STScI | SPACE TELESCOPE SCIENCE INSTITUTE

JWST ANALYSIS REPORT

Title: Verification of Plan to Solve for the Distortion Solution	Doc #: JWST-STScI-005361, SM-09.3 Date: 28 October 2016 Rev: -
Authors: Jay Anderson Phone: 410 338- 4982	Release Date: 20 September 2017

ReqID	Requirement Title and Text	Parent	Children	I	A	D	T	Verification Compliance Summary
SOC-2686	3.4.3.2.5 Field Distortion Uncertainty The DMS shall calibrate the field mapping of the SI coordinates relative to standard astrometric fields so that after calibration, the field distortion uncertainty within any SI and the guider does not exceed 0.005 arcsec, 1 sigma per axis.	GS-187 : S&OC Observatory Telemetry Calibration	DMS-268		A			Verified by analysis below.
DMS-268	3.2.1.5 Field Distortion Uncertainty The DMS shall calibrate the field mapping of the SI coordinates relative to standard astrometric fields so that after calibration, the field distortion uncertainty within any SI and the guider does not exceed 0.005 arcsec, 1 sigma per axis.	SOC-2686: Field Distortion Uncertainty	N/A		A			Verified by analysis below.

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

Abstract

This report documents the verification by simulation that the distortion solution within each SI detector and the guider will be able to be modeled to better than 5 mas, 1 sigma per axis.

1 Statement of Requirement

The Data Management Subsystem Requirements Document (JWST-STScI-002249), section 3.2.1.5 provides the requirement on the Field Distortion Uncertainty:

> The DMS shall calibrate the field mapping of the SI coordinates relative to standard astrometric fields so that after calibration, the field distortion uncertainty within any SI and the guider does not exceed 0.005 arcsec, 1 sigma per axis. [DMS-268]

This requirement is a direct flow-down of the parent requirement from the S&OC Element Requirment Document (JWST-STScI-0046), Section 3.4.3.2.5:

> The DMS shall calibrate the field mapping of the SI coordinates relative to standard astrometric fields so that after calibration, the field distortion uncertainty within any SI and the guider does not exceed 0.005 arcsec, 1 sigma per axis. [SOC-2686]

This Technical Report will show that a single short observation of the calibration field through each detector should be able to model the distortion within any SI or the guider to better than 5 mas, 1 sigma per axis. This analysis applies to both requirements.

Verification status: Passed.

2 JWST Calibration Field

A field in the LMC (Large Magellanic Cloud) has been chosen to serve as a calibration field for JWST. It is located in the CVZ (continuous viewing zone) and has a nice, flat distribution of stars that are almost ideal for JWST calibration. There are plenty of medium-brightness stars that can be used for (1) photometric calibration, (2) astrometric calibration, and (3) PSF characterization. In addition, there are not too many bright stars (which would saturate and be useless for calibration while at the same time introducing large-scale artifacts across the detector).

The JWST Calibration field covers 5′×5′ and is centered on α = 05:21:56.5 and δ = −69:29:54.1. It was observed in 2006 through F606W with ACS by program GO-10753 (PI-Diaz-Miller) and has been observed subsequently in 2009 through F606W with WFC3/UVIS (CAL-11444, PI-Dressel) and through F160W with WFC3/IR (CAL-14445, PI-Dressel).

One of the reasons this particular field was chosen is that it is far away, and internal proper motions should therefore be small. If LMC stars are moving with an internal dispersion of 50 km/s, then that would result in a PM dispersion of about 0.2 mas/year. Anderson & Diaz (2011) use the three-year baseline above to show that the internal motions are well below 0.03 ACS pixel per year, consistent with this. In 2018, the typical star will have moved about 2.5 mas

relative to the positions given in this catalog. Such motions will introduce a random error for each star, such that their impact on the accuracy of the distortion solution will be smaller by a factor of $\sqrt{N_{stars}}$, so that the impact of star motions will be well below the 5 mas requirement we are considering here and will therefore not prevent us from arriving at an adequate instantaneous correction for distortion.

The catalog for the JWST Calibration Field can be found at the following website: http://www.stsci.edu/~jayander/JWST_CALIBFIELD. The full catalog contains almost 550,000 stars, but there is a catalog of relatively isolated stars that contains 218,000 stars with photometry in V, J, H, and K, along with photometric estimates for the stars in bands where observed photometry is not available. Stacked images in V, J, H, and K are also available on the website. Figure 1 shows the stacked F606W image. The J, H, and K stacks are from the HAWK-I data (see Anderson 2008).

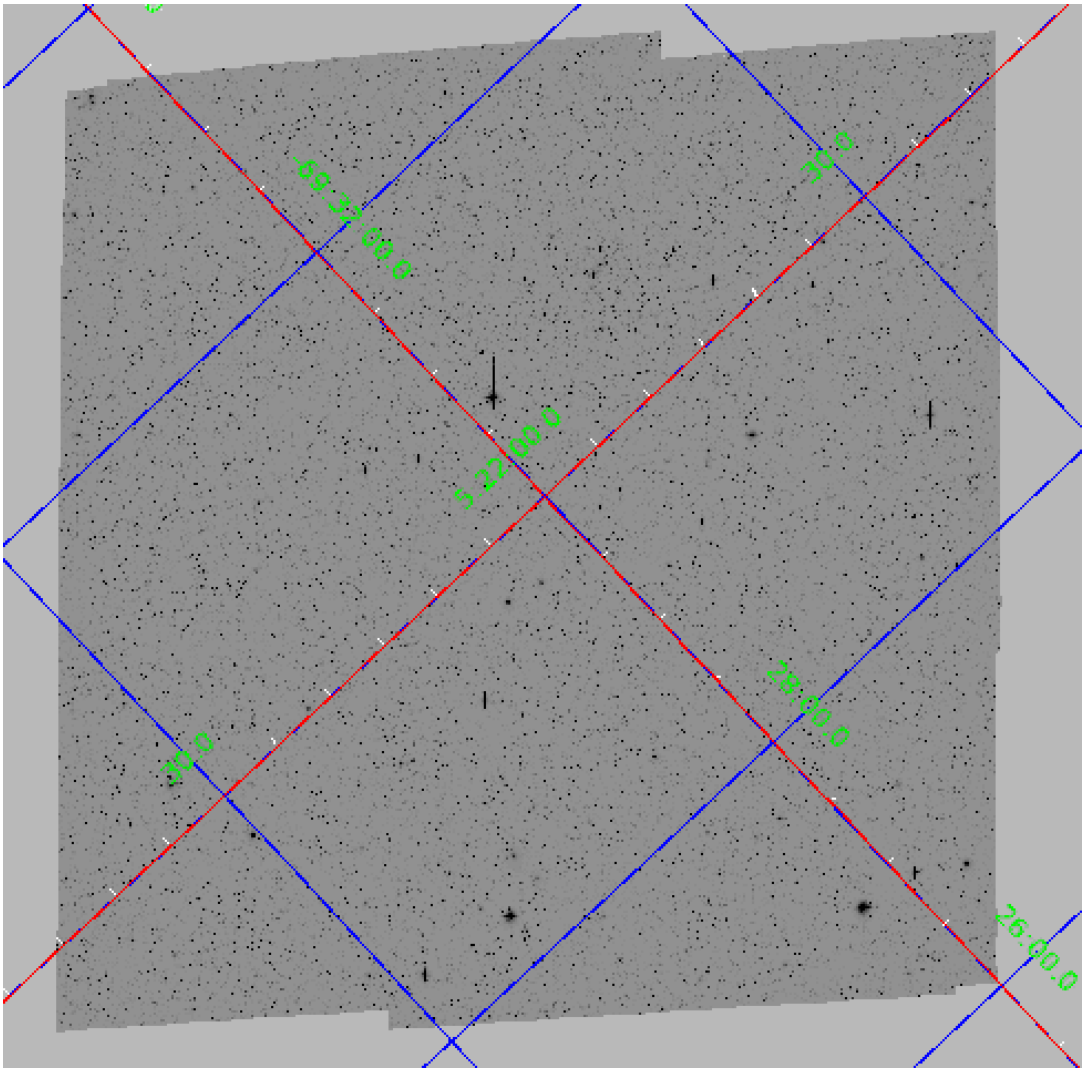

Figure 1: F606W stack of JWST Calibration Field, with lines of latitude and longitude drawn in.

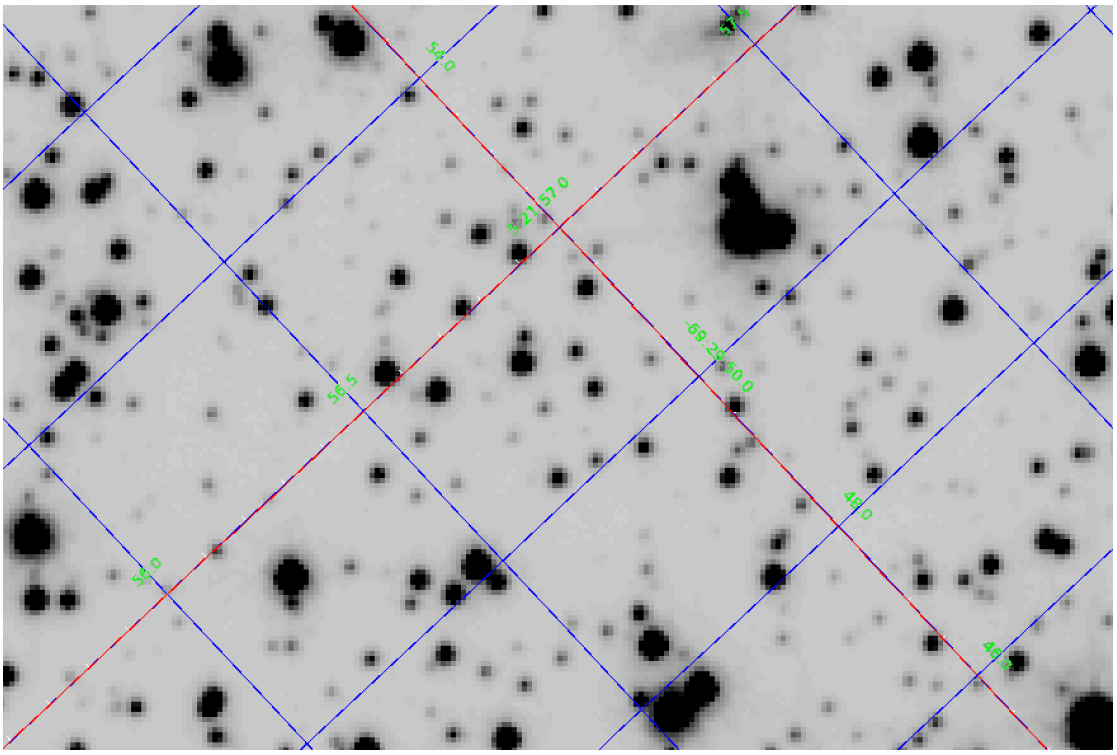

Figure 2: Close up of the center of central 10″×6″ of the calibration field, showing that there are plenty of bright isolated stars.

The 5′×5′ size of this HST-observed field is adequate to solve for the distortion within any SI, but it is not adequate to measure the offset between SIs and the guiders. The HAWK-I field (see Anderson 2008) extends beyond this field, but it has not yet been combined into a large-field catalog and extending the catalog is beyond the scope of this document.

To quickly fill in a region around the HST-observed catalog, we simply supplemented the catalog with about 16,000 2MASS stars, going out to about 15′ from the center of the field. In this way, we can put a 2MASS star at the center of one of the guiders in such a way as to place the target SI on the calibration field so that both the internal distortion can be measured as well as the location and orientation in the V2/V3 plane. 2MASS positions are accurate to only 100 mas (Skrutskie et al. 2006), but Gaia positions should be much better and should be available well before the time of calibration (if not already). If for some reason, the Gaia catalog is not found to be accurate enough to do the large-scale V2/V3 boresite mappings, we can extend the field with HAWK-I. Figure 3 shows the full HST+2MASS catalog we will work with here in this document.

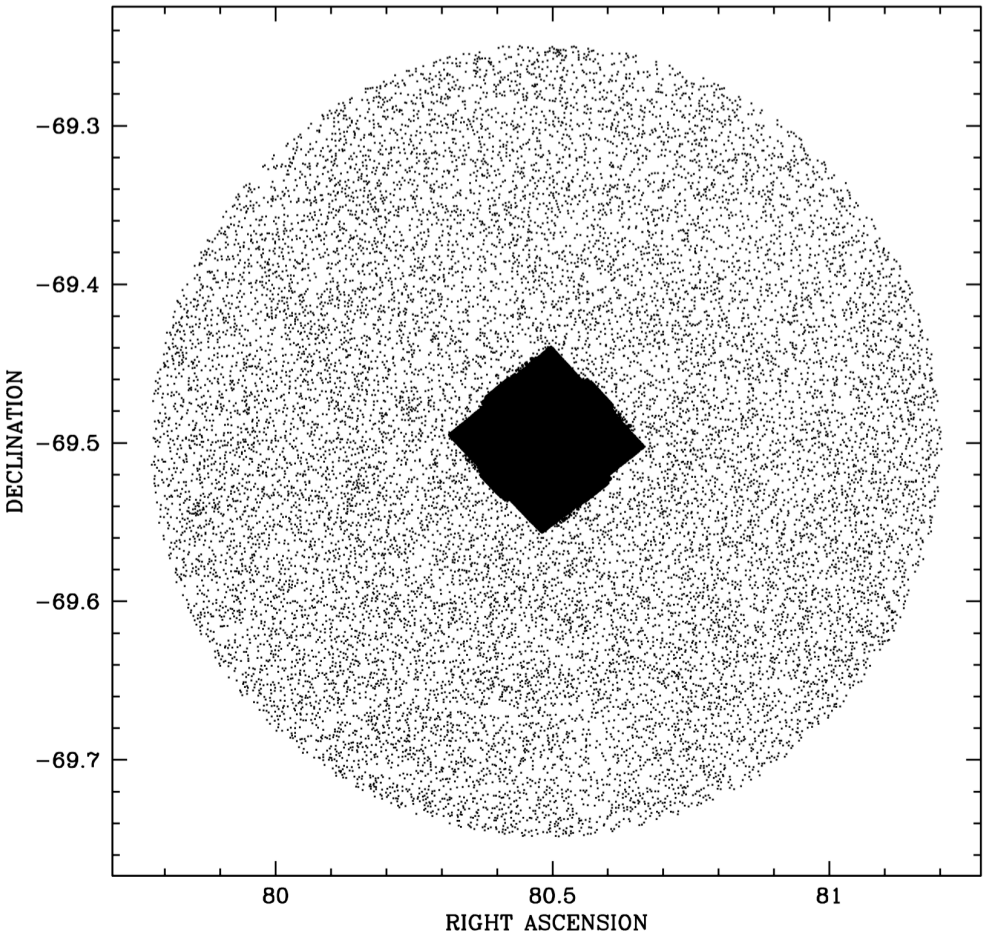

Figure 3: This shows the astrometric catalog we will be working with in this document. The inner 5′×5′ field is from the dense HST catalog, and the sparse outer field (going out to about 15′) is from 2MASS.

3 Showing that the solution can be extracted

The general astrometric-calibration plan during commissioning will be to fine-guide off of a star in one of the FGS detectors, while placing the target SI near the center of the calibration field. This can be done for several different target stars if we desire to get multiple realizations of calibration field in the detector for multiple constraints on the distortion solution. However, in this document we will examine only one placement per detector, since we have found that to be adequate to meet the 5-mas requirement. In addition, we will explore just one particular telescope orientation, though the penumbra around the calibration field should allow us to perform this calibration at *any* JWST orientation.

The location of the various instruments relative to the guiders is codified in the SIAF files, an early version of which has been provided to us by Cox *et al*. (2009). These files contain parameters that codify mappings to and from four different frames: (1) the raw detector frame, (2) the "science" frame, which is an offset (and possibly flipped) version of the detector frame, (3) the "ideal" frame, which is distortion-corrected and tangent-plane projected about the reference point, and finally (4) the v2/v3 plane, which is in a spherical coordinate system.

The reference point for each detector associates an (*i*, *j*) location on the detector, with the location (0.0″,0.0″) in the ideal frame, and with a particular location in (*v2*, *v3*) system. The transformations from the ideal to the science frame are given by 5th order polynomials for NIRCam, and MIRI. The transformation for the FGSs and NIRISS is 4th order. NIRSpec's transformation is more complicated and we will deal with it below.

These transformations represent a composite of two operations. First, a full-telescope ray-trace study by Ball maps the sky to the entrance apertures for each instrument. This was done by Ball and is well represented to extremely high precision by a 5th order polynomial (personal communication). This mapping is combined with a mapping from the entrance apertures to the detector plane, done with models of the camera optics. The composite 5th-order polynomial that maps the ideal frame to the science frame is given in the SIAF files. Even though these distortion models may not be perfect (the instruments could have some slight offsets, rotations, or mis-alignments from nominal, for example), these maps should be *representative* of the actual distortion present. So we assume that if we can solve for these maps directly from simulated observations, then we should be able to solve for whatever actual distortion is present, even if it differs in a quantitative (but not qualitative sense) from the ray-trace-based models.

We note that all of the HST instruments are well characterized to much better than 5 mas with such a fifth-order polynomial. Several instruments have detector-related distortions that are not well represented by polynomials — such as the manufacturing defects or filter distortion fingerprints in WFPC2, ACS, and WFC3/UVIS (see Anderson & King 2003, Anderson & King 2006, Kozhurina-Platais et al 2009), or the hysteresis in the HAWK-I detector (Anderson 2008) — but *all* these errors are well below the 5 mas threshold that we are aiming for to meet the formal requirement. Therefore, since we know of no manufacturing defects in JWST's detectors, and the errors we have seen in other detectors are considerably below the 5 mas level, it is reasonable to assume that if we can solve empirically for the 5th (or 4th) order polynomials that have been found to be adequate for the ray-trace simulations of the telescope and camera optics using the best available models, then we should be able to meet the calibration requirements with confidence.

Our procedure will be to go through each of the instruments one by one in sections 3.1 to 3.6 and generate a realistic simulation of a single short observation of the calibration field through that instrument and demonstrate directly that we can use similar observations to constrain the distortion. Note that here we investigate only the imaging modes of the instruments. Calibrating the spectral modes (such as the IFUs or NIRCam's grism, MIRI's MRS, NIRISS's Aperture-Masking Interferometry and NIRSpec's spectral modes) will not be discussed here, as they require additional calibration data than the Calibration Field can provide.

3.1 NIRCam's A1 Detector in the Short Wave Channel

To simulate what it would be like for JWST's NIRCam to observe the calibration field during commissioning, we assumed an arbitrary PA_V3 orientation for the telescope of −75 degrees. We then determined which star in the penumbra of the catalog in Figure 3 could be put in the center of FGS2 in order to place NIRCam's A module near the center of the calibration field.

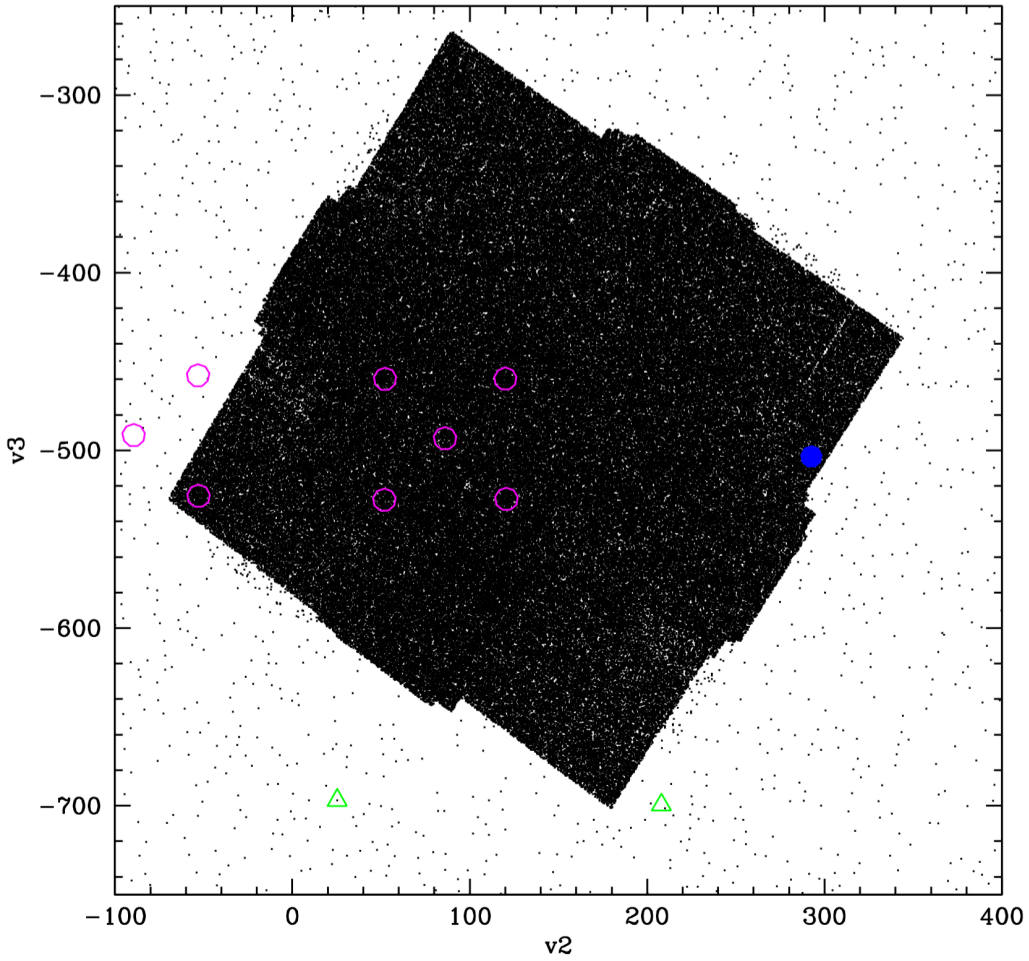

Figure 4: The JWST Calibration Field catalog along with the supplemental 2MASS penumbral stars projected into the (*v2*, *v3*) system as described in the text. The two green triangles correspond to the two FGSs (the guiding star was placed at the center of the left FGS, FGS2). The blue dot at (300,−500) corresponds to the center of the NIRSpec FOV and the open cyan circles to the NIRCam detectors (8 out of 10 shown; the target detector is in the lower right of the NIRCam points). NIRISS and MIRI are off the plot. Units are in arcseconds.

We simulated a pointing of JWST with the J=16 2MASS star at (80.63091, −69.54348) at the center of the FGS at (v2, v3) = (25.35″, −696.96″). To do this, we used standard spherical trigonometry to re-project the 2MASS catalog about the star's (α,δ) coordinate and rotated the projected coordinates by −75° (the assumed PA_V3). We then re-projected the transformed coordinates about (25.35″, −696.96″) to determine a (*v2*, *v3*) position for every star. Figure 4 shows the catalog in this (*v2*, *v3*) frame.

We then took the reference position for this particular detector from the SIAF file (*v2* = 120.6714, *v3* = −527.3827) and reprojected the full catalog about that point to get (Δu, Δv) for each star in the detector's ideal frame. Figure 5 shows all the catalog stars, with the stars that fall within the detector field of view highlighted in red.

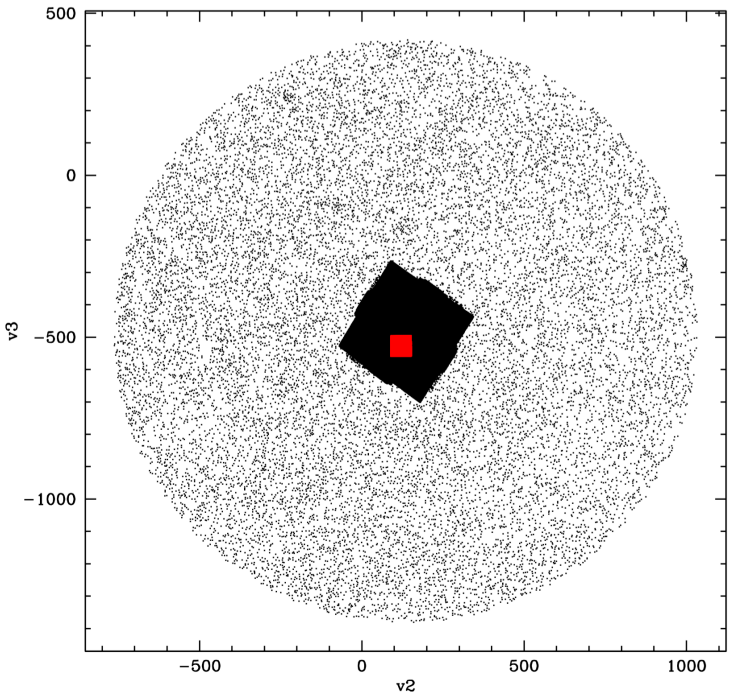

Figure 5: The projected catalog for the NIRCam-A1 detector, with the stars highlighted in red that land within the detector FOV. Units are in arcseconds.

Finally, we took the ideal-to-science transformation from the polynomial in the SIAF file and determined a position for each star in the science frame, which for this detector is identical to the detector frame. The polynomial contains an off-axis linear distortion term that corresponds to 0.05″ from one corner of the field to another. The second-order terms have amplitudes ranging from 0.03″ to 0.2″. The third-order terms range from 1.1 mas to 12 mas. The fourth-order terms are all below 0.13 mas, and the fifth-order terms are below 0.3 mas. It is clear that we could meet our 5-mas spec even if just considering the correction down to third order.

With these projected positions, we then simulated observations of the NIRCam detector with the F140W PSF from WebbPSF (Perrin *et al*. 2014). Using a NIRCam sensitivity spreadsheet, we found that a K=16.8 star should have 40,000 counts in a 10-exposure. So we simulated a 10-s exposure based on the K-band photometry in the calibration field list. We simulated both readnoise and a background of 100 electrons. We did not simulate flat-field errors for any of the detectors. Astrometric measurements are insensitive to L-flat-type errors, and pixel-to-pixel flat field errors of 1% will introduce at most a 0.01-pixel error in the astrometry. Such errors should be negligible for the purposes of calibration to 5 mas. Figure 6 shows the constructed field.

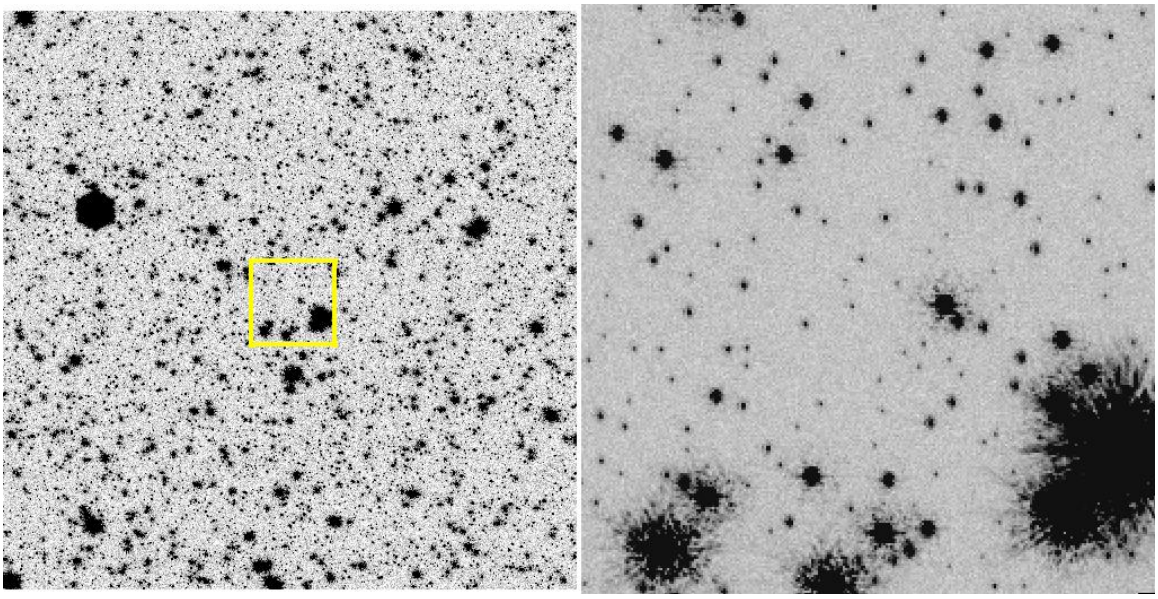

Figure 6: (Left) the full simulated 2048×2048-pixel field of the A1 chip of the NIRCam Short-Wave Channel. (Right) a close up of the highlighted region.

We ran a simple finding algorithm on the image, identifying all the unsaturated stars that had more than 1000 counts in their brightest pixels and were brighter than any neighbors within 20 pixels. We measured fluxes by simple 3×3-pixel aperture photometry (over sky) and obtained the (*x*, *y*) positions using simple centroids[1]. We blindly cross-identified this starlist with the input list, thus allowing us to associate a measured (x_{RAW}, y_{RAW}) position with a (Δu, Δv) position in the ideal frame.

This data is identical to what we will have during commissioning. If we take an observation with a particular guidestar at a particular location in an FGS, we can project the catalog in (*v2*, *v3*) space, then do a tangent-plane projection about the reference point to get accurate positions in the distortion-corrected "ideal" frame.

So, with these pairs of associated positions (*x*, *y*; *Δu*, *Δv*) we first need to purge the mis-associated positions (sometimes when cross-matching stars in two lists, there are sometimes two very close stars of similar brightness, and the wrong stars sometimes get associated with one another). We developed a simple procedure that went through the list star by star and identified the closest 25 neighbors for each star and used them to predicted the star's (*x*, *y*) position based on the star's (*Δu*, *Δv*) position and a linear transformation based on the neighbors' (*x*, *y*; *Δu*, *Δv*)

[1] Note that when the PSF is undersampled, centroid-measured positions are often biased with respect to the pixel phase of the star. The assumption inherent in the centroid is that flux transfer is linear with position shift, however, in the presence of undersampling the transfer of flux can be highly non-linear. These systematic errors are typically 0.01 to 0.05 pixel and will not affect our ability to achieve a chip-wide distortion solution good to 5 mas. It is also worth noting that variations in the intra-pixel sensitivity (Π) can also affect centroids. The WFC3/IR camera has a H1RG detector and a recent ISR (Anderson 2016) shows that Π is within one percent of being flat-and-flush.

position pairs. This local-transformation-based approach was necessary because we have not yet solved for the distortion and we do not want uncorrected distortion to impact the purging. We iteratively removed the worst association until the largest offset was less than 4.5 σ. Typically there were 5-10 stars removed this way. Note that if this purge stage is applied to the actual data during commissioning, it would remove any high proper-motion stars that are significantly displaced relative to their catalog positions.

The next step is to solve for the distortion solution. We put the pairs of associated positions into a least-squares routine that determined the 21 fifth-order coefficients to relate (Δu, Δv) to (x, y):

$$\Delta u = \sum_{i=0,5} \sum_{j=0,5-i} a_{ij} \hat{x}^i \hat{y}^j$$

$$\Delta v = \sum_{i=0,5} \sum_{j=0,5-i} b_{ij} \hat{x}^i \hat{y}^j,$$

Where $\hat{x} = (x - 1024.5)/1024$ and $\hat{y} = (y - 1024.5)/1024$. (For the FGSs and NIRISS, we used only fourth-order transformations.)

The next step is to evaluate the RMS error in the distortion correction. To do this, we determine the predicted (x,y) positions for a finely spaced array of (Δu, Δv) points using the official Idl2Sci mapping, then we use the (x,y) positions and our derived distortion solution to predict the input (Δu, Δv) values. The residuals are indicative of errors in the distortion correction. **The RMS in the recovered ideal positions was 0.099 mas in *x* and 0.087 mas in *y*.** One pixel in NIRCam's short wave detector is 32 mas, so this error corresponds to about 0.3% of a pixel. Since we had 1500 stars, each of which had a measured position good to about 0.03 pixel, it is not surprising that we were able to solve for the distortion to a factor of 10 times the single-star measurement precision.

This simulation is a clear demonstration that a single observation of the calibration field should be able to calibrate the distortion in NIRCam with an accuracy of much better than 5 mas during commissioning. Although we have solved for the distortion here, we have not yet solved for the SIAF parameters exactly[2]. If the mapping of the NIRCam reference point or the orientation of the detector in the (v2, v3) plane are not exactly as specified in the SIAF file, then our projection of the catalog into the ideal frame will be slightly off. This would mean that our distortion polynomials would a different center-point than those in the SIAF. It would be trivial to converge on the full SIAF solution through iteration, however that is beyond the scope of this document: the limited goal here is simply to calibrate the distortion *within* each SI.

We ran the same procedure on the A2 chip and found the distortion solution to be similarly well

[2] There is not a single, unique distortion solution that describes a detector. There is a four-parameter family of distortion solutions: distortion solutions can be transformed into another frame conformal transformations (2-D offset, rotation, and scale) without introducing distortion.

calibrated. Since the polynomial solution and the star density should be similar in all the NIRCam short-wave chips, it seems clear that we will be able to solve for the intra-SI distortion for all of them to well within the 5 mas requirement.

3.2 NIRCam's A5 Detector in the Long Wave Channel

NIRCam's LWC (Long-Wave Channel) has a different plate scale and a different PSF from the SWC, so it makes sense to validate it separately. We used the same projected (*v2*, *v3*) catalog as in the previous sub-section and mapped each star into the A5 ideal frame and identified which stars landed within the detector boundary. Figure 7 shows the stars that land in this detector footprint.

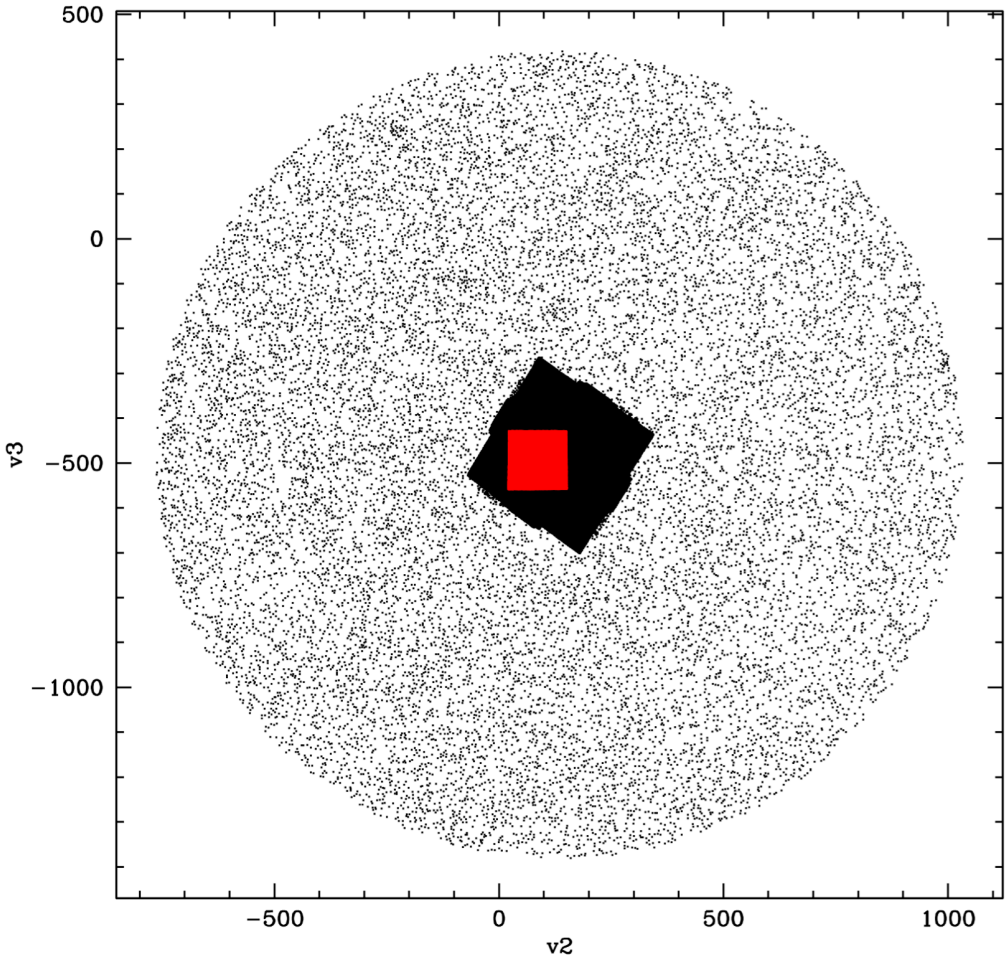

Figure 7: This shows the (v2,v3) projection for the catalog with the pointing appropriate for the NIRCam LWC Module A calibration. Units are in arcseconds.

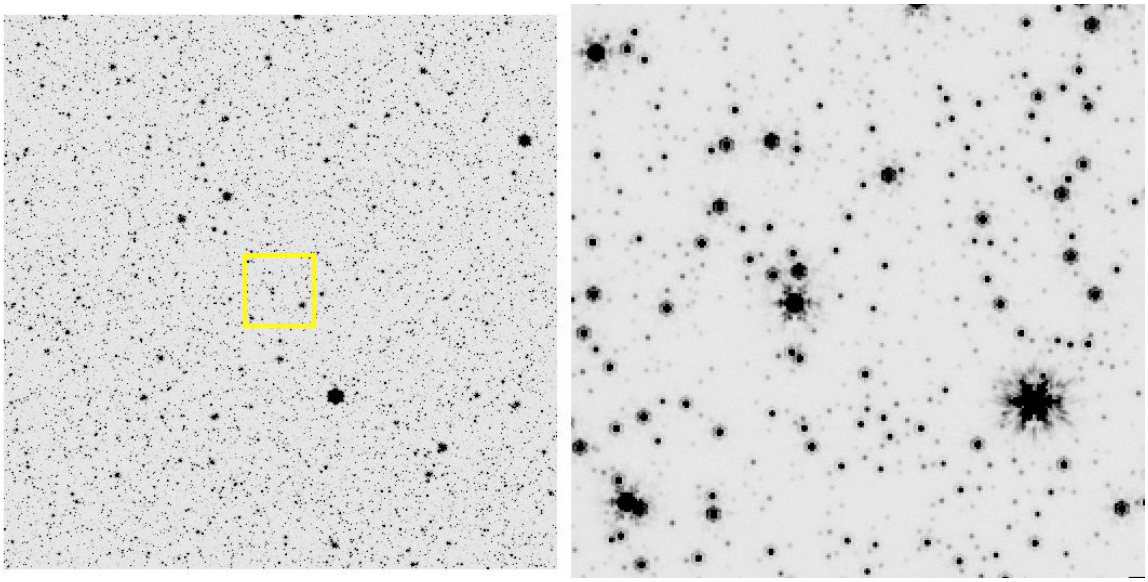

Figure 8: (Left) Full-frame of the Module-A LWC detector. (Right) 8x close-up of the boxed region, showing the star size.

As before, we simulated the field with a WebbPSF (using the one for the F356W filter), adding noise and a background of 100 electrons. We measured these stars using the same procedure as above: simple sky-subtracted aperture photometry and centroid positions. The "found" star list was cross-identified with the catalog and purged, as above for the A1 detector. This again gave us a list of associated positions ($x, y; \Delta u, \Delta v$), which allowed us to solve for the distortion solution, in a manner identical to that above.

We tested this distortion solution by comparing the mapping from a dense grid of points in ($\Delta u, \Delta v$) space to (x, y) using the original Idl2Sci mapping, and then mapped back to the ideal frame using our polynomials to get ($\Delta u', \Delta v'$). **We then compared our distortion corrected positions against the original positions and found that the RMS error in the mapping was about 0.24 mas, again well below the 5 mas threshold.**

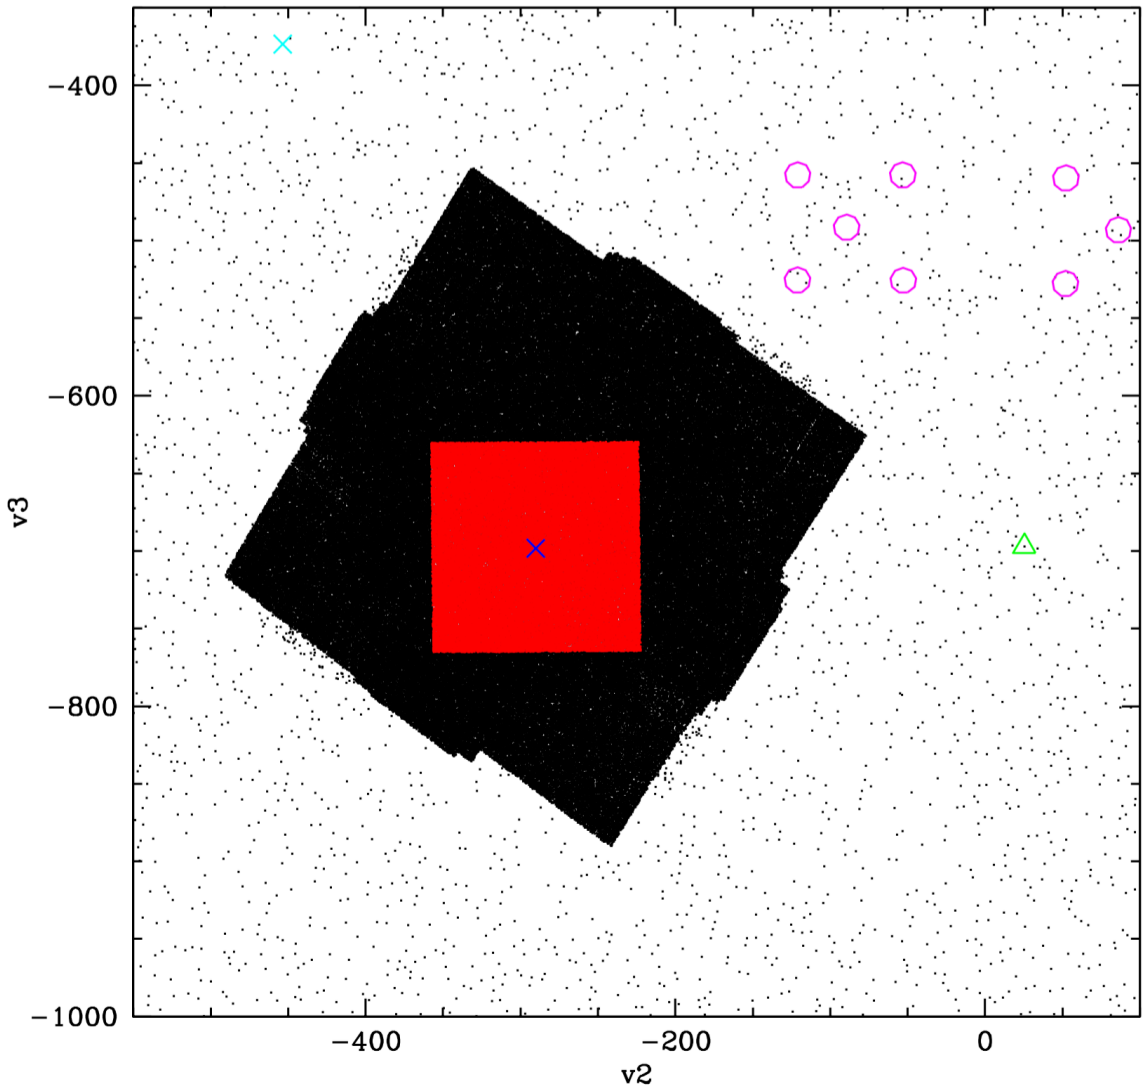

Figure 9: This shows the (v2,v3) projection for a pointing that puts a bright star on FGS2 (open triangle) and the NIRISS reference point (blue square) at the center of the calibration field. The red points are the stars that land within the NIRISS detector field of view. Units are in arcseconds.

3.3 The NIRISS Detector

The NIRISS detector covers the same area as the NIRCam long-wave channel, but can observe much shorter wavelengths, therefore the detector is considerably undersampled. Figure 9 shows the choice of a guidestar that places NIRISS in the calibration field for the chosen orientation (PA_V3 = −70°).

We used the PSF from WebbPSF for NIRISS's F200W filter to simulate the images. Again, we used the `Idl2Sci` transformations in the SIAF file to convert positions in the ideal frame into the detector frame. The polynomial for NIRISS goes down only to fourth order, so there were only 15 terms in the forward transformation.

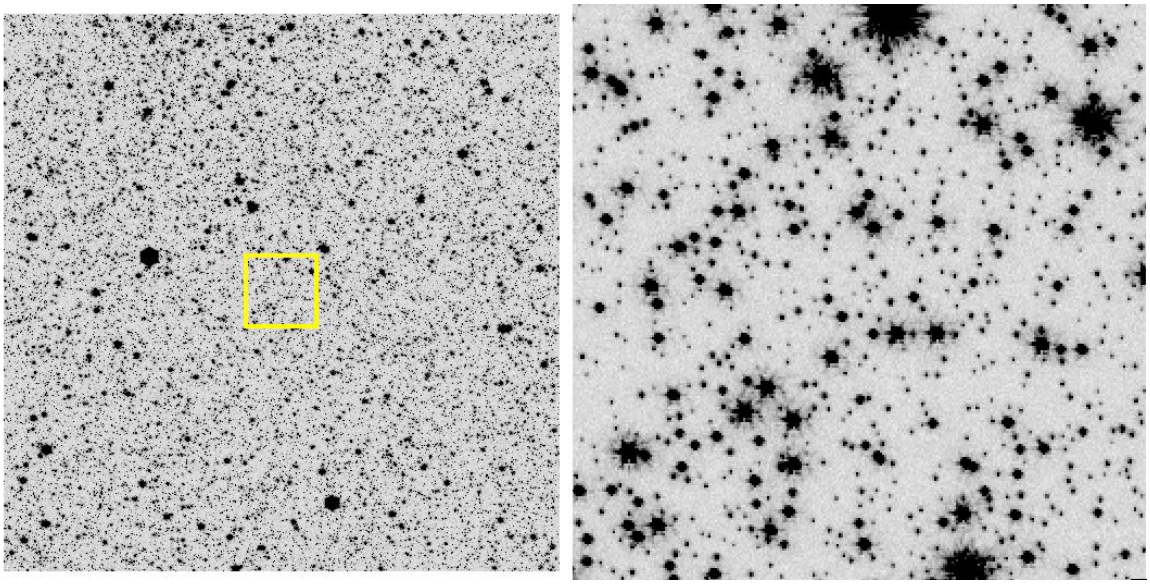

Figure 10: The NIRISS field.

We ran the same finding algorithm on the image and came up with a star list with positions and fluxes. This list was again blindly cross-identified with the input list, thus giving us pairs of associated positions (x, y; Δu, Δv), where (x,y) is measured and (Δu, Δv) comes from the catalog. We again purged this list of the inconsistent associations and ended up with 2858 stars. We then fit these stars with a fourth-order polynomial and again compared the original (Δu, Δv) positions against the ($\Delta u'$, $\Delta v'$) positions from the model. **We found the residuals to be 0.56 mas in *x* and 0.38 mas in *y*, again clearly adequate to meet the formal requirements.**

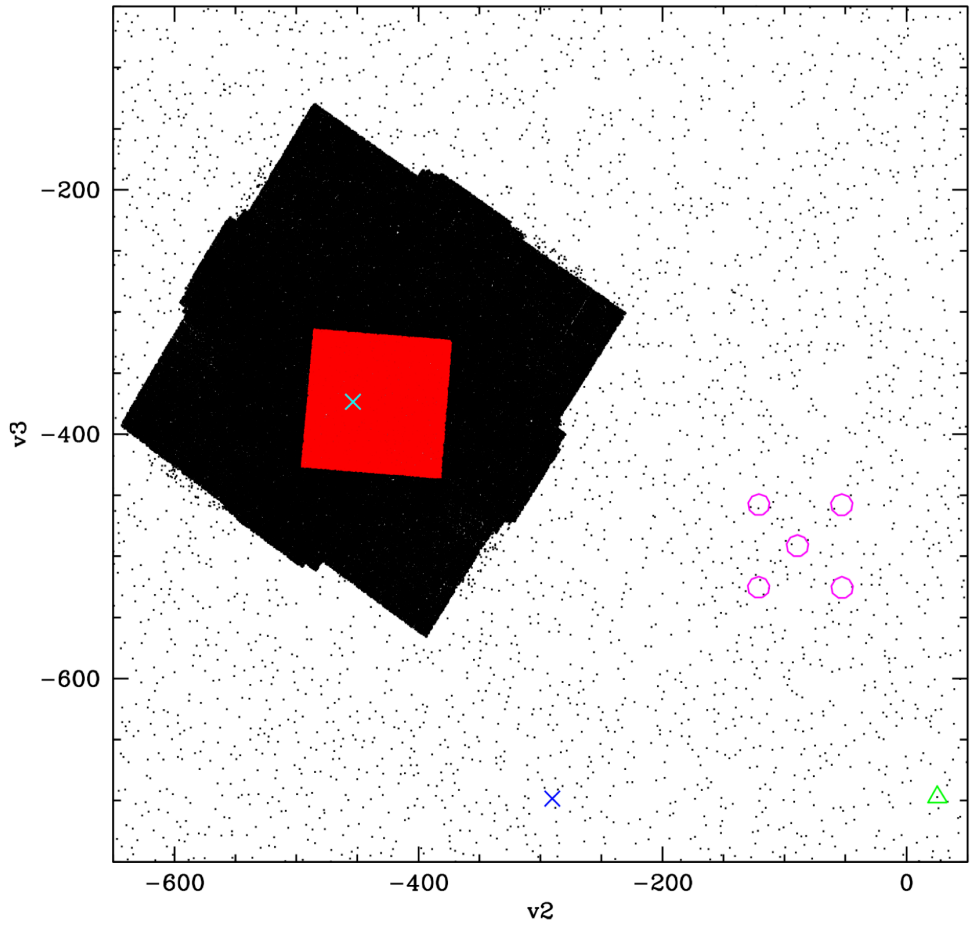

Figure 11: This shows the (*v2*, *v3*) projection for a catalog that places a bright star at the center of the FGS2 guider and MIRI at the center of the calibration field. Note that the MIRI reference point (cyan cross) is off-center from the center of the detector. This is because the SIAF file has the reference point at (693, 512), for the 1032×1024 detector, since a portion of the detector is obscured by the metering structure. Units are in arcseconds.

3.4 The MIRI Detector

MIRI is designed to probe wavelengths from 5μm to 25μm. MIRI is the only instrument that does not have a H2RG detector; its detector size is 1032×1024, rather than 2048×2048. As we did for the other detectors, we identified a star to place at center of the FGS so that the main MIRI field would be in the center of the calibration field. Figure 11 shows this. Note that here we are calibrating only MIRI MRS imaging mode.

We projected the catalog about the reference point to determine an ideal coordinate for each star and determined which stars were in the MIRI FOV. For each of those stars, we used the polynomial to convert the positions into detector coordinates and used the WebbPSF model for MIRI's F560W to simulate the images. As before, we ran an automated finding routine and identified a list of stars and have for each a centroid position and a flux. We again cross-identified these stars with the catalog, purging the inconsistent associations, and came up with a list of associated pairs of positions (x, y; Δu, Δv).

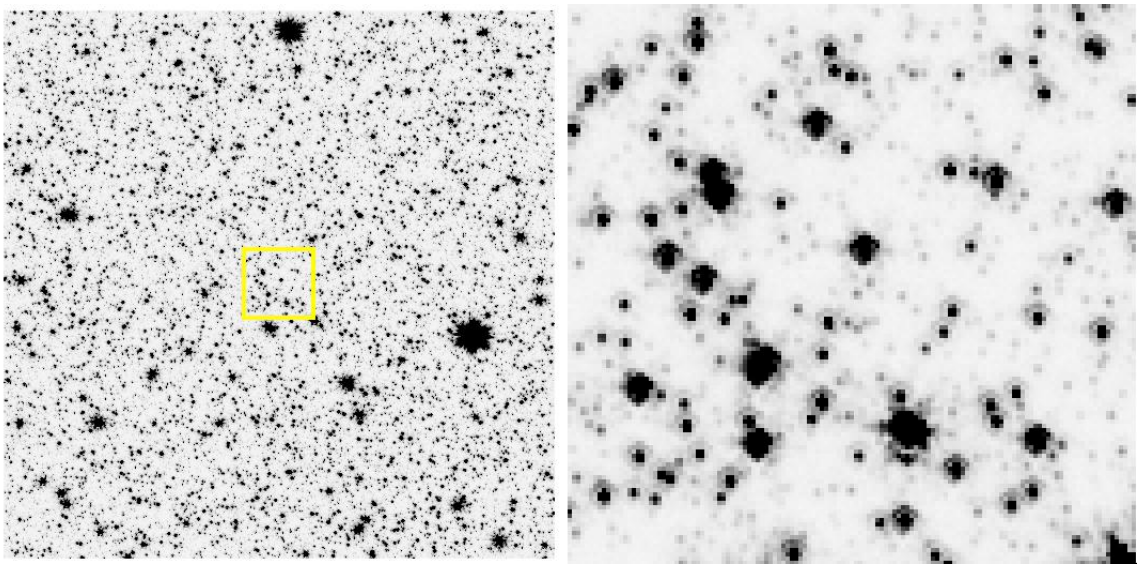

Figure 12: (Left) The full, simulated field of MIRI, (Right) ×8 close up of the center of the field.

There were only 660 stars this time, since the MIRI detector has four times fewer pixels than the other detectors, and our finding identified only extremely isolated stars.

Again, we fit a polynomial to the observed-to-catalog mapping to model the distortion. **We evaluated the quality of the correction for a finely spaced array of points across the detector and found it to be good to 1.4 mas in *x* and 2.6 mas in *y*.** It makes sense that the solution it would be worse in MIRI than in NIRCAM and NIRISS, since the PSF is broader and we have 4× fewer stars to use to constrain a fifth-order polynomial. Even so, we should be able to take a single observation of the calibration field with MIRI and use the catalog to back out the distortion solution.

The FGS Detectors

The FGS guiders will also need to have their distortions calibrated so that we can know the (*v2*, *v3*) focal-plane positions of stars that are observed at non-reference locations in the detectors. Figure 13 shows how we can place a bright star from the calibration field at the center of FGS2 to get the entire FGS1 detector within the main calibration field.

As before, we projected the (*v2*, *v3*) positions into the FGS1 frame and used the `Idl2Sci` transformations to determine, for the stars inside the FOV, where each star in the catalog would land on the detector. We then simulated an image as before. Since WebbPSF does not have any PSFs for the guiders, we simply used the NIRISS PSF, since its platescale is similar. The particular PSF used should have a negligible impact on the final precision, since we are using simple centroids to measure positions for stars.

We measured stars as before and cross-identified them with the catalog and purged the inconsistent matches. There were 2896 stars. Again we fit a polynomial to the observed (*x*, *y*) and catalog (Δu, Δv) positions in the ideal frame to empirically measure the distortion. **We then transformed a fine mesh of points in the (Δu, Δv) frame into the detector frame, and used the solved-for polynomial to transform them back and found that the residuals to be 0.4 mas in *x* and 0.6 mas in *y*, well below the requirement.**

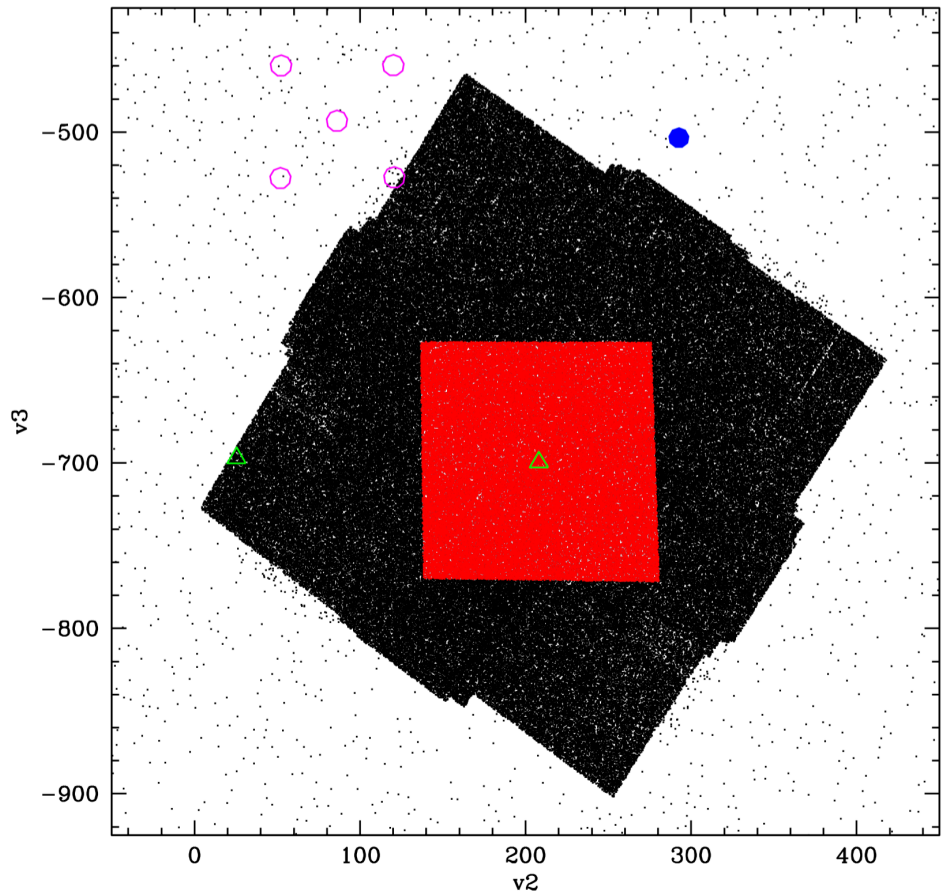

Figure 13: A bright stars is placed at the center of FGS2 (left triangle) and FGS1 images the calibration field.

Units are in arcseconds.

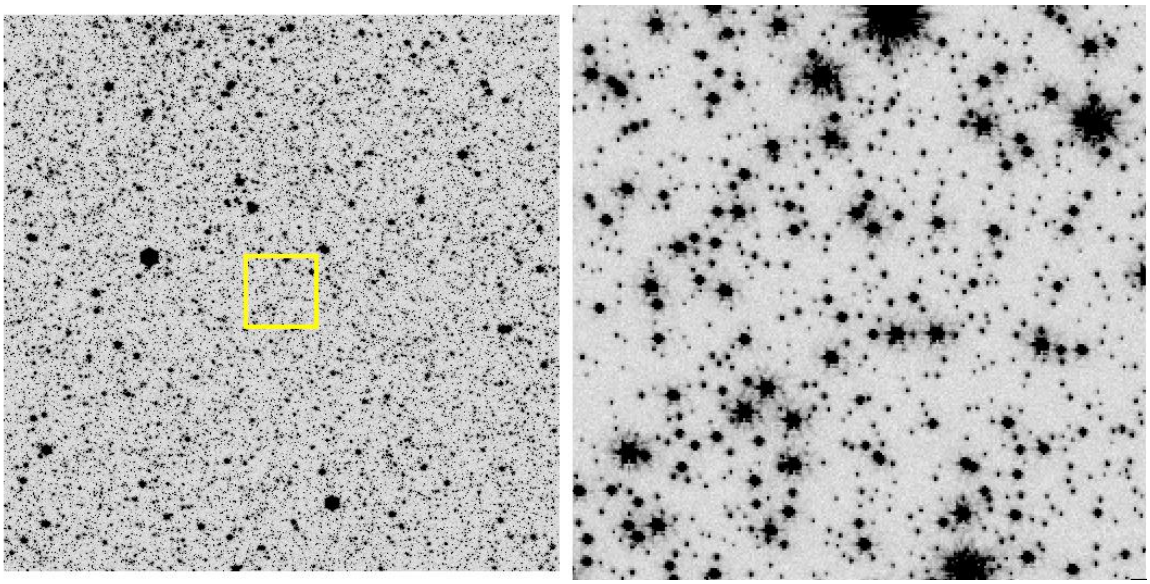

Figure 14: (Left) The entire FGS1 detector centered on the calibration field. (Right) a ×8 close-up of the central region.

3.5 The NIRSpec Detector

NIRSpec is much more complicated to deal with than the other detectors, primarily because its main focus is not imaging, but spectroscopy. Although it has an imaging mode, the imaging is done through the MSA. In imaging mode, the MSA allows light through to only part of the detector, and even in the open portions of the field, it imprints a zero-throughput grating on the field. In addition to having an obstruction in front of its imager, NIRSpec also has a more complicated distortion-solution formulation than the other instruments. I discuss all three of these issues in turn before turning to the solution demonstration.

3.5.1 The MSA

The MSA has a substantial effect on the imaging mode of NIRSpec. In addition to blocking some fraction of the light, there is also some diffraction around the shutter edges. In what follows, I will simulate the blockage of light by the shutters and will show that it is consistent with ground tests of impact of the MSA on astrometry.

To deal with the MSA, I generated a mask image that had 10× the pixel scale of NIRSpec so that I could simulate the sharp edges of the mask relative to the oversized 100-mas NIRSpec pixels. The total extent of each slit is 260 mas horizontally and 510 mas vertically, with 30 mas obstruction all around[3]. Figure 15 shows a close-up cartoon of the MSA along with the pixel grid.

[3] These numbers were used in the simulations. More recent measurements show that the actual shutter pitch is 260 mas along the dispersion axis (as used) but is 520 mas along the spatial axis, which is off by 10 mas from the 510 mas used here. This slight difference should not affect the validity of the simulations.

To validate our MSA simulation, we simulated a 10× supersampled NIRSpec scene with a F110W NIRSpec WebbPSF and integrated it over pixels, multiplying it by the MSA mask. I measured centroid positions for stars in a manner similar to that for the previous detectors and compared them with the original positions. Figure 16 shows the results. There is an elongation in the residual distribution, largely in the *x* direction. The *x* dispersion appears to be about 0.2 pixel in amplitude and that in the *y* direction is about 0.1 pixel. This is not surprising given that there are twice as many MSA boundaries in the *x* direction than in the *y* direction.

This simulation addresses the blocking effect of the MSA, but not the diffraction. To ensure that diffraction doesn't make our positions considerably worse, we analyzed a pair of images taken during Cryo Vac 3 testing at Goddard. One exposure was taken with a simulated set of about 200 point sources at one dither pointing, then the detector was shifted by (1.3, 2.5) pixels for the second exposure. I ran the same position-measuring software on both images and examined the difference, modulo the offset. I found the residuals to be about 0.2 pixel along one axis and about 0.08 pixel along another axis, extremely similar — both in amplitude and in nature — to those seen in Figure 16. For this reason, we can be confident that our MSA simulation is not qualitatively or quantitatively different from reality.

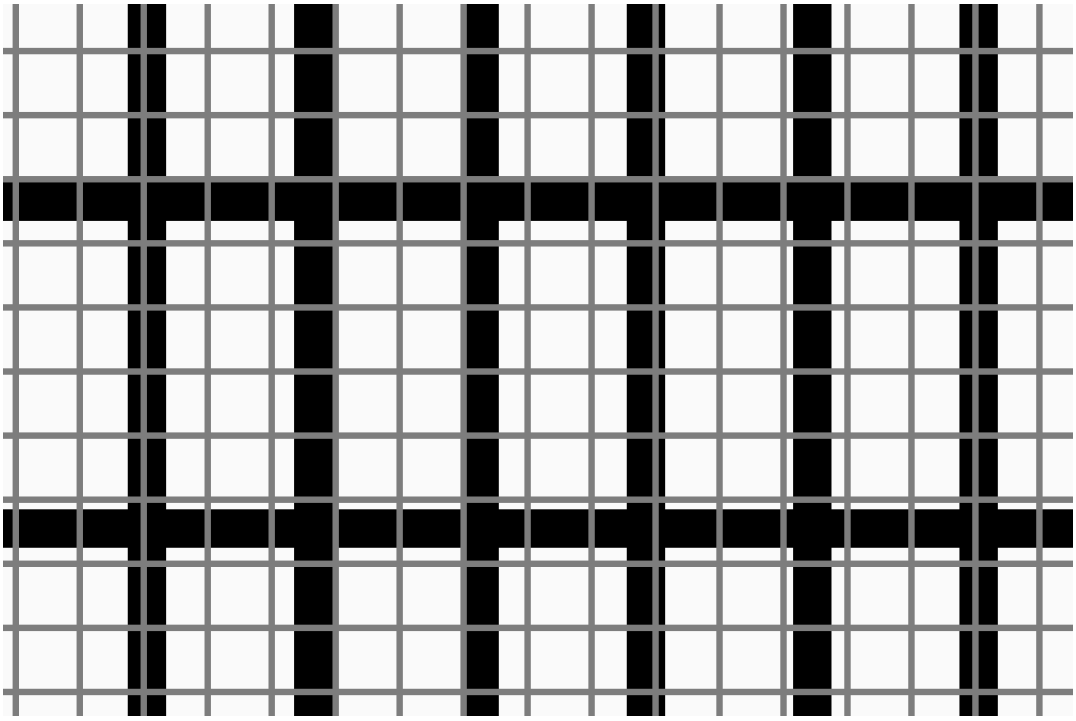

Figure 15: A cartoon showing our MSA model. The MSA shutters (black) are about 260 mas × 510 mas and the pixels (shown in gray) are about 100 mas × 100 mas.

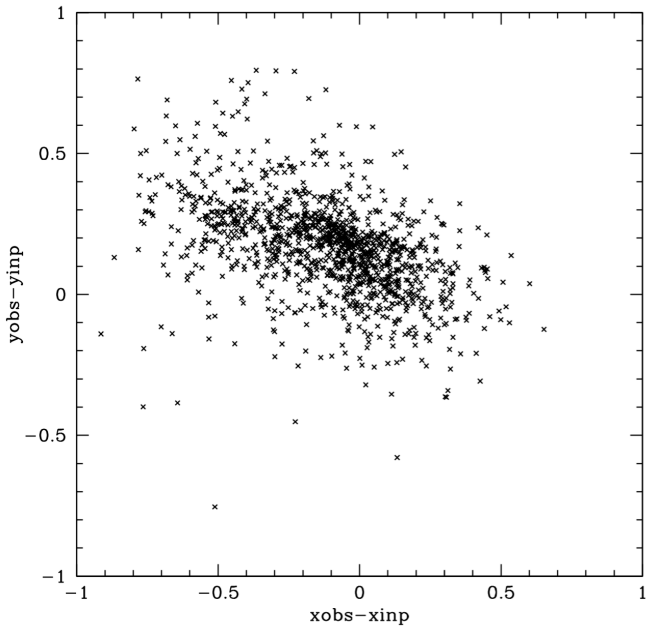

Figure 16: The difference between input and output position measurements to show the impact of the MSA on centroids. The offset from (0,0) is simply related to asymmetries in the PSF model and is not indicative of errors in position. The MSA introduces about twice as much dispersion in the *x* positions as in the *y* positions. This is not surprising given that there are twice as many shutter boundaries along *x* as along y.

3.5.2 Limited field

Since NIRSpec's main purpose is to be a spectrograph and disperse light horizontally across the detector, the imaging field is somewhat limited in extent. In fact, in imaging mode more than half of the detector is occulted: the outer half of each detector is receives no light and there is a large gap between the upper and lower MSAs for each detector. We have simulated this in our mask for each detector. Figure 17 shows the mask for the left NIRSpec chip. We simulated the MSAs as having a small orientation with respect to each other just to ensure that such tweaks to perfection would not cause problems in the distortion solution.

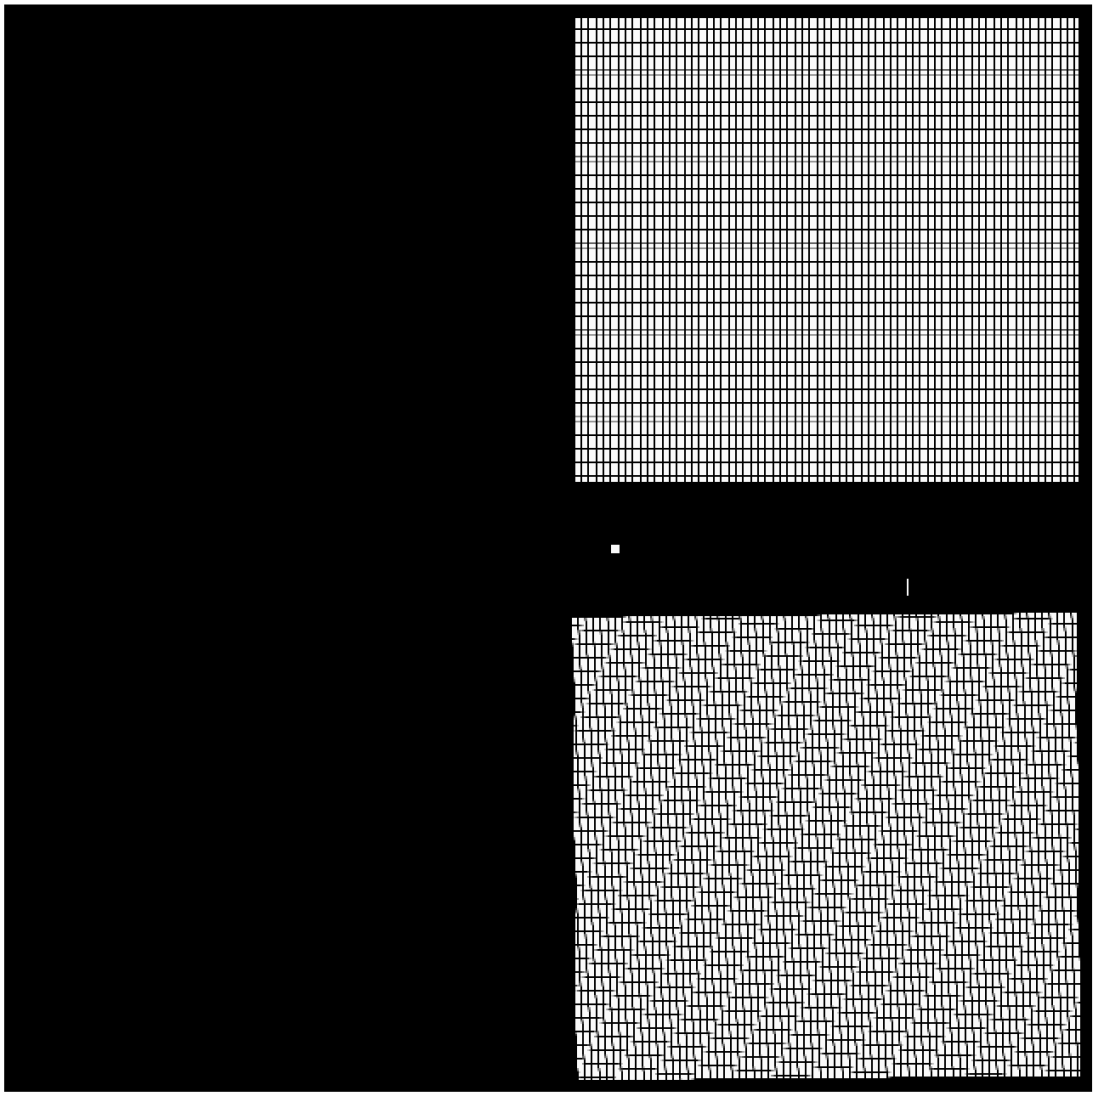

Figure 17: The full mask used for the left NIRSpec simulation, covering the 2048×2048 detector. The mask allows light through to only a part of the detector in imaging mode. The lower MSA is shown rotated very slightly with respect to the detector.

3.5.3 The net distortion solution

As mentioned above, NIRSpec also has a more complicated distortion-solution formulation than the other instruments. The SIAF files at present contain only part of the distortion-solution model: the path from the telescope to the first grating.

The post-grating half of the solution depends on the grating angle. In imaging mode, this angle should be fixed, such that the distortion should be constant (apart from a small dependence on filter coming from the FORE optics). The second half of the solution is provided in some NIRSpec documents (Giardino et al. 2014) and has been distilled into some preliminary working papers by the Telescopes group for three different filters (Colin Cox, personal communication).

The complicated multi-stage solution above can easily be validated by observation, but it would be very hard to reconstruct it from only a single set of on-sky observations. The reason for this is that the sky-to-grating transformation involves a fifth-order polynomial mapping, similar to what is done for the ideal-to-detector mappings for the other detectors. At the grating, there is a reflection at an angle, where the positions are reversed and transformed in a linear way before serving as the basis for a second fifth-order-polynomial transformation.

It would be extremely hard to tease out individually the two nested polynomials from observations of stars on the detector. Thankfully, all of the imaging will be done at a single setting of the grating and it should be possible to treat the entire transformation from sky to detector as a single mapping. This is the approach shown here to demonstrate that there will be enough stars to constrain the mapping. Once JWST is in orbit, in order to derive a common set of distortion polynomials for both imaging and spectroscopy, these on-sky observations will be combined with data acquired with NIRSpec's internal lamps, which is used to solve for the distortion from the detectors to the MSA plane. This includes the calibration to derive the exact position of the grating wheel. Once that part of the NIRSpec instrument is characterized, the distortion solution for the MSA-to-(v2,v3) mapping will be derived.

I have codified the multi-part solution for NIRSpec's F110W filter into a single mapping from the (*v2*, *v3*) plane to the detector. Using this mapping, I determined for every pixel in the detector its location is in the (*v2*, *v3*) plane. I created two 2048×2048 images (one for *v2* and one for *v3*) that contain the focal-plane location to which each pixel [*i*, *j*] maps. This also permits us to use an iterative process to determine the inverse mapping from (v2, v3) to detector (*x*, *y*). In some ways, this makes the NIRSpec simulation easier than the other detectors, since there is no need to deal with "ideal" frames and "detector" frames to generate the simulations.

3.5.4 The image simulations and analysis

As before for the other instruments, we determine which guidestar we could put at the center of FGS2 to place the NIRSpec FOV into the high-precision part of the calibration field. Figure 18 shows the pointing. The yellow points *map* onto the detector, but the red points actually *make it through* one of the MSAs to hit the detector.

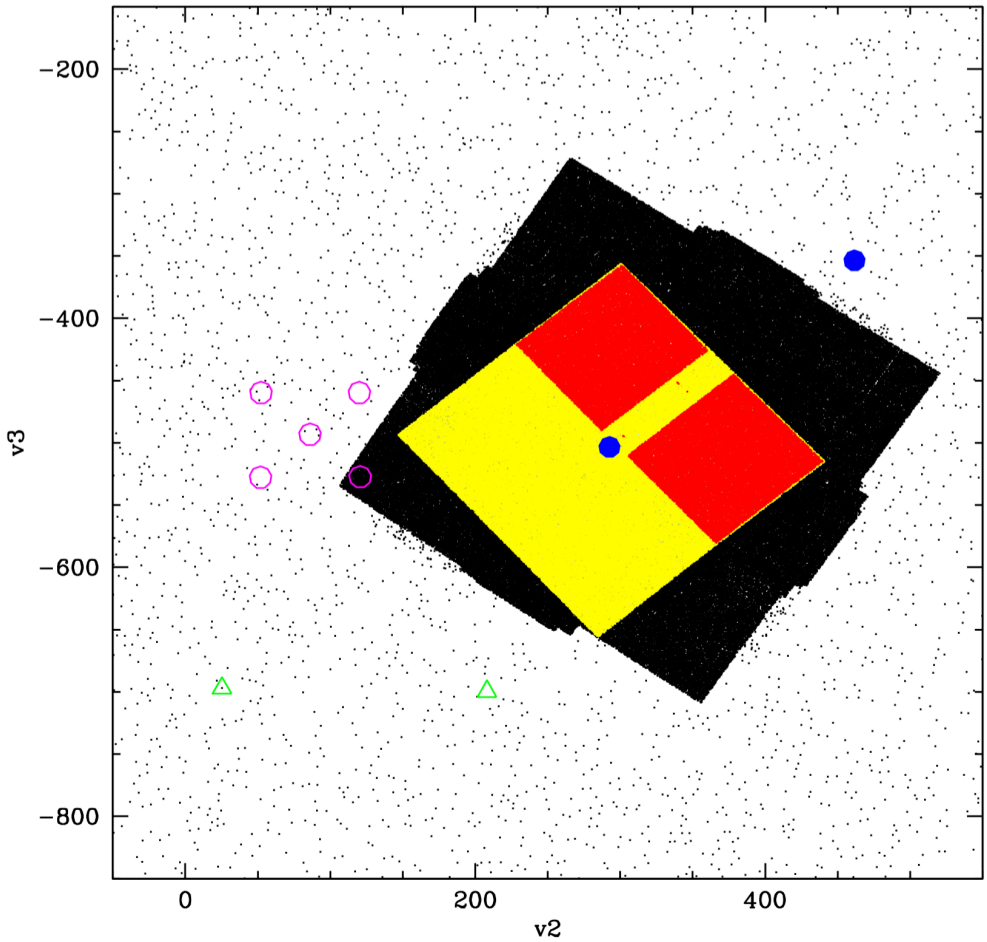

Figure 18: The mapping of the catalog into the (v2,v3) plane to put the left NIRSpec chip near the center of the calibration field. FGS2 (the left green triangle) is centered on a bright star in the 2MASS-filled penumbral region. The blue dot shows the center of the NIRSpec1 frame. The yellow points are the points that are formally within the detector are blocked and do not go through an MSA. The red points show the stars that should make it through the MSA (however imperfectly) to register on the detector. Units are in arcseconds.

We adopted a PSF from WebbPSF for NIRSpec's F110W filter (which is the filter corresponding to the post-grating distortion model we used) and simulated a NIRSpec observation. We super-sampled the scene by a factor of 10 and multiplied the scene by the MSA mask (which was 1 or 0 for any 0.1×0.1 pixel). We then integrated this super-sampled scene over the 10×10 sub-pixels to determine the value for every actual NIRSpec pixel.

Figure 19 shows the full image that was simulated. The plot on the left shows what would happen if all pixels could see the sky, and the middle plot shows the impact of the MSA and vignetting. The mask used is shown on the right.

Figure 20 shows close ups of all these things. The range of MSA throughput values goes from 25% to 100% transmission for each detector pixel (though the actual integration over the MSA is done at ×10 resolution).

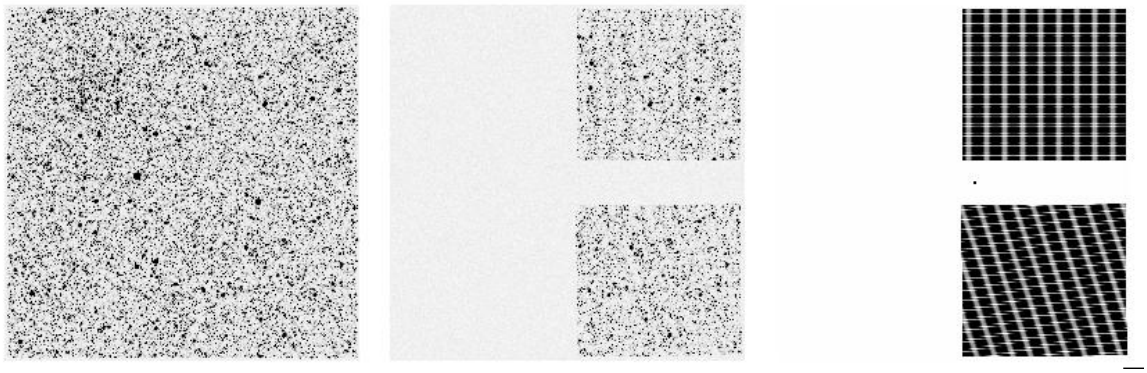

Figure 19: (Left) mapping of the calibration field onto the entire detecor. (Middle) After multiplication by the MSA-mask-throughput. Much of the detector does not get any direct photons from the sky. (Right) The MSA mask (the striping comes from aliasing in the display program).

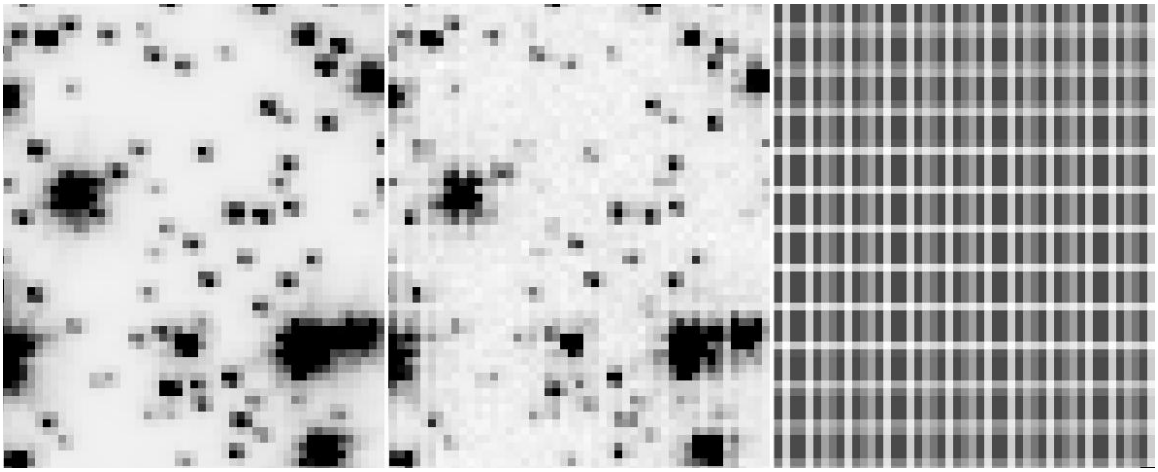

Figure 20: Close up of a portion of the images shown in Figure 18. The left image shows the "truth" super-sampled image integrated over NIRSpec pixels. The middle shows the same things but integrated over the MSA. The right image shows the integral of the MSA throughput for each NIRSpec pixel. The throughput of a pixel ranges from about 25% to 100%.

Now that we have all the elements of the simulation together: the star-field mapping to the detector, the modeling of the PSF, and the impact of the MSA shutters on the scene, we are finally in a position to evaluate how well we can measure the distortion. For the previous instruments, we were able to transform from the (*v2*, *v3*) frame to the official "ideal" frame for the detector and solve for the distortion between the ideal frame and the science frame, a solution that is largely just a scale change and distortion correction with very little rotation. Here, since we do not have an ideal frame, we will go directly from the (*v2*, *v3*) frame to the detector frame to show that some solution is possible. At some later point, this distortion solution can be put into a different format, but that will just involve a conformal transformation (offset, scale, and rotation).

With all this in hand, we went through the middle image in Figure 19 and applied the same finding algorithm that we had used previously, namely we identified a potential source to be in every pixel that was unsaturated, have more than 1000 counts over sky in its central pixel, and

had no brighter pixels with a radius of 20 pixels. We measured the position for every source with a simple centroid algorithm and the flux using a simple 3×3-pixel aperture. There were 1255 such sources. We then cross-identified the found stars with the input catalog and identified a potential match every time a catalog star had an observed source within 2.5 pixels. This gave us 1272 stars (some of which were clearly mis- or duplicate identifications). We next applied the same local purging algorithm as before and weeded the list down to 1248 stars, each of which gives us an observed position on the detector (x_{obs}, y_{obs}) and a known position in the telescope frame (*v2*, *v3*).

Since the imaging field does not cover the entire detector, we decided to solve for a polynomial using a more appropriate basis function. In our simulation, the *x* coordinate went from 1070 to 2023 and the *y* coordinate went from 25 to 2022, so we parameterized the polynomial with dimensionless parameters $\hat{x} \equiv (x-1547)/476$ and $\hat{y} \equiv (y-1547)/476$. By parameterizing the distortion this way, the amplitude of the various terms automatically tells us their maximum impact and found the polynomial that best mapped ($\hat{x}$, $\hat{y}$) to (*v2*, *v3*), where *v2* and *v3* parameterize the focal plane in arcseconds. The zero-order terms tell us the location of the center of the imaging area in the (v2, v3) plane. The first-order terms tell us the scale, orientation and any linear distortions (such as a non-orthorgonality of the axes, or a difference in scale between them). The higher order terms are all distortion. The second-order terms have amplitude of about 1 arcsecond (10 pixels). The third-order terms have an amplitude of about 0.04 arcsecond (0.4 pixel). The fourth-order terms have amplitudes of about 0.03 arcsecond (0.3 pixel). Finally, the fifth-order terms have amplitudes of 0.015 arcseconds (0.15 pixel). (Note that these polynomials are not Legendre-type polynomials, so the high-order terms have almost all of their impact at the edges of the field, so they are not indicative of the typical impact of these polynomials.)

We evaluated the quality of the distortion solution as before. We determined the mapping from (*x*, *y*) to (*v2*, *v3*) for a fine mesh of points covering the usable region of the detector using the current best-estimate ray-trace model and using our polynomial. **We found that our polynomial provided a distortion-corrected mapping from (*x*, *y*) to (*v2*, *v3*) with an RMS error of 1.6 mas in v2 and 3.1 mas in v3, well below the 5 mas requirement.** This fit was based on a single short exposure of the field.

We note that we have not simulated the impact of failed-shut MSA shutters or detector defects, so there are probably a few stars in our lists that would not in practice be available. This should not significantly affect our ability to meet the requirement, especially given that during commissioning, the NIRSpec team will likely take many dithered exposures of the calibration field to validate and improve on the solution. This investigation here was designed to show that even with just one observation, we should be able to measure the distortion to the required precision.

4 Conclusions

We have shown that a single short exposure of the JWST calibration field in the LMC should be adequate for solving for the distortion to better than 5 mas per coordinate in all of the detectors (NIRCam SWC and LWC, NIRISS, MIRI, the FGSs, and NIRSpec). Of course we will want to use more than one observation to validate the solution and improve it to considerably better than 5 mas, but this document clearly demonstrates that the minimum calibration requirements should be easy to achieve with a very simple imaging strategy.

We note that the current catalog of stars is based on observations taken in 2006, and by the time of commissioning in 2018, most of the stars will have moved by about 2.5 mas, each in an unknown direction. This small, unknown shift in positions will not prevent us from meeting the 5 mas per-coordinate accuracy in calibrating the internal distortion solution, since it will introduce a random error in each star. Nevertheless, having more precise positions could be useful if we desire a considerably more accurate solution or if we want to undertake more complex analyses (such as an immediate PSF reconstruction).

It is likely that a mosaic of NIRCam observations will enable us to construct a commissioning-epoch catalog soon after commissioning, but given how busy everyone will be, it might take a while to construct such a catalog. Some thought should be given as to whether a few HST orbits should be spent in the next year or so to allow us to update the catalog now, before JWST is launched. With two well-separated HST epochs, we can calibrate both the positions *and* the motions, such that the catalog should be valid for years to come.

5 References

Anderson, J. and King, I. R. 2003 PASP 115 113. *An Improved Distortion Solution for the Hubble Space Telescope's WFPC2*

Anderson, J. and King, I. R. ACS/ISR 2006-01. *PSFs, Photometry, and Astrometry for the ACS/WFC*

Anderson, J. and Diaz, R., JWST-STScI-002474, SM-12, *Validation of the Astrometry in the JWST Calibration Field*

Anderson, J., JWST-STScI-001378, SM-12, *IR Photometry of the JWST Calibration Field*

Anderson, J. 2008 JWST-STScI-001324, *A Hysteresis Effect in the Readout of the HAWAII-2RG Detectors*

Anderson, J. WFC3/ISR 2016-12, *Empirical Models for the WFC3/IR PSF*

Cox, C., Lallo, M., van der Marel, R. M., and Donaldson, T. JWST-STScI-001550, *Description and Use of the JWST Science Instrument Aperture File*

Giardino, G., Ferruit, P., Alves de Oliveira, C. NTN-2014-005. *Geometrical Transforms for NIRSpec Target Acquisition*

Kozhurina-Platais, V., Hammer, D., Dencheva, N, & Hack, W. WFC3/ISR 2013-14, *Astrometric Correction for WFC3/UVIS Lithographic-Mask Pattern*

Perrin, M. Sivaramakrishnan, A., Lajoie, C.-P., Elliot, E., Pueyo, L., Ravindranath, S, & Albert,

L. Proc. SPIE 9143. *Updated Point Spread Functions for JWST with Webb PSF*

Skrutskie, M. F., et al. 2006 AJ 111 1163 *The Two Micron All-Sky Survey (2MASS)*