

WFC3/IR Dark Current Stability

B. Hilbert and L. Petro
June 11, 2012

ABSTRACT

Examination of the dark current behavior of the WFC3/IR channel over a span of almost 1000 days shows that the mean dark current has remained unchanged since launch. Image-to-image variation in mean dark current is 16% of the total mean dark rate. This variation appears to occur in a spatially-dependent but consistent manner, with the bottom and right edges displaying dark rates which increase more quickly than those elsewhere on the detector. Pixel-to-pixel variations in the dark current across the dataset imply that the 1σ uncertainty associated with the standard pipeline dark current subtraction is roughly 20%-33% of the mean dark current rate. We have also identified a population of ~ 2500 pixels (0.25% of the science pixels) which are inconsistently hot over time. In a separate search, we find 2000-4000 pixels (0.2% - 0.4%) in most SPARS200 ramps which have dark rates more than 3σ from their nominal mean level. We find no correlation of mean dark rate with bias level, but an anti-correlation of the number of hot pixels with bias level. The measured dark current also does not appear to be correlated with instrument temperatures or voltages, or sun angle relative to HST.

Introduction

As of mid-April 2012, the WFC3 IR detector had spent 35 months on-orbit, during which time it has been subjected to the flux of high-energy, ionizing radiation in the low Earth orbit space environment. This exposure is especially increased during the daily passages of the South Atlantic Anomaly (SAA; Barker et al. 2010). In a typical IR channel observation, the results of cosmic rays impinging upon the detector are visible as

high-signal points or tracks, affecting many pixels. The cosmic rays also damage the lattice in the HgCdTe detection volume, which may change the functional characteristics of the pixels. The purpose of this study is to characterize the dark current signal of the IR channel over time, and to search for any possible long-term trends which may be the result of this time spent in low earth orbit.

Data

Our goal was to analyze a long timeline of high signal-to-noise dark current measurements in order to be sensitive to potentially small changes in dark current. We used a dataset of 252 DARK ramps taken during Servicing Mission Observatory Verification (SMOV) and Cycles 17 through 19. These ramps were collected over a span of 977 days, between July 2009 and March 2012. Of the 252 observations, 183 were collected using the SPARS100 sample sequence, with an exposure time of 1403 seconds per ramp. The other 69 files were collected with the SPARS200 sample sequence, which has an exposure time of 2803 seconds per ramp. Table 1 in the Appendix lists the individual files used in the analysis, along with the proposal numbers under which they were collected.

We used slightly modified versions of the *flt* files output from the *calwf3* data processing pipeline for this analysis. Each file was initially composed of 16 non-destructive readouts of the detector. During *calwf3* processing, each file had reference pixel and zeroth read subtraction performed, along with non-linearity correction and cosmic ray rejection. A best-fit line through the 16 reads of each pixel then transformed the dark current signal measurements to a signal rate. These calibration steps were accomplished by turning on the following *calwf3* calibration switches in the header of each file: DQICORR, ZOFFCORR, BLEVCORR, NLINCORR, CRCORR, and UNITCORR.

Upon completion of the pipeline, we manually performed two more calibration steps on each file. First, the signal rate was converted from units of data numbers (DN) per second into e^-/sec using a gain value of $2.38 e^-/\text{DN}$, which is the value currently used by *calwf3*. (Hilbert, 2008b)

We then attempted to remove persistence effects from the calibrated *flt* files. Initial examination of the data after processing by *calwf3* revealed the presence of low-level persistence in many of the dark current images. Details on persistence in the IR channel are given by Long et al. (2010 and 2011). This persistence could potentially cause two problems with our dark current analysis. Persistence resulting from point sources or small extended sources increased the measured signal rate in a small number of pixels such that these pixels could mistakenly be flagged as hot pixels. Secondly, in the case where persistence was wide-spread across the detector, as in several cases where flat-field observations were taken prior to the darks, the persistence artificially increased the

mean signal rate across the entire dark current image. In an attempt to remove persistence-induced signal and minimize the two contamination effects described above, we applied the persistence model and correction described by Long on the persistence website (http://www.stsci.edu/hst/wfc3/ins_performance/persistence). A link to this website is located in Section 6.9 of the WFC3 Data Handbook (Rajan et al. 2010). After correction, persistence levels in our dark current data were reduced to levels which minimized false identification of hot pixels.

Analysis

Mean Dark Current

Before beginning the dark current analysis, we first identified populations of pixels to exclude from our calculations. By reading in the data quality array associated with each file, we obtained a list of all pixels flagged as bad from the application of the bad pixel table as well as through various data quality checks within *calwf3*. For most of the analysis in this study, we chose to ignore any pixel flagged as “dead” or “unstable” (approximately 0.4% and 1.0% of the pixels on the detector, respectively). Details of the various flavors of bad pixels can be found in Hilbert & Bushouse (2010). While many (~3,000) of the “unstable” pixels are also flagged as hot due to high dark current levels, we excluded these pixels from our analysis due to their inconsistent behavior. As detailed below, we did make some preliminary checks into the behavior of these unstable-and-hot pixels. However, unstable pixels were excluded from most of the analysis reported in this discussion.

The first calculations made to track the dark current behavior involved simply plotting a histogram of the dark current rate for each of the input files. The top panel of Figure 1 shows the histograms for all 252 files overlaying one another on a log scale, with a different color for each file, for clarity. In order to assign a single, objective, precise value for the “typical” dark current of each exposure, we fit a Gaussian to the peak of each histogram and identified the peak value of the fit. The peak values are plotted versus time in the bottom panel of Figure 1. The colors of the data points are consistent between the top and bottom panels. For the remainder of this study, we will refer to these Gaussian-fit peak values as “mean” dark current values, for brevity.

The top panel of Figure 1 shows that the histograms of the ramps all have a roughly similar shape to one another, with the curve sliding as a whole left and right along the x-axis between files. Figure 2 shows the same dark current histograms, with each shifted horizontally by its mean value. This lines up the peaks of the histograms, allowing the more subtle changes in the shape of the distributions to be more visible. The majority of the histograms follow one of two distribution shapes. The taller, narrower distributions

correspond to the longer SPARS200 files, which have double the exposure time of the SPARS100 files that make up the shorter, wider distributions.

The color scheme for the 252 files is the same in both panels of Figure 1 as well as Figure 2, in order to more easily identify individual ramps within these plots. The plot of the “mean” dark current values as a function of time in the bottom panel of Figure 1 gives a clearer picture of the dark current’s overall behavior with time. From this plot, we see no long-term trend in the mean dark current level over the 977 days of the dataset. The mean value of these peak dark current levels is $0.046 \text{ e}^-/\text{sec}$, with a sigma-clipped standard deviation of $0.0075 \text{ e}^-/\text{sec}$. From the plot we see that ramp-to-ramp variation dominates any variation due to a long term trend in the mean level. The seven points in this plot that are circled in red indicate observations made shortly after the collection of internal flat field observations. This implies that these files are at the highest risk for severe, widespread persistence contamination. However, since these seven files all show mean dark current rates within the envelope defined by all of the other observations, we assert that the persistence-correction script was successful in removing most of the persistence signal in these files.

To better assess the differences between the dark current measurements, we created a variability map by calculating the sigma-clipped standard deviation of the signals for each pixel (except for the dead and unstable pixels noted above). Figure 3 shows a histogram of this variability map. The peak of the histogram shows that the 1σ variation in dark rate for most pixels is $0.0094 \text{ e}^-/\text{sec}$ in the SPARS200 ramps, and $0.015 \text{ e}^-/\text{sec}$ in the SPARS100 ramps, which is approximately 20% and 33% respectively, of the $0.046 \text{ e}^-/\text{sec}$ mean dark rate and 2.3 and 2.6 times the 1σ values associated with Poisson noise.

Variation in the dark current on the level suggested by Figures 1 and 2 implies potential errors in dark current calibration of science data, and therefore larger-than-anticipated uncertainties in faint source photometry results. Therefore, our next task was to investigate this dark current behavior further.

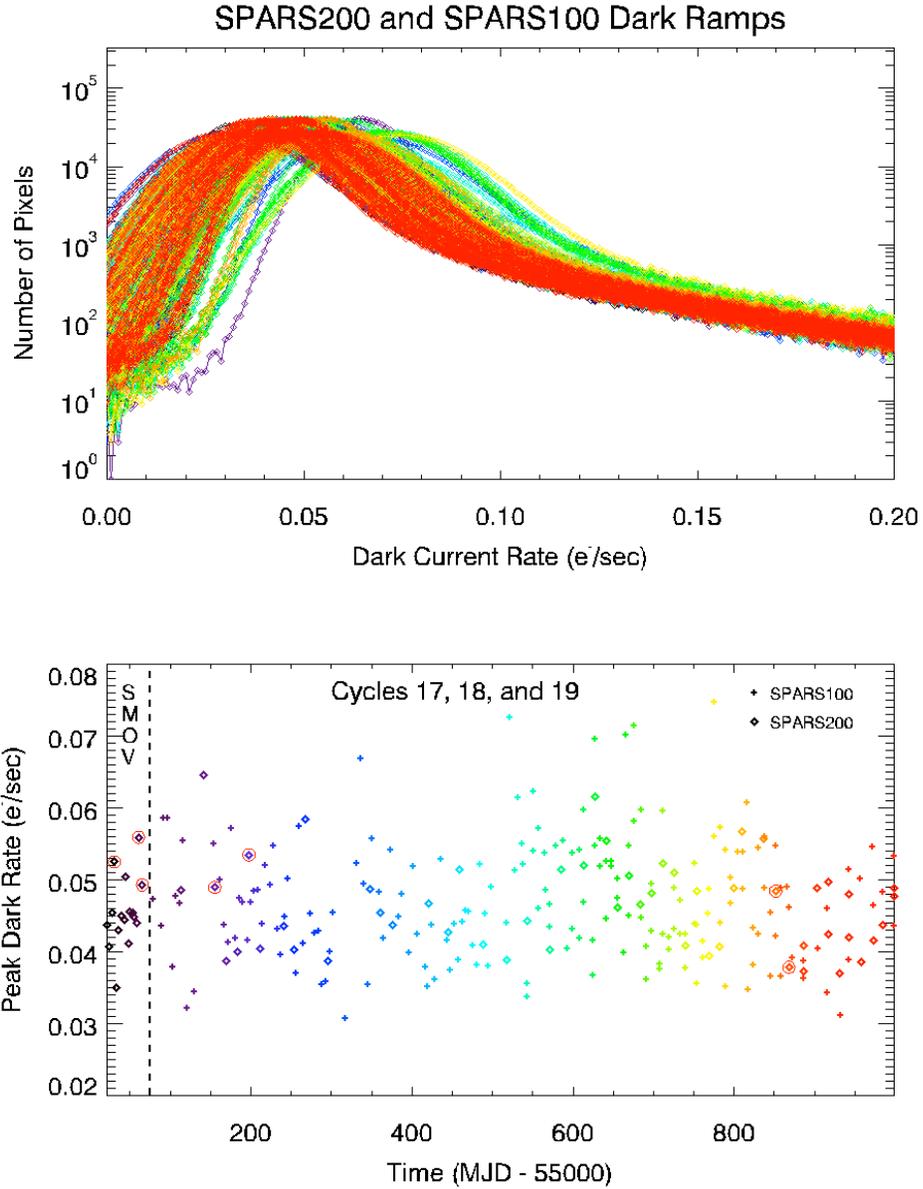


Figure 1: (*Top panel*) Histograms of the dark current for all input files. (*Bottom panel*) The values of the peaks for all histograms, plotted versus time (note that an MJD of 55000 corresponds to June 18, 2009). Diamonds correspond to SPARS200 ramps, each with an exposure time of 2803 seconds, while crosses show data for SPARS100 ramps, each with an exposure time of 1403 seconds. Colors correspond between the two panels, as well as with those in Figure 2. The sigma-clipped mean of the peak values is $0.046 \text{ e}^-/\text{sec}$, while the sigma-clipped standard deviation is $0.0075 \text{ e}^-/\text{sec}$. Points circled in red indicate observations made shortly after internal flat field observations, suggesting the possibility of wide-spread persistence contamination.

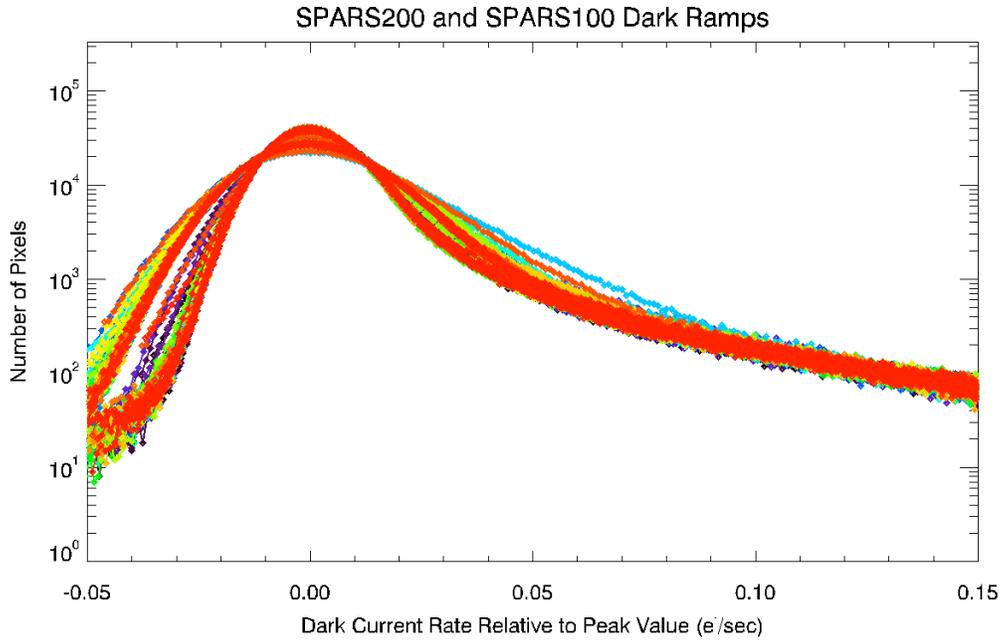


Figure 2: Same as the top panel of Figure 1, but with each histogram horizontally offset by its peak value. By lining up the peaks of all of the histograms, the more subtle differences in the shapes of the distributions are more evident. Focusing on the region containing the peaks (-0.01 to 0.01 e^-/sec), one can see that there are two main types of distributions. Those with the taller and more narrow peak are from the longer SPARS200 ramps, while the shorter, wider distributions are from the SPARS100 ramps, which have half the exposure time of the SPARS200 observations.

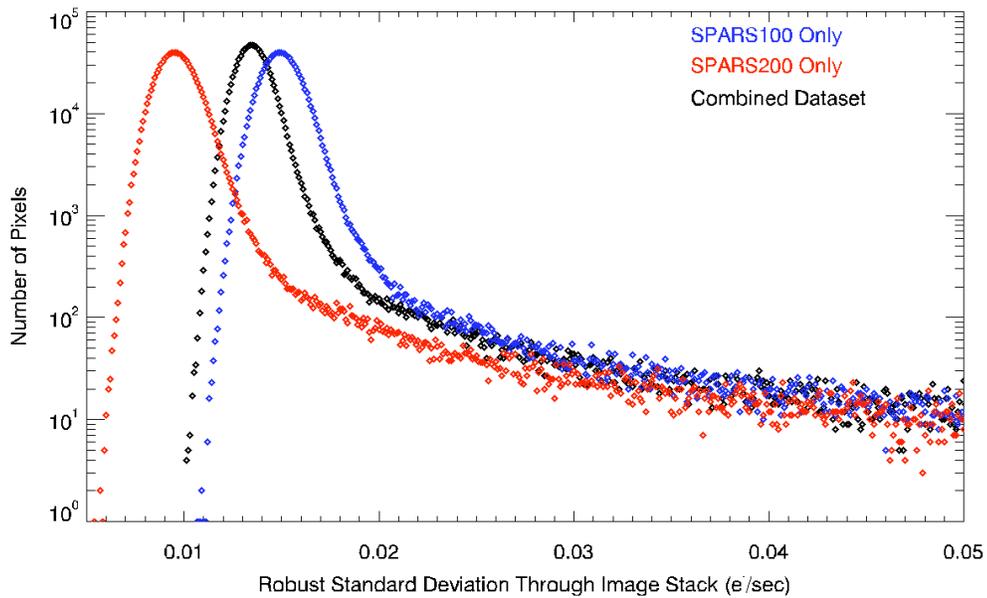


Figure 3: Histograms of the sigma-clipped standard deviation through the stacks of input images for each pixel. The red curve was constructed from the SPARS200 ramps, the blue curve from the SPARS100 ramps, and the black curve from the combined dataset.

Dark Current Variations

Galleries of Dark Current Observations

Figure 4 displays all SPARS200 persistence-corrected dark current images used in this analysis, placed in order of increasing mean brightness from left to right and bottom to top. There are several features to note in this mosaic beyond the obvious increase in mean signal level from image to image. First, there are residual persistence features in a number of the frames, including dithered point sources and extended sources, as well as over-corrected (i.e. black) grism sources. While the persistence correction scripts do contain tunable parameters, once we had confirmed that the correction had removed excess signal to the point that contamination of our hot pixel census would be minimized, we made no effort to identify optimal parameter values.

A second feature present in many of the frames is the “wagon wheel”. This well known feature is located in the bottom right corner of the detector and is a collection of pixels which up until now were thought to be non-nominal due only to their low quantum efficiency. We were initially surprised to see the wagon wheel in a dark current image, as it was not visible in ground-test dark current data, as presented in Figure 1 of Hilbert (2008-30). However, a re-examination of the ground-test data using a higher signal-to-noise image calculated by line-fitting the signals up the ramp for each pixel, rather than the last read minus first read image in Hilbert (2008-30), shows that the wagon wheel was indeed present in ground-test dark current images (Fig. 5). This implies that the pixels in the wagon wheel exhibit low dark current rates as well as low quantum efficiencies (Hilbert 2009-42).

Images highlighted in red correspond to the points circled in red in the lower panel of Figure 1. These dark current files were each collected shortly after one or more internal flat field observations and are therefore seen as having the greatest potential for a large enough persistence contamination that the detector-wide mean dark current rate could be affected. The fact that these images do not have the highest mean dark current levels out of the 69 files, but are instead mixed in among other images with equal or higher dark current suggests that the persistence correction was effective at removing the majority of the persistence in these files.

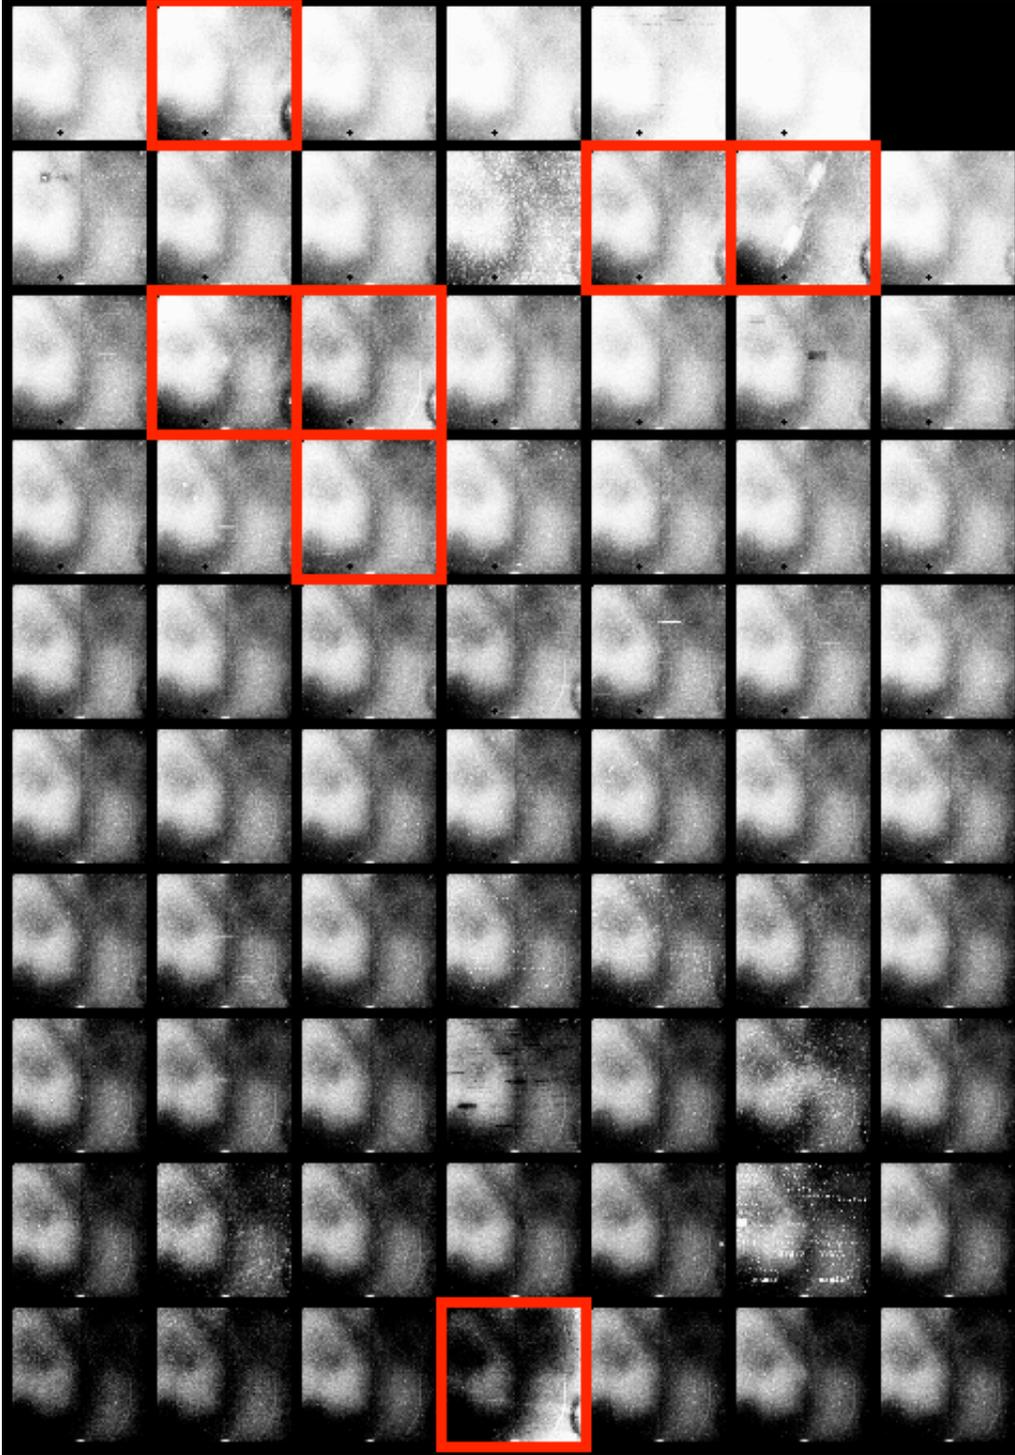


Figure 4: 7x10 grid showing all 69 SPARS200 dark current observations used in this analysis, placed in order of increasing mean brightness from left to right and bottom to top. Residual persistence is visible in some frames. This image has a histogram equalization stretch from 0 – 0.07 e/sec.

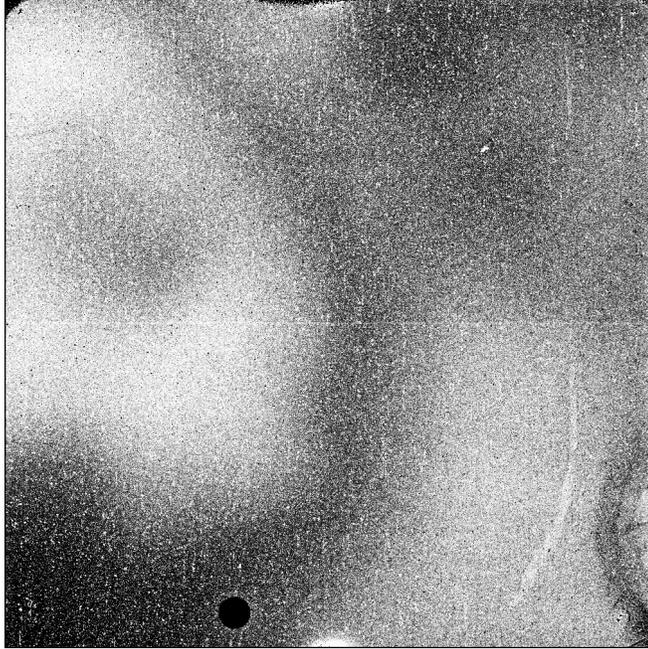


Figure 5: Image of the SPARS200 dark current reference file derived from ground-test data. The image is shown with a histogram equalization stretch between 0 and 0.07 e^-/sec , matching the scale in Figure 4.

Figure 5 shows a mean dark current image derived from ground-test data, and presented with an identical stretch to the images in Figure 4. Note the similar appearance to the images in the lower rows of Figure 4 (excluding residual persistence), indicating that to first order, the dark current behavior across the detector is unchanged.

To look more closely at spatial and temporal changes in the dark current signal, we needed to remove the base dark current from all of the files and examine the residual signals. Figure 6 is the same as Figure 4, but with the fifth dimmest dark current file (located third from the right on the bottom row) subtracted from each of the other darks. We chose not to use any of the lower dark current files for this subtraction for two reasons. First, one of these files (highlighted in red in the bottom row of Figure 4) was collected soon after internal flat field observations, and displayed a non-nominal signal morphology. Through iterative testing, we also found that the other three of the four dimmest files display non-nominal dark current morphologies across the detector as seen in Figure 8. We therefore settled on the fifth dimmest dark current image as the basis of comparison, and subtracted it from the other SPARS200 images.

The result of the subtraction, seen in Figure 6, is a mosaic of images displaying the extra signal measured by the detector above a base dark current map. The most obvious difference between Figure 6 and Figure 4 is that the images for most frames now appear relatively flat. The donut-shaped region on the left side of the detector is now gone, as is the area of elevated dark current in the lower right quadrant. Sources due to residual persistence are more visible than in Figure 4 (including 6 of the 7 darks with preceding flat field observations). There also appears to be a very slight depression in signal in the lower- and right-most two-thirds of the detector in many of the images.

The wagon wheel is still visible in a handful of frames, most notably those frames with bright lower right corners and edges. Assuming that the subtraction of the nominal dark current image from these images has successfully removed all features created by spatial differences in dark current, the appearance of the wagon wheel in these images suggests two possibilities. First, that a spatially uniform but time-variable external illumination is present on the detector, and the non-nominal QE of the pixels in the wagon wheel is to blame for its appearance. Another possibility is that the dark current of the detector is varying in time, but increases more along the right edge of the detector than elsewhere. We investigate the latter possibility in the section below.

Figure 7 shows the ground-test-derived dark current image with the same on-orbit dark subtracted. This image appears very similar to many of those in Figure 6 except for a left/right asymmetry which follows the quadrant boundaries. Note the slight brightening of the right edge of the detector, as well as the very faint appearance of the wagon wheel.

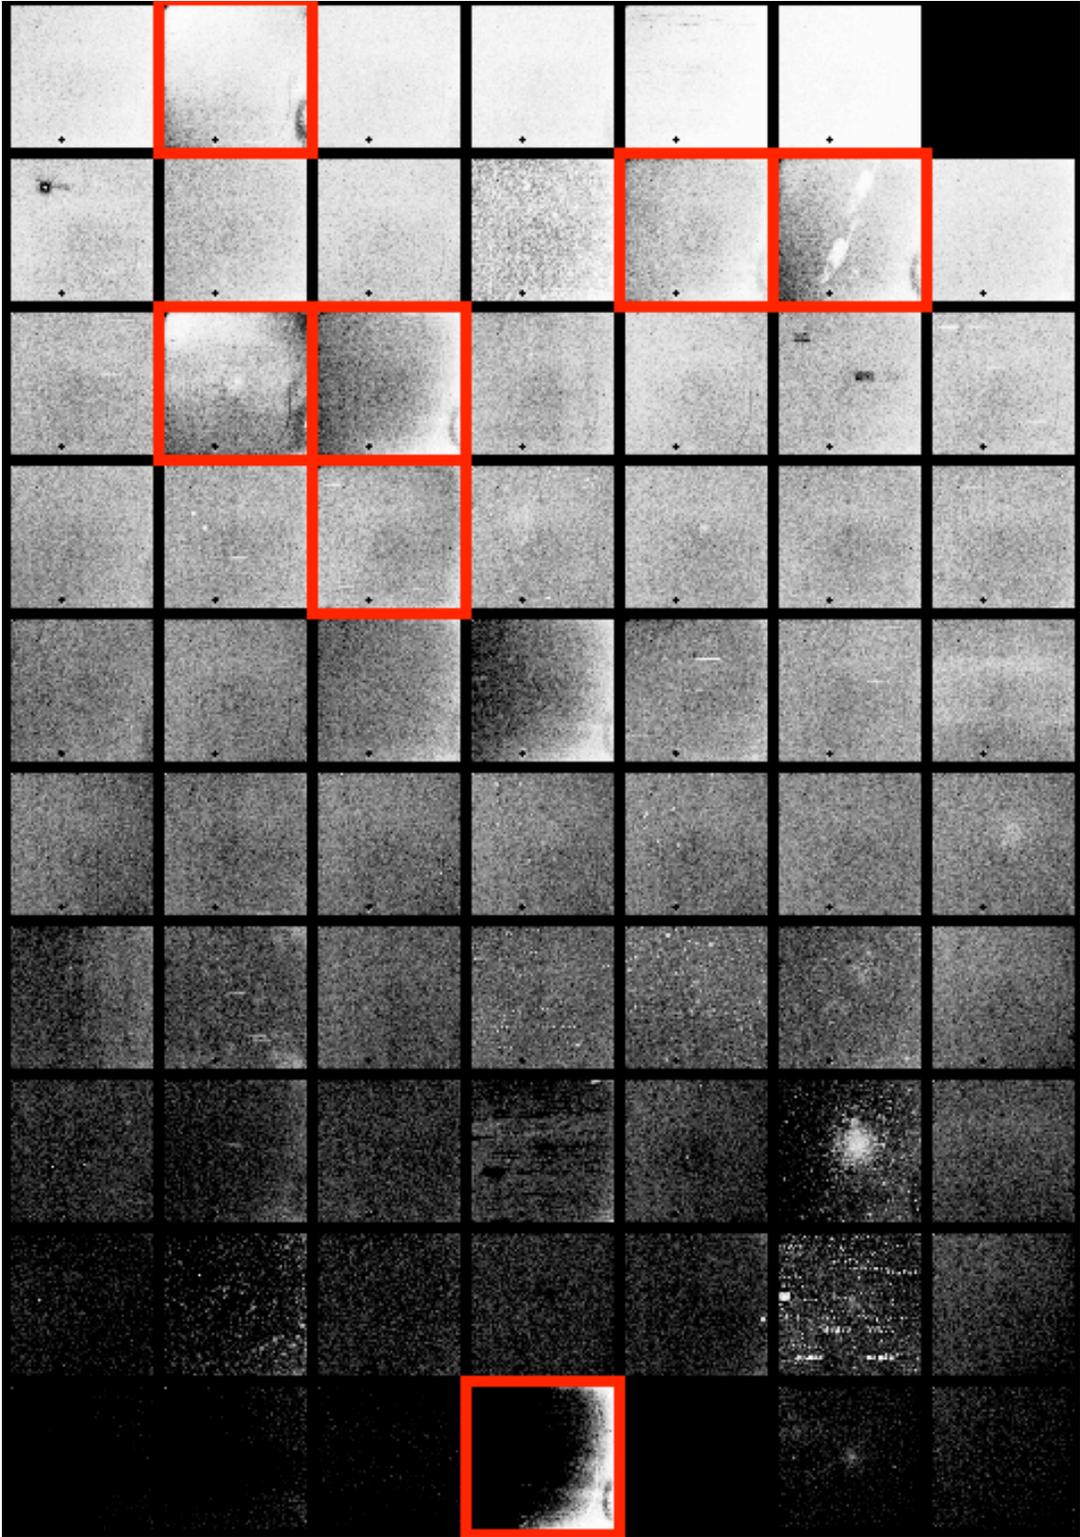


Figure 6: Gallery of dark current images, similar to Figure 4, where each image has had a “nominal” dark current image subtracted from it. The “nominal” frame is the third frame from the right in the bottom row. The figure is shown with a global histogram equalization stretch from 0 to 0.07 e'/sec . Files with preceding flat field observations are highlighted in red.

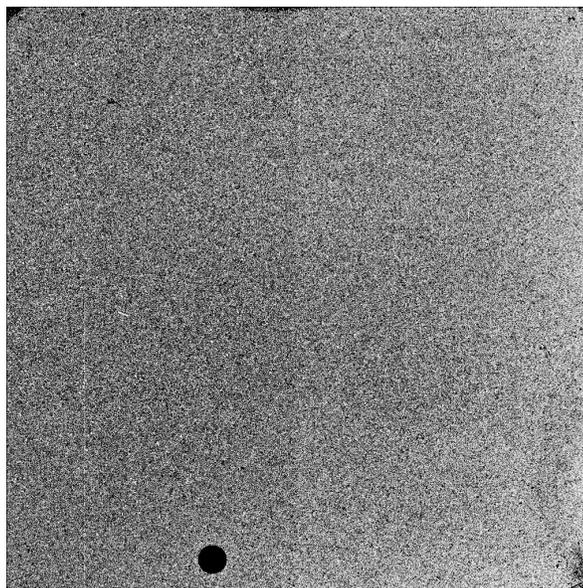


Figure 7: TV3 dark current reference file with the “nominal” dark current image subtracted, similar to what is shown in Figure 6. This image is presented at the same stretch as Figure 6, for ease of comparison.

To emphasize smaller-scale differences in dark current signal, we performed another subtraction, with the goal of highlighting any intra-quadrant regions with small dark current differences. For this, we took the images in Figure 6 and calculated the mean signal in each of the four quadrants of each image. We then subtracted each of those four mean values from the appropriate quadrant of each image, with the results displayed in Figure 8. The mean of each quadrant after this subtraction was zero, meaning that signal differences at the quadrant boundaries were potentially artificially inflated. Examples of this include six of the seven frames with preceding internal flat fields, and highlighted in red.

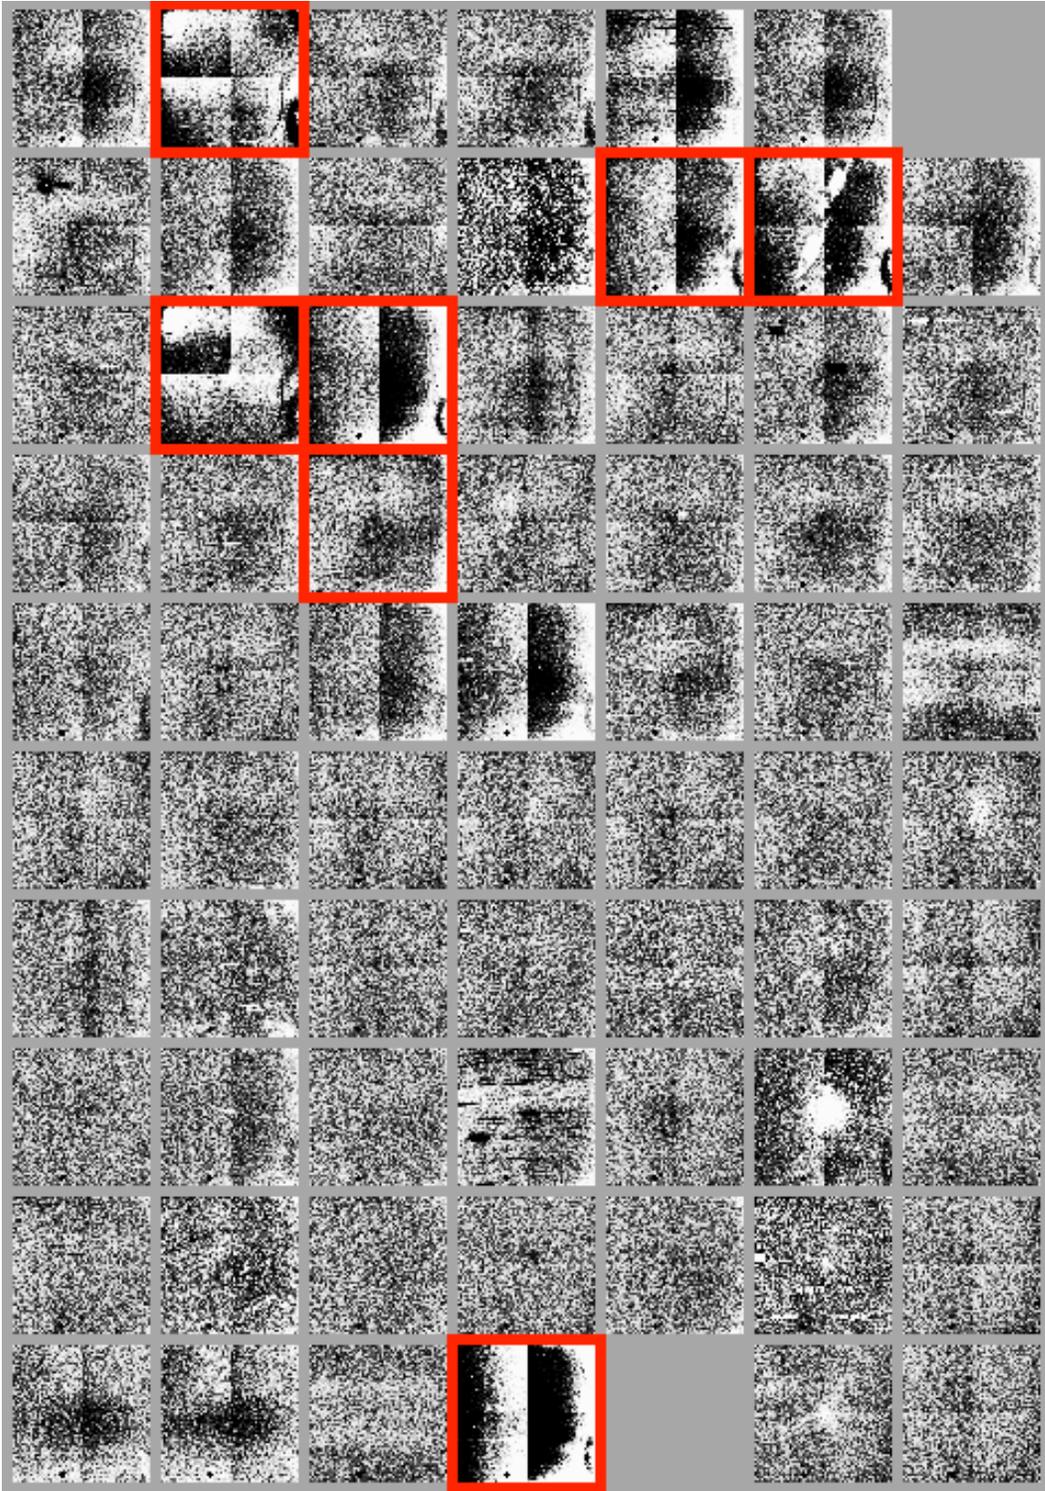


Figure 8: Similar to Figure 6, but after the mean value of each quadrant is subtracted from that quadrant. This enhances the visibility of high frequency variations in signal. This figure is shown with a histogram equalization stretch from -0.005 to $+0.005$ e/sec (i.e. about 12% of the nominal dark current rate-measured in SMOV data). Files corresponding to the five circled points in the lower half of Figure 1, which had very high persistence signals from prior flat field are highlighted in red.

Figure 9 shows the same quadrant by quadrant subtraction performed on the TV3-derived dark current reference file. Again, the brightening on the right edge is present, but subtle. This implies that the source of the bright right edge was present in ground testing as well as on orbit.

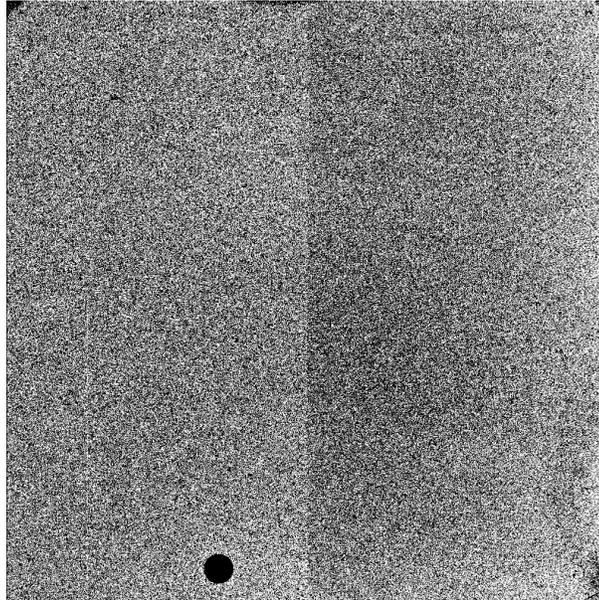


Figure 9: Similar to Figure 8, but using the TV3 dark current reference file. The “nominal” dark current frame from the on-orbit dataset has been subtracted, as has the mean signal value in each quadrant. This image appears flatter than most in Figure 8. A slight brightening of the right edge is apparent, as is the wagon wheel.

Figure 8 clearly shows the sporadic appearance of a bright right edge of the detector, as was first noted in Figure 6, with the area of elevated signal level growing as one moves towards the upper and lower corners of the detector. This elevated signal appears to be present, often with the wagon wheel visible inside it, in about a third of the SPARS200 files. Positive and negative residuals from persistence are also visible.

Dark Current Signal Contributors

Consideration of Figures 4, 6, and 8 suggests that the following signals are present in WFC3/IR dark current data:

- 1) the intrinsic, spatially variant dark current
- 2) residual persistence
- 3) a temporally and spatially variable extra signal

Intrinsic Dark Current

Images in the lower few rows of Figure 4 as well as Figure 5 clearly show the large-scale dark current features. This includes the higher signal “doughnut” on the left side, with the lower signal area below it, along with the higher signal lower right quadrant, and the lower signal, blotchy appearance of the upper right quadrant.

Persistence

As mentioned previously, we did not attempt to find optimal parameters for persistence subtraction. The residual persistence sources in Figures 4, 6, and 8 show that persistence signals were not completely removed from the dark current images. However, Figure 6 shows that in almost all cases, persistence affects a small fraction of the total pixels on the detector. Also, the extreme stretch (from -0.005 to $+0.005$ e^-/sec , or roughly $\pm 10\%$ of the mean dark rate) shows that the magnitudes of the residual persistence signals after persistence correction are small. These factors suggest that persistence caused by external point and extended sources should not significantly affect the detector-wide mean dark rates seen in Figure 1. The only persistence-derived signals which could potentially affect enough of the detector to contaminate the mean would be persistence following flat field observations. Some of the dark current observations in this dataset were made within several hours of flat field observations (highlighted in red in Figures 4, 6, and 8). However, given that the measured mean signals of the (corrected) flat-field-contaminated files (points circled in red in Figure 1) are not significantly above the mean dark rate measured in the remainder of the data, we believe that persistence cannot be the primary cause of the variability of the measured dark rate in Figure 1. Further support for this hypothesis comes from the presence of the relatively high dark rate seen in frames where the detector was not used for many hours prior to the dark current observation. This is true for several images in the top few rows of Figure 6.

Figure 10 below shows the effects of the persistence correction on all of the input files. The plot shows the mean dark current measured in both the uncorrected and persistence-corrected *flt* files versus the signal accumulated in the last minus first reads of the corresponding raw ramps (i.e. including persistence signal). With two versions of each *flt* file included in this plot, the points occur in pairs. For example, the light blue diamonds show the signals measured in the uncorrected SPARS100 files. Just below each light blue diamond there is a dark blue diamond, corresponding to the persistence-corrected version of the same file. Since the persistence correction is performed on the *flt* file but not the raw, the correction has the effect of moving points vertically downwards in this plot. The same is true for the orange (uncorrected) and red (corrected) crosses of the SPARS200 files. (Note that the raw signals of the SPARS200 files have been scaled to the SPARS100 exposure time, in order for both datasets to fit on the plot.)

The mean distance between the uncorrected and corrected points across all of the input files is 0.0003 e^-/sec , as expected given that the components of many of the light

blue/dark blue and orange/red pairs of points lie atop of one another. This small difference in mean dark current rate resulting from the persistence correction is not a significant contributor to the variations seen in the lower panel of Figure 1.

Finally, the seven files collected after internal flat field observations are shown by pairs of light and dark green asterisks. Note the much larger vertical offset between each pair of points compared to those in the blue and red points. Curiously, with several of the darker green asterisks, representing the persistence-corrected *flt* files, falling well below the main population of points on the plot, it appears that the default parameters in the persistence correction algorithm tend to over-estimate the level of flat field-induced persistence. This may lead to a small underestimate in the number of hot pixels in these images, although the level of over-subtraction seen below is still small compared to the threshold value for hot pixels ($0.4 \text{ e}^-/\text{sec}$).

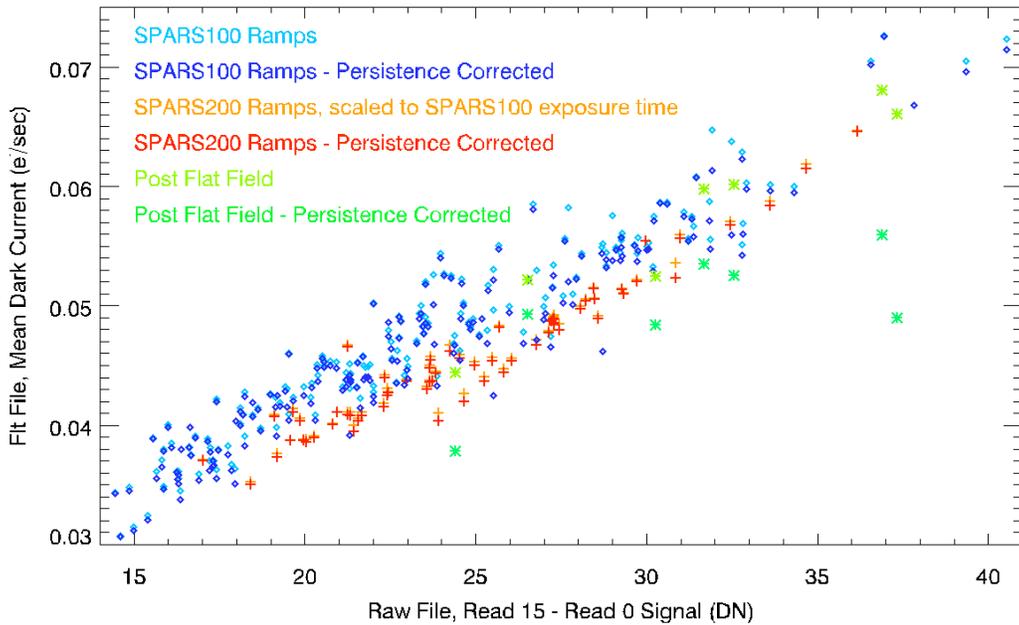


Figure 10: Mean dark current in the persistence-corrected and uncorrected *flt* files versus the signal measured up the ramp in the raw, non-persistence-corrected files. The raw signal values in the SPARS200 files were scaled to match the SPARS100 exposure time, in order to easily compare the two sets of points. Persistence correction will move points vertically downwards in this plot, since the persistence correction is not performed on the raw files. Note how each dark blue diamond is directly (vertically) below a corresponding light blue diamond. Similarly, each red cross is vertically offset from an orange cross. The amount of offset between the pairs of points shows the magnitude of the change in the mean dark current due to the persistence correction. Asterisks mark the seven ramps which were collected shortly after internal flat fields, resulting in varying levels of persistence. Note the larger difference between corrected and uncorrected points for these seven files.

Extra signal

The idea of the dark current increasing in a non-uniform manner across the detector prompted a different look at the data in Figure 3. While subtracting a base dark current image from all other images removed large-scale dark current features and left behind the “extra” signal, it did not give a clear indication of how the dark current increased from image to image. By examining the relative rate at which the dark current increased from image to image across the detector, we were able to create a map that gives a better idea of how the dark current changes. Note that the image and plots in Figures 11 and 12 were constructed using the SPARS100 data, as the larger number of SPARS100 images yielded lower uncertainties. We repeated the analysis using the SPARS200 data and found the results to be similar.

Using all 182 of the **SPARS100** *flt* files, we began by binning each *flt* into boxes that were 26 pixels on a side, resulting in an array of 39x39 boxes. Each box was given the value of the sigma-clipped mean of the pixels it encompassed. This gave us a stack of 182 images, each binned into a 39x39 array. For each box, we plotted the values through the stack of 182 images versus the corresponding mean dark rates (calculated from the histogram peaks, like those in Figure 1). This produced a plot similar to those shown in Figure 11. The slope of a best-fit line to this plot provided a measure of how much the dark current within this box increased relative to the amount which the overall mean dark current increased. For example, a slope of 1.0 would indicate a box where the image-to-image increase or decrease in the dark current matched that in the overall detector. A slope of 1.10 would indicate a box where the dark current increased 10% more from image-to-image than the overall mean dark rate. Conversely, a slope of 0.90 would indicate an area where the mean dark rate increased 10% more slowly than the overall mean dark rate.

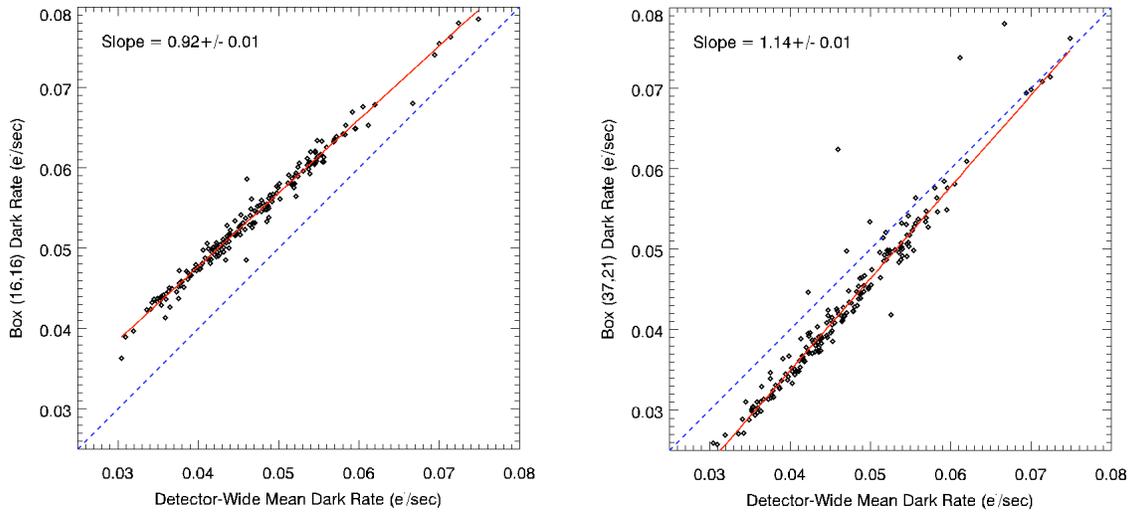


Figure 11: Plots of the dark current within two 26x26 pixel boxes versus the detector-wide mean dark current. The plot on the left shows a box where the dark current increases more slowly than the overall mean, while the plot on the right shows a box located along the right edge of the detector, where the dark current increases more quickly than the overall mean.

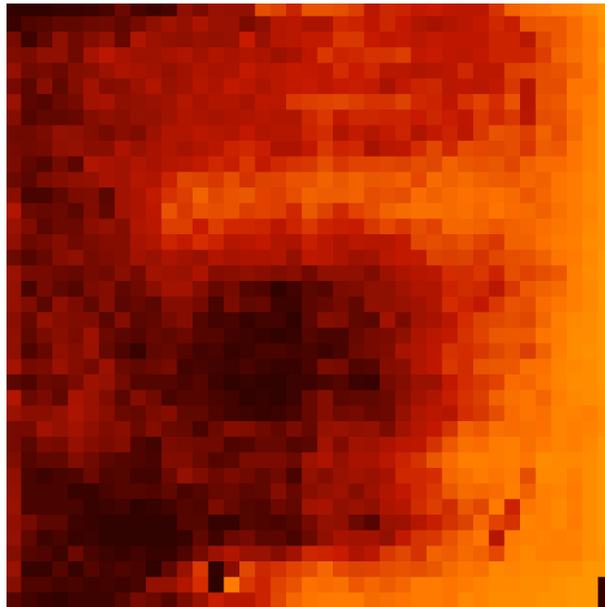


Figure 12: Map of relative dark current changes within the SPARS100 files. Bright areas show where the dark current on the detector increased more than average, while in dark areas, the dark current increased less than average. The figure is shown in histogram equalization scale, from 0.8 to 1.35. (ie, the brightest areas had dark current increase 35% more than the detector-wide mean between the dimmest SPARS100 file and the brightest SPARS100 file).

The result of the line-fitting to all of the 39x39 boxes is shown in the map in Figure 12. This figure is shown with a histogram equalization stretch to bring out the differences in the best-fit line slopes across the detector. The values in this image range from 0.8 to 1.35. This map shows that the signal increases most quickly along the right and bottom edges of the detector, as well as within a horizontal bar stretching most of the way across the detector, roughly two thirds of the way up from the bottom. Beneath these features, there also appears to be a subtle left/right asymmetry in the map, with quadrants 3 and 4 (lower and upper right) having dark current values that increase slightly faster than those in quadrants 1 and 2.

Examination of the data for each of the boxes revealed plots very similar to those shown in Figure 11, with the points behaving in a very linear fashion for the entire length of the plot. **This behavior suggests that the spatial pattern of signal increases seen in the map above are permanent features.** If the bright edge was the result of a time-variable light leak or external illumination of the detector, then we would not expect to see such consistent, linear behavior across all mean dark current levels.

It also appears that there is no correlation between the mean dark rate in a given area of the detector and its' amount of change from image to image, i.e. there's no obvious spatial correlation between the map in Figure 12 and the maps in Figures 4 and 5. So the changes from image to image do not appear to be either a multiplicative or additive increase in the overall dark current pattern. This would argue that whatever is causing the fluctuations is perhaps not directly associated with the regular dark current itself.

Investigation of Potential Sources of Extra Signal

Figure 12 shows that after removing the large scale features associated with the true dark current, the observations show an extra signal which appears to increase in a non-uniform way across the detector. There are a number of possibilities besides an unexpected illumination of the detector (i.e. light leak) for the source of such a signal. Changes in detector temperature or voltages associated with the IR channel could cause a change in dark current. Unanticipated signals in the reference pixels during a ramp could cause errors during the reference pixel subtraction. However, all of these possibilities seem unlikely, given the morphology seen in Figure 12. Reference pixel and voltage-induced effects are most likely to be uniform across the detector, or possibly variable on a quadrant-by-quadrant basis. Non-uniform temperature across the detector would appear to be the most likely culprit. We examined many of these possibilities and show the results below.

Instrument Temperatures, Voltages, and Pointing Direction

Using the dark current values from Figure 1, we searched for a correlation of dark current rate with a large number of instrument temperatures and voltages. The complete list of telemetry monitors examined in this comparison are given in Table 2 in the Appendix. We were able to locate only one telemetry item which showed a correlation with the measured dark rate. This was the voltage input to WFC3 from the HST bus. During ground testing, Bushouse (2009) found a correlation between the bus voltage and the measured dark rate. We find a similar correlation here, with a lower bus voltage corresponding to a lower measured dark current rate. Figure 13 shows that while this general trend is present, the ramp-to-ramp scatter in the dark current rate is large compared to the variation due to the changing bus voltage.

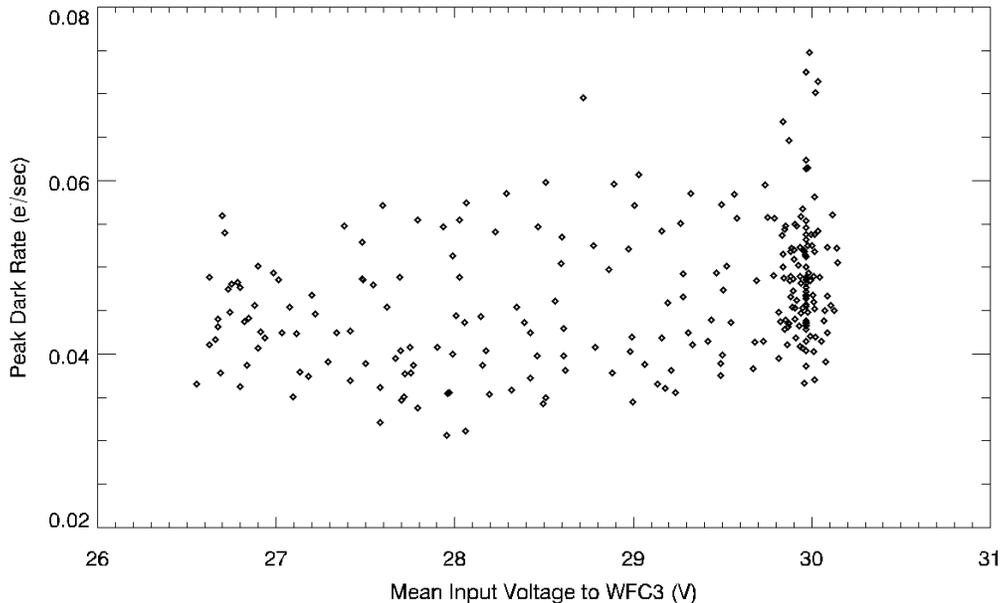


Figure 13: Mean dark current rate versus the voltage input to WFC3 from the spacecraft bus. While the mean dark current increases slightly with increasing voltage, the ramp-to-ramp scatter still dominates the dark current variations.

One possibility for the muted level of correlation could be the fact that the input voltage varies on the timescale of an orbital period, as seen in Figure 14. The SPARS100 and SPARS200 dark current ramps in this analysis have exposure times equivalent to roughly 25% and 50% of an orbit, meaning that the reported dark current rate for a given file is potentially averaged over large changes in input voltage. We attempted to work around this fact by collecting statistics about the bus voltage behavior within each observation. For each ramp, we collected the maximum, minimum, and mean values of

the input voltage, as well as the maximum, minimum, and mean of the first derivative of the voltage.

Note that the voltage as plotted in Figure 14 is ‘change-only’ data. This means that the only points present in the data are at locations where the voltage level changed. For example, in Figure 14 the voltage is listed as 29.9 volts at a time of 128 minutes. The next point to the left shows the voltage at 30.1 volts at a time of 146 minutes. Despite the lack of points between these two values, measurements of the voltage were in fact made during this time period. However, after the voltage level was measured at the 128 minute mark, it remained constant at 29.9 volts up until minute 146, when it changed to 30.1 volts. In order to minimize the data volume in the resulting file, only the points where the voltage level changed were saved. Before computing derivatives or comparing voltages to measured dark rates, we re-sampled this data into an array with a 1-minute time resolution, effectively re-inserting the previously mentioned points which were thrown out, in order to calculate accurate time-averaged values.

Figure 15 shows a plot of the dark current rate versus the mean of the first derivative of the input voltage. Points near zero delta voltage indicate observations made largely during the flat peak of an orbital cycle like those in Figure 14. Positive values of delta voltage correspond to observations made while the input voltage was increasing, and negative values of delta voltage indicate times when the input voltage was decreasing. The plot shows that observations made while voltage was increasing tended to have slightly lower dark current values than observations made when the voltage did not change, or was decreasing. This agrees with the plot in Figure 13 above, where higher voltage corresponds to higher dark rates. At times when the voltage was increasing the most quickly, the absolute value of the voltage was low, as seen in Figure 14, and the corresponding dark current was also low. This correlation appears slightly more significant than that of mean dark rate versus absolute value of the input voltage, as the magnitude of the overall trend here is comparable to the ramp-to-ramp scatter of the points.

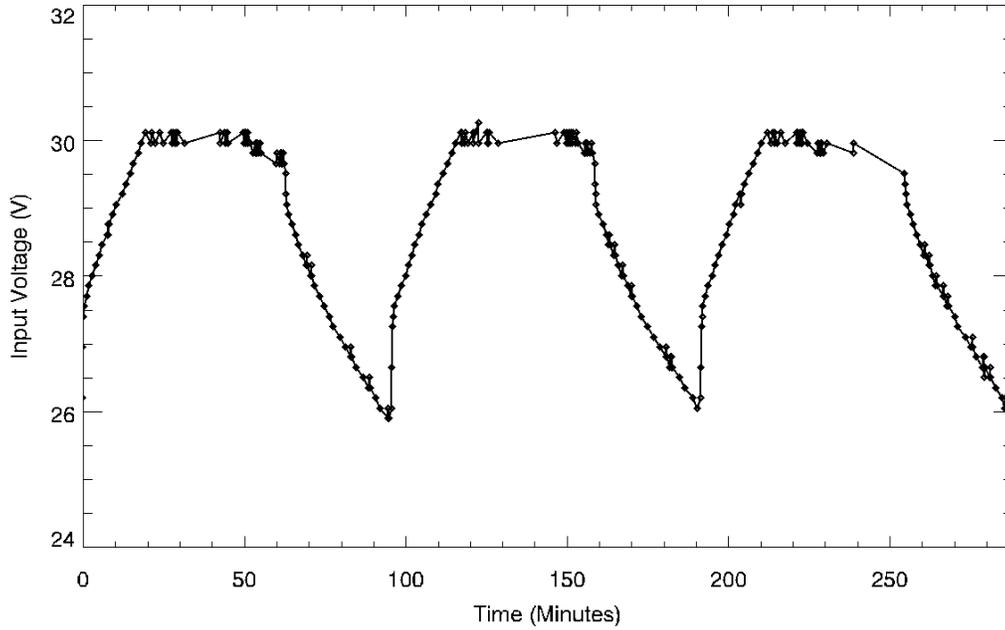


Figure 14: Voltage input to WFC3 over 3 orbits.

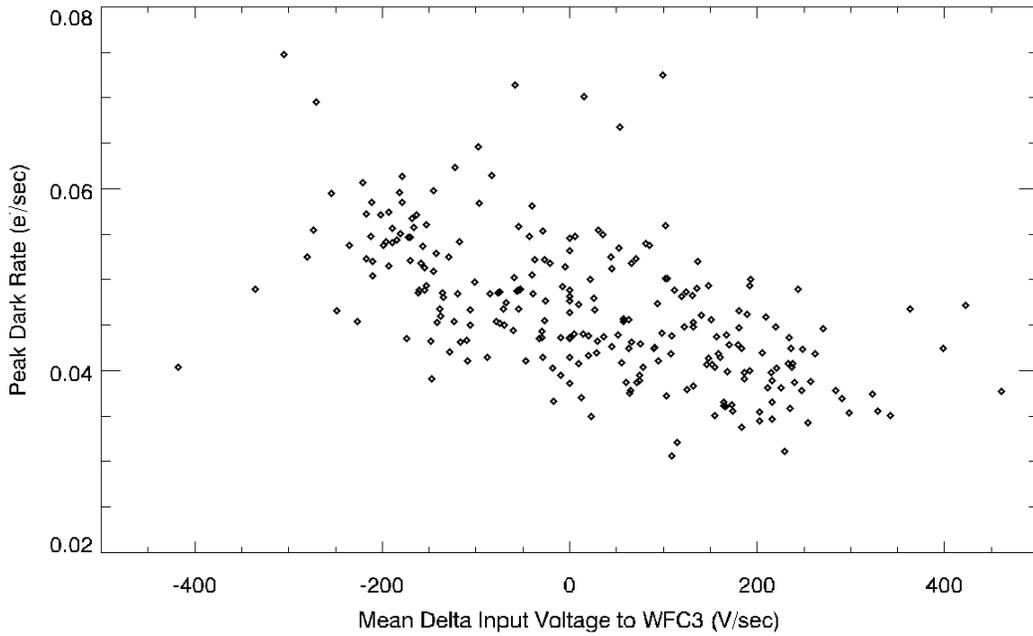


Figure 15: Mean dark current rate plotted against the first derivative of the voltage during the exposure. Large positive values of delta voltage correspond to observations made while the voltage was increasing during the orbit. Values close to zero indicate that most or all of the observation was made while the voltage hovered near 30 V during the peak of a cycle like those shown in Figure 14.

We also repeated the analysis above and examined the number of hot pixels versus voltage and the change in voltage. Corresponding plots are shown in Figures 16 and 17. In this case, there appears to be no clear relationship involving the number of hot pixels. One might expect the number of hot pixels to increase as the mean dark rate increases, in which case Figures 16 and 17 would appear similar to Figures 13 and 15, but this appears not to be true. Further analysis on the hot pixels is presented in the Hot Pixel Population section below.

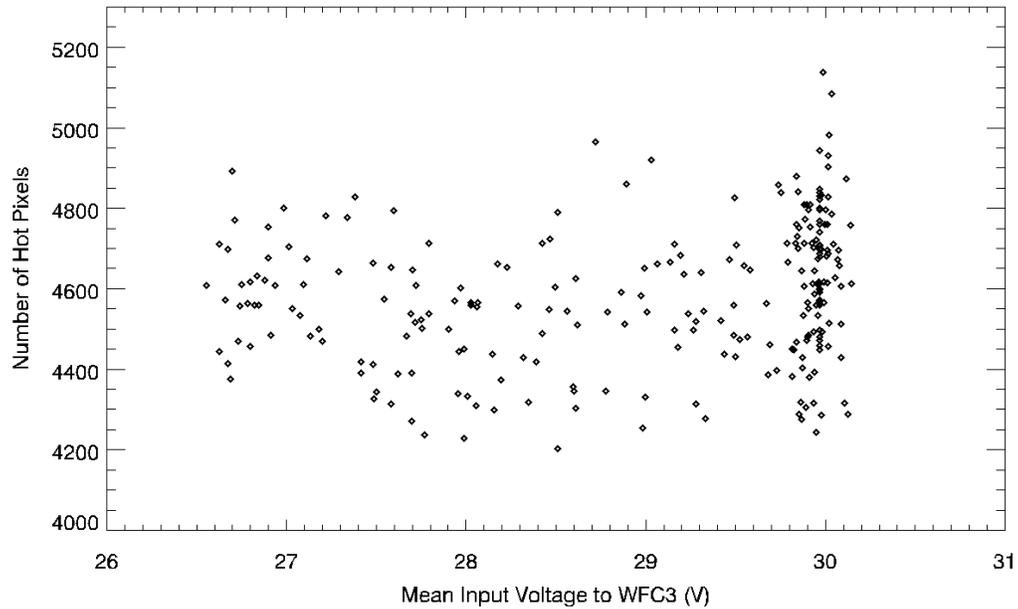


Figure 16: Number of hot pixels in each ramp, plotted versus the input voltage.

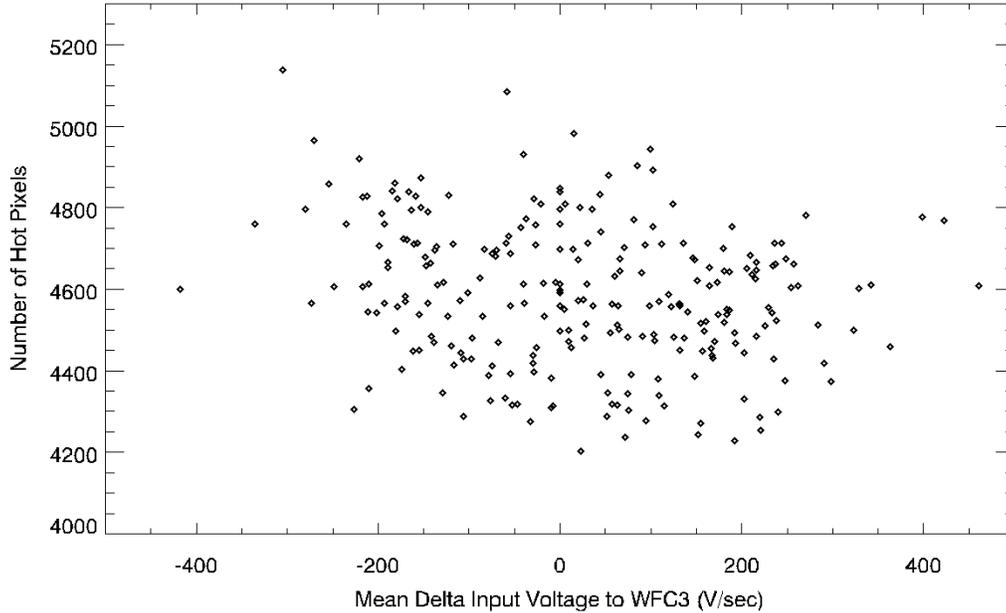


Figure 17: Number of hot pixels plotted against the mean first derivative of the input voltage.

Finally, we also compared the mean dark rate in the input images to the angle of the sun relative to the V1 direction, which points essentially out along the optical axis of the spacecraft. No correlation was found with mean dark rate, implying that a light leak coming down the optical axis of the telescope is not the cause of the varying mean dark current level.

Bias Level

Moving away from the search for primary causes of the dark current variation, we focused on correlations of the dark current with other detector measurements. First, we compared the mean bias level in each frame to its measured dark current rate. To find the bias level, we calculated the sigma-clipped mean signal value in the vertical inboard reference pixels in the 0th read of the raw (prior to any data reduction steps) version of each file. In order to avoid any effects of persistence, we do not present bias levels derived from the science pixels.

Figure 18 shows a plot of the mean bias level in the reference pixels of quadrants 1 (in color) and 3 (in black) in each ramp versus the mean dark rate, where the colors match those shown in the dark current plots in Figure 1. This figure is very similar to Figure 8 in Hilbert (2012), which presents a much more detailed study of reference pixel behavior over time. Most obvious in this plot are the several large (~80-100 DN) jumps in the bias level. While there is a hint of a change in the peak dark current in Figure 1 at about the same time as the large bias jump at MJD 55325, the change in the dark current is much

smaller than the point-to-point scatter. Our calculations reveal a change in the mean dark current of $0.0001 \text{ e}^-/\text{sec}$ between observations with the higher bias level versus those with the lower bias level. We do not consider this level of change to be significant. Further, a scatter plot of dark current versus bias level, shown in Figure 19, reveals no clear relationship.

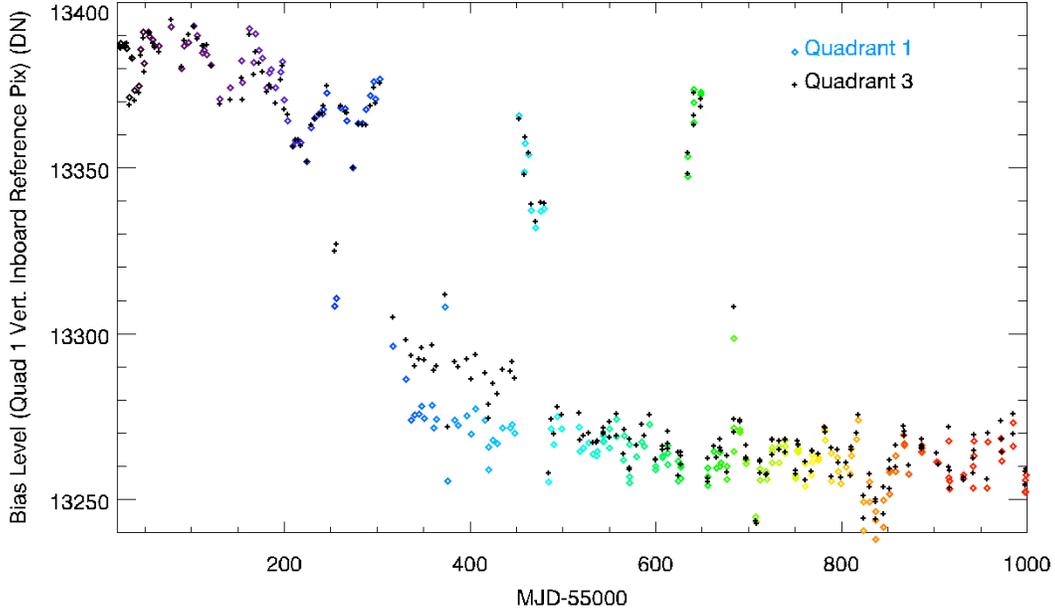


Figure 18: The mean reference pixel signal in the zeroth read for quadrants 1 (in color) and 3 (in black). The colors in this plot correspond to those in Figure 1. Note the collection of light blue points centered around MJD 55400 where the quadrant 3 values are elevated compared to the quadrant 1 values.

A second behavior to note in Figure 18 is the relatively large difference between the quadrant 1 and quadrant 3 reference pixels in the MJD 55325 – 55450 time period. Quadrant 3 is the lower right quadrant of the detector while quadrant 1 is the upper left, indicating a left/right asymmetry in reference pixel behavior. Given the bright right edge of the detector seen in Figure 12, we compared the reference pixel signals with those from the science pixels to determine if the cause of the bright science pixels was also affecting the reference pixels. Figure 19 shows the scatter plot of the dark current rate in the right-most 100 columns of science pixels versus the signal in the reference pixels. The large jump in bias signal at MJD 55325 separates the points into two populations in the x-direction. However, the plot reveals no clear relationship between the dark rate and the bias signal. The red points in Figure 19 correspond to the images taken between MJD 55325 and 55450 which are seen (in blue) in Figure 18 to have relatively large differences between quadrant 3 and quadrant 1 reference pixel signals. While these red points show the increased reference pixel signal (relative to the black points in the low-

reference-pixel-signal population), there is no clear relationship with dark current rate. This suggests that the variations in the signals of the reference pixels do not translate into variations in the overall mean dark current, nor into the bright science pixels along the right edge of the detector.

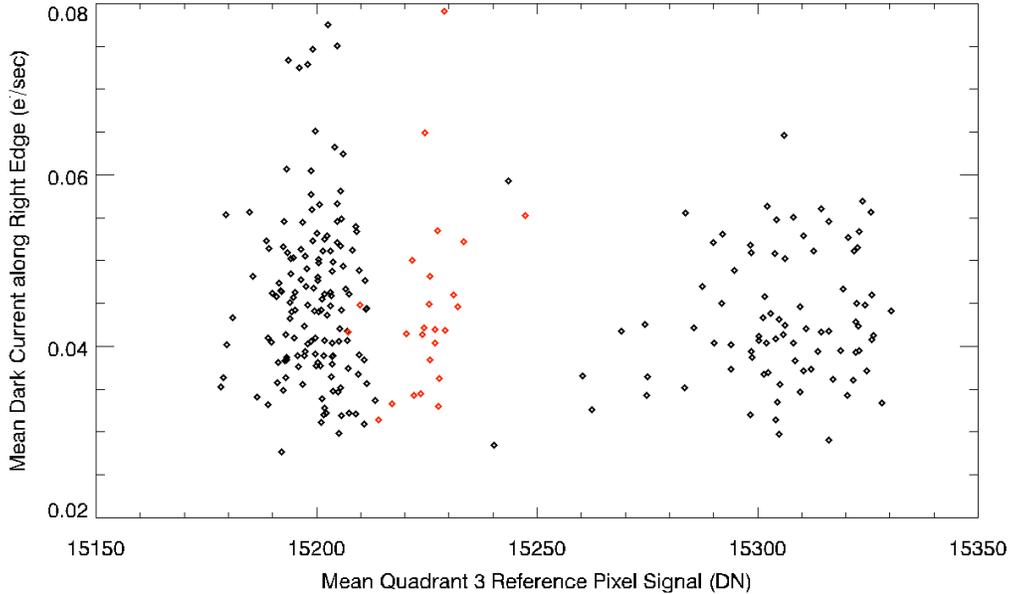


Figure 19: Plot of the mean dark current measured in the right-most 100 columns of science pixels, versus the mean signal in the quadrant 3 reference pixels. The two populations of points on the left and right are due to the large bias jumps seen in Figure 18. Red points correspond to the observations between MJD 55325 and 55450, where the quadrant 3 reference pixels exhibited elevated signal relative to quadrant 1. No correlation is evident in this plot. Plots using reference pixels from the other quadrants, as well as the overall mean dark rate rather than that in only the right edge appear similar, in spite of the differences between quad 1 and quad 3 reference pixels seen in the blue points in Figure 18. SPARS100 and SPARS200 files are mixed evenly throughout the two populations.

Intra-Ramp Reference Pixel Effects

While the mean bias level as measured by the reference pixels in the 0th read appears uncorrelated with the dark rate, we also wished to examine the behavior of the reference pixels in all of the reads going up the ramp. If the reference pixels were somehow accumulating signal during the ramps, then reference pixel subtraction could be introducing a false detector-wide signal into the science pixels. Figure 20 shows that this is not the case.

This figure contains three curves. The central red curve, which corresponds to the red axis on the right side of the figure, shows the same mean dark current rates as seen in the

lower panel of Figure 1, but reordered by increasing brightness. This was done to facilitate comparison with Figures 4, 6, and 8. The upper black curve shows the detector-wide mean signals measured in the science pixels in the original raw files. These raw files have not had reference pixel correction, zeroth read subtraction, nor persistence correction performed. We manually subtracted the zeroth read of each raw file from the final read, and calculated the detector-wide mean. The overall trend of this line, with the mean signal increasing with file number, agrees with the red curve from the processed data. However, with no persistence correction performed on the raw data, the effects of persistence can be seen, most notably in the green points, which match the points circled in red in Figure 1, and indicate files with significant persistence from prior flat field observations. For several of these points, we see the mean signal from the raw file is significantly higher than that of the surrounding files. However after the persistence correction (red curve) this excess signal is not present, suggesting a successful removal of persistence signal. The bottom blue line in Figure 20 shows the mean signal level for the reference pixels in the final read minus zeroth read difference. The lack of large peaks in this curve at the location of the green points shows that wide-spread persistence from prior flat field observations was not seen by the reference pixels.

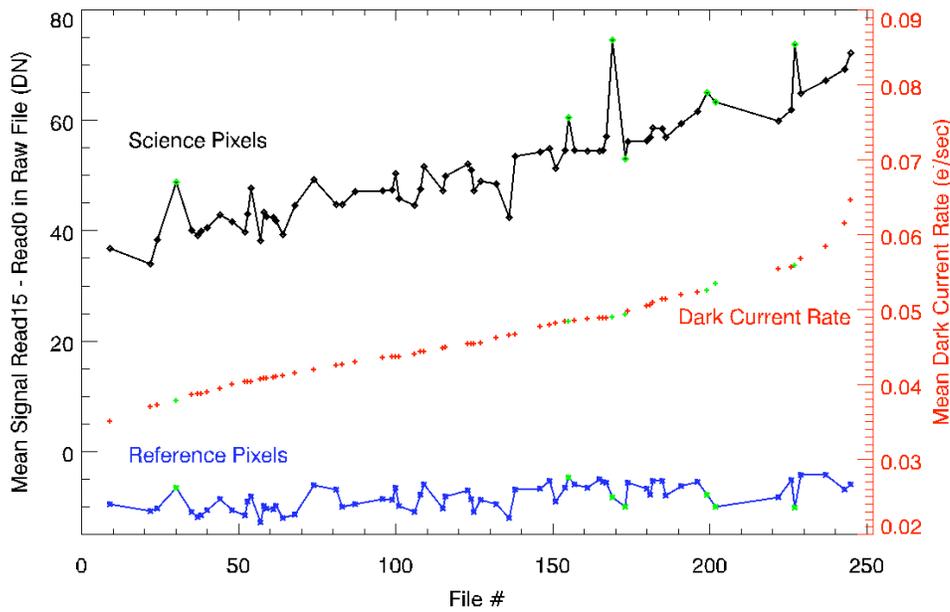


Figure 20: Mean dark rate measured in the *flt* files (red curve, corresponding to red axis), along with the mean signal in the corresponding raw ramps (final read minus zeroth read) prior to *calwf3* and persistence correction. The bottom curve shows the mean reference pixel signal from the raw file (final read minus zeroth read). Peaks in the raw file data (such as for file #1) are due to persistence, as the reference pixels show no corresponding peaks, which would imply a shift in bias level. The lack of a peak at the same spot in the red curve means that the persistence correction algorithm brought the mean signal level down to expected levels.

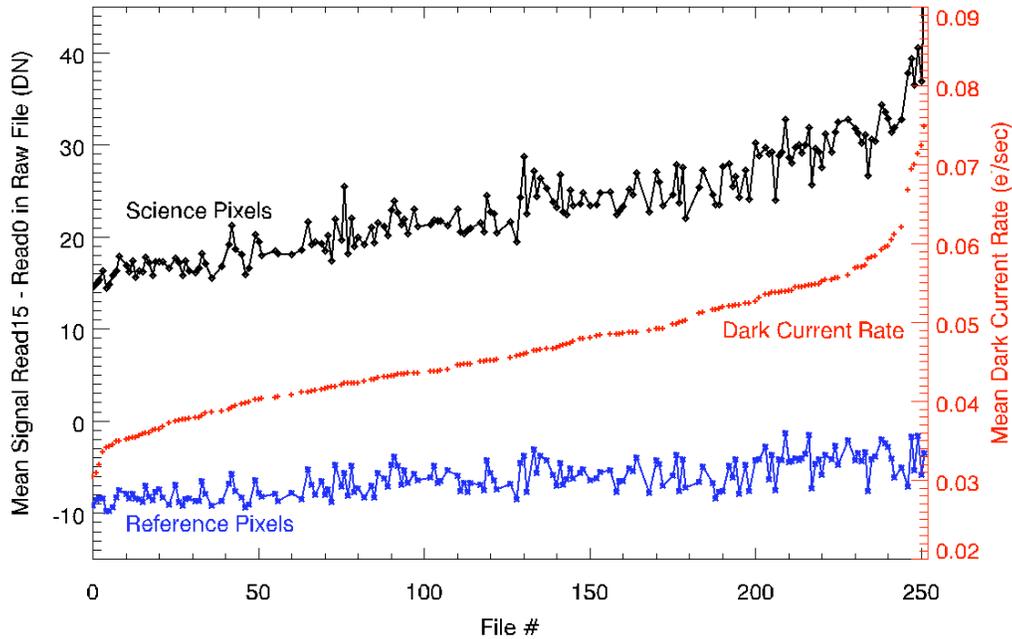


Figure 21: Same plot as Figure 20 above, but for the SPARS100 data. With more than twice as many observations compared to the SPARS200 data, this plot gives a fuller picture of the science pixel/reference pixel/dark current relationship.

Beneath the Extra Signal

Under the assumption that the map of signal rate increases shown in Figure 12 represents the true, intrinsic dark current behavior of the IR detector, we wished to see if there were any other signals buried beneath those associated with the map. In order to accomplish this, we scaled the map by the mean dark current of each of the individual SPARS100 dark current images (seen in Figure 1), added the map of best-fit intercept values (also created from plots such as those in Figure 11) and subtracted this scaled image from the *flt* image. Similar to Figures 6 and 8, Figures 30 through 32 in the Appendix show a gallery of the original *flt* images for all of the SPARS100 darks, and a gallery of the extra signal-subtracted darks, respectively. The gallery of subtracted images appears similar to Figure 8. There are several recurring patterns in the images, most notably a residual right/bottom edge difference from the rest of the detector. Unlike in the other galleries, the right/bottom edge appears alternately as both brighter and dimmer than the other areas of the detector in the images. Beyond this pattern, other effects which are visible include persistence, subarray “banding” (defined in Dulude et. al 2011), and left/right asymmetries. The stretch of the images in Figure 32 is $\pm 0.005 \text{ e}^-/\text{sec}$ (identical to that in Figure 8), or roughly $\pm 10\%$ of the overall mean dark current rate, implying that even if the scaling and removal of signal in Figure 12 was successful at removing the intrinsic dark current signal, that large spatial scale variations of roughly $\pm 10\%$ of the mean dark current rate remain.

Hot Pixel Population

The next step in characterizing the dark current behavior was to monitor the population of hot pixels present on the detector. We remained consistent with past dark current studies by using $0.4 \text{ e}^-/\text{sec}$ as the threshold value for defining a hot pixel (Hilbert and McCullough 2009). We measured the number of hot pixels present in each file (ignoring for the moment any pixel flagged as unstable) and plotted this number versus time in Figure 22. At first glance, it appears that the number of hot pixels increased with time up to roughly MJD 55700, and then remained roughly constant, or even decreased. However, closer inspection reveals step functions in the plot. The most obvious of these occur at MJDs 55325 and 55450. Curiously, these dates match up with the jumps in bias level seen in Figure 18. Further, there appear to be abrupt jumps in the number of hot pixels coincident with the dates of all of the bias level jumps in Figure 18. **A plot of the number of hot pixels versus the bias level, shown in Figure 23, confirms that there is a relationship between the quantities, with lower bias signal corresponding with larger counts of hot pixels.** The physical reason behind this relationship is not clear, especially given the lack of relationship between the bias level and the mean dark current, seen in Figure 19.

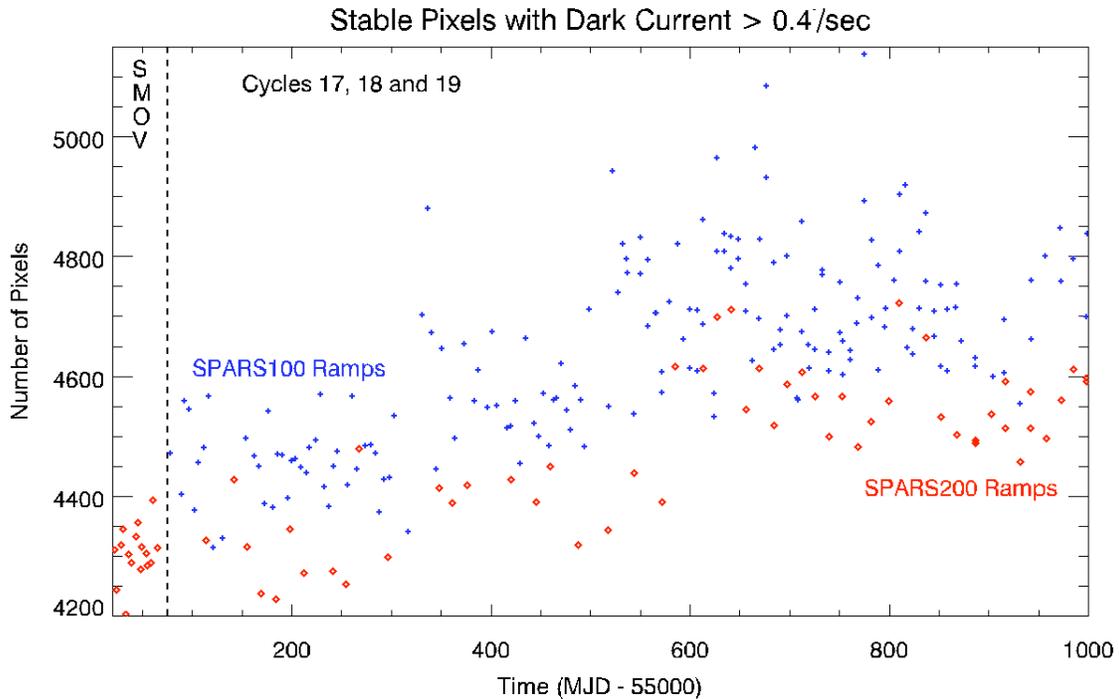


Figure 22: Plot of the number of hot pixels present in each of the input images.

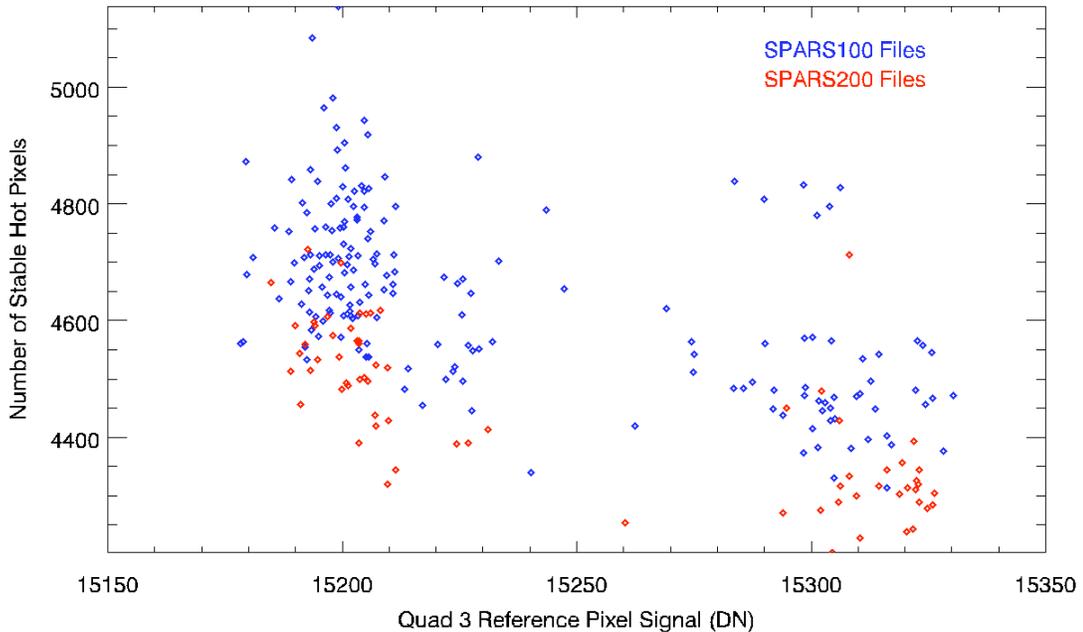


Figure 23: Number of stable, hot pixels plotted against the bias level for all of the SPARS100 and SPARS200 files. Considering the populations associated with the two sample sequences separately, there does appear to be a relationship between hot pixels and bias level. This is contrary to Figure 19, which shows no relationship between dark current and bias level (SPARS100 and SPARS200 points are homogeneously mixed within Figure 19).

Another important detail to note is that while Figures 22 and 23 show a minimum of 4200 hot pixels across the dataset, the truth is that there are not 4200 hot pixels, to which several hundred more are added in each observation. Of the roughly 4,200-5,000 hot pixels, **only 3,195 are hot in every one of the 252 dark current files**. Note that this number is well below the range of hot pixels plotted in Figure 22, implying that for each observation at least ~ 1000 of the hot pixels are not hot all the time. There are several behaviors wrapped up within this population of inconsistently hot pixels. Some pixels initially have dark current rates below the hot pixel limit and then become hot as time passes, while others have been seen to be hot and then have their dark rate decrease below the hot pixel threshold over time. Others still have overall mean dark rates very close to the $0.4 \text{ e}^-/\text{sec}$ hot pixel signal rate limit, and rates in individual ramps which wander above and below the hot pixel limit from ramp to ramp. Closer examination of the data revealed that in each one of the 252 dark current images, there were between 5 and 10 pixels which were hot in only that image and no others.

Up to this point, we had only considered good pixels (i.e. no flags present in the data quality arrays) for our hot pixel analysis. With a large percentage of the unstable pixels (data quality flag of 32) on the detector also flagged as hot, we repeated the measurements above on the population of unstable pixels, in an effort to see if the behavior was similar to that seen in the stable pixels. Figure 24 shows the number of hot pixels versus time for the unstable pixel population.

As with the stable hot pixels, there appear to be discontinuities in the curve corresponding to the times of the large bias level jumps, particularly at MJD 55325, 55450, and 55625. Focusing on the blue SPARS100 points, it appears that between bias signal jumps, the hot pixel population slowly decreases with time, and then quickly increases when the bias signal jumps. Looking at the red SPARS200 points, it also appears that the unstable hot pixel population was at a minimum during SMOV. After SMOV the population was suddenly larger, and has been slowly decreasing over time, eventually reaching the level seen in SMOV.

As was the case when looking for the cause behind the extra signal in the dark current images, we found no correlation between the number of hot pixels with any of the instrument temperatures and voltages or angle of the sun relative to the spacecraft.

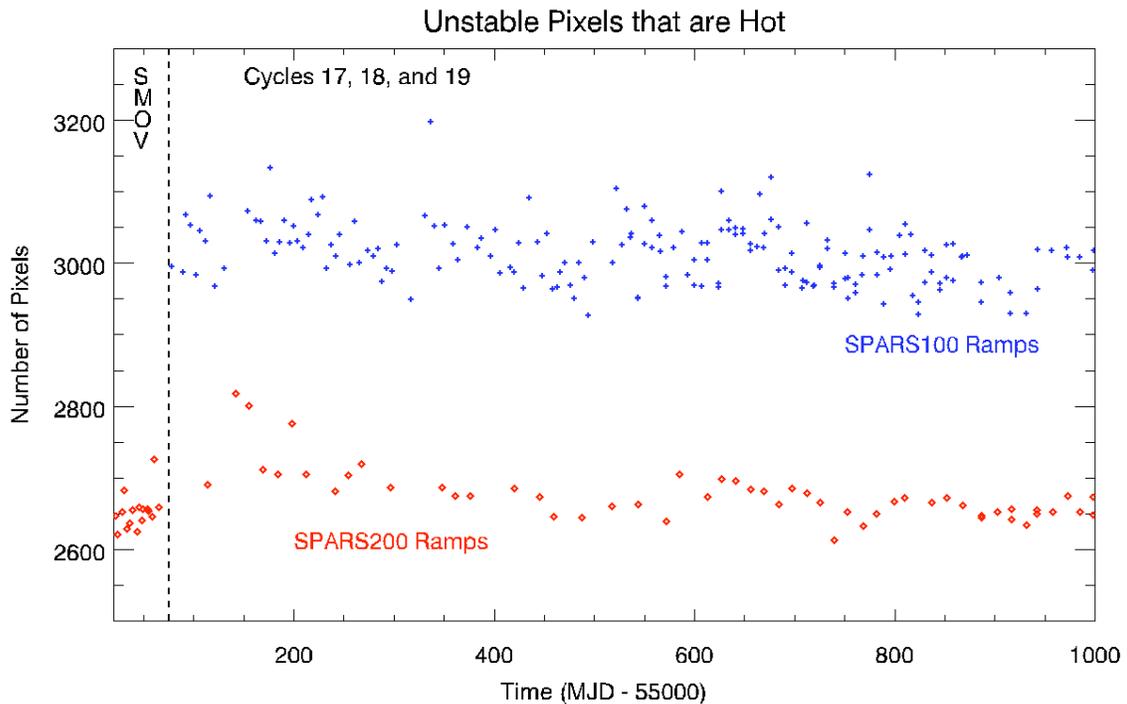


Figure 24: Plot of the number of unstable pixels which are hot in each input image.

Looking more closely at the frequency with which individual pixels were hot, we created a histogram of the number of hot pixels versus the percentage of ramps in which those pixels are hot. The results are shown in Figure 25, with separate curves for the SPARS200 and SPARS100 sample sequences, due to the different exposure times (and corresponding noise levels) and number of ramps for each. Focusing on the SPARS100 (blue) results, we see on the right edge of the plot that 3,290 pixels are hot in all 183 ramps. On the left side, 700 pixels are hot only once within the 183 ramps (corresponding to 0.5% of the ramps on the x axis). In the center of the plot, 7 pixels are hot in half of the input ramps (92 out of the 183 ramps). Note that there is no information here about which 92 ramps those 7 pixels are hot in. Each of those 7 pixels could be hot in a different set of 92 ramps. The idea behind this plot is to characterize the variability of the hot pixel population, and give an idea of the number of pixels displaying dark rates which are variable to the point that those pixels are flickering above and below the hot pixel limit, and how often they are doing it. Excluding the pixels which are hot in all input ramps, and totaling the number of pixels in all other points in Figure 25, we find that 4849 pixels do not remain exclusively above or below the hot pixel limit for the duration of the SPARS100 data set. In the case of the 69-ramp SPARS200 dataset, this number is 3546 pixels. As mentioned previously, many of these pixels have mean dark rates which are very close to $0.4 \text{ e}^-/\text{sec}$ and need only to vary by a small amount from ramp to ramp in order to jump above and below the $0.4 \text{ e}^-/\text{sec}$ threshold, creating the appearance of instability.

However, other pixels in this population have dark current rates which vary wildly from ramp to ramp, while others have been observed to have a nominal dark current rate ($\sim 0.04 \text{ e}^-/\text{sec}$) for almost all of the ramps, but then jump up to a dark rate above the $0.4 \text{ e}^-/\text{sec}$ threshold in only one or two ramps. These latter two populations describe pixels which could be considered unstable, depending on the magnitude of their variations. Details and descriptions of unstable pixels can be found in Hilbert and Bushouse (2010). This analysis suggests an upper limit of 4849 pixels which are potentially unstable but not in the current bad pixel table (u4c1709ri_bpx.fits). The true number will be less, with some pixels simply the victims of the noise levels shown in Figure 2. **With all previously flagged unstable pixels excluded from this analysis, Figure 25 suggests that unstable pixel searches should be performed regularly in order to identify new unstable pixels.** Examination of the location of the pixels tallied in Figure 25 reveal that they are spread evenly across the detector.

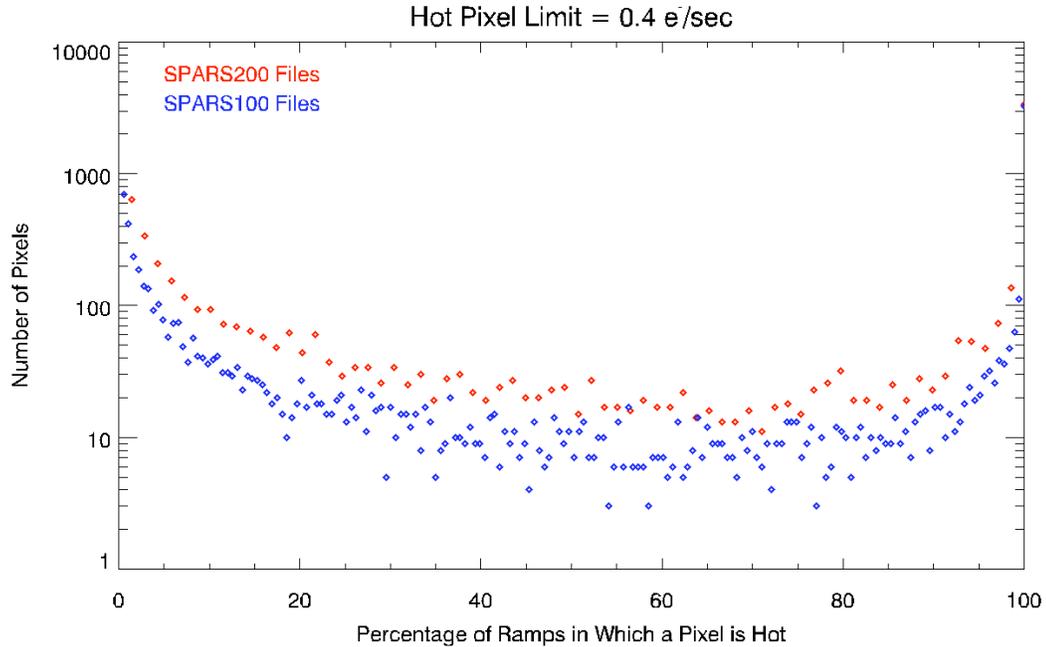


Figure 25: Plot of the number of pixels which are hot in a given percentage of the dark current ramps. We plot curves for SPARS100 and SPARS200 datasets separately due to the different noise characteristics of the two. We also plot these numbers versus the percentage of input ramps in order to fit both curves on the plot given the fact that the SPARS100 dataset contains 2.5 times as many files as the SPARS200 dataset. While not visible on this plot, the point at $x = 0$, corresponding to the number of pixels which are never hot (in any of the input ramps), would show that in the SPARS100 data, 1,020,007 pixels (out of a total of 1,028,196 or 99.2%) always have dark current rates below the hot pixel limit. 3,290 (0.3%) of the pixels are consistently hot in 100% of the SPARS100 files (seen on the right axis, far from the main distribution of points). 700 pixels are hot only once in the SPARS100 dataset (seen in the left-most blue point). Points at x -values above 0% and less than 100% indicate pixels which are inconsistently hot. This can be caused by a truly unstable pixel, or by a pixel with a dark current rate very close to the hot pixel limit, such that shot noise will at times push the measured dark current value over $0.4 e^-/\text{sec}$.

Motivated by the inconsistently hot pixels seen in Figure 25, we looked more closely at the pixel-by-pixel noise levels in the dark current data, how they compare to the (Poisson only) noise levels of an ideal detector, and how the difference between the two affects the population of inconsistently hot pixels.

Figure 2 shows histograms of the 1σ measured noise levels across the detector for the dark current images used in this analysis. By plotting this noise level versus the mean dark current level for all pixels, we are able to more easily examine the noise properties of this dataset. Figure 26 shows this plot, where the noise values are the same as those used to create the red (SPARS200) histogram in Figure 2, and the dark current values are

the sigma-clipped mean values for each pixel, through the stack of SPARS200 files. (The version of this plot created from the SPARS100 data appears nearly identical).

There are two narratives presented within this plot. First, the measured noise levels across the detector can be compared to the theoretical (Poisson) noise level for a given dark current rate by comparing the location of the points to the three dashed blue curves. These blue curves represent the noise levels associated with 1-, 3-, and 5 times the Poisson noise for the dark current rates along the x-axis. In a perfect detector, all pixels would exhibit only Poisson noise and would fall along the 1σ dashed blue line towards the bottom of the plot. Using the SPARS200 data, we see that this 1σ line acts as a lower boundary on the measured noise levels of the pixels, as one would expect. Comparing the measured noise levels to the set of Poisson noise-associated blue lines, we find that 96.5% of the pixels have measured noise levels that are between 1 and 3 times the ideal, Poisson-only noise levels. 2.6% (27,200) of the pixels have measured noise levels that are between 3 and 5 times the ideal, Poisson-only noise, and 0.8% (8445) have noise levels more than 5 times higher than the Poisson noise level.

The second story being told in Figure 26 involves the red, orange, and yellow points. In this case, we worked only with the **measured** noise of the pixels (i.e. no more comparisons to ideal, Poisson-only noise). The yellow, orange, and red points show pixels that would become inconsistently hot (like those seen in Figure 25) across a dataset upon a change in their **measured** dark rates of 1, 3, and 5σ , respectively. (Note that the yellow points correspond to a 1σ change, the orange and yellow points together are for a 3σ change, and the yellow, orange, and red points together are for a 5σ change).

The population of yellow points show pixels that, given their measured mean dark current and measured noise, would show up as inconsistently hot in a dataset if their dark rate changed by 1σ or more in one or more of the files in the dataset. For example, take the case of a normally unflagged pixel with a mean dark rate of $0.3 \text{ e}^-/\text{sec}$ and a measured noise level of $0.1 \text{ e}^-/\text{sec}$. If this pixel displays a signal in an individual dark current image that is elevated by 1σ above its mean, this implies a signal of $0.3 + 0.1 = 0.4 \text{ e}^-/\text{sec}$, meaning that the pixel would be flagged as hot in that image. An elevation in the dark rate of less than 1σ would keep the pixel below the $0.4 \text{ e}^-/\text{sec}$ threshold, and the pixel would not be flagged as hot. Similarly, for a pixel with a mean dark rate of $0.49 \text{ e}^-/\text{sec}$ and a measured noise level of $0.1 \text{ e}^-/\text{sec}$, meaning that it is normally flagged as hot, if in an individual *flt* image the pixel's signal is 1σ below its mean ($0.49 - 0.10 = 0.39 \text{ e}^-/\text{sec}$), that pixel would not be flagged as hot, making this pixel appear inconsistently hot.

Adding the orange points to the yellow points, we see the population of pixels where a variation of the dark rate in an input image of up to 3 times their measured noise levels would result in inconsistently hot pixels. And adding both the yellow and orange points to the red, we see pixels where a 5σ variation in signal would cause them to be seen as inconsistently hot.

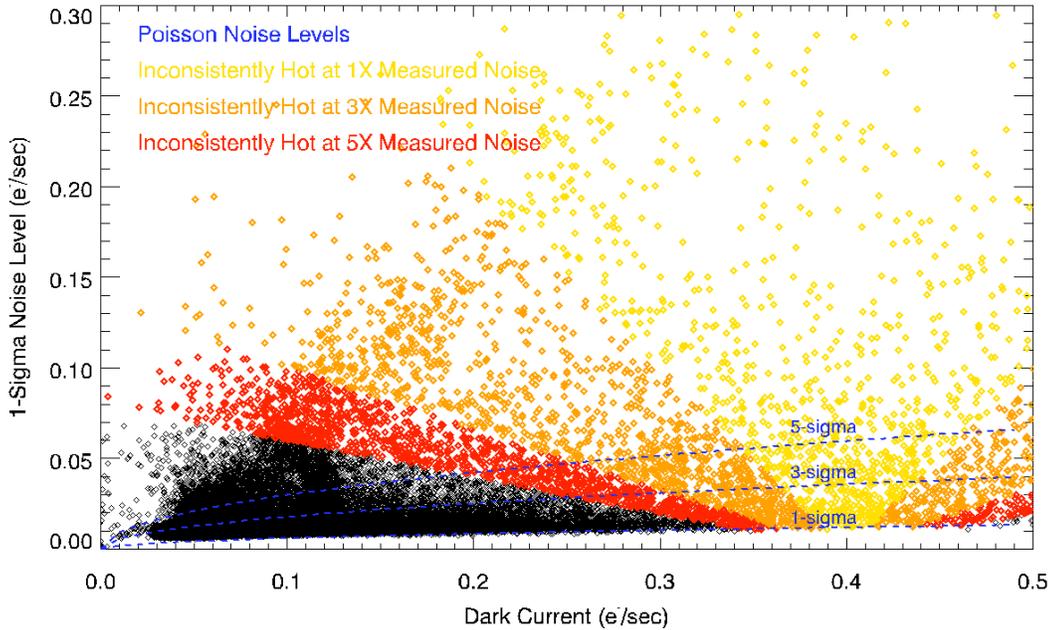


Figure 26: Plot of noise versus signal rate for the SPARS200 mean dark current image. A histogram of these noise values is shown in the red curve in Figure 2. Dashed blue lines show multiples of theoretical (Poisson) noise corresponding to the dark current on the x axis. Red, orange, and yellow points show pixels that will be seen as inconsistently hot across the input flt files when their signal varies by 1, 3, or 5 times their measured noise level, respectively. Note that the yellow points are a subset of the orange points, and both are a subset of the red points. Out of the total of 1,028,196 pixels, 1252 pixels are yellow, 2274 are orange, and 2280 are red.

Totaling up the number of yellow, orange, and red points, we find that in the SPARS200 data, for signals shifted 1σ , 3σ , and 5σ from their means, that 1252, 3526, and 5806 pixels become inconsistently hot, respectively.

Next we repeat the calculation, but using Poisson noise rather than the measured noise, in order to calculate how many inconsistently hot pixels we would expect to see in an ideal detector with the same mean dark current values. This calculation shows that signal shifts of 1, 3, and 5 times the Poisson noise would result in 225, 686, and 1135 inconsistently hot pixels. Comparing these numbers to those calculated using the measured noise values, it appears that we see roughly 5 times as many inconsistently hot pixels compared to what we would expect from Poisson noise alone. Also, the 3546 inconsistently hot pixels found experimentally in the SPARS200 data (and shown in Figure 25) is very close to the number of inconsistently hot pixels predicted from the mean dark current image where the dark current signals are shifted 3σ from their means (3526 orange points of Figure 26). Given that our input dataset for these calculations consisted of only 69 SPARS200 files, and that the probability of a given pixel being 3σ from its mean at some point in the dataset is only $0.3\% * 69 = 20\%$, it is unlikely that

Poisson noise alone is responsible for driving all of the 3546 observed inconsistently hot pixels beyond 3σ . This result is not surprising. Figure 27 shows examples of dark current rates over time for inconsistently hot pixels. These plots show that there are obviously other effects at work besides Poisson noise, with large, sudden jumps in the dark rate and in the case of the right panel, a slowly increasing dark current rate even between regions of highly elevated values. Despite the relatively low signal rates, readnoise also cannot account for these variations. With a correlated double sampling (CDS) readnoise value of $21e^-$, we expect the readnoise contribution in an *flt* image, to be $21e^-/\sqrt{16 \text{ reads}} = 5e^-$. This corresponds to $0.002 e^-/\text{sec}$ in a SPARS200 ramp, and $0.004 e^-/\text{sec}$ in a SPARS100 ramp, which are small compared to the variations in Figure 27.

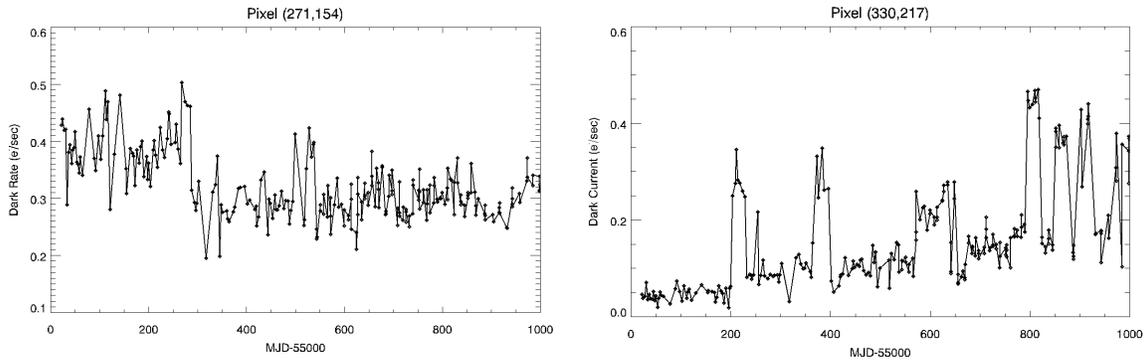


Figure 27: Plots of measured dark current versus time for two individual pixels, both of which are flagged as hot in some images, but not others. Note that Poisson noise is not the dominant noise source for either pixel. Large jumps in the measured dark rate indicate other, unknown effects are larger than the variations due to Poisson noise.

Plots like those in Figure 27 for all pixels across the detector show similar sudden jumps in dark current rates between consecutive ramps, although typically with much smaller magnitudes than those seen below. In some cases, these jumps are due to persistence but in others, the dark current itself appears to change.

Using plots like these for all pixels, we were able to note the maximum dark current difference between consecutive ramps for all pixels. (For example, for the pixel shown in the right panel, we recorded $0.35 e^-/\text{sec}$ as the largest change between any two consecutive ramps. Our goal in creating a map of these values was to look for any coherent, large spatial scale areas on the detector that were prone to large, sudden changes in dark current. The map (not shown) contained a small amount of contamination from persistence, but otherwise appeared relatively flat, with salt and pepper noise across the entire detector, and a peak value (in a histogram) of $0.057 e^-/\text{sec}$. This implies that the average pixel on the IR detector can potentially more than double its

nominal dark rate from one ramp to the next. Keep in mind that this value is from the map of the maximum dark current differences and therefore represents a worst-case scenario. A similar map of the mean (rather than maximum) change in dark current between consecutive ramps also appears flat across the detector, and has a peak value of zero. As shown in Figure 3, the 1σ variation in the dark current of the average pixel can be as high as $0.015\text{ e}^-/\text{sec}$, which is roughly 33% of the nominal mean dark current rate. With these maps containing only pixel-to-pixel noise and no large scale structures, this appears to be a fundamental behavior of the detector and **implies a limit to the accuracy of the pipeline dark current subtraction of $0.015\text{ e}^-/\text{sec}$ and in extreme cases, as much as $0.057\text{ e}^-/\text{sec}$.**

Variable Pixels

Our final analysis task, inspired by the inconsistent pixels seen in Figure 25, was to monitor the population of pixels in each image which were far from their mean dark current level. For each pixel, we fit a line to the measurements of that pixel's dark rates across the 252 input images. We then subtracted this best-fit line from the measured dark rates, in order to remove any long-term trend in the data and set the expected value to 0 in all images. Using the sigma-clipped standard deviation map (the histogram of which is shown in Figure 2), we recorded the images in which the pixel's dark rate was more than 3 times its standard deviation. Results are shown in Figures 28 and 29.

The blue points in Figure 28 show the number of pixels with dark rates more than 3σ from their expected values for the SPARS100 data. The red points show the results for the SPARS200 images.

Under the assumption that the measured dark current rate from image to image in a given pixel follows poisson statistics, we would expect to find 3,084 pixels in any given image more than 3σ from their mean value ($= 1,028,196\text{ total pixels} * 0.3\%$). Note that this is a lower limit, as it does not account for the effects of readnoise which, for a pixel with the mean dark rate of $0.046\text{ e}^-/\text{sec}$ in a SPARS100 ramp, can cause variations of $\sim 10\%$ in the measured signal rate. Nevertheless, this predicted population is consistent with the number of pixels flagged in most of the SPARS200 images in Figure 28. For the SPARS100 ramps, the measured population of highly variable pixels is well above the predicted lower limit.

Examination of a separate map of these pixels for each of the input images revealed that contamination from residual persistence caused at least a small number of pixels to be erroneously identified as having dark current more than 3σ from their means.

Figure 29 shows the correlation of dark rate with the number of pixels more than 3σ from their mean dark rate. The blue points, derived from the SPARS100 data, show a

well-defined relationship, with a minimum in the number of highly variable points at a dark rate of roughly 0.043 e-/sec. The relationship in the red points, derived from the SPARS200 files, is less clear, with a general trend of fewer points more than 3σ from their mean as the overall mean dark current rate decreases. As with previous figures, the points circled in green indicate dark current observations made shortly after internal flat fields and potentially subject to wide-spread persistence. With all but one of these points falling in line with the rest of the data, it appears that the number of pixels driven more than 3σ from their mean as a result of persistence is minimal. Repeating the calculations with a threshold of 5σ reveal plots that appear very similar to Figures 28 and 29, but with the number of pixels about an order of magnitude less than for the 3σ case.

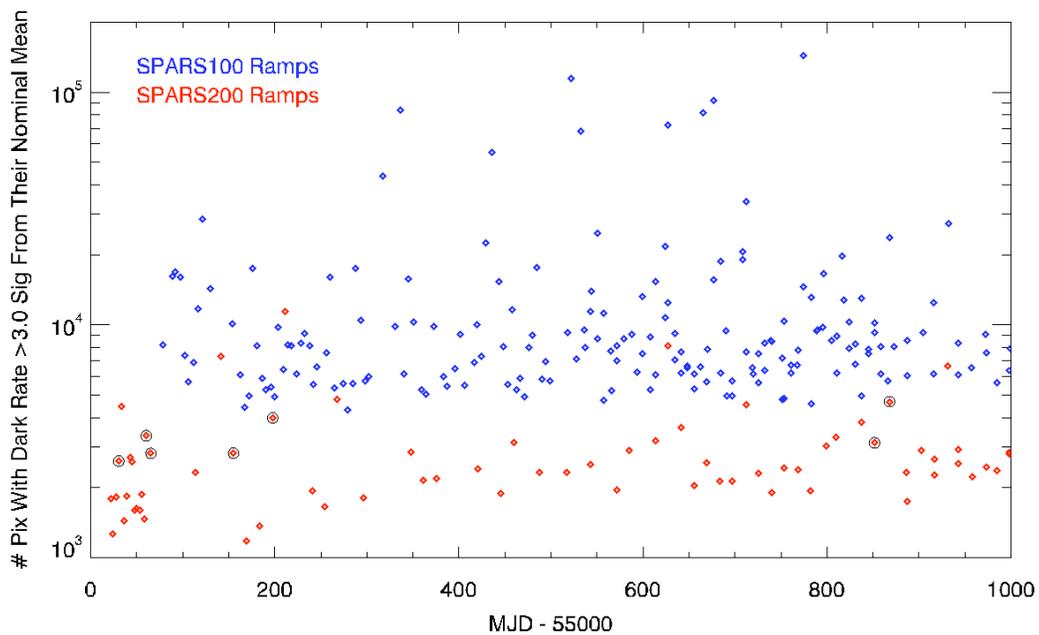


Figure 28: Points show the number of pixels in each image displaying a dark rate more than 3σ from their mean dark current value. Blue points are for SPARS100 data, while red points are for SPARS200 data. Points circled in green indicate observations made shortly after internal flat field observations.

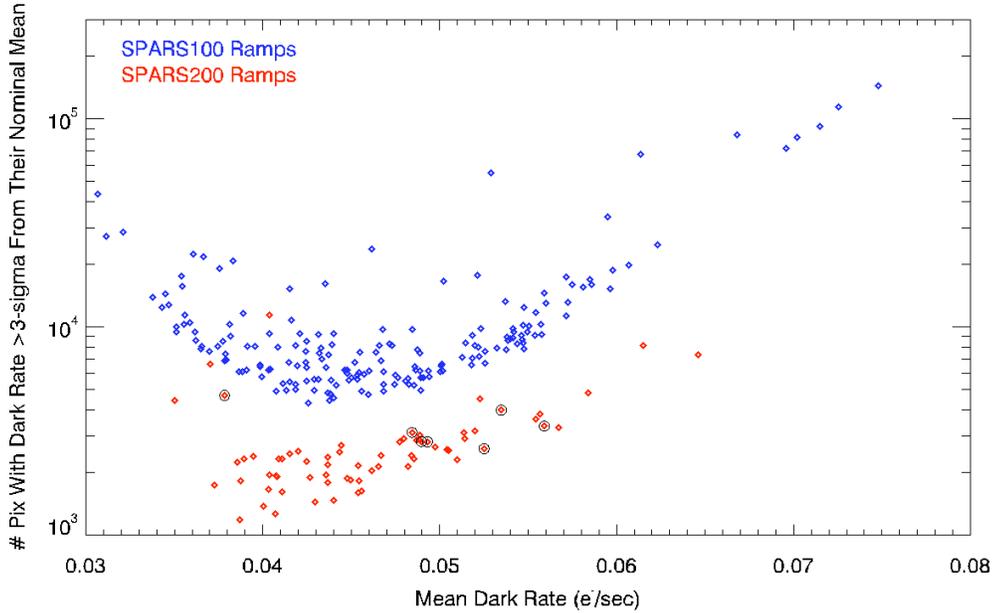


Figure 29: Plot showing the correlation between mean dark rate and number of pixels with abnormal dark current.

Conclusions

Our analysis of long exposure time WFC3/IR dark current data collected over a 977 day period indicates that there are some changes in the detector’s behavior over time, mostly relating to populations of hot and unstable pixels.

The mean dark rate does not appear to have changed during WFC3’s time on orbit. Ramp to ramp variation in the detector-averaged mean dark rate is 0.0075 e⁻/s (~16% of the overall mean dark rate). On a pixel by pixel basis, the majority of pixels show a 0.0094 e⁻/sec variation in dark rate in the SPARS200 dataset, which is 20% of the overall mean dark rate. This value increases to roughly 33% of the mean dark rate for the SPARS100 ramps. These variations, which do not display any significant spatial structure across the detector, imply a fundamental limit on the accuracy of the dark current subtraction performed by the *calwf3* pipeline.

Attempts to find correlation between the mean dark current level and various instrument temperatures, voltages, and HST pointing direction revealed no obvious connections.

Examination of the relative rates of dark current increases across the detector reveal spatial patterns, with the dark current increasing more quickly along the right and bottom edge, as well as a horizontal bar about one-third of the way down from the top edge, compared to other areas on the detector.

We find that the number of hot pixels varies with time and across the detector, and that this number is anti-correlated with the bias level measured in the reference pixels.

Abrupt changes in the bias level match up with abrupt changes in the number of hot pixels. However, the mean dark current rate does not appear to be correlated with bias level.

Counting the number of pixels displaying dark current rates more than 3σ from their mean value reveals that the number of flagged pixels is minimized for images with a mean dark current of roughly $0.043 \text{ e}^-/\text{sec}$.

Recommendations

Periodic searches for new unstable and hot pixels are recommended, given the observed pixel populations in Figures 22 and 28. Further analysis into the nature of the signal variations in the dark current images is also recommended.

References

Barker, E. A., McCullough, P., and Martel, A. R., 2010. WFC3 ISR 2009-40. “*WFC3 IR SAA Passage Behavior*”.

Bushouse, H. , 2009. WFC3 ISR 2009-02. “*WFC3 TV3 Testing: Orbital Cycling Effects on IR Images*”.

Dulude, M., Bagget, S., Bushouse, H., and Hilbert, B., 2011. WFC3 ISR 2011-04. “*WFC3/IR Banding*”.

Hilbert, B. , 2012. WFC3 ISR 2012-05. “*WFC3/IR Reference Pixel Characterization #1: Comparison of Bias Subtraction Methods*”.

Hilbert, B. , 2009. WFC3 ISR 2009-42. “*WFC3 SMOV Program 11433: IR Internal Flat Field Observations*”.

Hilbert, B. , 2008. WFC3 ISR 2008-30. “*WFC3 TV3 Testing: IR Channel Dark Current*”.

Hilbert, B. , 2008b. WFC3 ISR 2008-50. “*WFC3 TV3 Testing: IR Gain Results*”.

Hilbert, B., and Bushouse, H., 2010. WFC3 ISR 2010-13. “*WFC3/IR Bad Pixel Table: Update Using Cycle 17 Data*”.

Hilbert, B. , and McCullough, P., 2009. WFC3 ISR 2009-21. “*WFC3 SMOV Results: IR Channel Dark Current, Readnoise, and Background Signal*”.

Long, K. , Wheeler, T. and Bushouse, H., 2011. WFC3 ISR 2011-09. “*IR Detector Timing and Persistence*”.

Long, K., Baggett, S., Deustua, S., and Riess, A., 2010. WFC3 ISR 2010-17. “*WFC3/IR Persistence as Measured in Cycle 17 using Tungsten Lamp Exposures*”.

Rajan, A. et al. 2010, “WFC3 Data Handbook”, Version 2.0, (Baltimore:STScI).

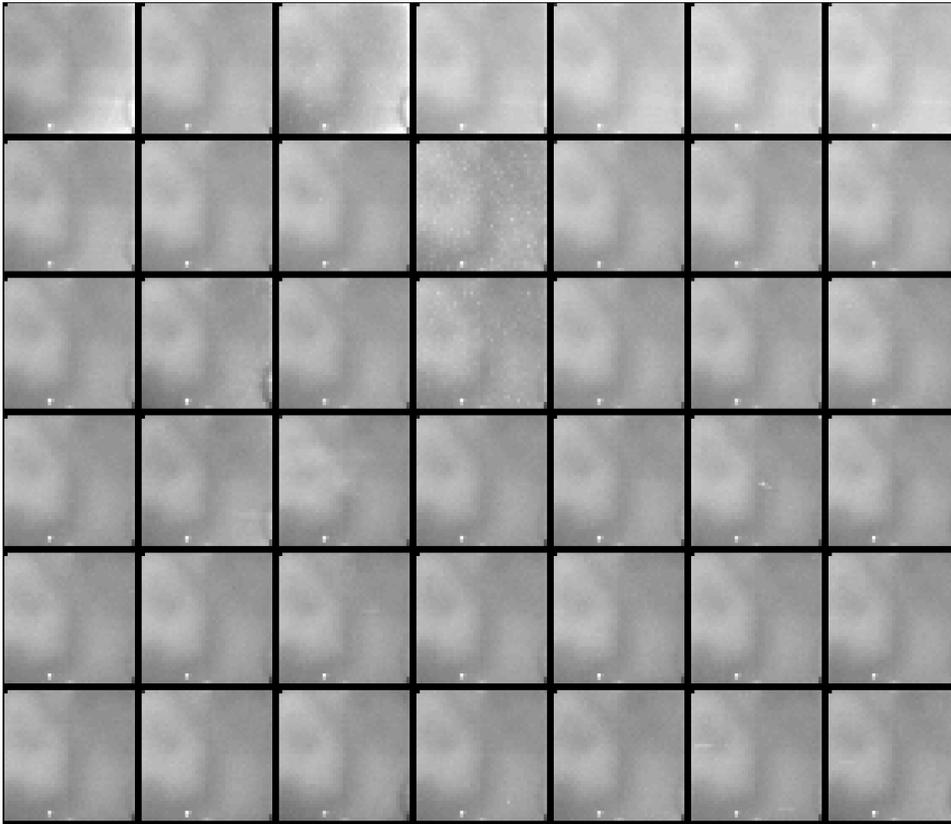
Acknowledgements

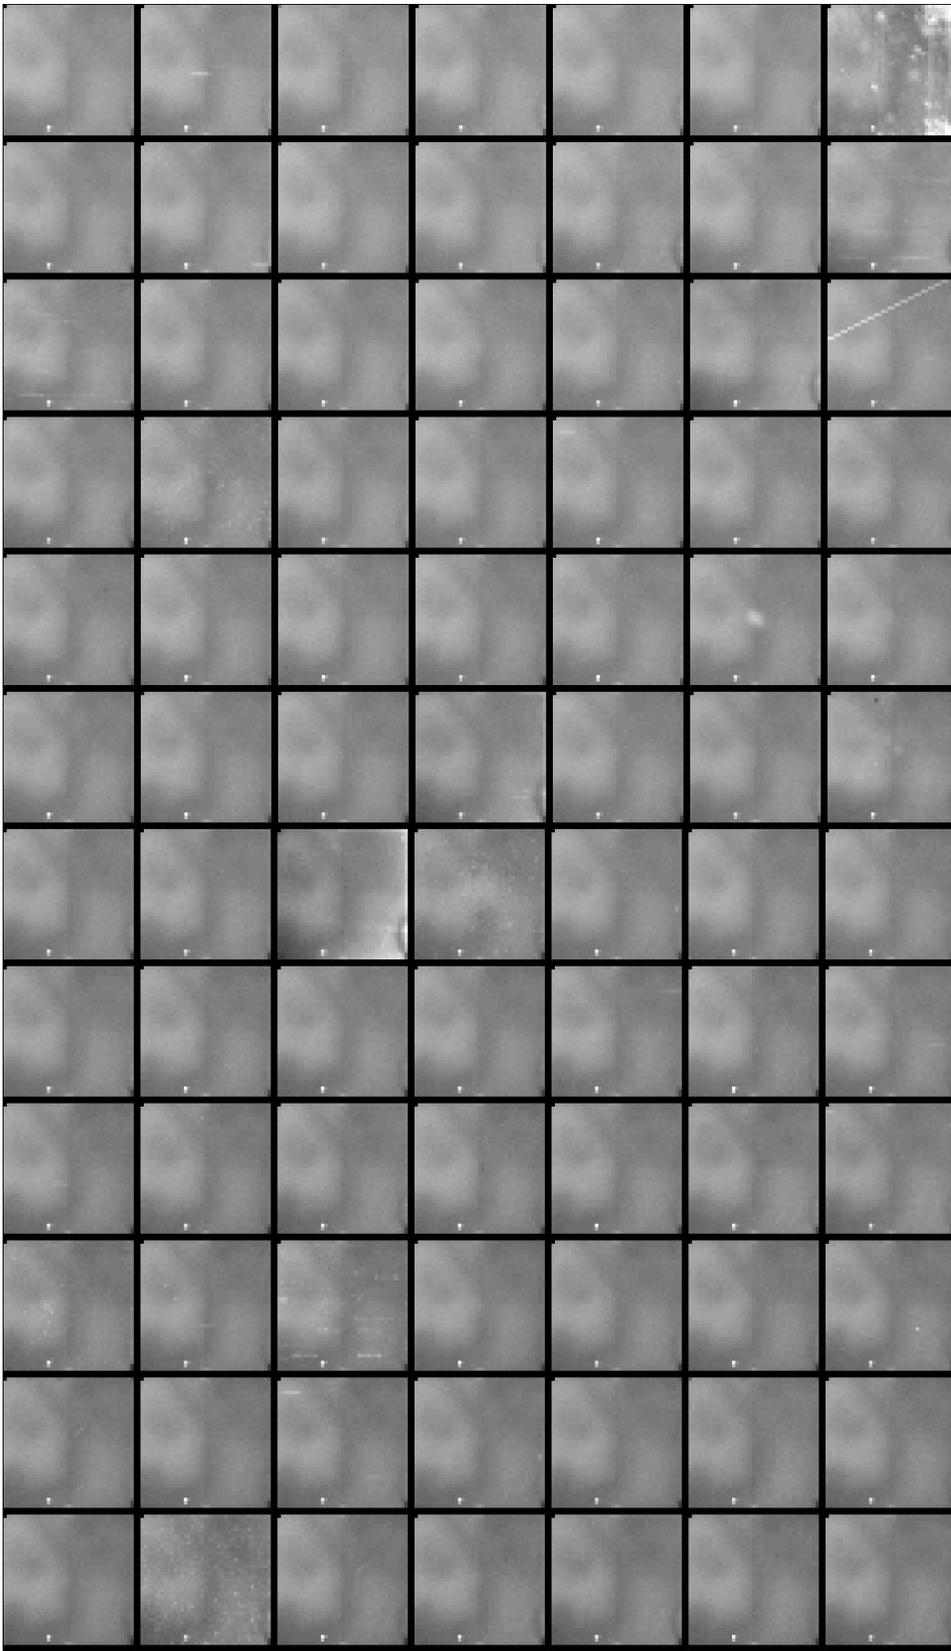
The authors would like to thank Sylvia Baggett for many discussions and analysis ideas involving the dark current data, and Howard Bushouse for many useful comments.

Appendix

Appendix 1

Figure 30: SPARS100 dark current gallery. Linear stretch from 0 to 0.1 e⁻/sec.





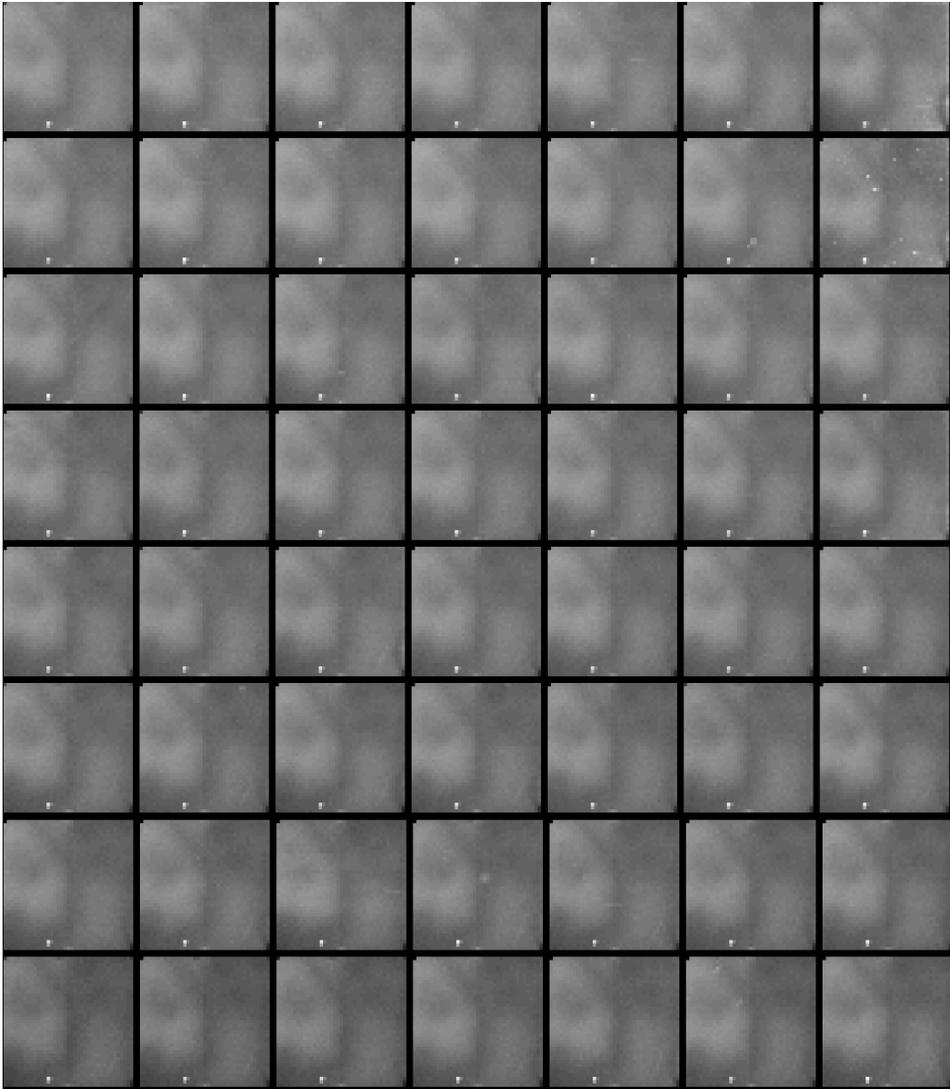
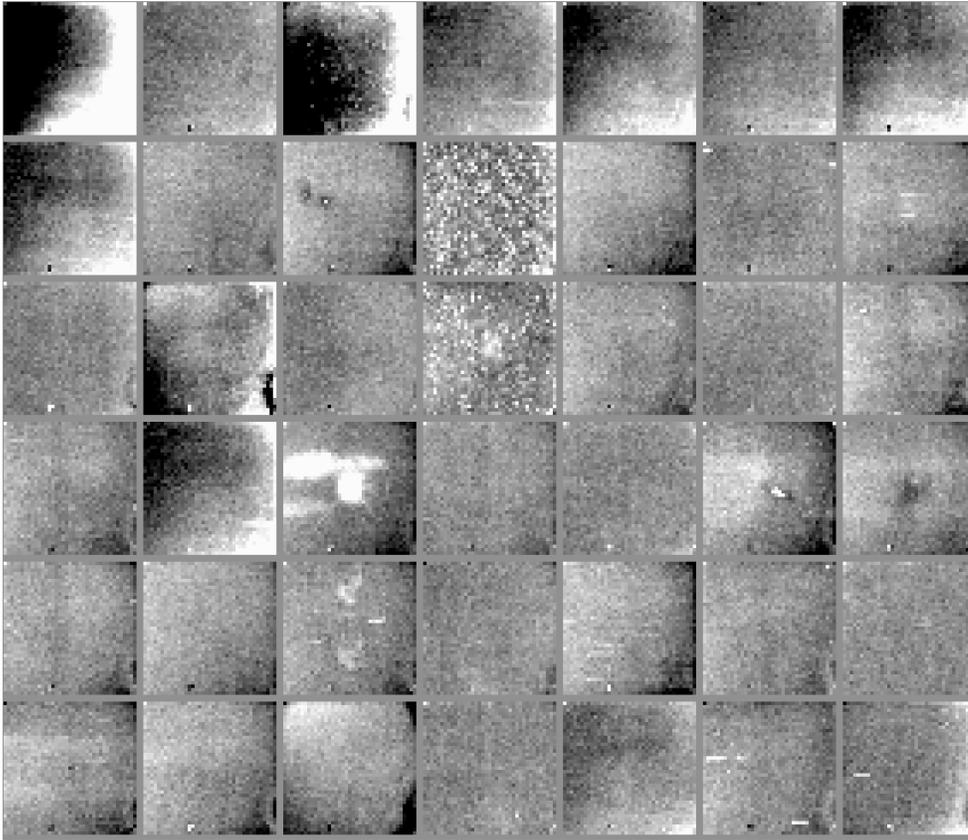
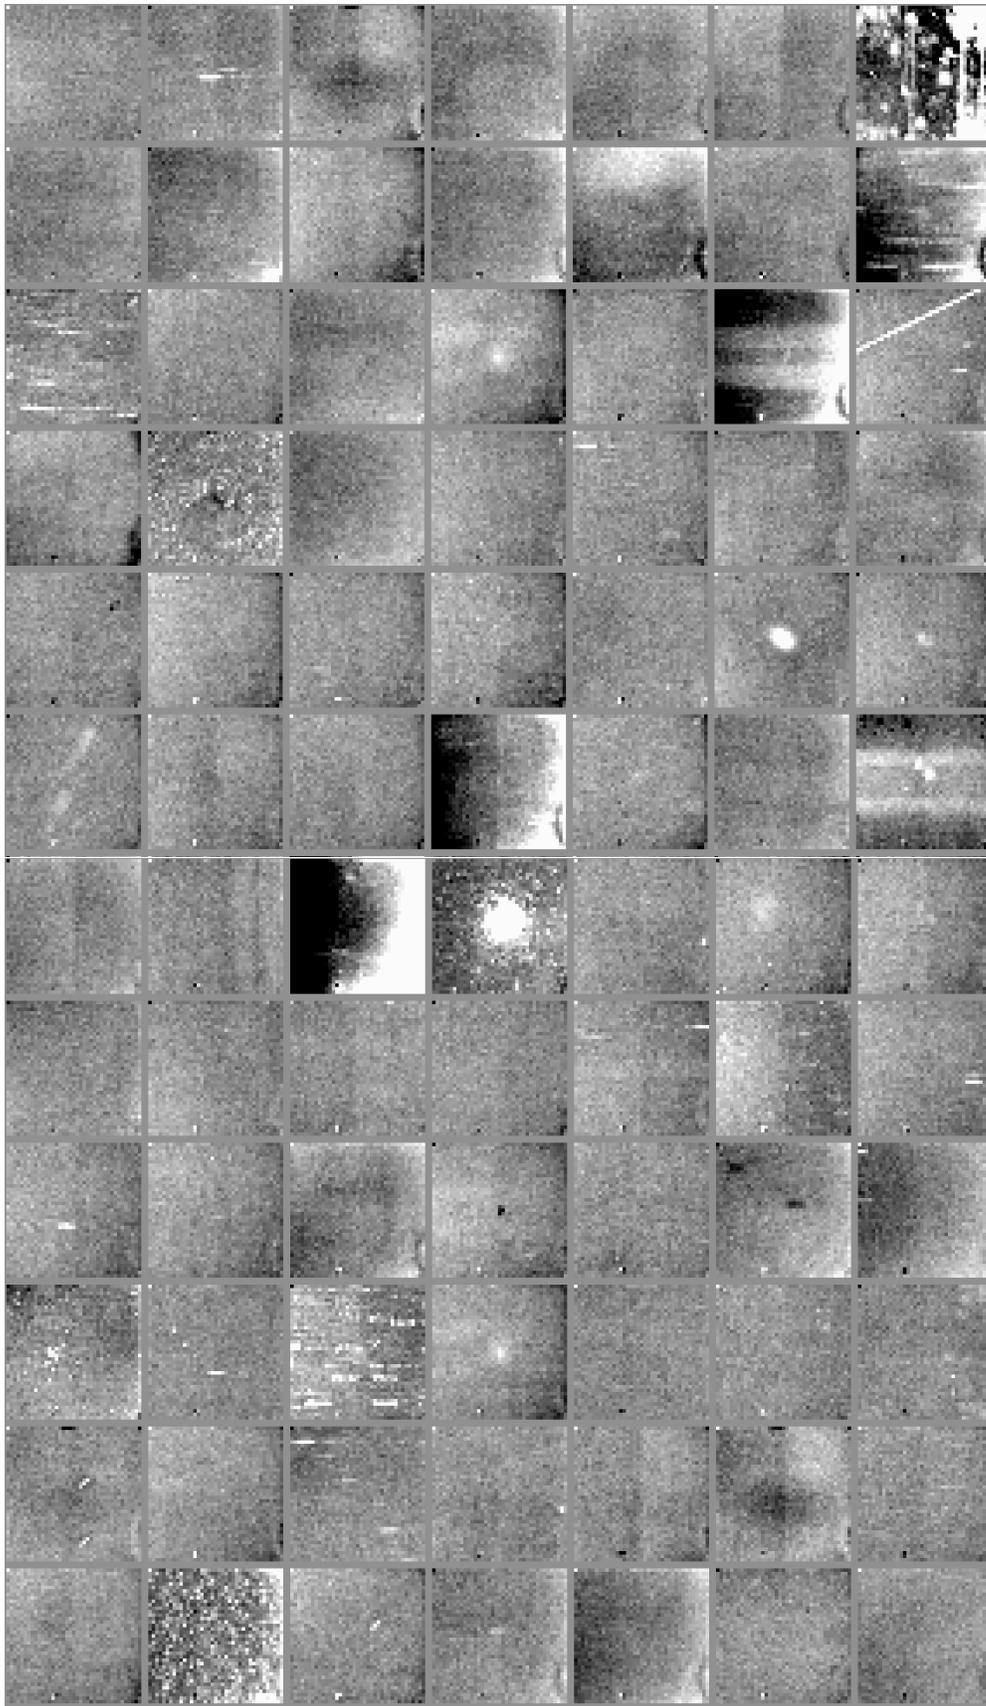


Figure 31: SPARS100 dark current gallery after subtraction of the scaled line-fit signal contribution (Figure 12). Linear stretch from -0.005 to $+0.005$ e^-/sec , which is roughly $\pm 10\%$ of the mean dark rate.





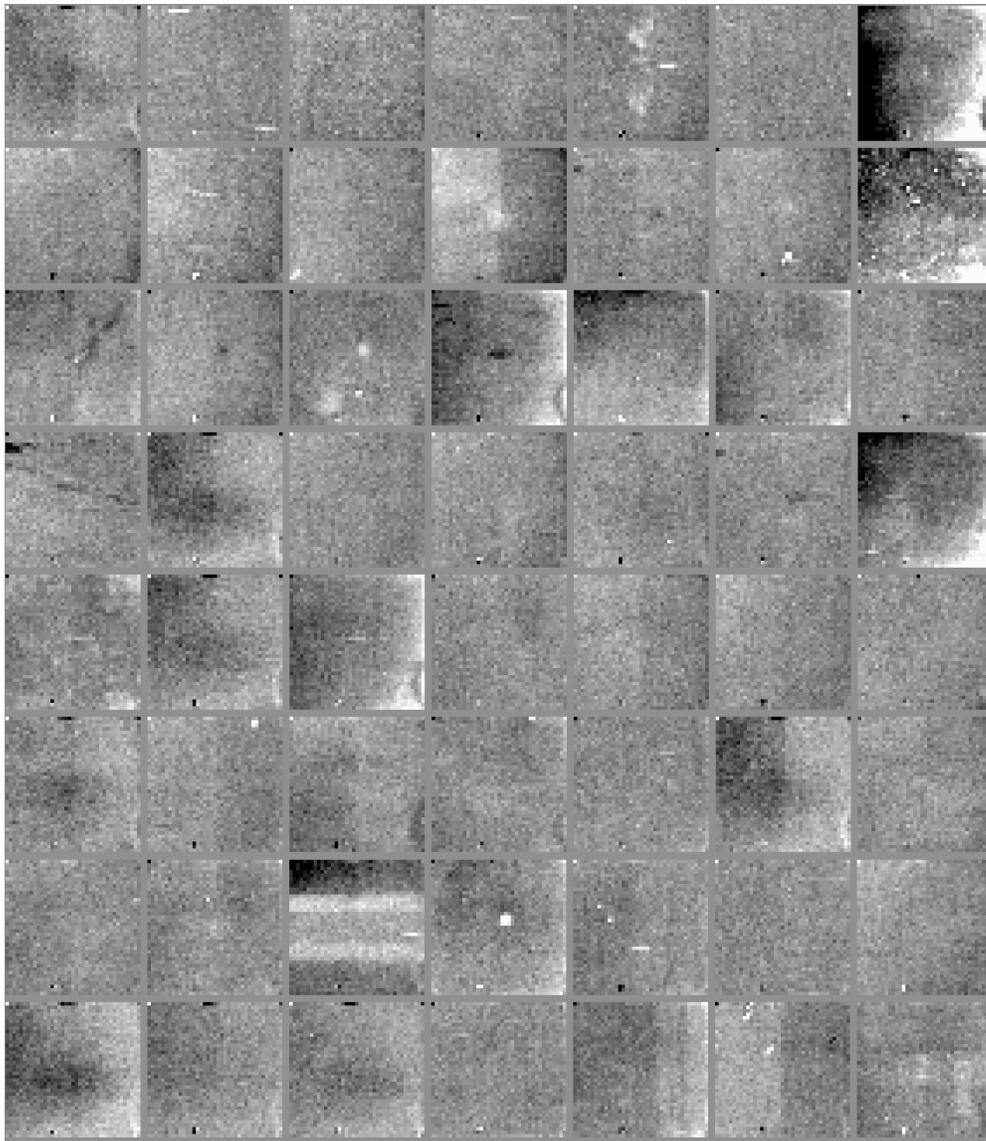
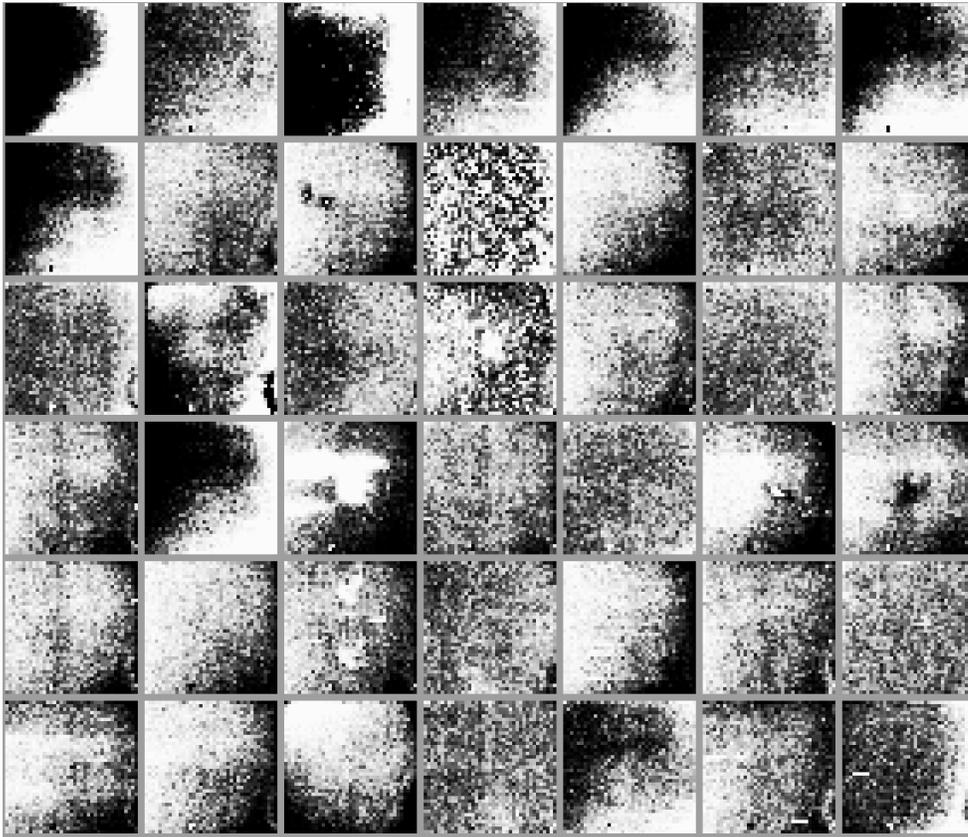
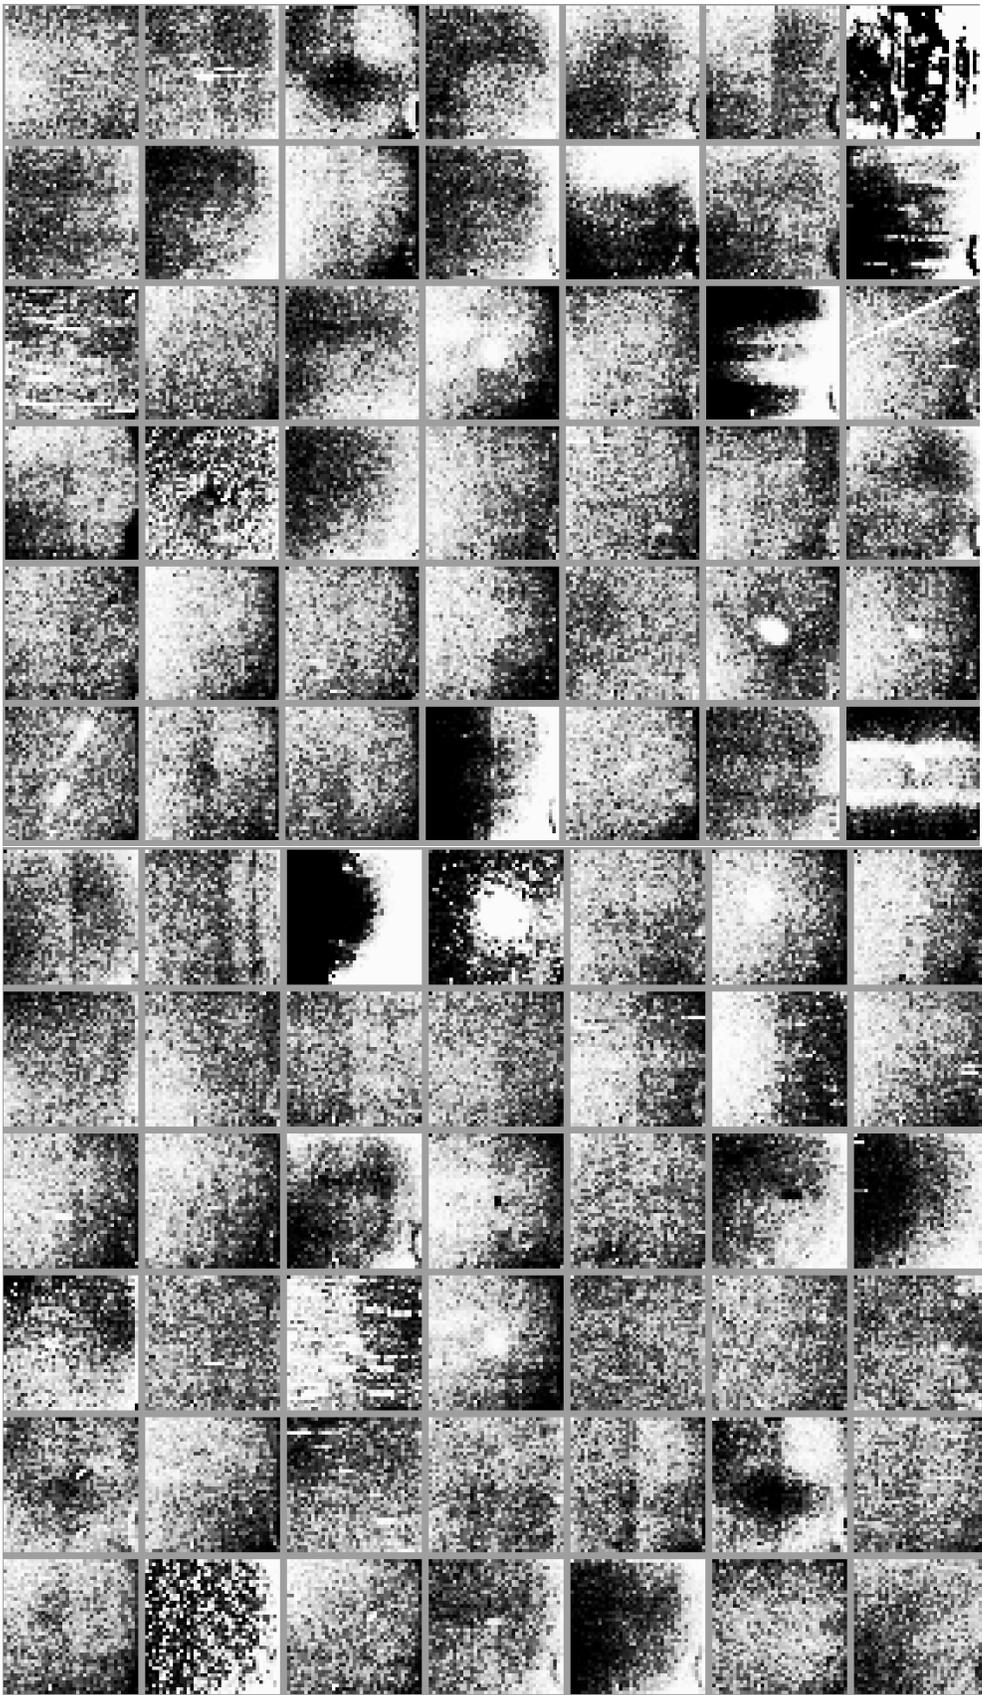


Figure 32: Same as above, but with a histogram equalization stretch, to accentuate residual features in the gallery.





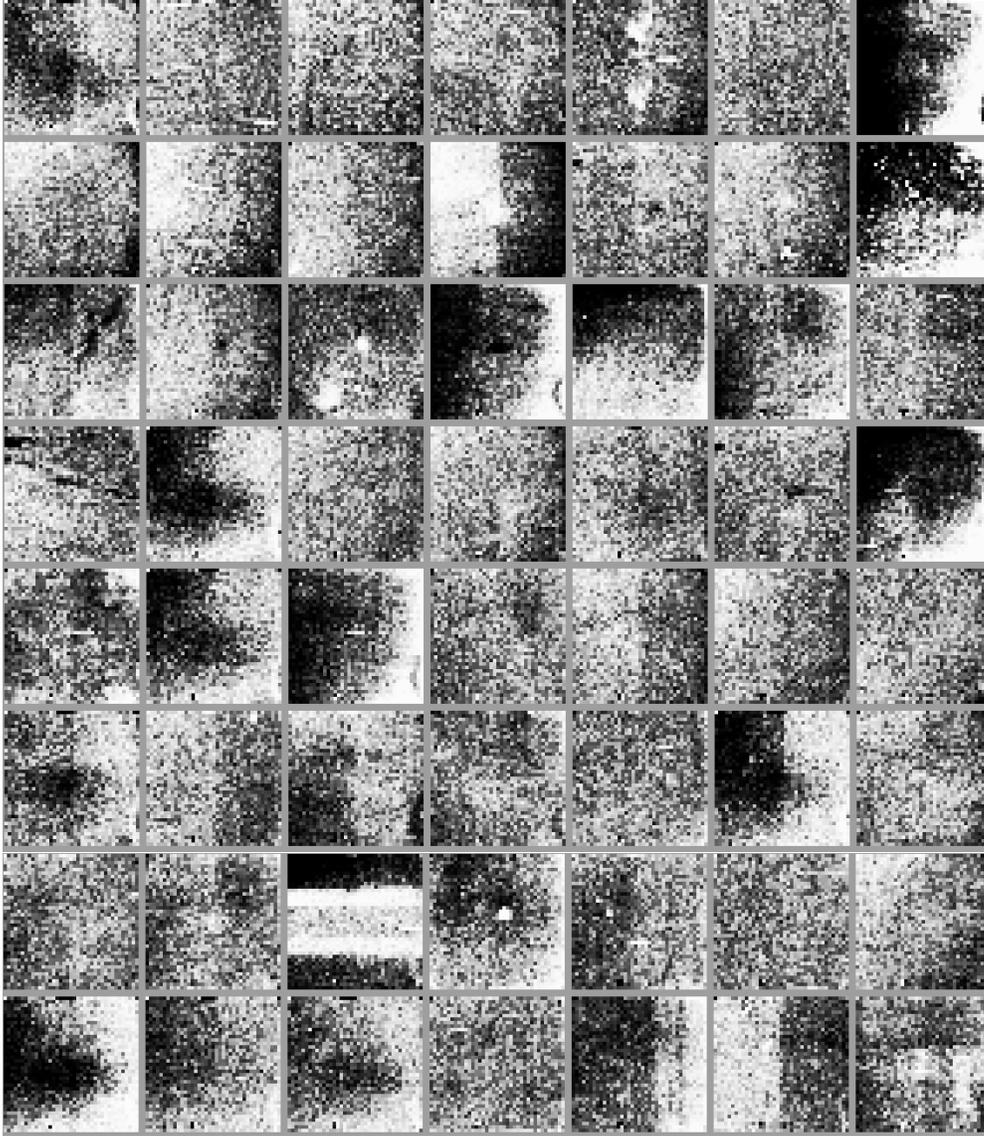


Table 1: List of dark current ramps used in this analysis.

Sample Sequence	Exposure Time (sec)	Proposal Number	Files
SPARS100	1402.9	11929	ibcu0oamq_flt.fits, ibcu10jqj_flt.fits, ibcu15ejq_flt.fits, ibcu1amjq_flt.fits, ibcu1berq_flt.fits, ibcu1gx8q_flt.fits, ibcu1noj_q_flt.fits, ibcu1odoq_flt.fits, ibcu1tctq_flt.fits, ibcu22sfq_flt.fits, ibcu23h5q_flt.fits, ibcu28ikq_flt.fits, ibcu2aq3q_flt.fits, ibcu2bdjq_flt.fits, ibcu2gyrq_flt.fits, ibcu2oa1q_flt.fits, ibcu2th7q_flt.fits, ibcu35z5q_flt.fits, ibcu36dqq_flt.fits, ibcu3alvq_flt.fits, ibcu3gxlq_flt.fits, ibcu3nm3q_flt.fits, ibcu3oe7q_flt.fits, ibcu3tvxq_flt.fits, ibcu41gpq_flt.fits, ibcu49c5q_flt.fits, ibcu4al2q_flt.fits, ibcu4bchq_flt.fits, ibcu4gzfq_flt.fits, ibcu4nfyq_flt.fits, ibcu5abiq_flt.fits, ibcu5bzuq_flt.fits, ibcu5nffq_flt.fits, ibcu5oqqq_flt.fits, ibcu6adeq_flt.fits, ibcu6bqkq_flt.fits, ibcu6nd7q_flt.fits, ibcu6tkbq_flt.fits, ibcu74s1q_flt.fits, ibcu7ad0q_flt.fits, ibcu7gagq_flt.fits, ibcu7neuq_flt.fits,

			<p>ibcu7owmq_fit.fits, ibcu80xxq_fit.fits, ibcu871eq_fit.fits, ibcu88cmq_fit.fits, ibcu8akxq_fit.fits, ibcu8bw3q_fit.fits, ibcu8ge3q_fit.fits, ibcu8ne7q_fit.fits, ibcu8otrq_fit.fits, ibcu93wtq_fit.fits, ibcu9aknq_fit.fits, ibcu9bvkq_fit.fits, ibcu9ghaq_fit.fits, ibcu9nktq_fit.fits, ibcu5gdxq_fit.fits, ibcubhaaq_fit.fits, ibcubgfrq_fit.fits, ibcub9axq_fit.fits, ibcub4axq_fit.fits, ibcuagfq_fit.fits, ibcuajdfq_fit.fits, ibcuayyq_fit.fits, ibcuadcyq_fit.fits, ibcu6a6q_fit.fits, ibcu11keq_fit.fits, ibcu0pxq_fit.fits, ibcu9terq_fit.fits, ibcu9oseq_fit.fits, ibcu02nfq_fit.fits, ibcu0aodq_fit.fits, ibcu0bc3q_fit.fits, ibcu0gxxq_fit.fits, ibcu0nrfq_fit.fits, ibcu0ttyq_fit.fits</p>
SPARS100	1402.9	12349	<p>iblt44ihq_fit.fits, iblt4ah4q_fit.fits, iblt4hl2q_fit.fits, iblt4ij4q_fit.fits, iblt4pfjq_fit.fits, iblt4qdiq_fit.fits, iblt4xf5q_fit.fits, iblt4yf8q_fit.fits, iblt51j9q_fit.fits, iblt59q6q_fit.fits, iblt5fenq_fit.fits, iblt5gapq_fit.fits, iblt5nldq_fit.fits, iblt5olnq_fit.fits, iblt5vdgq_fit.fits, iblt5wb6q_fit.fits, iblt68kiq_fit.fits, iblt6dnfq_fit.fits, iblt6enhq_fit.fits, iblt6lj1q_fit.fits, iblt6me3q_fit.fits, iblt6tqlq_fit.fits, iblt6ugjq_fit.fits, iblt75dhq_fit.fits, iblt76d9q_fit.fits, iblt7bjtq_fit.fits, iblt7cjoq_fit.fits, iblt7jj6q_fit.fits, iblt7kj2q_fit.fits, iblt7rdnq_fit.fits, iblt7skjq_fit.fits, iblt7zwoq_fit.fits, iblt83k6q_fit.fits, iblt84k1q_fit.fits, iblt8hk0q_fit.fits, iblt8iizq_fit.fits, iblt8pb4q_fit.fits, iblt8qndq_fit.fits, iblt8xhwq_fit.fits, iblt8yhhq_fit.fits, iblt91cjg_fit.fits, iblt92c8q_fit.fits, iblt99lpq_fit.fits, iblt9fe0q_fit.fits, iblt9gewq_fit.fits, iblt9ncdq_fit.fits, iblt9obyq_fit.fits, iblt9vj5q_fit.fits, iblt9wiyq_fit.fits, iblt93j3q_fit.fits, iblt94iyq_fit.fits, iblt9bf5q_fit.fits, iblt9achsq_fit.fits, iblt9ajahq_fit.fits, iblt9akmq_fit.fits, iblt9bmdlq_fit.fits, iblt9ldeq_fit.fits, iblt9mcqq_fit.fits, iblt9tjfq_fit.fits, iblt9uwiwq_fit.fits, iblt95mbq_fit.fits, iblt96jbq_fit.fits, iblt93bf9q_fit.fits, iblt9cea_fit.fits, iblt9jcaq_fit.fits, iblt93kclq_fit.fits, iblt93rymq_fit.fits, iblt93sywq_fit.fits, iblt93zhaq_fit.fits, iblt943jfq_fit.fits, iblt903feq_fit.fits, iblt904r1q_fit.fits, iblt90akrq_fit.fits, iblt90hdtq_fit.fits, iblt90idmq_fit.fits, iblt90pd1q_fit.fits, iblt90qcrq_fit.fits, iblt90xdqq_fit.fits, iblt90ydoq_fit.fits, iblt911drq_fit.fits, iblt912daq_fit.fits, iblt919clq_fit.fits, iblt91fdeq_fit.fits, iblt91gczq_fit.fits, iblt91ngqq_fit.fits, iblt91og5q_fit.fits, iblt91vo3q_fit.fits, iblt91wsgq_fit.fits, iblt920cqq_fit.fits, iblt927crq_fit.fits, iblt928f8q_fit.fits, iblt92dg8q_fit.fits, iblt92ehlq_fit.fits</p>
SPARS100	1402.9	12695	<p>ibvi08exq_fit.fits, ibvi84f4q_fit.fits, ibvi75iqq_fit.fits, ibvi66m2q_fit.fits, ibvi65raq_fit.fits, ibvi56ftq_fit.fits, ibvi47dkq_fit.fits, ibvi46f8q_fit.fits, ibvi37z5q_fit.fits, ibvi28m7q_fit.fits, ibvi27m9q_fit.fits, ibvi18iiq_fit.fits, ibvi09elq_fit.fits, ibvi85e0q_fit.fits</p>
SPARS200	2802.9	11447	<p>iabz16h4q_fit.fits, iabz06geq_fit.fits, iabz09edq_fit.fits, iabz10ghq_fit.fits, iabz18suq_fit.fits, iabz07axq_fit.fits, iabz17q3q_fit.fits, iabz11acq_fit.fits, iabz19e2q_fit.fits, iabz12d4q_fit.fits, iabz05xqq_fit.fits, iabz13lfq_fit.fits, iabz14ljq_fit.fits, iabz08ojq_fit.fits, iabz15h7q_fit.fits,</p>

			iabz20u1q_fit.fits
SPARS200	2802.9	11929	ibcu7yxnq_fit.fits, ibcu85v8q_fit.fits, ibcu1yi2q_fit.fits, ibcu72buq_fit.fits, ibcu8yd7q_fit.fits, ibcu98b1q_fit.fits, ibcu2ldjq_fit.fits, ibcub1tyq_fit.fits, ibcu59fkq_fit.fits, ibcu2ywtq_fit.fits, ibcu33onq_fit.fits, ibcu9ldzq_fit.fits, ibcuabmaq_fit.fits, ibcu3yvvq_fit.fits, ibcu6lq3q_fit.fits, ibcu0ldxq_fit.fits, ibcu5ydaq_fit.fits, ibcu0yvqq_fit.fits, ibcu5ldbq_fit.fits
SPARS200	2802.9	12349	iblt5jf5q_fit.fits, iblt5zasq_fit.fits, iblt63chq_fit.fits, iblt0lcxq_fit.fits, iblt3nbyq_fit.fits, iblt8ldzq_fit.fits, iblta7eas_fit.fits, iblt6plcq_fit.fits, iblt15c6q_fit.fits, iblt47hbq_fit.fits, iblt1bdcq_fit.fits, iblt4dgwq_fit.fits, iblt1rfoq_fit.fits, ibltanamq_fit.fits, iblt95btq_fit.fits, iblt4tetq_fit.fits, iblt7fg2q_fit.fits, iblt2hfxq_fit.fits, iblt7vc7q_fit.fits, iblt2xfqq_fit.fits, iblt9rbkq_fit.fits
SPARS200	2802.9	12695	ibvi92deq_fit.fits, ibvi16d1q_fit.fits, ibvi44x7q_fit.fits, ibvi82faq_fit.fits, ibvi35trq_fit.fits, ibvi34tmq_fit.fits, ibvi63npq_fit.fits, ibvi25vdq_fit.fits, ibvi54ciq_fit.fits, ibvi73s2q_fit.fits, ibvi91caq_fit.fits, ibvi53cfq_fit.fits, ibvi15feq_fit.fits

Table 2: List of instrument temperatures and voltages to which the mean dark rate was compared in a search for trends.

Header Keyword Name	Location	Description
IRFPATMP	spt file, extension 2	FPA Temperature
IREVAPT	spt file, extension 2	Baseplate Evaporator Saddle Temperature
IRENEVPT	spt file, extension 2	Cold Enclosure Evaporator Temperature
IOBIRM1T	spt file, extension 2	Temperature: Optical Bench Near M1 Mount
IDRIRCST	spt file, extension 2	Temperature: Detector Radiator on IR side by condensor saddle
IDRIRT	spt file, extension 2	Temperature: Detector Radiator on IR side away from saddle
IRBDACHI	spt file, extension 2	Digital Analog Converter output for vbias signals
IRDP15	spt file, extension 2	Switched +15V for detector ANALOG to buffer PWA
IRDACREF	spt file, extension 2	Digital to analog converter +5V reference
IR1SP15	spt file, extension 2	Input +15V from active SPB, quads 1,2
IR1ANBP5	spt file, extension 2	Analog +5V applied to detector, quads 1,2
IR1FB1	spt file, extension 2	Fixed bias to detector substrate,

		quads 1,2
IR1FB2	spt file, extension 2	Fixed bias to detector bias generator, quads 1,2
IR12OFST	spt file, extension 2	Offset removal for signal processing chain quads 1,2
IR1GND	spt file, extension 2	Analog ground reference for buffer PWA quads 1,2
IR1VB	spt file, extension 2	Programmable bias to photodiodes thru MUX quads 1,2
IFSMTMP	spt file, extension 2	FSM motor temperature
IFSMMECT	spt file, extension 2	FSM housing temperature
IR6TECV	spt file, extension 4	IR FPA 6-stage TEC voltage
IRCETECV	spt file, extension 4	IR cold enclosure TEC voltage
IRDBTECV	spt file, extension 4	IR detector baseplate TEC voltage
I28VIN2	Not in observation files	Voltage input to WFC3 from the spacecraft bus