

Space Telescope Science Institute
JAMES WEBB SPACE TELESCOPE MISSION
SCIENCE AND OPERATIONS CENTER

An Investigation of Optimal Dither Strategies for JWST

Authors:
Anton M. Koekemoer
Kevin Linday

Released: September 15, 2005

REVISION HISTORY

Revision	DESCRIPTION	DATE
-	Initial Release	29-Jul-05

Prepared By:

Electronic Signature
Anton M. Koekemoer

07/29/05

Electronic Signature
Kevin Lindsay

07/29/05

Reviewed by:

Electronic Signature
Gerald Kriss

08/09/05

Electronic Signature
Knox Long

08/12/05

Electronic Signature
Stefano Casertano

8/26/05

Electronic Signature
Rusty Whitman

8/22/05

Approved by:

Electronic Signature
David Hunter
JWST Project Manager

9/14/05

Electronic Signature
Peter Stockman
JWST Project Scientist

9/13/05

Space Telescope Science Institute
Baltimore, Maryland

An Investigation of Optimal Dither Strategies for JWST

Anton M. Koekemoer and Kevin Lindsay

Executive Summary

This study is aimed at an initial examination of defining and using metrics to describe and quantify the relative merits of different dither patterns for the JWST instruments, allowing the development of optimal patterns that trade off scientific benefits against costs such as increased overhead. An important component of this study is the detailed definition of quantitative metrics that are defined in a sufficiently general way that they can be used to evaluate the merits of any specific dither pattern or observing scenario, given sufficient details about the instrumental properties and the observing overheads.

The most detailed treatment in this study is for the NIRC*am* instrument, where the instrumental and observational characteristics are sufficiently well defined that simulated images could be produced for a range of different dither patterns and filters, and were subsequently combined using the MultiDrizzle program in the same way as is currently done for dithered HST observations. This allowed quantitative measurements of the properties of simulated objects on these images, specifically their morphological, astrometric and photometric properties, in a way that allowed direct comparison between different dither patterns and the ideal case of infinitely well sampled data. For a variety of assumed observing scenarios, drawn generally from considerations in the DRM, these measurements allow a quantitative comparison between the benefits obtained from dithering against the fractional increase in overhead time associated with each observing scenario. The general results are that, for NIRC*am*, short and medium observing programs may benefit most from 2 to 4-point sub-pixel dither patterns, while long observing programs may benefit more from 9-point sub-pixel dither patterns with a relatively small additional cost in overhead.

The other instruments are also considered in this study, although in a more generalized fashion since the available code was not yet capable of generating simulated data for these instruments. For the FGS-TF, similar considerations apply as for NIRC*am* except that the FGS-TF pixels are more undersampled at the short-wavelength end, therefore short programs may generally need at least a 4-point sub-pixel dither while both medium and long observing programs would likely benefit sufficiently from using 9-point sub-pixel patterns to offset the increased costs in overhead. The MIRI imaging mode is better sampled than NIRC*am* or FGS-TF, hence much of the discussion for MIRI instead centers around its IFU where it is possible to define a combination of slice widths and dither steps to provide optimal sub-pixel sampling in all four wavelength channels simultaneously with a single set of telescope offsets. NIRS*pec* is unique in that its MSA provides an additional set of requirements for dithering; the general motivation is to improve photometric fidelity by moving sources to a number of locations within the shutter apertures, to mitigate photometric uncertainties.

Table of Contents

1	Introduction.....	6
2	Scope of this Study.....	6
3	Scientific Drivers for Dithering	7
3.1	Improving the sub-pixel PSF sampling.....	8
3.2	Sampling vs. resolution.....	9
3.3	Improving photometry and astrometry by mitigating intra-pixel sensitivity variations.....	10
3.4	Improving the S/N by averaging over flatfield and background variations.....	11
3.5	Detector gaps, blemishes, bad pixels and persistent flux.....	12
3.6	Dithering to produce mosaics.....	13
4	Defining Metrics to Quantify Dither Patterns.....	14
4.1	Morphological properties.....	15
4.2	Astrometric centroiding precision	16
4.3	Photometric precision	17
4.3.1	Intra-pixel sensitivity	18
4.3.2	Background level uncertainties.....	19
4.4	Overhead time spent in dithering maneuvers.....	20
4.5	Area coverage lost near the field edges	20
5	Evaluation of Dither Patterns for JWST Instruments.....	21
5.1	NIRCam	21
5.1.1	Creation of Simulated Images	23
5.1.2	Analysis of simulated images	27
5.1.3	Results from Image Analysis.....	30
5.1.4	Final Evaluation of Metrics for NIRCam Dither Patterns.....	32
5.2	FGS-TF	35
5.3	MIRI.....	37
5.4	NIRSpec	41
6	Summary and Conclusions	42
7	References.....	43

An Investigation of Optimal Dither Strategies for JWST

1 Introduction

An important observational technique in imaging and spectroscopic observations involves the use of “dithering”, or offsetting the target by varying amounts across the detector for different exposures. In particular, the scientific quality of observations obtained with the JWST instruments may benefit from dithering in a number of ways, including the ability to improve the Point Spread Function (PSF) sampling by sub-pixel dithering, the use of large-scale offsets to help mitigate background effects across the detector, as well as the ability to account for a range of temporally and spatially variable effects arising from the instruments and the telescope.

The design of any dither pattern contains a significant number of variables, and these may be different depending upon the science goals of the observations and the limitations of the particular instrument that is being used to obtain the observations. Some of the parameters that describe a given dither pattern include the number of dither pointings, their relative location across the detector, and the degree to which they sample the sub-pixel space. Other issues also need to be taken into consideration, such as the degree of distortion in the camera (which may produce variations in the degree of sub-pixel sampling across the detector) as well as the possibility of obtaining multiple images at each dither location, either in the same filter or in different filters. Finally, the presence of temporal changes in the detector background or the telescope PSF may influence the temporal frequency and spatial distribution of exposures at different dither positions.

2 Scope of this Study

The purpose of this study is to provide a preliminary examination of optimal dither patterns for the JWST instruments. The study is preliminary in its scope since the design for all the JWST instruments has not yet been finalized, and they are all at different stages in their design and development process. Thus, the dither patterns are evaluated as “optimal” only in the context of a relatively limited and well-defined range of conditions.

More specifically, various aspects of the instrument design that would be prerequisites for a full discussion of optimal dither patterns are still in the early stages of characterization for some of the instruments (for example, the camera distortion and intra-pixel sensitivity variations). Nevertheless, it is possible to use the current knowledge of the instruments, together with extensive previous experience from HST, to carry out a relatively robust investigation of a number of aspects of dither patterns that will be relevant for JWST. Moreover, the results from this study may prove useful in evaluating trade-offs during the instrument design and development process. As the design and characterization of the instruments progresses, a more detailed follow-on study of dither patterns may be appropriate.

Therefore, the present study aims to address the fundamental issues that need to be considered when designing different types of dither patterns for the JWST instruments. This is achieved by breaking down the study into the following elements:

- Discuss the principal scientific drivers for different types of dither patterns, in the context of the observational design for the science goals described in the JWST Design Reference Mission (DRM, Donahue et al. 2002);
- Define metrics to quantify the relative advantages and disadvantages of different dither patterns, for particular types of science;
- Examine the properties of different dither patterns for the JWST instruments and calculate their metrics, to be used as a figure of merit in evaluating which patterns are optimal;
- Discuss the expected usage of different types of dither patterns, based on various different observational programs in the DRM. For observations involving multiple filters, discuss the relative operational and efficiency implications of changing the filters at each dither location, as opposed to executing a complete dither pattern before changing filters.

The study does not cover the small angle maneuvers carried out during initial target acquisition, since the images obtained in such a sequence are not considered part of the science data but rather form part of the process of moving the telescope to a desired location. Instead, this study is concerned with dither patterns used as part of the observational design of the science programs, and not during the acquisition phase.

The results from the present study are presented in the form of metrics that describe how well different dither patterns are able to improve the scientific quality of different types of scientific programs.

3 Scientific Drivers for Dithering

In a previous study, “Science Drivers for NGST Small-Angle Maneuvers” (Ferguson et al. 2000), the general science drivers for using dither patterns were described in detail and are summarized here.

1. Improve PSF sampling by non-integer dither offsets, particularly valuable in the short-wavelength regime where the PSF is often undersampled by the detector;
2. Improve S/N by averaging variations in pixel sensitivity, flatfield, background;
3. Mitigate detector-based artifacts (blemishes, bad pixels and persistent cosmic rays, dark current, bias level) which remain at residual levels after calibration;
4. Allow mosaicing of large regions using dithers comparable to the detector size.

Each of the instruments on JWST has somewhat different instrumental limitations and considerations, which are discussed in the following sections along with those issues that are common to the various instruments. Note that the instrumental specifications assumed in this study are derived from the JWST Mission Operations Concept Document (JWST-OPS-002018, Revision A, September 2, 2004).

3.1 Improving the sub-pixel PSF sampling

A number of the DRM programs require the highest possible spatial resolution (e.g., faint white dwarfs in globular clusters, the AGN/galaxy connection, and weak lensing). However, all the JWST instruments are undersampled to different degrees, as follows:

- The FGS-TF is the most severely undersampled instrument, since it has detector pixels of 65 mas and thus does not provide Nyquist sampling for any wavelengths in its short-wavelength channel (nominally 1.0 – 2.1 μm); however, the long-wavelength channel (2.1 – 4.9 μm) is Nyquist sampled above about 4 μm .
- NIRC*am* also has a short- and long-wavelength channel, but the short-wavelength pixels are half the size of those in the long-wavelength channel (31 mas and 65 mas respectively), therefore the two channels are Nyquist sampled at 2 μm and 4 μm respectively. However, the wavelength coverage of the short-wavelength channel may possibly extend all the way down to 0.6 μm , in which case undersampling would be a significant issue across much of its wavelength range.
- The MIRI imaging mode is well sampled across most of its 5 – 27 μm wavelength range, with 110 mas pixels that provide Nyquist sampling for all wavelengths above 7 μm . The spectroscopy mode uses two pairs of integral field spectrographs, one in the range 5 – 12 μm (with 200 mas pixels), and the other in the range 12 – 27 μm , with a range of spatial and spectral scales for the pixels. In addition, the instrument has coronagraphic masks and a fixed-slit low-resolution spectrometer (LRS).
- The NIRSpec detector pixels are 100 mas in size, which undersample the PSF across its entire wavelength range of 0.6 – 5 μm . More importantly, however, it makes use of a micro-shutter array (MSA) consisting of 750x350 shutters, each of which is 0.2''x0.45'' in size, with 0.25''x0.5'' spacing. Therefore, dither patterns for NIRSpec are concerned with subsampling the detector pixels as well as the MSA apertures, and also with accounting for light losses associated with the shutters.

An important consideration for all the undersampled imaging modes is to ensure that dither patterns are able to recover the original PSF as closely as possible, within the constraints of limited observing time. For the FGS-TF, this may require dither patterns as finely spaced as 1/3 of a pixel near the blue end, while for NIRC*am* and MIRI it may be sufficient to obtain sub-pixel sampling to 1/2 pixel (possibly 1/3 pixel near the blue end of the NIRC*am* short-wavelength channel). These finely dithered images would then be used to construct an output image with a pixel scale corresponding to the size of the dither steps, thereby allowing both the PSF sampling and resolution to be improved.

3.2 Sampling vs. resolution

It is important at this point to draw a distinction between improving *sampling* and *resolution*, to help avoid potential confusion between these two concepts. Each pixel in an exposure can be thought of as a point sample of the “true image” of the sky which has been convolved by the telescope PSF and the detector pixel response function. Obtaining multiple exposures of undersampled images with different sub-pixel shifts therefore provides additional information on sub-pixel scales that cannot be recovered from a single exposure. The use of a linear image reconstruction technique, such as “drizzle” (Fruchter & Hook 2002) allows the images to be combined in a way that can improve both sampling and resolution, but in separate ways, described as follows:

- **Sampling:** if the exposures have been obtained in a way that provides sufficient sampling of the sub-pixel space (for example, a 4-point 2x2 or a 9-point 3x3 sub-pixel grid) then enough information is present to enable all these exposures to be meaningfully combined to an output image with a smaller scale. Typically the output image pixel scale will be smaller than the original pixel size by an amount corresponding to the degree of subsampling (e.g., 1/2 or 1/3 of the original detector pixel size), and the resulting image will have more spatial information content than would have been the case if no sub-pixel dithering had been applied.
- **Resolution:** The way in which separate exposures are combined holds the key to how much the resolution can be improved. If the input exposures are simply shifted and co-added (“shift-and-add”) onto a finer output pixel scale, then the output PSF is convolved a second time by the input pixel size (as well as by the output pixel size) since each input pixel covers several output pixels. This convolution can be avoided if the input pixels are shrunk down to delta functions before mapping onto the output plane (“interlacing”). Moreover, if each output pixel is large enough to contain the central positions of a large number of input pixels, then this produces an additional convolution by the output image pixel size, which can be mitigated by reducing the output pixel scale to a level comparable to the dither offset between images. The “drizzle” algorithm provides a continuum between these different types of behavior. The sharpest resolution that can therefore be recovered using “drizzle” or other standard linear reconstruction algorithms consists of a single convolution of the input PSF by the input detector pixel scale, with no additional convolutions.

While linear reconstruction algorithms such as “drizzle” or interlacing can at best recover the convolution of the PSF with the input detector pixel scale, there also exist other techniques that can in principle improve the resolution further. These “image restoration” techniques are conceptually different from image reconstruction, in that they aim to improve on the resolution by enhancing high-frequency components, while at the same time attempting to preserve signal-to-noise ratio, for example the Richardson-Lucy method (Richardson 1972; Lucy 1974), maximum entropy, or full Fourier reconstruction (Lauer 1999). Unfortunately there is always a trade-off between S/N and resolution with these techniques, in addition to limitations arising from distortion in the camera, and as a result they have not been as widely applied to the regime of undersampled HST observations. While they may be considered for specialized applications, the present

study concentrates on the use of linear reconstruction techniques such as “drizzle”, where no sacrifice for S/N is made and where the behavior of errors and other systematics remains well understood.

For any sub-pixel dither pattern, it will be essential to determine the precise relative positions of the images being combined to a small fraction of a pixel. The most stringent constraint in this regard comes from the short wavelength channel of NIRCам which has the smallest pixels – the FWHM of the PSF near the blue end becomes smaller than the pixel scale, and would require 1/3 pixel subsampling to fully recover. This corresponds to 10 mas offsets, thus the current baseline pointing accuracy of 5 mas r.m.s. would be sufficient to ensure that the offsets are reliably executed.

If additional pointing uncertainties are present, then non-optimal sub-pixel dither positioning will result. Moreover, it may not always be feasible to measure shifts *a posteriori* from the images directly, if there are insufficient objects to permit reliable registration (for example, in narrow-band observations of faint sources). Therefore, it remains extremely important to maintain a pointing and repeatability accuracy of no worse than 5 mas during a given set of exposures.

For spectroscopic modes, specifically the MIRI IFU and the NIRSpec MSA, it is also important to be able to subsample the Line Spread Function (LSF), which determines the degree of spectral resolution of the data. The reasoning is quite similar to that which is applied in the imaging domain, namely that the LSF should be at least Nyquist sampled by the detector pixels. There are essentially three spectral resolution regimes: the prism mode ($R \sim 100$) and the low- and high-resolution first-order gratings ($R \sim 1000$ and 3000). If these do not provide adequate sampling of the LSF by the detector pixels, then additional dither strategies will be needed to recover the spectral resolution.

3.3 *Improving photometry and astrometry by mitigating intra-pixel sensitivity variations*

Another effect that can be mitigated by sub-pixel dithering is the variation in sensitivity across each individual pixel. Several of the DRM programs require very precise photometry to be obtained, for example halo population studies and stellar measurements in globular clusters. However, in undersampled images, one of the dominant sources of uncertainty and error in detecting compact sources comes from the sensitivity variations across pixels, i.e. the intra-pixel sensitivity response function. Depending on where the peak of the unresolved emission is located on the pixel, the resulting flux can vary by as much as 0.4 magnitudes in the HST NIC3 images (Lauer 1999). In addition to photometric uncertainties, these variations can also introduce centroiding errors, which in turn directly impact DRM programs requiring precise astrometry. Furthermore, such centroiding errors also cause problems generally when attempting to register images based on the measured positions of sources, and this in turn degrades the resolution and signal-to-noise ratio of the final image.

Intra-pixel sensitivity variations are likely to have the strongest effect in observations of compact objects with detectors where the pixels are the most undersampled, namely the

short wavelength channel of the FGS-TF, and the blue end of the FGS-TF long wavelength channel as well as the blue end of both wavelength channels of NIRCам. Spectroscopy with NIRSpec and the MIRI IFU may also be impacted in cases where the spectral resolution is not well sampled by the detector pixels.

Photometry measurements of extended sources are less likely to be affected by intra-pixel sensitivity variations, since the flux is distributed over more detector pixels and the integrated flux is not dominated by a single peak pixel, as in the case of undersampled unresolved sources. Nevertheless, some types of science can be affected, for example color gradient measurements across galaxies, or surface brightness fluctuation studies. Therefore, studies of extended sources can be considered to provide similar requirements for mitigating intra-pixel sensitivity variations as those derived from unresolved sources, and the resulting dither patterns can be used for both types of science.

3.4 Improving the S/N by averaging over flatfield and background variations

Since the counts in many of the DRM observational programs will be dominated by sky background by up to a few orders of magnitude, this means that even small uncertainties in the level of this background can lead to serious errors in the detection and photometry of sources. The quality of photometric measurements ultimately depends upon an accurate knowledge of the background level around the target objects, and is thus dominated by uncertainties in this background. There are three fundamentally different contributions to such uncertainties:

- 1) Temporal variations in the level of the background – in exposures at different locations, this can arise from different amounts of scattered light from bright objects (either the target itself, or other sources on the detector or outside the field of view). In addition, time-dependent changes in the detector electronics (dark current, bias level, thermal glow from various components in the instrument or from the telescope) can also contribute to a temporally variable background.
- 2) Uncertainties in the large-scale flatfielding calibration, even if the background level remains unchanged from one exposure to the next. Such uncertainties can arise if different parts of the detector or filter slowly evolve with time after launch, or if there are fundamental limitations in the method used to obtain the flatfield (for example, if color-dependent flatfield variations are imperfectly accounted for).
- 3) Uncertainties in the pixel-to-pixel flatfield (as opposed to the intra-pixel sensitivity variations discussed earlier). Each pixel has a slightly different total response compared with other pixels, and these can only be well characterized by obtaining exceptionally high S/N flatfields. If pixel sensitivities evolve slowly after launch and are not re-measured on a regular basis to the same level of accuracy, then this eventually corresponds to a degradation in the quality of the flatfield, which in turn adds a noise component to the measurement of the background emission.

The strategy to mitigate all three of these effects is to dither on scales of at least several pixels, as opposed to the sub-pixel dithering discussed previously; note that combining

sub-pixel shifts with integer pixel shifts is commonly done to achieve a number of different improvements simultaneously. The shifts in this case would ideally need to be on scales large enough to sample a set of independent background regions on the detector, and the improvement in S/N of the background determination would essentially scale in quadrature with the number of pointings, until it becomes insignificant relative to the photon noise from the background itself. One caveat also exists in the mid-IR – if the dithering itself introduces thermal changes in the telescope as a result of slightly different pointings, then it may be better not to dither by large amounts, although currently it is not expected that such thermal changes would be significant.

An additional advantage of large-scale dithers is that they potentially allow self-calibration to be achieved, for example to actually improve the detector sensitivity calibration on a given field by comparing measurements of the same sources across different parts of the detector (Arendt et al. 2000; Fixsen et al. 2000). With a sufficient number of dither positions, such self-calibration can in principle be achieved to levels substantially more accurate than the original uncertainties inherent in the background.

In practice, background uncertainties exist on a range of scales – not only may the large-scale flatfielding uncertainties and scattered light components be present at different spatial scales, but other effects due to the detector electronics (eg, bias level or dark current effects) can have completely different signatures, such as ripples or waves across the detector. Hence, this suggests the need for the dither patterns themselves to sample the full range of spatial scales across the detector, in all directions.

However, one final point to consider is that dithers on scales comparable to the detector size are more appropriate when constructing a mosaic consisting of multiple pointings, and this is discussed in more detail in a subsequent section. For observations that are not aiming to carry out mosaicing, the loss in sensitivity around the edge of the field should be taken into consideration when choosing the dither pattern. Moreover, while the considerations outlined here are applicable to both the imaging and spectroscopic modes of the JWST instruments, the spectroscopic modes have additional considerations (such as loss of wavelength coverage) that may limit the maximum step size of large dithers.

3.5 Detector gaps, blemishes, bad pixels and persistent flux

The common theme underlying these issues is that they present some form of loss of information on the detector that requires dithering to different positions in order to recover. Although not all the DRM programs require contiguous coverage in instruments that have gaps between the detectors (for example, the 2x2 array of SCAs in the short-wavelength channel of NIRCam), it is likely that imaging observations will often make use of dithering simply to fill in such gaps and thereby simplify analysis steps such as cataloging and object photometry. For example, many of the current HST ACS/WFC observing programs include dither steps large enough to cover the 2" gap between the two chips, even if the observations are not of large objects that are impacted by the gap, and this is generally in order to provide a single contiguous image that will simplify the analysis.

Moreover, if the detectors have blemishes from the manufacturing process (or develop them over time) then dithering will be needed in order to move these around on the sky and ensure that a given target of interest does not fall on such a region during all the exposures. Similar considerations apply to bad pixels that may have anomalously high dark current, since the information on such pixels is effectively lost and they need to be moved around the sky in order to recover the flux for any object that falls onto them.

The dither patterns for covering both chip gaps and detector blemishes can be developed ahead of time and will likely remain unchanged, since such features are not likely to change significantly during the course of a given set of observations (although it is possible that blemishes may change over timescales comparable to the life of the observatory). Since the considerations are similar to those of flatfielding variations, it is likely that these dither patterns will form part of the same class of large-scale patterns.

A different category of artifacts is presented by the presence of persistent flux from bright pixels in previous exposures, either astronomical targets or cosmic rays. Persistence is a known problem on the HST/NICMOS cameras and also the instruments on board Spitzer, and the residual emission can persist for many subsequent exposures while also decaying on timescales comparable to the timescale of the observations. This is one of the most difficult types of effects to compensate for by means of dithering, since its scope is effectively unknown at the start of a new sequence of observations. For a given field, however, the anticipated effect of persistence in subsequent exposures from objects in that field can be estimated quite reliably, and the dither pattern can therefore be adjusted accordingly before the observations execute. In general, this will likely consist of a modification of one of the pre-defined large-scale patterns with shifts appropriate to ensure that any persistent emission on the detector is moved sufficiently around the field.

The NIRSpec MSA presents its own unique set of requirements when considering dither patterns to mitigate gaps on the detector. Because of the significant gaps between each shutter, the DRM spectroscopic programs may require dithering to move these gaps to different parts of the sky. Since the shutters are in a rectangular grid pattern, this allows the design of dither patterns that move the detector by corresponding amounts and obtain spectra for contiguous regions along the sky.

3.6 Dithering to produce mosaics

Some of the DRM programs, for example the dark matter weak lensing program and rich cluster galaxy programs, require contiguous imaging of regions much larger than the individual detectors, each containing 100 or more pointings. This is achieved by means of large-scale dithers comparable to the scale of the detectors. These programs will most likely make predominant use of NIRCам and MIRI, and the dithering patterns are thus driven by the instrumental specifications of these cameras.

Each pointing may be observed with a number of large-scale plus sub-pixel dithers to mitigate the effects that have been described in the previous sections. In addition, however, shifts comparable to the size of the detector scale will be added to the

patterns. The primary pointing requirement is to provide sufficient accuracy to ensure contiguous coverage and avoid gaps due to inaccurate placement, while at the same time desiring to minimize overlap in order to cover as much area as possible. The baseline acquisition accuracy of < 1 arcsecond should be sufficient to achieve this, while limiting overlap to less than a few percent of the detector scale. Programs such as COSMOS with HST/ACS (600 separate pointings) have demonstrated that this technique can be used extremely successfully in practice.

Finally, an important concern is the trade-off in overhead between filter changes and dither steps, along with the degree of stability of the telescope in order to ensure uniformity in the resulting dataset. Thus, several different dither patterns may be required, depending upon the number of filters used, the total size of the mosaic, and the extent to which visits need to be repeated in a given filter to search for transients.

4 Defining Metrics to Quantify Dither Patterns

The science drivers in the previous section lead to three different regimes for constructing dither patterns, and a given observational program may make use of any or all of them in various combinations:

- Sub-pixel sampling
- Intermediate-scale dithering, larger than a few pixels but smaller than the detector
- Mosaicing on scales significantly larger than the detector

In order to examine which patterns are “optimal”, their relative merits are quantified by means of cost/benefit metrics that are defined in the context of specific observational science questions. Based on the considerations presented by the various science drivers, the metrics can be divided into the following categories:

- 1) Morphological properties (e.g., FWHM)
- 2) Astrometric centroiding precision
- 3) Photometric precision
- 4) Overhead time spent in dithering maneuvers
- 5) Area coverage lost near the edges

The first three categories describe scientific improvements in the output image as a result of the use of dither patterns. The metrics defined in this section provide a means of quantifying these improvements, but it is still expected to be generally true that more detailed dither patterns will continually produce improvements in the resulting images, asymptotically approaching the recovery of the ideal input image as the number of dithers becomes large. Thus, in order to meaningfully describe a pattern as “optimal”, its

advantages must therefore be traded off against disadvantages. The fourth and fifth categories achieve this by providing a metric to describe quantitatively the cost of a given dither pattern, which can then be evaluated against the metrics that describe the improvement in scientific quality.

For each of the different instruments, the relative merits of different types of dither patterns can be examined quantitatively, either by means of simulations or by using analytical calculations. The fundamental aim is to examine how well different dither patterns are able to recover the intrinsic characteristics of sources, given the fact that the desire for improved observational precision needs to be traded off against limited observing time, and that overheads from small angle maneuvers should not become excessive relative to the total amount of observing time. While other, additional metrics may also be included, for the present study we begin by discussing the above set, which represent a large range of scientific issues involved in obtaining dithered observations.

In this section, the various metrics are defined along with analytical and numerical procedures that can be used to quantify them for different dither patterns. These metrics then provide a quantitative means by which observational improvements from dithering can be traded off against the costs required to obtain them.

4.1 *Morphological properties*

We begin by describing mathematically the convolutions applied to the true astronomical image when observed with a series of dithered exposures, adapting the formalism in Fruchter & Hook (2002) to the specific cases in the present study. The true distribution of flux on the sky, S , is initially convolved by the telescope optics O and subsequently forms an image $S \otimes O$ at the detector focal plane, where \otimes is the convolution operator. The flux measured by each pixel of the detector can then be thought of as a point sample of this image convolved by the electronic pixel response function E , thus we obtain the convolved image $I_C = S \otimes O \otimes E$. Here the quantity E would ideally describe the intra-pixel sensitivity response function as well as other effects that may extend beyond the pixel, such as the “charge diffusion” kernel in the HST/WFPC2 cameras, although in HST/NICMOS this kernel is smaller than the physical pixel. The convolved image I_C therefore represents the sharpest image that can be recovered from sub-pixel dithered images using standard image reconstruction techniques, and forms the baseline for comparison with measured morphological quantities from combined dithered images.

The scenario that is considered here for producing combined images from separate input images is to use the “drizzle” software, which allows input pixels to be “shrunk” to an arbitrarily small size P prior to being mapped onto the output grid. This provides a continuum of behavior ranging from shift-and-add (if $P =$ original pixel size) to interlacing ($P = 0$), and can therefore be adapted to the specific requirements of different dither patterns. In addition, the scale of the output pixel grid G can also be adjusted: if each output pixel contains the centers of many input pixels, then the final convolved image is $I_C \otimes P \otimes G$. On the other hand, if the output pixel grid is made small enough that each output pixel contains flux from only one input pixel, then the convolution with G disappears, yielding just $I_C \otimes P$, and this in turn reduces simply to I_C if $P = 0$.

Thus, in principle, the use of “drizzle” can allow the original convolved image I_C to be recovered when combining a series of dithered images, if the dither pattern is sufficiently well designed and if other adverse effects such as pointing errors are sufficiently small. Therefore this metric is simply defined as the ratio of morphological properties measured for sources on the output image (e.g., FWHM) relative to their values in the original convolved image I_C . How closely a given dither pattern allows these values to be recovered is thus the figure of merit that is used in comparing different dither patterns. A wide range of morphological parameters can be considered in principle, including size measurements for extended objects, scale length, ellipticity, and stellarity or degree of concentration. However, FWHM is the simplest to examine in this preliminary study, and also relates directly to the properties of the PSF when describing unresolved sources.

4.2 *Astrometric centroiding precision*

The precision with which object positions can be recovered in a drizzled, combined image is a function of the original degree of undersampling of the exposures, the extent of subsampling achieved by the dither pattern, as well as the final pixel scale and the signal-to-noise of the source. While good astrometry is necessary for all science programs, the DRM programs that place the most stringent requirements on astrometry are those that involve positions of unresolved sources (for example, stars or unobscured active galactic nuclei). Therefore, for the purposes of the present study, unresolved sources are used as the basis for defining a metric to describing the quality of astrometric registration for different dither patterns.

A general rule in the case of a well-sampled PSF with an approximately Gaussian core is that the $1\text{-}\sigma$ error in the centroiding accuracy scales linearly with the signal-to-noise ratio SNR, as follows:

$$\sigma \approx \text{HWHM} / \text{SNR}$$

where HWHM is simply 1/2 of the FWHM of the PSF. In the case of significantly undersampled images, this rule applies to the PSF that results from the convolution of the telescope optics with the detector pixel kernel (I_C in the preceding section). Thus, the accuracy with which any centroids can be determined is limited by the input pixel detector scale, in the regime of sufficient subsampling and sufficiently small output pixels.

Thus, a relatively robust metric to describe the ideal accuracy with which astrometric centroids can be recovered consists of determining the FWHM of an unresolved source in the convolved image I_C , and then measuring for different dither patterns the actual accuracy with which source positions are recovered. The figure of merit for describing each dither pattern is then defined as the amount of dispersion of these position measurements, as compared with what is expected for the ideal image I_C , as a function of SNR of the sources.

4.3 Photometric precision

Two important effects that can degrade the precision of photometric measurements are intra-pixel sensitivity variations and uncertainties in the measured background level. Both of these can be mitigated by means of appropriate dithering strategies, on the sub-pixel scale in the first instance and on a range of larger scales for the second case. In each case, a metric for quantifying the figure of merit of different dither patterns can be defined by considering the degree of uncertainty introduced by either the intra-pixel sensitivity variation or the background level.

The general expression for the variance of the photometry of a source in an aperture of n_{phot} pixels, with total integrated object count-rate c_o , observed for a total exposure time t divided over N_{dth} separate dither positions, with N_{exp} read-out exposures at each dither point, and sky count-rate per pixel c_s , dark current rate per pixel c_d , read-out noise per pixel r , and using n_{bkgd} pixels around the object to estimate the background, is given by:

$$\begin{aligned} \sigma_{\text{tot}}^2 = & [c_o + n_{\text{phot}}(c_s + c_d)]t + N_{\text{dth}}N_{\text{exp}}[r^2n_{\text{phot}} + r^2n_{\text{phot}}^2/n_{\text{bkgd}}] \\ & + [\delta_p(c_o + n_{\text{phot}}c_s)t + \delta_s c_s t n_{\text{phot}}^2/n_{\text{bkgd}}]/N_{\text{dth}} \end{aligned} \quad (1)$$

where the different contributions to the variance are as follows:

$(c_o + n_{\text{phot}}(c_s + c_d))t$ is the noise contribution from the source photon statistics, the background sky, and dark current, in the photometric aperture being used;

r^2n_{phot} is the variance from read-noise in the photometric aperture, per exposure;

r^2n_{bkgd} is the variance from read-noise in the background aperture, per exposure, which is scaled by $(n_{\text{phot}}/n_{\text{bkgd}})^2$ to the photometric aperture being used;

$\delta_p(c_o + n_{\text{phot}}c_s)t$ is from intra-pixel variations in the photometric aperture

$\delta_s c_s t n_{\text{phot}}^2/n_{\text{bkgd}}$ is from systematic uncertainties in the background estimate, which is scaled by $(n_{\text{phot}}/n_{\text{bkgd}})^2$ to the photometric aperture being used;

Note that the variance from both read-out noise and photon noise associated with the background measurement decreases as the ratio $(n_{\text{phot}}/n_{\text{bkgd}})^2$, representing the fact that a larger background region, relative to the photometric aperture, provides improved statistics in estimating the background level.

The quantities δ_p and δ_s represent the fractional contribution to the total variance from intra-pixel variations and systematic uncertainties in the background sky estimate respectively, at a single dither position. Thus, for example, $\delta_p = 0.05$ would represent an additional contribution of 5% of the total counts r.m.s. as a result of intra-pixel sensitivity variations in all the pixels in the photometric aperture. Similarly, $\delta_s = 0.1$ would imply an additional contribution of 10% of the background sky r.m.s. obtained in a larger region around the object of interest, for example resulting from large-scale fluctuations in the flatfield accuracy or other systematic structures across the background aperture.

It should also be noted that the variance contributions from δ_p and δ_s do not depend on the number of read-out exposures at each dither position, but only on the total number of dither positions. Their overall contribution decreases for an increased number of dither positions, under the assumption that the dither positions are placed such that they sample independent regions for these effects. This assumes that these systematic effects remain the same for multiple exposures at the same dither location, which is likely to be the case for intra-pixel sensitivity variations as well as uncertainties in the large-scale flatfield calibration.

Since the contributions from these two terms δ_p and δ_s to the total variance are quite independent, and are mitigated by two different types of dither strategies (sub-pixel and large-scale), this metric is divided into two sub-categories, treating these variance contributions separately in each case.

4.3.1 Intra-pixel sensitivity

The precise intra-pixel sensitivity distribution of the detectors in the JWST instruments is not yet known, since none of the final flight detectors have yet been procured. Nonetheless, a relatively robust metric can be defined in terms of the total amount of intra-pixel sensitivity variation integrated over a kernel corresponding to the convolved PSF, thus representing the level of uncertainty in the measured flux of an unresolved source in the high count-rate limit.

From equation (1) we see that, in the regime where the object and sky counts dominate over the dark current and read-out noise – i.e. $(c_o + c_s) \gg c_d, r$ – the variance reduces to:

$$\sigma_{\text{tot}}^2 \approx t(c_o + n_{\text{phot}}c_s)(1 + \delta_p/N_{\text{dth}}) \quad (2)$$

Thus, if intra-pixel sensitivity variations contribute $N\%$ to the total variance at a single dither position, then this can to first order be reduced below 1% by obtaining N dithers across the sub-pixel space. It is worth noting that this conclusion applies regardless of how bright the source is relative to the background – in other words, it applies equally well in the regime of bright sources as well as faint sources where the variance is dominated by the counts from the background in the photometric aperture.

This formalism makes no statement about how the dithers should be distributed across the sub-pixel space, other than the assumption that they are distributed in a way that provides independent samples of the intra-pixel sensitivity variation across the pixel. If such variations occur on more or less the same spatial scales across the pixel, then the optimal dither pattern for sampling them would consist of a regular square grid of steps in sub-pixel space.

Hence, the metric for assessing how well a particular sub-pixel dither pattern mitigates the effect of a particular distribution of intra-pixel sensitivity variation involves a comparison between the expected variance and the measured variance for objects on the combined image from a set of sub-pixel dithered exposures, in the limit where counts from the object dominate the overall variance.

4.3.2 Background level uncertainties

The precision with which the background level can be measured around an object of interest is another dominant source of error in the quality of the photometric data. Uncertainties in the background distribution across the detector can arise from a number of different effects, including an imperfect knowledge of the flatfield calibration, optical effects such as scattered light from sources in the field or thermal glow from the telescope or instrument, electronic effects arising from sources in the field (e.g., cross-talk), or changes in the detector electronics (e.g., ripples or other structures arising from the bias level or even the dark current). All these effects can be considered together as introducing a net uncertainty in the background measurement around an object in a given exposure, and are mitigated by obtaining multiple large-scale dither positions.

Assuming that intra-pixel variations are already accounted for (e.g., by adding sufficient sub-pixel dither steps to the larger-scale dither pattern), then in the regime where the variance is dominated by object and sky counts – i.e. $(c_o + c_s) \gg c_d$, r – the total variance in equation (1) reduces to:

$$\sigma_{\text{tot}}^2 \approx t[c_o + n_{\text{phot}}c_s(1 + \delta_s n_{\text{phot}}/(n_{\text{bkgd}}N_{\text{dth}}))] \quad (3)$$

This assumes that the large-scale dither pattern samples the background variations sufficiently well at independent regions, in order to decrease the uncertainty from the systematics. The decrease in this case scales in the same way as increasing the number of pixels that are used to obtain a measurement of the background, which is applicable in the regime where the region to measure the background is smaller than the spatial scale of the background variations. As in the case of intra-pixel variations, this formalism places no constraints on the dither pattern other than the requirement that it provide independent samples of the background variation on all the relevant spatial scales.

For a given science program, with certain target sizes that determine both n_{phot} and n_{bkgd} , together with a given object count-rate relative to the sky, c_o/c_s , the contribution of δ_s to the total variance for each target can then be immediately determined based on the number of pointings in the dither pattern. For example, if the systematic uncertainty in the background determination contributes $N\%$ to the total variance at a single dither position, then this can to first order be reduced below 1% by obtaining $N \times n_{\text{bkgd}} / n_{\text{phot}}$ large-scale dithers across the detector (thus, for example, obtaining N dithers if the aperture used to measure the background n_{bkgd} contains the same number of pixels as the photometry aperture n_{phot}).

Hence, the metric for quantifying the relative performance of different large-scale dither patterns also consists of the comparison between the expected variance and the measured variance, as it was for the intra-pixel case, but in addition it depends upon the ratio of aperture sizes used for photometry and background determination. For different science cases in the DRM, this metric can then be specifically quantified to provide a figure of merit for different large-scale dither patterns.

4.4 Overhead time spent in dithering maneuvers

One of the principal costs of dithering is the overhead time lost to moving the telescope. This leads to a relatively straightforward metric, defined as the fraction of time lost to overhead f_o when using a dither pattern of N dithers, each of exposure time E , where the time for a dither offset consists of a slew time S_i (which depends on the dither distance) and a constant component C for each dither step, representing the time required for telescope settling, guide star reacquisition (if necessary) and other related overheads:

$$f_o = \frac{NC + \sum_{i=1}^N S_i}{N(C + E) + \sum_{i=1}^N S_i} \quad (4)$$

Different DRM science programs may have different requirements for the maximum tolerable value of f_o ; in general, HST observational programs have aimed to keep this quantity below the level of 5-10%.

In the special case of uniform dither steps from one exposure to the next, the expression for the fraction of dithering overhead time simplifies to:

$$f_o = \frac{C + S}{C + S + E} \quad (5)$$

and thus becomes independent of the number of dither steps, depending only upon their distance and the reacquisition/settling timescale overhead for each step. This is useful because even relatively complex dither patterns can be generally arranged in an order that minimizes the steps from one position to the next, thus approximating the constant-dither step regime and allowing the simplified version of this metric to be applied. The exposure time required to keep f below a given value is then simply:

$$E \geq (C + S) \frac{(1 - f_o)}{f_o} \quad (6)$$

This simplifies further in two limiting regimes: small dither steps where the constant overhead time is dominant ($C \gg S$), and large-scale steps where the slew time becomes dominant ($S \gg C$). Thus, given a required total observation time, together with some minimum and maximum constraints on the length of the time for each exposure read-out, as well as the time spent at each dither position, the fraction of time lost to dither-related overheads can be quantified for any given dither pattern, and hence provides a figure of merit for its cost.

4.5 Area coverage lost near the field edges

Another cost of dithering is the necessary loss in area that is covered to the full depth of all the exposures, resulting in some area around the edges that is covered to a lower

depth. This consideration is generally most important for observations of a single field, where the coverage is comparable to the detector size. For mosaicing observations of a contiguous area covering a large number of fields, the effect becomes less important since the shallower regions around the edges of each field can be overlapped with corresponding regions from adjacent fields to recover the full depth.

For the case of an observation of a single field, consider a series of dither steps, each with a pixel location (x_i, y_i) . In this case a metric f_A can be defined to represent the ratio of the area with full coverage (i.e., smaller than the detector size D pixels) relative to the more extended area that includes the shallower edges:

$$f_A = \frac{(\max(x_i) - \min(x_i))(\max(y_i) - \min(y_i))}{(2D + \min(x_i) - \max(x_i))(2D + \min(y_i) - \max(y_i))} \quad (7)$$

For mosaicing observations aiming to produce contiguous coverage, it can be shown that the full depth is recovered in the shallow edges around each of the individual fields (except those along the outside edge) by repeating the single-field dither pattern at large steps of exactly D pixels along the x and y directions. This would yield a mosaic with uniform coverage to the full depth across its entire extent, except for shallower coverage around its edge. In this case, for a mosaic consisting of $N_x \times N_y$ fields, the expression for the fraction of full depth to total area can be generalized to:

$$f_A = \frac{((N_x - 1)D + \max(x_i) - \min(x_i))((N_y - 1)D + \max(y_i) - \min(y_i))}{((N_x + 1)D + \min(x_i) - \max(x_i))((N_y + 1)D + \min(y_i) - \max(y_i))} \quad (8)$$

Thus, the above expressions provide a direct figure of merit for any dither pattern that can be considered as a cost against the benefits obtained from the pattern.

This metric is most likely to be useful in evaluating large-scale dithers of observations of a single field, where the benefits of improved background subtraction and photometry resulting from the large-scale dithers must be traded off against the loss in area covered to full depth. In addition to the loss of exposure time depth, other issues associated with this cost include less effective cosmic ray removal and less accurate measurement of the background for sources along the edges.

5 Evaluation of Dither Patterns for JWST Instruments

5.1 NIRCam

NIRCam has a total of 10 sensor chip assemblies (SCAs), each 2048x2048 pixels, mounted in two identical optical modules. Each optical module has a short wavelength channel (0.6 – 2.3 μm) consisting of a 2x2 grid of SCAs with 31 mas pixels, together with a single long-wavelength SCA (2.4 – 5 μm) which has 65 mas pixels. Thus the two wavelength channels are Nyquist sampled at 2 μm and 4 μm respectively, although the short wavelength channel remains undersampled for much of its wavelength coverage,

and this is particularly severe towards its blue end at 0.6 μm . Thus a significant component of this study involves the investigation of sub-pixel dithers to recover the PSF. In addition, the 2x2 array of SCAs in the short wavelength channel are each separated by a 5'' gap, thus many observers will likely make use of dither patterns that cover the gap to provide a seamless image, even for observations of a single field.

At the time of writing, the distortion model for NIRCcam was not yet finalized, therefore this study instead focuses on the gains achieved with sub-pixel sampling on a single target, neglecting the fact that other locations on the chip may have different degrees of subsampling as a result of distortion. This should be addressed in a future study.

We carried out simulations of sub-pixel dither patterns that range from a simple 2-point pattern to a full 9-point (3x3) sub-pixel grid, and also examined ways in which patterns can simultaneously achieve sub-pixel sampling in the two wavelength channels. Dither patterns were created for all the broad-band filters in each channel: F070W, F110W, F150W and F200W for the short wavelength channel, and F270W, F357W and F444W for the long-wavelength channel. These patterns are listed in Table 1, together with the drizzle parameters that were used to construct the resulting combined images. The drizzle parameters are chosen to include parameters designed to optimize resolution for each pattern, as well as allowing intercomparison between different dither patterns.

Table 1. NIRCcam dither patterns and drizzle parameters

Pattern	Pixel shifts (x,y)	Output scale	Output pixfrac
2-point: Short	(0,0) (0.5,0.5)	0.5	1.0
Long	(0,0) (0.5,0.5)	0.5	1.0
3-point: Short	(0,0) (1.333,1.333) (2.667,2.667)	0.5	1.0
Long	(0,0) (0.656,0.656) (1.313,1.313)	0.5	1.0
4-point A: Short	(1.0,0.5) (1.5,1.5) (0.0,0.0) (0.5,1.0)	0.5	0.8
Long	(1.0,0.5) (1.5,1.5) (0.0,0.0) (0.5,1.0)	0.5	0.8
4-point B: Short	(2.500,1.000) (3.500,3.500) (0.000,0.000) (1.000,2.500)	0.5	1.0
Long	(1.231,0.492) (1.723,1.723) (0.000,0.000) (0.492,1.231)	0.5	1.0
9-point: Short	(0.000,2.667) (1.333,2.667) (2.667,2.667) (0.000,1.333) (1.333,1.333) (2.667,1.333) (0.000,0.000) (1.333,0.000) (2.667,0.000)	0.5	0.8
Long	(0.000,1.313) (0.656,1.313) (1.313,1.313) (0.000,0.656) (0.656,0.656) (1.313,0.656) (0.000,0.000) (0.656,0.000) (1.313,0.000)	0.5	0.8

Note that the patterns are chosen to provide optimal subsampling in the short-wavelength channel and near-optimal subsampling in the long-wavelength channel; pattern 4A provides a comparison with optimal subsampling in the long-wavelength channel.

5.1.1 Creation of Simulated Images

The first step in creating the simulated dithered NIRCam images consisted of creating individual exposures at each dither location, for each NIRCam filter. These exposures utilized a model of the telescope and detector optics that had initially been created as a means of testing and establishing the limits on routine wavefront sensing (Sivaramakrishnan et al.). The simulation utilizes the currently anticipated NIRCam detector specifications, as follows:

- Gain = 1.3 e⁻/adu
- Readnoise = 15 e⁻ RMS per read
- Saturation level = 100,000 e⁻
- Dark current = 0.01 e⁻/second
- Pixel scale (short wavelength detector) = 0.0318" per pixel
- Pixel scale (long wavelength detector) = 0.0647" per pixel

The simulation software used as input the multi-wavelength PSF produced by the telescope optics, which was obtained from S. Casertano and represents the nominal JWST PSF expected during science operations. Simulated NIRCam images using this PSF were subsequently produced in each of the broad-band filters, where each broad-band filter was described by a top-hat transmission function as shown in Table 2.

Table 2. NIRCam filter throughput functions used in the simulations

Filter	Low λ (μm)	Central λ (μm)	High λ (μm)
F070W (short)	0.625	0.70	0.775
F110W (short)	1.00	1.10	1.20
F150W (short)	1.35	1.50	1.65
F200W (short)	1.80	2.00	2.20
F270W (long)	2.40	2.70	3.00
F357W (long)	3.17	3.57	3.97
F444W (long)	3.99	4.44	4.89

For each dither pattern, a series of dithered exposures were created for each filter, of an artificial star field consisting of a grid of 1024 (32x32) stars across the detector. Since distortion information could not yet be incorporated into the simulations, the star field is instead used to provide a Monte-Carlo simulation of how well the PSF can be recovered using different dither patterns, for stars located at a range of positions in sub-pixel space. The principal grid pattern for the stars consisted of 64-pixel increments, with a secondary random offset added for each star ranging from 0 to 1 pixel in both x and y. In addition, the stars were chosen to uniformly sample a range of count rates between 0.01 and 10

counts/second, to cover a broad dynamic range while avoiding saturation.

The above procedure defined a grid of simulated stars on the sky, which were then moved around the detector for each exposure, according to the dither pattern being investigated, to produce a set of simulated dithered images in each filter for each dither pattern. The total exposure time corresponded to 9,000 seconds, with exposure time evenly divided among the pointings in each dither pattern. Each simulated exposure contained a realistic detector noise model, including dark current and read noise corresponding to the total exposure time, together with the number of samples and the number of reads per sample.

Once a set of simulated dithered images had been produced for each dither pattern and filter, the next step consisted of combining them using the MultiDrizzle software (Koekemoer et al. 2002). For the simulated JWST NIRCам images to be successfully processed and combined using MultiDrizzle, it was necessary to modify the image headers to include specific header keywords, keyword values. The appended file extensions include: SCI, ERR, DQ, SAMP, and TIME, corresponding to the current HST/NICMOS data structure. The added/modified header keywords include: instrume, nextend, camera, photplam, unitcorr, idctab, pa_v3, ra_targ, dec_targ, adcgain, filter, origin, iraf-tlm, date, detname, exptime, numutr, numsamrd, crpix1, crpix2, crval1, crval2, ctype1, ctype2, cd1_1, cd1_2, cd2_1, cd2_2, orientat, instrume, nextend, camera, photplam, unitcorr, idctab, ra_targ, dec_targ, adcgain, and pa_v3.

For each dither pattern, an output scale of 0.5 was used, in order to afford a direct comparison of the fidelity with which different dither patterns allow the input properties to be recovered. For the half-pixel dither patterns, the natural output pixel scale is 0.5 (relative to the input pixel size), since the half-pixel steps provide sufficient information on this scale. For finer patterns, in particular the 9-point pattern, slightly smaller scales could in principal be investigated, but in practice there is little to be gained by going much below 0.4 pixels since the PSF is generally Nyquist sampled at smaller scales, therefore the 9-point patterns were also drizzled to a scale of 0.5 in order to provide a direct comparison with the other patterns. The other primary parameter was “pixfrac” which is simply the amount by which input pixels are “shrunk” before being mapped onto the output pixel plane, thus smaller values can reduce unnecessary convolution of the final image by the input pixel size. Choosing too small a value of pixfrac can lead to undesirable large variations in the output weight map, therefore this parameter is generally chosen between about 0.5 and 1.0. In these simulations, the value of pixfrac was set to 1.0 for the 2-point and 3-point dither, as well as one of the 4-point dither patterns, and was set to 0.8 for the other 4-point dither pattern and the 9-point pattern.

Initially, a test run was conducted using images simulating a 2-point dither, displaying a grid of 16 stars with a uniform count rate based on the NIRCам short wavelength detector, using the F070W filter (centered at 0.7 μm with a width of 0.15 μm). Once the code had been fully verified to work with this limited set of data, the full set of simulated images were created, for all the filters and dither patterns, and were all successfully combined with MultiDrizzle to produce a set of drizzled output images for all the filters and dither patterns. A total of 154 such simulated 2048x2048 exposures were generated, for all 7 filters and the various different dither patterns investigated in this study, each

with 1024 stars distributed across it.

In Figure 1 we show one of the input exposures, for the NIRCcam F270W filter, showing the full 2048x2048 array with the grid of 1024 stars spread across the chip, sampling count rates between 0.01 counts/second and 10 counts/second. We also show the combined drizzled image after having combined a series of such dithered exposures with the MultiDrizzle software, which produces an image covering 4096x4096 pixels (at a scale of half the original input pixel size), thereby improving the sampling and resolution of the original data.




Figure 1. (Left) One of the simulated input exposures, for the NIRCcam F270W filter, showing one of the SCAs (2048x2048 pixels), with the grid of 1024 stars distributed across the detector. The stars are uniformly selected from a range of different count-rate values between 0.01 counts/second and 10 counts/second. Detector-related properties such as read-noise and dark current are appropriately modeled and added to the image. (Right) The combined drizzled output image, after having used MultiDrizzle to combine a set of dithered input images such as those on the left. The MultiDrizzle output image has been drizzled to a scale of 0.5 the original pixel scale, thus it has a size of 4096x4096 pixels. Since it is the combination of several input exposures, its background r.m.s. noise is lower, and its smaller pixel size also provides better sampling of the PSF than each of the original input exposures.

In Figure 2 we show an close-up example of the input dithered exposures, together with the combined output drizzled image, for a simulated star imaged using a 1/2-pixel 4-point dither for the NIRCcam long-wavelength channel F270W filter. The 4 input exposures in this particular example were drizzled onto a combined output image using an output scale of 0.5 input pixels, and a pixfrac value of 0.8, thereby recovering some resolution by reducing the amount by which the output image is convolved by the input pixel size.




Figure 2. Example of a simulated 4-point dithered exposure of a star imaged with the F270W filter on the NIRCAM long-wavelength channel, taken from 4 different input exposures such as those shown in Figure 1. The top 4 panels show the 4 individual exposures, including detector-related effects such as read-noise and dark current, as well as object photon statistics. The bottom two exposures show the output drizzled image (at half the input pixel scale), at two different greyscale stretches to reveal the extended large-scale PSF structure (bottom left) as well as the improved sampling in the PSF core (bottom right).

5.1.2 Analysis of simulated images

Once all the drizzled output images had been created for each set of input dither patterns and NIRCcam filters, the properties of the stars were measured on these images using the IRAF DAOPHOT package (Stetson 1987) as well as the SExtractor software (Bertin & Arnouts 1996).

The three principal quantities to be measured consisted of the positions, magnitudes and FWHM values of the stars on the original dithered exposures as well as the final combined, drizzled images. The SExtractor software was used to provide FWHM measurements for all the stars in each dataset, which it calculates based on an assumed Gaussian profile for the stars. Although the JWST/NIRCcam PSF has significant non-Gaussian components, the FWHM is nonetheless a useful parameterization for the purpose of evaluating the relative performance of different dither patterns, where the primary concern is to be able to provide a quantitative relative comparison.

The DAOPHOT package was used to measure the centroid and photometry values of stars, given the input PSF size for each filter. The reason why DAOPHOT was chosen for this instead of SExtractor is because the latter is written for general sources (including extended emission), whereas DAOPHOT is optimized for unresolved sources and therefore provides a more reliable measure of the centroid and photometry values of stars. However, DAOPHOT does not allow the FWHM to be fitted for, which is why SExtractor was used to measure that quantity.

For each dither pattern, and each filter, the astrometry, photometry and FWHM values were measured in the individual exposures as well as the final combined drizzled images, and were compared to the input values that were specified for the 1024 simulated stars in each set of exposures. An example of these comparisons is shown in Figure 3, for the F270W filter on the long-wavelength channel of NIRCcam, using a 2-point dither pattern. A total of 35 such sequences of plots were generated, covering all 7 filters and all the dither patterns.

The example in Figure 3 shows the difference in measured positions of objects on the exposures relative to their input positions, in units of image pixels. The histograms in the top of the figure, together with the plots below it, show that the measured positions of most of the 1024 stars agree with their input positions to below 1 pixel, with a few percent of the very faintest stars scattering to larger deviations. This can also be seen from the plots of x and y offset as a function of count-rate, where most of the brighter stars (above about 0.1 counts/second) show good astrometric behavior, while larger errors only appear for the faintest sources. The primary measurement from this set of plots is to provide the degree of dispersion of the measured astrometry in the post-drizzled images, as measured for different dither patterns and different filters.

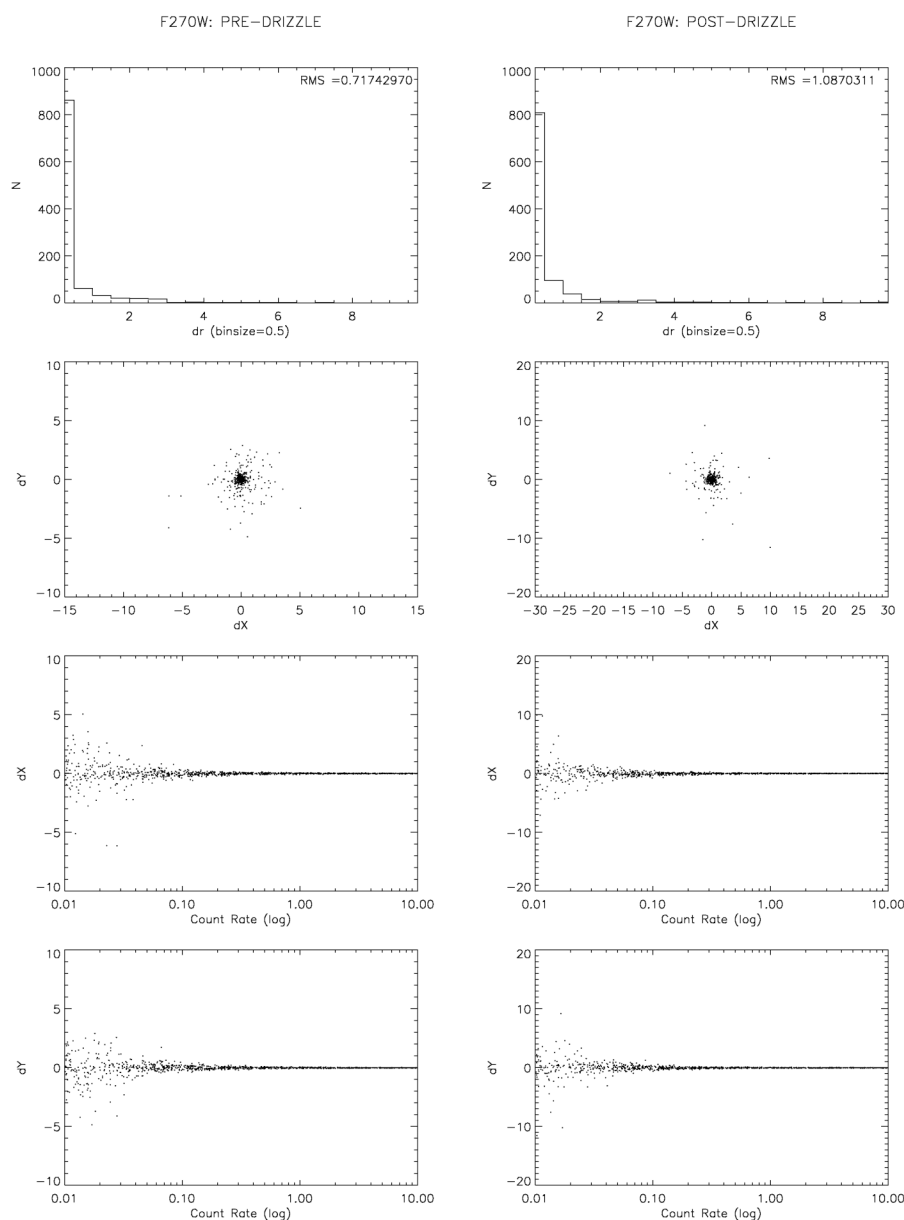


Figure 3a. Results from measuring the x and y positions, in image pixels, for a 2-point dither pattern using the F270W filter on NIRCcam. Note that the drizzled image has a pixel scale 0.5 times that of the input exposure, thus the dX and dY scales are different by a factor of 2. The scatter plot (second row) shows that the distribution is tighter in the drizzled image, which is expected since it has better sampling than the input exposures.

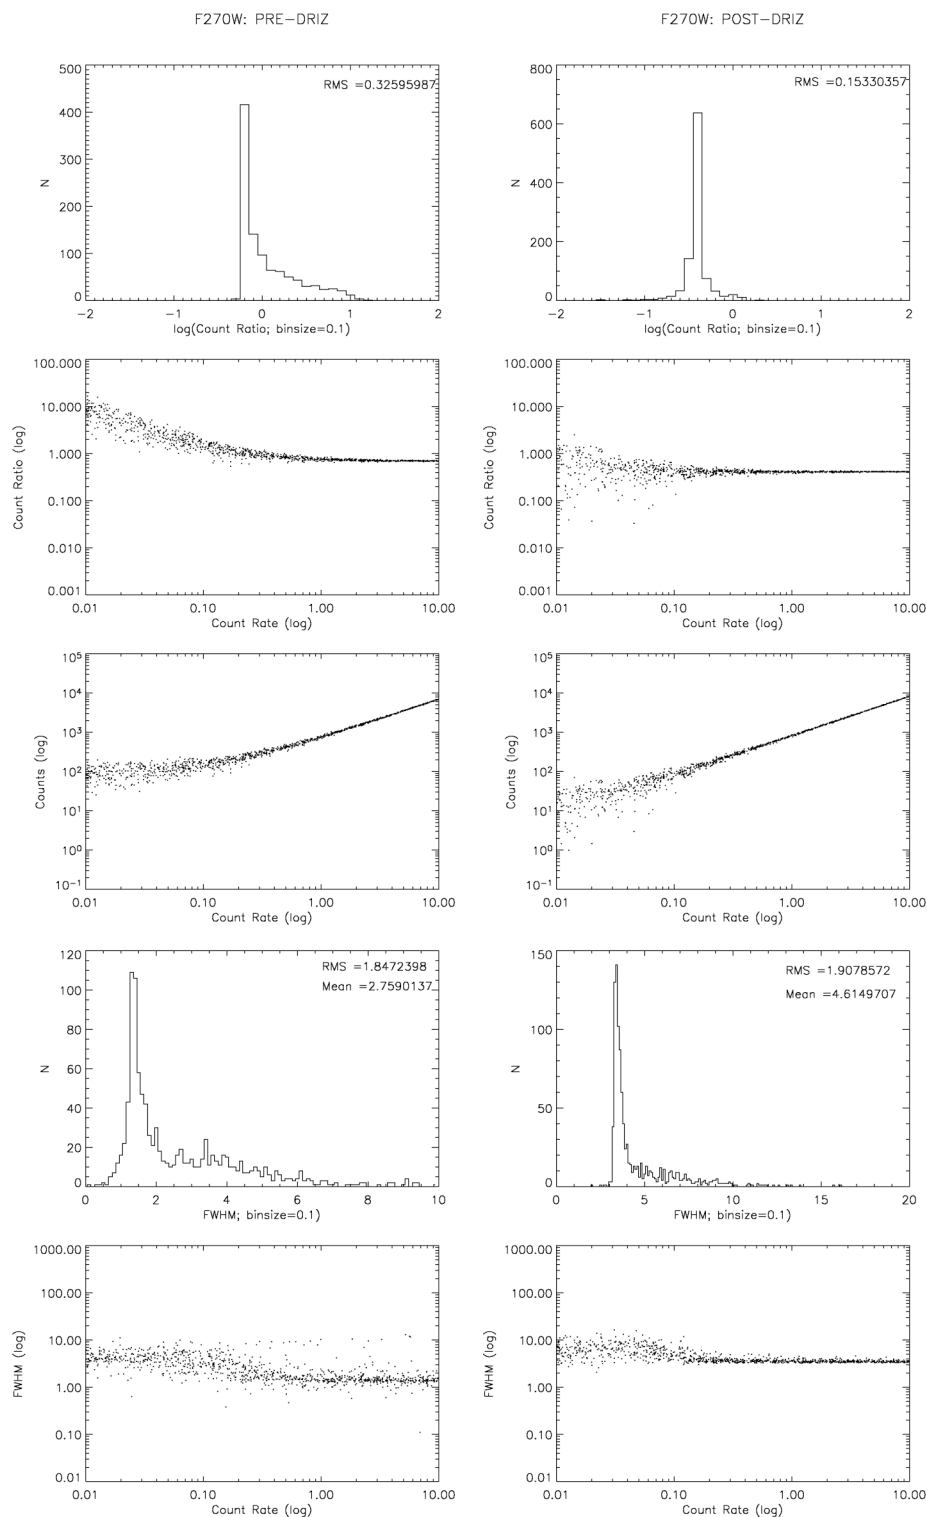


Figure 3b. As for Figure 3a, but showing the measured photometry (top three rows) and FWHM values (bottom two rows). Note that some of the fainter sources display certain noise-related systematics, therefore only sources above 0.1 counts/sec are included in the remainder of this analysis.

Similarly, the panels in the top three rows of Figure 3b show the ratio of measured count-rate relative to the input count-rate that was specified for the stars. At the faintest end there are some systematics related to the behavior of DAOPHOT, and in addition there are some differences in normalization that are simply related to the way in which the count-rates are defined. However, the primary measurement from these panels is to provide the degree of dispersion of the histograms for the brighter sources, which are taken to be those with count-rates above 0.1 counts/second. This dispersion was measured for all the drizzled images, for all the dither pattern and filter combinations.

The final of the three metrics is shown in the bottom two panels of Figure 3b, and represents the measured FWHM of the sources relative to the theoretical optimal PSF FWHM. Again, there are effects at the faint end that cause the measured FWHM to broaden significantly since pixel noise begins to become dominant, but for brighter sources the measurements are generally well behaved and can thus be used as a reliable metric of output FWHM in the drizzled images for different drizzle parameters, filters and dither patterns.

5.1.3 Results from Image Analysis

For each of the different dither patterns and filters, the metrics defined earlier for FWHM, astrometry and photometry could be derived directly from the histograms presented in plots such as those shown in Figure 2, avoiding stars at the faintest end where noise becomes dominant, but being able to use the majority of stars where the count-rate was sufficiently high (typically above 0.1 counts/second) to enable reliable measurements to be obtained.

For the FWHM, the average values measured in the drizzled images were compared with the theoretical FWHM values expected on the basis of the knowledge of the original wavelength-dependent PSF FWHM together with the subsequent resampling by the detector pixels in either the short-wavelength or long-wavelength channel, as appropriate. This therefore provides a direct measurement of the metric described in sub-section 4.1, for each dither pattern and filter specification.

For the astrometric and photometric measurements, the metric of choice is defined slightly differently as the degree of dispersion of the measured astrometry and photometry values for sources in the drizzled image, relative to what would be expected in the limiting case of the ideal convolved image I_C , as described in sub-sections 4.2 and 4.3, where the only error contribution is from the source photon statistics. Thus, the measured width of the histograms of astrometry and photometry in the drizzled images, derived from plots for all the filters and dither patterns such as the example shown in Figure 3, was compared with the histogram that would be expected if the measurement errors were purely the result of photon noise.

The resulting values for the metrics describing FWHM, astrometry and photometry, for all the broad-band NIRCcam filters, are shown in Figure 4 as a function of dither pattern. This plot summarizes a total of 35 plots similar to Figure 3, which were each derived from measurements of the 1024 stars in the various dither pattern / filter combinations.

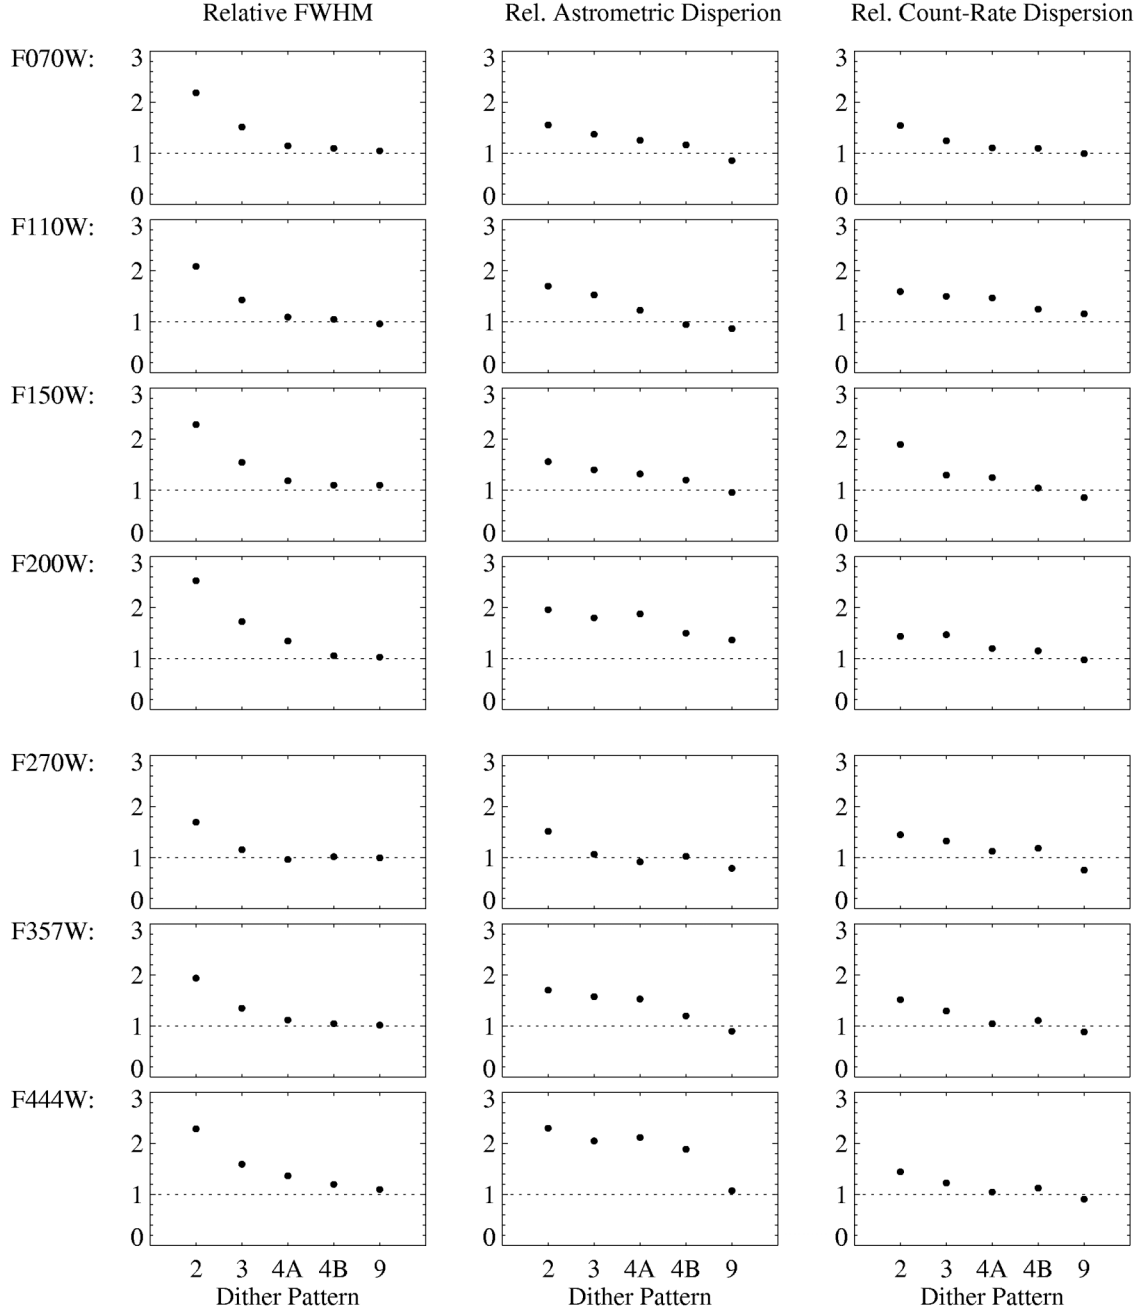


Figure 4. Final results for the three metrics of FWHM, astrometry and photometry, as a function of all the sub-pixel dither patterns investigated in this study, covering all seven broad-band filters in the NIRCcam short-wavelength and long-wavelength channels. The first column shows the measured FWHM relative to the ideal FWHM that can be recovered. The second and third columns show the dispersion of the measured astrometry and photometry in the drizzled image, relative to the dispersion that would be obtained in the ideal case of infinitely well-sampled data, drizzled using the same drizzle parameters. It can be seen that the 4-point sub-pixel dither patterns generally perform better than 2 or 3 point patterns and are close to optimal, being only slightly outperformed by 9-point patterns.

Figure 4 yields slightly different optimal patterns depending upon whether FWHM resolution or photometric fidelity are the most important quantity. For FWHM, it is clear that a significant improvement is generally obtained from 2-point to 3-point dither patterns, with further improvement by going to a 4-point pattern which essentially yields near-optimal results, for the drizzle parameters under consideration; a 9-point pattern generally does not yield any further significant gains for recovering FWHM. This can essentially be interpreted as stating that the information necessary to recover the Nyquist-sampled FWHM is effectively obtained by a 4-point pattern with half-pixel subsampling, and that a 9-point pattern with 1/3 pixel subsampling does not provide significant additional gains when attempting to recover the original FWHM.

However, for photometric fidelity (i.e., the third column in Figure 3), the improvements from more finely spaced dither patterns do not approach the ideal case quite as quickly as for FWHM, with the result that for a few filters the 9-point pattern provides an improvement over the 4-point pattern. In other words, the relative dispersion of the measured photometry in the drizzled images only begins to approach the ideal dispersion when using a 9-point pattern. Most likely this reflects the fact that photometric fidelity is driven not so much by Nyquist sampling as by pixel-to-pixel effects. At the same time, however, the relative count-rate dispersion with the 4-point pattern is generally no more than about 5-10% higher than that obtained with the 9-point pattern. In other words, an observation sequence divided into a 4-point dither pattern will reach a photometric r.m.s. that is about 5-10% higher than that achieved with a 9-point pattern, assuming ideal pointing offsets are achieved in both cases. Thus, for example, if a given observation sequence would achieve 1% photometric precision with the 9-point dither pattern, then a slightly lower precision of 1.05 – 1.1% would be achieved by dividing the same exposure time into a 4-point dither pattern instead.

5.1.4 Final Evaluation of Metrics for NIRCcam Dither Patterns

Having quantified the benefits achieved by different dither patterns, the final step in the process of evaluating them is to determine the costs associated with different observing strategies. Here we assume three representative types of observing programs that are described in Section 9 of the NIRCcam Operations Concept Document (OCD), Version 1, October 20, 2003 (P. McCullough), and calculate the overheads for filter wheel moves and dither offsets for each type of program, for the different dither patterns investigated.

The three types of programs described in the NIRCcam OCD are categorized as “short”, “medium” and “long”. In each program, the dither pattern is repeated at three different locations in order to cover the gaps between the SCAs. The short program assumes a visit with a total elapsed observing time of 3,000 seconds, using 2 filters, which might be typical for a survey of bright stars, for example. The medium program assumes a visit with a total of 20,000 seconds elapsed observing time, using 5 filters, which might be representative of programs obtaining multi-wavelength coverage of moderately faint extragalactic targets. The long program assumes a total of 100,000 seconds elapsed observing time but is otherwise the same as the medium program. Here we use the same assumptions described in the NIRCcam OCD, namely that filter wheels and pupil wheels move in series and require 30 seconds each, and that the dither time overhead between

exposures is 10 seconds. For each observing program, we assume two scenarios “A” and “B”: in Scenario A, the filters are cycled through at each dither pattern, as described in the NIRCcam OCD, while in Scenario B the filters are changed as infrequently as possible, hence the dither pattern is repeated for each filter. In Table 3 we summarize the number of filter wheel and pupil wheel moves for each of these scenarios, for an N-dither sub-pixel pattern, together with the total overhead time for all the mechanism moves.

Table 3. Mechanism moves and exposure/overhead times for N-point dither patterns

	Total elapsed observing time (sec)	Number of filters	Number of dither moves	Number of filter wheel moves	Number of pupil wheel moves	Total overhead time (sec)
Short - A	3,000	2	3xN	2xN	2xN	150xN
Short - B	3,000	2	2x3xN	2	2	120 + 60xN
Medium - A	20,000	5	3xN	5xN	5xN	330xN
Medium - B	20,000	5	5x3xN	5	5	300 + 150xN
Long - A	100,000	5	3xN	5xN	5xN	330xN
Long - B	100,000	5	5x3xN	5	5	300 + 150xN

Finally, we calculate the overheads required for the various different scenarios, and represent them as a fraction of total elapsed time (including all overheads plus observing exposure time, for all the filters combined) for each of the different dither patterns. In addition, we calculate the average relative FWHM, relative astrometric r.m.s. and relative photometric r.m.s. as was presented in Figure 4. The results of these comparisons are shown in Table 4.

Table 4. Final Cost/Benefit Metric Evaluations for NIRCcam Dither Patterns

	2-point	3-point	4-point	9-point
Benefits:				
Average relative FWHM	2.14	1.47	1.08	1.03
Average relative astrometric r.m.s.	1.76	1.54	1.27	0.97
Average relative photometric r.m.s.	1.55	1.34	1.14	0.94
Costs: overhead as a fraction of total elapsed time				
Short - A	10.0%	15.0%	20.0%	45.0%
Short - B	8.0%	10.0%	12.0%	22.0%
Medium - A	3.3%	5.0%	6.6%	15.0%
Medium - B	3.0%	3.8%	4.5%	8.3%
Long - A	0.7%	1.0%	1.3%	3.0%
Long - B	0.6%	0.8%	0.9%	1.7%

Several trends are apparent from the values in Table 4. Firstly, there are noticeable improvements in all three of the measured observational metrics (FWHM, astrometry and photometry) between the 2-point, 3-point and 4-point patterns. In addition, the relative FWHM (i.e., as measured in the drizzled image, compared with the ideal measurements) approaches the ideal case very quickly in the 4-point dither pattern, and not much extra is

gained by the 9-point pattern. By contrast, the relative dispersion of astrometric and photometric measurements continues to improve beyond the 4-point pattern, with the 9-point pattern yielding significantly better results.

The overhead costs behave in two different ways, depending on whether the filters are changed at each dither pointing (Scenario A) or whether instead all the dither patterns are cycled through before changing a filter (Scenario B). In the first scenario, the overhead time for filter change becomes a fixed fraction of the overhead time for a given dither offset, and therefore the fractional overhead increases linearly as a function of the number of dither positions given that the total elapsed observing time is constrained to remain the same. In the second scenario, the overhead also increases for a larger number of dither positions, as expected, but it increases more slowly and the total overhead is lower because the overhead required for each dither offset is less than the overhead for a filter change, and the filters are changed less frequently.

It is immediately clear that certain observing strategy / dither combinations are *not* optimal, for example a short program using 9 dither points and changing the filters at each dither location, which would use up 45% of the observing time in overhead. On the other hand, a short program using a 4-point dither would benefit by almost a factor of 2 in overhead (12% vs. 20%) if the filters were changed only once during the visit (i.e., Scenario B), instead of at every dither position. Thus, for short programs, if a 4-point or 3-point pattern is being used, then changing filters infrequently and repeating the dither pattern for each filter would provide significant gains in overhead as compared with changing the filter at each dither position. For a short program using a 2-point dither, the difference is less significant, therefore filters could be changed at each dither position without much additional penalty if there was sufficient need, for example to improve the degree of uniformity of data taken with different filters if there are significant time-dependent effects during the observation.

For medium and long programs using 2, 3 or 4-point dither patterns, the relative difference in overhead between Scenarios A and B is generally less than 20-30%, and the total overhead remains below ~5-6% at most, therefore for such programs Scenario A might generally be preferred, again to help provide more uniformity between data taken with different filters. For medium and long programs, a 4-point pattern would generally appear to provide more benefits over the 2- and 3-point patterns, given the relatively small additional overheads.

When considering the 9-point pattern, the overheads are about a factor of 2 worse than for 4-point patterns; for medium programs this can be significant (8.3% – 15%, depending on whether Scenario A or B is used) and therefore it is likely that most medium programs would not obtain much additional benefit from the 9-point pattern, relative to the increased cost, thus for medium programs the 4-point pattern appears optimal, under the assumptions made about the observing program details.

On the other hand, long programs only pay 1.7% – 3% in overheads for the 9-point pattern, whereas the benefits obtained from it would most likely generally be considered to compensate for this, therefore for long programs the 9-point pattern appears optimal.

5.2 FGS-TF

The FGS-TF is the most severely undersampled instrument, since it has detector pixels of 65 mas and therefore does not provide Nyquist sampling for any wavelengths in its short-wavelength channel (1.0 – 2.1 μm); the long-wavelength channel (2.1 – 4.9 μm) is Nyquist sampled above about 4 μm . The FGS contains a total of four 2048x2048 SCAs, two of which will be used for guiding while the other two correspond to the short- and long-wavelength channels for imaging light from the tunable filter, which is separated into the two wavelength channels by a dichroic.

For the purposes of discussing dither patterns, the principal differences between the FGS-TF and NIRCам are as follows:

- FGS-TF only has 65 mas pixels, whereas NIRCам also has 31 mas pixels;
- NIRCам has an array of 2x2 SCAs with a gap between them, while FGS-TF has only one short-wavelength SCA and one long-wavelength SCA that image the same region of sky in different wavelength ranges, therefore there is no gap that needs to be covered;
- Since the tunable filters have a resolution $R \sim 100$, the FGS-TF background will be much less severe than with NIRCам and read-noise may dominate, thereby potentially impacting the motivations for dithering since flat-field systematics will have less effect on photometric accuracy than with NIRCам;
- The FGS-TF will likely have much more pronounced optical ghosts as a result of the reflective optics in the tunable filter, therefore large-scale dithers designed to mitigate these will form an important part of the overall FGS-TF dither strategy.

At the time of writing, the code used to generate the simulated images for NIRCам was not yet capable of generating similar images for FGS-TF, therefore quantitative measurements of the relative merits of different dither patterns could not yet be performed. Hence it is also not yet feasible to perform the quantitative comparison of overhead times for different dither and filter-changing strategies. However, some general conclusions can still be derived about likely dither patterns for FGS-TF, based on the current knowledge of the instrument.

The fact that both the short- and long-wavelength channels in FGS-TF have the same pixel scale provides a significant simplification in designing dither patterns, since the same degree of sub-pixel sampling can be achieved in both channels simultaneously. Since the short-wavelength channel is so severely undersampled, it is likely that most or all sub-pixel dithering in FGS-TF will be driven by the need to optimize the spatial sampling of the short-wavelength channel. It is likely that the metrics for the 4-point and 9-point dither patterns will tend to favor the 9-point 1/3-pixel pattern more strongly than for NIRCам, specifically for medium-length observations. Thus, while the 4-point pattern appeared optimal for the prototypical medium-length NIRCам observing strategy

discussed in Section 5.1, it is likely that a similar prototypical FGS-TF medium-length observation might instead favor the 9-point pattern.

Similarly, short observations (less than a few thousand seconds) with FGS-TF may achieve sufficient benefits from the 4-point 1/2-pixel pattern that these will outweigh the increased overhead costs in obtaining them, if the overheads for FGS-TF are comparable to those for NIRCам.

It should be noted that the distortion across the FGS-TF detectors is expected to be on the level of $\sim 2\%$, thus comparable to that of NIRCам. Therefore if the dither patterns are kept to less than a few pixels in their total extent, then the degree of subsampling achieved will remain comparable across the detector. For example, a 1/2-pixel 4-point dither covering 2.5 pixels at the center of the chip will correspond to about 2.55 pixels at the corner, a difference of about 3.3 mas and less than other uncertainties related to the pointing accuracy. Thus, both the 1/2-pixel 4-point dither and the 1/3-pixel 9-point dither should provide nearly uniform sub-pixel sampling across the entire detector if the total size of the dither pattern is kept below a few pixels.

Therefore, in summary, it is likely that FGS-TF observations will favor dither patterns that provide one additional level of subsampling as compared with similar length observations on NIRCам, in the sense that short observations will likely benefit from 1/2-pixel 4-point dither patterns, while medium- and long-timescale observations will both benefit from 1/3-pixel 9-point dither patterns. Once the simulation code used in this study is capable of producing simulated FGS-TF images, it may be worthwhile performing a follow-on study similar to that carried out for NIRCам, to specifically measure the quantitative improvements in FWHM, astrometry and photometry that would be provided by different dither patterns.

5.3 MIRI

The MIRI instrument provides four general types of observations, using apertures as shown in Figure 5:

- Direct imaging, provided by the MIRI imager (MIRIM), making use of a single 1024x1024 pixel sensor chip assembly (SCA) covering a 113" field of view with 110 mas pixels across a 5 – 27 μm wavelength range, Nyquist sampled above 7 μm ;
- Coronagraphic masks located along one edge of the imager, each covering a field of view of 26" – 30"
- Low-resolution spectroscopy (LRS), obtained with a 0.6" x 5" aperture along the same side as the coronagraphic masks;
- Integral Field Unit (IFU) spectroscopy which uses two 1024x1024 SCAs sensitive from 5 – 27 μm , with the wavelength range divided such that one SCA covers the range 5 – 12 μm (with 196 mas pixels) while the other covers 12 – 27 μm (with 245 – 273 mas pixels).

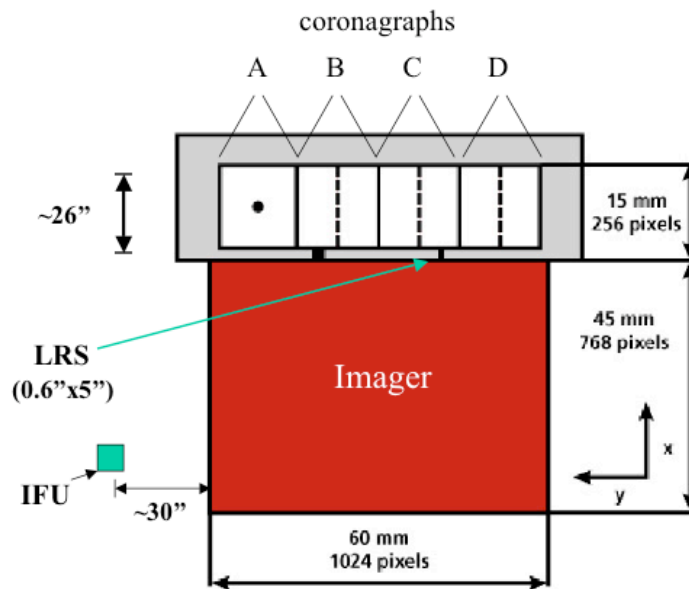


Figure 5. MIRI focal plane layout (from Nelan et al. 2005). The IFU, LRS, and 4 coronagraphic masks are shown along with MIRI's Imager (MIRIM). Note that the coronagraphs and LRS occupy one edge of the 1024x1024 SCA used for imaging, thereby reducing its field of view by about 20%. The IFU feeds light to two different 1024x1024 SCAs, not indicated here.

The code used in this study was not yet capable of generating simulated images for any of the MIRI observational modes; therefore quantitative metrics cannot yet be measured or compared for different dither strategies. Therefore, the discussion of dither patterns in this section is instead aimed at a more general description of the relative expected benefits and costs of various strategies.

The imaging mode is relatively well sampled across most or all of its 5 – 27 μm wavelength range, since the 110 mas pixels provide Nyquist sampling for all wavelengths above about 7 μm . Therefore, it is expected that sub-pixel sampling will be a less dominant driver for the MIRI dither patterns than for NIRCам or FGS-TF. If sub-pixel sampling is required for the imaging mode, it will likely consist of 1/2-pixel subsampling obtained using either a 2-point or 4-point pattern, depending on whether the observations are short or long, and will yield the most benefit at wavelengths shortward of 7 μm . It is unlikely that there would be any strong need for sub-pixel scales finer than 1/2-pixel subsampling (such as the 1/3-pixel subsampling discussed earlier for NIRCам and FGS-TF). Otherwise, dithering on small scales will be primarily used to mitigate bad pixels and other detector defects.

An important additional driver for dithering in MIRI imaging mode is the ability to remove background fluctuations, which may vary both as a function of time and location across the detector, on timescales comparable to typical exposures. These can arise from thermal emission inside and outside MIRI, stray light reflected from the radiation screens and the telescope, and ghost images from the filters. Observational techniques to mitigate these will consist of large-scale dithers comparable to the size of the SCA. Large-scale dithers, comparable to the size of the detector, would also be used to mosaic or map large regions of the sky, and would make use of similar considerations to those discussed for the other two imaging instruments, NIRCам and FGS-TF.

A more complex scenario is presented by the MIRI IFU, where the spectral window is divided into a total of four wavelength channels as follows:

- 1A: 5.00 – 7.71 μm , 3.7"x3.7" FOV, 176 mas slice width, 196 mas pixels
- 1B: 7.71 – 11.89 μm , 4.5"x4.5" FOV, 277 mas slice width, 196 mas pixels
- 2A: 11.89 – 18.35 μm , 6.1"x6.1" FOV, 387 mas slice width, 245 mas pixels
- 2B: 18.35 – 28.30 μm , 7.7"x7.7" FOV, 645 mas slice width, 273 mas pixels

The 1A and 1B channels are imaged on two halves of one of the SCAs, while the 2A and 2B channels are imaged by the second SCA. Furthermore, each of the four wavelength channels is divided into three spectral bands by means of dichroics and gratings. Therefore, a single exposure will obtain 1/3 of the spectral range for each of the four wavelength channels 1A, 1B, 2A, and 2B, imaged on the two SCAs. Full wavelength coverage would be obtained by obtaining three separate exposures, with each one sampling one of the three spectral bands. An example of this is shown in Figure 6 for a single exposure, showing only the SCA that covers the 1A and 1B wavelength channels.

The wavelength coverage for 1A in this example ranges from 5.00 – 5.79 μm , and for 1B ranges from 7.71 – 8.93 μm , thus 1/3 of the total spectral coverage of each.

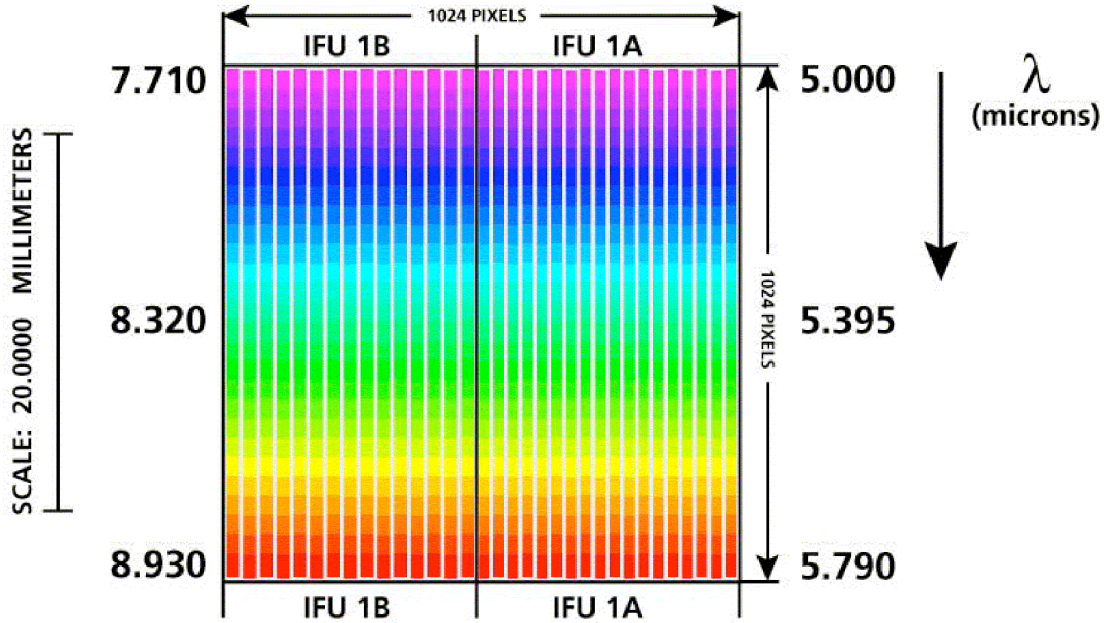


Figure 6. Spectral and spatial layout on one of the two SCAs of the MIRI IFU from the MIRI OCD (31 May 2005, M. Meixner), showing an example of the spectral coverage obtained in one exposure in the 1A and 1B wavelength channels. The spectral coverage in each case is 1/3 of the total spectral coverage available for these wavelength channels, therefore full spectral coverage would be achieved by obtaining three separate exposures with different wavelength coverage selected. The spatial information from the IFU is spread out into adjacent rows with the information from each slice separated by a small gap, while the spectral information runs along the columns as shown. The other SCA covering wavelength channels 2A and 2B is similar.

The four IFU slice widths are chosen to match the PSF FWHM at the shortest wavelength end, thus channel 1A has a slice width of 176 mas corresponding to the assumed PSF FWHM at 5 μm . The slice width of the subsequent channels is therefore chosen to scale as a function of the wavelength, by factors of 11/7, 11/5 and 11/3 in order to provide the best sampling using a 3-point dither sequence. The optimal 3-point dither sequence for the IFU is defined and described in a detailed treatment in MIRI-TN-0001-ATC (Issue C, December 8, 2003, A. Glasse), and therefore only the relevant points are summarized here.

In order to carry out optimal spatial sampling across each slice, the telescope must be moved by $N_i+1/2$ slice widths, where N_i may represent a different integral slice width for each of the four channels. Operationally, the most efficient strategy is to be able to achieve optimal sampling for all four channels simultaneously with a single telescope offset X_{tel} , and this can be achieved by choosing the slice widths a_i for each of the four channels such that $a_i(N_i+1/2) = X_{tel}$. This therefore provides the values for the four slice widths, such that optimal sub-pixel sampling in all four channels can be achieved with a

three-exposure dither sequence where the second dither position is 0.95'' across the slice and the third dither position moves an additional 1'' along the slice.

If the characteristics of either the telescope PSF or the IFU optics change significantly, then the above dither pattern would need to be revised, along with a likely revision to the slice with properties. However, it should in general remain possible to find a combination of slice width and dither offset size that will provide optimal or near-optimal sampling across all four wavelength channels simultaneously.

Finally, MIRI contains a low-resolution spectrometer (LRS) that has a slit with 0.6'' width and 5'' length, providing $R \sim 100$ spectroscopy across the 5 – 10 μm range. The spectrum is spread over 380 pixels on the SCA, and should therefore provide sufficient sampling of the spectral LSF that sub-pixel dithering along the spectral direction should generally not be necessary. Along the slit direction, the pixels provide a spatial sampling of 110 mas while the PSF FWHM ranges between 176 mas at 5 mm to 350 mas at 10 mm. Thus, at the shortest wavelength end of the spectrum, the spatial dimension of the slit is not quite Nyquist sampled and observers may wish to make use of a dither pattern that provides 1/2-pixel subsampling along the slit, along with integer pixel offsets to help mitigate bad pixels and cosmic rays. Generally, however, the spatial distribution of the spectrum along the slit will likely be co-added into a one-dimensional spectrum, therefore it is not anticipated that sub-pixel dithering will be an important observational strategy for the LRS mode.

5.4 NIRSpec

A significant difference between NIRSpec and the other JWST instruments is its micro-shutter array (MSA) consisting of 750x350 shutters, each of which is 0.2"x0.45" in size, with 0.25"x0.5" spacing. The NIRSpec detector pixels are 100 mas in size, which undersample the PSF across its entire wavelength range of 0.6 – 5 μm . Therefore, dither patterns for NIRSpec are concerned not only with sub-pixel sampling but also with accounting for variations in light due to where the targets fall within the MSA shutter apertures, and whether science data is desired for a single object or for a distribution of targets across the field.

A detailed discussion of potential dither strategies for NIRSpec is presented in STScI-JWST-TM-2005 (Draft 2.0, 12 July 2005, M. Regan), therefore only the most relevant aspects of this treatment are summarized here. The baseline observing scenario involves placing a target at the "sweet spot" across the shutter, where an error in positioning from target acquisition would yield less than a 10% error in flux. However, in this scenario, larger pointing uncertainties decrease the size of the sweet spot, and as a result an alternative observing strategy is discussed in the aforementioned document that makes use of dither patterns to mitigate potential pointing uncertainties by obtaining a series of dithered exposures with the target at different locations across the shutter.

The primary advantage of using dither patterns to sub-sample the MSA shutters is that it provides a more accurate flux calibration since several samples are obtained at different locations across the shutter response function. An analogy in the imaging domain would be the case of severely undersampled pixels with an unknown intra-pixel response function that could introduce errors into the photometry, that are best mitigated by moving the target around to different locations within sub-pixel space. One disadvantage of this approach is that the overall throughput is about 20% lower since some of the exposures will be obtained near the edges of the shutter response function. In addition, the MSA may need to be reconfigured in some cases to ensure that the correct shutters are opened after certain dither offsets are performed.

Nevertheless, the improved photometric accuracy from such dithering may offset its disadvantages for many science programs, including programs with multiple targets across the field of view. In programs with sufficient numbers of sources across the field of view, it is possible that about 90% of the targets may be observed without requiring any MSA reconfigurations, which would provide a significant advantage in terms of overhead cost savings. At the time of writing, the simulation code used in the present study was not yet capable of producing simulated NIRSpec data, but if this were incorporated in future then it may be worthwhile extending the optimization study to provide quantitative measurements of the relative advantages of different dither strategies under this scenario.

6 Summary and Conclusions

This study is aimed at an initial examination of defining and using metrics to describe and quantify the relative merits of different dither patterns for the JWST instruments, allowing the development of optimal patterns that trade off scientific benefits against costs such as increased overhead. An important component of this study is the detailed definition of quantitative metrics that are defined in a sufficiently general way that they can be used to evaluate the merits of any specific dither pattern or observing scenario, given sufficient details about the instrumental properties and the observing overheads.

The most detailed treatment in this study is for the NIRC*am* instrument, where the instrumental and observational characteristics are sufficiently well defined that simulated images could be produced for a range of different dither patterns and filters, and were subsequently combined using the MultiDrizzle program in the same way as is currently done for dithered HST observations. This allowed quantitative measurements of the properties of simulated objects on these images, specifically their morphological, astrometric and photometric properties, in a way that allowed direct comparison between different dither patterns and the ideal case of infinitely well sampled data. For a variety of assumed observing scenarios, drawn generally from considerations in the DRM, these measurements allow a quantitative comparison between the benefits obtained from dithering against the fractional increase in overhead time associated with each observing scenario. The general results are that, for NIRC*am*, short and medium observing programs may benefit most from 2 to 4-point sub-pixel dither patterns, while long observing programs may benefit more from 9-point sub-pixel dither patterns with a relatively small additional cost in overhead.

The other instruments were also considered in this study, albeit in a more generalized fashion since the available code was not yet capable of generating simulated data for these instruments. For the FGS-TF, similar considerations apply as for NIRC*am* except that the FGS-TF pixels are more undersampled at the short-wavelength end, therefore short programs may generally need at least a 4-point sub-pixel dither while both medium and long observing programs would likely benefit sufficiently from using 9-point sub-pixel patterns to offset the increased costs in overhead.

The MIRI imaging mode is better sampled than NIRC*am* or FGS-TF, hence much of the discussion instead centers around its IFU where it is possible to define a combination of slice widths and dither steps to provide optimal sub-pixel sampling in all four wavelength channels simultaneously with a single set of telescope offsets.

NIRSpec is unique in that its MSA provides an additional set of requirements for dithering; the general motivation is to improve photometric fidelity by moving sources to a number of locations within the shutter apertures, to mitigate photometric uncertainties.

Finally, the present study was conducted using very preliminary information about many of the instruments, and a more detailed follow-on study may be worthwhile in future when the instrumental configurations are more stable.

7 References

- Arendt, R. G., Fixsen, D. J., Moseley, S. H. 2000, “*Dithering Strategies for Efficient Self-Calibration of Imaging Arrays*”, *Astrophys. J.*, Vol. 536, p. 500
- Bertin, E. & Arnouts, S. 1996, “*SExtractor: Software for Source Extraction*”, *Astron. & Astrophys. Supp.*, Vol. 117, p. 393
- Donahue, M. et al. 2002, “*Science Objectives and Requirements / Design Reference Mission*”, NGST-RQMT-000804 (Draft, 26 July 2002)
- Ferguson, H. C. 2000, “*Science Drivers for NGST Small-Angle Maneuvers*”, STScI-NGST-R-002A (Issue A, 19 April 2000)
- Fixsen, D. J., Moseley, S. H., Arendt, R. G. 2000, “*Calibrating Array Detectors*”, *Astrophys. J. Supp.*, Vol. 128, p. 651
- Fruchter, A. S. & Hook, R. N. 2002, “*Drizzle: A Method for the Linear Reconstruction of Undersampled Images*”, *Proc. Astron. Soc. Pac.*, Vol. 114, p. 144
- Glasse, A. 2003, “*The Choice of Slice Width and Field of View for the MIRI IFU*”, MIRI-TN-0001-ATC (Issue C, 8 December 2003)
- Lauer, T. 1999, “*Combining Undersampled Dithered Images*”, *Proc. Astron. Soc. Pac.*, Vol 111, p. 227
- Long, K. 2004, “*JWST Mission Operations Concept Document*”, JWST-OPS-002018 (Revision A, 2 July 2004)
- Lucy, L. B. 1974, “*An Iterative Technique for the Rectification of Observed Distributions*”, *Astronomical Journal*, Vol 79, p. 745
- Koekemoer, A. M., Fruchter, A. S., Hook, R. N., Hack, W. 2002, “*MultiDrizzle: An Integrated Pyraf Script for Registering, Cleaning and Combining Images*”, 2002 HST Calibration Workshop, eds. S. Arribas, A. M. Koekemoer, B. Whitmore (Baltimore: STScI), p. 337
- McCullough, P. 2003, “*NIRCam Operations Concept Document*”, JWST-OPS-002843 (Version 1, 20 October 2003)
- Meixner, M. 2005, “*MIRI Operations Concept Document*”, JPL D-25632 (Initial Release, 31 May 2005)
- Nelan, E. 2005, “*JWST Science Instrument Target Acquisition Concepts*”, STScI-JWST-R-2005-0001 (Issue A, 3 January 2005)
- Regan, M. 2005, “*An Alternative Observing Strategy for NIRSpec and its Effect on NIRSpec Target Acquisition*”, STScI-JWST-TM-2005 (Draft 2.0, 12 July 2005)
- Richardson, A. 1972, “*Bayesian-based Iterative Method of Image Restoration*”, *J. Opt. Soc. Am.*, Vol 62, p. 55
- Stetson, P. B. 1987, “*DAOPHOT – A Computer Program for Crowded-Field Stellar Photometry*”, *Proc. Astron. Soc. Pac.*, Vol 99, p.191

