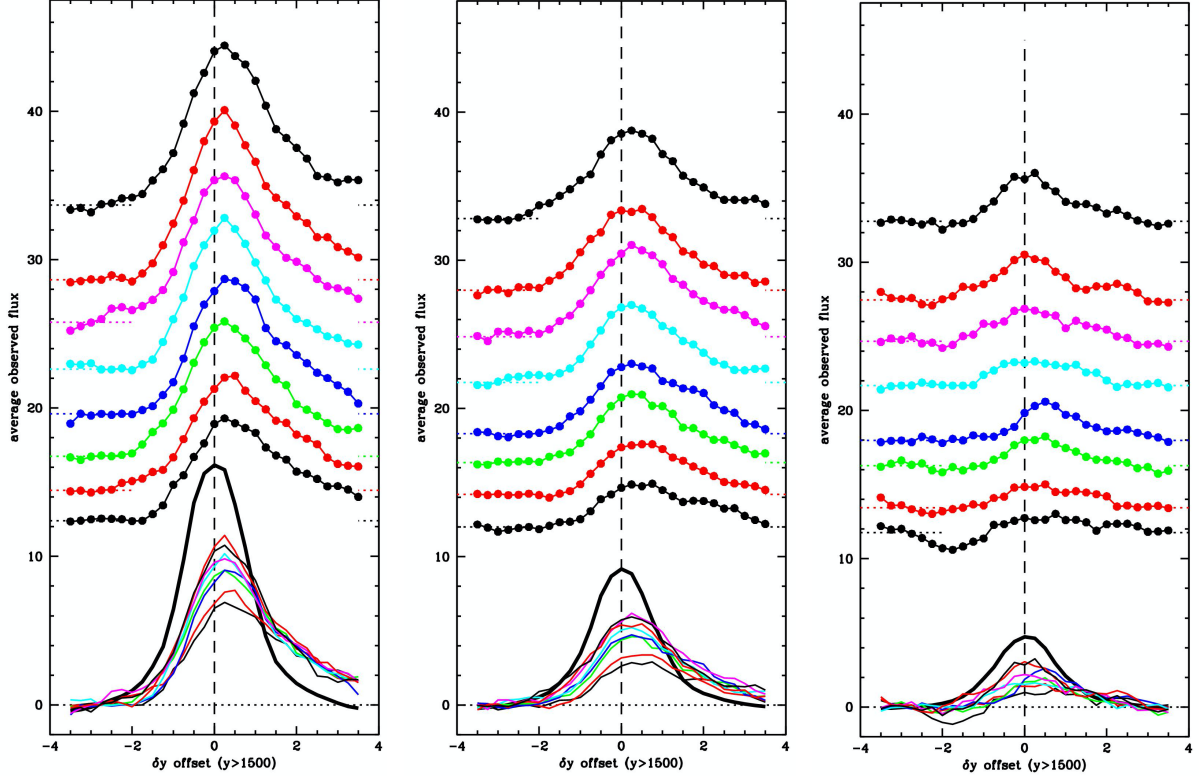
Figure 17: **The rightmost panel in Figure 15 is repeated here on the left, for a star far from readout with 15 electrons in its central pixel. The middle panel shows a source with half the flux in its central pixel (7.5 electrons), and in the rightmost panel, we show a source with half again as much flux (3.75 electrons) in its central pixel. At the bottom, with a sky-subtracted baseline, we show all the curves together along with the profile (black) we expect based on the scaled-down deep image.**

In **Figure 17** we show the same results as in **Figure 16**, but for stars that are 2× and 4× fainter. The faintest stars at the top of the detector (y>1500) nearly disappear at the lowest backgrounds.

It is worth reiterating that the profiles shown in these plots reflect a composite of many, many stars, which has enabled us to sample the faint-star profile every 0.25 pixel. A single star in a single exposure would have only one sample per pixel, amounting to just 5 samples across the vertical Δy range shown here. Each pixel would have a readnoise of 3.25 e⁻ in addition to Poisson noise, so the faintest stars examined here will be marginal S/N~1 detections even without CTE losses. When we include CTE losses, these stars are almost completely wiped out when the background is too low.

The process we have used here to stack the profiles of many stars into a single composite is of course not typical of any science program. However, the process *does* resemble what one might do to find faint objects in a deep set of observations of a field. Typically, in such programs, many long exposures are stacked together to allow extremely low signal-to-noise objects to be found. Such objects cannot be reliably found and measured in individual exposures, so the typical star-based CTE-loss studies (Kuhn and Bajaj 2021) are not able to make an assessment of what to expect. The type of analysis we have used here is the only way to assess losses at the extremely low-signal end.
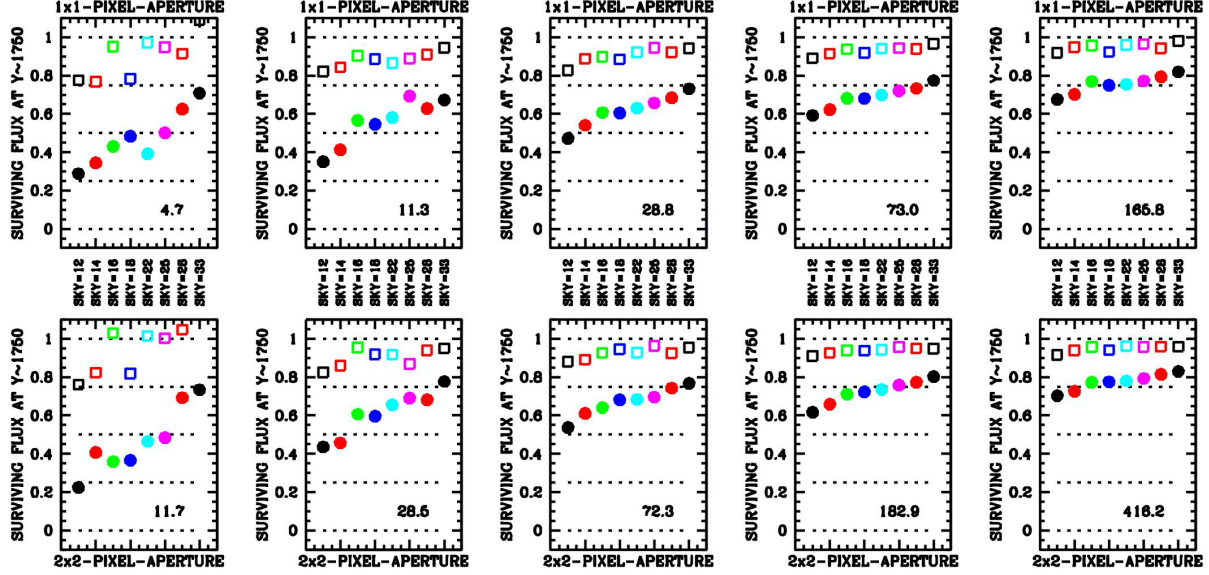
Figure 18: Surviving stellar flux as a function of sky background (similar to the WP analyses shown in Figures 14 and 15. The upper panels show the results for a 1×1-pixel aperture, and the lower panels show the results for a 2×2-pixel aperture. The columns from left to right show the fraction of recovered flux as a function of background sky for five magnitude bins, from faint to moderate flux. The true flux over sky is listed in each plot at the lower right. The 2×2-pixel aperture contains about 2.5 times as much star flux as the 1×1-pixel aperture. The filled circles correspond to stars in the top quarter of the detector (y>1500), farthest from the readout amplifier; and the open squares correspond to stars in the bottom quarter of the detector (y<500), closest to the readout amplifier.

The profiles shown here represent the first chance we have had to evaluate how faint objects on typical backgrounds survive the readout process, whether stars that are close to the background might be shielded from some losses due (perhaps) to upward noise fluctuations in downstream pixels or due to self-shielding. It is now clear from the rightmost panel in **Figure 17** that *upward background noise fluctuations do not significantly shield faint stars from significant CTE losses*.

6.3 Predicting and Correcting Photometry

The vertical profiles shown in **Figures 16** and **17** give a *qualitative* picture of how CTE impacts faint stars. If we hope to use these simulations to estimate the S/N for a prospective source observation —or to correct an observed object for CTE losses — then we need to distill the profile results into the context of photometry.

In **Figure 18**, we show two different ways of measuring these faint sources on the various backgrounds. The first way is to use a simple 1-pixel aperture. This is obviously impractical, but it maximizes signal to noise for faint sources close to the background. We considered the profiles above for sources with total fluxes of 25, 63, 156, 391, and 977 electrons (a spacing of one magnitude). Since **Figure 17** shows that the centroid sometimes gets dragged in the +y direction, we determined for each composite star the Δy offset that would maximize the signal in the aperture. (This is typical of what researchers would do when they find a source of interest:

they don't know its "true" location; they find it where it is peaked.) We then measured the flux-over-sky for the composite star in each magnitude bin. We did this for stars in the lower 500 pixels (open squares) and for the stars in the upper 500 pixels (filled circles). The results are shown in the top row of panels. The average true flux of the stars in each bin is reported in the bottom-right corner of each panel.

The filled circles show that when the faintest stars undergo 1500+ parallel shifts on a background of 12 e⁻, they lose almost 70% of their original counts. More and more flux is preserved as the background increases to 20 e⁻ (50% lost) and 30 e⁻ (35% lost). It is interesting to note that these trends track almost identically with the results in **Figure 14** for the warm pixels. (More on this below in **Section 6.4**.)

Brighter and brighter stars show lower and lower fractional losses. This is not surprising, as the brightest pixel in the brighter stars has an electron cloud that is large enough to experience a lower trap density. The trends with background appear to be very nearly linear, and become flatter as the star gets brighter. *This indicates that losses for bright stars are not as sensitive to the background as are the losses for fainter stars*.

The bottom row of panels shows the results for a 2×2-pixel aperture, which is a more practical way to measure faint stars. These results are quite similar to the results in the top row. This is a good sign that the correction is not particularly dependent on exactly how we measure the star, which should make it easier to come up with a useful correction. The numbers given in the lower right of each panel show the average true number of counts for each star in the 2×2 aperture. This tends to be about 2.5× the number in the upper panel, which is consistent with expectations from the F606W PSF[3], which predicts that about 19% of a star's light will fall in its central pixel when it is centered on a pixel, and about 12% in each central pixel when the star lands on in the corner-intersection of four pixels (12×4 = 48, and 48/19 ~ 2.5).

6.4 Stars versus Warm Pixels

The trends seen in **Figure 17** for the stars are remarkably similar to the trends seen in **Figure 14** for the warm pixels (WPs), even in a quantitative sense. At a background of 12 e⁻ losses in 2021 are ~60%, at 20 e⁻ losses are ~50%, and at 30 e⁻ losses are ~30%. This goes for both stars and WPs.

We expected that stars would show a shallower trend than WPs, on account of the self-shielding they must experience: the leading side of the star should fill in traps that the trailing side does not have to suffer. This can be seen by the loss of flux on the left side and the gain of flux on the right side of the profiles in **Figures 16** and **17**.

Why, then, are we seeing essentially *no* benefit from the shielding? This could be due to the super-short release trails that are present for the lowest electrons ($q < 70$ e⁻). The right side of **Figure 5** shows that the trails for traps that affect small electron packets release over 50% of the trapped flux in the first upstream pixel and over 80% in the first two pixels. For this reason, there is very little difference between stars and warm pixels. If the trap-release length were

---

[3] https://www.stsci.edu/hst/instrumentation/wfc3/data-analysis/psf

longer, then traps would be filled by the downstream pixels and would stay filled, letting some upstream pixels pass through unscathed.

The short release time also likely makes it harder for Poisson noise in the background to provide natural mitigation for low S/N objects. One could imagine that if the trail length were ~10 pixels and the background is 25 e⁻ (with a Poisson noise of 5 e⁻), then some 16% of the pixels downstream might be over 30 e⁻, and thus could provide some shielding for faint stars that are ~1-σ detections in each individual exposure. However, if the trail length is very short, then this natural shielding would be considerably less effective.

Whatever the explanation, we do not see much "self-shielding" for stars. The fractional losses we see for stars appears to be very similar to what we see for warm pixels.

7.   Advice for Pre-Observation Planning

Pre-observation CTE mitigation is critical. Once the observation has been taken, there is no way to improve the signal to noise within a given exposure. Our improved understanding of the 2021 charge-transfer properties of the WFC3/UVIS detector should help us to take better observations.

While it is unfortunate that the low-background observations that once were WFC3/UVIS' strong suit are now severely impacted by CTE losses, that does not mean that nothing can be done. In many cases, it is still possible to tailor observing strategies to preserve enough signal to do good science. There is, however, no one-size-fits-all strategy: an observing plan should be tailored for the specifics of the targets. Observers can adopt one or several of the following strategies.

Option#1: Subarrays

If the target is small, then using a subarray (or using a full frame, but placing the main target at a CTE-optimizing aperture) can minimize the number of transfers the target undergoes and thus cut in half, a quarter, or even an eighth, the CTE losses it would suffer at the top of the detector (near the chip gap of a full frame). Subarrays worth considering are: UVIS2-C1K1C-*SUB* and UVIS2-C512C-*SUB*. One could also use full-frame apertures UVIS2-C1K1C-*CTE* or UVIS2-C512C-*CTE* to image a larger field while focusing CTE-protection on a target. The C512C aperture location places the target at row ~256, which endures about an eighth of a full-frame transfer. Additionally, one could use POS-TARGs to move a particularly small target even closer to the readout amplifier.

Option#2: Fewer Deeper Exposures

If the observing plan requires many orbits of the same field in order to reach the faintest objects, it is best to take as few exposures per orbit as is reasonable. This benefits CTE in several ways.

First, the more signal from the source per exposure, the higher the charge-packet transmission will be during readout. To illustrate this, the bottom panel of **Figure 18** shows that for a 2×2-pixel aperture on a 20-e⁻ background, a faint source at the top of the detector with 12 electrons in

the aperture will lose about 50% of its flux. If the source has 30 electrons, then it will lose 35%, and if it has 75 electrons, it will lose only 30%. *So, the more flux per exposure, the better the transmission rate*.

Second, by taking fewer exposures, more of the background can come from natural noise, thus requiring less added noise from post-flash. Minimizing the noise added will clearly improve the net signal-to-noise achieved. As an added benefit, fewer exposures will similarly minimize the impact of readnoise.

Several programs[4] have even taken the bold step of taking full-orbit exposures in order to maximize signal and minimize added noise. Of course, there are twice as many pixels impacted by cosmic rays (CRs) in full-orbit exposures of ~2600s than in half-orbit exposures of ~1300s.

CRs in CCD images are much more problematic than CRs in CMOS detectors. When a CR strikes a pixel in an up-the-the-ramp CMOS detector, we simply lose a couple of reads, but in general the slope for the pixel can still be accurately determined. Unfortunately, when a CR strikes a CCD pixel, then that pixel is unsalvageable. About 1% of the pixels in a 900s dark exposure are impacted directly by CRs that impart 50 e$^-$ or more. In a ~2600s full-orbit image, this will amount to about 3% of the pixels.

Another downside of CCDs is that, thanks to imperfect CTE, cosmic rays can impact several times more pixels indirectly than directly on account of their trails. A single-pixel CR of 1000 electrons on a background of 20 e$^-$ can lose 300 electrons to CTE and, therefore have 60 electrons deposited in the first upstream pixel, 30 electrons in the next upstream pixel, 20 in the next, etc. Thankfully, the pixel-based correction typically does an excellent job removing these CR CTE trails.

Figure 19 shows the subtraction residuals from trails behind cosmic rays in dark exposures: the new correction is quite effective at removing the trails with net-zero bias. The old correction did a decent job but tended to over-subtract the trails.

Option#3: Add post-flash

The final option available to WFC3/UVIS observation planners is to add post-flash. How much post-flash to add will depend significantly on what brightness of star is being targeted. Stars that are faint lose a larger fraction of their flux when the background is low. At the same time, though, these faint stars are also more affected by background noise. Determining the optimal background will involve a compromise between preserving signal and adding noise.

Stars that are brighter receive less benefit from a higher background, but they are also less impacted by background noise. The optimal background for a project will likely be dictated by the faintest sources of interest.

---

[4] See, for example, GO-15940, GO-15949, GO-16165, GO-16252, and GO-16255.
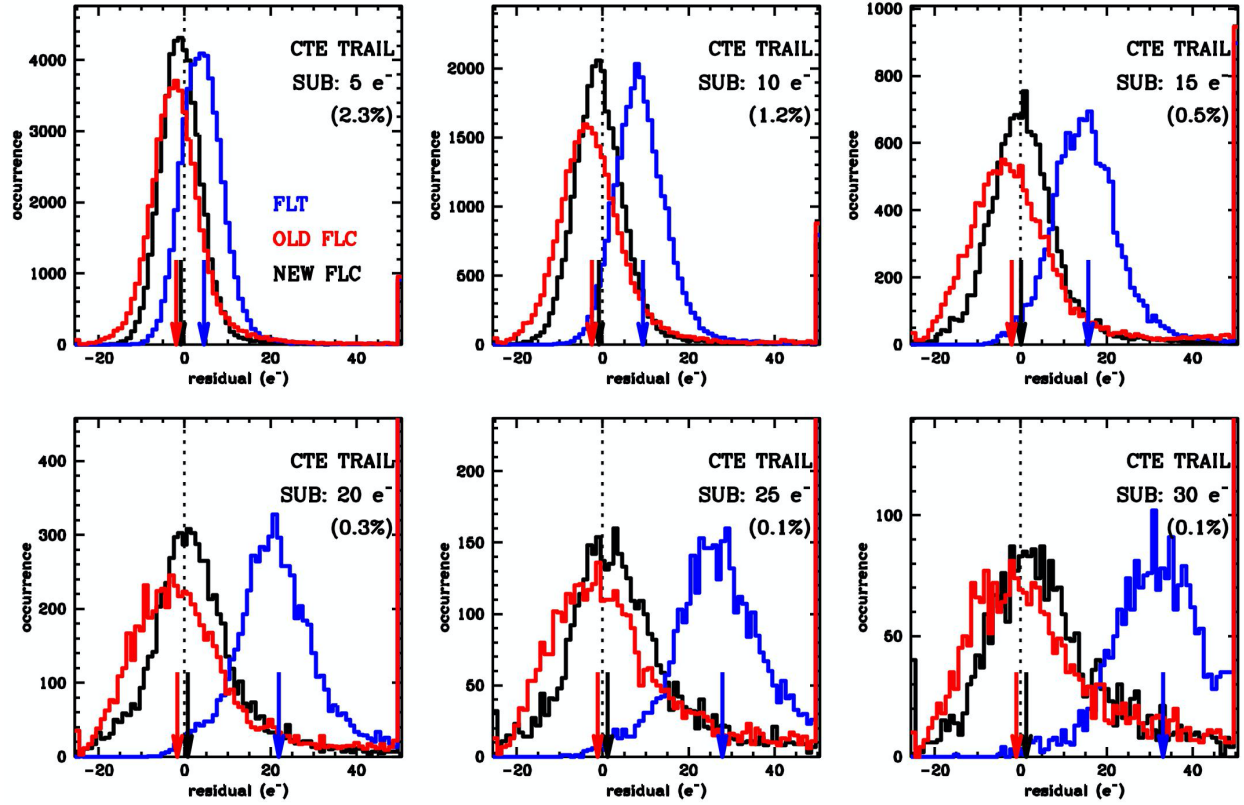
Figure 19: Effectiveness of the v1.0 and v2.0 pixel-based CTE corrections in removing CTE trails from behind cosmic rays. In each panel, the blue histogram shows the difference (in electrons) between a pixel value in a particular uncorrected dark exposure and the corresponding dark stack (which should reflect the true background underneath the CR trail). The red histogram shows the subtracted residual for the v1.0 pixel-based CTE correction, and the black histogram shows the subtraction for v2.0. The six panels show pixels for which the v2.0 model predicted a trail correction of 5 e⁻, 10 e⁻, 15 e⁻, 20 e⁻, 25 e⁻, and 30 e⁻. .

Option#4: Consider Large Dithers

Observers who are greatly concerned about understanding how CTE is impacting their observations could perform large (e.g., ½ chip-height shifts) so that they get both far-from and close-to the amp measures of their targets. Since the two UVIS chips read out in opposite directions, observing with a full-chip shift will cause observations with maximum losses in the first exposure to have minimal losses in the second. Such a strategy is used in calibration to provide a definitive assessment of CTE losses for stars (see Kuhn & Bajaj 2021a). The same could be done for non-stellar objects.

Observing the same field at different roll angles might also help to make the CTE losses more evenly distributed across the field. A 90-degree roll will put some of the sources that were at the top of the detector at the bottom, and vice-versa. Observers should keep in mind, though, that it is not always possible to observe every field at arbitrary roll angles.

8.   Advice for Post-Observation Analysis

The fact that the pixel-based correction can no longer be expected to correct the lowest-value pixels for imperfect CTE introduces unavoidable complications into our measurements. The pixel-based correction has worked well for a long time, and this is clearly not a welcome change. Unfortunately, though, the noise-amplification that the pixel-based correction introduces at low pixel values makes the corrected pixels near the sky level almost useless for science. If we are going to have noise in an image, that noise may as well be from real background electrons that helped to mitigate CTE, rather than artificial noise added by the reconstruction algorithm on a low background that failed to preserve the signal from our sources.

Once the observation has been taken, it is worth underlining that there is no way to improve the signal to noise that is present in the `raw/flt` images. We can add noise, but we cannot add any signal. This has always been true, but it is even more of a concern now that it is so easy to amplify the noise. The best thing that can be done is to measure the flux in the best way possible, then correct it for the CTE losses.

One option for observers to explore is tuning the CTE correction readnoise parameter for their science case. The default correction parameters in the pipeline are set to avoid read-noise amplification at all costs. Observers less concerned about this have the option of re-running either the old version of the pipeline (calwf3 v3.5) or the current pipeline (calwf3 v3.6) with a lower `PCTERNOI` threshold than the conservative value adopted in the pipeline (see Kuhn 2021b). Lowering the noise-mitigation thresholds can provide some correction for fainter sources, but the correction will not be perfect.

8.1 Bright Sources

For the purposes of this discussion, "bright" stars have detectable flux over the background in several of their central pixels, such that the star resembles a PSF and we can measure both a flux and a position for it. This typically implies S/N ~ 15, or ~250 electrons total.

Figure 20 shows the photometric residual trends for four different reductions: the uncorrected `_flt` images (**red**), the v1.0 pipeline (**blue**), the v2.0 pipeline (**green**) and a supplemental v2.0 run with a fixed RN = 4.0 electrons (**cyan**), which does somewhat more correction at the expense of more noise amplification. The trends shown here come from an independent analysis of the data from Cal-16401 (PI-Kuhn), which are being written up separately (Kuhn et al., in prep).

The photometry goes from the faintest stars that can be measured in individual exposures (at the left of each plot) to just below saturation (at the right). The stars have been measured by fitting the central 5×5 pixels with a PSF that is normalized such that the flux reported corresponds to the total flux-over-sky within a 10-pixel aperture. Note that astrometric losses (in pixels) are similar numerically to photometric losses (in magnitudes). A future ISR will quantify both loss astrometric and photometric losses as a function of source brightness and background.

In the visible bands for UVIS, about 18% of a star's flux will land in its central pixel when the star is centered on that pixel (less when the star lands on an edge or corner). To aid in application of these plots, we provide the number of counts in a star's central pixel along the top
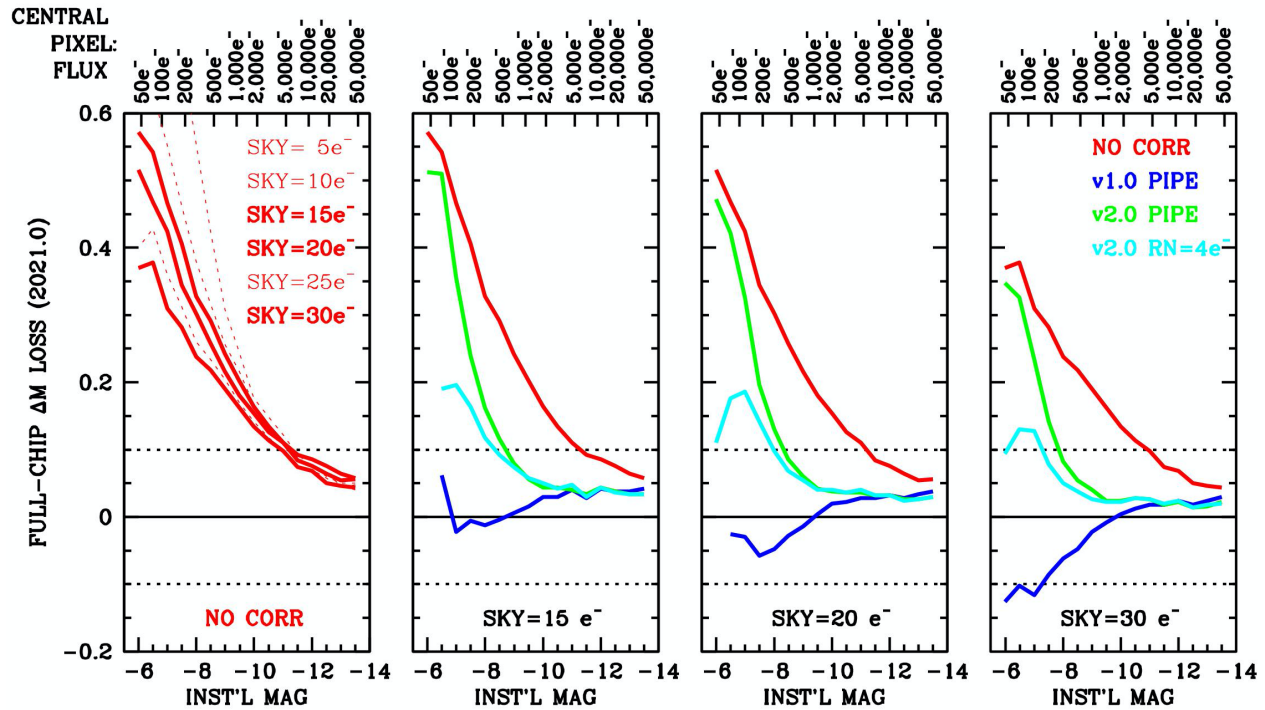
Figure 20: (Left panel) CTE losses as a function of instrumental magnitude in the uncorrected `_flt` images for various sky backgrounds. The panels to the right show the raw trends (red), the v1.0 pipeline correction (blue), the v2.0 pipeline correction (green) and a standalone run of the new correction (cyan) that does more correction while allowing more noise. Sky background levels are noted at the bottom of the rightmost three panels. These results correspond to epoch 2020.98 for 2048-pixel parallel transfers.

axis. A star with an instrumental magnitude of −13.5 (250,000 total counts) is near saturation with about 50,000 e^- in its central pixel. A star with an instrumental magnitude of −6 has about 50 e^- in its central pixel; this is the faintest star that can be reliably measured with PSF fitting in a single exposure.

The left panel focuses on the uncorrected losses (i.e., the `_flt` images) for a variety of sky backgrounds. Note that CTE losses for bright stars are not sensitive to the background. This is because, in an absolute sense, there are more traps that affect large electron packets than affect small packets. **Figure 9** shows that there are ~10 times more traps that affect charge packets larger than 500 e^- than affect those smaller. Thus, for bright stars, the losses due to the traps near the background level are a small fraction of the stars' total losses: as **Figure 20** shows, the losses for a bright star far from the readout amplifier are 4% for a background of 30 e^- and 6% for a background of 15 e^-.

However, even though most of the traps affect larger packets, the *per electron* losses are still greater for smaller packets (see **Figure 8**). The faintest stars shown in **Figure 20** suffer ~40% losses (0.35 magnitude) for a background of 30 e^-. And the losses go off the charts as the background gets lower. Losses are indeed severe for faint stars, but note that even bright stars near saturation on high backgrounds far from the readout amplifier, lose about 3% of their flux from a 5×5-pixel aperture.

The rightmost three panels in **Figure 20** show how well the various flavors of correction do for three selected backgrounds: 15 e$^-$, 20 e$^-$, and 30 e$^-$. The first background level (15 e$^-$) in the second panel is below the current recommendation of 20 e$^-$, hence the images suffer significant noise amplification. The v1.0 pipeline appears to do best for this background. For the brighter stars (m < −10), all the corrections do equally well. The cyan curve corresponds to the v2.0 correction with read-noise mitigation hard-set to 4.0 electrons (via the `PCTERNOI` keyword). The green curve corresponds to the v2.0 pipeline reduction, where the correction parameters have been set to prevent noise amplification at the expense of under-correcting faint sources. The green curve sensibly asymptotes to the red curve for faint stars (where there is no correction) and to the cyan curve for the brighter stars (where there is the best correction).

At the currently recommended background of 20 e$^-$, the v2.0 cyan curve does the best overall, though it is worth noting that none of the curves does a great job. Indeed, one of the reasons we are now leaning towards recommending reducing the flt images and correcting the post-measurement results for CTE is that *none* of the available corrections provides photometry without bias one way or the other. One could imagine perhaps averaging the v1.0 and the v2.0 corrections, but that would be *ad hoc*, without any theoretical justification.

The reason that the v1.0 correction does well at some backgrounds and not others is because it has a flat trap-density profile at typical background levels. **Figure 3** shows that the old model over-predicts warm-pixel losses between 15 e$^-$ and 40 e$^-$ electrons background and as a consequence, the model over-corrects the images. The v2.0 corrections are less dependent on the background than the v1.0 corrections, but they are still not perfect. We took some new calibration data in Cal-16440 that should allow us to use the direct-pinning method all the way up to electron clouds of ~1000 e$^-$, which should further improve the pixel-based model. Nevertheless, we recognize that the trade-off between pixel-based correction and noise amplification will continue to limit the effectiveness of even a perfect pixel-based model, and this trade-off will become worse as radiation damage continues to accumulate. Finally, it is worth noting that the v1.0 and v2.0 algorithms do not completely correct the brightest stars. Both algorithms under-correct all stars that are far from the readout amplifier that are brighter than S/N ~ 100 (m_{inst} ~ −10) by about 3%. It is not clear how the model could be improved to ameliorate this. WFC3/UVIS has very few high intensity warm pixels making it hard to calibrate the pixel-based model for the largest electron clouds. But the under-correction could also be related to incorrect assumptions we have made about how downstream charge packets shield upstream packets. That would be much harder to improve in the model, since we do not have direct ways to examine this aspect of the model.

All of these issues suggest that correcting measurements made on the `flt` images might be the safest way to correct point-source observations for CTE losses. The measurements could then be corrected for CTE losses by using the empirical fits available as a function of background level, source brightness, and distance from the readout amplifier (see Kuhn & Bajaj 2021a or Kuhn 2021b). It is certainly more complicated than simply reducing the pixel-corrected `flc` images, but the result is cleaner and more reliable. We will provide additional information on how to do this in future ISRs, including both formula-based and table-based corrections. In this report, we have focused on photometry, but corrections will be provided for astrometry as well.

8.2 Faint Sources

Figures 16 and **17** graphically show the impact of imperfect CTE on the profiles of faint point sources. An ideal model for the PSF would account for the observed shape and intensity changes with CTE loss, since it could operate on the non-corrected `flt` pixels. But such a model would be a complicated function of (1) the number of parallel transfers (i.e., distance of the source from the readout amp), (2) the date of observation, (3) the total flux of the star, and (4) the sky background level, as well as the usual other PSF parameters, including detector location (x, y). Furthermore, it is not clear that the pixel-based forward model could be used to take a pre-readout PSF and accurately adjust it for the CTE-blurring process for significantly affected stars. While this approach might work, it is unclear whether the current pixel-based model, which is based solely on delta-function WPs, does a suitable job handling the details of self-shielding of stars in the large-loss regime. We would need to compare predictions against the profiles observed in **Figure 16**, and if the model and observations did not happen to agree perfectly, it is not clear how to improve the model.

8.3 Work in Progress for Faint Sources

The good news is that **Figure 16** shows that the photometric correction does not depend critically on the aperture: we see the same loss fractions for the 1×1-pixel aperture as for the 2×2-pixel aperture. As such, the correction should boil down to interpolation of a two-dimensional table of sky and flux, such as the example tables we present below.

Table 1 below gives the fraction of flux lost as measured in a 2×2-pixel aperture as a function of sky and flux for the faint stars explored in **Figure 17**. It can be used to help plan observations or calibrate losses on existing observations. The table value corresponds to a source observed at ~1750 rows from the register at the 2020.95 epoch. To adjust for a different location on the detector or date, the loss fraction should be multiplied by:

$$\frac{t - 2009.35}{2020.95 \ - \ 2009.35} \times \frac{j}{1750}$$

where j is the row number of the star's position.

Table 1: Fractional flux loss for a source with the true 2×2-pixel flux as listed on the left and the sky value as listed along the top.

		sky (e⁻)				
		12	16	20	25	30
	11.7	0.77	0.65	0.55	0.45	0.30
	28.5	0.57	0.45	0.40	0.32	0.27
true 2×2 flux (e⁻)	72.3	0.47	0.35	0.32	0.27	0.24
	183	0.38	0.30	0.27	0.23	0.21
	417	0.30	0.25	0.22	0.20	0.18

Table 2 below provides a direct way to map an observed 2×2-pixel-aperture flux for a point source at $j \sim 1750$ on a particular background to the true point-source flux. Since the trends seen in **Figure 17** appear linear, then linear interpolation should be sufficient for Tables 1 and 2 as well. The PSF models on the UVIS website[5] are normalized to have a flux of 1.00 within a radius of 10 pixels. The F606W model predicts that about 18% of a star's flux will land in a pixel if the star is centered on that pixel, and about 48% of the flux will land within the central 2×2 pixels. So, one could scale up these corrected numbers by 1.00/0.48 = 2.08 to obtain a flux that is normalized to that of a PSF-fit star.

Table 2: Flux in electrons observed for a source in a 2×2-pixel aperture for the true 2×2-pixel flux as listed on the right and the sky value listed along the top. This table corresponds to the 2020.95 epoch for a source at $j \sim 1750$.

sky (e⁻)						
12	**16**	**20**	**25**	**30**		
2.7	4.1	5.3	6.4	8.2	**11.7**	
12.3	15.7	17.1	19.4	20.8	**28.5**	**true 2×2 flux (e⁻)**
38	47	49	53	55	**72.3**	
114	128	134	141	145	**183**	
292	313	325	334	342	**417**	

Many sources at the faint limit of these tables cannot be observed with high significance in a single exposure. It takes many exposures stacked together to add up to a significant detection. If this correction is used on a stacked set of images, then it is only reliable if the individual exposures are at the same pointing (apart from small dithers), such that each would suffer the same CTE losses. If we have a mosaic of pointings where a particular source lands near the readout amplifier in some exposures and far from the amplifier in other exposures, the stacking will mix high-signal observations with low-signal observations. It is unclear how best to reduce data taken in such a situation, as the various exposures would have very different signal-to-noise qualities and photometric corrections.

And there are some complicating factors. First, any added postflash is routinely subtracted from `flt` and `flc` images in the calwf3 pipeline. This is so that the true astronomical sky background will be reflected in the output products. As a consequence, the background seen in a `drz` or `drc` image is not representative of the background that the source experienced on its path to the readout. One can reduce the data without subtracting the postflash to obtain an image with full background (i.e., set the `FLASHCORR` keyword to "Omit" then run calwf3); to do so will require downloading the raw files as well as all necessary reference files and tables. The calwf3 pipeline does record the mean post-flash level removed from the image (MEANFLSH) but the illumination pattern varies by ±20% across the field of view and is slightly different (by a few percent) depending on which shutter blade was used for the flash (see Martlin & Baggett 2017)

---

[5] https://www.stsci.edu/hst/instrumentation/wfc3/data-analysis/psf

The 2×2-pixel aperture discussed above is in the flt-frame, so if there is any oversampling done in the drizzle process, the aperture size would have to be transformed as well. An additional complication with correcting the source flux measured on a stacked drizzle image is that the location of the source in the drizzle stack does not directly map to a particular number of parallel transfers. One might have to identify its location in an `flt` image to determine the number of parallel shifts experienced.

Imperfect CTE has an impact on both photometry and astrometry. The empirical correction provided above does allow for a slight astrometric shift of the location of the aperture to maximize the flux in the adopted aperture; this is typical of how we would find and measure faint stars in image stacks. We do not report the shifts here, but we are planning to provide both photometric and astrometric tables for correcting bright and faint stars in `flt`-type images in a future report.

The fact that `flt` and `flc` images will both require some kind of post-measurement correction, it seems easier to make measurements on the `flt` images and no longer use the pixel-based correction, since the sky background has to be quite high for it to provide a correction that can be used without concerns about noise amplification or source non-restoration. However, one reason that the pixel-based correction is still useful is that the pixel-based correction removes the trails from cosmic rays and warm pixels. **Figure 18** shows that it does this quite well, thus restoring a large fraction of the detector for analysis. This is particularly relevant for long exposures. Perhaps in the future, it may be beneficial to come up with `flt`-type images that have the CR trails removed, but nothing else changed, to get the best of both worlds.

8.4 Extended Sources

We have not yet taken any calibration observations to evaluate the impact of CTE on extended sources. Now that we have a good handle on how CTE impacts point sources, it may be worthwhile to evaluate explicitly how it impacts barely resolved and resolved objects.

We saw above that faint stars and warm pixels appear to suffer the same photometric losses, so it makes sense to think that barely resolved objects would suffer similar fractional losses. However, at some point, self-shielding has to have some effect, since the background itself provides significant mitigation.

The vertical profiles shown in **Figures 16** and **17** show that sources do get broadened by CTE. It is not clear, though, how to quantify this for science applications.

9. Conclusions

The key points presented in this report can be summarized as follows:

- A new pixel-based CTE correction (v 2.0) is available for WFC3/UVIS observations. The new model has a more accurate distribution of charge traps at the low and bright ends. By default, all data processed through MAST as of Apr 2021 have this correction applied.

- We describe the development of the model. The correction, for both the original and new version, is based on a two-parameter model, $\phi(q)$ and $\tau(\Delta j, q)$, where the first represents the number of traps in the typical 2048 pixel column that affect charge packet sizes $> q$ electrons, and the latter describes the release profile of the charge traps, as a function of row number j, that absorb the qth charge (see **Figure 7**).

- The original WFC3 pixel-based CTE correction, in use until Apr 2021, performed well for many years. As long as CTE losses were at a perturbation level (up to ~25%), the correction properly restored middle-intensity sources albeit slightly over/under-correcting faint and bright sources. With increasing radiation damage, the CTE losses for faint sources are currently well beyond a perturbation. As of early 2021, marginal CTE losses are over 60% for faint sources in images with backgrounds of 12 electrons (see **Figure 14**). This required a new approach.

- The WFC3 warm pixel (WP) trails exhibit sharply-dropping profiles for faint WPs (< 50e-) while bright WPs (> 70e⁻) exhibit a more gradual trail typical of brighter WPs (**Figure 5**). We note this is different from the WPs in ACS, where profiles for WP's 20-90 e⁻ and brighter are roughly the same. Also, while the WFC3 and ACS CTE models are similar at the bright end, for faint WP's, WFC3 experiences significantly higher losses (**Figure 9**). It is unclear how these differences are related to the mini-channel manufactured into the WFC3/UVIS CCDs.

- The approach to pinning the WFC3 model has changed. For the v 1.0 pixel-based correction (calwf3 v 3.5 and earlier; prior to Apr 2021), the model was pinned based on CTE losses assessed via the trails behind warm pixels (WPs) in zero-background dark frames, i.e., and indirect method. The v 2.0 correction (calwf3 v 3.6 and later, after Apr 2021), is pinned based on CTE losses measured by directly evaluating warm pixel levels in short darks relative to their levels in long exposure time darks (i.e. 'truth'), at a variety of image background levels.

- The new model is better able to simulate the actual empirical losses than the old model. For example, for hot pixels (~75e⁻), the old model systematically under-predicted the losses by 15-20% when image backgrounds are 15-20e⁻/pix (Fig 4). Only at extremely high image backgrounds (70-100 e-/pix) did the two models perform equally and in those cases, both under-predicted the CTE loss. We note there does not appear to be signficant 'self-shielding' in stars: the fractional losses for stars are similar to those for the WPs used for pinning the model (Fig 17).

- A new key feature of the v2.0 correction is a significant reduction in added noise. This approach works extremely well and keeps the pixel-to-pixel noise in v2.0 corrected images at the `flt` level, in stark comparison to the noise in the v1.0 correction (Fig 12). **However, the penalty of this noise mitigation approach is that the fainter the source, the less CTE correction applied (Figure 19).** The fact that the WFC3 CTE losses are no longer perturbations means that the pixel-based correction cannot be relied on to correct faint sources on these moderately low backgrounds without the addition of a significant amount of noise.

- Optimum CTE mitigation begins at the observation planning stage. Observers should (1) place targets close to the amplifier when possible (2) take fewer but deeper exposures, and (3) add post-flash to ensure a total image background (dark+sky+flash) of at least 20e-/pix. Observers greatly concerned about the effects of CTE loss may wish to

perform large dithers in order to alternate which portions of their target are close to the readout amplifier, to allow a direct assessment of the CTE losses on their specific targets.

- Post-observation CTE mitigation at this stage in the instrument lifetime has become complex. For observers with relatively bright targets (S/N > ~30) on image background of at least 20e-/pix, the v2.0 pixel-based correction provides a reasonable correction, restoring sources to within ~5% or better. Observers with faint sources (S/N < ~30) will need to consider other approaches, such as: (a) tuning the CTE correction readnoise parameter (PCTERNOI) to their particular science case and rerun calwf3 (a process illustrated in a Jupyter notebook, Kuhn 2021b), (b) runing the v2.0 correction then correct faint sources afterwards, using either the tables provided in this report (Tab 1,2) or the full correction formulae (Kuhn 2021), or (c) foregoing the pixel-based CTE correction and applying only formulaic corrections.

- We continue to work on alternate methods for measuring accurate fluxes for stars of all brightnesses both photometrically and astrometrically. These will be documented in upcoming reports, advertised in the STANs (Space Telescope Newsletter[6]) and posted on the WFC3/UVIS CTE repository[7].

- This document has been focused on CTE's impact on photometry. CTE also impacts astrometry. An upcoming ISR will tabulate both the astrometric shifts and the photomeric losses as a function of source brightness and background.

Acknowledgements

10. References

Anderson, J. & Baggett, S. 2012 WFC3/ISR 2012-12: "WFC3/UVIS Sky Backgrounds"

Anderson, J. & Baggett, S. 2014 WFC3 ISR 2014-13 "Sink Pixels and CTE in the WFC3/UVIS Detector"

Anderson, J. & Bedin, L. R. 2010 PASP 122 1035, "An Empirical Pixel-Based Correction for Imperfect CTE. I. HST's Advanced Camera for Surveys", aka **AB10**

Anderson, J. & Ryon, J. 2018 ACS ISR 2018-04 "Improving the Pixel-Based CTE-Correction Model for ACS/WFC"

Biretta, J., & Baggett, S., 2013 WFC3/ISR 2013-12: "WFC3 Post-Flash Calibration"

---

[6] https://www.stsci.edu/hst/instrumentation/wfc3/documentation/stsci-analysis-newsletter-stan

[7] https://www.stsci.edu/hst/instrumentation/wfc3/performance/cte

Kuhn, B. and Bajaj, V. 2021a WFC3/ISR 2021-03: "WFC3/UVIS: External CTE Monitoring 2009 – 2020"

Kuhn, B. 2021b, Jupyter notebook for reverting to the old processing: https://www.stsci.edu/files/live/sites/www/files/home/hst/instrumentation/wfc3/performance/cte/_documents/calwf3_oldcte.zip

Kuhn, B. 2021c, WFC3/ISR in progress, "Formula Corrections for CTE in WFC3/UVIS Images"

Massey, R. et al., 2010 MNRAS 401 1, "Pixel-Based Correction for Charge Transfer Inefficiency in the Hubble Space Telescope Advanced Camera for Surveys"

Martlin, C. and Baggett, S. WFC3/ISR 2017-03, "Long-Term Stability of the Post-Flash LED Lamp"

Montes-Quiles, M. 2021 WFC3/ISR in prep. "Sink Pixels in the Modern Era"