TECHNICAL REPORT

Title: Dither Patterns for NIRCam Imaging		Doc #: JWST-STScI-001738, SM-12 Date: 30 June 2009 Rev: -
Authors: Jay Anderson	Phone: 410-338-4982	Release Date: 5 November 2009

1.0 ABSTRACT

Dithering will accomplish two major objectives for NIRCam: it will allow us to mitigate the effects of bad pixels and gaps between the detectors. and it will also allow us to improve the sampling of the astronomical scene. The current expectation is that all but a very small subset of NIRCam observations will be dithered as much as possible, meaning that no two exposures will be taken at the same pointing through the same filter. This strategy will allow the calibration pipeline to generate the most useful reduced data products both for the primary observer and for the archive.

Dithers in NIRCam will be divided into two classes: primary and secondary. The primary dithers will be larger and will mitigate the effects of detector gaps and low-frequency flat-field errors ("L-flat" errors) on the reconstruction of the astrometric scene. The secondary dithers will involve a pattern of small dithers superposed on each primary dither. These will be useful to improve sub-pixel sampling, while providing additional mitigation of bad pixels and pixel-to-pixel flat-field errors. An observer will select the combination of primary and secondary pattern that best addresses the aims of his or her scientific program.

This document presents a set of primary and secondary dither patterns that should be flexible enough to accommodate all currently anticipated scientific needs.

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

2.0 INTRODUCTION

A single exposure at a single pointing gives us a limited perspective on the astronomical scene. Bad pixels and gaps between detectors give us no information about certain parts of the field. In addition, if the detector pixels are too large, they may not sample all the detail of the scene that has been delivered by the telescope to the detector. Good dithering allows us to overcome all of these shortcomings. Each dithered observation constitutes a completely independent realization of the scene. By taking multiple exposures with offsets between them, we can regain information that was not available in a single exposure.

Dithering is commonly used on ground-based telescopes. In the optical it is used largely to mitigate the effects of bad pixels or bad columns. In the infrared it has historically been used to "chop" the scene and sky in an effort to mitigate the effects of the large and variable background. Since most ground-based detectors are carefully matched with the telescope to critically sample the scene, there is generally no need to dither to improve the sampling. Furthermore, since the seeing at most ground-based observatories varies from exposure to exposure, it is difficult to combine dithered ground-based observations into a better-sampled whole, since each exposure results from a different point-spread function (PSF), and therefore samples a different effective scene.

It was not until HST that the full usefulness of dithering became apparent. The stability of its undersampled PSF made it possible to use dithering to improve pixel sampling, as well as to remove artifacts. `Drizzle` (Fruchter & Hook 1997, 2002) is a popular routine that was designed to combine multiple dithered observations of the same scene into a single, more finely sampled representation of the scene. It has become a standard part of the HST pipeline, and is likely to be a standard part of the JWST pipeline, as well.

For all of the above reasons, it is currently envisioned that virtually all NIRCam observations will be dithered as much as possible. This means that no two exposures will be taken at the same pointing through the same filter. The only NIRCam observations that will not benefit from dithering are planetary-transit observations, which require the highest precision time-series differential photometry. All other observations will benefit from following an optimized dithering strategy.

The dithering prescription presented here will be two-fold. We will use large dithers ("primary dithers") to image over the gaps between detectors, and we will use small dithers ("secondary dithers") at each of the primary-dither locations to improve the sub-pixel sampling. Observers will carefully choose how many primary and secondary dithers to use in order to optimize their observations. The total number of dithers will be the multiplicative product of the number of primary and secondary dithers, $N_{TOT} = N_P \times N_S$, so we will need to allow a lot of gradation in each parameter to ensure flexible patterns that do not require an excessive total number of exposures.

The report is organized as follows. In Section 3, we will briefly contrast the NIRCam dithering strategy with what is currently in use on HST. In Section 4, we will discuss three different kinds of primary dither patterns that will dither over the detector gaps and move targets around the detectors. Section 5 will cover the secondary dithers, which will be executed at each of the primary-dither locations to provide better pixel-phase sampling, so that sub-pixel structure in the scene can be recovered. Section 6 briefly touches on the larger context of how primary and secondary dithers interact and how the primary-dither patterns can be used to tile large regions, including the special case of NIRSpec pre-imaging. Finally, we summarize the results in Section 7.

3.0 HOW JWST DITHERING WILL DIFFER FROM HST DITHERING

With HST, the user is free to specify any number of exposures with any specified `POS_TARG` (offset) relative to the nominal target location. This has given observers great freedom in tailoring complicated dither patterns to accomplish the specific scientific aims of their program. The downside is that without a clear dithering mandate, many HST programs have been taken without careful attention to dithering, and as a result, many observations now have limited value in the archive.

JWST will do things differently. To ensure the highest-quality data products for the GO and for the archive, we will include dithering as an integral part of the JWST planning tool. Observers will specify the readout pattern and integration time for each exposure, then they will specify the number of exposures to take by selecting the primary and secondary dither patterns. For example, if they choose a dither pattern with 3 primary and 4 secondary dithers, they will get 12 exposures. There will not be a standard mechanism for selecting the number of exposures independently of the number of dithers: every exposure will be taken with a different dither. Since different programs will benefit from very different dithering strategies, the allowed primary and secondary patterns will be designed with the flexibility to let users optimize their observations.

In an effort to preserve filter-mechanism lifetime, it will be general practice to take all the dithers through one filter before switching filters. It remains to be determined whether this includes dither offsets that span multiple visits. The answer to this will likely depend on how long it takes to change filters, compared to the time it takes to switch guide stars, and the extent to which frequent filter changes may impact the filter-mechanism lifetime.

4.0 PRIMARY DITHERS

Primary dithers are offsets of more than a few pixels and are designed to accomplish two goals: (1) averaging over any low-frequency errors in the flat fields and (2) ensuring that all regions of the field of interest are covered by a sufficient number of exposures.

Figure 1 shows NIRCam's unusually shaped field. It has two modules, A and B, with a 40″ gap between them. Each module has a dichroic that allows the same field to be imaged simultaneously by a short-wave channel (SWC) and a long-wave channel (LWC). The SWC is made up of a 2×2 array of 2048×2048 SCAs (sensor-chip assemblies) with a pixel scale of 32 mas/pixel. The LWC covers the same region with a single 2048×2048 chip, with a pixel scale of 65 mas/pixel. The only difference in field coverage between the SWC and LWC comes from the fact that the SWC is made up of four chips that do not abut perfectly, and thus there are gaps between the chips where no photons will be detected. These small gaps are about 5″ in size. The LWC modules have no internal gaps, but the LWC will still have the 40″ gap between modules.

The footprint in Figure 1 represents the field of view covered by a single SWC exposure using both modules. The gray areas represent the regions that are not covered by detectors. The LWC is designed to cover the footprint of the 4 SWC SCAs to within 5% (the aim is to do it within 1%). When we work with large mosaics and tiling, we will consider the NIRCam footprint to extend to the boundaries of the gray area (extended by 5″ in the vertical dimension and 40″ in the horizontal), so that we will be able to work with a symmetric region about each module.

The pattern in Figure 1 has two-fold vertical and horizontal symmetry, and can be reduced to a 2 × 2 pattern of the field shown in Figure 2 below. Considering this

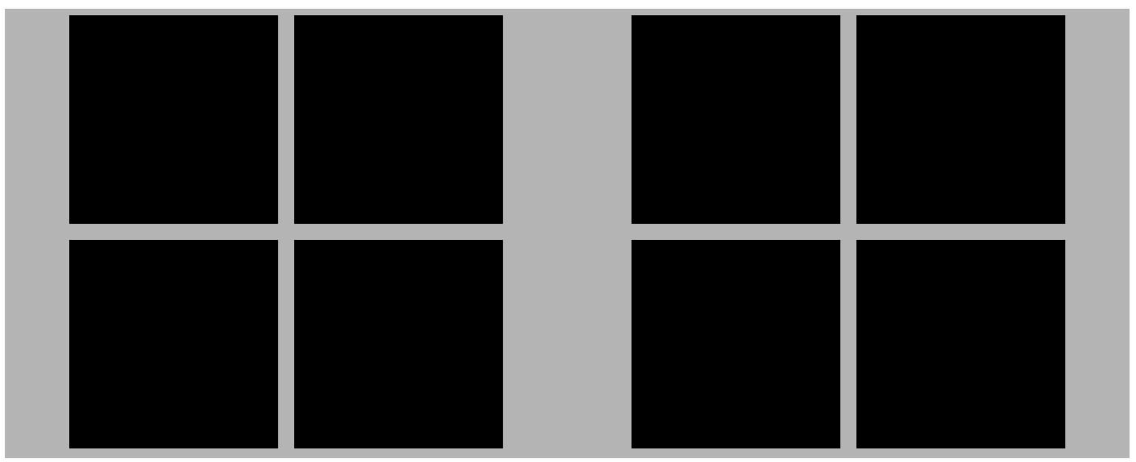

Figure 1: The NIRCam field of view for the SWC. The separation between SCAs is 5″, and the separation between modules is 40″. Each SCA in the each LWC module will cover the 2 × 2 SWC SCAs, so the LWC will not have the small gaps shown here. The full extent of this field is 350″ wide by 140″ tall.

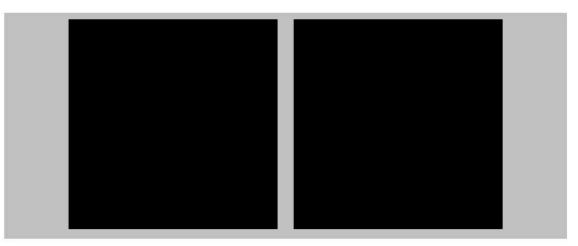

Figure 2: The irreducible pattern in Figure 1. The horizontal extent is 175″ and the vertical extent is 70″. The gap between the SCAs is 5″, and the horizontal padding at the left and right ends is 20″ on each side, and the vertical padding is 2.5″.

symmetry, we can dither NIRCam observations by focusing on patterns that cover these smaller units evenly, knowing that the larger pattern will also be evenly covered, by symmetry. For example, Figure 3 below shows how the NIRCam FOV can be tiled with the repeatable blocks of the pattern in Figure 1.

The main goal of primary dithers will be to cover gaps in the field. There are two kinds of gaps we must concern ourselves with: the 5″ vertical and horizontal gaps between the short-wave SCAs and the 40″ vertical gap between the modules.

In this section, we will construct three kinds of primary dithers. An observer must choose one of these for each set of observations. The first primary-dither strategy will focus on covering as much of the entire NIRCam field as evenly as possible, taking into account all the gaps. The second strategy will be for the case when there is no need to cover the gap between modules. Finally, the third strategy will be focused on small-sized sources that we want to move around a single SCA, but where we do not care about the gaps.

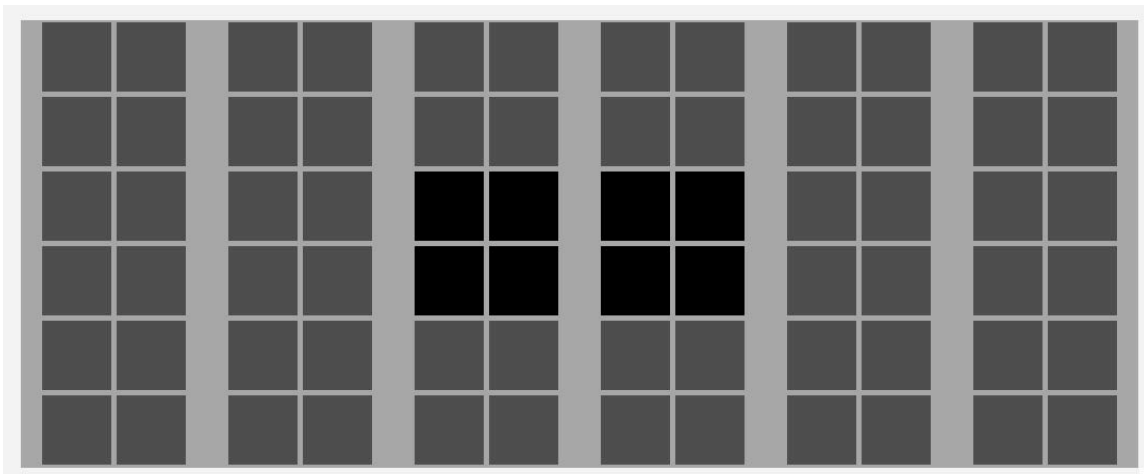

Figure 3: How the NIRCam FOV tiles larger regions. The entire central NIRCam FOV is highlighted. It is made up of a 2×2 array of the "irreducible" unit, which is clearly the basis for the tiling. So this tiling has a 3×3 array of full 350″×140″ FOVs, and a 6×6 array of the smaller 175″×70″ unit.

4.1 Primary Strategy I: Full-Field Dithers

These large primary dithers will be designed to cover the entire NIRCam field as evenly as possible. This coverage will be necessary for targets that are larger than a single NIRCam module (135″×135″). If the target of interest is larger than a single NIRCam FOV (350″×150″), then we will tile several dithered pointings together to achieve even coverage over a larger region. Much smaller targets may require only one module or only a single SCA (see the following sub-sections).

4.1.1 The complicated shape

The complicated shape of the NIRCam field can be seen in Figure 1. In this section, we will define the boundary of "NIRCam field" to be the grey region that surrounds the eight

SWC detectors; this region is 350″ wide by 150″ tall. As mentioned above, by including this grey border around the outside, we give the region horizontal and vertical symmetry, which makes it easier to tile and cover large fields more evenly.

While we consider the entire FOV here, it is still worth bearing in mind that the field can be broken down into a 2 × 2 array of a smaller template, shown in Figure 2. We really need to dither evenly only over this small region, and the rest of the pattern will take care of itself by symmetry. This has the advantage that the size of the needed dithers will be less than half the extent of the entire NIRCam FOV, which will save in slew time and minimize the need to change guide stars.

The main goal of these full-field dithers will be to image large regions of the sky as uniformly as possible. Given this general goal, having regions of the field with inadequate coverage will be bad, and having regions of the field with more coverage than average will not be particularly beneficial, thus it is clear that the evenest possible coverage is preferable. An added benefit of large dithers is that a given point in the field will land on very different locations in the detector, helping us to average out any possible L-flat-type errors.

Finally, it is worth keeping in mind at the outset that since that there are both horizontal and vertical gaps in the field coverage, there is no way to cover the entire field with only two dithers. A minimum of three dithers will be required to cover each point in the large contiguous central region at least once. In the following sub-sections, we will construct primary patterns for 3, 6, 9, 15, 21, 27, 36, and 45 dither positions.

4.1.2 "Dithers" versus "mosaics"

In the standard scheduling terminology, "dithers" represent offsets between exposures that can be achieved within one visit, and "mosaics" represent offsets that require new guide stars, and consequently different visits. The NIRCam field of view is 350″ × 140″ and is much larger than the FGS field (140″ × 140″). Even the "irreducible" NIRCam unit in Figure 2 is 175″ wide, considerably wider than the FGS. This means that it will not in general be possible to use a single guide star when we execute the full-field primary dithers.

The baseline plan is that observations that require switching guide stars will be broken up into different visits, and the offsets are then not formally considered "dithers" but rather part of a "mosaic". In this document, however, we will treat them all as dithers, since this is likely how they will be specified in APT (the Astronomer's Proposal Tool). It is worth noting that since APT is already set up to handle tiling over large regions of the sky, it is already expected to package the many activities associated with a single observation into different visits. It remains to be determined whether the detailed guide-star assignments will be done during the APT stage or afterwards.

An additional issue about switching guide stars is that it will take time; and some estimation of this time will have to be figured into the observational overhead "tax" in some statistical way. This determination will involve some assessment of how long it might take to acquire a guide star after a ~1′ dither where a field-identification step may not be necessary, or at the very least, may be expected to be quicker than a guide-star acquisition after a long slew. These issues will figure into the "overhead" tax assessed to various observations, but they are beyond the scope of this document.

4.1.3 Two 3-point full-field patterns

It was mentioned above that since there are both vertical and horizontal gaps in the NIRCam footprint, we need a minimum of three dithers to ensure that each point in the field will be covered in at least one pointing.

It is easier to think about the dithering by separating the horizontal and vertical considerations. Dithering over the horizontal SCA gap is easy: we simply need to ensure that each dither in *y* places the horizontal gap at a different place in the field. The complications arise from the vertical gaps. The vertical gap between the modules is 40″ wide. The SWC SCAs are 65″ wide, so we cannot dither over the module gap more than once before the gap between the SCAs also comes into play.

After some experimentation, we found that we could get even coverage in the horizontal direction by stepping by 58″, one third of a 175″ module. Each of these horizontal dithers is accompanied by a vertical dither to ensure that no part of the field falls in the horizontal SCA gap in more than one pointing. There are two ways to do this. It can be done compactly, dithering vertically just enough so that the horizontal gap is imaged over. Alternatively, the vertical dithers could be spaced out evenly to ensure that each part of the field is covered by a very different part of the detector.

Figure 4 below shows these two vertical schemes, which are designated as `3-point tight` and `3-point tile`. In the case where the observation is not part of a larger tiling, some observers will prefer the `tight` vertical dither, to ensure the maximum contiguous area covered by the largest number of exposures. In the case where the observation is part of a larger tiling pattern, it will be preferable to use the `tile` vertical spacing, since this allows even coverage when large regions are tiled in a grid. The prescribed horizontal spacing is the same for both schemes.

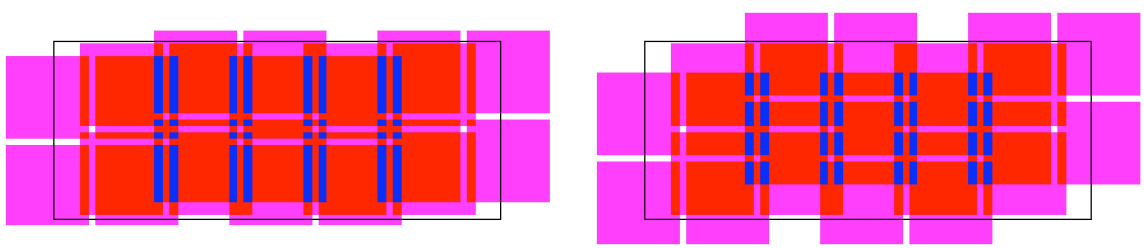

Figure 4: (Left) The 3-point tight pattern; (right) the 3-point tile pattern. The magenta regions are covered by one pointing, the red regions by two, and the blue regions by three. The black outline represents the fiducial 350″×140″ NIRCam field.

The `3-point tight` pattern, shown on the left, features good 2-exposure coverage of most of the central 285″ × 114″ region. The `3-point tile` pattern, on the other hand, provides a good, contiguous amount of double coverage only when it is part of a larger tiling. When visits are tiled together — as shown in Figure 5 below — about 71% of the field is covered in two pointings; 18% gets three-fold coverage, and only 11% is covered in one pointing.

A table containing all the dithers for these patterns in arcseconds will be provided at the end of the section. The horizontal dithers go from −58″ to +58″, so the total extent is 116″. The FGS is about 140″ on a side, so while it will sometimes be possible to find a guide star that can go from one edge of the FGS field to the other, it is improbable that this set of dithers will be doable with a single guide star. It will therefore likely require separate visits.

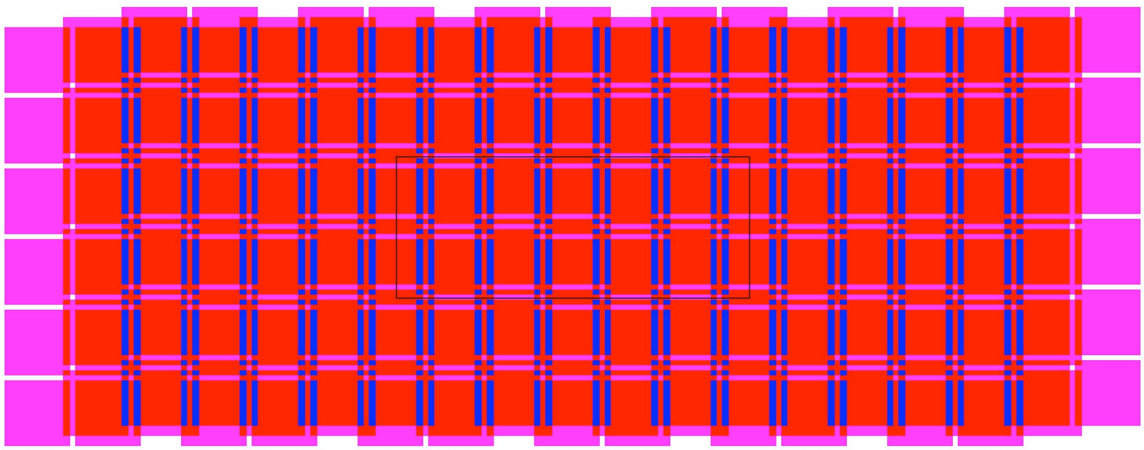

Figure 5: The result of a 3×3 tiling of the `3-point tile` pattern shown on the right of Figure 4, with the same color-coding for coverage. The coverage is as even as possible over a large region. The black lines outline the fiducial NIRCam footprint.

It is worth noting that all the dithers prescribed in this document are in the frame of the detector. Before we can implement the dithers in this frame, we will need to know the transformation from the detector-based frame into the V2-V3-based frame (the frame to which the FGS guide-star movements are tied). We can probably approximate the V2-V3 plane based on telescope models, and we will develop a better model of the focal-plane metrology during the ground tests at Goddard and JSC. But only after we open up on the sky and do a full astrometric calibration will we be able to translate offsets in the detector frame into accurate offsets in the V2-V3 frame, and finally into specific commands for guide-star placement. Fortunately, this refinement will simply require modification of numbers in tables; the general structure of the dithering scripts will be unchanged.

4.1.4 A 6-point full-field pattern

The 3-point pattern above ensures that each point in the field will be covered by at least one exposure. There is no guarantee that this point on the detector will be free of blemishes and artifacts, so when possible, it will be good to take additional observations to improve the coverage. For this, we have designed the 6-point and 9-point dither patterns.

The 6-point pattern was designed with the same ideas as the 3-point pattern. Again, the vertical and horizontal dithers can be treated independently. For the vertical spacing, there is not much benefit in doing separate “tight” and a “tile” dithers; a single vertical set should suffice. The horizontal spacing is more complicated, but can be thought of as a superposition of two of the 3-point patterns, with an offset between them so that the regions that received 3-fold coverage (the darkest regions in Figure 4) will not land on top of each other: it is clearly better to have more 4-, or 5-point coverage than to get some 6-point coverage.

Figure 6 shows the coverage of the 6-point pattern. The legend shows that within the tiled region, the minimum coverage we get is 3 exposures and the maximum is 5. Our aim was to get a distribution of coverage that was as uniform as possible. Within the tiled region we get: 16% of the field covered by 3 exposures, 55% by 4 exposures, and 29% by 5 exposures.

The exact dithers for the `6-point tile` pattern will be given in §4.4, at the end of the primary-dither section. The leftmost horizontal dither is −72″ and the rightmost dither is +73″, for a total span of 145″. This is larger than the 140″extent of a single FGS footprint, so it is clear that we will have to execute this with two different guide-star acquisitions, and hence two visits.

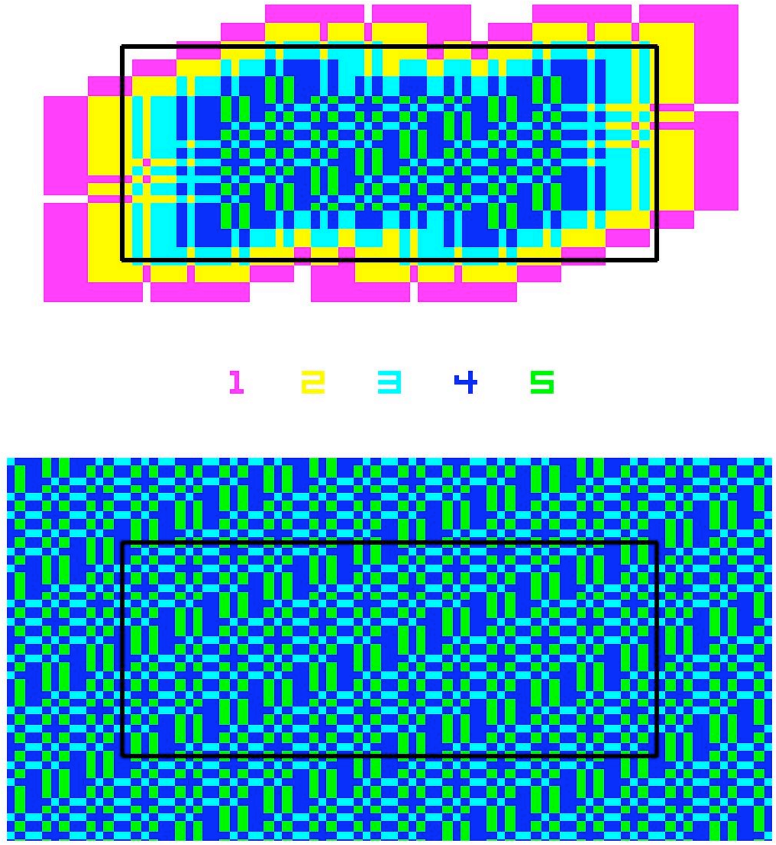

Figure 6: (Top) A single region covered with the 6-point pattern. (Bottom) The 6-point pattern as part of a large-field tiling. The legend in the middle shows that within the tiled region, we never get fewer than three exposures at any given point in the field.

Note that since the horizontal and vertical dithers are independent, we can order them in an optimal way to minimize total slew time. Here, we order them both in ascending *x* and *y* in an effort to minimize the number of guide stars needed. For example, if we had gone from −72 to +73 and back again to −43, we would need three different guide star acquisitions. Since the extent of the vertical dithers is 60″, less than half the size of the FGS FOV, it might be possible to order the vertical components of the dither differently to achieve an outer coverage that does not have the upward slant that is seen in the upper panel of Figure 6. But the ability to do this will depend on having a guide star that is centrally located within the FGS FOV. This will not happen all the time, so to be safe, we will order the *x* and *y* dithers monotonically.

4.1.5 A 9-point full-field pattern

The next full-field pattern we will consider here is the 9-point pattern. This is the largest pattern that allows us to comfortably dither the vertical gaps without overlapping them. The SCA gap is 5″ tall and the SCA height is 65″, so at most we can fit in 12 dithers (11 plus the initial dither) before we end up covering the same part in the field twice with a horizontal SCA gap. It is not wise, though to space the SCA gaps so tightly, since regions around the detector edges are the most likely to have bad pixels. So, 9 is the largest number of vertical dithers we can comfortably fit, giving us a ±1″ margin above and below each gap.

As before with the 6-point dither, we duplicate the 3-point dithers, spacing them out horizontally and vertically in order to achieve coverage that is as even as possible. Figure 7 shows the coverage achieved. It is clear that in the fully covered region, each point in the field is covered by between 5 and 7 pointings. The breakdown is: 13% of the field in 5 pointings, 52% in 6 pointings, and 35% in 7 pointings — very even coverage.

The horizontal dithers here range from −78″ to +78″. As with the 6-point dither, this range can only be achieved with two different guide-star acquisitions. Once again, we order the dithers in increasing *x* and *y* so that the slew time and the number of guide-star changes is a minimum.

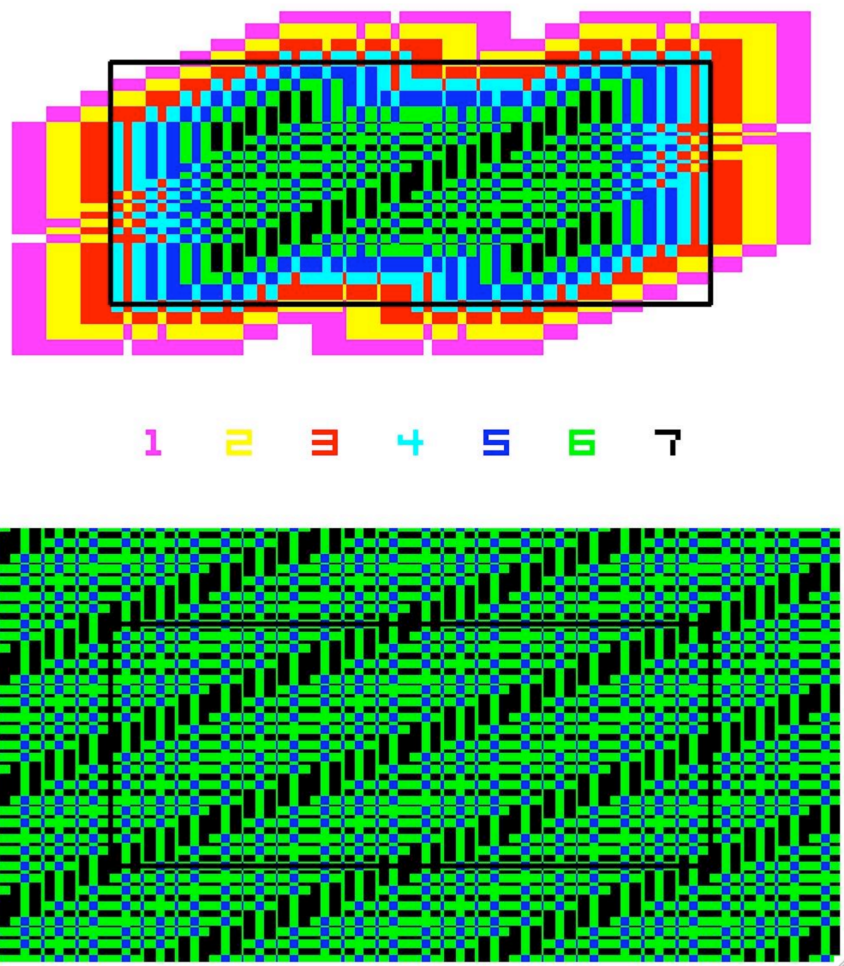

Figure 7: The coverage achieved by the 9-point mosaic dither. The top shows a single set of dithers, and the bottom shows the effect when tiled. All parts of the field get covered by between 5 and 7 exposures.

4.1.6 Additional full-field dithers

The 3-, 6-, and 9-point patterns above provide good coverage over a central region and can be paired (as seen below) with secondary dithers to dial in any number of total exposures. Unfortunately, taking secondary dithers will not improve the field coverage. Even with the 9-point primary dither, we will have 14% of the field in only 5 exposures, meaning that a large fraction of the field will get covered in only 56% of the exposures, 13% less than the average coverage depth of 68% (the RMS variation in the coverage depth is about 10%).

Large programs will involve many more than 9 exposures, so it makes sense to provide dither patterns that allow the evenest possible coverage. To this end, we also add dither patterns for 15-point, 21-point, 27-point, 36-point, and 45-point primary patterns. These patterns were constructed by combining the 3-, 6-, and 9- patterns at various offsets in order to optimize the flatness of the coverage: maximize the minimum number of exposures covering any given point, and minimize the maximum number of exposures covering any given point. The new patterns are shown in Figure 8 below, which also reports the fraction of the field covered by various exposure depths. The 45-point pattern has a variation of only 3.8% in depth of coverage across the field.

The NIRCam Users' Manual will contain recommendations for how to select dither patterns and APT will facilitate interaction with tiling programs, so we will not focus on these issues in this document. However, it is worth pointing out that all these patterns fit nicely together in a large-area tiling mosaic. Figure 9 reports how much field is covered by what how many exposures for each of the full-field dither patterns for 1×1, 2×2, and 3×3 tiling arrays.

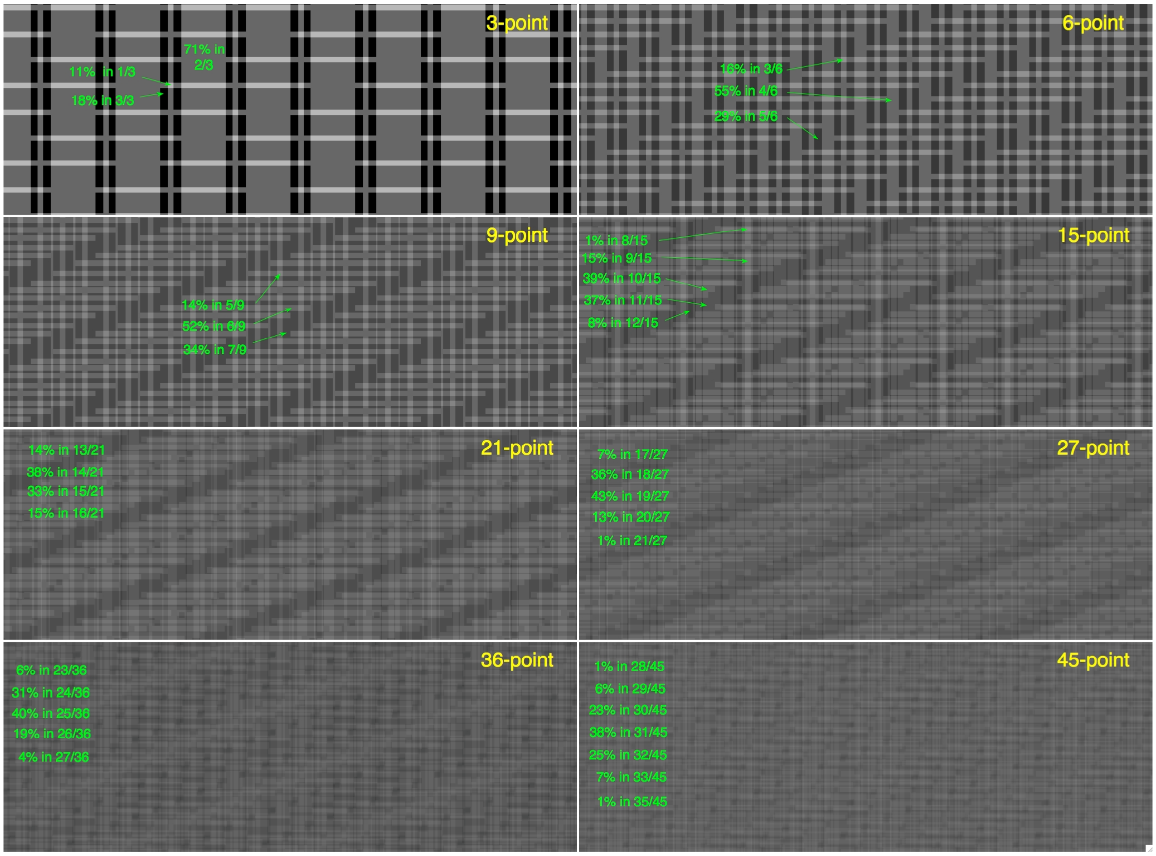

Figure 8: Depth of coverage for the additional dither patterns (as marked in yellow). The depth of coverage is given in green. The average coverage is 68.97%, and the patterns approach this with a smaller and smaller variation about this mean.

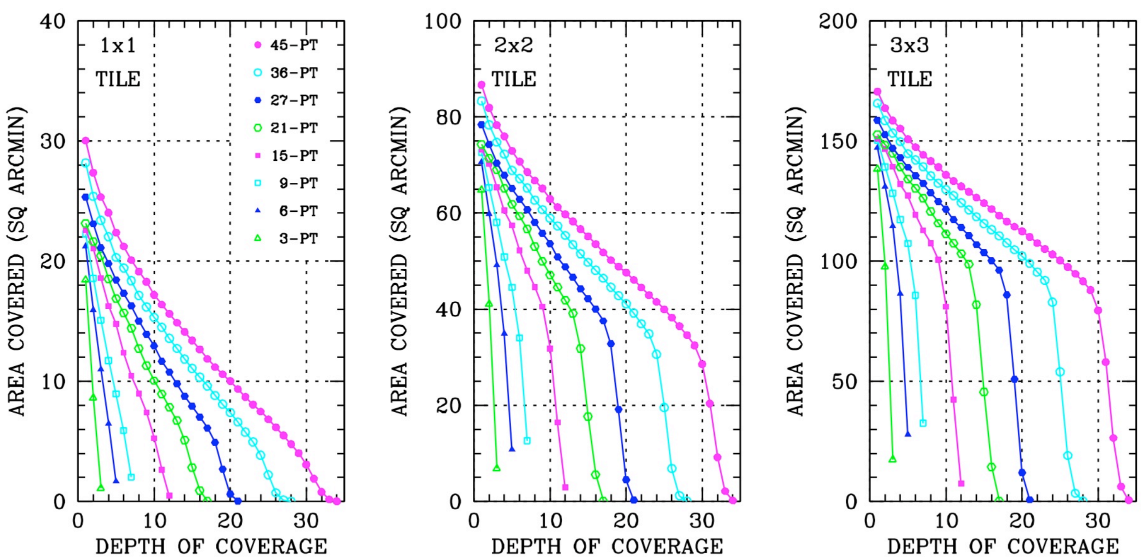

Figure 9: The area covered to various depths for 1×1, 2×2, and 3×3 mosaic-tilings of the various full-field primary dithers.

4.1.7 Mitigating chip-edge effects

It is well known that the H2RG detectors frequently have more bad pixels at the edges and corners, either because of detached Indium bonds or other manufacturing issues. We have spaced out the dithers in our patterns as much as possible so that not only will a patch of the sky fall in a gap a minimal number of times, but it will also fall near the detector edges a minimal number of times. For example, in the 9-dither case, with the nominal SCA size we achieved 13%, 52%, and 35% coverage with 5, 6, and 7 exposures, respectively. If we shrink the SCAs by 1″ (31 pixels) all around, then these numbers change to: 24%, 55% and 21% — still very good, and each point is still covered a minimum of 5 times. Only when we shrink the SCAs by another 1″ do we get a small region (less than 2%) that has coverage in only four pointings. The 6-point and 3-point patterns are robust against even more shrinkage, meaning that we will be more robust against larger bad areas near the detector edges. The 15+-point dithers are too tightly packed to allow any edge margin, though the variation of the dithering should ensure that a point in the field impacted by the edge in one exposure will likely not be impacted by that edge in other exposures.

4.2 Primary Strategy II: Dithers to cover the SCA gap

The above "full-field" primary dither patterns were generated in an effort to cover a large region on the sky as evenly as possible. Covering a large, contiguous field will not be critical for every project. This might be the case if the program is focused on the 136″×136″ field covered by a single module. Or, it might be the case if a program does need to cover a large area, but can tolerate a gap in the middle (such as doing a statistical survey of galaxies in a large cluster where not every single galaxy must be counted).

The gaps between the SCAs are about 5″ in width and height, so we must dither by at least this amount to ensure that the parts of the field that land in the gap in one pointing will be covered in another. To be safe, we will dither by 7″ to mitigate the prevalence of bad pixels near the detector edges.

The dither pattern we construct will be a compromise between the quantity and the quality of coverage. Whenever we dither two fields, we end up increasing the total area covered by any one pointing, while at the same time decreasing the total area covered by both pointings. The larger the dither is, the smaller will be the size of the region with multiple coverage. Furthermore, in this case, if the extent of the dithers is larger than the 40″ module gap, then we may as well just go with the full-field dithers discussed in the previous section. Therefore, we should clearly keep the extent well below this. This means that we need a dither pattern that is as compact as possible, but which still has decent coverage for the stars impacted by the SCA gaps.

As we saw in the previous section, since the detector has both vertical and horizontal gaps, three separate dithers will be required if we want cover the entire field in at least one exposure. If we have four separate dithers, then we can ensure that each part of the inner field will be covered by a minimum of two exposures. In general, $N_{min} = N_{dith} - 2$, so long as the all the dithers are separated in x and y by more than the gap size. Figure

10 shows how much of the field is covered by different numbers of pointings for between one and four dithers. If we have five dithers, we could in principle get a minimum of three exposures everywhere, but we start running into other problems: the size of the high-coverage field starts to shrink. Furthermore, the dithers begin to reach the size of the module gap, defeating the purpose of the single-module focus.

So, for all these reasons we will focus here on ensuring good coverage within the inner 114″ × 114″ of the field — the region that bounds the magenta parts of the left panel of Figure 10. This corresponds to more than 70% of the area spanned by a single module (135″ × 135″). The total extent where the dither pattern has at least single-exposure coverage here is 156″ × 156″, 34% larger than the footprint of a single module.

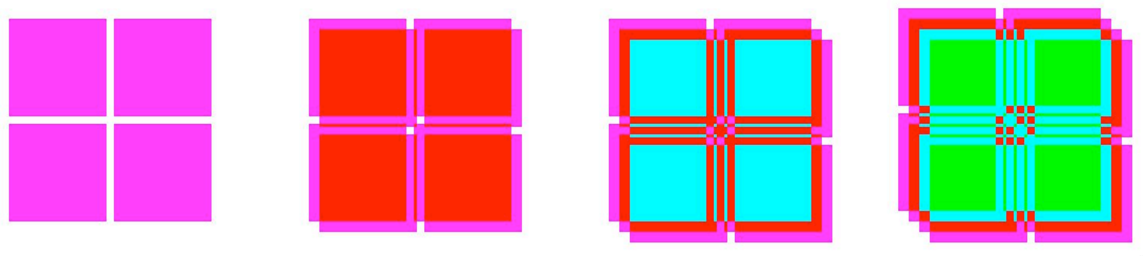

Figure 10: The depth of single-module coverage for one, two, three, and four dithers. The colors magenta, red, cyan, and green show how which parts of the field are covered by 1, 2, 3, and 4 pointings, respectively.

We have constructed a 16-point dither pattern that covers the targeted region evenly and densely. This pattern has been designed such that it can be executed with any number of points from 3 to 16, and good even coverage will still be achieved across the field. A table with the actual dither offsets in arcseconds will be given at the end of this section.

Figure 11 shows the coverage achieved for the dither pattern for $N_{dith} = 1$ through $N_{dith} = 16$. Note that it is only when $N_{dith} \geq 3$ that we cover the entire inner region with at least one pointing (as can also be seen in Figure 10). For $N_{dith} = 4$, we cover all points in the field in at least 2 exposures. For $N_{dith} = 8$, we get all points in at least 5 exposures. For 12 dither positions, we get each point at least 6, and typically more than 9 pointings. Finally, for 16 dither positions, we get each point in a minimum of 9, and typically get more than 12. Table 1 shows the percentage of the field with different levels of coverage as a function of the number of dither positions.

The dithers presented in this section will involve slews of at most 21″, considerably smaller than the dithers in the previous section. These dithers should have no problem using the same guide star, so they can all be scheduled as a single visit. If the observing plan requires only one module, then the target will likely be placed at the center of the module; the coverage there will not be quite as deep as in the semi-cardinal directions, but it will be symmetric about that point. If the plan involves using both modules, but sacrificing the region in the module gap, the coverage will be as shown in Figure 12. An object at the center of a the NIRCam footprint will get no coverage.

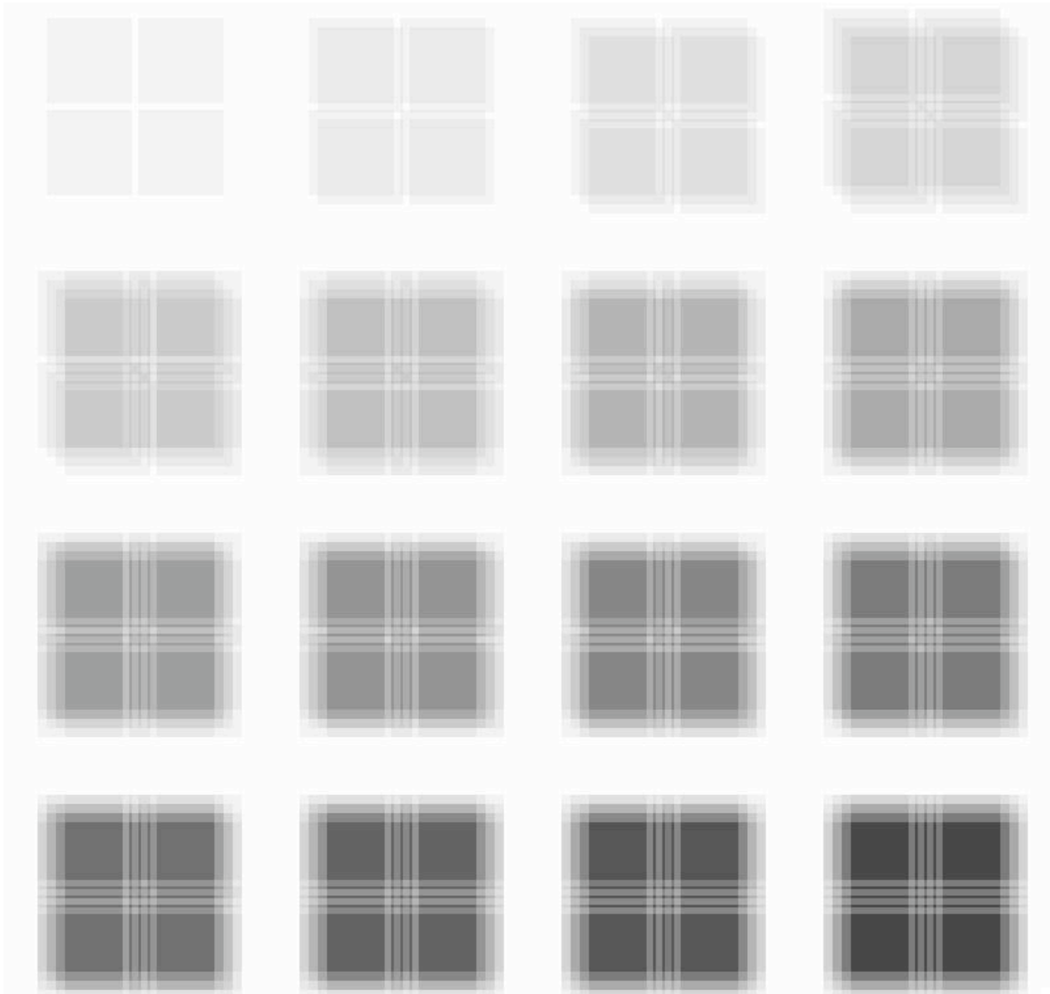

Figure 11: This shows the intra-module dither pattern for N=1 to N=16. The aim was to achieve the most uniform possible coverage over the largest possible region.

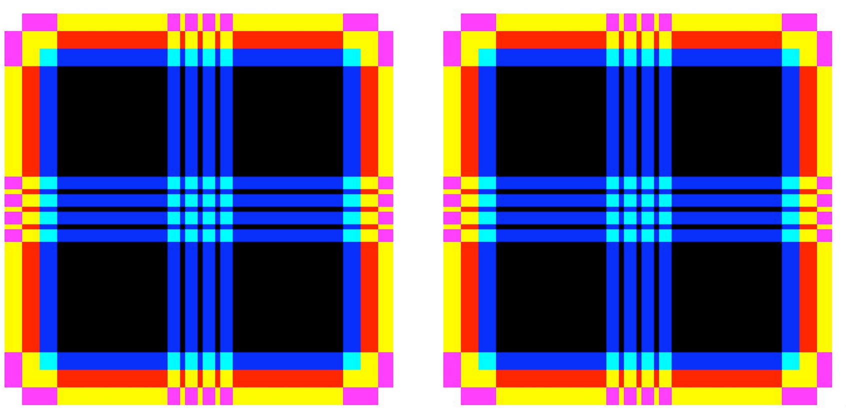

Figure 12: The intra-module pattern when executed for two modules. The inner square of diminished coverage is about 20″ on a side. The magenta, yellow, red, cyan, blue and black regions have, respectively, coverage of 2, 4, 8, 9, 12, and 16 exposures.

Table 1: For various numbers of points in the intra-module dither, we report the percent of the 114″×114″ field that gets the listed depth of coverage.

NUMBER OF DITHERS	PERCENT OF FIELD WITH THE LISTED DEPTH OF COVERAGE																
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	9	91															
2	1	16	83														
3	0	1	25	74													
4	0	0	3	30	67												
5	0	0	1	9	23	67											
6	0	0	0	3	16	14	67										
7	0	0	0	1	2	24	7	67									
8	0	0	0	0	2	2	30	0	67								
9	0	0	0	0	1	2	8	23	0	67							
10	0	0	0	0	0	1	3	14	15	0	67						
11	0	0	0	0	0	0	2	1	23	7	0	67					
12	0	0	0	0	0	0	1	3	0	30	0	0	67				
13	0	0	0	0	0	0	0	2	1	7	23	0	0	67			
14	0	0	0	0	0	0	0	1	2	1	14	16	0	0	67		
15	0	0	0	0	0	0	0	0	2	1	0	23	7	0	0	67	
16	0	0	0	0	0	0	0	0	0	3	0	0	30	0	0	0	67

4.3 Primary Strategy III: Dithers within an SCA

The above "full-field" and "intra-module" dither patterns have been designed to mitigate the effects of the gaps in the detector coverage. If the target is confined to a single SCA, we will not need to consider gaps, yet we may still have good reasons to dither within the SCA. For instance, if the target is smaller than an SCA (65″ for the short-wave channel, 135″ for the long-wave channel), then we can move it around the SCA without losing coverage. This can help to mitigate any possible L-flat errors. Such observations of a well-chosen field can also allow us to *solve* for the distortion solution or the corrections to the L-flats.

Our aim with the intra-SCA pattern will be to move a target around the chip as much as possible without moving it off the chip. There are two parameters to consider: the size of the target and the number of dithers desired. We will characterize the size of the target in terms of being "small", "medium", or "large", depending on whether its diameter is about 25% the size of the detector (~ 16″), 50% the size of the detector (~ 33″), or 75% the size of the detector (~ 50″). This will tell us how close the center of our object is allowed to get to the edge of the detector. Figure 13 shows graphically how the extent of the pattern depends on the size of the target.

As for the number of dithers, the user will select the number of dithers, and we will space them out to fill the available area on the detector. We will accommodate up to 25 dithers, in a 5 × 5 pattern. As with the intra-module pattern in §4.2, the intra-SCA pattern will be designed to execute a good dither for any number between 1 and 25 by going sequentially through the list and stopping after the requisite number is completed. The full 25-point dither pattern is shown on the left of Figure 14 in terms of execution order. The patterns executed for selected dithers of 1, 5, 9, 13, 17, and 25 are shown on the right.

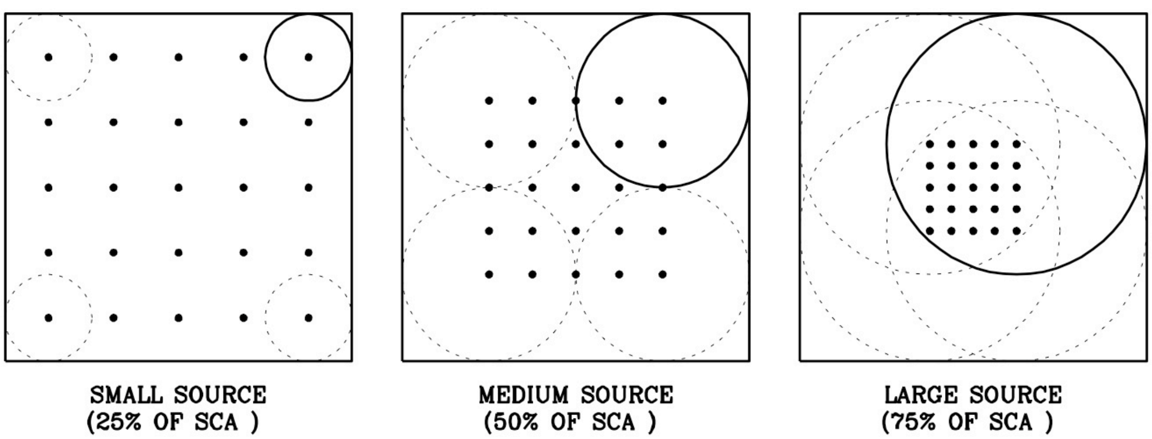

Figure 13: The panels from left to right show the full 25-point dither pattern achieved for small, medium, and large sources, respectively. The heavy circle shows the placement of the target at the upper right extreme dither point. The dotted circles show the target at the other three extreme dither positions.

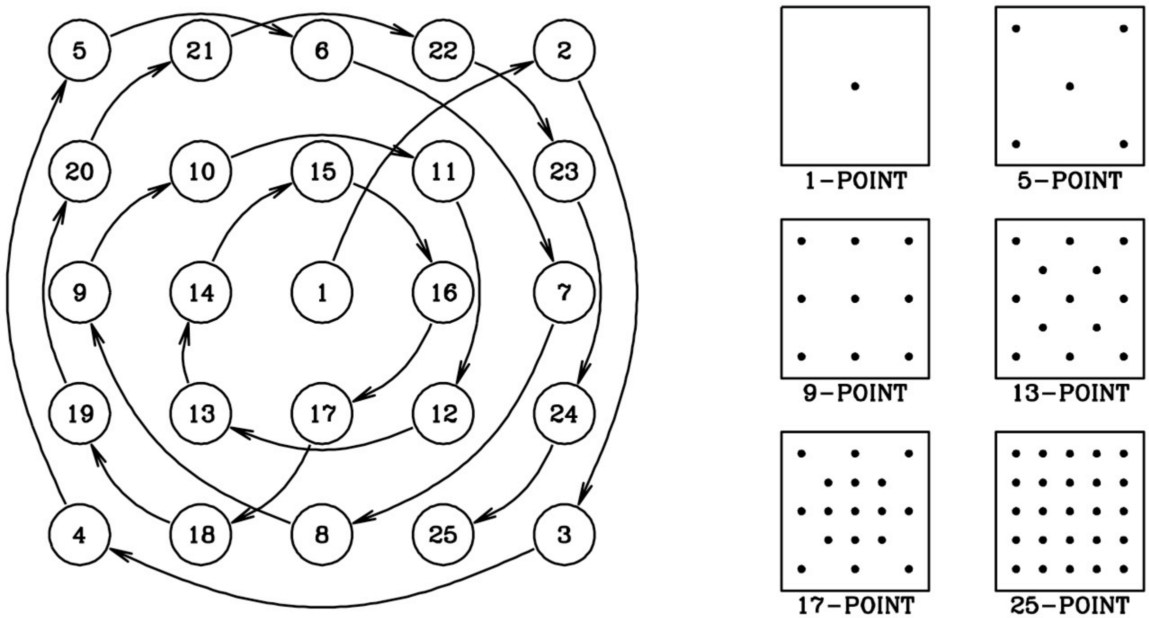

Figure 14: (Left) the order of execution of the 5 × 5 array of dither points. (Right) The dither pattern achieved for 1, 5, 9, 13, 17, and 25 dithers. Any number between 1 and 25 can be executed.

This dither pattern will be centered with the target at the center of the chosen SCA and will come in three sizes, for the three sizes of target. Smaller targets can naturally afford a more spread-out pattern. Figure 13 shows the patterns for the three sizes of target. This pattern can of course be executed for any aperture (the location on the detector where the target is placed with no dither offset), but we anticipate that these patterns would only be executed with apertures centered on SCAs.

The patterns discussed above are given in a table at the end of this section. The numbers given apply to a single SCA in the short-wave channel. It may also be useful to do the same in the long-wave channel. To get the LWC dithers, we simply need to multiply each element in the SWC pattern by two. Since the LWC object sizes will be relative to that detector, the three corresponding target sizes will be 33″, 66″, and 100″.

4.4 Summary of primary patterns

We have developed three major types of primary dither patterns, which are summarized in Table 2 below. The first pattern was designed to deal with the 40″ gap between the modules, so that we can tile large regions with coverage of the field that is as even as possible. The second pattern covers the 5″ gaps between SCAs but not the gap between modules. The third pattern focuses on a single SCA and moves a target of a given size around the chip as much as possible.

These patterns are designed to be used in conjunction with secondary dithers, which will be executed at each of the primary dither locations, but the primary patterns are not designed to be used in conjunction with each other. Tables 3 through 5 provide the actual dithers in arcseconds for each of the patterns. The intra-module and intra-SCA dithers will all start with dither #1, and can execute as many dithers as desired (up to 16 and 25, respectively). The full-field dithers are designed to be executed with all the points in the pattern.

Table 2: Summary of the various primary dither patterns.

Type	Primary Pattern	Description
Full-field dither	`3-point tight`	3-point pattern that maximizes depth of 3× coverage
	`3-point tile`	3-point pattern that maximizes inner area of 2× coverage
	`6-point tile`	6-point pattern that maximizes area of 4× coverage (minimum 3×-coverage everywhere)
	`9-point tile`	9-point pattern that maximizes area of 6× coverage (minimum 5×-coverage everywhere)
	`15-,21-,27-,36-, 45-pt tile`	Specific patterns to take optimally flat coverage
Intra-module dither	`N-point intra-module`	N-point dither to mitigate the effects of the SCA gaps
Intra-SCA dither	`N-point intra-SW-SCA (small)`	N-point dither within a SW SCA, small target (< 16″ diameter)
	`N-point intra-SW-SCA (medium)`	N-point dither within a SW SCA, medium target (< 33″ diameter)
	`N-point intra-SW-SCA (large)`	N-point dither within a SW SCA, large target (< 50″ diameter)
	`N-point intra-LW-SCA (small)`	N-point dither within a LW SCA, small target (< 32″ diameter)
	`N-point intra-LW-SCA (medium)`	N-point dither within a LW SCA, medium target (< 66″ diameter)
	`N-point intra-LW-SCA (large)`	N-point dither within a LW SCA, large target (<100″ diameter)

Table 3: The offsets for the full-field primary dither patterns.

3-point tile full-field		
1	−58″	−23.5″
2	0″	0″
3	+58″	+23.5″

3-point tight full-field		
1	−58″	−7.5″
2	0″	0″
3	+58″	−7.5″

6-point tile full-field		
1	−72″	−30″
2	−43″	−18″
3	−14″	−6″
4	+15″	+6″
5	+44″	+18″
6	+73″	+30″

9-point tile full-field		
1	−78″	−33″
2	−58″	−24″
3	−38″	−15″
4	−20″	−8″
5	0″	0″
6	+20″	+8″
7	+38″	+15″
8	+58″	+24″
9	+78″	+33″

15-point tile full-field								
1	−77″	−32″	6	−19″	−7″	11	39″	16″
2	−72″	−33″	7	−14″	−9″	12	44″	15″
3	−57″	−23″	8	1″	0″	13	59″	25″
4	−43″	−20″	9	15″	3″	14	73″	27″
5	−37″	−14″	10	21″	9″	15	79″	32″

21-point tile full-field								
1	−77″	−36″	8	−19″	−5″	15	39″	18″
2	−75″	−30″	9	−17″	−7″	16	41″	17″
3	−70″	−21″	10	−12″	2″	17	46″	27″
4	−55″	−23″	11	3″	0″	18	61″	27″
5	−48″	−18″	12	10″	5″	19	68″	34″
6	−41″	−12″	13	17″	11″	20	75″	29″
7	−35″	−12″	14	23″	12″	21	81″	34″

27-point tile full-field								
1	−84″	−41″	**10**	−26″	−16″	**19**	32″	7″
2	−77″	−37″	**11**	−19″	−12″	**20**	39″	11″
3	−71″	−20″	**12**	−13″	5″	**21**	45″	28″
4	−64″	−32″	**13**	−6″	−9″	**22**	52″	16″
5	−57″	−28″	**14**	1″	−5″	**23**	59″	20″
6	−51″	−11″	**15**	7″	12″	**24**	65″	37″
7	−44″	−23″	**16**	14″	0″	**25**	72″	23″
8	−37″	−19″	**17**	21″	4″	**26**	79″	27″
9	−31″	−2″	**18**	27″	21″	**27**	85″	44″

36-point tile full-field								
1	−83″	−54″	**13**	−25″	−29″	**25**	33″	-6″
2	−78″	−27″	**14**	−20″	−2″	**26**	38″	21″
3	−77″	−37″	**15**	−19″	−12″	**27**	39″	11″
4	−72″	−10″	**16**	−14″	15″	**28**	44″	38″
5	−63″	−45″	**17**	−5″	−22″	**29**	53″	3″
6	−59″	−18″	**18**	0″	5″	**30**	58″	30″
7	−57″	−28″	**19**	1″	−5″	**31**	59″	20″
8	−52″	−1″	**20**	6″	22″	**32**	64″	47″
9	−43″	−36″	**21**	15″	−13″	**33**	73″	10″
10	−38″	−9″	**22**	20″	14″	**34**	78″	37″
11	−38″	−19″	**23**	21″	4″	**35**	79″	27″
12	−32″	8″	**24**	26″	31″	**36**	84″	54″

45-point tile full-field														
1	−84″	−61″	**10**	−50″	6″	**19**	−17″	−5″	**28**	22″	21″	**37**	55″	10″
2	−81″	−47″	**11**	−44″	−43″	**20**	−12″	22″	**29**	23″	11″	**38**	60″	37″
3	−76″	−20″	**12**	−41″	−29″	**21**	-6″	−29″	**30**	28″	38″	**39**	61″	27″
4	−75″	−30″	**13**	−36″	−2″	**22**	−3″	−15″	**31**	32″	−13″	**40**	66″	54″
5	−70″	−3″	**14**	−35″	−12″	**23**	2″	12″	**32**	35″	1″	**41**	72″	3″
6	−64″	−52″	**15**	−30″	15″	**24**	3″	2″	**33**	40″	28″	**42**	75″	17″
7	−61″	−38″	**16**	−26″	−36″	**25**	8″	20″	**34**	41″	18″	**43**	80″	44″
8	−56″	−11″	**17**	−23″	−22″	**26**	14″	−20″	**35**	46″	45″	**44**	81″	34″
9	−55″	−21″	**18**	−18″	5″	**27**	17″	−6″	**36**	52″	−4″	**45**	86″	61″

Table 4: The dither offsets for the intra-module dither.

Dither	x-offset	y-offset
1	−3.75″	+3.75″
2	+3.75″	−3.75″
3	+11.25″	−11.25″
4	−11.25″	+11.25″
5	+11.25″	+11.25″
6	−11.25″	−11.25″
7	−3.75″	−3.75″
8	+3.75″	+3.75″

Dither	x-offset	y-offset
9	−11.25″	+3.75″
10	+11.25″	−3.75″
11	+3.75″	−11.25″
12	−3.75″	+11.25″
13	−11.25″	−3.75″
14	+11.25″	+3.75″
15	−3.75″	−11.25″
16	+3.75″	+11.25″

Table 5: The dither patterns for the three intra-SCA dithers for the SWC. The medium- and large-object dithers are just scaled down from the small-object dither. This pattern can be executed with any number of dithers, from 1 to N, and a good N-point dither will result. The LWC patterns would be scaled up by a factor of two.

Dither number	Small object		Medium object		Large object	
	x-dither	y-dither	x-dither	y-dither	x-dither	y-dither
1	+0.00″	+0.00″	0.00″	+0.00″	+0.00″	+0.00″
2	+24.56″	+24.56″	+16.38″	+16.38″	+8.19″	+8.19″
3	+24.56″	–24.56″	+16.38″	+16.38″	+8.19″	–8.19″
4	–24.56″	–24.56″	–16.38″	–16.38″	–8.19″	–8.19″
5	–24.56″	+24.56″	–16.38″	–16.38″	–8.19″	+8.19″
6	+0.00″	+24.56″	+0.00″	+16.38″	+0.00″	+8.19″
7	+24.56″	0.00″	+16.38″	+0.00″	+8.19″	+0.00″
8	+0.00″	–24.56″	+0.00″	–16.38″	+0.00″	–8.19″
9	–24.56″	+0.00″	–16.38″	+0.00″	–8.19″	+0.00″
10	–12.28″	+12.28″	–8.19″	+8.19″	–4.09″	+4.09″
11	+12.28″	+12.28″	+8.19″	+8.19″	+4.09″	+4.09″
12	+12.28″	–12.28″	+8.19″	–8.19″	+4.09″	–4.09″
13	–12.28″	–12.28″	–8.19″	–8.19″	–4.09″	–4.09″
14	-12.28″	+0.00″	–8.19″	+0.00″	–4.09″	+0.00″
15	+0.00″	+12.28″	+0.00″	+8.19″	+0.00″	+4.09″
16	+12.28″	+0.00″	+8.19″	+0.00″	+4.09″	+0.00″
17	+0.00″	–12.28″	+0.00″	–8.19″	+0.00″	–4.09″
18	–12.28″	–24.56″	–8.19″	–16.38″	–4.09″	–8.19″
19	–24.56″	–12.28″	–16.38″	–8.19″	–8.19″	–4.09″
20	–24.56″	+12.28″	–16.38″	+8.19″	–8.19″	+4.09″
21	–12.28″	+24.56″	–8.19″	+16.38″	–4.09″	+8.19″
22	+12.28″	+24.56″	+8.19″	+16.38″	+4.09″	+8.19″
23	+24.56″	+12.28″	+16.38″	+8.19″	+8.19″	+4.09″
24	24.56	-12.28	16.38	-8.19	8.19	–4.09
25	12.28	-24.56	8.19	-16.38	4.09	–8.19

5.0 SECONDARY DITHERS

The primary dithers we considered in the previous section were designed to accomplish two things: cover the selected field as evenly as possible while at the same time placing the same sources at different locations on the detectors to mitigate the effects of low-frequency flat-fielding errors.

Even for the smallest intra-SCA dithers (4.09″), the dither achieved at the center of the detector is about 128 pixels in the SWC. If the distortion is 2% (the upper limit of the requirement), then this means that the shift achieved at the edge of the detector could be more than 130.5 pixels (128 × 1.02), and in other places on the detector could be as low as 126.5 pixels (128 × 0.98). This means that if we attempt to execute (say) the intra-SCA dithers with carefully tuned sub-pixel offsets to improve the pixel sampling, we may achieve good pixel-phase coverage at the center of the chip, but distortion will cause the pixel-phase effort to be incoherent across most of the detector chip. Furthermore, it is not even clear that the FGS will be able to dither with good sub-pixel accuracy over such large distances (MO-162, MOCD). For these reasons, we cannot hope to achieve good pixel-phase coverage with only large dithers.

The solution is to take a set of small dithers at each large-dither location. These dithers will be small enough to ensure coherent pixel phases across the detector. Better pixel-phase coverage across the field will lead to a more accurate, more complete reconstruction of the astronomical scene.

It is beyond the scope of this document to recommend dithering strategies for specific programs. An early discussion of JWST dithering strategies can be found in Koekemoer and Lindsay (2002), which was written before some of the specifics of the detectors were finalized. A new treatment with more up-to-date PSF models will be presented in a follow-up document that will study how various sampling strategies will enable (or not) typical science use-cases, such as distinguishing between resolved and unresolved sources and fitting of morphological shapes and profiles. A discussion of these issues will also figure prominently in the NIRCam Users' Manual, so that proposers can wisely choose the dithering strategy that best matches the scientific goals of their program. In the present document, the discussion will be confined to a determination of how to achieve the pixel-phase coverage that is necessary for NIRCam.

This section will first set out what is meant by the "astronomical scene", and then will describe how we can reconstruct this scene from multiple dithered observations. We will then touch on some specific issues related to NIRCam, such as how much we expect distortion to impact pixel-phase coverage and how to execute dithers simultaneously in the SWC and the LWC. Finally the last two sections will describe two approaches to secondary dithering, with the actual dithers provided in tables at the end of this section.

5.1 The delivery of the astronomical scene to the detector

There are several processes that go into the generation of an astronomical image. We start with the perfect scene on the sky, S(x,y), which can be thought of as having infinite resolution and dynamic range. This scene then goes through the telescope and camera optics, producing a scene on the detector that is essentially the perfect scene convolved with the instrumental point-spread function (PSF), $\psi(\Delta x, \Delta y)$. This can be seen intuitively in that if the scene is made up of just one star, then the scene has a delta-function at the star's location, and we would expect a simple instrumental PSF to arrive at the detector. If we decompose the scene into a very finely spaced array of delta functions, then it is clear that what the detector sees is the scene convolved with the PSF.

This PSF-convolved scene is then collected in the detector pixels and read out for processing. The collection into pixels can also be thought of as a convolution process, similar to the convolution due to the optics, but this time the convolution function is a top-hat, $\Pi(\Delta x, \Delta y)$, or some other function that describes the pixel-response function.

Since convolution is commutative, the two convolutions can be thought of as a single convolution with a function that is the instrumental PSF convolved with the pixel-response function. This new function is called the "effective" point-spread function (see Anderson & King 2002), because it represents the PSF that we are actually able to sense with the detector:

$$\psi_E(\Delta x, \Delta y) = \psi(\Delta x, \Delta y) \otimes \Pi(\Delta x, \Delta y).$$

When the scene is convolved with this composite PSF, we call it the "effective" scene:

$$S_E(x,y) = S(x,y) \otimes \psi(\Delta x, \Delta y) \otimes \Pi(\Delta x, \Delta y)$$
$$= S(x,y) \otimes \psi_E(\Delta x, \Delta y).$$

This effective scene is a smooth, continuous, two-dimensional function, since both the perfect scene and the "effective" PSF are smooth, continuous, two-dimensional functions.

The scene on the sky of course has infinite resolution, and can have structure on any spatial scale. The instrumental PSF $\psi(\Delta x, \Delta y)$ can have structure on angular scales as small as λ/D, where λ is the wavelength of light and D is the diameter of the telescope. The pixel-response function has structure on very fine scales, since the edges of the pixels can often be thought of as being quite sharp. In practice, the resolution scale of the effective image is set by the broadest scale in the convolution product, so it is set by the instrumental PSF. This means that the effective image cannot vary with infinite resolution, but is limited to being no sharper than the instrumental PSF (λ/D). Thus there is no point in probing structure in the "effective" scene at finer scales than the PSF's spatial filter will let through.

We are usually accustomed to thinking of pixels as being light buckets that collect light over a specified area, but when we work in the "effective" domain described above, the pixel integration has already been included. So instead of thinking of each pixel as integrating over a region of the field, it is now more appropriate to think of each pixel as sampling this effective scene at one point in the field — the location where the pixel is centered. The array of pixels in an image then provide an array of point-samplings of this scene at a grid of locations spaced one pixel apart.

If the effective image has structure on a finer scale than the angular size of one pixel, then it is said that the image "undersamples" the scene: there is structure in the scene that changes faster than the pixels can sample. In NIRCam, the short-wave pixel scale is 32 mas/pixel. The PSF at 2 microns has a FWHM of λ/D = 64 mas, so there are two pixel samplings for every resolution element and the detector is considered to be just critically sampled at this wavelength. At shorter wavelengths, the detector is undersampled, and it becomes badly undersampled at wavelengths of less than 1 micron. The bluest filter NIRCam has is F070W at 0.7 microns, and it is undersampled by a factor of about three.

Fortunately, there are ways to improve the sampling of undersampled detectors. If we take one exposure at one pointing, then shift the detector by half a pixel in the x direction and take another exposure, then we double the sampling frequency of the pixels along that axis, and we can recover structure that was lost on account of the large pixels. This is the strategy we will pursue with the secondary dither patterns.

This document is focused on providing dither patterns that allow us to sample astronomical scenes adequately for various science goals. It is not focused on how to analyze those samples to either reconstruct an image or to do science directly on the set of raw images. It is assumed that if the scene is sampled adequately, software such as `Drizzle` or something else, can be used for tha analysis, but if the scene is not sampled adequately, no amount of post-analysis can recover the necessary information. Thus, our focus here is to ensure that our dithering strategies will sample the scene adequately for high-precision science.

5.2 Some special considerations

The secondary-dither patterns discussed below will be designed with several different considerations in mind. The small-dither offsets will be large enough to place sources in different pixels, to mitigate flat-field errors, but not so large that the pattern loses pixel-phase coherence across the detector. In addition, the offsets will be designed to accomplish a good pixel-phase dither for the short-wave and long-wave channel simultaneously. Finally, the patterns will be designed as flexibly as possible so that users can get exactly the number of dithers they want. These considerations will be discussed below.

5.2.1 Whole-pixel dithers and distortion

The main goal of the secondary dithers is to improve the sub-pixel sampling. But so long as we are taking the time to reset the exposure and execute a small slew, it makes sense to dither more than a fraction of a pixel, so that the source will land on completely different pixels. This will help average out any pixel-to-pixel flat-fielding errors and will also mitigate the influence of bad pixels.

Thus, the plan will be to execute sub-pixel dithers on top of dithers of a few whole pixels. Distortion puts a limit on how large these whole-pixel offsets can be. Distortion in NIRCam could be as large as 2% at the edges of the detector. This means that the plate scale at the edge of the field could differ by 2% from that at the center of the field, so that a dither of 25.5 pixels at the center of a detector would result in a dither of 26.0 pixels at the edge. Thus, a dither pattern designed to improve the sub-pixel sampling would be effective for the central part of the chip, but would become ineffective towards the edges if it spans too many pixels.

Figure 15 provides a simple example of this. In both panels, we have executed a 4-point pattern to improve the sampling by a factor of two in each dimension. We have simulated 2% distortion in the field. The left set of panels shows the achieved sub-pixel dither for the case where the whole-pixel offset is 5 pixels, i.e., offsets of (0.0, 0.0), (0.0, 5.5), (5.5, 0.0), and (5.5, 5.5). A reasonable dither is achieved across the entire detector. The right set of panels shows the dither achieved for whole-pixel offsets of 20 pixels between dithers, i.e., dithers of (0.0, 0.0), (0.0, 20.5) (20.5, 0.0) and (20.5, 20.5). The pattern is good at the detector center (as constructed), but becomes progressively worse as we go towards the edge of the detector, until finally at the corner almost all of the samplings land nearly on top of one another, and the dither does not improve our sampling of the scene. Our aim in constructing sub-pixel dithers will be to keep the pattern smaller than about 10 pixels.

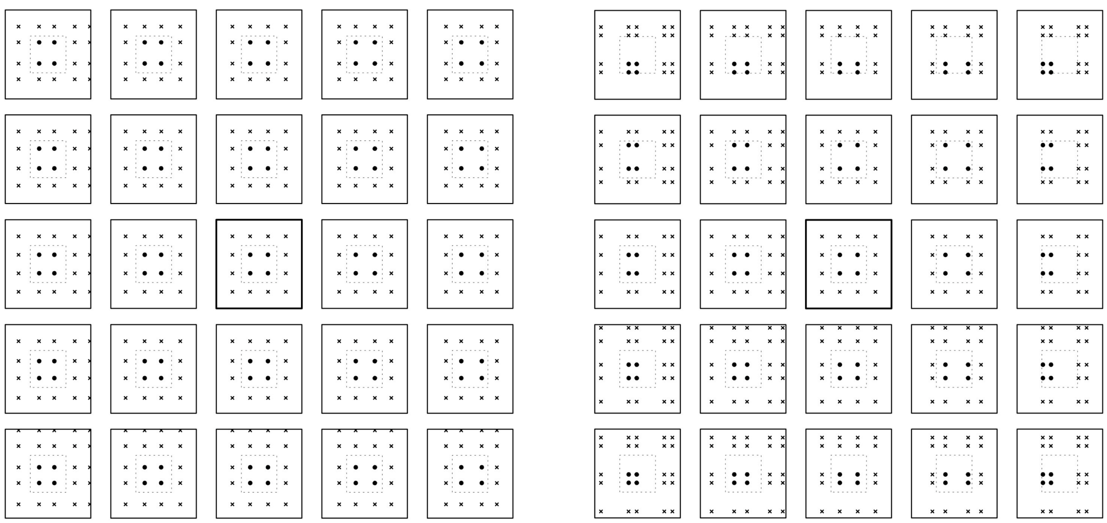

Figure 15: How distortion changes a perfect 2x2 dither pattern at various locations across the chip. The central panel shows the dither achieved at chip center, and the others show the pattern as seen at a 5 × 5 array of locations across the detector. The distortion reaches 2% at the edge of the chip (the outer panels). The left panels shows the dither pattern achieved with only 5-pixel offsets between dithers. The right panels shows the pattern achieved with a 20-pixel offset between dithers. The filled dots are the sampling that occur within the fiducial pixel (the dotted square), the other points are there simply to show the entire set of observations.

5.2.2 Achieving a good dither in both channels.

Sampling will be a major issue for the short-wave channel, which is critically sampled at 2 μm, but can reach all the way down to 0.7 μm. Sampling will also be a moderate concern for the long-wave channel, as that channel will be critically sampled at 4 μm, but will go down to 2.5 μm or so.

The SWC and LWC observe the same field simultaneously, so we will have to find dither patterns that do a good job simultaneously in both channels. Our strategy will be to optimize the sub-pixel patterns for the short-wave channel (since it is at its worst 3× times undersampled, compared with the LWC's worst case of 1.6×), but we will choose the whole-pixel offsets to ensure a satisfactory dither in the long-wave channel as well.

The SWC pixels are 32 mas in size, and the LWC pixels are 65 mas, and the ratio of the scales is very nearly a factor of two. This is in fact a similar situation to what was encountered when we tried to simultaneously generate a good dither in the PC chip (45 mas/pixel) and the WF chips of WFPC2 (100 mas/pixel). In that case, users often optimized the sub-pixel dither for the PC, then added whole PC pixels to the offsets to accomplish a reasonable dither in the WF chips. Here, the two scales are even more commensurate, and as a consequence it will be harder to use whole-pixel offsets to accomplish a perfect dither for both detectors.

5.2.3 Keeping the total number of exposures reasonable

The total number of dithers executed will be the product of the number of secondary and the number of primary dithers: $N_{TOT} = N_P \times N_S$. This means that if both primary and secondary dithering is to be done, there will likely be a large number of exposures taken, which will invariably require a lot of telescope time. Some mitigation of the total time can be accomplished by changing the duration of each exposure to get a certain number of dithers into a fixed amount of time, but the exposure time will likely be best optimized by considering the brightness of the target and background and cannot be arbitrarily changed without impacting data quality.

It is clear that we need to have as much flexibility as possible in selecting the number of primary and secondary dithers, so that we can find a primary/secondary combination that will that give us a total number of exposures (and hence, total exposure time) that is reasonable. For most of the primary patterns discussed in the previous section, we allowed the user flexibility in choosing any number of primary dithers, from 1 to 64. Only the "full-field" dithers were constrained to have 3, 6, or 9 primary pointings.

We will aim to do the same for the secondary dithers: we will provide patterns for any number of dithers that will make the most of the number chosen. We will treat dithers from 1 to 9 as special cases which will be optimized one by one, and patterns with more than 9 dithers will be treated in a more regularized way.

5.3 Secondary Strategy I: Patterns tailored for 1 to 9 dithers

Table 6 at the end of this section will list the dither pattern constructed for the 1- to 9-point patterns. These dithers were designed by hand with three priorities: (1) an even pixel-phase distribution in the SWC, (2) large dithers that space out the pointings over a 10×10-pixel region, so that no pointing will lie too close to any other, and (3) fine-tuning the large dithers to ensure a satisfactory pixel-phase distribution for the LWC.

Figure 16 shows the pixel-phase portion of these patterns for the SWC, the first priority. The filled circles indicate the sample spacing achieved within the typical pixel. The open circles are shown to indicate how this sampling plays out over a larger region. A reconstructed scene would be sampled at all these locations.

The sampling pattern shown in the figure (which often resembles a regular grid that is pitched at an angle to the image's coordinate axes) does not represent the way an image would be reconstructed. In general, when we combine multiple dithered pointings to create a higher-resolution composite stack, the output pixels will not correspond exactly with any of the input pixels. The exception to this are the N = 4 and N = 9 cases, where simple interlacing can be used to generate a high-resolution image from the multiple exposures (see Lauer 1999 for an explanation). However, even these "easy" cases are not so trivial when we must simultaneously correct for distortion. Some resampling will invariably be necessary. This can be done with Drizzle or another image-stacking routine.

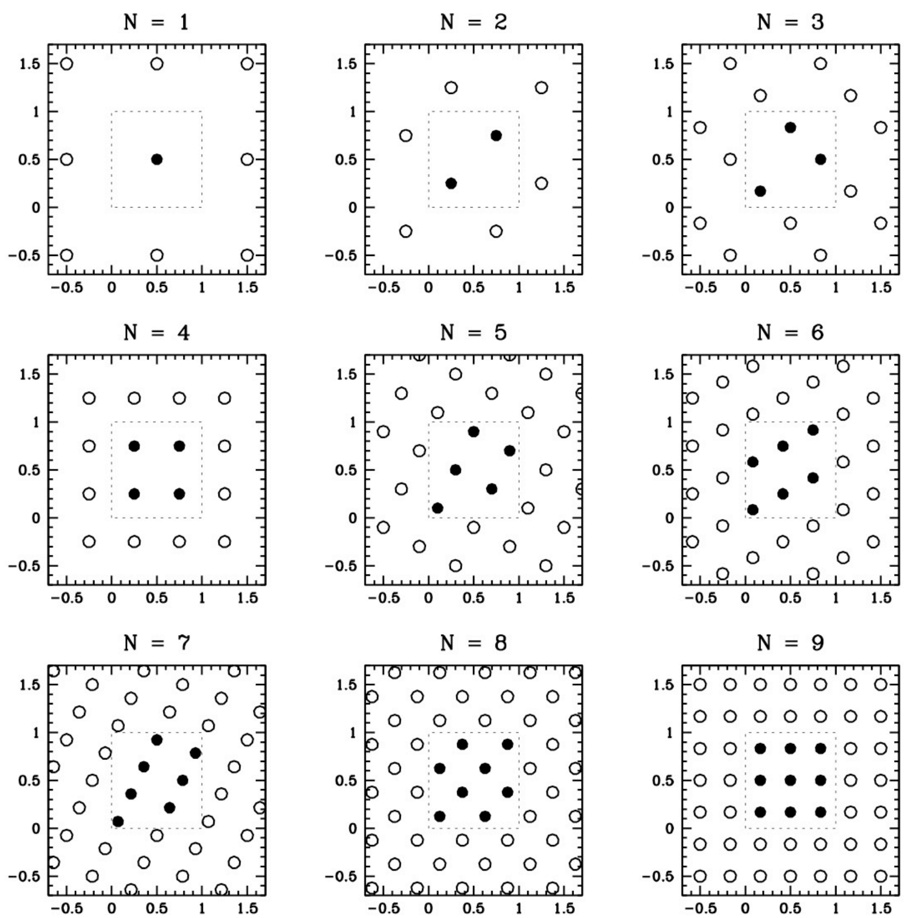

Figure 16: The dither patterns for N = 1 to N = 9, optimized for the short-wave channel. The filled circles indicate the dithers that fall the bounds of the fiducial pixel, and the open circles correspond to the dithers that fall in neighboring pixels. The first dither in the tables provided at the end of the section is at (0,0), but for clarity we offset all the dithers here to place them symmetrically relative to the fiducial pixel.

The pattern in Figure 16 shows the sub-pixel pattern as constructed for the SWC. The actual dither pattern will be constructed by adding whole pixels to this fractional-pixel pattern in order to effect a good dither in the long-wave channel and to space out the SWC dithers so that a star will occupy different pixels in different dithers. When possible, we will also try to construct our dithers so that stars are not observed multiple times in the same column or row, in an effort to avoid bad rows or columns. The whole-pixel shifts can span up to 10 pixels without introducing too much pixel-phase incoherence due to distortion.

Figures 15 and 16 show the entire dithers (whole pixel plus fractional pixel) achieved for the 2-point to 9-point dithers. The whole-pixel offsets spread the dithers out over the inner 10 × 10 SWC pixels, corresponding to 5 × 5 LWC pixels. The rightmost panels show that the LWC pixel phases are nicely spread out, as designed.

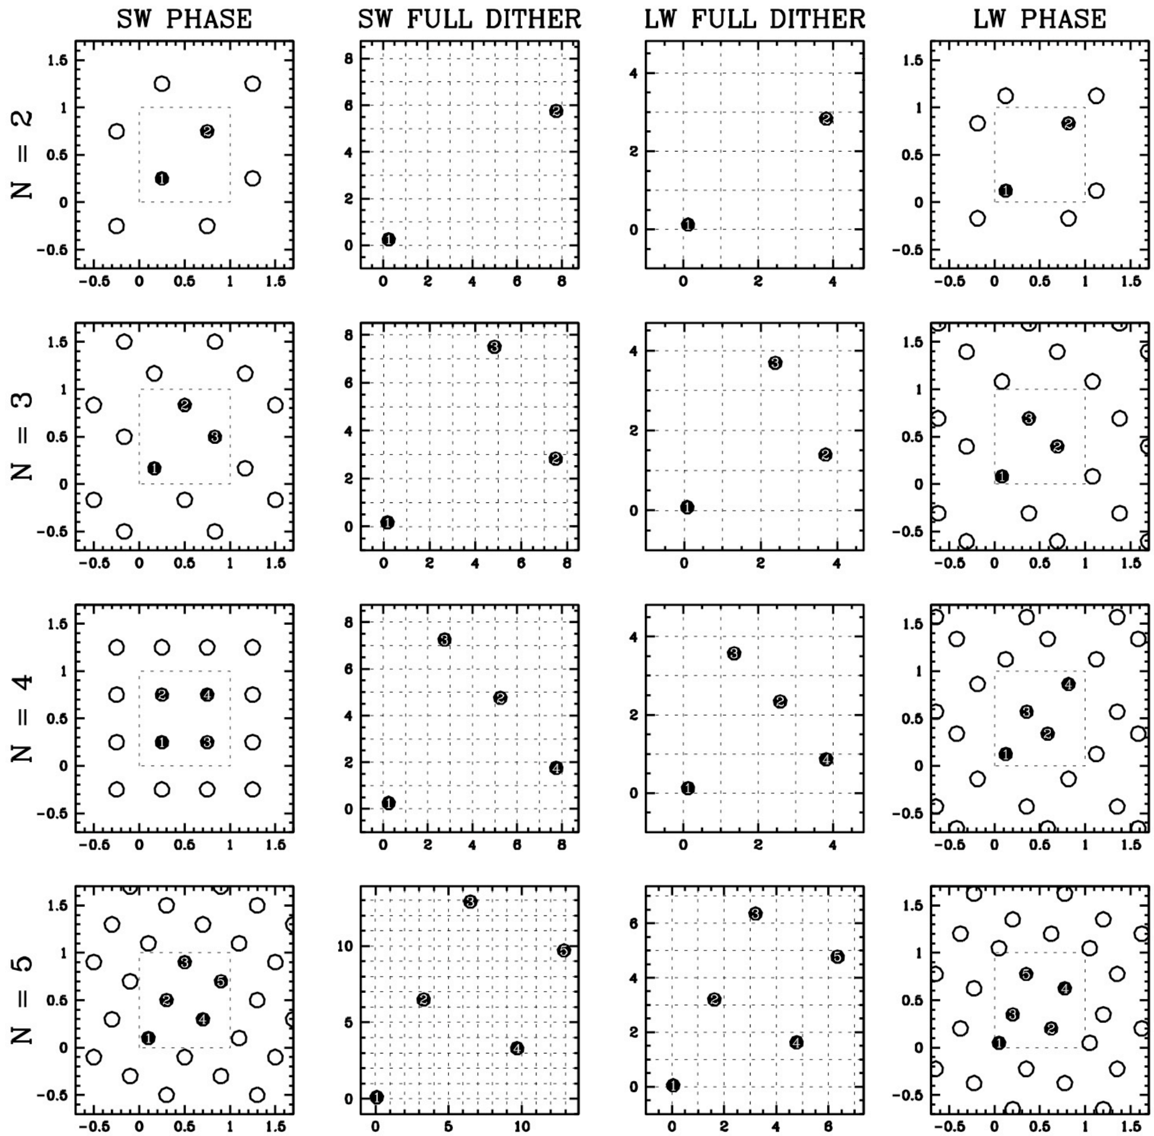

Figure 17: The first column shows, for each dither pattern, the pixel phases for that dither in the short-wave channel. The next column shows the full dither, whole pixel plus fractional pixel, for in the SWC frame. The third column shows the dither in the LWC pixel frame. Finally, the fourth column shows the pixel-phase dither achieved for the LWC.

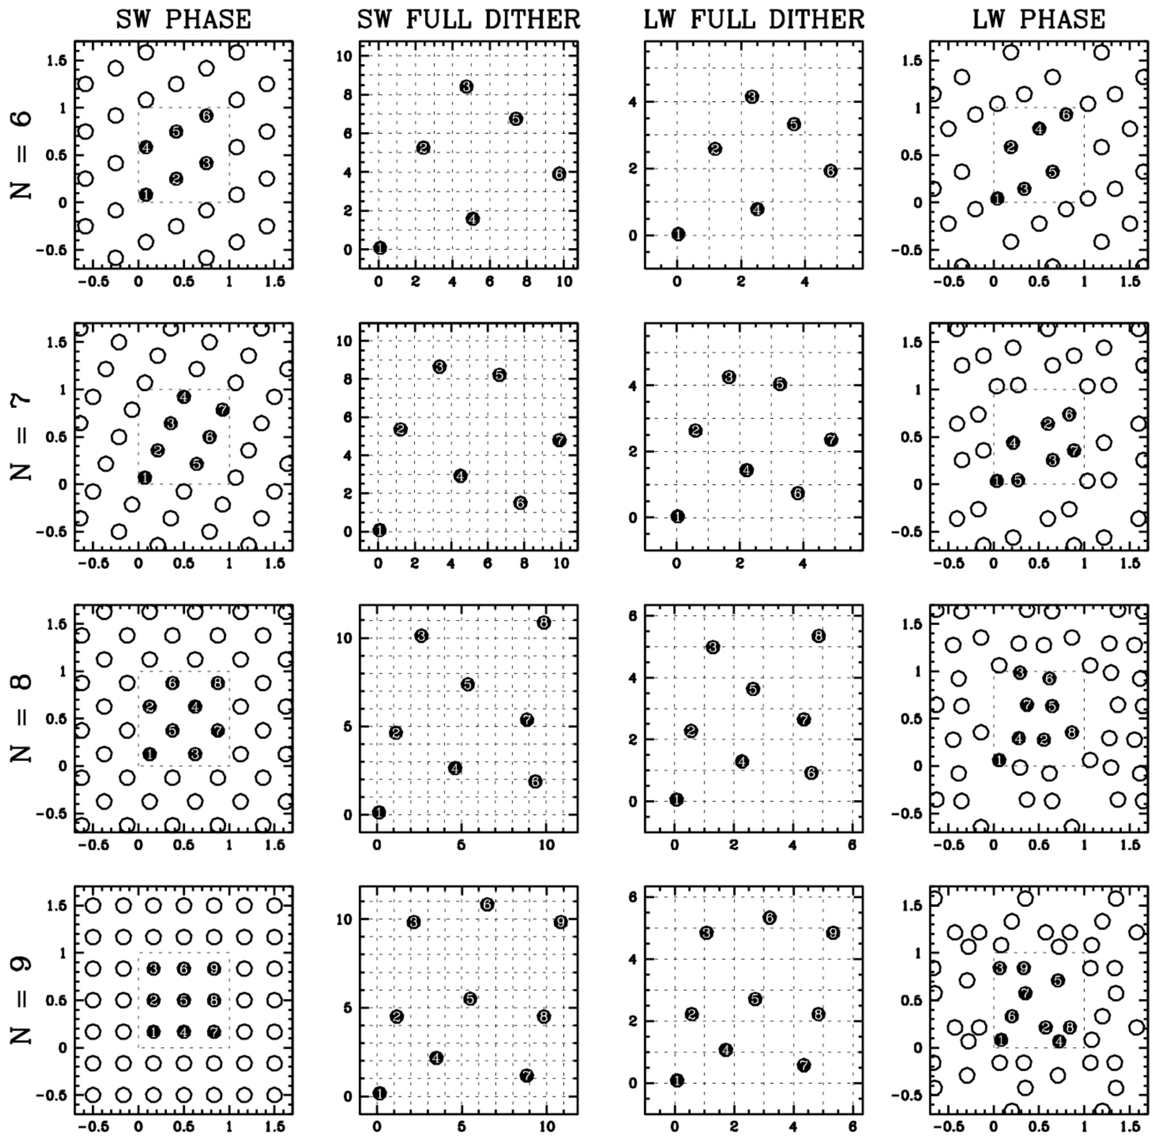

Figure 18: Same as Figure 17, but for the 6-point to 9-point patterns.

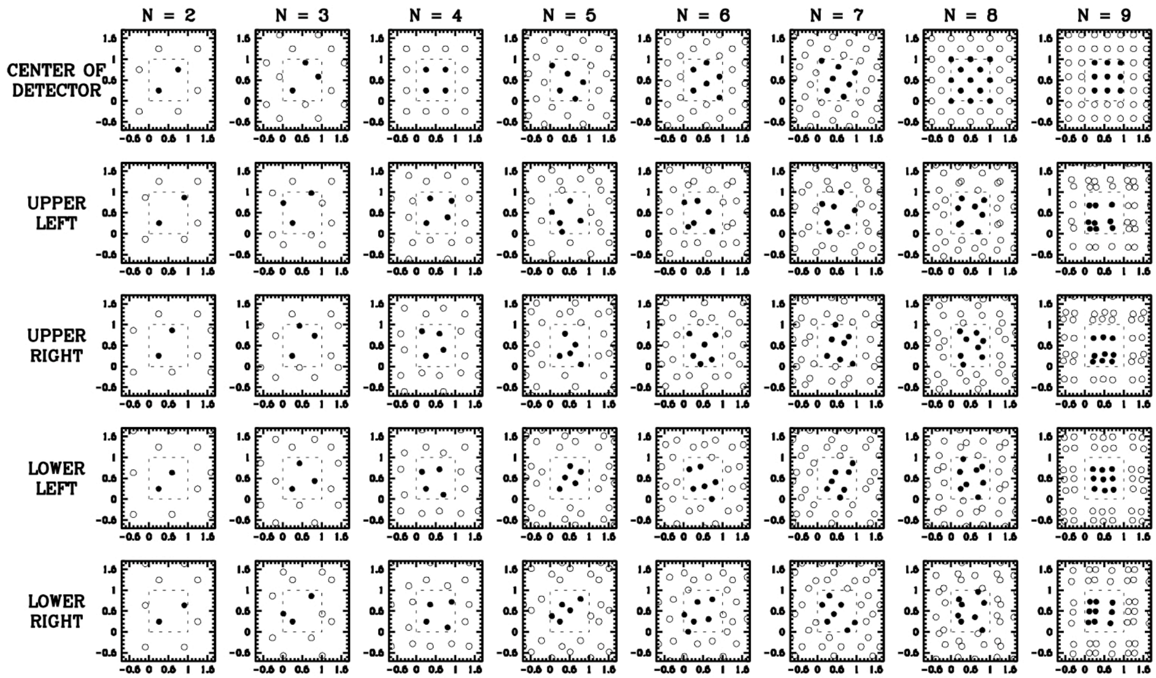

Figure 19: Upper row shows the pixel phases as constructed in the previous two figures. This is valid at the center of the detector. We simulate 2% distortion at the four corners of the field in the next four panels.

Figure 19 above shows how these patterns will be affected by distortion. The top row of panels show the "perfectly executed" dithers at the center of the chip, and the other rows show different manifestations of 2% distortion[1]. It is clear that the coverage becomes less optimal with distortion, but the pixel-phase coverage is still satisfactory for all examples shown. The goal is not to cover the pixel-phase space perfectly, but to have reasonable coverage so that all the structure present in the scene can be inferred. If we had more distortion (or, equivalently, larger dithers), then the coverage would become bad for many of the patterns.

It is unfortunate that in the LWC the dithers are all clustered within a 5 × 5 pixel region, and that there will not always be a full PSF separation between them when the wavelength is long. This means that the same star will land on some of the same pixels in different dithers, which will not completely mitigate bad pixels or pixel-to-pixel flat-field errors. Unfortunately, this cannot be avoided if we want to provide a 9-point pattern that improves the sampling by a factor of three along each axis, which will be necessary for the bluest filters of the SWC. One additional concern that arises when stars will be dithered onto nearby pixels is image persistence or latency. When we understand persistence better, we will be in a position to assess how this will impact tight dithers, such as those presented here. Also, when we get a more realistic handle on the actual amount of distortion present, we may find that it is much less than the 2% worse-case

[1] The upper left corner has an *x* scales and *y* scales that are 1.02 times that at the center, the lower left corner has an *x* scale that is 1.02 times that at the center, but a *y* scale that is 0.98 times that at the center, etc.

scenario managed here; as such, we can space these small dithers out more, thus ameliorating the persistence and LWC issues.

In the next sub-section we will provide a different sort of secondary dither that will ensure that all the dithers will be sep arated by at least a FWHM of the broadest PSF, so that the secondary dithers themselves can help to average over flat field errors and control for bad pixels.

5.4 Secondary Strategy II: A general pattern for up to 64 dithers

The 9-point pattern above was designed to sample the pixel-phase space every third of a pixel. More coverage than this is not critical for NIRCam, since even the bluest filter (F070W), which should have the sharpest PSF, will not have structure on finer scales than this. Nevertheless, it may be useful to have available more than 9 secondary dithers for the cases where we want to pursue a large number of independent observations over a small region.

So, we have constructed a 64-point dither pattern from 4-dither sets. Each 4-dither set was executed within a 10 × 10 pixel region. The dithers are spread out over this region sufficiently to ensure that a given star's image will land on a different set of pixels in each exposure, but the region is still small enough, that even with 2% distortion, the pixel phases within each 4-dither set will remain coherent all the way to the edge of the detector.

Larger regions than this will not maintain the coherence of the pixel phases, which is why we restricted ourselves to this region even for the 9-point dithers of the previous section. Nonetheless, many projects will want secondary patterns that allow more dithers and better spacing between dithers, in order to improve the bad-pixel and flat-field-error mitigation. So to do this, we will repeat the 4-dither pattern at multiple 20-pixel offsets. The pixel phases will not be coherent from 4-dither chunk to 4-dither chunk, but within each 4-dither chunk they will coherently supersample the pixel phase by a factor of two, so that we will be ensured this much sampling at a minimum. HST's Ultra-Deep-Field (UDF) observations were taken with a similar strategy, and turned out to be a magnificent example of what can be done with a good, deliberate dithering strategy.

The upper left panel in Figure 20 shows the 64-point pattern, which samples the face of a short-wave pixel in an even 8×8 array of super-samplings. Although sampling this finely will surely be overkill, we may as well spread them out to remove any phase-dependent artifacts in the reconstruction procedure. (An even sampling will make the output pixels less correlated with each other.)

The bottom left panel shows that the entire pattern covers a region that is about 60 × 60 pixels in extent. At the edge of the chip, distortion will make these fine pointings

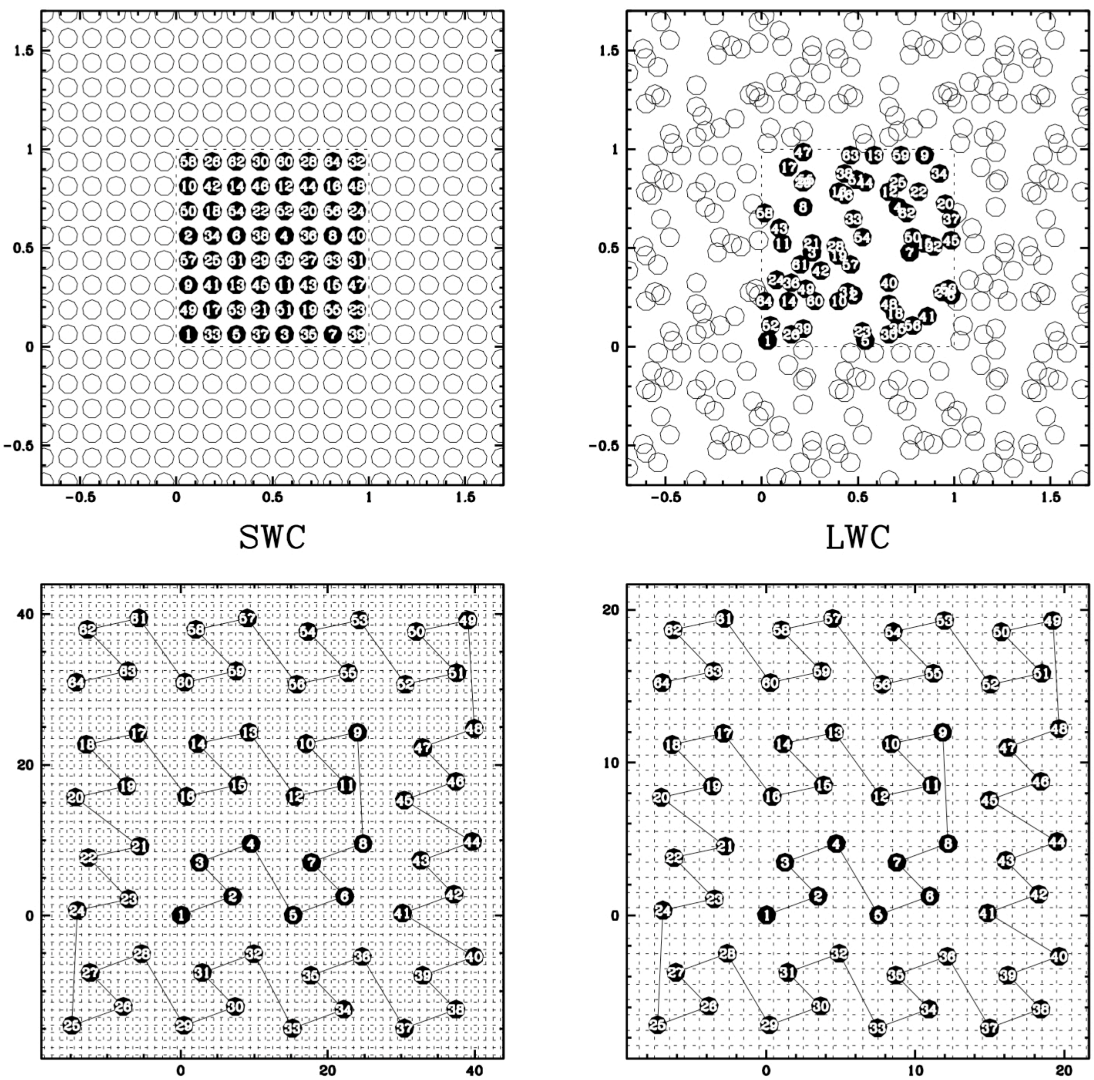

Figure 20: The entire 64-dither secondary pattern for the SWC (left) and the LWC (right). The pixel-phase distribution is shown in the top panel, and the actual dither (in pixels) is shown at the bottom.

incoherent, but for subarrays[2] and at the center of the pattern, the pixel face should be very evenly probed.

It is worth bearing in mind that the FGS may not be able to execute this dither pattern with the precision implied here; the relative pointing error on small slews could be up to 5 mas in each coordinate. For comparison, the 16×16 samples in Figure 20 are separated

---

[2] If the dithers are designed with the "average" scale of the detector in mind, then a subarray that is in a location of the field that has this "average" scale will get a nicely phased dither pattern. Subarrays at different locations will likely suffer from distortion, since we do not envision tailoring this 64-point pattern for each subarray.

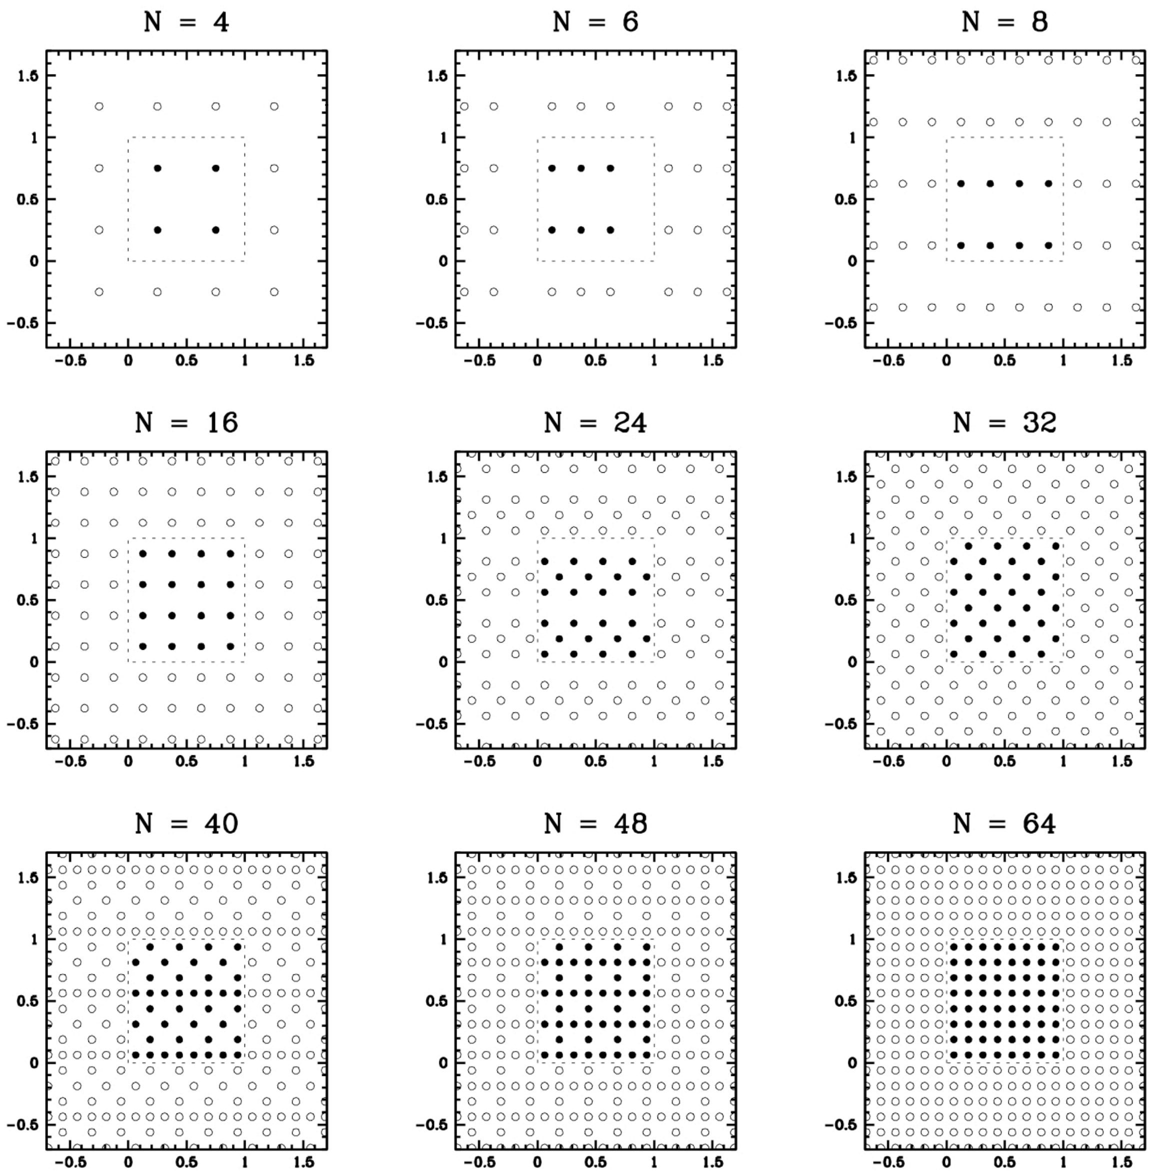

Figure 21: The SWC pixel-phase distribution for the N-dither secondary pattern for a variety of values of N.

by 4 mas. Still, we may as well aim for flat coverage; otherwise we will only be introducing more bias in the pixel-phase distribution.

This 64-dither pattern can be executed for any number of dithers from 1 to 64. The pattern executed will cover the face of the pixel as evenly as possible for the desired number of dithers, and will be as compact as possible, while never putting the same star on the same pixels more than once. Figure 21 shows how this pattern is built up, and that for any number of dithers, the pixel phase coverage will be as even as possible.

Table 7 in the next section gives the exact offsets (in arcseconds) for this pattern.

5.5 Tables for the secondary dithers

This section contains the tables with the offsets for the "tailored" and "general" secondary dither patterns.

Table 6: The prescribed offsets for the 1- to 9-point "tailored" secondary dithers. Note that the points shown in Figures 12 through 14 have been offset (for clarity) to center the pattern on the fiducial pixel. The units of these offsets are NIRCam short-wave-channel pixels, which are 32 mas on a side.

1-point	Δx	Δy
1	0.000	0.000

2-point	Δx	Δy
1	0.000	0.000
2	7.500	5.500

3-point	Δx	Δy
1	0.000	0.000
2	7.333	2.667
3	4.667	7.333

4-point	Δx	Δy
1	0.000	0.000
2	5.000	4.500
3	2.500	7.000
4	7.500	1.500

5-point	Δx	Δy
1	0.000	0.000
2	3.200	6.400
3	6.400	12.800
4	9.600	3.200
5	12.800	9.600

6-point	Δx	Δy
1	0.000	0.000
2	2.333	5.1667
3	4.666	8.333
4	5.000	1.500
5	7.333	6.667
6	9.666	3.833

7-point	Δx	Δy
1	0.000	0.000
2	1.143	5.286
3	3.286	8.571
4	4.429	2.852
5	6.571	8.143
6	7.714	1.429
7	9.857	4.714

8-point	Δx	Δy
1	0.000	0.000
2	1.000	4.500
3	2.500	10.000
4	4.500	2.500
5	5.250	7.250
6	9.250	1.750
7	8.750	5.250
8	9.750	10.750

9-point	Δx	Δy
1	0.000	0.000
2	1.100	4.333
3	2.000	9.667
4	3.333	2.000
5	5.333	5.333
6	6.333	10.667
7	8.677	1.000
8	9.667	4.333
9	10.667	9.667

Table 7: The prescribed offsets for the "general" secondary dither pattern, which can have up to 64 points. This pattern ensures at least ×2 super-sampled pixel-phase coverage for all parts of the detector with no star landing on the same pixels more than once. The units of these offsets are SWC pixels, which are 32 mas on a side.

N	Δx	Δy
1	0.000	0.000
2	7.000	2.500
3	2.500	7.000
4	9.500	9.500
5	15.250	0.000
6	22.250	2.500
7	17.750	7.000
8	24.750	9.500
9	24.000	24.250
10	17.000	22.750
11	22.500	17.250
12	15.500	15.750
13	9.250	24.250
14	2.250	22.750
15	7.750	17.250
16	0.750	15.750

N	Δx	Δy
17	−5.875	24.125
18	−12.875	22.625
19	−7.375	17.125
20	−14.375	15.625
21	−5.625	9.125
22	−12.625	7.625
23	−7.125	2.125
24	−14.125	0.625
25	−14.875	−14.625
26	−7.875	−12.125
27	−12.375	−7.625
28	−5.375	−5.125
29	0.375	−14.625
30	7.375	−12.125
31	2.875	−7.625
32	9.875	−5.125

N	Δx	Δy
33	15.125	−15.000
34	22.125	−12.500
35	17.625	−8.000
36	24.625	−5.500
37	30.375	−15.000
38	37.375	−12.500
39	32.875	−8.000
40	39.875	−5.500
41	30.125	0.250
42	37.125	2.750
43	32.625	7.250
44	39.625	9.750
45	30.375	15.250
46	37.375	17.750
47	32.875	22.250
48	39.875	24.750

N	Δx	Δy
49	39.000	39.125
50	32.000	37.625
51	37.500	32.125
52	30.500	30.625
53	24.250	39.125
54	17.250	37.625
55	22.750	32.125
56	15.750	30.625
57	9.000	39.375
58	2.000	37.875
59	7.500	32.375
60	0.500	30.875
61	−5.750	39.375
62	−12.750	37.875
63	−7.250	32.375
64	−14.250	30.875

6.0 ADDITIONAL CONSIDERATIONS

Now that we have set out the Primary and Secondary dithering strategies separately, it is worthwhile to consider how they might interact with each other, and with efforts to image even larger fields, and to accomplish other specific tasks such as MSA pre-imaging and self-calibration.

6.1 The Interaction between Primary and Secondary dithers

Figure 15 and Figure 19 showed how distortion can affect the distribution of pixel phases achieved by the same dither in different places of the detector — even for dithers as small at 10 pixels. We have mentioned before that dithers much larger than this (in the presence of 2% distortion) will essentially randomize the pixel phases achieved. So, all the primary dither patterns will end up placing each source at a random location with respect to the pixel grid.

Getting a series of observations with randomized pixel phases will certainly be better than having no pixel-phase coverage at all. But the coverage we get with random pixel phases is considerably worse than the coverage we get with carefully placed sub-pixel dithers. The obvious way to ameliorate this is to take a set of secondary dithers at the location of each of the primary dithers. These secondary dithers should have coherent pixel phases, thus ensuring that many of the dithers will be well space out.

An upcoming report (Anderson 2010) will aim to discuss in more detail how image reconstruction and point-source identification are impacted when the pixel-phase coverage is inadequate. It will look into what the trade-offs are between (for example) taking a 3-point primary pattern with 3 secondary dithers at each one, and taking a 9-point primary pattern. In general taking fewer primary and more secondary observations will result in better pixel-phase coverage, but will be more subject to flat-field errors and possibly regions of the field with substantially less coverage. The report will also use model PSFs to show how much sub-pixel structure there is in NIRCam images taken through different filters, so that users can know how much pixel-phase coverage is beneficial. A general discussion of primary-vs-secondary dithers, with a focus on scientific objectives, will also be a part of the NIRCam Users' Manual.

6.2 Tiling

The primary dither patterns discussed here are designed to fit well into patterns that will mosaic larger fields. Figures 4, 5, and 6 showed how the $350'' \times 140''$ NIRCam field fits nicely into a tiled pattern.

This report has not touched on how these tiling sequences are to be generated, but it is likely that they will be constructed in APT with an efficient way to slew from one field to the next. The dithers in this study can be executed in any order, so it may be worthwhile to consider whether when tiling a field in one direction, it may be more efficient to reverse the order that the primary dithers will be taken.

It is worth noting that, since the NIRCam field has a four-fold symmetry (see Figure 2), the primary dithers extend only over about half of the NIRCam footprint in each dimension. This means that the spacing between the tilings will be much greater than the primary dithers, so there is no chance of using the same guide star for different tiles.

6.3 Pre-imaging for NIRSpec MSA observations

In order to use NIRSpec's MSA (Micro-Shutter Array) efficiently, we will need to know the location of each source within the NIRSpec field to well within the size of an MSA shutter (0.2″ × 0.46″). The current expectation is that positions will have to be known to about 5 mas, so that after target acquisition the achieved locations of targets within the shutters will be known to better than 15 mas. This precision is necessary to estimate expected throughput through the MSA shutters and calibrate the resulting spectrum. In many cases, this will require pre-imaging by NIRCam.

The NIRCam pre-imaging will typically be taken a few months before the spectroscopy, since it will take a some time to analyze the NIRCam images, select the sources, then optimize the MSA configuration before submitting the new observations into the long-term plan. Targets that are not in the CVZ will have visibility periods only two times a year, often for only 50 days at a time, so it is likely that the pre-imaging will be done during one period of visibility and the spectroscopy will be done at the next one, 3 or 6 months later.

Since the goal of pre-imaging is to get source positions with a precision of about 5 mas (0.15 NIRCam pixel), it shouldn't be necessary to take a set of images with a careful sub-pixel dither. A single exposure covering each target of interest should be sufficient for most programs. For this reason, it is worthwhile to devise a bare-bones dither strategy that accomplishes this with a minimum of overhead.

Figure 22 shows the outline of the NIRCam MSA footprint on the field covered by NIRCam (magenta). This is the nominal orientation of the two detectors. For targets at higher ecliptic latitudes, it will be possible to observe the field in NIRCam with a different orientation than is achieved with NIRSpec, but targets at low ecliptic latitude can a achieve at most a 10° relative roll between the two (modulo 180°). Our plans here will be optimized for this no-roll worst-case scenario. It is clear, though, from the figure that even if we could orient the two parallel to each other (at 45° to nominal), it would still require two NIRCam tiles to adequately cover the MSA footprint: NIRCam is only 135″ tall, while the NIRSpec MSA covers 204″ in its shortest dimension. These two tiles would clearly require different guide stars, so the minimum number of guide stars required for pre-imaging will be two, irrespective of orientation considerations.

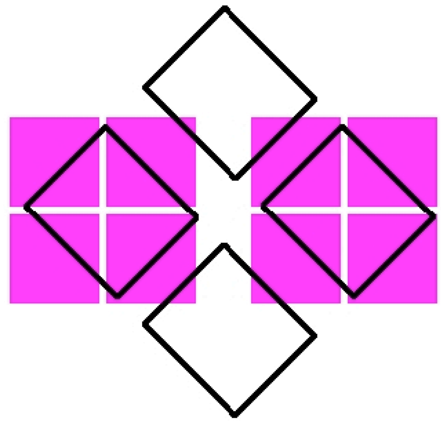

Figure 22: The overlap between the NIRSpec and NIRCam footprints (black outline and magenta, respectively).

The full-field primary patterns devised in Section 4.1 all require a minimum of two guide-star acquisitions per tile location. We could certainly use these standard patterns to achieve the aims of NIRSpec pre-imaging. Figure 23 shows how a 1 × 2 tiling with the standard "full-field" patterns produces full coverage over the entire MSA field.

The tight 3-point pattern in the upper right provides the best match to the NIRSpec FOV, but the looser 3-point tile pattern in the upper left surveys a slightly larger region, allowing perhaps more flexible pointing of the MSA for the spectroscopic follow-up. We show the 6- and 9-point dithers in the bottom panels since some observers may want to accomplish more than simple source identification with the pre-imaging. Some observers may want to survey the field with multiple exposures through multiple filters to allow for an empirical photometric cross-calibration for their targets between NIRCam and NIRSpec. Naturally, larger NIRCam tiling patterns can be used to survey larger regions for programs where the target is larger and multiple MSA placements are anticipated.

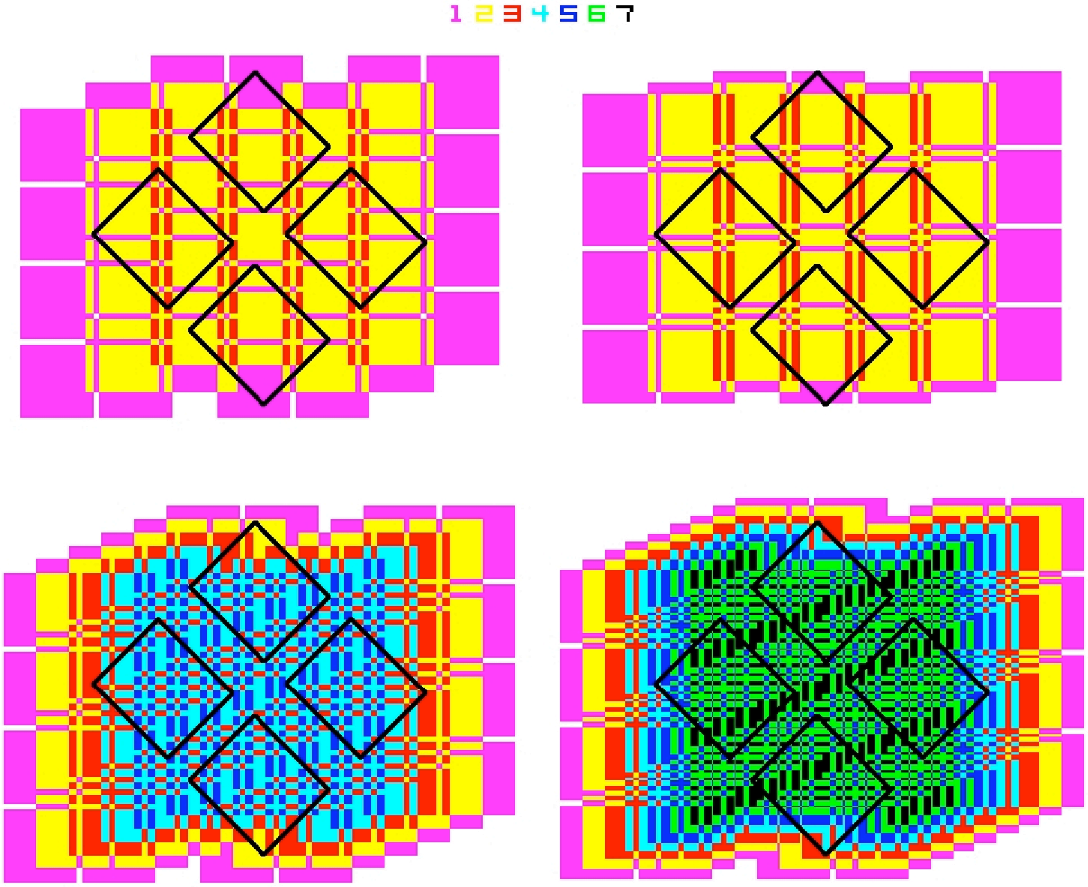

Figure 23: The effect of doing a 1×2 tiling of the NIRSpec field with the four full-field primary dither patterns discussed in Section 4.1. Upper left is the 3-point tile pattern, upper right is the 3-point tight pattern, lower left is the 6-point tile, and lower right is the 9-point tile. The colors indicate how many exposures each point in the field is covered by. The legend is given at the top.

The patterns shown in Figure 23 will likely all require four guide-star acquisitions, two for each tiling. Each guide-star acquisition will add about 2 minutes to the program overhead. It would be possible to do the entire mosaic with only two guide stars (one per tiling), if the user is willing relax the requirement of covering every single point in the field in at least one exposure. This strategy would involve dithering over the 40″ gap, but not covering all the SCA gaps. The consequence would be only a few small holes in coverage.

Figure 24 shows the coverage achieved with such a strategy. There are small regions with no coverage (totaling only about 2.5% of NIRSpec's total MSA footprint), but most of the field is covered by at least one exposure. This could be acceptable if there are a large number of sources that need not all be observed with the MSA. Such a minimal strategy could also be acceptable even in the case where there are sources that cannot be missed, if we can ensure that the small gaps in the pre-image coverage can be placed at harmless locations in the field.

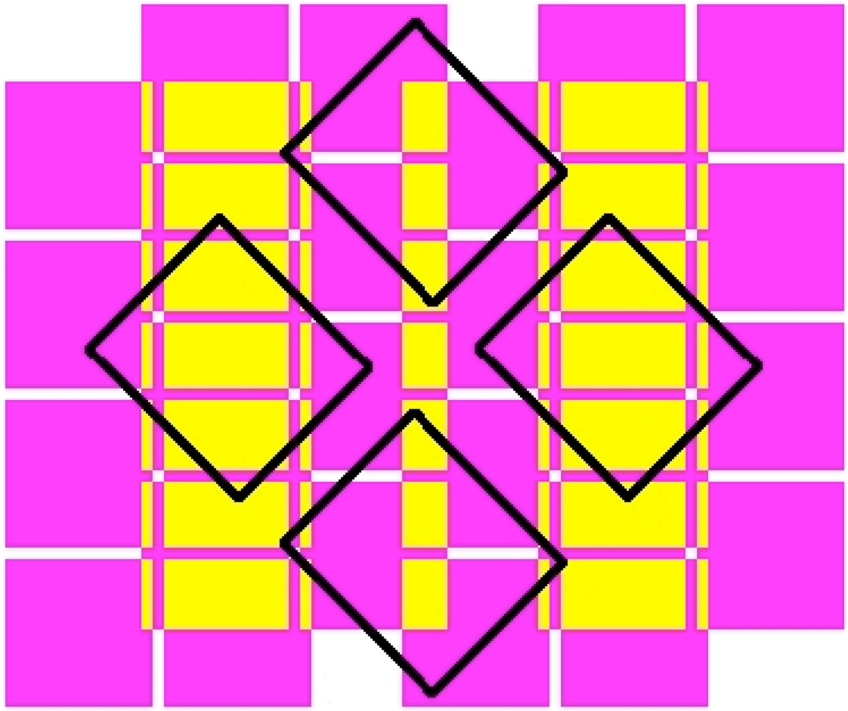

Figure 24: This shows the coverage achieved by a mosaic with 4 pointings: two dithers at each of two tiles. The magenta indicates regions of the field covered in one exposure, and the yellow indicates regions covered in two.

It is unclear whether this pattern should be characterized as one set of dither patterns or two, since it will formally require two visits. In Table 8 we will treat it as one dither pattern, with the offsets as given relative to the center of the field.

Table 8: The dither pattern for "minimal" NIRSpec pre-imaging.

N	ΔX	ΔY
1	–30″	–87″
2	+30″	–53″
3	–30″	+53″
4	+30″	+87″

6.4 Dithering for self-calibration

The dither patterns developed here have been focused on scientific applications, under the assumption that the detectors are already calibrated and understood as well as possible, and all that remain are random-type calibration errors, such as pixel-to-pixel flat-field errors and random region-to-region L-flat type errors.

These dither patterns should also be useful for many calilbration purposes, both formal L-flat calibration programs using standard fields or self-calibrations within scienfic programs. While scientific programs will be restricted to these patterns, calibration programs will not be: special templates will allow them to dial in any pattern they find useful. Much of the formal L-flat-type calibrations will come from sky flats, which will not involve dithering. Nevertheless, to the extent that calibration will be done (or verified) using dither patterns, it is valuable to consider how well the standard observing patterns presented here can be used for this purpose.

Arendt, Fixsen & Moseley (2000) discuss effective dithering strategies for sky-based self-calibration. They evaluated several different dithering strategies (VLA-type arrays, geometric patterns, Reuleaux patterns, random patterns, etc) and found that the patterns that provided the most accurate self-calibrations were those that had the fewest redundant dither baselines. In their least-squares formalism, the quality of each pixel's calibration depends on how many other pixels its value can be compared directly against; if there are redundant baselines in the dither pattern, there are fewer independent pixel-to-pixel comparisons possible.

The 45-point dither discussed in § 4.1.6, has 45×44 different baselines, and only about 10% of these are non-redundant. (One reason for the redundancy is that the pattern is made up of five 9-point patterns.) This is not optimal for self-calibration, but it is also not terrible: each pixel still has over 200 pixels it can be compared against in a completely independent way. Still, if it becomes necessary to do sky-based self-calibration by means of dithering, we may end up constructing specialized patterns for this purpose. At this juncture, though, the plan is to do L-flats by either dithering about a star-filled calibration field (which would not face the issues investigated by Arendt et al. above) or by observing blank sky in parallel.

7.0 CONCLUSIONS

All but a very small subset of NIRCam observations will be dithered. The expectation is that no two exposures will be taken through the same filter with the same pointing. The reason for this is that each independent image we have of the scene increases our information much more than does an additional exposure at the identical pointing: a second pointing can help to identify and remove cosmic rays, image defects, and can ameliorate issues related to undersampling.

This mandatory dithering is a break from the HST tradition, where the user has been free to design whatever dithering strategy he or she sees fit, using a detailed prescription of `POS_TARG`s, a canned dither strategy, or no dithering at all. Sadly, much of the HST archive is of compromised value because many users chose no dithering. The no-dither option makes cosmic-ray removal, image stacking, and source collation slightly easier in a bookkeeping sense, but it makes more sophisticated analyses and evaluation of systematic-type errors more difficult, if not impossible. Therefore, in an effort to make the observations and the archive as valuable as possible, dithering will be done automatically for almost all observations, using a pre-canned dithering strategy that is optimized for the scientific goals of the program.

This report has considered the science goals of various projects that will use NIRCam. There are a wide variety of projects, such as: high-precision astrometry and photometry of stellar fields, extra-galactic studies that aim to measure detailed morphological parameters for barely resolved galaxies, lensing studies that require the highest possible resolution and shape information, searches for first-light objects that will have to determine whether an object is barely resolved or unresolved, and surveys that will have to map large areas of the sky in as uniform a way as possible. Each of these projects will have different requirements as far as field coverage, resolution, and flat-field mitigation are concerned, and will benefit from very different dithering strategies.

This report presents sets of primary and secondary dither patterns to address the needs of the various programs. The primary patterns consist of larger dithers and both mitigate large-scale flat-fielding errors and ensure even coverage of all parts in the field. The secondary patterns are designed to mitigate the effect of undersampling: by placing sources at a variety of locations with respect to the pixel grid, we can infer more information about their shape than a single exposure can provide.

The primary patterns discussed here come in three varieties. The first variety is called `full-field dithers`, and involves large slews to cover both the 5″ gaps between SCAs and the larger 40″ gap between modules. The dithers come in 3, 6, and 9-pointing options. The `intra-module dithers` can be executed for anywhere between 1 and 16 dithers, and cover the plus-shaped gaps between the SCAs. The `intra-SCA dithers` can be executed with anywhere between 1 and 25 dither pointings and will move the target of a specified size around the SCA.

The secondary patterns will come in two varieties: tailored and general. The tailored patterns will have a tight configuration for up to 9 dithers, so that they will maintain pixel-phase coherence across the detector. The general dithers will have decent pixel-phase coverage, but will ensure better separation between the dithers. Up to 64 general secondary dithers can be executed.

These dithering options should be easy to specify in APT. Under "Primary Dithers" there will be ten options, and each option will have a number of dithers that can be specified via a sub-menu. The "Secondary Dithers" will have two options, and will be even easier to specify. Table 9 summarizes the primary and secondary dithering options.

Table 9: The options and sub options available for the primary and secondary dithers.

Type of Primary Pattern	# Dithers
Full-field tight	3
Full-field tile	3, 6, 9, 15, 21, 28, 36, or 45
Intra-module	1, 2, 3, … , 16*
Intra-SCA: SWC-small	1, 2, 3, … , 25*
Intra-SCA: SWC-medium	
Intra-SCA: SWC-large	
Intra-SCA: LWC-small	
Intra-SCA: LWC-medium	
Intra-SCA: LWC-large	
Minimal NIRSpec pre-imaging dither	4

Type of Secondary Pattern	# Dithers
Tailored	1, 2, 3, … 9
General	1, 2, 3, … 25*

* These N-point patterns can be achieved by executing the first N positions in a single list.

In general, most programs will use some combination of primary and secondary dithers to accomplish their objectives. The secondary pattern will be executed at each of the primary-dither locations, so the total number of dithers will simply the product of the number of primary and secondary dithers: $N_{TOT} = N_P \times N_S$. All but the full-field primary patterns allow for almost any number of dithers, so that we can essentially dial in any non-prime number of N_{TOT} dithers by varying N_P and N_S.

The total exposure time will simply be the single-exposure integration time, multiplied by N_{TOT}. The amount of time the entire observation will take will be longer by the time it takes to execute the dithers. Most of the secondary-dither patterns can be executed with less than 10 seconds slew-time between dithers. The primary patterns will take longer to execute. The primary dithers will take about 30 seconds slew time.

The horizontal span of the "mosaic" primary dithers will be 116″ for the 3-point dither and 156″ for the 9-point dither. The FGS field of view will be 140″ × 140″. Even though the 3-point pattern fits within the FGS field, it is not likely we will find a guide star that can go all the way from one side of the FGS to the other. Therefore, all observations taken with a full-field primary dither will require a guide-star reacquisition between them, which will necessitate that it be broken into two visits. The dithers have been ordered in such a way as to require no more than two guide stars per tile. There is no general way to know at which point in the pattern execution we will need to switch guide stars, so there will be no way to divide the set of observations *a priori* into visits. Either APT or something further downstream will have to do this.

Another issue that comes up with the changing of guide stars during a visit is that the pointings after the guide-star switch may contain a systematic offset relative to the first set, due to inaccuracy in the guide-star positions. For HST, this has been as large as 1.5″ and has made it difficult to combine mosaics taken at different pointings in an automated way. It is anticipated that the guide-star positions should be good to 0.1″ by the time JWST is launched, so this will introduce only a shift of about 3 SWC pixels, which is still a lot when the aim is generating a well-sampled stack. For this reason, it may be necessary to develop better strategies for empirical image co-registration. The primary patterns discussed in Section 4 were designed with some tolerance for either pointing errors or bad pixels at the edges of detectors, so guide-star errors should not impact our ability to cover the field.

One dither pattern that has been suggested, but was not included here, includes dithering along the coronagraphic wedge. The fiducial location of a source along the wedge will be specified by the wavelength of the filter, but it may be useful to shift that position by a few pixels to average over flat-field errors, and possibly to improve any pixel-phase sampling along the wedge direction.

8.0 References

Anderson, Jay & King, Ivan R. 2000 PASP 112 1360 "Towards High-Precision Astrometry with WFPC2 I. Deriving an Accurate Point-Spread Function"

Anderson, Jay 2010a. JWST - STScI – 000NNM "A Least-Squares Approach to Combining Undersampled Dithered Images"

Anderson, Jay 2010b. JWST - STScI – 000NNN "How Specific Dither Patterns Allow Us to Reconstruct the Astrometric Scene"

Arendt, Richard G., Fixsen, D. J., & Moseley, S. Harvey. 2000 ApJ 536 500. "Dithering Strategies for Efficient Self-Calibration of Imaging Arrays"

Fruchter, Andrew & Hook, Richard N. 1997 Proc SPIE Vol. 3164, p. 120-125, "Novel Image Reconstruction Method Applied to Deep Hubble Space Telescope Images."

Fruchter, Andrew & Hook, Richard N. 2002 114 144. "Drizzle: A Method for the Linear Reconstrution of Undersampled Images"

Koekemoer, Anton & Lindsay, Kevin 2005 JWST-STScI-000647. "An Investigation of Optimal Dither Strategies for JWST"

Lauer, Tod 1999 PASP 111 227. "Combining Undersampled Dithered Images"

Mather, John 2005 JWST-OPS-002018 Rev B. "Mission Operations Concept Document" (MOCD)