Instrument Science Report STIS 2018-05

Performance of the STIS CCD Dark Rate Temperature Correction

Doug Branton[1]

[1] Space Telescope Science Institute, Baltimore, MD

October 10, 2018

ABSTRACT

Since July 2001, the Space Telescope Imaging Spectrograph (STIS) onboard Hubble has operated on its Side-2 electronics due to a failure in the primary Side-1 electronics. While nearly identical, Side-2 lacks a functioning temperature sensor for the CCD, introducing a variability in the CCD operating temperature. Previous analysis utilized the CCD housing temperature telemetry to characterize the relationship between the housing temperature and the dark rate. It was found that a first-order 7%/°C uniform dark correction demonstrated a considerable improvement in the quality of dark subtraction on Side-2 era CCD data, and that value has been used on all Side-2 CCD darks since. In this report, we show how this temperature correction has performed historically. We compare the current 7%/°C value against the ideal first-order correction at a given time (which can vary between ~6%/°C and ~10%/°C) as well as against a more complex second-order correction that applies a unique slope to each pixel as a function of baseline dark rate and time. This second-order correction has shown that it can remove up to an additional ~5% dark counts than can the first-order correction, with the degree of improvement dependent on the temperature of the observation. Additionally, we present initial evidence suggesting that the variability in pixel temperature-sensitivity is significant enough to warrant a temperature correction that considers pixels individually rather than correcting them uniformly.

Contents

- Introduction (page 2)
- Temperature Dependence in Time (page 2)

Operated by the Association of Universities for Research in Astronomy, Inc., for the National Aeronautics and Space Administration.

- Second-Order Correction (page 6)
- Pixel-by-Pixel Correction (page 11)
- Conclusions (page 14)
- References (page 16)

1. Introduction

On May 16, 2001, STIS experienced a malfunction when its primary "Side-1" set of electronics short circuited. Ever since, STIS has operated on its redundant "Side-2" set. While nearly identical, Side-2 lacks the ability to monitor the temperature of the CCD and consequently cannot control the current to hold the CCD at a stable detector temperature. With a variable CCD temperature, Side-2 is affected by a variable dark rate, dependent on the detector temperature at the time of the observation. While STIS lacks a direct measurement of the detector temperature, it is able to track the CCD housing temperature (OCCDHT). The housing temperature can be used as a suitable proxy temperature measurement, with typical values ranging anywhere from ~15°C to ~23°C.

Shortly after the switch to Side-2, the temperature dependence of the STIS CCD dark rate was characterized [1]. This analysis found a 7%/°C relationship between the dark rate (at a reference temperature of 18°C) and the CCD housing temperature, which has been used as the correction factor for the STIS CCD temperature variability for all Side-2 observations. Now, after almost two decades of operation on Side-2, the quality of this correction is unknown. If the temperature dependence of the STIS CCD has evolved or changed in time, the temperature correction may need to change as well. In this report, we present findings from a historical study of Side-2 temperature correction. Using 1100s darks from October 2001 to May 2017, we characterize the temperature dependence as a function of time and assess the quality of the 7%/°C correction over time. We explore a second-order temperature correction, a different method that assigns a unique scale value to each pixel dependent on the baseline dark rate of the pixel and time, and evaluate any improvement in the quality of STIS CCD dark subtraction. Additionally, we present initial evidence suggesting that the variability in pixel temperature-sensitivity is significant enough to warrant a temperature correction that considers pixels individually rather than correcting them uniformly.

2. Temperature Dependence in Time

2.1 Current Correction

The current temperature correction was derived from a set of darks from the first three anneal periods (roughly corresponds to the first three months) of Side-2 operation. For

each anneal period, two dark frames taken at temperatures closest to 18°C (the chosen reference temperature) were compared against the rest of the dark frames in the anneal period. The ratio of dark rates as a function of temperature difference revealed a fractional change in dark rate per °C scaling relationship, referred to as a "scale value" hereafter. A key result of this original analysis was that the scale value varied depending on the dark rate itself, though this never made it into the reduction pipeline. Figure 1 shows a reproduction of this original result (Brown, et al.), plotting the scale value against dark rate.

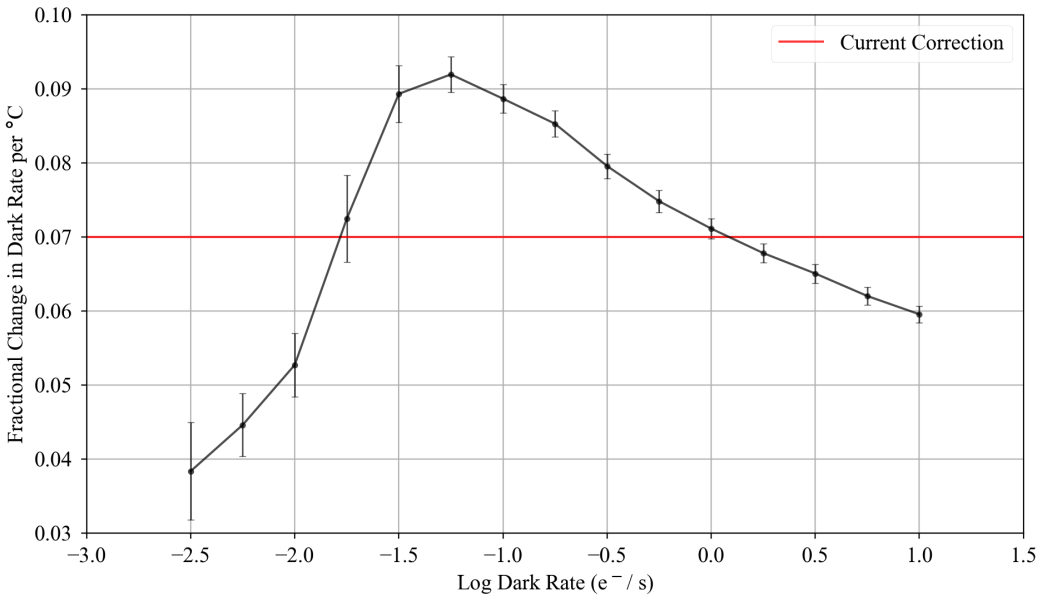

Figure 1.: Scale value curve of the STIS CCD dark rate when Side-2 operations began. Each point represents a bin of dark rates that is 0.25 dex wide centered on the point. Pixels exhibiting baseline dark rates within that bin were used to calculate the scale value for the center of the bin.

This curve shows that the scale value can range from ~4%/°C to ~9%/°C depending on which dark rate bin the pixel resides within. The current correction does not consider this dark-rate dependence as it applies a flat correction to the entire detector, referred to as a "first-order" correction. It is also important to note that the 7%/°C value was chosen because it reflects the midpoint of the range. One might expect the ideal first-order correction to center on the scale value associated with the median dark rate of the CCD instead, as the scale value would be tailored to the "typical" pixel. In this approach, the first-order scale value would be ~5%/°C instead. Since the median dark rate of the CCD has been slowly growing over time, this decision may have been an attempt at a temperature correction that remained stable with time.

2.2 Time Dependence

There are two possible time-dependent effects that would impact the effectiveness of the first-order correction:

The first is the linear growth of the median dark rate in time. Generally, we'd expect the majority of pixels to slowly move along the scale value curve, adopting the scale value corresponding to their growing dark rate over time. In the lifetime of Side-2, the log median dark rate has evolved from ~-2.3 to ~-1.75, which, assuming that the scale value curve has remained static, implies that the first order correction has stayed within ~2% of the ideal scale value thus far.

The second effect is the change in the scale value curve over time. While the first-order correction is inconsiderate of variations in the scale value of specific dark rate regions, the change in temperature sensitivity of pixels at that dark rate affects the quality of temperature correction for that subset of pixels. Additionally, shifts of the entire scale value curve would have an obvious impact on the ideal first-order scale value. In order to see how the scale value curve has changed in time, we've performed a historical survey of the majority of darks taken during the lifetime of Side-2. Figure 2 shows the change in scale value over time for a set of different dark rate bins.

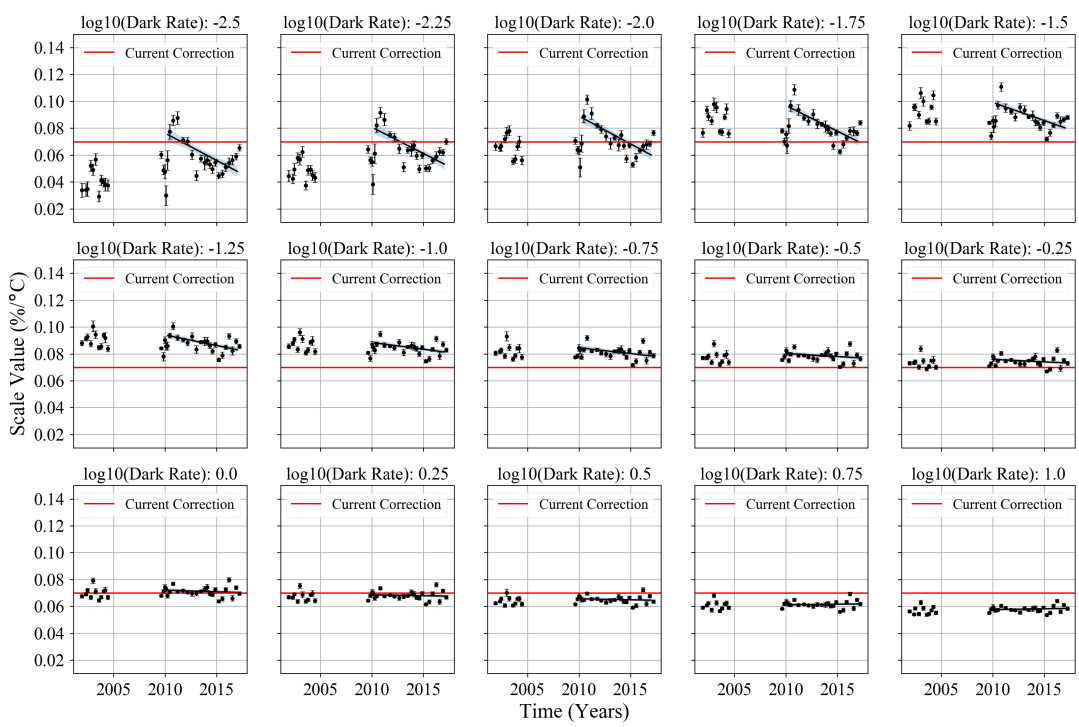

Figure 2.: Time evolution of each dark rate bin. Each bin is centered on the labelled log value and is 0.25 dex wide. A linear fit is added between 2010 and 2017 for visual aid.

Figure 2 indicates that the scale value curve's time evolution is dependent on the baseline dark rate itself. Generally, the larger the dark rate, the less change in scale value that rate has experienced in time. However, as the log median detector dark rate is currently ~–1.75, Figure 2 implies that the majority of pixels have been subject to changes in scale value over the lifetime of Side-2. Comparing each bin against the current temperature correction, we see that the discrepancy has been as large as 4%, once again primarily affecting cooler pixels.

An unexpected finding of this survey is that at some point in 2010, the CCD experienced a sudden rise in temperature sensitivity for cooler pixels. The reason behind this spike is not currently known. The temperature sensitivity since then has decayed linearly and leveled off at values measured prior to the spike. Figure 3 shows the evolution of the scale value after the 2010 spike. We once again see the difference in time-dependence between different dark rates, as cooler dark rates vary much more dramatically in time compared to warmer rates. Furthermore, we see that the dark rate corresponding to the peak scale value is slowly evolving in time.

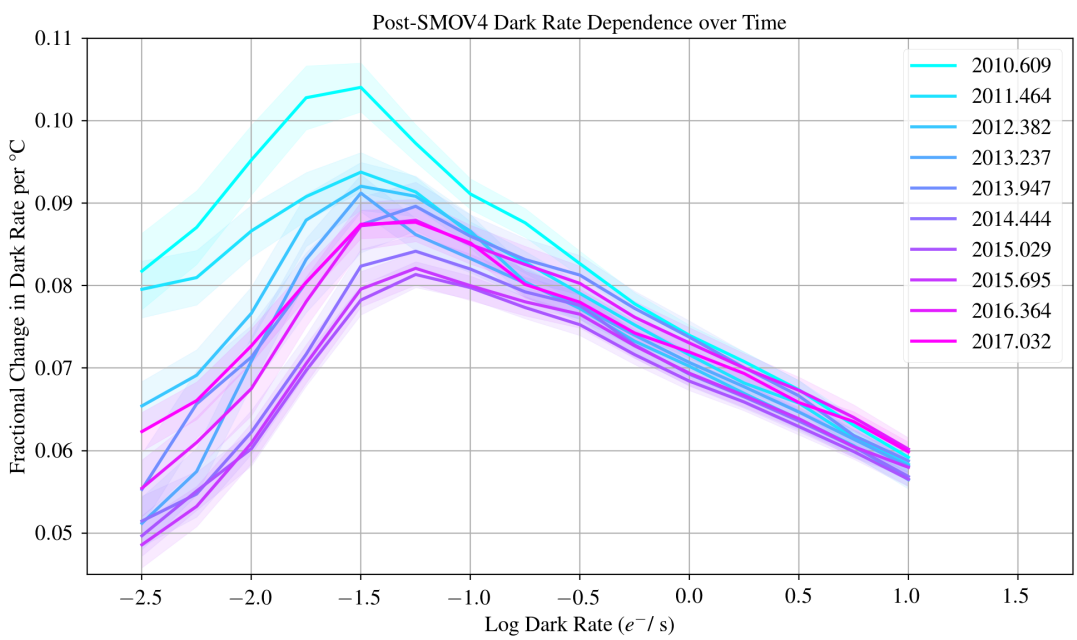

Figure 3.: Time evolution of the scale value curve post-SMOV4.

2.3 First-Order Correction in Time

Since there is a clear time-dependence to the temperature sensitivity of the CCD, the optimal first-order correction has likely fluctuated in time. With this in mind, we took a look at the ideal first-order correction in time against the current correction. The ideal first-order correction in this case is defined as the scale value corresponding to the

median dark rate on the scale value curve at a particular point in time, accounting for both the changing dark rate in time and the changing scale value curve in time. In Figure 4, this comparison over time is shown.

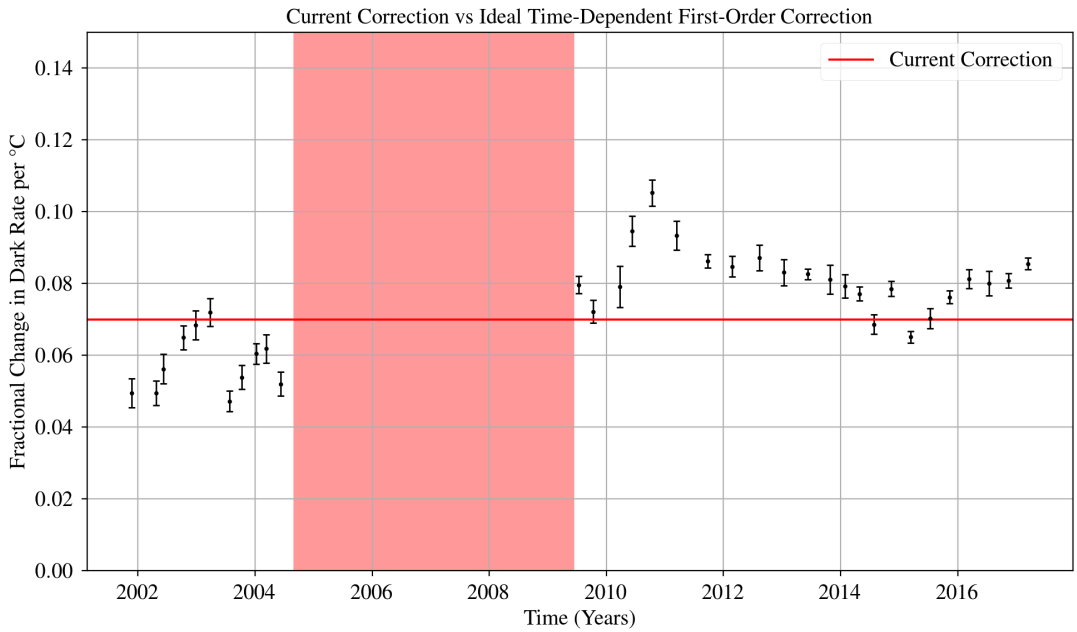

Figure 4.: Ideal first-order correction in time.

While the current correction is off by as much as 3%/°C at some points in time, this plot shows that the current 7%/°C correction has been as good of a choice as one could make for a singular correction value over the lifetime of Side-2. It is worthwhile to point out again that at the beginning of Side-2 the correction value was chosen to be 7%/°C despite the median dark value at the time corresponding to a 5%/°C correction. It's clear to see that if 5%/°C was chosen, this correction would be much less optimal for post-SMOV4 data. However, this figure shows that a simple improvement to the temperature correction would be to use a different scale value for pre-SMOV4 and post-SMOV4, using values that better reflect the temperature sensitivity of the median dark rate for each era.

3. Second-Order Correction

3.1 The Second-Order Correction Method

While previous analysis resulted in a first-order correction, a method where a best-fit linear temperature scale correction is used for all pixels on the detector, they also considered a second-order correction [1]. As originally detailed, a second-order correction is where pixels are temperature corrected individually based on their baseline dark rate.

Essentially, it uses the value on the scale value curve corresponding to the dark rate of a given pixel. On paper, a second-order method should yield a better result than a first-order method as each pixel is getting a scale value tailored to its dark rate. In practice, the original results were underwhelming as the second-order method failed to demonstrate noticeable improvement on the first-order method [1]. However, correcting for the apparent time-dependence of the detector (see Figure 2) offers an opportunity to revisit the second-order method. With the addition of a time-dependent component, the combined adjustments for both dark rate and time may result in a noticeable improvement over the current first-order correction.

In order to generate a unique scale value for a given dark rate and time, the scale value curves are bi-linearly interpolated to the observation date of the image and the log dark rate of the pixel. Figure 5 gives a sense of the scale value extracted for any region in rate-time space. Contrary to the original analysis, we only apply the second-order correction when scaling the superdark to the temperature of the science frame. The original analysis creates a second-order superdark, where each of the darks used to create the superdark is scaled to the reference temperature using the second-order method. We opted to instead use the first-order method when creating the superdark, which means that the superdark used in both the current method and the second-order method is identical. The reason for this is that the second-order superdarks were found to adversely affect the quality of the dark correction, more on this in §3.2.

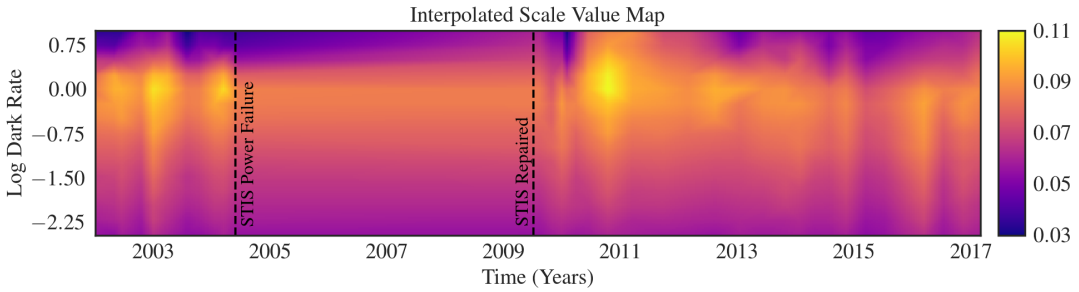

Figure 5.: Map of scale values generated by linearly interpolating scale value curve survey data. Region between Power Failure and Servicing Mission 4 is included for simplicity.

3.2 Performance Comparison

In dark correction, the purpose of the temperature correction is to scale the superdark to the temperature of the science observation, so that the dark subtraction removes as much of the temperature-altered dark level as possible. If the temperature correction is not accurate, it will affect the quality of the dark subtraction, either under-correcting or over-correcting the science image for the dark level. If the science image is instead a dark frame, the error in dark correction will present itself directly in the residual. Thus, dark frames are ideal for comparing the first-order and second-order temperature correction methods against each other as the magnitude of the residuals can be directly compared. In order to perform this comparison, a selection of darks were taken from a variety of post-SMOV4 anneal periods. For each dark, the corresponding superdark is scaled to the "science" temperature in both the first-order and second-order method. The percent difference between the dark and each scaled superdark for a 1024x7 pixel cutout is calculated, with this cutout representing the size of the STIS CCD spectroscopic extraction region. This cutout is located near the top of the detector and allows us to not be concerned with Charge Transfer Inefficiency (CTI), though there appears to be hardly any variation in temperature sensitivity with detector y-coordinate from initial investigation. Figure 6 compares the performance of each method by measuring the percentage of the dark level that remains after dark correction.

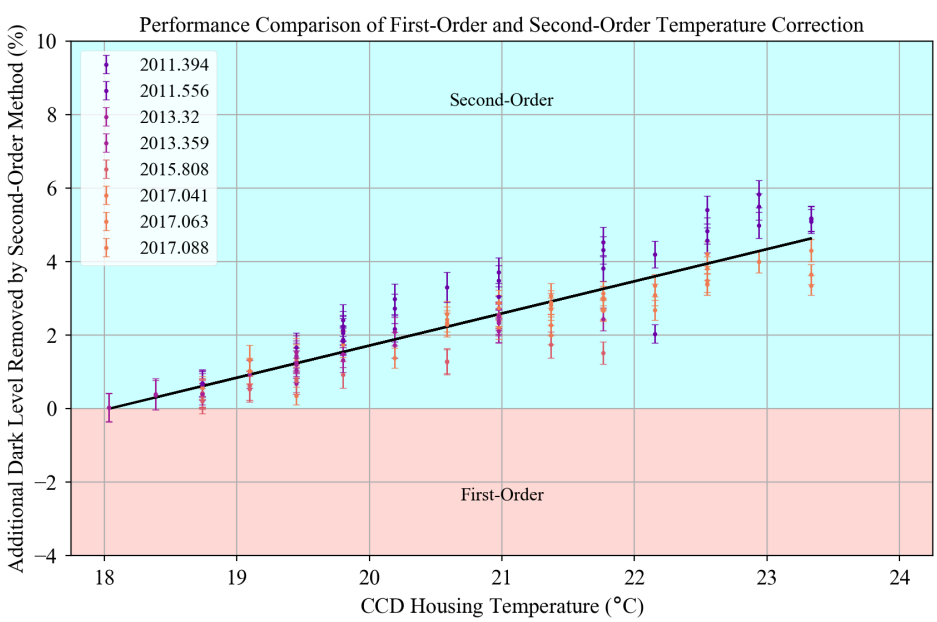

Figure 6.: Performance comparison of the first-order and second-order correction on a 1024x7 pixel slit. Positive percentages indicate improvement in favor of the second-order correction (e.g 2% would mean that the second-order method removes an additonal 2% of the dark level that the first-order method does not).

As the CCD housing temperature of the dark moves further from the reference temperature (18°C), the performance benefit of the second-order method grows. The percent improvement appears to trend linearly with temperature, at a rate of ~1%/°C. The difference in performance between anneal periods occuring at different times is also visible, as we see that in 2011, when the current 7%/°C correction was furthest from ideal (see Figure 4), the percent improvement of the second-order method is largest. Whereas in 2015 and 2017, when the current correction was closer to ideal, the percent improvement of the second-order method is less. However, at the highest temperatures, the discrepancy between anneal periods at each of these times is only ~1% compared to the ~5% total percent improvement, which indicates that the dark rate correction component of the second-order method is the primary driver for the overall improvement of the second-order method.

When following the original vision for a second-order method, applying the second-order scaling both in superdark creation and in dark subtraction, little improvement was observed from it in comparison to the first-order method. This was consistent with the conclusion that was arrived at in the original analysis [1]. However, this changed when we tried a blend of the two methods, applying the first-order method when creating superdarks, and applying the second-order method in the scaling of these superdarks to the science temperature during dark subtraction. This adjustment yields a considerable boost in performance, indicating that the second-order method does a poorer job of creating superdarks than does the first-order method. Its not clear why the first-order method would produce better superdarks, but the obvious empirical benefit to using first-order superdarks motivated using them over their second-order counterparts for this analysis. For the sake of potential further exploration of this issue, we'll detail the procedure for creating these second-order superdarks. Since the second-order method derives scale values based on the dark rate of a given pixel with respect to a given reference temperature (18 °C in this analysis), all of the component darks used in superdark creation need to be scaled down to the reference temperature first to have their dark rates sampled. This, in a somewhat self-defeating manner, requires the use of the first-order method to do this initial scaling. Once the reference temperature dark rates have been measured, second-order method scale values are calculated for each pixel based on these dark rates and are then applied to the original unscaled component darks to scale them down to the reference temperature. Again, this method has empirically performed worse in our analysis, though if one were determined to try to improve its performance, they might consider trying an iterative version of this method, where the reference temperature scale value obtained after the second-order scaling is used to calculate an adjusted scale value for that pixel. The second-order scaling with the adjusted scale values would be performed again, and this process would be repeated until the solution converges.

3.3 Scientific Impact

While it is clear that the second-order method offers a performance boost compared to current methodology, the practical science impact of this performance boost remains to be seen. An improvement of a few percent may be a negligible reduction in dark noise for the typical observation. In order to get a physical sense of the performance difference, we'll consider the worst case scenario. A temperature correction will have its largest effect when furthest from the reference temperature and when the exposure time of the observation is large. The exposure time doesn't factor into the percent improvement, but physically a larger exposure time allows for more dark current to build on the detector. Figure 7 shows the residual distribution of pixels across the whole detector for an 1100s dark taken at 23.3308 °C, near the warmest temperature the CCD housing is typically seen at, for both 2011 and 2017.

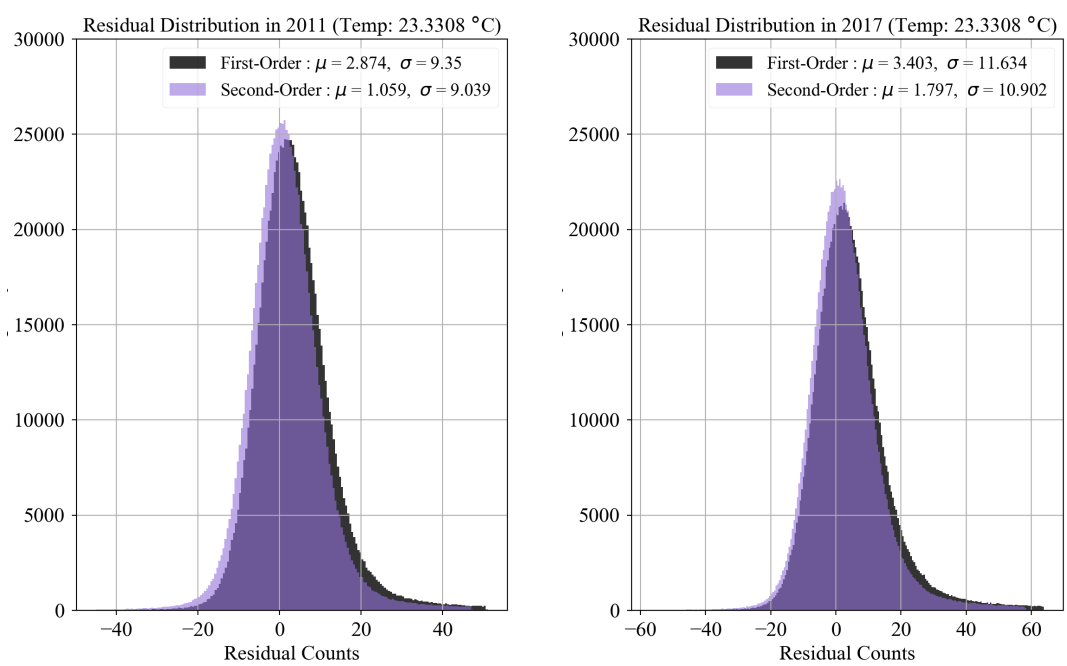

(a) Residual in 2011

(b) Residual in 2017

Figure 7.: Residual comparisons between correction methods

In 2011, the median dark rate was approximately 0.014 $e^-/s/pix$. For an 1100s dark, this means that for the typical pixel we expect ~15.4 e^-/pix, or counts at gain=1. Factoring in the temperature of the observation, this value is elevated to ~22 counts. We see in Figure 7a that correcting the dark noise using a first-order scaled superdark yields a typical pixel residual of ~2.9 counts, while the second-order scaled superdark yields ~1.1 counts. While we see that the typical pixel has had almost 2 more counts of dark

noise removed by using the second-order method, we also see that the spread of residuals has tightened.

Likewise, in 2017, the median dark rate was approximately 0.02 $e^-/s/pix$. For an $1100s$ dark at gain=1, this translates to ~31 counts after temperature scaling has been accounted for. We see in Figure 7b that correcting the dark noise using a first-order scaled superdark yields a typical pixel residual of ~3.4 counts, while the second-order scaled superdark yields half of that. We again see that the spread of residuals has tightened.

Removing an additional 2 counts of dark noise can be significant when considering low signal-to-noise data. However, this is not what the typical user would see as a benefit from switching to the second-order method. Again, this is the worst case scenario for temperature correction, where exposure time is large and the temperature of the observation is furthest away from the reference temperature. As of the publication of this ISR, only ~14% of all STIS CCD scientific observations were taken with exposure times larger than $500s$, though co-added observations are not considered in this statistic and will be affected by temperature at the same level. An even smaller subset of this data was taken when the CCD housing temperature was far away from the reference temperature, as the typical post-SMOV4 CCD housing temperature sits around ~20.5 °C. So despite the performance improvement of the second-order correction, the vast majority of STIS CCD data is not sensitive to the performance difference. It follows from this that, in general, it is hard to produce a temperature correction method that would improve the dark correction by an appreciable amount for the typical observation. However, as shown in Figure 7, the improvement from the second-order correction comes primarily in the reduction of the mean pixel residual dark level after dark correction. While the spread is also tightened, it's only by a small amount. Of the two parameters, tightening the spread is much more important as the average residual level will be removed in further processing (background correction), leaving only the residual dark level that is spread out from the average. Thus, for a temperature correction method to deliver a greater improvement to the dark correction, finding a way to reduce the residual spread is key.

4. Pixel-by-Pixel Correction

4.1 Pixel Scale Value Variability

At any given point in time, the dark rate of a pixel may be modeled as a function of temperature sensitivity and baseline dark rate (Figure 1). The main idea behind the second-order correction (§3.1) is to tailor a scale value to each pixel dependent on the measured dark rate of the pixel at the reference temperature. The inherent assumption made is that there is a high agreement in scale value for a group of pixels all measured near the same dark rate. Figure 8 shows the dark rate ratio for every pixel on the detector

between a dark frame at an elevated temperature and two dark frames at the reference temperature.

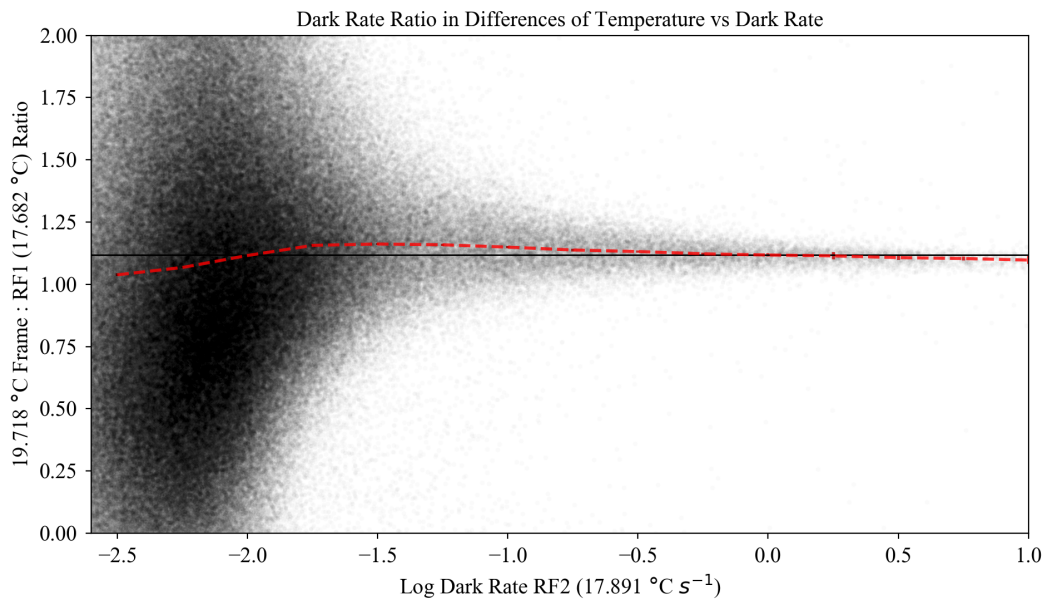

Figure 8.: Ratio of dark rate for a dark frame at an elevated temperature against two darks near the reference temperature. Median ratios are calculated across the dark rate space.

For a given dark rate, the spread in pixel ratios is very large (at least for cooler dark rates). If all of the pixels of a given dark rate had the same temperature sensitivity, we'd expect the pixel cloud to conform much more closely to the dashed median line. This variability suggests that while the second-order method delivers a better scale value than the first-order method, the scale value is still not ideal for a large subset of the pixels on the detector. In fact, for individual pixels the ideal scale value can vary significantly. Figure 9 shows a distribution of the ideal scale value calculated for individual pixels in a single anneal period.

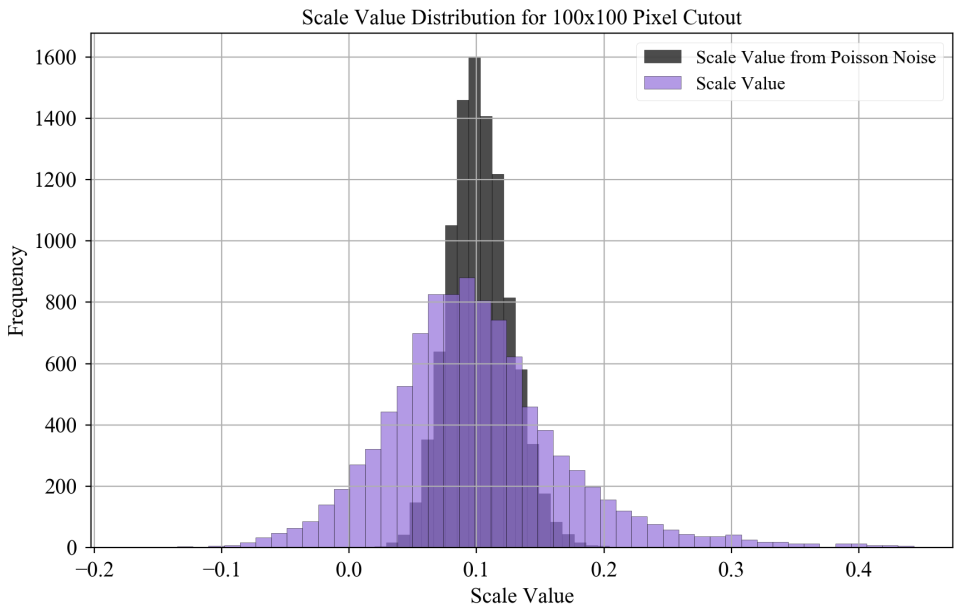

Figure 9.: Distribution of pixel scale values within a single anneal. A poissonian distribution about the median scale value is also plotted to show confidence that the variability in scale value is statistically significant.

Figure 9 sheds some light on why the improvement the second-order method offers over the first-order method is less pronounced. For a typical pixel, the second-order method may adjust the assigned scale value by only 1%/°C - 3%/°C. These adjustments are relatively small compared to the large range of scale values a pixel can exhibit within a given dark rate bin, which suggests that the second-order correction and first-order correction perform similarly for a large subset of the pixels.

This pixel-specific scale value variability suggests an even higher order correction method. A method that applies a unique scale value for each pixel based on that pixel's own historical temperature sensitivity. With this pixel-by-pixel method, it is possible that the temperature correction would yield a residual with much tighter spread, as it does not generalize the temperature sensitivity to any pixel characteristic and corrects for each pixel independently. Such a correction may be a natural product of a STIS CCD pixel history project, as done for the ACS and WFC3 CCDs [2], where the dark current of every pixel on the CCD is historically analyzed to track its behavior and stability.

4.2 Changes in Pixel Sensitivity

While correcting pixels individually based on historical temperature sensitivity is promising, considering pixels individually introduces some challenges not present when gen-

eralizing based on dark rate. For one, cosmic ray events can have discrete impacts on pixels on the detector. Damage caused by these events can affect the sensitivity and stability of a given pixel either temporarily or long term [2]. These sensitivity variations will likely affect the temperature sensitivities of these pixels as well.

Additionally, the annealing process likely impacts temperature sensitivity of individual pixels. As an anneal involves heating and cooling the detector in an attempt to release charge traps and correct hot pixels, there is a possibility that pixels adopt completely different temperature sensitivities from one anneal to the next. Figure 10 takes a look at two subsequent anneal periods and plots a cumulative distribution of the change in scale value for pixels across an anneal. From Figure 10, it's clear that the scale values derived

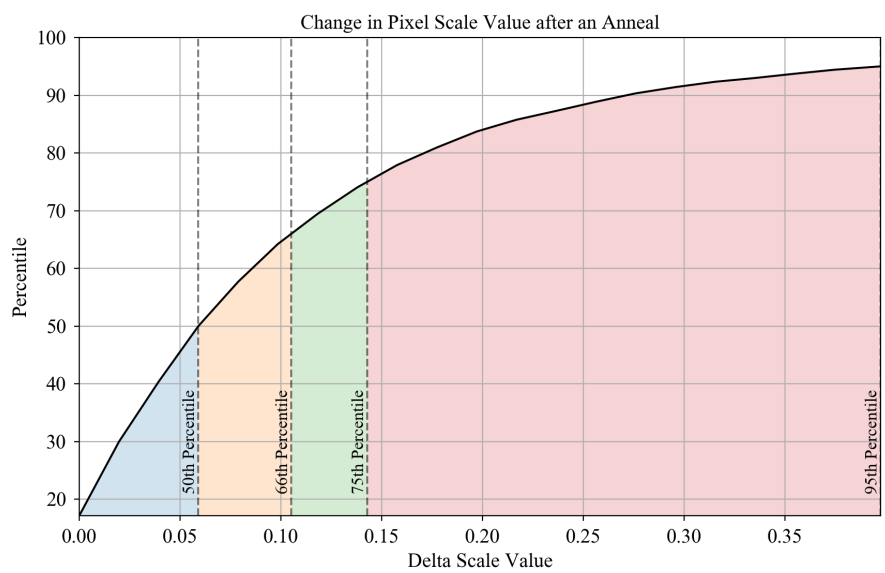

Figure 10.: Change in pixel sensitivity after an anneal.

for each pixel based on its behavior over an anneal period is not consistent between anneal periods. Almost half of the detector is 5%/°C different and a third of the detector is over 10%/°C different. This suggests that for each pixel, an entirely new scale value would need to be assigned for each anneal period. While doable, this introduces a significant amount of overhead to the reduction pipeline that the first-order method does not require.

5. Conclusions

When the STIS CCD switched to its Side-2 set of electronics, the inability to hold the detector at a fixed temperature introduced a temperature-dependent variability to the

dark rate measured on the detector. At a given point in time, the scaling relationship between temperature and dark rate has been dark rate dependent (see Figure 1), as the dark rate of the pixel at the reference temperature (18°C) affects the sensitivity of that pixel to changes in temperature. Additionally, the dark rate dependence of the temperature sensitivity has evolved in time, particularly in the case of cooler pixels (log dark rate $<$ -1.5) (see Figure 2). Overall, the scale value varies anywhere from 3%/°C to 11%/°C for a given dark rate and time. Despite these factors, the STIS pipeline has employed a flat "first-order" correction factor of 7%/°C.

While not the ideal temperature correction, the current first-order correction has performed at a satisfactory level in time. Historically, the ideal flat correction factor has varied between ~5%/°C and ~10%/°C (see Figure 4). If there could only be one correction value for the history of STIS Side-2, the current value 7%/°C ends up very close to the optimal choice. The application of a "second-order" method, a method that applies a unique scale value to each pixel dependent on the pixels dark rate and the date of observation (see Figure 5), yields up to a 5% performance boost compared to the first-order correction (see Figure 6). For an 1100s exposure, this 5% boost physically translates to ~1-2 counts of additional dark current per pixel being removed during dark correction. However, a 5% boost represents the best case scenario for the second-order correction, at temperatures furthest away from the reference temperature (see Figure 7). For the typical observation where the temperature is closer to the reference temperature, the difference between the performance of both methods is comparable.

If the scientific use case demanded the best possible temperature correction, the first place to look would be a "pixel-by-pixel" correction. This method treats pixels individually, assigning them their own unique scale value based on their own measured temperature sensitivity, either historically or for a specific anneal period. In support of this method, the range of scale values a given pixel can adopt is actually quite significant, with pixels scaling by as much as 30%/°C (see Figure 9). Pixels at this level would be significantly under-corrected by both the first-order and second-order method, meaning that a pixel-specific method may offer a significant improvement in the overall dark correction. However, considering pixels individually presents its own challenges. Namely, each time the CCD is annealed, a large subset of pixels on the detector experience large changes to their temperature sensitivity (see Figure 10), indicating that a new scale value would need to be generated for each pixel every single anneal period. This would introduce a significant amount of overhead to the reduction pipeline that is not present with the current first-order correction method. Based on this, a pixel-by-pixel correction likely never replace the first-order correction, but would instead be offered as a standalone tool for the select user cases that require a better temperature correction.

If you are concerned about this effect for your low-SNR CCD observations, while we (at least at the time of writing this ISR) don't have concrete plans for implementing a higher-order temperature correction at this time, we can make a recommendation for

mitigating most of this effect. It may be worth considering dithering your observations along-slit, as this will minimize the systematic errors (fixed pattern noise) that result from the application of a non-optimal temperature correction.

Acknowledgements

Special thanks to John Debes, Sean Lockwood, and Allyssa Riley for providing helpful discussion and suggestions, as well as the STIS team in general for providing valuable feedback over the course of the project.

Change History for STIS ISR 2018-05

Version 1: October 10, 2018- Original Document

References

[1] Brown, T., 2001, STIS Instrument Science Report 2001-03

[2] Borncamp, D., Grogin, N., Bourque, M., & Ogaz, S., 2017 ACS Instrument Science Report 2017-05