

When there is a discrepancy between the information in this technical report and information in JDox, assume JDox is correct.



STScI | SPACE TELESCOPE
SCIENCE INSTITUTE

JWST TECHNICAL REPORT

Title: NIRSpec DARK Reference Files and Plans for Cycle 4 DARK Calibration Programs	Doc #: JWST-STScI-009226 Date: 4 November 2025 Rev: -
Authors: Karakla, D., Glidic, K., Hayes, C., Böker, T., and Muzerolle-Page, J. Phone: 410-338-4947	Release Date: 18 December 2025

1. Abstract

Master Bias and Dark reference files for FULL frame NIRSpec observations have been delivered to CRDS in March 2025. Separate darks for each of the detectors and modes (Traditional, and IRS2 FULL frame) were created for a range of dates from the end of Commissioning until the end of September 2024. Subarray darks were similarly created for traditional mode and cover the period of Cycle 2.

2. Introduction

In March 2025, the Bias and Dark reference files delivered to the Calibration Reference Data System (CRDS) marked a transition from the European Space Agency's (ESA) Commissioning Analysis Procedure (CAP) code to the James Webb Space Telescope (JWST) Science Calibration Pipeline (which we reference as ST code) for producing these reference files. There are several processing differences between the ESA CAP code and the JWST pipeline, most notably the inclusion of linearity correction in the latter, as well as minor differences in handling bad pixels; these are discussed in detail in Section 4.

In Table 1, we summarize the calibration programs from Commissioning (Cycle 0) through Cycle 3, detailing the date ranges covered, program IDs (PIDs), and any significant differences in processing for the dark and bias reference files delivered to CRDS. Each full-frame traditional mode (NRS, NRSRAPID) input file is 88 groups in length, while each IRS2 mode (NRSIRS2, NRSIRS2RAPID) input file is 200 groups long. All subarray data use the traditional readout mode; however, the number of groups varies between subarrays and cycles because calibration programs were designed based on the longest requested exposures for each subarray within a cycle.

Operated by the Association of Universities for Research in Astronomy, Inc., for the National Aeronautics and Space Administration under Contract NAS5-03127

Check with the JWST SOCCER Database at: <https://soccer.stsci.edu>

To verify that this is the current version.

Table 1. Full-Frame and Subarray Bias/Dark Calibration Programs for Cycles 0-3

Date Range	PID	Reference File Generation and Delivery
Commissioning (December 2021 – June 2022)	1121 (full-frame) 1130 (subarray)	Completed at the end of Commissioning with the ESA CAP code. Untouched.
Cycle 1 (July 2022 – June 2023)	1484 (full-frame) 1495 (subarray)	Full-frame TRAD and subarray darks completed at the end of Cycle 2 with the ESA CAP code. Corrected full-frame IRS2 darks to remove artifacts and re-delivered to CRDS with the new ST code which now includes linearity correction.
Cycle 2 (July 2023 – June 2024)	4455 (full-frame) 4456 (subarray)	New TRAD and IRS2 full-frame and subarray bias and dark reference files processed with the ST code and delivered to CRDS.
Cycle 3 (July 2024 – June 2025)	6633 (full-frame) 6634 (subarray)	New TRAD and IRS2 full-frame bias and dark reference files processed with the ST code and delivered to CRDS.

3. Pre-Filtering of Input Files to Omit Artifacts

To generate reference files for each cycle, we first created a list of darks within defined date ranges (listed in Table 1), which excludes those taken in the MIRROR configuration because they can have undispersed light leaks.

We identified additional artifacts in files processed to rate files that required further culling of the input list of dark exposures:

Dark observations taken after a switch from IRS2 to TRAD readout mode, and vice versa, often exhibit a residual a bias structure (“tree rings” and “picture frame” effects) as shown in Figure 1. By examining a list of ASIC temperatures from the engineering database measured before, during, and after the exposures, we noticed a clear division between files (Table 2). Files with ASIC temperature *differences* ($T_{\max} - T_{\min}$) $> 0.02\text{K}$ were omitted as a means of flagging input files containing bias structure that likely occurred after a mode switch.

Far fewer subarray observations were taken, so we accepted a wider range of ASIC delta-temperature range (0.13K) than for full-frame darks) to preserve enough input files to produce a deliverable dark.

- As shown in Figure 2, there is an apparent glow or light leak from the edge of one of the Micro-Shutter Assembly (MSA) quadrants that appears in some input dark rate files, causing diagonal stripes. These features were more noticeable in darks taken with the PRISM and also in the IRS2 mode where the 1/f noise is lower. We suspect this is a glow from the edges of the MSA (or it could be a short-bandwidth light leak) that gets dispersed onto the detector.

Check with the JWST SOCCER Database at: <https://soccer.stsci.edu>
To verify that this is the current version.

PRISM exposures have the lowest dispersion of all dispersers and show the effect more prominently. We decided to omit all darks taken with the PRISM from our final lists of input files for full-frame, for both traditional and IRS2 readout mode.

- Finally, we found and omitted additional dark images showing obvious persistence in the fixed slits and other areas of the MSA from previous observations of very bright sources, (e.g. lamp exposures). For reference, many looked similar to the dark in Figure 25 taken after a MOS S-flat observation, though not all have the bright linear features caused by the long slit configuration since persistence can happen in any MOS configuration when the sky is bright.




Figure 1. Residual bias artifacts. A TRAD dark for detector NRS1 showing the “picture frame” and “tree-ring” artifacts.

Table 2. ASIC Temperature distribution of input dark files

	A	B	C	D	E	F	G	H	I	J	K	L	M
	2024-11-06	15:16:26,421	stpipe - INFO	-	jw04455071001_	320	1_0000	1_nrs1	2023-11-28 PRISM	43.39	43.34	43.48	0.14
64	2024-11-06	15:16:26,421	stpipe - INFO	-	jw04455080001_	220	1_0000	1_nrs2	2024-01-29 PRISM	43.15	43.10	43.23	0.14
65	2024-11-06	15:16:26,423	stpipe - INFO	-	jw04455031001_	420	1_0000	1_nrs1	2024-02-09 PRISM	43.58	43.47	43.62	0.14
66	2024-11-06	15:16:26,418	stpipe - INFO	-	jw04455023001_	430	1_0000	1_nrs1	2024-02-25 G395H	43.58	43.48	43.62	0.14
67	2024-11-06	15:16:26,421	stpipe - INFO	-	jw04455094001_	330	1_0000	1_nrs1	2024-05-05 MIRROR	43.40	43.35	43.49	0.14
68	2024-11-06	15:16:26,421	stpipe - INFO	-	jw04455051001_	420	1_0000	1_nrs2	2023-07-12 G395M	43.13	43.08	43.21	0.13
69	2024-11-06	15:16:26,416	stpipe - INFO	-	jw04455064001_	330	1_0000	1_nrs2	2023-10-10 G395M	43.12	43.06	43.19	0.13
70	2024-11-06	15:16:26,420	stpipe - INFO	-	jw04455066001_	220	1_0000	1_nrs1	2023-10-24 G395H	43.40	43.35	43.49	0.13
71	2024-11-06	15:16:26,420	stpipe - INFO	-	jw04455068001_	420	1_0000	1_nrs2	2023-11-05 MIRROR	43.14	43.09	43.22	0.13
72	2024-11-06	15:16:26,415	stpipe - INFO	-	jw04455069001_	220	1_0000	1_nrs1	2023-11-13 MIRROR	43.41	43.36	43.49	0.13
73	2024-11-06	15:16:26,419	stpipe - INFO	-	jw04455074001_	220	1_0000	1_nrs1	2023-12-20 G395M	43.40	43.35	43.49	0.13
74	2024-11-06	15:16:26,417	stpipe - INFO	-	jw04455075001_	220	1_0000	1_nrs2	2023-12-24 G395H	43.14	43.09	43.22	0.13
75	2024-11-06	15:16:26,416	stpipe - INFO	-	jw04455080001_	220	1_0000	1_nrs1	2024-01-29 PRISM	43.41	43.36	43.49	0.13
76	2024-11-06	15:16:26,419	stpipe - INFO	-	jw04455081001_	420	1_0000	1_nrs1	2024-02-11 G395H	43.40	43.35	43.48	0.13
77	2024-11-06	15:16:26,422	stpipe - INFO	-	jw04455036001_	230	1_0000	1_nrs1	2024-03-15 G395H	43.58	43.48	43.61	0.13
78	2024-11-06	15:16:26,415	stpipe - INFO	-	jw04455039001_	730	1_0000	1_nrs1	2024-04-03 G395M	43.58	43.49	43.62	0.13
79	2024-11-06	15:16:26,421	stpipe - INFO	-	jw04455072001_	220	1_0000	1_nrs1	2023-12-04 G395M	43.40	43.35	43.47	0.12
80	2024-11-06	15:16:26,418	stpipe - INFO	-	jw04455095001_	420	1_0000	1_nrs2	2024-05-12 G235H	43.13	43.09	43.20	0.12
81	2024-11-06	15:16:26,421	stpipe - INFO	-	jw04455051001_	420	1_0000	1_nrs1	2023-07-12 G395M	43.39	43.35	43.46	0.11
82	2024-11-06	15:16:26,417	stpipe - INFO	-	jw04455064001_	330	1_0000	1_nrs1	2023-10-10 G395M	43.38	43.34	43.45	0.11
83	2024-11-06	15:16:26,417	stpipe - INFO	-	jw04455068001_	420	1_0000	1_nrs1	2023-11-05 MIRROR	43.40	43.35	43.47	0.11
84	2024-11-06	15:16:26,421	stpipe - INFO	-	jw04455075001_	220	1_0000	1_nrs1	2023-12-24 G395H	43.39	43.35	43.46	0.11
85	2024-11-06	15:16:26,417	stpipe - INFO	-	jw04455095001_	420	1_0000	1_nrs1	2024-05-12 G235H	43.39	43.35	43.45	0.10
86	2024-11-06	15:16:26,422	stpipe - INFO	-	jw04455004001_	220	1_0000	1_nrs1	2023-08-08 G235M	43.63	43.62	43.64	0.02
87	2024-11-06	15:16:26,418	stpipe - INFO	-	jw04455011001_	320	1_0000	1_nrs1	2023-09-21 G140H	43.62	43.61	43.63	0.02
88	2024-11-06	15:16:26,418	stpipe - INFO	-	jw04455013001_	330	1_0000	1_nrs2	2023-10-10 G395M	43.36	43.35	43.36	0.02
89	2024-11-06	15:16:26,424	stpipe - INFO	-	jw04455027001_	520	1_0000	1_nrs1	2024-01-10 G395M	43.63	43.62	43.64	0.02
90	2024-11-06	15:16:26,424	stpipe - INFO	-	jw04455024001_	520	1_0000	1_nrs2	2024-02-09 G395H	43.37	43.36	43.38	0.02
91	2024-11-06	15:16:26,422	stpipe - INFO	-	jw04455035001_	520	1_0000	1_nrs2	2024-03-10 MIRROR	43.38	43.37	43.39	0.02
92	2024-11-06	15:16:26,422	stpipe - INFO	-	jw04455047001_	220	1_0000	1_nrs1	2024-05-30 G395H	43.63	43.62	43.63	0.02
93	2024-11-06	15:16:26,422	stpipe - INFO	-	jw04455047001_	220	1_0000	1_nrs2	2024-05-30 G395H	43.38	43.37	43.39	0.02
94	2024-11-06	15:16:26,422	stpipe - INFO	-	jw04455049001_	220	1_0000	1_nrs1	2024-06-26 G140H	43.63	43.61	43.63	0.02
95	2024-11-06	15:16:26,417	stpipe - INFO	-	jw04455049001_	220	1_0000	1_nrs2	2024-06-26 G140H	43.38	43.37	43.39	0.02
96	2024-11-06	15:16:26,416	stpipe - INFO	-	jw04455001001_	330	1_0000	1_nrs1	2023-07-12 G395M	43.63	43.62	43.63	0.01
97	2024-11-06	15:16:26,424	stpipe - INFO	-	jw04455001001_	330	1_0000	1_nrs2	2023-07-12 G395M	43.38	43.37	43.39	0.01

The ASIC temperature minimum, maximum, average and difference during dark exposures taken during Cycle 2 (for example). The table was constructed using data from the engineering database. Columns K and L contain the minimum and maximum ASIC temperatures. Column M contains the difference ($T_{diff} = T_{max} - T_{min}$) temperature. Tables like this for each Cycle helped to identify those files with steep temperature changes that show obvious artifacts. There is a clear division in T_{diff} (column M). Those files with $T_{diff} > 0.02K$ were omitted from final full-frame darks.

Check with the JWST SOCCER Database at: <https://soccer.stsci.edu>
To verify that this is the current version.




Figure 2. Residual light leaks seen in NRS1 darks due to exposures taken with the PRISM (upper panels), which causes over-subtraction in science data reduced with these darks (lower panels). These data have been processed to remove the reference pixel striping normally seen in IRS2 data. (Figure provided by E. Bergeron).

Creating the final dark and bias reference files was an iterative process that included initial pre-filtering and processing to rate files, followed by removing problematic rate files with strong bias artifacts (e.g., tree rings and picture frame), and re-combining. Tables 3 and 4 summarize the finalized full-frame reference files delivered to CRDS.

Note: Just prior to the start of Cycle 3 (late May 2024), we implemented a change to the scheduling system to prevent darks from being taken immediately after a mode switch. Additionally, at the start of Cycle 3 (July 2024), we changed our monitor program to only obtain dark exposures in G140H (high dispersion) to mitigate light leaks and improve monitoring efficiency. These changes apply to both full-frame and subarrays.

Both changes have increased the number of useful observed darks; however, the mode switch restriction resulted in fewer opportunities to schedule TRAD darks due to the large preference for IRS2 observations. This factor impacted our calibration plans for Cycle 4 Dark monitoring (see Section 7).

4. Description of Dark Processing

4.1. General Description of the Dark Pipeline Code

Until April 2024, the Commissioning Analysis Procedure (CAP-008) code developed by ESA (Birkmann et al. 2018, and Birkmann, 2022), provided the framework for producing NIRSpec biases and darks, as well as updating the bad pixel masks to include new hot pixels. This code, however, depended on the ESA processing pipeline, which is no longer supported. To keep the code up-to-date and make use of continuing developments in the JWST pipeline, we converted the ESA CAP code into a routine that is stored in a Grit project for NIRSpec calibration codes called “*nirspec_reffiles*”. We began running this new code in February 2025 to generate new Cycle 1-3 darks (some of which replaced existing darks for the same periods but which included known light leaks). The new routine, “*nrs_bias_dark*,” depends on the JWST pipeline and, therefore, differs from the CAP-008 code in a few ways, including pre-processing of the input darks and how DQ flags are set.

The *nrs_bias_dark* routine reads a list of input uncalibrated dark files and processes them to create bias, dark, and mask reference files. This process is broken up into five steps, described in more detail in the following sections:

1. Making a bias – *make_bias_step*
2. Calibrating input darks – *prepare_dark_pipeline*
3. Producing an average dark ramp – *average_dark_step*
4. Synthesizing a dark cube – *synthesize_dark_step*
5. Analyzing the bias, dark, and mask reference files – *analyze_bias_dark*

4.1.1. *make_bias_step*

The *make_bias_step* uses uncalibrated darks as input to produce a new bias reference file. This step first copies and performs a reference pixel correction on the first and second groups of each input dark, preserving the average bias level of the first group, so that it is not removed in the reference pixel correction. The first and second groups are averaged (per pixel) across all input files, to create average group 1 and group 2 count images. The bias reference file is then generated by projecting the group 1 and 2 counts back to time = 0, i.e., $\text{bias} = 2 * \text{group 1} - \text{group 2}$. Uncertainties on the bias are calculated by combining the standard deviation of the counts of the first and second groups across all input darks. Pixels with bias uncertainties that are 5s above the mean are flagged as “UNRELIABLE_BIAS” in the bias reference file DQ plane.

4.1.2. *prepare_dark_pipeline*

The second step is to process the uncalibrated input darks to produce calibrated dark ramps and dark rate images. This step of the processing is done using the following JWST pipeline *calwebb_detector1* steps:

1. *group_scale* – corrects reported on-board group averaging
2. *dq_init* – Initialize a bad pixel mask (hot pixels are unflagged in the bad pixel mask so that they can be processed)
3. *saturation* – flag saturated pixels
4. *superbias* – correct for the detector bias
5. *refpix* – apply reference pixel corrections
6. *linearity* – perform linearity corrections
7. *jump* – detect jumps due to cosmic rays and “snowballs”
8. *ramp_fit* – fit slopes to calibrated ramps to measure count rates

While some of these steps differ slightly from how they were implemented in the ESA calibration pipeline (e.g., reference pixel corrections) there are two significant differences in the pipeline processing here. 1) The order of linearity correction and dark subtraction is different in the two codes: the JWST pipeline requires linearity-corrected darks because it performs dark subtraction after linearity correction. 2) In addition to basic jump detection to identify cosmic rays, the JWST pipeline also includes the flagging of “snowballs” in individual dark frames (extremely energetic cosmic rays that saturate the detector and can affect the ramps of nearby pixels), which reduces the number of artifacts in the final dark reference files.

4.1.3. *average_dark_step*

Using the calibrated dark ramps created from individual exposures, the *average_dark_step* produces an average dark ramp and flags the following types of pixels: resistor-capacitance (RC), random telegraph noise (RTN), and bad reference pixels. Before calculating the average dark ramp, the code for this step first masks any groups that occur after a jump in each of the input ramps. For each pixel, an average dark ramp is then calculated for each group by taking a 3s-clipped mean of the counts across all input darks. Similarly, the uncertainties on the counts for each group are calculated by taking the standard deviation of the good counts across each of the input darks.

The pixel flagging that is performed for RC, RTN, and bad reference pixels is described as follows:

RC Pixels: RC pixels are a type of hot pixel with an extremely high dark rate that follows an exponential curve as seen in resistor-capacitance circuits (Schlawin et al. 2020; and references

therein). Due to the high, nonlinear dark rates, these pixels are often flagged with saturation or have their entire ramp flagged with jump detections in the JWST pipeline. In either case, these pixels typically do not have their ramps fit and cannot be detected as hot pixels by inspecting their dark rates, which is done for more typical hot pixels (see Section 5.1.4). Instead, to detect these pixels and flag them, we analyze the calibrated dark ramps. RC pixel candidates are first identified in each calibrated ramp by selecting any pixels in the first group with counts that are more than 10σ above the mean counts across all pixels in group 1. To remove cases of large cosmic rays that happen in the first group, we remove any candidates whose neighbors are also 10σ above the group 1 mean. We also require that RC candidates are flagged with JUMP (due to their large non-linear ramps) or DO_NOT_USE (e.g., if the large dark rate causes saturation) in the next two groups. Any RC candidates that are not flagged as JUMP or DO_NOT_USE in subsequent groups will have their ramps fit and can otherwise be detected via the hot pixel flagging done using the calculated dark rates in the following step. Pixels that are flagged as an RC candidate in more than half of the input dark ramps are flagged as RC and DO_NOT_USE (due to their high dark rate) in the DQ array of the final dark reference file.

RTN Pixels: RTN pixels are pixels which toggle between two or more signal levels as shown in Figure 3. The frequency of the toggling and the difference between signal levels ranges amongst RTN pixels (Rauscher et al. 2007). RTN pixels that toggle with a very low frequency may simply appear as a single jump in an observed ramp, and those with sufficiently small level differences may simply be lost in the read noise. Therefore, it is important to flag RTN pixels, though we do not currently flag them as DO_NOT_USE. The defining feature of RTN pixels is that they toggle between states and provide a rare case of negative jumps (i.e., when toggling from a higher state to a lower one), which can be used to identify these pixels. To identify RTN candidates, the *average_dark_step* counts the number of negative jumps across all pixels in a calibrated dark ramp and flags any pixels that have more negative jumps than 5σ above the mean. If that pixel is identified as an RTN pixel in more than one input dark ramp, it is flagged with TELEGRAPH in the final dark reference file DQ array.

This detection method is slightly different from that used in the ESA processing, which marked any pixels with N jumps (averaged over all input dark exposures) $> (1 - 0.9973) * (N \text{ groups} - 1) * 4$. Comparing the two methods for a similar set of input dark exposures, the threshold for RTN pixel flagging set by the ST based dark production is more conservative. For example, we compared the IRS2 mode CRDS dark reference files *jwst_nirspec_dark_0445.fits* (generated with the ESA pipeline using dark data taken between 01/01/2023 – 06/30/2023) to an updated reference file generated with the ST pipeline, *jwst_nirspec_dark_0498.fits* (using many of the same input darks taken over a similar time frame 02/13/2023 – 06/30/2023). Between the two reference files, the ESA pipeline used for *jwst_nirspec_dark_0445.fits* reports about 2.2% of pixels as RTN pixels, whereas the ST pipeline flags only 0.61% of pixels as RTN, which is closer to the $<1\%$ estimate identified in earlier H2RG tests done by Rauscher et al. (2007).

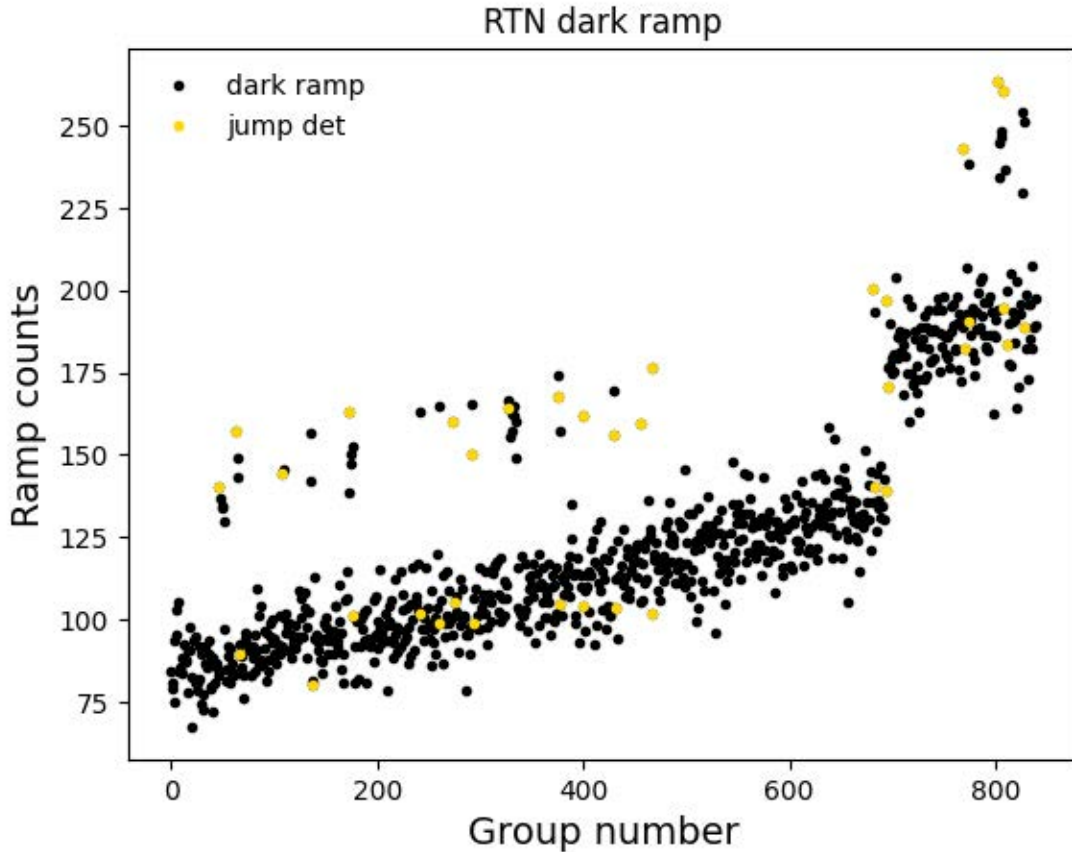


Figure 3. Calibrated dark ramp for a random telegraph noise (RTN) pixel. The ramp of this pixel shows a gradual increase in counts due to the dark rate in addition to toggling between two states. The two states of this pixel can be seen as the scatter around the ramp starting around 80 counts and the more sparsely populated ramp beginning around 140 counts. Groups where a jump (positive or negative) has been detected by the JWST pipeline are shown in yellow. Note that a cosmic ray occurs around group 700, highlighting the benefit of identifying RTN pixels based on negative jumps.

Bad Reference Pixels: Bad reference pixels are identified by averaging the counts and errors in each reference pixel across all groups and identifying any reference pixels where the average counts and errors are more than 4σ away from the median count/error across all reference pixels. Any such pixels are flagged with `BAD_REF_PIXEL` and `DO_NOT_USE` in the final dark DQ array.

4.1.4. *synthesize_dark_step*

The *synthesize_dark_step* first averages the individual calibrated dark rate images, to produce mean dark rates for each pixel. The average dark rates are used to flag hot pixels, defined as pixels with dark rates > 1 DN/s, which corresponds to more than 100 times the average dark rate in full frame exposures. The input bad pixel mask reference file is updated to add hot pixels to the DQ plane and written out into a new reference file. Because some hot pixels lead to false positive detections as RTN pixels, the RTN flags that were set in *average_dark_step* are removed for any hot pixels.

To create the final dark cubes, for most pixels, we project the average dark rates to calculate the expected dark counts per group. Averaging the dark rates for individual pixels over many exposures provides higher quality estimates of the dark current than using the averaged dark

Check with the JWST SOCCER Database at: <https://soccer.stsci.edu>
To verify that this is the current version.

ramps for a given pixel, because ramp fitting reduces the group-to-group noise that is present in averaged ramps. Because the SNR ratios on the final dark cubes can be relatively low (with a threshold of $\text{SNR} = 2.5$ defined as the minimum criteria for traditional and subarray readout modes), some pixels are reported with negative dark rates. Typically, about 1% of pixels have negative rates. Negative dark rates are intrinsically unphysical, and these measurements typically appear to be noisy rate measurements (i.e., the scatter to the low-end tail of the dark rate noise distribution), so pixels with negative dark rates use the frame-averaged dark rate to project their dark cubes.

The approach described above is valid because dark ramps are typically linear (at least after linearity correction), however not all pixels show linear dark ramps. For example, neighbors of hot pixels often show a nonlinear ramp in the first several groups due to inter-pixel capacitance (IPC) coupling between the pixel and the neighboring hot pixel, as shown in Figure 4.

Therefore, for any pixel flagged as `DO_NOT_USE` or any pixel that neighbors a hot or RC pixel, this step fills the dark cube with the average dark ramp for that pixel using the product of *average_dark_step*.

Once the final dark cubes are created, they are written out to a new dark reference file. This reference file also contains an error array, a pixel and group DQ plane, and a copy of the average dark rate image and its uncertainties. The error array for the dark cube uses the group-by-group errors calculated in the average dark ramps. The pixel DQ plane records any pixels that were flagged as RC, RTN, or bad reference pixels, as well as a set of `DO_NOT_USE` flags (particularly for RC and bad reference pixels). The pixel DQ plane also flags a more general group of pixels with `UNRELIABLE_DARK`. This includes 1) pixels that were flagged with any DQ flags in the pixel DQ planes across all calibrated dark ramps (i.e., pixels that have pipeline processing data quality concerns across all groups in all inputs), but does not include pixels affected by cosmic ray jumps which are recorded to the group DQ plane, 2) pixels flagged as HOT, RC, or `BAD_REF_PIX`, and 3) pixels with high dark rate uncertainties (i.e., errors that are more than 3s above the mean error across all pixels). Finally, a “DRK” extension contains the average dark rate image that was used to produce the dark cube, and a “UNC” extension containing the uncertainties on those dark rates.

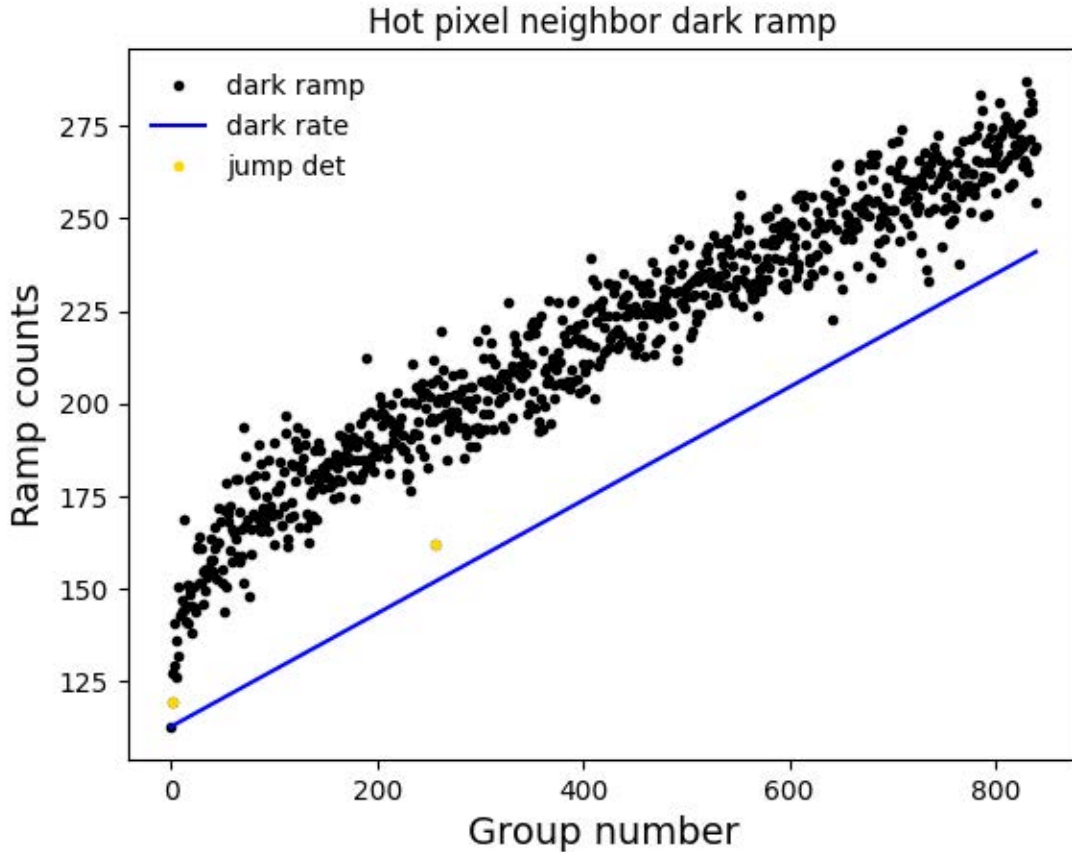


Figure 4. Calibrated dark ramp for a pixel neighboring a hot pixel. The ramp of this pixel shows a non-linear increase in counts in the first several groups due to IPC coupling of this pixel to a neighboring hot pixel. The counts for each group are shown as black points, with groups that have a jump detected flagged in yellow. The blue line shows the projection of the linear dark rate calculated for this pixel (beginning from the counts in the first group), which under-predicts the dark counts at in all groups.

4.2. *analyze_bias_dark*

Finally, the last step of the *nrs_bias_dark* pipeline analyzes the output reference files and compares them to the most recent CRDS reference files. This step produces diagnostic plots of the superbias, dark rate image, and the difference between the new superbias compared to the CRDS reference file. It does so by plotting images of each of these, as well as histograms of the counts, count rates, or differences, respectively. In addition to diagnostic plots, this step also reports and logs basic statistics on several quantities, such as the sigma-clipped average bias level, average dark rate, the dark S/N, counts of the hot pixels, etc. It also reports whether the reference files pass the various quality thresholds, which may be useful for evaluating whether to delivering a new reference file.

In the end, we decided that the only criterion for delivery that is important is the SNR in the final dark. We adopted 2.5 as the minimum acceptable SNR for traditional mode (NRS, NRSRAPID), and 5.0 for IRS2 mode (NRSIRS2, NRSIRS2RAPID). The SNR is calculated as the ratio of the dark rates to uncertainties for each pixel, then applying a 3-sigma clip of the median value across all pixels. The adopted limits attempt to balance the need for short-period pixel masks to correct for the evolving hot pixel population, with the number of observed input darks unaffected by persistence and other artifacts. The short-period darks shown in Tables 3 and 4 met this criterion

and were accepted for delivery to CRDS. For the files delivered to CRDS in March 2024 that are discussed in this document, results were produced using the *nrs_bias_dark* routine with jwst version 1.16.0, and CRDS context jwst_1293.pmap.

5. The Creation of Full-frame Temporal Reference Files and Comparison Master Reference File

We created full-frame “temporal” (short-period) dark reference files taken over periods ranging from 6 to 14-months, to monitor changes in the dark level and spatial distribution, and to track the evolution of hot pixels. These were created by collecting enough good files (without evident artifacts) over the shortest time periods possible to achieve the minimum signal-to-noise (SNR) criterion that we adopted, as described above. They are labelled as “T0”, “T1”, etc. in Figure 5, and Tables 3 and 4.

No significant temporal differences were seen in the darks over the entire period, apart from the growth in the number of hot pixels with time. (Böker et al., 2025)


Tables 3 and 4 present a summary of the delivered (temporal) dark reference files, and notes the growth in hot pixels. The PEDIGREE column in Table 4 is the period over which input darks were collected. The USEAFTER date in the table is the start date selected for the application of a given reference file, which is chosen to be the middle of the PEDIGREE period.

We also produced long-term (“master”) darks spanning the period from the end of commissioning to mid-January 2025, for comparison with the period darks. They include some taken after the (slight) ASIC heater adjustment made on September 4th, 2024.

Given that the hot pixel population increases over time, with roughly 100 new hot pixels per month (Böker et al. 2025), there are inherent problems with the application of dark reference files that have a mismatch of hot pixels with the science data that it is being applied to:


- 1) For science data observed near the reference dark USEAFTER date, which is in the middle of the period over which the input darks are collected (the “PEDIGREE” in Table 4), the dark reference file may include many hot pixels that appeared after the science data were taken, resulting in the over-correction of good science pixels.
- 2) For science data observed after the reference dark PEDIGREE period, there may be hot pixels in the science data that were not present when the input darks were taken, resulting in the under-correction of science pixels.

Since these issues are exacerbated for the long-period master darks, the master darks were not delivered to CRDS as they are deemed not appropriate for science data reduction.

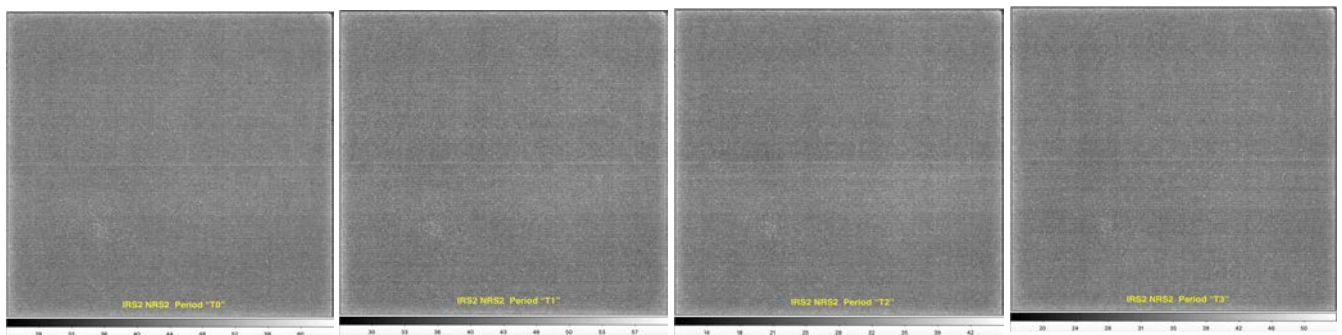


5a. New Traditional Dark Reference Files – NRS1 (top panel) NRS2 (bottom panel)

Check with the JWST SOCCER Database at: <https://soccer.stsci.edu>
To verify that this is the current version.



5b. New IRS1 NRS1 Dark Reference Files



5c. New IRS2 NRS2 Dark Reference Files

Figure 5. New Temporal Darks. Temporal (period) darks are labelled “T0”, “T1”, etc. and correspond to the date ranges shown in Tables 3 and 4. All figures show the last group of the dark cube.

Full frame bias and dark reference files were created by combining the useable input biases and darks from the periods shown in Tables 3 (the column labelled “CRDS Pedigree”, which contains actual dates of the first and last input darks in a given period), and 4 (the column labelled “Time Range”, which denotes the dates over which we searched for appropriate darks).

Table 3. Summary of Dark Reference Files Delivered to CRDS.

Label	Cycle	CRDS PEDIGREE	CRDS USEAFTER	CRDS dark	CRDS bias	Timestamp
TRAD-NRS1						
T0	1 - 2	INFLIGHT 2022-07-26 2023-09-26	2023-02-25T00:00:00	0491	0513	20250103T2011
T1	1 - 2	INFLIGHT 2023-06-05 2024-03-31	2023-11-01T00:00:00	0487	0509	20241108T1957
T2	2 - 3	INFLIGHT 2024-02-29 2024-09-26	2024-06-14T00:00:00	0488	0517	20241112T1924
TRAD-NRS2						
T0	1 - 2	INFLIGHT 2022-07-26 2023-09-18	2023-02-21T00:00:00	0497	0514	20250103T2358
T1	1 - 2	INFLIGHT 2023-06-05 2024-03-31	2023-11-01T00:00:00	0486	0508	20241108T1856
T2	2 - 3	INFLIGHT 2024-03-17 2024-09-26	2024-06-22T00:00:00	0512	0531	20250107T0009
IRS2-NRS1						
T0	1	INFLIGHT 2023-02-13 2023-06-30	2023-04-23T00:00:00	0498	0512	20241114T2313
T1	1 - 2	INFLIGHT 2023-04-17 2023-09-21	2023-07-05T00:00:00	0494	0511	20241110T1847
T2	2	INFLIGHT 2023-07-02 2024-02-09	2023-10-21T00:00:00	0489	0507	20241114T1707
T3	2 - 3	INFLIGHT 2024-01-10 2024-07-16	2024-04-14T00:00:00	0495	0516	20241116T1604
IRS2-NRS2						
T0	1	INFLIGHT 2023-02-13 2023-06-30	2023-04-23T00:00:00	0499	0518	20241124T1344
T1	1 - 2	INFLIGHT 2023-04-25 2023-09-21	2023-07-09T00:00:00	0496	0505	20241124T1945
T2	2	INFLIGHT 2023-07-12 2024-04-15	2023-11-28T00:00:00	0492	0510	20241125T0214
T3	2 - 3	INFLIGHT 2024-02-09 2024-08-26	2024-05-19T00:00:00	0490	0515	20241125T1212

Check with the JWST SOCCER Database at: <https://soccer.stsci.edu>
To verify that this is the current version.

Table 4. New Full-frame dark rates, SNR, and hot pixel populations.

Label	Time range	Months	SNR	# good input files	Med SCI Dark Rate DN/s	New Hot Pixels (Comparison)	Time stamp
TRAD-NRS1							
T0	07/01/22-10/01/23	14	5.25	16	0.0083	9022 (6914 mask_0051)	20250103T2011
T1	04/01/23-04/01/24	12	4.42	11	0.0084	9530 (8397 mask_0074)	20241108T1957
T2	02/01/24-10/01/24	8	5.96	19	0.0086	10111 (8397 mask_0074)	20241112T1924
TRAD-NRS2							
T0	07/01/22-10/01/23	14	3.15	14	0.0049	4184 (2613 mask_0049)	20250103T2358
T1	04/01/23-04/01/24	12	2.88	11	0.0051	4823 (4160 mask_0073)	20241108T1856
T2	02/01/24-10/01/24	8	2.98	12	0.0050	5573 (4160 mask_0073)	20250107T0009
IRS2-NRS1							
T0	01/01/23-07/01/23	6	11.61	11	0.0076	8207 (7096 mask_0055)	20241114T2313
T1	04/01/23-10/01/23	6	13.36	15	0.0074	8428 (7337 mask_0071)	20241110T1847
T2	07/01/23-03/01/24	8	12.02	13	0.0076	8715 (7337 mask_0071)	20241114T1707
T3	12/01/23-08/01/24	8	11.79	13	0.0074	9235 (7337 mask_0071)	20241116T1604
IRS2-NRS2							
T0	01/01/23-07/01/23	6	5.96	12	0.0036	3866 (3196 mask_0054)	20241124T1344
T1	04/01/23-10/01/23	6	6.44	12	0.0038	4155 (3490 mask_0072)	20241124T1945
T2	07/01/23-05/01/24	10	6.05	12	0.0036	4490 (3490 mask_0072)	20241125T0214
T3	12/01/23-09/01/24	9	5.33	14	0.0035	5155 (3490 mask_0072)	20241125T1212

6. The Creation of SUBARRAY Reference Files

Given the large number of subarrays and limited scheduling time, we observe only a fraction of subarrays each cycle. Therefore, unlike full-frame mode, where multiple files are delivered, we typically plan to deliver just one dark and one bias reference file per requested subarray each cycle, assuming the reference file meets our SNR criterion. The bad pixel mask derived from full-frame analysis does not account for interpixel capacitance (IPC); therefore, we update subarray darks primarily to reflect the increasing population of hot pixels and their impact on neighboring pixels caused by IPC. Table 5 lists all subarray darks delivered to CRDS through the end of Cycle 3.

Note: Entries marked *n/a* do not always mean the observation was not planned, sometimes we did not meet the required SNR for delivery. Additionally, Table 5 lists USEAFTER dates only for delivered reference files in CRDS, as there are too many file names to list them all.

Table 5: Summary of Subarray Dark/Bias Reference Files in CRDS

Subarray (and appropriate dispersers)	Cycle 0	Cycle 1	Cycle 2	Cycle 3
ALLSLITS	2022-01-01 T00:00:00	2023-01-01 T00:00:00	Epoch 1: 2023-10-15 T00:00:00 Epoch 2: 2024-03-16 T00:00:00	2024-12-26 T00:00:00
SUB2048 (G140M, PRISM)		n/a	2023-12-29 T00:00:00	n/a
SUB2048 (G140H, G235M)		n/a		2025-01-05 T00:00:00
SUB2048 (G235H)		2023-05-01 T00:00:00	2024-03-01 T00:00:00	n/a
SUB2048 (G395M)		n/a	2024-02-01 T00:00:00	2025-01-13 T00:00:00
SUB2048 (G395H)		2022-06-30 T00:00:00	2024-05-10 T00:00:00	n/a
SUBS200A1 (G140M, G140H, G235M, G235H, G395H, PRISM, MIRROR)		n/a	2023-08-18 T00:00:00	2025-02-02 T00:00:00
SUBS200A1 (G395M)		2023-05-10 T00:00:00	2023-10-02 T00:00:00	2025-02-28 T00:00:00
SUBS200A2/GEN (G140M, G140H, G235M, G235H, G395H, PRISM, MIRROR)		n/a		2025-04-21 T00:00:00
SUBS200A2 (G395M)		n/a		2022-01-01 T00:00:00
SUBS400A1/GEN (G140M, G140H, G235M, G235H, G395H, PRISM, MIRROR)		2022-07-01 T00:00:00	2023-11-22 T00:00:00	n/a
SUBS400A1 (G395M)		n/a		
SUB512/GEN (G140M, PRISM)		n/a	2024-06-05 T00:00:00	n/a
SUB512S/GEN (G140M, PRISM)		n/a		

Note: Subarrays for SUB1024A and SUB1024B are not included in this table as no deliveries have been made since Commissioning.

7. Analysis of Bias and Dark Reference Files

7.1. Comparison with Earlier Reference Files

We examined each bias and dark reference file produced with the new ST code after artifact filtering. For full-frame data, we compared the temporal files to each other, to the previous ESA product, and to the master reference files produced for the entire period from the start of Cycle 1 to the end of the first quarter of Cycle 3, to detect any temporal variations that would signal a change in the instrument or its electronics. For the subarrays, a representative ST product was compared to a previously delivered ESA file to assess for any temporal changes and differences in processing methods.

Until we have identified a different way of tracking the evolving population of hot pixels, we must keep producing short-period darks and their associated masks to correct the IPC in pixels neighboring a hot pixel in the science data from different time periods. There are no significant global changes seen in the darks, so if we can somehow accomplish the tracking of hot pixels in a different way, we might be able to instead make high SNR darks to apply to science data.

7.1.1. Full-frame Reference File Comparisons

Figures 6-9 show data for traditional (NRS, NRSRAPID) mode, and Figures 10-13 show IRS2 (NRSIRS2, NRSIRS2RAPID) mode results. The figures are formatted in a common way and presented in pairs – even-numbered figures show bias comparisons, while odd-numbered figures show dark comparisons. In the top panel of each figure, the last group of the master bias (or dark) reference file is shown on the left, and the difference image with respect to a given “period” bias (or dark) reference file is shown on the right. In the bottom panel of each figure, we show a plot of the median column from the difference image to illustrate the magnitude of any spatial variations. In some cases, we also show a plot of the median row to illustrate the magnitude of the $1/f$ noise of the difference image for comparison. In the plots, areas shaded blue are within 1σ of the median spatial difference seen in the “master” bias (or dark), which is reported at the top of the plot, along with the median value. In traditional mode darks, the median column differences are noticeably smaller than the $1/f$ noise differences. In the IRS2 darks, the $1/f$ noise residuals are significantly reduced, as expected, and therefore, are not presented. In some of the images and median column plots, slight amp-to-amp differences can be seen. This is particularly noticeable in Figure 12.

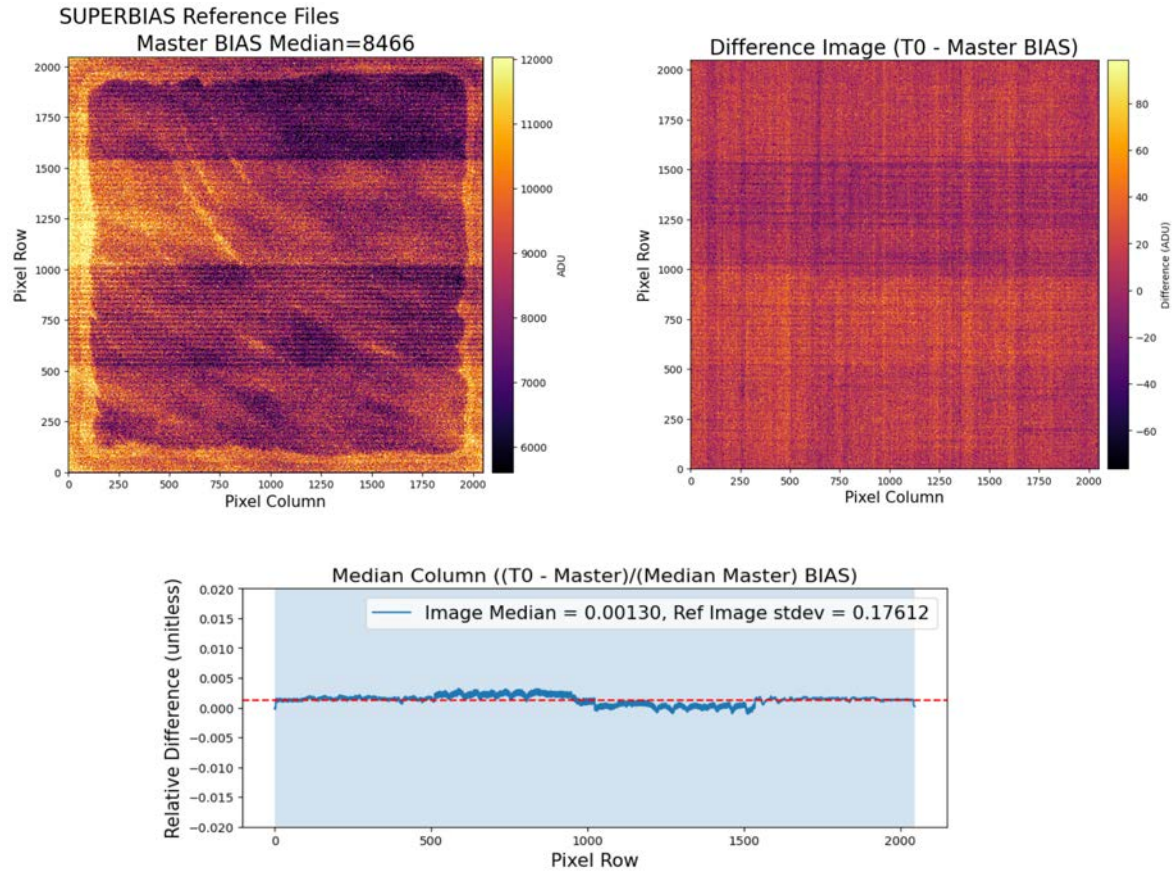


Figure 6: NRS1 TRAD Temporal (T0 Period) Bias Image Comparison with the Master Bias. Typically, bias differences between the master and temporal biases are very small (\ll stdev of the master bias). There are consistent weak features (like 1/f noise), and amp level differences, and row-to-row differences which are worse in some amps. The plot shows the median column of the difference image to illustrate that the differences are small compared to the 1-sigma value of the master bias image.

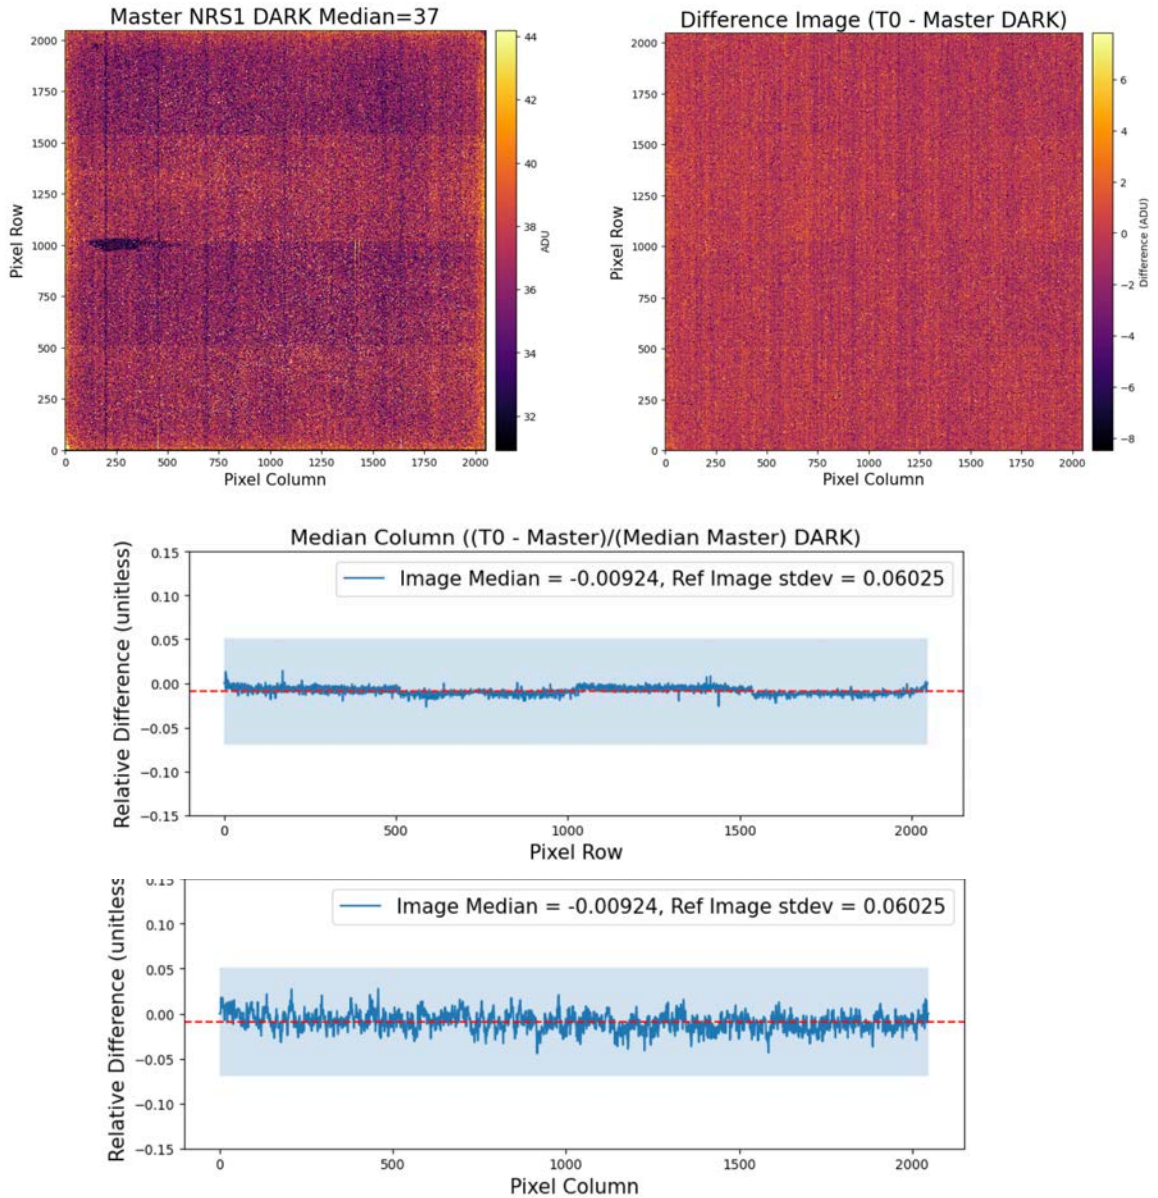


Figure 7: NRS1 TRAD Temporal (T0 Period) Dark Image Comparison with the Master Dark. Dark differences are small (\ll stdev of Master dark) and less than differences due to $1/f$ noise, as seen in the lower plot of the median row of the difference image.

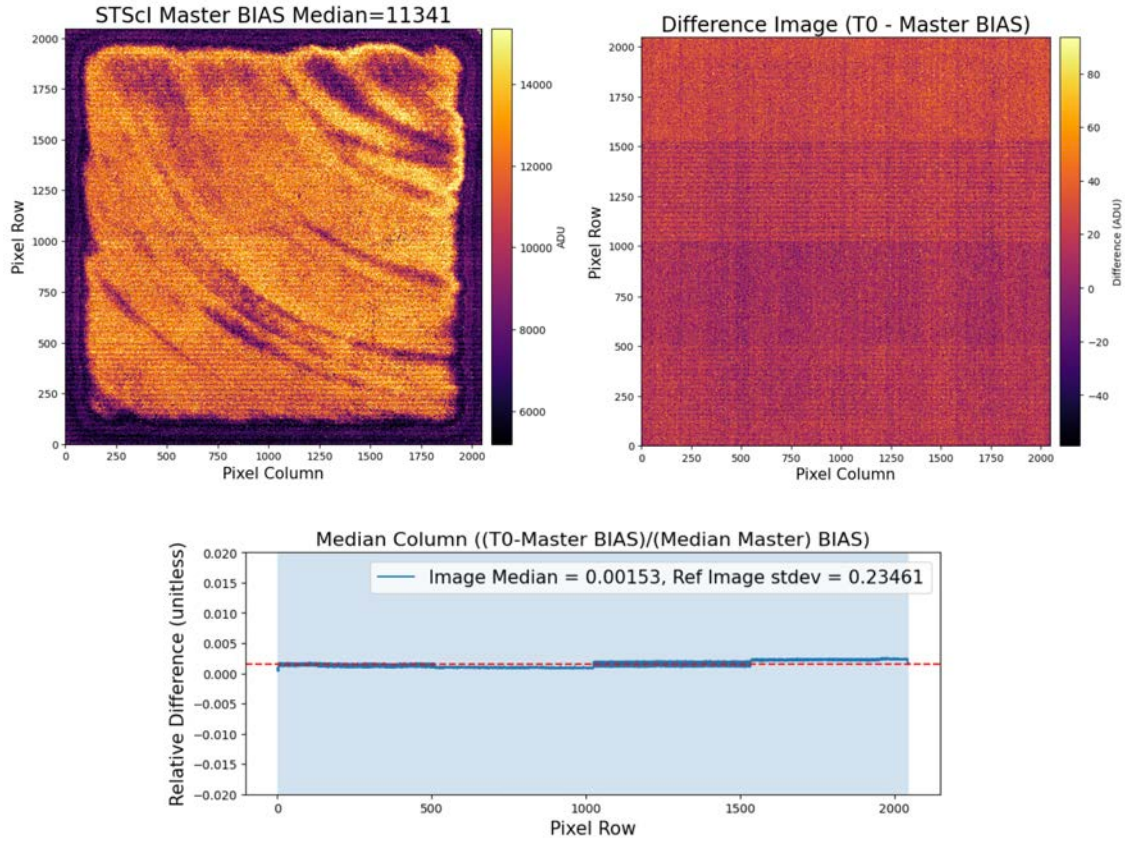


Figure 8: NRS2 TRAD Temporal (T0 Period) Bias Image Comparison with the Master Bias. There can be small differences in the bias over time: there are amp level differences, and odd-even row differences which are worse in some amps. All are \ll stdev (Master Bias).

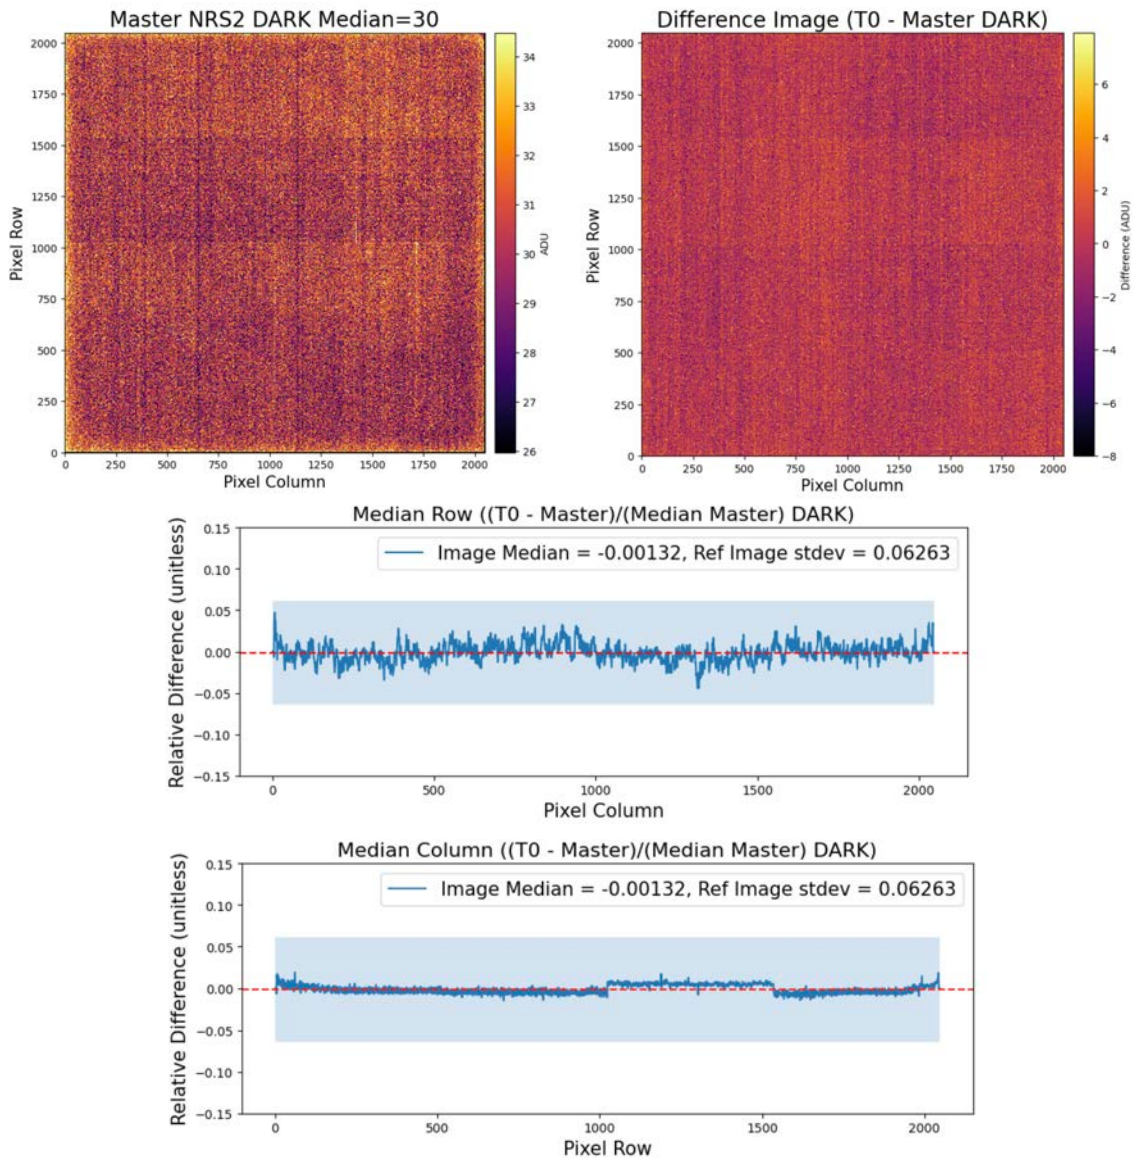


Figure 9: NRS2 TRAD Temporal (T0 Period) Dark Image Comparison with the Master Dark. Less than 2% differences are seen (\ll stdev of Master dark). Differences are within 1/f differences, as seen in the median row plot.

Check with the JWST SOCCER Database at: <https://soccer.stsci.edu>
 To verify that this is the current version.

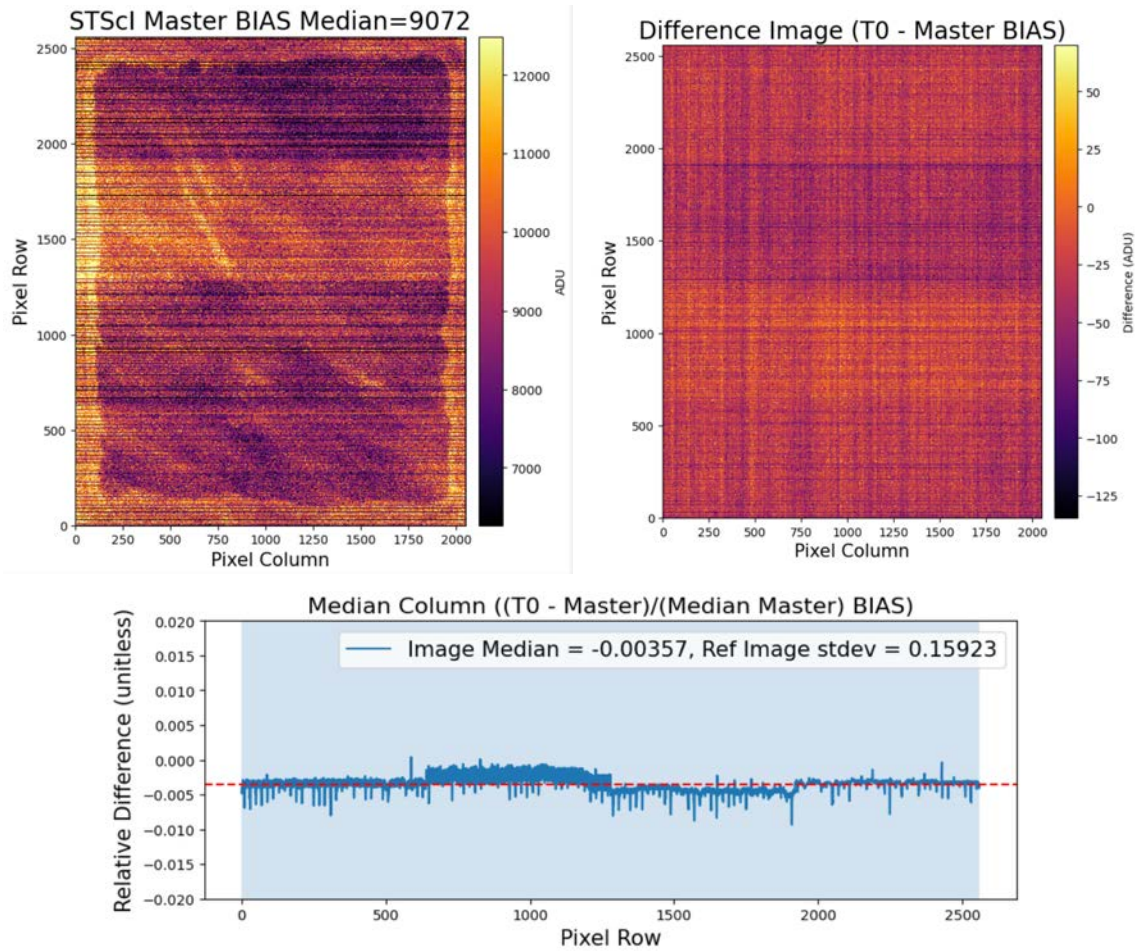


Figure 10: NRS1 IRS2 Temporal (T0 Period) Bias Image Comparison with the Master Bias. Temporal bias from period T0 compared to the master bias. Spikes seen in the plot are caused by the reference pixels.

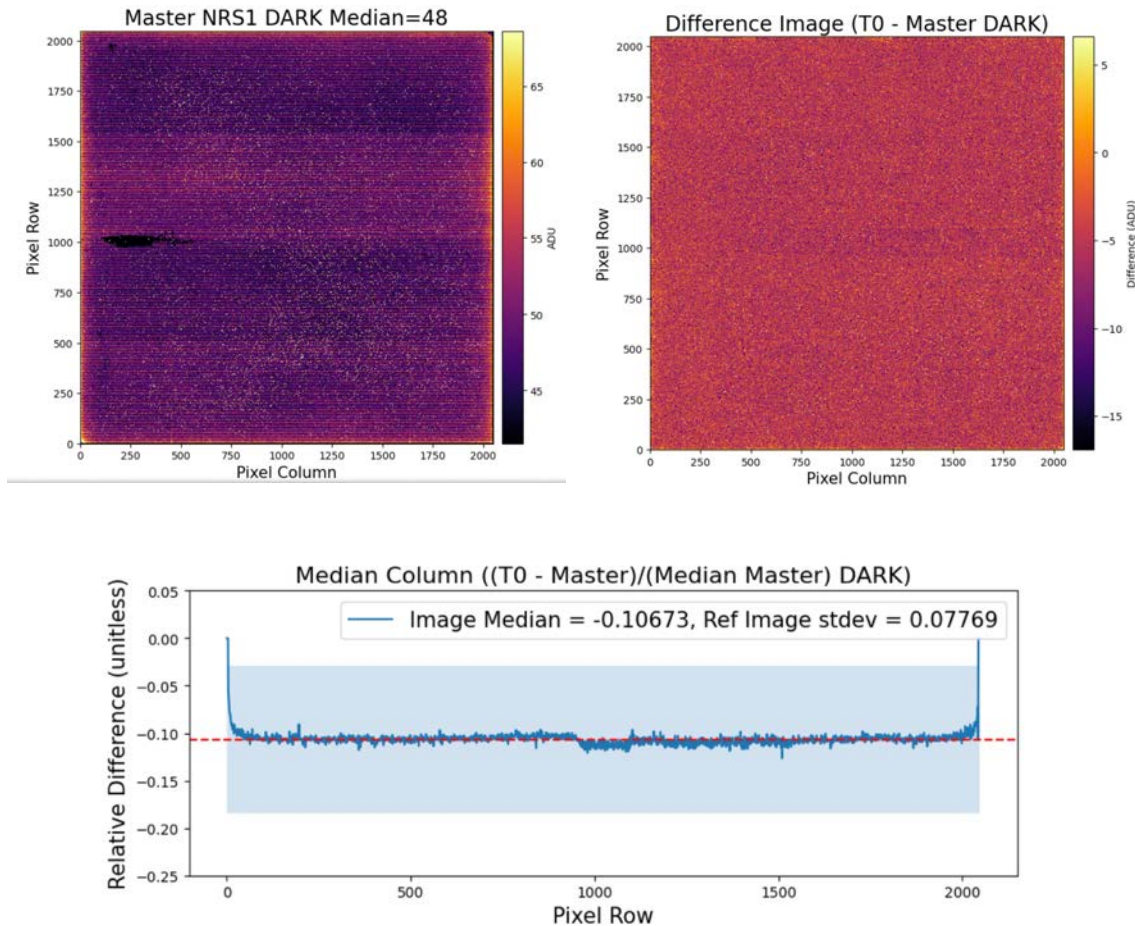


Figure 11: NRS1 IRS2 Temporal (T0 Period) Dark Image Comparison with the Master Dark. Dark level differences are fairly flat (1/f is mitigated with this mode.) T0 may have residual bias or glow at the edges from the cosmic ray floor (which is very hard to remove). The Dark levels for the different time periods can vary, but differences are less than 1 stdev of the Master dark.

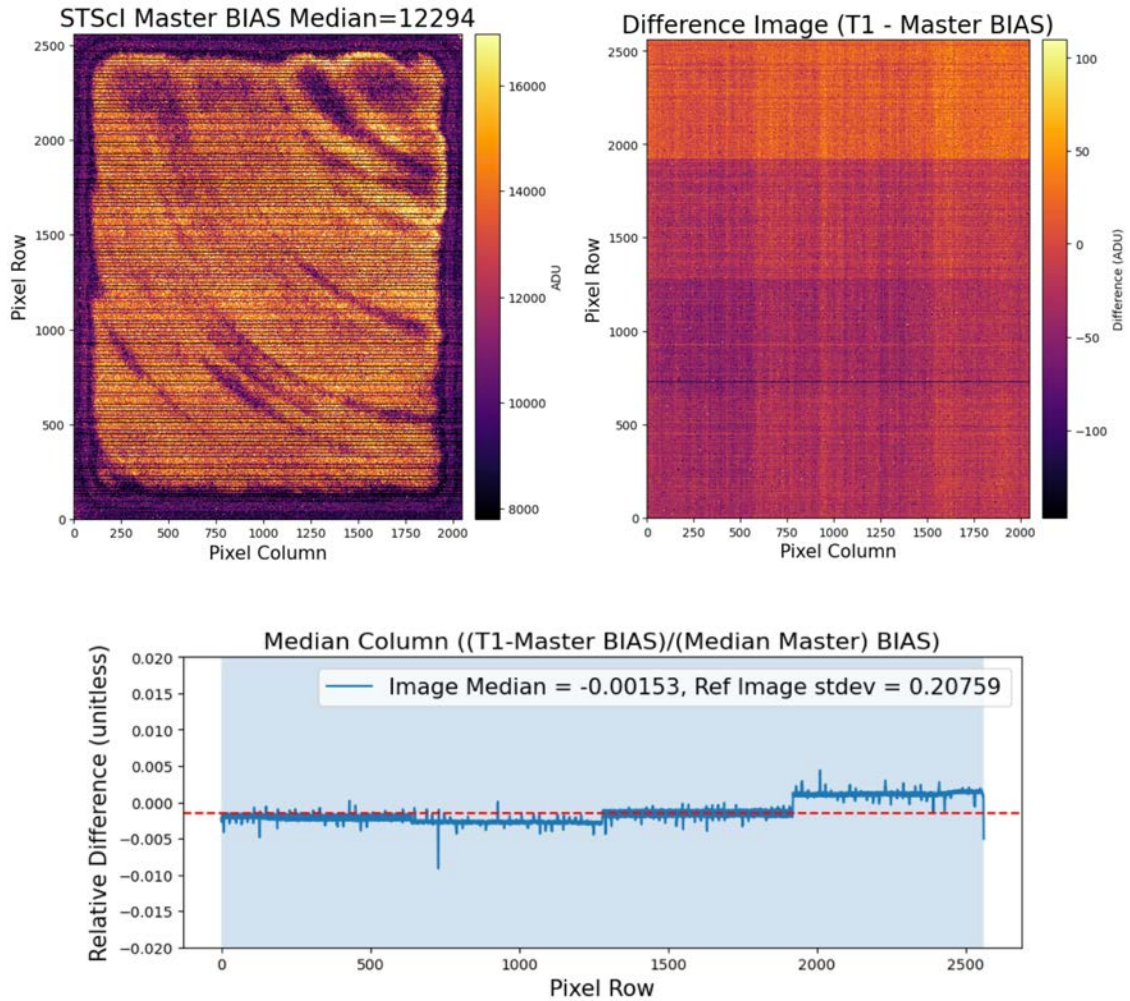


Figure 12: NRS2 IRS2 Temporal (T1 Period) Bias Image Comparison with the Master Bias. Temporal bias from period T0 compared to the master bias. Spikes seen in the plot are caused by the reference pixels. Amp level differences are about $\pm 0.3\%$, much less than one stdev of the Master NRS2 bias.

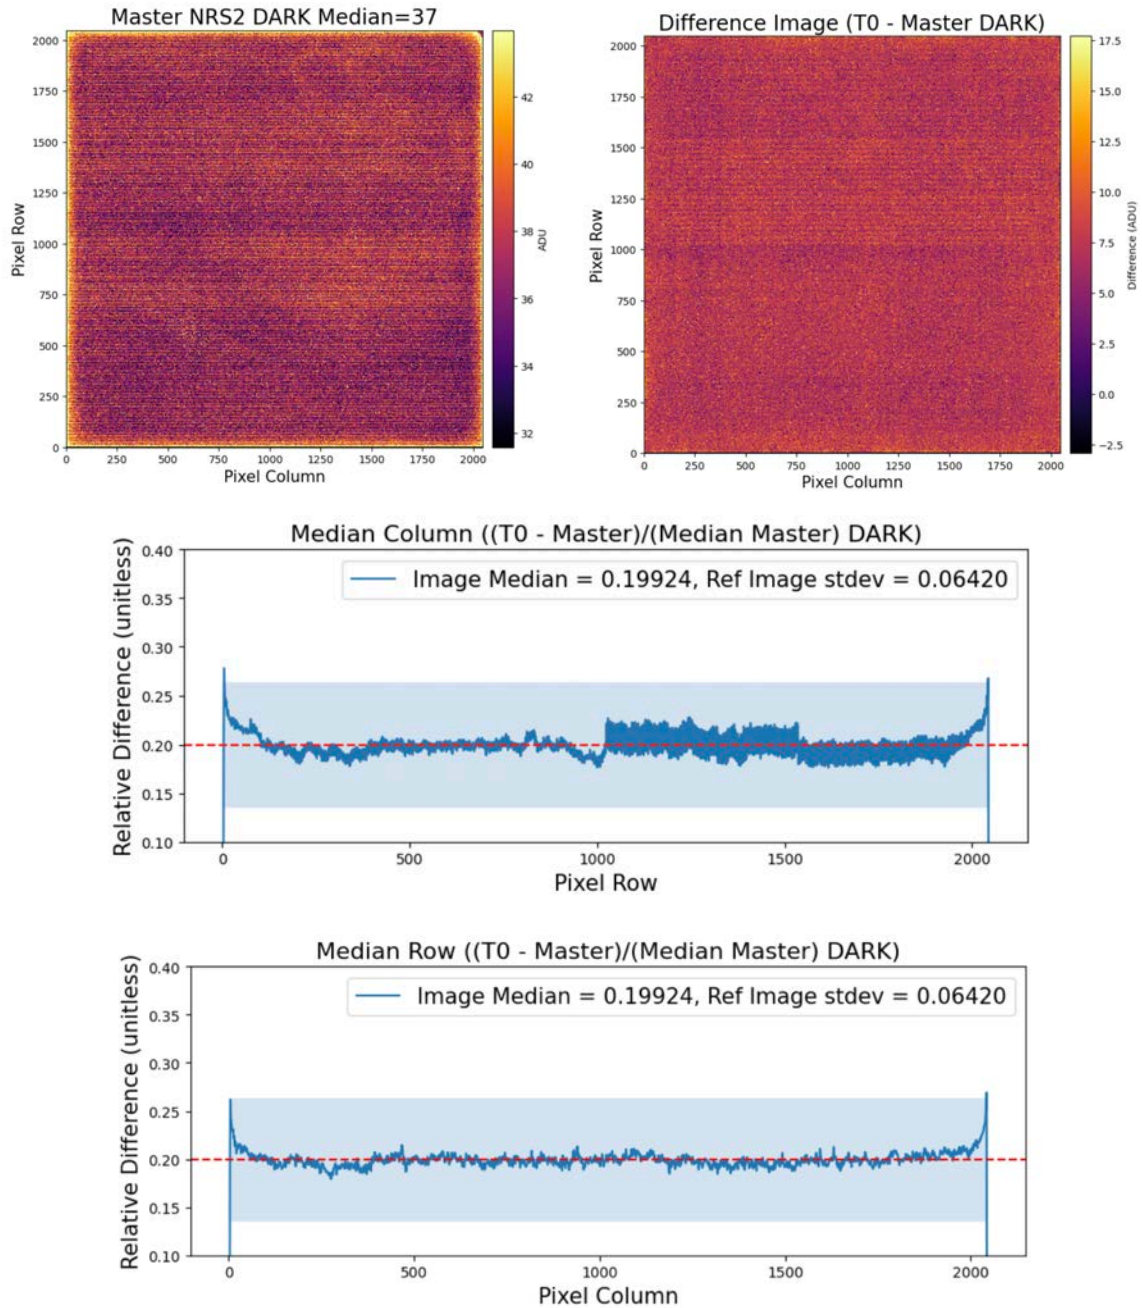


Figure 13: NRS2 IRS2 Temporal (T0 Period) Dark Image Comparison with the Master Dark. Largest differences in median column plot appear to be related to odd-even row and amp differences. Differences are typically small compared to 1 stdev of the Master dark, except perhaps at the edges. The median row plot shows that $1/f$, as expected, is mitigated in this mode.

Check with the JWST SOCCER Database at: <https://soccer.stsci.edu>
To verify that this is the current version.

7.1.2. Subarray Reference File Comparisons

For the subarrays listed in Table 5, since we typically have only one data point per Cycle per subarray, we compare the latest ST-delivered dark and bias reference files from Cycle 2 to the last ESA-delivered dark and bias files from Cycle 1 available in CRDS. Across many of these comparisons, we find no significant global differences in the darks or biases that we cannot explain. Some differences do appear, primarily due to differences in processing between the ESA and ST pipelines, such as the application of linearity correction and variations in $1/f$ noise, but these differences are minimal. Figures 14 and 15 show a representative example using the ALLSLITS subarray, which encloses all other subarrays.

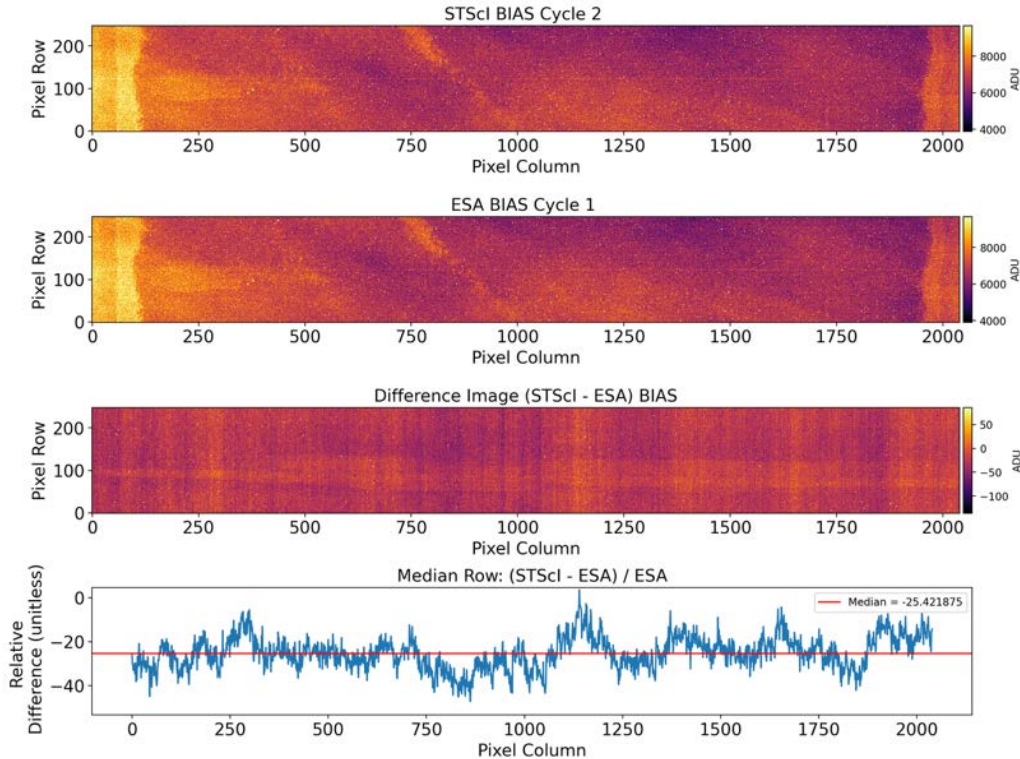


Figure 14: Comparison of Cycle 1 and Cycle 2 BIAS files for the ALLSLITS subarray on NRS1. The top panel shows the Cycle 2 BIAS processed with the ST code, while the second panel displays the Cycle 1 ESA-processed BIAS. The third panel presents the difference image between the two, and the bottom panel shows the median row of the relative differences, with the median value marked by a red line. Although some differences appear between the two BIAS files, they are small compared to the typical bias level. Some fluctuations can be attributed to differences in the $1/f$ noise, which is not corrected here.

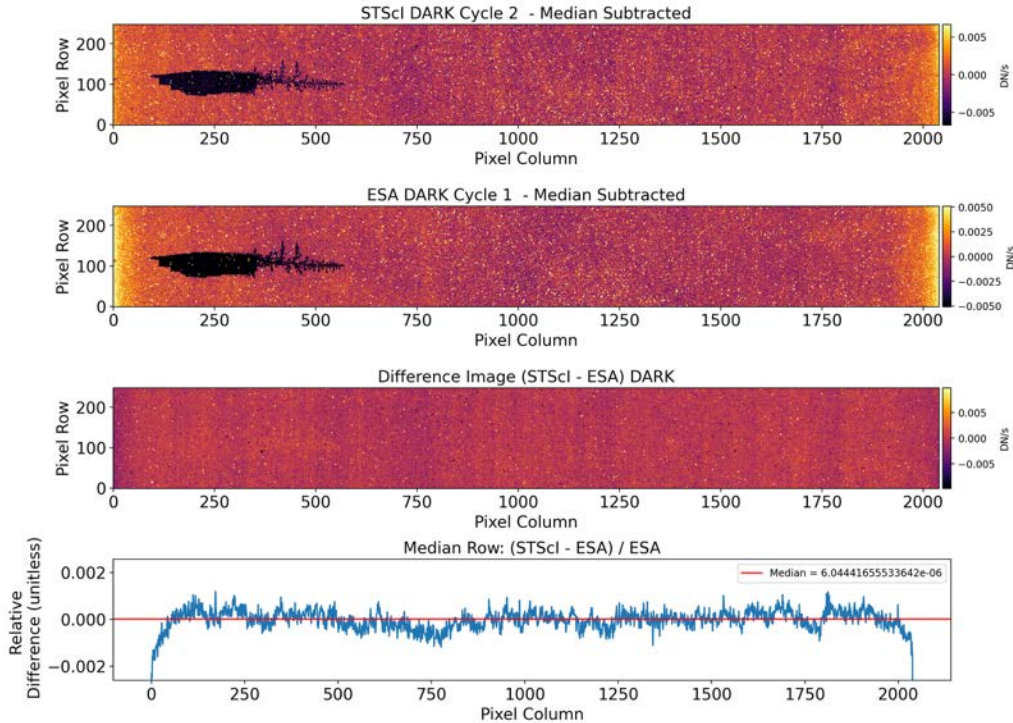


Figure 15: Comparison of Cycle 1 and Cycle 2 DARK files for the ALLSLITS subarray on NRS1. The top panel shows the Cycle 2 DARK processed with the ST pipeline, while the second panel displays the Cycle 1 DARK processed with the ESA pipeline. The third panel presents the difference image between the two, and the bottom panel shows the median row of the relative differences, with the median value marked by a red line. While some differences are apparent between the two DARK files, they are generally small, with the most significant variations near the edges. In the ESA product there is larger glow near the edges likely due to residual picture frame in the input files that may not have gotten filtered out. Some differences can be attributed to $1/f$ noise differences, reference pixel treatment, and the linearity correction applied in the ST pipeline, but not in the ESA pipeline.

8. Science Data Comparisons

In this section, we present the results from comparisons made between reducing science data processed using the master dark, bias, and mask reference files versus the period dark, bias, and mask reference files. For the subarrays we compare science reductions using the ST and ESA-created darks, bias, and mask reference files. The goal of the comparisons was to assess the impact on the science products arising from small differences in the reference files over time.

8.1. Full-frame Science Data Comparisons

8.1.1. Traditional Mode (NRS, NRSRAPID)

For the TRAD mode comparison, we selected an IFU program. The science data are from PID 2747 Observation 7 (observed June 29, 2023) which uses the PRISM/CLEAR configuration (total exposure time ~ 1084 sec). Prism spectra from the IFU fall entirely on NRS1. These are data from a faint extended source, an Oort Cloud comet (Bolin, B. 2023), and the signal in all slices appears to be dominated by the sky background. Part of the light from the source falls into the always open fixed slits. NRS2 comparisons below (Figure 17) show only differences in the failed open shutters, as well as in the S200B1 slit.

For each detector, in Figures 16 and 17 below, we show the (1) “scirate” science data (processed using the master dark), (2) a difference image between science data processed with the master reference files, and that from period T1. (3) a median collapsed column plot (spanning the 4 amps) from the difference image of the MOS science data processed both ways. Deviations are tiny compared to the standard deviation of the science data measures in an empty area of the science frame, and much less than the visible $1/f$ noise seen in the difference image.

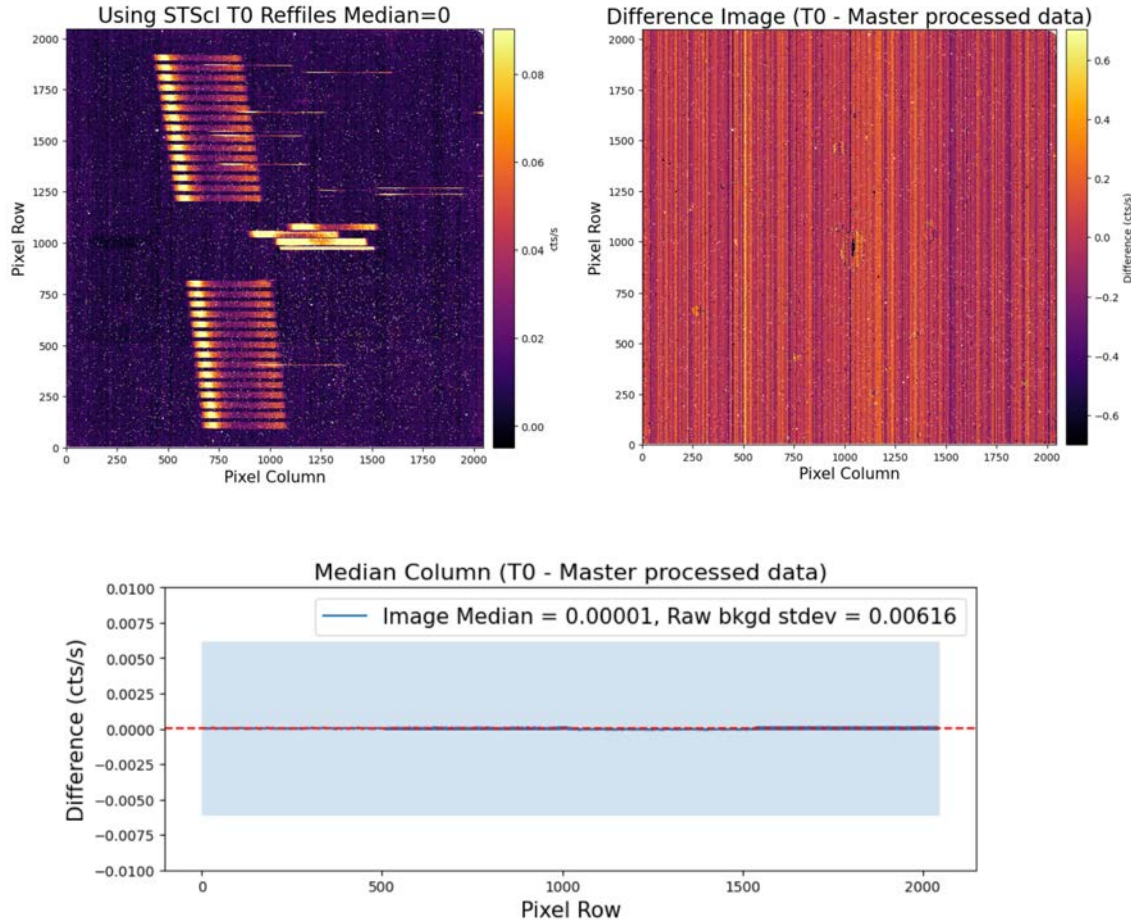


Figure 16: NRS1 Science Data Comparison for Traditional Mode. Comparison of IFU prism science data from program 2747 processed through dark subtraction using the T0 period bias and dark reference files compared to the same products processed using the “master” reference files from the entire period since the end of commissioning until the end of September 2024. Slight differences between data reduced with Master and T0 bias and darks. The $1/f$ features are noticeable across the detector in the difference image. Also seen are differences due to snowball masking, but the differences in the median column are small relative to one stdev (measured in a background area of the science data).

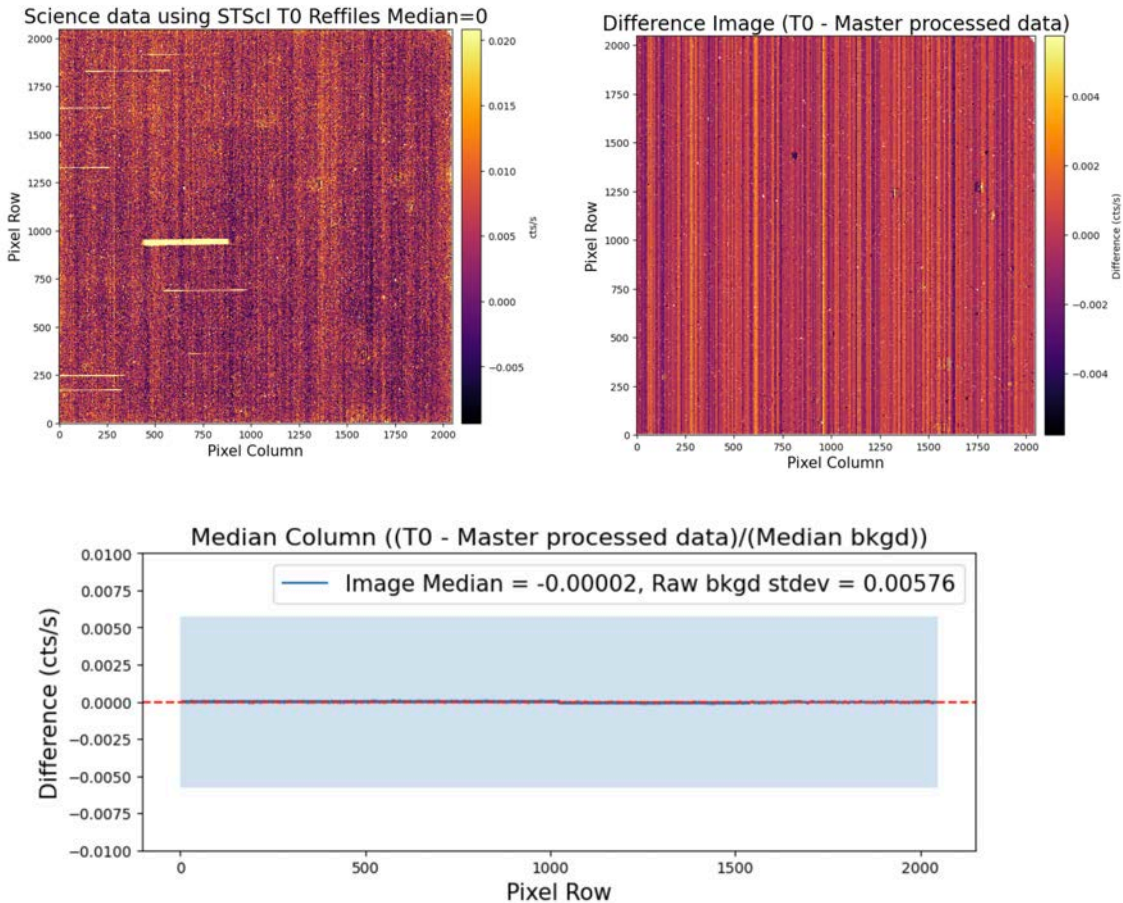


Figure 17. NRS2 Science Data Comparison for Traditional Mode. Comparison of IFU prism science data from program 2747 processed through dark subtraction using the T0 period bias and dark reference files compared to the same products processed using the “master” reference files from the entire period since the end of commissioning until the end of September 2024. Again, the differences are much smaller than one standard deviation of the value measured in a “background” area of the science data devoid of spectra (marked by the blue shaded area in the plot).

The NRS1 science data were further processed through the Stage 2 pipeline to produce an X1D extracted spectrum of the observed source. The extracted spectrum, processed with the reference files from the same time period T0, is plotted in Figure 18, along with the extracted spectrum from the long-term “master” dark. The X1D spectrum is from an extended source but the spectrum is apparently dominated by the sky spectrum. There are several points in the spectrum at which the spectral brightness drops sharply. This is likely due to unflagged and uncorrected bad pixels in the science data (e.g. unflagged negative flat-field pixels). The JWST pipeline includes a number of improvements over the ESA pipeline, and we continue to make improvements in the correction of bad pixels since this analysis.

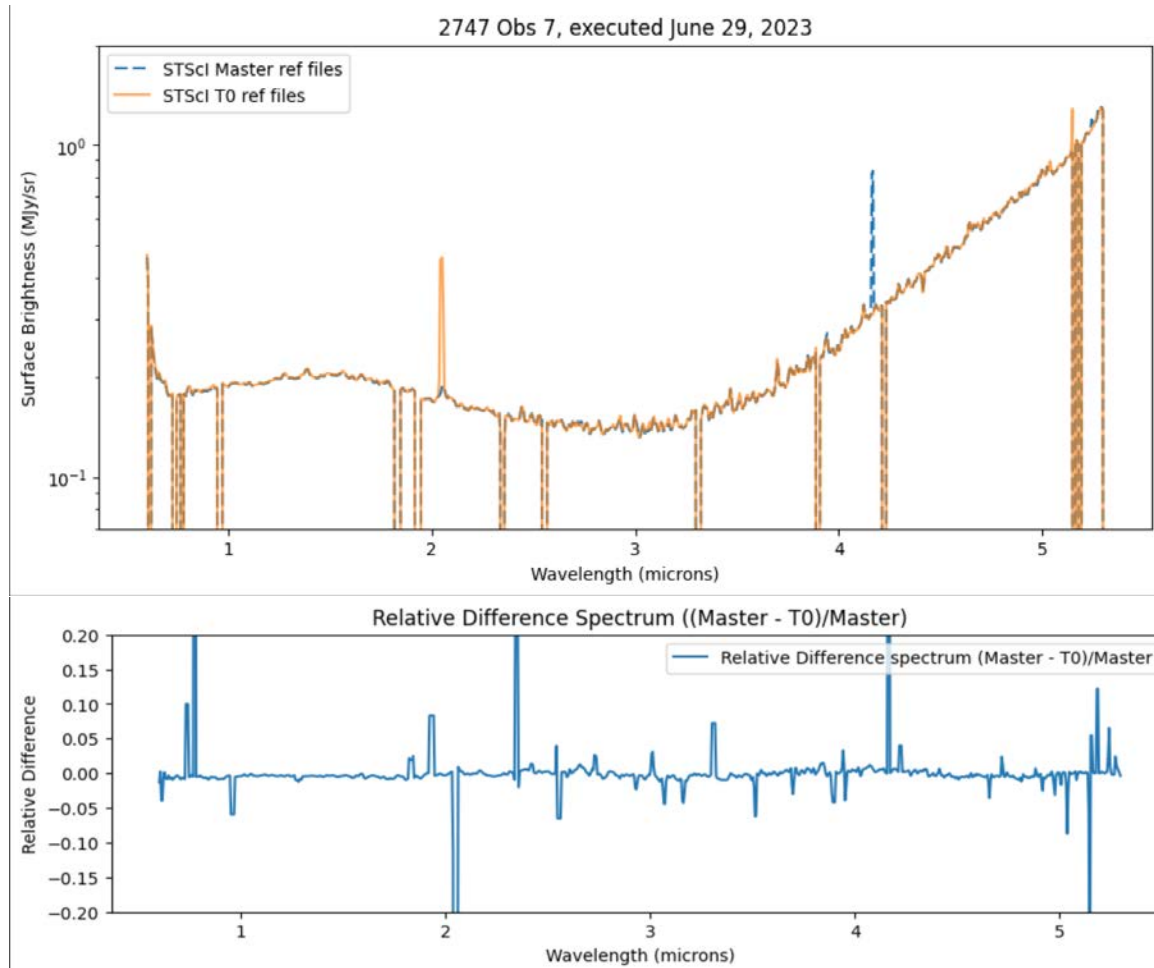


Figure 18: (Top) Extracted X1D plot of science data processed with the Master and the Temporal (T0 Period) TRAD reference files. (Bottom) Relative Difference Spectrum

8.1.2. IRS2 Mode (NRSIRS2, NRSIRS2RAPID)

To check the new darks and bias reference files for IRS2 mode, we selected science data from PID 1747 for both detectors. Figures below show the processing of MOS science data with the T2 period dark and bias and comparisons with the master dark and bias for **IRS2 mode**. Stage 1 products (“rateints” files) are compared, followed by a comparison of stage 2-processed X1D spectra.

PID 1747 is a MOS program that uses the PRISM/CLEAR combination (exposure time 817 sec). It’s a study of the IGM seen in several bright galaxies from the BORG field (Roberts-Borsani 2025). In the stage1 product comparison, we show:

For each detector, in Figures 19 and 20 below, we show the (1) the science “scirate” data (processed using the master dark), (2) A difference image between science data processed with the master reference files, and that from one of the temporal periods. (3) A median collapsed column (spanning the 4 amps) from the difference image of the MOS science data processed both ways. As for the traditional mode comparison, deviations are tiny compared to the standard deviation of the science data measures in an empty area of the science frame.

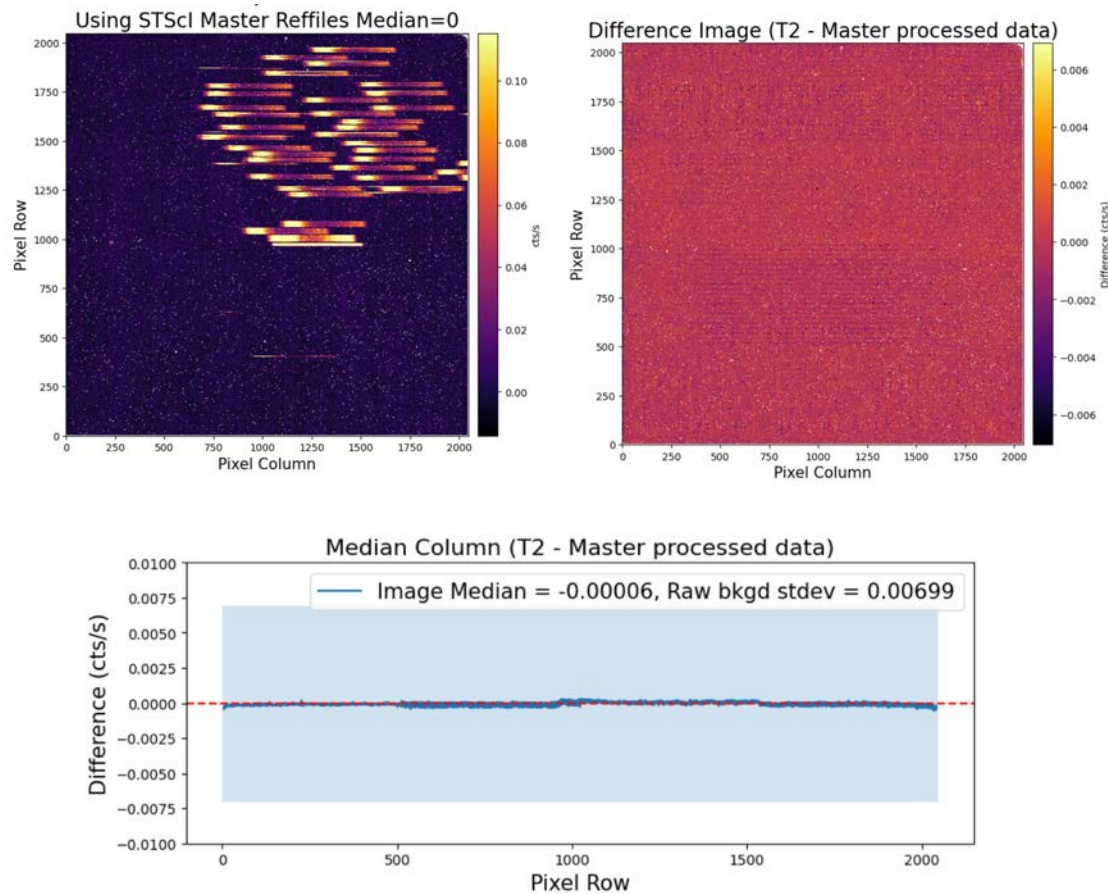


Figure 19: NRS1 Science Data Comparison for IRS2 Mode. Comparison of MOS prism science data from program 1747, Obs 16 (executed Nov 8, 2023) processed through dark subtraction using the T2 period bias and dark reference files compared to the same products processed using the “master” reference files from the entire period since the end of commissioning until the end of September 2024. There are slight differences between data reduced with Master and T2 bias and darks. The differences in the median column are small relative to one stdev (measured in a background area of the science data processed with the Master reference file). One can clearly see the improvement in the reduction of 1/f noise in this mode.

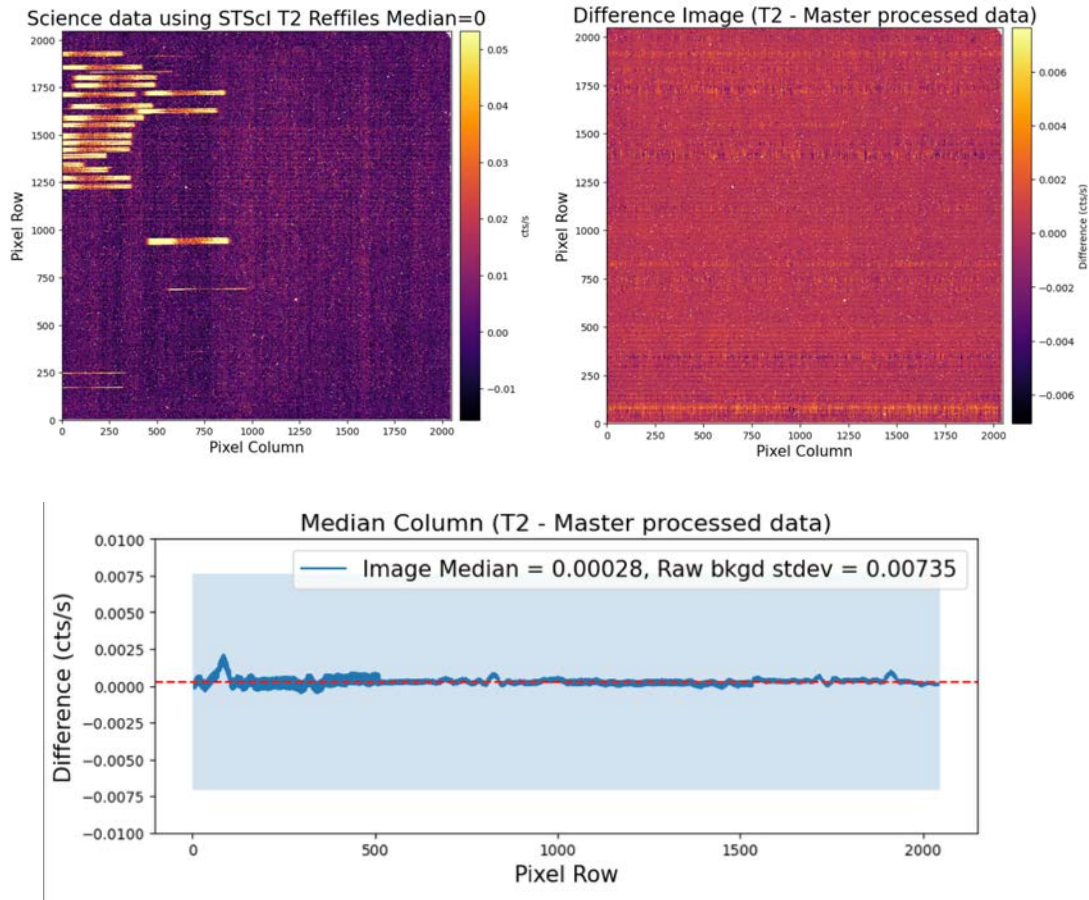


Figure 20: NRS2 Science Data Comparison for IRS2 Mode. Comparison of MOS prism science data from program 1747 processed through dark subtraction using the T2 period bias and dark reference files compared to the same products processed using the “master” reference files from the entire period since the end of commissioning until the end of September 2024. Again, the differences seen are much smaller than the standard deviation of a background area of the image processed using the “Master” reference files (marked by the blue shaded area in the plot). The horizontal banding seen in the difference image and appearing as small peaks in the median column plot is likely due to the different average reference pixel values in the two reference files.

In the Stage 2 comparisons of X1D spectra extracted from these science data, Figures 21 and 22, we show a few sample spectra from PID 1747. The top figure in each pair is the reference (in this case, the “T0” period reference files) processed extracted image for a single slit, followed by the “T3” processed extracted image, and the difference of the two. The plot below is the X1D extracted spectrum of a science source processed with each of the sets of reference files. The differences in the 1D extracted spectra are typically much lower than 10% (shown in the bottom plot of residual differences). In the figures below are examples of a point source spectrum (in Flux units), and an extended source spectrum (in surface brightness units). Large relative differences are occasionally seen. These appear to be due to pixel-level differences in regions where the noise is high, or the flux is low. Large differences may also be due to individual hot pixel flagging differences in the masks from the two different time periods.

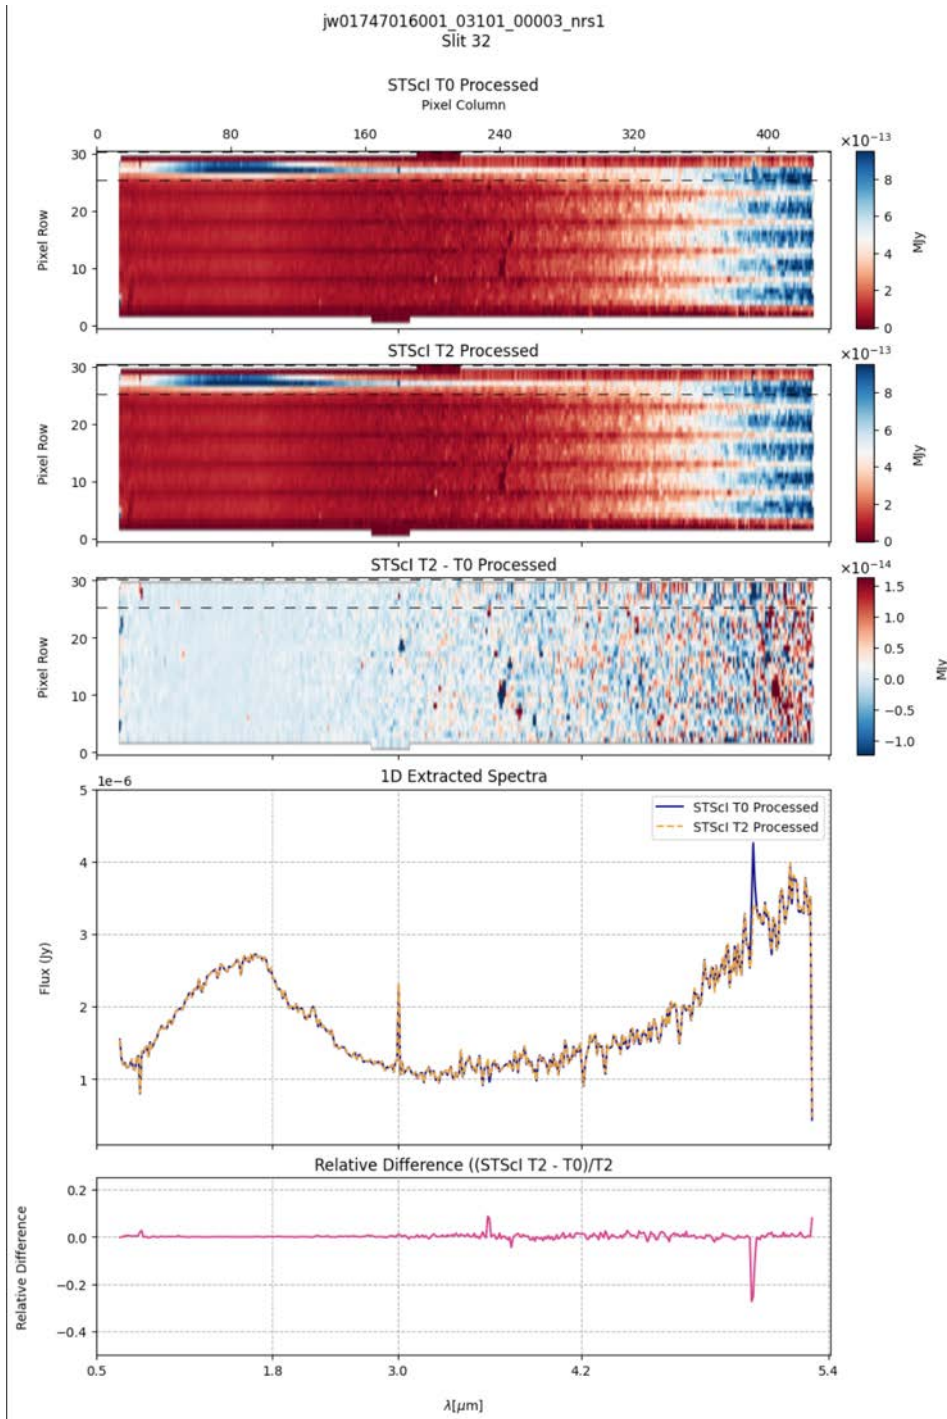


Figure 21: Extracted 2D dataset and X1D plots of science data processed with the IRS2 T0 and the T2 reference files.

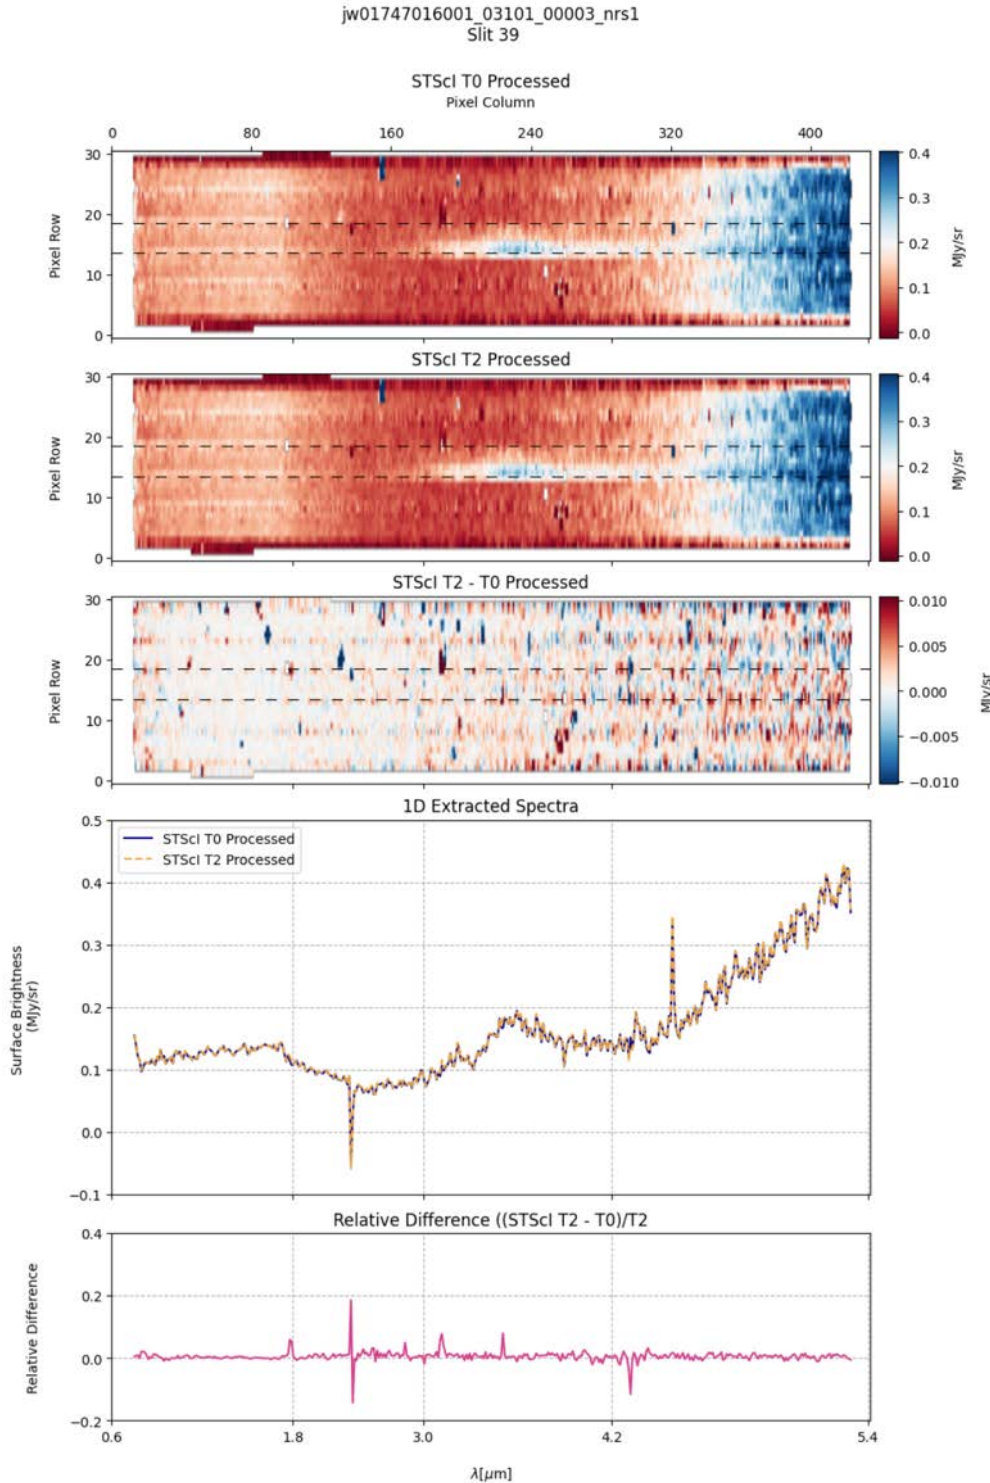


Figure 22: Extracted 2d dataset and 1D plots of science data processed with the T0 and the T2 IRS2 reference files.

8.1.3. Subarray Science Data Comparisons

For these comparisons, we selected a fixed slit observation from program PID 1292, Observation 1, which used the F070LP/G140H with the ALLSLITS subarray. Figures 23 and 24 display data from the S200A1 and S1600A1 slits, which contain a point source and an extended background

Check with the JWST SOCCER Database at: <https://soccer.stsci.edu>

To verify that this is the current version.

source, respectively. Both slits fall on detector NRS1. In each case, we compare the differences in spectral products processed with the respective dark and bias reference files. We observe minor differences, with the ST-processed data consistently showing slightly higher flux compared to the ESA-processed data. This increase is consistent with the expected effect of the newly implemented linearity correction in the ST pipeline.

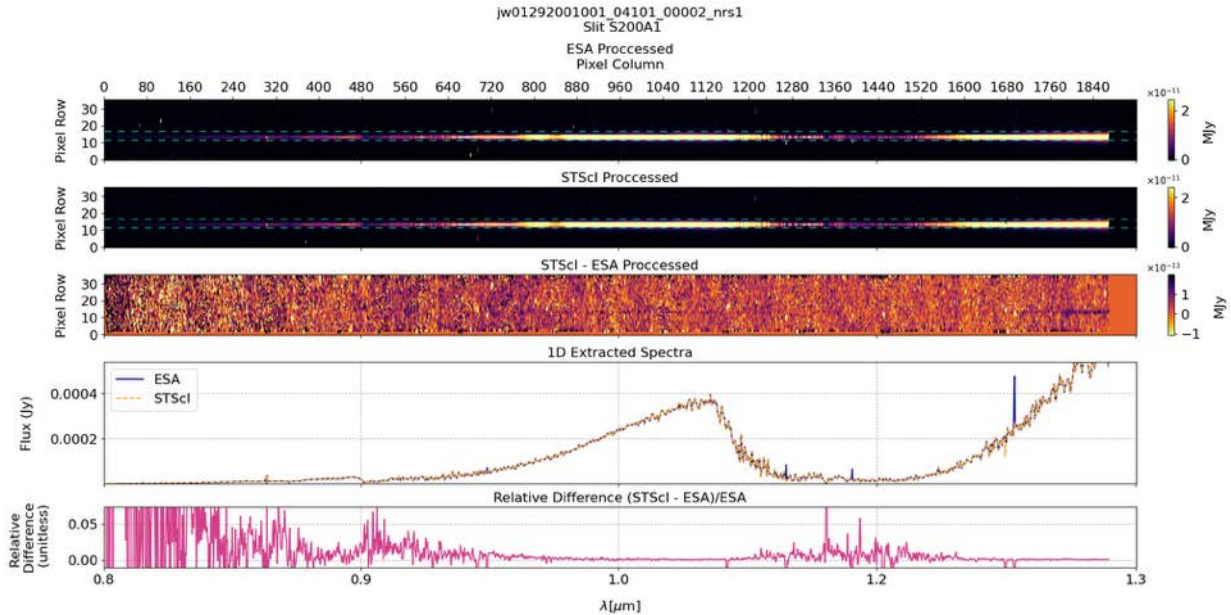


Figure 23: 2D and 1D spectral products from the S200A1 slit in the ALLSLITS subarray. The source is a point. The top two panels show the 2D spectra processed with the ST and ESA files, respectively. The third panel displays the difference image between the two. The fourth panel presents the 1D extracted spectra for both versions, and the bottom panel shows the relative differences between them.

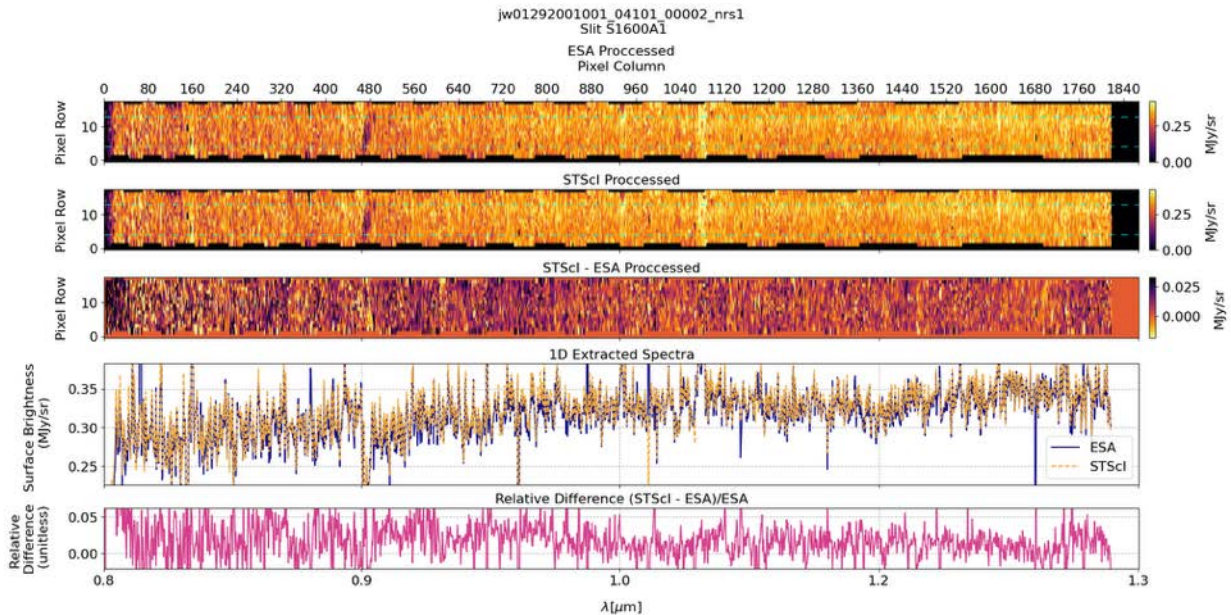


Figure 24: 2D and 1D spectral products from the S1600A1 slit in the ALLSLITS subarray. The source is background (extended). The top two panels show the 2D spectra processed with the ST and ESA files, respectively. The third panel displays the difference image between the two. The fourth panel presents the 1D extracted spectra for both versions, and the bottom panel shows the relative differences between them.

Check with the JWST SOCCER Database at: <https://soccer.stsci.edu>
To verify that this is the current version.

9. Implications for Cycle 4 and Beyond Bias/Dark Calibration Monitoring Programs

The numbers of requested, observed, and usable full-frame input darks is summarized in Table 5. These were used to update the strategy that will be used for the Full frame dark monitor program for Cycle 4 (NRS-CAL-407, PID 9273).

9.1. Observed Fraction of Full-frame Dark Files

For **IRS2 and TRAD**, the fraction of those observed to those requested generally increased through the first 3 cycles, as shown in Table 5. The dip in Cycle 3 TRAD numbers indicates slots lost due to the mode switch scheduling requirement that we put in place in May 2024, even with an increased rate of requested darks relative to previous cycles.

9.2. Usable Fraction of Observed Full-frame Dark Files

The number of observed FULL frame dark exposures are further reduced due to ASIC temperature filtering, and filtering for other artifacts: Of those successfully observed, the number of those accepted as input files for making the dark reference frames is shown in the last column of Table 5. Many files were lost due to ASIC temperature filtering in Cycles 1 and 2. The ASIC temperatures may have been out of range, or the change in temperature during the exposure was too large, which is an indicator of a mode switch. By the start of Cycle 3, we put in new restrictions on scheduling, and we started observing all darks in G140H (in May 2024). In Cycle 3, of those that survived the preliminary temperature-limit filtering, only 1 TRAD file (and 0 IRS2 files) was lost due to the delta-temperature value being too high, indicating that our strategy is helping. However, new persistence artifacts were found in the Cycle 3 IRS2 data, requiring further culling of IRS2 input frames. See Figure 25.

Generally, the numbers of artifact-free files for both TRAD and IRS2 increased over the cycles, mainly due to the mode switch scheduling change. However, persistence artifacts account for a large number of additional frames that are unusable. In preliminary analysis of more recent Cycle 3 data, this trend worsened a little. We are finding just 49 – 50% of the observed IRS2 frames are usable (compared to the 65% reported here for early Cycle 3). This demonstrates how the final numbers are completely dependent on when the darks are scheduled, and what observations precede them.

Table 5: Some Important Numbers for Future Dark Monitor CAL Programs

	APT requested	Observed (% of Requested)	T-filtered (% of Requested)	GOOD % of Observed
IRS2				
Cy 1 PID 1484	100	67 (67%)	64 (64%)	18 (27%)
Cy 2 PID 4455	50	46 (92%)	42 (84%)	19 (41%)
Cy 3 PID 6633	84 (49 as of Jan 15)	48 (97%)	42 (86%)	31 (65%)
TRAD				
Cy 1 PID 1484	45	41 (91%)	39 (87%)	10 (24%)
Cy 2 PID 4455	50	50 (100%)	48 (96%)	15 (30%)
Cy 3 PID 6633	60 (35 as of Jan 15)	27 (77%)	20 (57%)	19 (70%)

Check with the JWST SOCCER Database at: <https://soccer.stsci.edu>
To verify that this is the current version.




Figure 25: NRS1 RATEINTS file showing artifacts seen in some Cycle 3 IRS2 Full-frame darks. These nearly vertical line artifacts are caused by persistence from the zeroth order of S-flat calibration lamp observations taken in MOS long slit configuration. Persistence in the fixed slits is also obvious as wide spectral traces located about half-way between the top and bottom of the image.

9.3. Factors Affecting the Cycle 4 Full-frame Darks Monitoring Program

TRAD Mode: In Cycle 3, after implementing a scheduling restriction around mode switches, we have seen a *lack of scheduling opportunities* for TRAD darks since many more science observations use IRS2 mode. Still, in Cycle 3, after implementing our new scheduling strategy, only ~70% were useful. Several files were eliminated because of observed artifacts. Therefore, we can expect only about 70% of those observed to be useful for making TRAD reference files due to ASIC temperature filtering and filtering of other visual artifacts. About 10-11 clean files are needed to make either a TRAD or IRS2 dark reference file with acceptable SNR (We adopted 2.5 for TRAD and 5.0 for IRS2).

In cycle 4, we can hope to observe only ~30-40 TRAD mode darks due to available slots. We will request 40. Of these, we can project that ~70% will be useful, which would allow us to make at most 3 TRAD dark reference files for the cycle (with SNR ~ 2.5).

IRS2 Mode: We need about 10 clean files to attain a minimum SNR~5 in NRS2, and about 12 clean files to attain an SNR~6 (NRS2 is the limiting detector due to its lower dark count rate). We had a 63% success rate (after correcting for those with artifacts). More recent examination of the entire set of Cycle 3 data yielded fewer usable files (~50%). This value is unfortunately dependent on which observations immediately precede the darks, and outside of campaign mode, we have little control over how these darks are scheduled.

Check with the JWST SOCCER Database at: <https://soccer.stsci.edu>
To verify that this is the current version.

Therefore, in cycle 4, and based on the more optimistic success rates we determined from the examination of cycles 1 to 3 data, with the goal to make quarterly IRS2 dark reference files (with SNR~6, which requires about 12 input files) we planned about 76 files. (By comparison, we requested 84 IRS2 files in Cycle 3.)

9.4. Recommendations for Cycle 4 Cal Program (PID 9273):

- Plan for quarterly IRS2 darks: 4 epochs of 19 files each (76 files total). Each epoch of IRS2 darks takes 19.26 hours.
- Plan for enough TRAD darks to make 4 TRAD dark reference files in Cycle 4: 4 epochs of 10 darks each (40 files total). Given the reduced rate of scheduling opportunities to observe in this mode, and the fact that they could be affected by persistence despite our attempts to reduce this occurrence, we expect to be able to make just 1-2 TRAD darks over the cycle.
- The draft APT program has 72.1 hours of science time and a charged time of 92 hours. Cycle 3 program was 83.9 hours of science time and 107.7 hours of charged time, so this is a slight savings.

10. Conclusions

The period T0 reference files in Table 4 have been delivered to CRDS, superseding the earlier Cycle 1 versions made with the ESA CAP code. Those older files, which included PRISM input darks (which are now excluded in the ST version), remain in CRDS but will no longer be automatically used. During reprocessing, the pipeline will use the newly delivered files, which incorporate the many improvements outlined in this report. The other period reference files labeled T1, T2, and T3 for each mode and detector were delivered to CRDS in March 2024 and are being used in the pipeline reduction of science data from that period to the present. Additionally, while preparing this report, we recently delivered new bias and dark reference files for the entirety of Cycle 3 to CRDS.

The reference file comparisons, and science data comparisons reduced with the different sets of reference files demonstrate that there is no significant change in the dark current over time. The relative differences in extracted science data reduced with different reference files appear to be entirely due to differences in the detection and correction of hot or other bad pixels.

There is a continuing need to monitor for any evolution in the dark current and the spatial morphology of the darks. Most importantly, though, there is a need to deliver hot pixel masks regularly, to more closely match and flag the time-dependent hot pixels in the science data calibration pipeline. Hence, the Cycle 4 and future monitoring programs should request enough IRS2 darks to ideally create a set of independent reference files on a quarterly basis. Recent improvements in the scheduling of darks should enable this. However, most NIRSpec science observations are now in IRS2 mode. Due to that and the scheduling restriction to find parallel slots which do not change the mode, we will reduce requests for TRAD darks in Cycle 4 and beyond (to ~ 40 files). As for the subarray observations, Cycle 4 has been planned to meet the required minimum SNR for the final dark reference files for the requested subarrays. Since all subarray darks are taken in TRAD mode and the number of available scheduling slots continues to decrease, we will maintain rather than reduce the current requested rate of subarray dark observations, with the hope that these observations can be scheduled.

References

- Birkmann, S. 2022, NIRSpec Calibration Data package Delivery Note - Bias, Dark, and Read Noise Reference Files, ESA-JWST-SCI-NRS-TN-2022-005
- Birkmann, S. et al. 2018, Generating NIRSpec Reference Files for Dark Current, Master Bias, and Hot Pixel Masks, ESA-JWST-SCI-NRS-NTN-2017-008
- Birkmann, S. 2014, Description of the NIRSpec bias and dark reference files, ESA-JWST-TN-20072 / NIRSpec Technical Note NTN-2013-003
- Böker, T. et al. 2025, Trending of hot pixels on the NIRSpec detector, JWST-STScI-009188
- Böker, T. et al. 2022, PASP, Vol. 135, p. 1538 In-orbit Performance of the Near-infrared Spectrograph NIRSpec on the James Webb Space Telescope
- Bolin, B. et al. 2023, in Asteroids, Comets, Meteors Conference 2023 (LPI Contrib. No. 2851)
- Rauscher, B. et al. 2007, PASP, Vol. 119, p. 768 Detectors for the James Webb Space Telescope Near-Infrared Spectrograph. I. Readout Mode, Noise Model, and Calibration Considerations
- Roberts-Borsani, G. et al. 2025, Ap. J. 983:18 (16pp)
- Schlawin, E. et al. 2020, *AJ*, Vol. 160, p. 231 JWST Noise Floor. I. Random Error Sources in JWST NIRCам Time Series
- Zeidler, P. 2022, NIRSpec Bad Pixel Mask Reference Files, ESA-JWST-SCI-NRS-TN-2022-004 CAP-011
- [NIRSpec Detector Performance JDOx article](#)