JWST TECHNICAL REPORT

Title: Trending of hot pixels on the NIRSpec detectors	Doc #: JWST-STScI-009188, SM-12 Date: 2 October 2025 Rev: -
Authors: T.Böker, K.Glidic, C.Hayes, D.Karakla, T.Keyes, N.Kumari Phone: 410-338 2467	Release Date: 17 November 2025

Abstract

We report the results of the first three years of in-orbit trending of the 'hot pixel' population of the NIRSpec detectors, i.e. pixels that display a higher than normal dark current. The analysis shows that the number of hot pixels grows slowly but steadily at a rate of ~100 new hot pixels per month (depending on readout mode and detector). This translates to an increase of no more than 0.03% of total pixels per year. We discuss the implications of this small but non-negligible growth of the hot pixel population for science data processing and – especially - the onboard target acquisition procedures. We conclude that the onboard files used to mask hot pixels during target acquisition need to be updated regularly, i.e. on a roughly yearly cadence.

1. Introduction

It has been known for some time that NIR detectors in general, and the NIRSpec ones in particular, develop a growing population of 'hot' pixels over time, i.e. pixels that exhibit a higher-than-normal dark current level. The exact mechanism for a normal pixel to become 'hot' is unknown, but the most plausible cause is radiation damage, which is more likely to occur during in-orbit operations.

The correct identification and masking of hot pixels is critical for the quality of the science products generated by the JWST science calibration pipeline, as hot pixels must be flagged (in the bad pixel mask reference file used during the *dq_init* pipeline step) so that during the data processing flow, their values can be ignored or replaced with appropriate estimates. In addition, and more importantly for seamless JWST science operations, hot (or otherwise bad) pixels must be recognized by the onboard scripts that analyze the target acquisition (TA) exposures, as hot pixels can be wrongly identified as acquisition targets. For this reason, so-called 'TA flat field' files are stored in onboard RAM. These files not only provide the relative pixel response variations but also flag any known hot pixels such that their values can be ignored, which minimizes the likelihood of TA failures caused by the centroiding algorithm ‘finding’ a hot pixel.

While frequent updates to the bad pixel masks are obviously important for high-quality pipeline products, updating the onboard TA flat field files requires a resource intensive, real-time activity during contact with the JWST spacecraft which should only be done when necessary. A good characterization of the hot pixel population and its growth over time are needed so that informed decisions can be made about the required frequency of TA flat field updates.

This report describes the results of the hot pixel monitoring over the first three years of JWST science operations and uses them to make recommendations for the maintenance of the onboard TA flats. The pre-launch procedure to create and upload these files is outlined in Proffitt & King (2021), and a separate report with an updated procedure for the science operations phase is currently being prepared (Keyes et al. 2025).

2. Hot pixel trending since launch

The exact definition of what dark current level requires the masking of a pixel as ‘hot’ depends on what the pixel is used for. The median dark current level for the two NIRSpec detectors is less than 0.01 e-/s in either readout mode (TRAD or IRS2, Birkmann et al. 2022). When processing NIRSpec data through the pipeline, a dark current level of $1e^{-}/s$ (i.e. more than 100 times the median rate) is considered ‘hot’, based on its non-negligible contribution to the typical flux in ‘faint’ science observations, which are often long and contain low signal levels.

On the other hand, TA exposures are short and contain comparatively bright sources, so the definition of ‘hot’ can be more lenient. So far, a dark current level of 3 e^{-}/s has been used to define a pixel as unusable for the analysis of TA exposures (Proffitt & King 2021). It is therefore useful to monitor the hot pixel population at different dark current levels.

Figures 1 and 2 show the in-orbit (i.e. post-launch) evolution of the hot pixel population as a function of time, for the two detector readout modes (IRS2 and TRAD) and for various dark current thresholds. We collect this data from full-frame dark reference files available in the Calibration Reference Data System (CRDS). Specifically, we look at the DRK extensions, which contain the average dark current rate for each pixel, and counted all pixels above a given threshold. As a consistency check, we also compare the number of pixels above $1e^{-}/s$ derived in this way to the number of pixels flagged as 'HOT' in the DQ extension of the bad pixel mask reference file – as expected, the numbers match.

The pronounced ‘jump’ in the hot pixel numbers in mid-2023 is caused by a change in the way these reference files are produced: during early JWST operations, we continued to use code created by the ESA NIRSpec team to create dark reference files which measured dark current in e/s. Since late in Cycle 1, the reference files are being produced with Python code created by STScI and are in units of DN/s to make them more directly compatible with the JWST pipeline. We note that because the detector gain in full frame mode is very close to 1, the same thresholds values (1 DN/s vs 1 e^-/s, etc.) were maintained, and the unit change has only a very small effect on the total number of hot pixels. More importantly, the code transition also introduced a change in how hot pixel neighbors are being treated: the elevated flux in a hot pixel often 'spills over' into the four neighboring pixels, forming a cross-like pattern. While the ESA code flagged only the central pixel, the STScI code also flags the adjacent pixels if their count rate exceeds the given threshold (even though they are not themselves producing the signal). This likely explains the observed 'jump' in the hot pixel counts for the STScI products.

As can be seen, the trends derived from the suite of dark current reference files indicate a nearly constant growth rate over time, regardless of which code version was used to produce them. Every month, about 100 additional pixels (using a conservative upper limit) on each NIRSpec detector become 'hot' with a dark current above 1 DN/s, and roughly half of those have a dark current higher than 3DN/s. These numbers translate to an increase of about 0.03% of total pixels per year (0.02% for the higher threshold of 3DN/s). Table 1 summarizes the results of the linear fits for the various thresholds and readout modes.

	1 DN/s	**3 DN/s**	**10 DN/s**	**100 DN/s**
NRS1 (TRAD)	64	48	25	5
NRS2 (TRAD)	79	64	33	3
NRS1 (IRS2)	65	35	13	5
NRS2 (IRS2)	86	44	13	4

Table 1 Growth rates of the hot pixel population (in pix/months) of the two NIRSpec detectors for various dark current thresholds and the two readout modes, based on linear fits to the data shown in Figs. 1 and 2

3. Consequences for Science Data Processing

Failing to update the DARK reference files and the corresponding bad pixel masks frequently results in a higher number of uncorrected hot pixel outliers to remain in science spectra. These can be tedious to identify and clean, can easily be confused with emission lines, and – once manually flagged - require reprocessing through the pipeline. Another potential consequence of unflagged hot pixels is that background subtraction (if applied) can introduce large negative outliers in the final spectra. Updating the DARK reference file alongside the bad pixel mask is also important because hot pixels can cause 'overflow' signal into their immediate neighbors due to Inter-pixel Capacitance.

It is therefore important to update the dark current and bad pixel mask reference files as often as possible. As discussed in Karakla et al. (2025), the planned cadence for these updates is limited by the availability of scheduling slots for (parallel) dark current calibration exposures. Based on the analysis of data acquired in JWST cycles 1 to 3, the frequency and quality of in-orbit calibration darks is sufficient to allow the creation of new DARK reference files every 3-4 months.

Based on the blue curves in Figure 1 and considering that the USEAFTER date of a DARK reference file falls in the middle of the time period over which the input darks were collected, a quarterly update cadence would imply up to 1.5 months worth of unmasked hot pixels. In other words, up to ~150 unmasked hot pixels (with dark rates above 1DN/s) might appear on each NIRSpec detector for a 'worst case' science exposure taken just before the DARK update. Depending on the observing mode and detector area used, this may result in a non-negligible number of outliers in NIRSpec spectra, although dithering should, in principle, allow to remove most of these. More frequent updates of the DARK reference files are not realistic, and a few outliers in NIRSpec spectra are therefore to be expected.

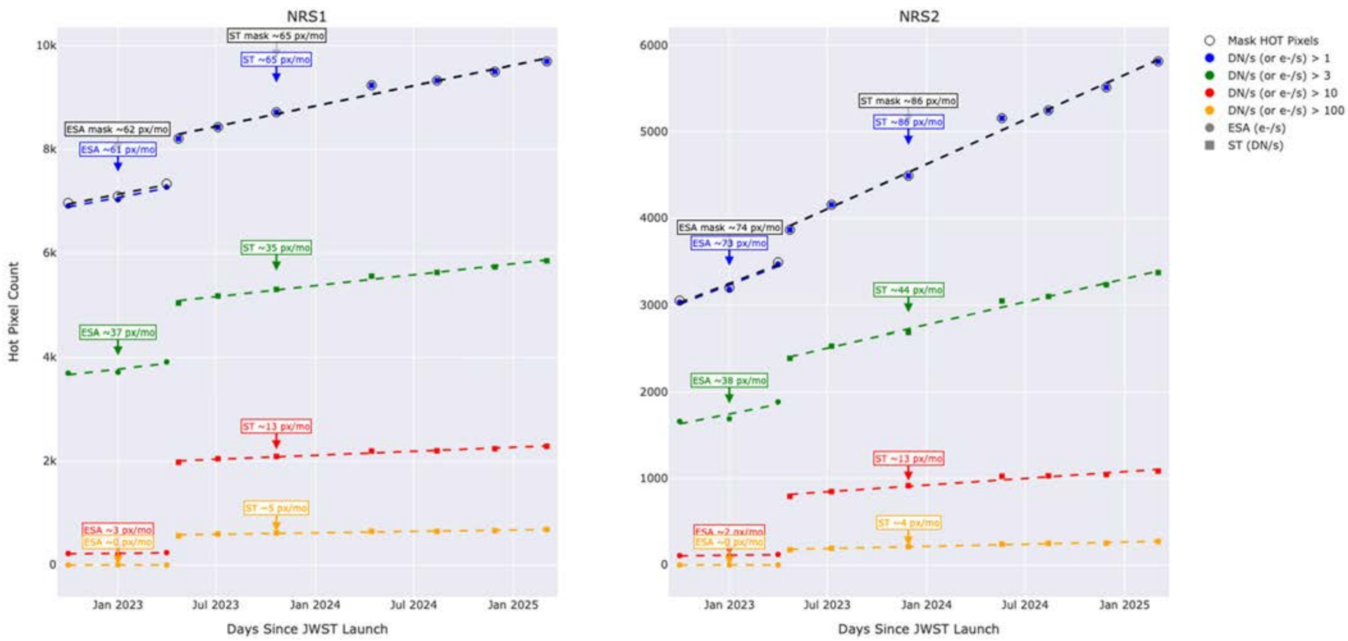

Figure 1: Results of the hot pixel trending since the start of JWST science operations (July 1 2022) for the NRS1 (left) and NRS2 (right) detector in IRS2 mode, tracked for dark current rates of 1, 3, 10, and 100 DN/s. At all levels, and for both NIRSpec detectors, the hot pixel population growth is well described by a linear trend, i.e. a constant rate of ~100 new hot pixels per month at the pipeline threshold of 1 DN/s. As discussed in Section 2, the apparent 'jump' in the hot pixel numbers around June 2013 can be explained by a code change affecting the treatment of hot pixel neighbors.

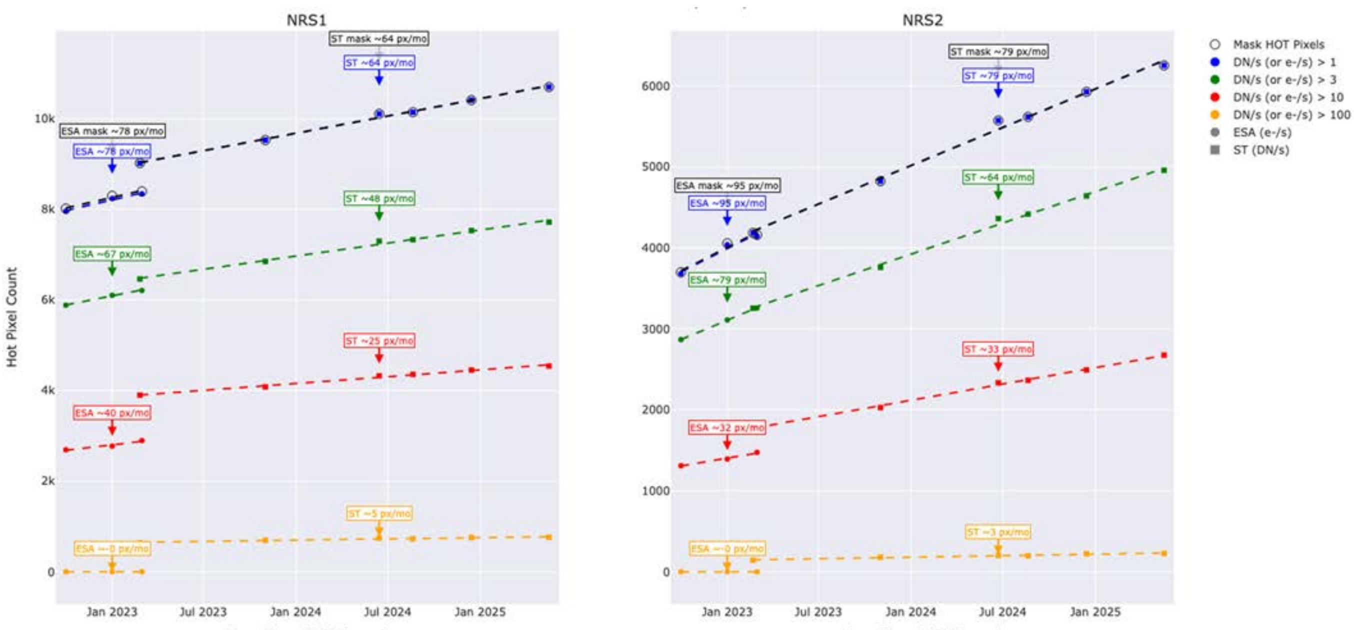

Figure 2: as Figure 1, but for TRAD readout mode. While the total numbers are different from those in IRS2 mode, the growth rates are similar.

4. Consequences for Target Acquisition

4.1 Background

In order to ensure accurate target placement in NIRSpec's narrow slit apertures, observers can employ two different Target Acquisition (TA) procedures[1], both of which use onboard code belonging to the Operations Script Subsystem (OSS) to autonomously compute a corrective slew to place the science target(s) into their intended slit aperture(s).

The first of these methods is called Wide Aperture TA (WATA) and is typically used for single-object observations in the Integral-Field Spectroscopy (IFS), Fixed Slit Spectroscopy (FSS), or the Bright Object Time Series (BOTS) modes. For WATA, the target is placed in the 1.6″ x 1.6″ wide S1600 slit, which is always imaged onto the same area on the NRS1 detector. Therefore, the prompt masking of newly identified hot pixels in this detector area is especially critical, as otherwise they will adversely affect the centroiding of all WATA targets.

The other TA routine is called Micro-Shutter Array TA (MSATA), and the details of its implementation are described in Keyes et al. (2018). MSATA is used exclusively for Multi-Object Spectroscopy (MOS) mode. For MOS observations, precise control of the observatory roll angle is important, which is why the MSATA algorithm must determine the location of multiple reference stars (typically 5 to 8) to compute the corrective slew and roll adjustment. Because the MSATA computation is done by averaging the centroiding results of all reference stars (with some built-in tolerance for outliers), the presence of hot pixels is less critical than for WATA. Nevertheless, the tracking and masking of hot pixels is also important for MSATA, given that the reference stars can fall anywhere on the detectors.
For both methods, the onboard TA code uses a centroiding algorithm on a number of pre-determined areas of the detector (so-called 'postage stamps'), each of which should - after the initial ('blind') pointing of the observatory - contain a point-like 'reference target' (a star, or in some cases a compact galaxy). It is therefore important that any 'hot' detector pixels inside those postage stamps are known to the TA algorithm, so that they can be ignored/replaced when locating the brightest target inside a postage stamp, in order to avoid impacting the calculated centroid of the TA target or the erroneous identification of a hot pixel as the intended TA target. To this end, dedicated onboard reference files (one per detector) are used, which contain the relative response of all pixels and identify/flag all known hot pixels.

The pre-launch versions of these files, commonly referred to as the 'TA flat field', were created from ground-test data, and are described in Proffitt & King 2021. During JWST commissioning, the pre-launch TA flat was replaced in May 2022 with a version based on in-orbit detector data (see Proffit 2021 for details). Given that the hot pixel population has been continuously growing since then, it is prudent to ask whether (and how often) the TA flat field files need to be updated to ensure smooth TA operations.

[1] For details, see the JWST User Documentation for WATA and MSATA

It is important to note that NIRSpec TA exposures are always obtained in the TRAD readout mode (because the interleaved reference pixels in IRS2 mode currently prevent the creation of ‘postage stamps’ in the onboard OSS code). The following discussion therefore uses the hot pixel trending results of Figure 2.

4.2 Impact of hot pixel growth on WATA

The WATA procedure is always performed on the same (32x32) pixel area (the SUB32 subarray), and any newly developed hot pixel(s) within that window will therefore be evident in all WATA images. They can only be masked once detected, i.e. on an ‘as needed’ basis. In other words, there is nothing one can do to pre-emptively avoid hot pixel-induced WATA failures. This is different from the MSATA procedure, as discussed in the next section.

4.3 Impact of hot pixel growth on MSATA

In contrast to the WATA case, MSATA uses different detector pixels every time, and therefore, the impact of a growing hot pixel population can only be assessed via a statistical analysis. For a first order estimate of the likelihood that MSATA will be affected by unmasked hot pixels, we compare the total number of pixels involved in the MSATA procedure to the expected number of new hot pixels per year. In the following, and as explained in Section 2, we will define a pixel as ‘hot’ for TA exposures if its dark current rate exceeds 3DN/s.

Based on the green curves in Figure 2, and using the NRS2 number as a conservative estimate, the number of pixels with dark rates above 3DN/s grows by ~70 pixels per month, or about 0.02% of total pixels per year. Assuming that the spatial distribution of (new) hot pixels is random, their measured growth rate of 0.02% of total pixels per year implies that after one year of not updating the onboard TA flat field file, a randomly selected pixel has a 0.02% chance of being ‘hot’ and unmasked.

The ‘postage stamps’ used for MSATA are (32 x 32) = 1024 pixels in size. Therefore, the likelihood of a MSATA postage stamp including at least one unmasked hot pixel is $1 - (0.9998)^{1024} = 0.19$. In other words, every 5th postage stamp will be affected by at least one unmasked hot pixel. If we also assume that on average 8 reference stars are used for MSATA, the likelihood that a given MSATA activity uses at least one postage stamp with an unmasked hot pixel (which could affect the centroid measurement) is $1 - (0.81)^8 = 0.81$, or 81%. For comparison, allowing two years between TA flat updates would result in a 34% chance of a given postage stamp being affected, and a 96% chance that one of eight is. Allowing only 6 months between TA flat updates, the corresponding numbers would be 10% and 57%, respectively.

Given the built-in robustness of the MSATA algorithm against outliers, this is not as serious as it may sound. However, a sizeable fraction of MSATA procedures (esp. in sparse fields) is confined to using fewer than 8 reference stars. In cases with only 5 reference stars (the minimum required), any one outlier caused by an unmasked hot pixel can potentially degrade the MSATA performance, or even cause MSATA to fail completely. Therefore, it is clearly desirable to maintain the onboard flat field masks on at least a yearly cadence.

5. Summary

We have reported trending measurements of the hot pixel population of the two NIRSpec detectors, which demonstrate a steady growth rate of no more than 100 new hot pixels per month and detector, or 0.03% of total pixels per year. We have discussed the implications of this result on the desirable cadence for updating the onboard TA flat field files in order to minimize the risk of TA failures caused by unmasked hot pixels.

While the first three years of JWST science operations did not cause any TA failures due to unmasked hot pixels, the potential impact of failing to update the onboard TA flats was highlighted in a couple of MSATA procedures with compromised accuracy. These cases could be traced to anomalous centroid measurements caused by unmasked hot pixels.

We conclude that the onboard TA flat field files should be updated on a roughly yearly cadence, following the procedure described in Keyes et al (2025, in prep.).

6. References

Birkmann, S. et al, 2022, SPIE Proc., Vol. 12180, id. 121802P

Karakla, D. et al., 2025, Technical Memorandum STScI-xxxxx (in prep.)

Keyes, C. D. et al., 2018, Proc. SPIE, 10704, 107041J-1

Keyes, C. D. et al., 2025, Technical Memorandum STScI-xxxxx (in prep.)

Koehler, R., 2016, "OSS Flat Field Format," Technical Memorandum JWST-STScI-004955

Proffitt, C. and King, T., 2021, "Preparation of NIRSpec Target Acquisition Flat Field files for use by OSS" Technical Memorandum STScI-007851

Proffitt, C., 2021, "Creating and Distributing the Onboard Flat-Field Reference Files for Target Acquisition", ESA Commissioning Analysis Procedure CAP-069

Shane, N., 2025, "Delivery of JWST Target Acquisition Files to OSS" Technical Memorandum JWST-STScI-008826