The spectro-photometric calibration of the JWST NIRSpec instrument

T. Böker[a], S. Birkmann[a], G. de Marchi[a], P. Ferruit[a], G. Giardino[a], M. Sirianni[a], T. Beck[b]
[a]European Space Agency, Keplerlaan 1, 2200 AG Noordwijk, The Netherlands; [b]Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA

ABSTRACT

NIRSpec is the main near-infrared spectrograph on board the James Webb Space Telescope, offering multi-object capabilities as well as an integral field unit and a number of fixed slits for studies of individual objects. In this paper, we describe the unique challenges in calibrating this complex instrument, and the approach taken to deal with them, both in terms of operational procedures and via automated processing of NIRSpec data. We provide a high-level description of the sequence of processing steps required for NIRSpec science data, and the necessary on-ground calibration files. We focus our discussion on the case of a typical multi-object observation with the MSA, in which adjacent micro-shutters are used to sample the science object and the sky background in an alternating way. This dithering strategy is particularly well suited for faint targets, but its guiding principles also apply to other NIRSpec modes.

Keywords:

1. INTRODUCTION

Over the past two decades, the astronomical community has grown accustomed to, and nowadays expects, a high level of automated data reduction for space-based observatories. The general idea is that, after downlink, astronomical raw data are fed through a "data pipeline" which is continuously updated to reflect up-to-date knowledge about the telescope and its various focal plane instruments. This approach, which has been pioneered by the Hubble Space Telescope, ensures that all observers can share optimal reduction strategies, and can benefit from all "lessons learned" promptly and in a common manner.

The general goal of such a data pipeline is the spectro-photometric calibration, i.e. the accurate conversion of counts registered in the detector electronics (measured in digital numbers, DN) to the actual signal received from the astronomical source (measured in physical units, e.g. W m^{-2}). This conversion must correct for all photon losses incurred throughout the entire optical path between the telescope primary mirror and the instrument detector. For each observing mode (i.e. each combination of filter, disperser, and detector), the correction must take into account the throughput differences with wavelength and position in the field of view (FOV). In the case of a multi-object spectrograph (MOS) in general, and NIRSpec in particular, there are a number of effects that make an automated "pipeline" approach to data calibration particularly challenging:

1. In contrast to traditional long slit spectrographs, in which a given wavelength always falls onto the same detector area, every pixel in a MOS must be calibrated for its response over the entire spectral range. This is because spectra from various aperture locations can illuminate the same detector pixel, albeit at different wavelengths. For a MOS, therefore, the "flat field" response map is a three-dimensional data cube that "stacks" the two-dimensional pixel response maps along the wavelength dimension. In principle, the wavelength axis must be sampled at the spectral resolution of the instrument (up to R~2700 in the case of NIRSpec). Possible simplifications of this approach, which would otherwise result in rather large and impractical calibration reference files, are discussed in this paper.

2. The FOV of NIRSpec, combined with the choice of using reflection gratings, leads to a mechano-optical design in which light hits the dispersive elements at rather large and varying incidence angles. This leads to a significant curvature of the projected slit apertures on the NIRSpec detectors. While each individual shutter aperture is essentially tilted with respect to the dispersion axis, the tilt angle varies significantly across the field of view and with wavelength. This complicates calibration and data reduction because each shutter spectrum must be rectified using different parameters.

Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave, edited by Mark C. Clampin,
Giovanni G. Fazio, Howard A. MacEwen, Jacobus M. Oschmann, Jr., Proc. of SPIE Vol. 8442, 84423F

3. Unlike for ground-based telescopes (which have a spatial resolution limited by atmospheric seeing), the point spread function (PSF) of a telescope in space is a strong function of wavelength. As a result, the full width at half maximum (FWHM) varies nearly by a factor of 10 between the blue and red ends of the NIRSpec spectral range. On the other hand, NIRSpec uses only a single imaging scale for its entire wavelength range, and the Micro-Shutter Assembly (MSA) offers only a single physical slit width. Although the shutter width (200 mas) has been carefully selected to optimise throughput and resolution across the entire spectral range, the fraction of light that passes through an MSA shutter (or fixed slit) is nevertheless a strong function of wavelength. In principle, the slit throughput can be accurately modelled, but this approach requires exact knowledge of the source shape and position within the shutter. Given the faintness of targets and the limited accuracy of the target acquisition process, such a priori knowledge is not generally available. Note that the slit loss variations with wavelength are most pronounced for a point source, while in the case of a source that uniformly illuminates the entire shutter or slit, only a small correction for wavelength dependent diffraction loss is needed.

4. Another effect to consider are diffraction losses due to the finite size of the NIRSpec optics behind the MSA plane, in particular the gratings. The diffraction pattern depends on wavelength, source shape and position within the slit, and thus the amount of light that is diffracted outside the NIRSpec optical train is difficult to predict. Because both these effects - the geometric light loss due to “cutting off” the PSF at the shutter/slit, and the diffraction losses in the optical train behind the MSA - have similar consequences, they are conceptually combined into a single term, the *chromatic slit loss*.

5. Because of the large field of view, use of a purely reflective optical train, and deployment behind a fast, off-axis telescope, optical distortion is significant in NIRSpec. The resulting plate-scale variations across the aperture provide additional challenges for the spectro-photometric calibration because one must account for the varying solid angle seen by both detector pixels and slits.

The complexity of the effects discussed above, combined with the limited testing facilities on ground and stringent efficiency requirements in orbit, leads to an approach for NIRSpec calibration which relies heavily on optical modelling of the instrument. Wherever possible, parametric descriptions of the instrument features will be constructed that are tested, verified, and optimized at every stage of the NIRSpec development, from sub-system level to instrument level to in-flight operations. This makes it possible, for example, to predict the precise location, shape, and wavelength calibration of individual micro-shutter spectra of faint sources without the need for a dedicated lamp exposure through the same shutter.

2. OBSERVING WITH NIRSPEC

The baseline NIRSpec multi-object usage scenario is illustrated in Figure 1a: each science target (assumed to be point-like, i.e. smaller than ~ 0.2") is positioned in a dedicated “slitlet” consisting of three adjacent opened micro-shutters. During a given exposure, one shutter contains the target, and two neighboring shutters in the cross-dispersion direction measure the adjacent background. In this manner, using small-angle maneuvers of the spacecraft in the direction along the slitlet, all objects can be simultaneously moved up or down in the slitlet by one shutter, as illustrated in Figure 1b.

This approach not only improves the accuracy of the background subtraction, but also dithers the object spectra on the detector between exposures, so that detector blemishes can be corrected. Moreover, this observing scheme can be easily implemented in an automated manner without reconfiguring the MSA. The remainder of this document discusses the data reduction steps that are required to properly combine data taken in this (or a similar) manner in order to achieve the highest possible signal-to-noise ratio.

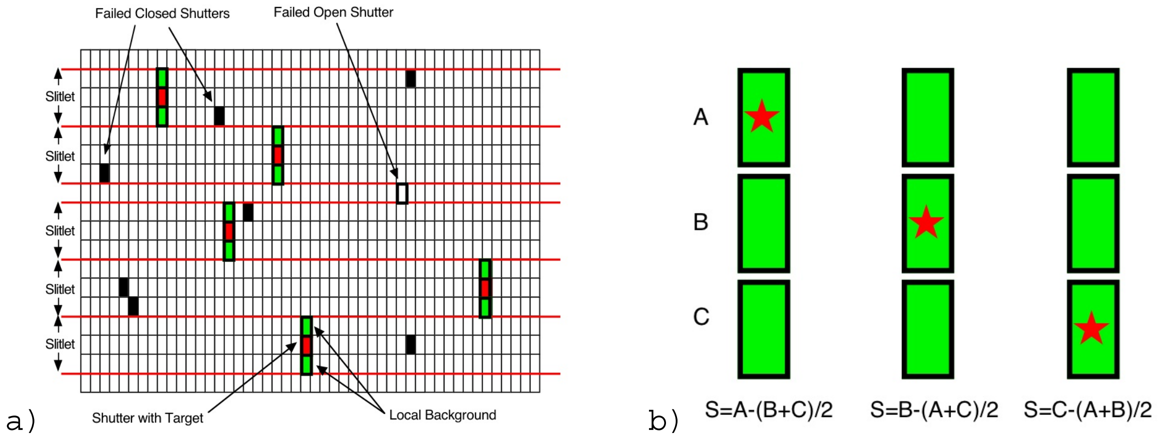

Figure 1: a) Illustration of the "slitlet" concepts for observing (small) NIRSpec science targets in three adjacent micro-shutters **b)** By dithering the telescope, the sky background can be sampled by the same detector pixels in subsequent exposures.

3. TWO KINDS OF DATA PIPELINE

The pipeline is a series of operations that need to be carried out on detector raw data to remove instrumental signatures and to convert, as a function of wavelength, DNs to absolute flux units. For NIRSpec, and MOS instruments in general, it is convenient to split these operations in two groups: i) those that depend solely on pixel coordinates on the detector, and ii) those that depend on pixel coordinates, wavelength of the incident light, and location of the (undispersed) source in the field of view.

The first group includes corrections for dark current, non-linearity, and bias level, all of which are independent of the disperser and filter used. We will refer to this first group of operations as the *pre-processing pipeline*. Operations belonging to the preprocessing pipeline are applied to the "raw" data files which, for the purpose of this paper, are the (2048 x 2048 x n_f) FITS files produced by so-called "up-the-ramp" sampling, i.e. a sequence of n_f non-destructive readouts (or "frames") of the Hawaii-2RG near-infrared detectors used in NIRSpec.

The second group of operations includes corrections for sensitivity variations (on both global and local scale), geometric distortion (affecting the flux, spatial scale and wavelength of each pixel), and absolute flux calibration. These steps are more logically executed at the level of individual spectra because they depend on which specific slit apertures (FS, IFU, or micro-shutter) are used. The following sections describe separately the specific operations in each of the two pipeline parts.

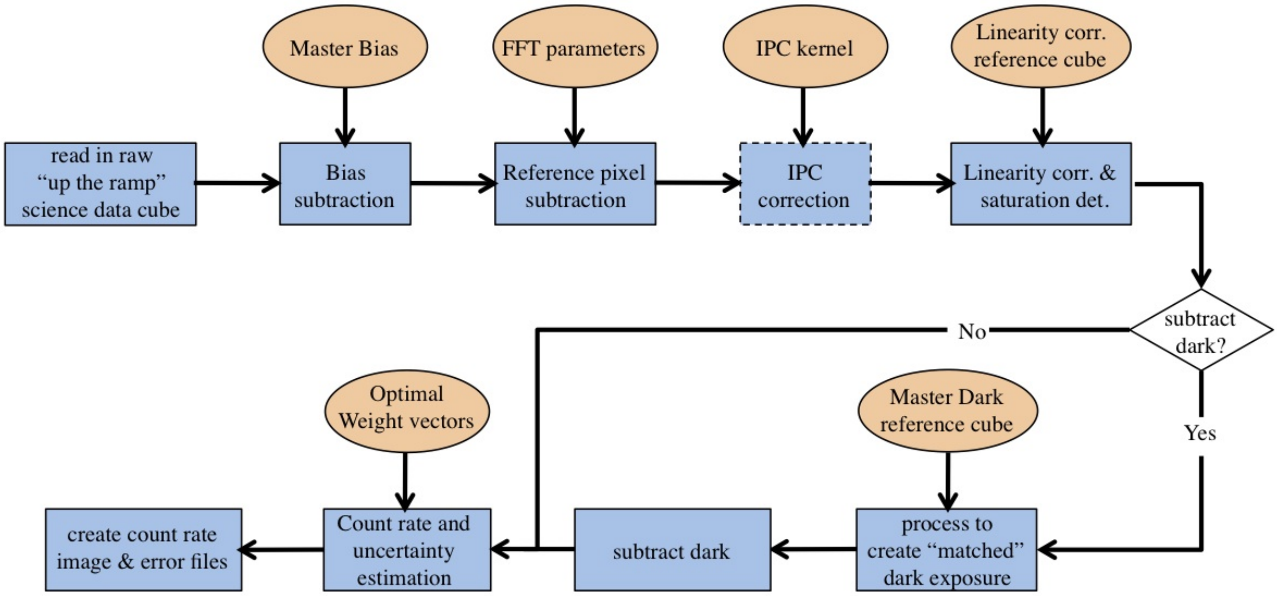

Figure 2: Schematic outline of the NIRSpec pre-processing pipeline

4. PRE-PROCESSING PIPELINE

The currently planned sequence of operations in NIRSpec pre-processing pipeline is illustrated in Figure 2. The individual steps are as follows:

1) Bias subtraction (default)

A ”master" bias frame is subtracted from every frame in the input file, removing all systematic (but not time-dependent) pixel-to-pixel variations in the reset level, as well as any large-scale gradients. Following this step, normal pixels should begin their integration ramps at close to zero DN under dark conditions. This is important for reference pixel subtraction and linearity correction, and therefore this step is the first to be performed.

2) Reference pixel subtraction (default)

This step is performed to remove possible time-dependent bias drifts and/or 1/f noise from individual readout frames, and consists of two parts. First, the four rows of reference pixels at the top and bottom of the detector (i.e. along the fast readout direction) are used to compute eight σ-clipped mean reference values for each frame: one each for the odd- and even-numbered columns for each of the four video outputs. These reference values are then subtracted from all corresponding pixels in the respective frame.

Then, the four reference columns to the left and right of the chip (i.e. along the slow readout direction) are used to remove any 1/f noise that is common to all outputs. This is done by first averaging the eight pixel values in each row row, yielding one reference value for each of the 2048 columns. Then, an FFT based smoothing is applied to this reference vector, and the result is subtracted from every row (see Moseley et al. 2010 for details).

3) IPC correction (optional)

Due to the inter-pixel-capacitance (IPC), charge that is accumulated in one pixel is partly (to the 1% level) detected in neighboring pixels. In principle, this effect can be corrected by convolving the 2D image (on frame by frame basis) with an appropriate (3x3 or 5x5) kernel. The kernel has negative values in the wings, and thus moves charge back onto the central pixel. At the moment, the baseline for the NIRSpec pipeline is not to perform any IPC correction.

4) Linearity correction and saturation detection (default)

Even under stable illumination conditions, up-the-ramp data is not completely linear, because the integrating capacitance

of each individual pixel changes (increases) with accumulated charge. NIRSpec ramps are linearized by applying a fourth order polynomial to the data, using individual coefficients for each pixel. The details on how the reference file is created will be described in a separate document, but the principle follows the recipe of Robberto (2011).
For a proper linearity correction, both bias and reference pixel subtraction must have been performed, and the pipeline will skip the linearity correction if these steps have not been executed before. As part of the linearity correction, saturated pixels will be flagged and the saturated part of the ramp will not be used for calculating the slope.

5) Dark current subtraction (optional)
When selected, the NIRSpec pipeline subtracts a "master dark" current reference file from the science exposure. This is done at cube level. The master dark cube is created by averaging a large number of "up the ramp" dark exposures, and thus its noise is much smaller than that of the science exposure. The pipeline reads and processes the master dark in exactly the same way as the science exposure, i.e. it performs bias subtraction, reference pixel subtraction, IPC correction (if switched on) and linearity correction *before* the dark is subtracted on a frame-by-frame basis. This approach improves the subtraction of hot pixels that suffer from field-enhanced emission and therefore non-stable dark currents.

6) Slope and uncertainty estimation (default)
In principle, once the up-the-ramp data is linearized, one could fit a straight line to the data and determine the slope. However, using a linear least squared fit with uniform weights is not optimal for high signal-to-noise data. Therefore, we adopted the algorithm described by Fixsen et al. (2000), which makes use of pre-calculated weighting vectors for 15 signal-to-noise bins for the given number of groups. In the low SN limit, the weights are uniform (damping the read noise as much as possible), while in the high SN limit, the weights are zero everywhere except for the two endpoints of the ramp (maximizing the effective integration time).

5. EXTRACTION PIPELINE

After the pre-processing pipeline has removed all effects that are caused by the detector system, a number of operations need to be carried out to remove the signature of the telescope and instrument optics, and to convert the measured count-rates to absolute units. Because these corrections are strongly wavelength dependent, and also vary across the field of view, they can only be executed on a spectrum-by-spectrum basis.

The required operations are conceptually not different from those normally carried out with traditional ground-based spectrographs, and include: removing large scale variations in the detector response (usually referred to as L-flat); tracing the spectrum and ensuring that dispersion and cross-dispersion axes are orthogonal (rectification); assigning a wavelength to each pixel (wavelength calibration); correcting the measured count rate for field- and wavelength-dependent variations in the throughput such as the grating blaze function (relative flux calibration); and comparing the measured count rates with those of a standard star (absolute flux calibration).

The currently planned sequence of operations in NIRSpec extraction pipeline is illustrated in Figure 3, the individual steps are as follows:

1) Extraction window definition
The first step towards deriving scientifically useful spectra from the count rate images produced by the pre-processing pipeline is to "cut out" a so-called *extraction window* for each astronomical target. As illustrated in Figure 4, an extraction window is a rectangular portion of the detector containing the area illuminated through each open slit aperture, i.e. one of the fixed slits, a micro-shutter, or the virtual slits created by the NIRSpec integral-field unit (IFU). The pipeline extracts one window for each open slit aperture, and from this step onwards, ignores the regions of the detector that are not exposed to light. The extraction window must be large enough to sample all usable light. Because of the unavoidable curvature of the spectra, extraction windows for neighbouring slit apertures may overlap.

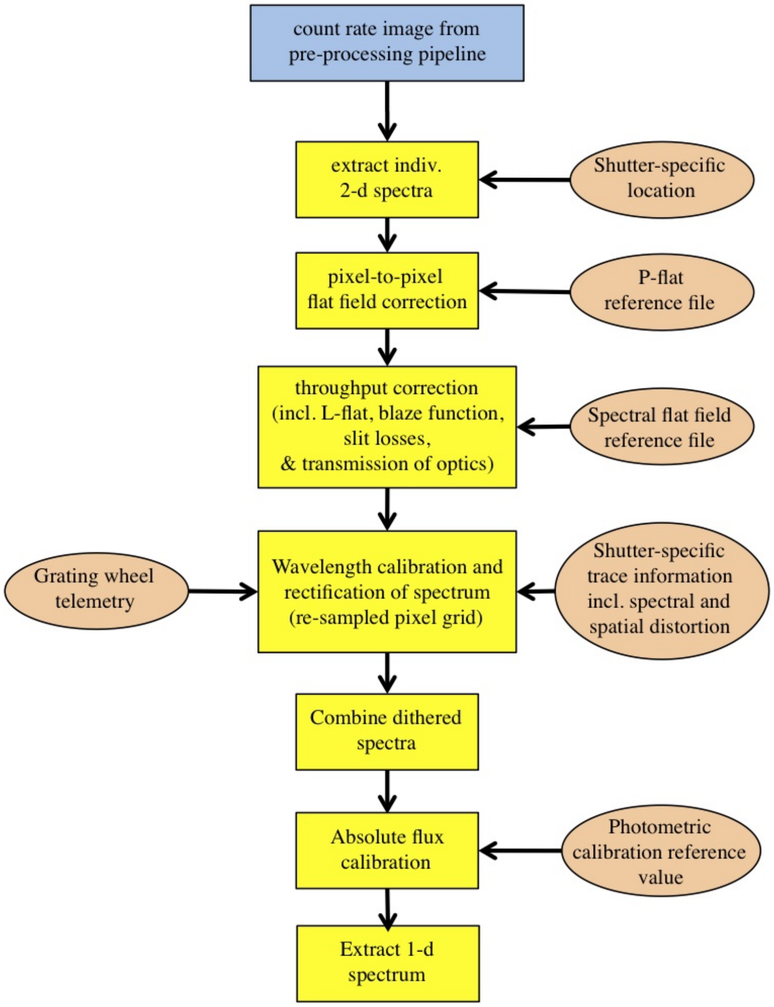

Figure 3: Schematic outline of the NIRSpec extraction pipeline

For each instrument mode (i.e. each combination of filter and disperser), the location of any extraction window can be accurately predicted, because it is, to first order, defined by the grating/prism equation and geometric distortion of the NIRSpec optics. This information is contained in the optical model of the NIRSpec instrument, which is tested and calibrated during the instrument-level test campaign via a limited number of spectra through selected apertures. For the general case of arbitrary combinations of open micro-shutters, the extraction pipeline then relies on the instrument model to define the location and size of each extraction window, using information in the FITS header about the selected aperture(s) and instrument mode[1].

In principle, the wavelength falling onto each pixel in the extraction window is fully determined by the instrument model. However, at this stage, neither the wavelength scale nor the spatial scale of each pixel is uniform across the extraction window because of spatial and spectral distortion. This is the reason why one cannot simply average extraction windows from different detector regions (obtained, e.g., from dithered observations) at this stage, but rather has to rectify and rebin them in order to arrive at a common spatial and spectral scale (see Step 4 below).

---

[1] In principle, the positional repeatability of the grating wheel mechanism may affect the precise location of the spectrum on the detector. However, tests have shown that this effect is minor, and can be ignored for the extraction window definition (but not for the wavelength calibration, see de Marchi et al. in these proceedings).

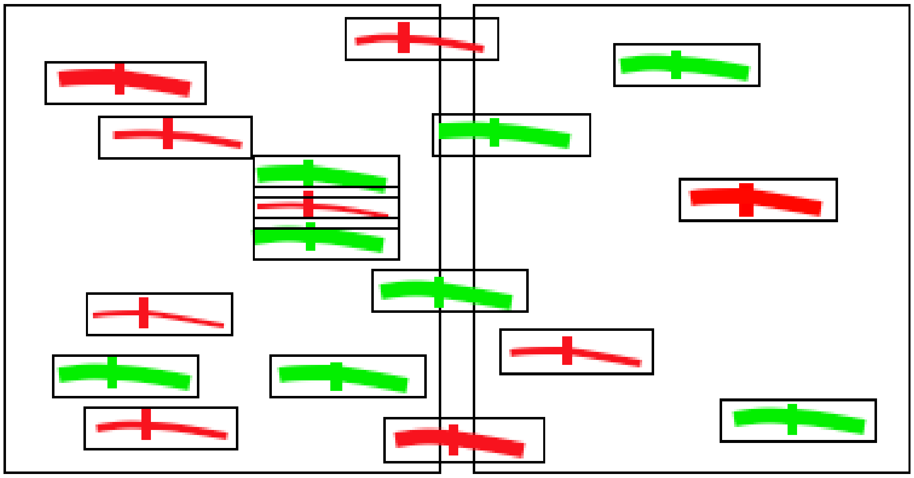

Figure 4: Schematic view of a NIRSpec count rate map. Each open aperture (micro-shutter or fixed slit) produces a spectrum on the detector (red for science targets, green for sky background). The precise location of the spectra, their length as well as their curvature vary with the dispersing element used. The extraction window associated with each spectrum must be large enough to contain all target light within the useful wavelength range of the selected disperser.

2) Pixel-to-pixel flat field correction

As part of the characterization of the NIRSpec detector system, the pixel-to-pixel response variations (the so-called P-flat reference files) have been obtained at ~20 wavelengths across the NIRSpec wavelength range. These tests have shown that variations in the P-flat with λ are small, but not completely negligible. For this reason, it is desirable to apply the P-flat correction only after it is known which wavelength is falling on the respective pixel.

Depending on the characterization of the final NIRSpec flight detectors, and on the wavelength range of the science spectrum, it may be sufficient for the pipeline to simply cut out the relevant pixel area from the P-flat reference file closest in wavelength, and to apply this "P-flat extraction window" to the science data. A slightly more complex approach would be to construct a P-flat reference cube from the detector characterization data, and to apply a correction that is more finely sampled in λ.

3) Correction for instrumental throughput variations

As described in the introduction, the complexities of a MOS and the practical difficulties for using "standard" calibration techniques with the MSA in the NIRSpec optical train, have serious consequences on the approach to calibrate NIRSpec, both on the ground and in orbit. In particular, some of the calibration steps that are usually performed separately in standard long-slit spectrographs, cannot easily be disentangled and thus must be performed simultaneously for NIRSpec.

An example is the correction for large-scale variations in the detector response (the so-called "L-flat") and instrumental throughput variations, i.e. the transmission of all optical elements along the NIRSpec optical path, incl. the filter and the disperser. Both effects depend on field angle (or pixel position) as well as wavelength, and both are characterized as much as possible at sub-system level. However, after integration of the full NIRSpec instrument, the dependence of the transmission of the NIRSpec optical train on field angle and wavelength is intimately tied to the response of the NIRSpec detectors, and the latter can no longer be measured independently due to the presence of the MSA in the optical path.

The approach taken, therefore, is to measure the combined effect of throughput and detector response. Here, too, time constraints do not permit such measurements to be obtained for each individual shutter. Instead, a suitable subsample of open shutters spread out over the NIRSpec field will be used, each uniformly illuminated with the same known input spectrum. To this end, NIRSpec is equipped with a set of internal continuum lamps with calibrated spectra that uniformly

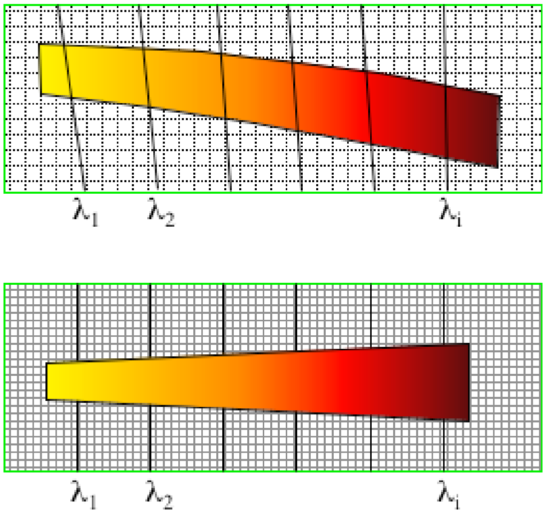

Figure 5: Schematic view of an extraction before (top) and after (bottom) rectification.

illuminate the MSA. These lamps can be used to measure, during thermal vacuum and in flight, the total instrument throughput as a function of wavelength and position.

For each NIRSpec mode, the resulting data set forms a sparsely populated three-dimensional array, with X and Y along the detector's coordinates and Z scaling with wavelength. Since any variations in the overall instrument response with position or wavelength are expected to be rather smooth, they can be interpolated via a low-order polynomial. In principle, therefore, it is only necessary for the pipeline to know the parameters of this polynomial in order to apply the proper correction for any micro-shutter. The advantage of this approach is that it can be easily repeated in flight, for a subset of the exposures, to verify the temporal stability of the throughput correction.

4) Rectification of spectra and wavelength calibration

As mentioned before, distortion in the NIRSpec optics causes the spectrum within each extraction window to be curved. In addition, neither the wavelength, nor the spatial pixel scale is constant across the window. In practice, this implies that the dispersion direction is not aligned with the pixel rows, and therefore the iso-wavelength lines are not parallel (see Figure 5.

In order to carry out extraction operations such as co-addition, averaging, or subtraction in a direction perpendicular to that of the dispersion – all rows in the extraction window must be brought to the same wavelength scale, i.e. the spectrum must be *rectified*. The pipeline does so by re-sampling all rows and columns onto a common, three to four times finer wavelength and spatial scale, in a way that conserves the flux and does not introduce additional noise. This can be achieved using linear reconstruction techniques or, preferably, an approach similar to the *drizzle* technique normally used for undersampled HST images (Fruchter & Hook 2002).

The resulting over-sampled grid is defined in such a way that wavelength increases linearly along rows and the wavelength scale is uniform, i.e. the sampling $\delta\lambda$ is the same for all grid elements within an extraction window, and for all extraction windows obtained with the same NIRSpec observing mode. Similarly, the plate-scale (in arcsec/element) in the cross-dispersion direction is also the same for all grid elements and for all extraction windows, allowing direct comparison of objects in different areas of the FOV which is essential for the next step of combining spectra.

5) Combination of dithered data and background subtraction

Nearly all NIRSpec observations will place the astronomical target at various places within the field of view in order to allow correction of detector blemishes. These spectra must be averaged in order to obtain the maximum signal-to-noise ratio. Similarly, the sky background that must be subtracted in order to obtain a "clean" source spectrum may be measured in a different field position. The addition or subtraction of individual spectra is only possible once they have been brought onto a common spatial and spectral sampling.

The challenge for the extraction pipeline at this point is to keep track of which spectra must be combined. In other words, the pipeline must group together all spectra of the same source and the corresponding sky background in so-called *associations*. It should be feasible to implement an automated data combination for spectra that are acquired with a pre-defined dither pattern. On the other hand, it may not be possible to automate the combination of MSA spectra acquired with a "one-off", user-selected dither pattern. At this time, the details of associations and the corresponding data combination process are still to be defined.

6) Absolute flux calibration

The *relative* flux calibration of each pixel within each window and across all extraction windows has been performed by the combined throughput corrections as described ion step 3) above. As a result of those operations, the values of the grid elements in the oversampled and re-binned extraction window are correct in a relative sense, but are still in arbitrary physical units. Therefore, the conversion into absolute physical units simply requires multiplying all values by a constant which is obtained - for every NIRSpec observing mode - by observing a suitable spectro-photometric standard star. These calibration constants are regularly monitored and are centrally stored, so that the pipeline always uses the best available calibration.

7) Extraction of 1-d spectra

The final step in producing scientifically useful 1-dimensional spectra of the science target requires the "collapse" of the re-sampled, 2-dimensional extraction windows. The details of this step depend very much on the source shape and the quality of the rectified spectra. Ideally, optimal extraction methods will be used, but a straight collapse may be necessary if the source is too faint for optimal extraction. At the end of the full reduction process, it should be possible to automatically combine multiple spectra of the same target observed with different MSA configurations or as part of different pointings or visits, in order to increase the signal-to-noise ratio.

REFERENCES

[1] Fixsen, D. J., et al. "Cosmic-Ray Rejection and Readout Efficiency for Large-Area Arrays", PASP, 112, 1350 (2000)

[2] Fruchter, A. & Hook, R., "Drizzle: A Method for the Linear Reconstruction of Undersampled Images", PASP, 114, 144 (2002).

[3] Moseley, S. H. et al., "Reducing the read noise of H2RG detector arrays: eliminating correlated noise with efficient use of reference signals", SPIE Conf. Series, Vol. 7742, 36 (2010).

[4] Robberto, M., "Implementation of a new algorithm for the correction of non-linearity in JWST detectors", STScI Technical Report JWST-STScI-002344, STScI (2011)