

Dithering for ACS and WFC3

Primes and Parallels

Jay Anderson and Norman Grogin
October 18, 2023

ABSTRACT

The dither patterns available in APT were designed with only one instrument in mind — the instrument that is “prime”. We explore here how effective the prime-instrument-based “box” patterns are for observations taken in parallel. To this end, we develop a metric to describe good and bad pixel-phase coverage. Not surprisingly, we find that a pattern that has been optimized for one detector observed in prime is often quite poor for another detector observed in parallel. We construct some additional patterns in the form of POS-TARGs that achieve a good sub-pixel dither for both prime and parallel observations for ACS/WFC and the two WFC3 cameras. It is worth noting that on account of distortion, there are sometimes tradeoffs between achieving good pixel-phase coverage and mitigating artifacts (bad pixels, blobs, persistence, bad columns, etc). In the process of this exploration, we discovered that the then-current ACS box dither likely got corrupted by post-SM4 changes in the SIAF files. We have since corrected those dither specifications to provide the intended sub-pixel phase sampling. The current document now provides 2-point and 3-point dithers in Appendix B that are good in prime/parallel instruments, in addition to the 4-point dithers. Users can group N dithers into sets of 2, 3 or 4 to achieve a good N -pt dither in both prime and parallel.

1. Introduction

Achieving a good dither is an important aspect of observing with detectors that undersample the point-spread function (PSF). In addition to mitigating the impact of hot pixels and bad pixels, dithering can improve the sampling of the astrophysical scene. Without good sub-pixel dithering there can be ambiguity between the location of a source within a pixel and the source’s true size (see Figure 1 in Anderson & King 2000). Good sub-pixel dithering can break this degeneracy.




Figure 1: The HST focal plane. (Taken from Figure 3.1 of the ACS Instrument Handbook.)

The ACS/WFC, WFC3/UVIS, and WFC3/IR cameras are all undersampled to different degrees, and all benefit from sub-pixel dithering. The available “canned” dither patterns¹ have been designed with only one instrument in mind —the instrument that is “prime”. If observations are taken in parallel, then it is unknown how good or bad the dither will be for the instrument in parallel. **Figure 1** shows the HST focal plane. It is clear that the pixel axes in ACS/WFC and WFC3 are oriented at roughly 45 degrees with respect to one another. Furthermore, each of the instruments has a different plate scale from the others, so one might expect an essentially “random” sub-pixel dither pattern for instruments observed in parallel.

In this ISR, we explore in detail how effective the current patterns are for observations taken in parallel. This can be done empirically by examining the image shifts observed in the parallel instrument when observations are taken with the recommended 4-point box pattern in the prime instrument. It can also be done theoretically by constructing an empirical linear relationship between dither shifts in the two WFC3 cameras and ACS’s WFC so that any potential dither pattern can be evaluated for the prime and parallel instrument.

Once we are able to relate dithers in a prime instrument to dithers in a parallel instrument, we must evaluate the quality of the achieved dither. We construct two different metrics to evaluate the efficacy of dithers and then use these metrics to evaluate the quality of the dither achieved in the parallel instrument.

There are four combinations of prime-parallel instruments:

1. ACS/WFC prime, with WFC3/UVIS parallel
2. ACS/WFC prime, with WFC3/IR parallel
3. WFC3/UVIS prime, with ACS/WFC parallel
4. WFC3/IR prime, with ACS/WFC parallel

¹ See e.g. <https://www.stsci.edu/hst/instrumentation/acs/proposing/dither-strategies> and <https://hst-docs.stsci.edu/hpiom/chapter-7-pointings-and-patterns/7-3-wfc3-patterns>

We evaluate the standard “box” dithers currently available in APT and find that three out of the four combinations result in quite unsatisfactory dithers in the parallel instrument. We find this both empirically by evaluating achieved dithers and theoretically by developing a mapping between POS-TARGs in prime instrument and the resultant pixel shifts in prime and parallel instruments.

We devise a scheme for constructing new 4-point dithers that are “perfect” in the prime instrument and as “good as possible” in the parallel instrument. This must be done separately for each of the four combinations of prime-parallel. We present the details for each of the four combinations along with some caveats.

This ISR is organized as follows. [Section 2](#) presents the standard “box” dithers for the three instruments in prime. We use actual observations to show the quality of the dither achieved in the prime and parallel instruments. In [Section 3](#) we examine how non-linear distortion can cause any sub-pixel dither pattern constructed at the center of a detector to lose coherence across the detector.

In [Section 4](#), we analyze a large set of dithers made with the four prime/parallel combinations and construct linear relationships between the commanded POS-TARGs and the achieved dithers in the prime and parallel pixel space. In [Section 5](#), we take a step back and consider what makes one dither pattern better than another and come up with two different metrics for evaluating dither quality. [Section 6](#) evaluates all possible “perfect” dithers in each prime instrument and determines which produces the best possible dither in the parallel instrument. The best prime-parallel dither patterns are presented in tables. Finally, [Section 7](#) reiterates that on account of distortion, it is sometimes not possible to simultaneously achieve a good sub-pixel dither and other aims of dithering, so we provide users with some ideas of how to make the most of these new patterns.

2. Evaluating standard “box” dithers in parallel instruments

The standard 4-point box dithers available in APT were designed to achieve several objectives. First, they aimed to image the scene at four well-spaced sub-pixel locations to allow an optimal super-sampled reconstruction of the scene. It is conventional wisdom that $(0.00, 0.00)$, $(0.50, 0.00)$, $(0.00, 0.50)$, and $(0.50, 0.50)$ constitutes the best possible sub-pixel dither, since that makes it possible in principle to interlace the pixels to get $2\times$ improved resolution in each dimension². All of the standard box dithers are designed to achieve this sub-pixel coverage.

In addition to having this critical sub-pixel spacing, the dithers have whole-pixel offsets that ensure that a particular bad pixel or bad column will impact a given small source in only one exposure out of the four. Since these dithers were devised for the prime instrument with no consideration given to parallels, the first step for us here was to explore what dither was achieved in the parallel instruments from the standard 4-point “box” dithers in the prime instruments.

² Note that it is not really possible with HST to use an “interlace” dither pattern and directly construct a $\times 2$ supersampled image, as described in Lauer 1999. This is because image distortion causes the distorted pixels to creep relative to an undistorted frame, similar to the issue of pixel-phase decoherence.

2.1 ACS Box dithers

One of the first things we noticed when examining the recent box dithers for ACS was that the box dither pattern was not as good as we had expected in terms of pixel-phase coverage. We examined dataset `jdxk02a[1,2,4,6]q` taken in 2019 under program GO-15645. It used a standard ACS box dither, which has POS-TARGS: (0.0000,0.0000), (0.24862,0.08266), (0.12456,0.22920), and (−0.12406,0.14654).

We measured stars in all four exposures and compared the positions for the stars in exposures 2 through 4 against the positions in the first exposure. **Figure 2** shows the full dither in raw detector pixels on the left, and in terms of pixel phase on the right³. The pixel phase coverage is clearly sub-optimal. We were very surprised to find this, since the ACS dither box has been used extensively for many years, in particular for the original HUDF program, which achieved an excellent dither strategy.




Figure 2: ACS Box Dither pattern as of the beginning of this study (early 2022). (Left) The full dither in pixel coordinates; (Right) The dither in pixel phase. The first dither at (0, 0) serves as the reference and is shown in black. The second dither is shown in red, the third in green, and the fourth in blue. The data represented here were taken by ACS/WFC in 2019 with program GO-15645.

The goal of this exploration was to evaluate the quality of dithers in the parallel instrument, so we were surprised to uncover problems with the dithers in a *primary* instrument. However, given the sub-optimal dither seen above, we decided to explore earlier observations to see whether the ACS box dither has always been sub-optimal. To this end, we downloaded the `jb5d01n[n,p,s,w]q` data from GO-11724 taken in 2009 through F606W. **Figure 3** shows the details of the dither achieved. It is clear that this dither provides nearly perfect pixel-phase coverage. It is represented by the following set of POS-TARGs: (0.0000, 0.0000), (0.24794, 0.09354), (0.12374, 0.23334), and (−0.12420, 0.13980). It looks like the specification of the Dither-Box POS-TARG changed sometime after the SM4 repair. We speculate that these slight post-SM4 changes may have been motivated by the revised estimate of the

³ Pixel phase simply corresponds to the fractional part of the dither: $\phi_x = \Delta x - \text{floor}(\Delta x)$. Good pixel-phase coverage for a dither means that the pixel phases are spread out across the face of a pixel. Practically speaking, it means that a given source will land in different parts of the pixel in different exposures. In one dither a star might land at the center of a pixel, while in other dithers it might land on an edge or a corner. This “pixel phase diversity” allows routines such as *astrodrizzle* to reconstruct a higher resolution version of the scene than is possible from a single dither.

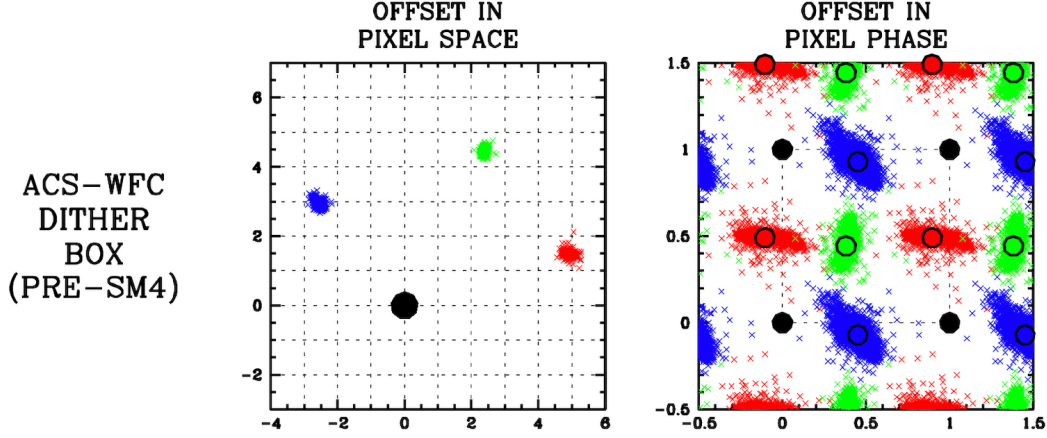


Figure 3: This figure shows the total and sub-pixel dither achieved from the ACS/WFC box dither from observations taken in late 2009, just after SM4. Note that not all stars experience the same dither due to the impact of non-linear distortion (see text).

WFC geometric distortion⁴. Some combination of these factors is the likely explanation for the degradation in WFC sub-pixel phase sampling observed between 2009 and 2019.

Once we discovered the sub-optimal ACS box dithers, we quickly evaluated the POS-TARG-to-pixel transformations and constructed a new set of offsets. The new POS-TARG offsets are: (0.0000,0.0000), (0.24668,0.09321), (0.12326,0.23183), and (−0.12342,0.13862). The new pattern was fast-track implemented by the APT team and is currently being used by observers.

Figure 4 shows this new pattern. The data correspond to images `jflg01t[o,p,r,t]q` from GO-17231 taken in mid 2022.




Figure 4: The full and sub-pixel dithers for the newly improved ACS/WFC box dithers. The individual x's correspond to individual stars.

⁴ The WFC geometric distortion has been shown by Hoffman & Kozhurina-Platais (2020) to be gradually time-varying.

Note that there is considerable spread in (say) the individual blue points about the average. This can be seen in all three figures (2 through 4). In fact, the exposures that went into the generation of [Figure 3](#) had so many reference stars that one can see that the distribution in pixel phase about each dither looks like a parallelogram. If we were to zoom in on the full-dither points in the left plot, they would look the same. The reason for this is that there is considerable non-linear distortion in the ACS/WFC camera that causes a dither at the edge of the detector to be different from a dither at the center of the detector. This is discussed in detail in [Section 3](#).

2.2 WFC3/UVIS Box dithers

The top panels of [Figure 5](#) show the WFC3/UVIS box dither as realized in the F606W images `ib4403e[n,p,s,e]q` of program 11684 taken in 2010. The box dither has the following POS-TARG prescription: (0.00000, 0.00000); (0.15819, 0.07005); (0.09867, 0.16492); (−0.05951, 0.09488).

These four WFC3/UVIS images were taken in parallel with ACS/WFC images `jb4403e[o,q,t,x]q`. The dither achieved in the parallel observations is shown in the bottom panels. It is clear that the pixel phase space is very poorly covered in the parallel observations (bottom right panel). The goal of this ISR is to come up with a pattern of shifts that produces an equivalently good dither in the prime instrument, but a much better dither in the parallel instrument.

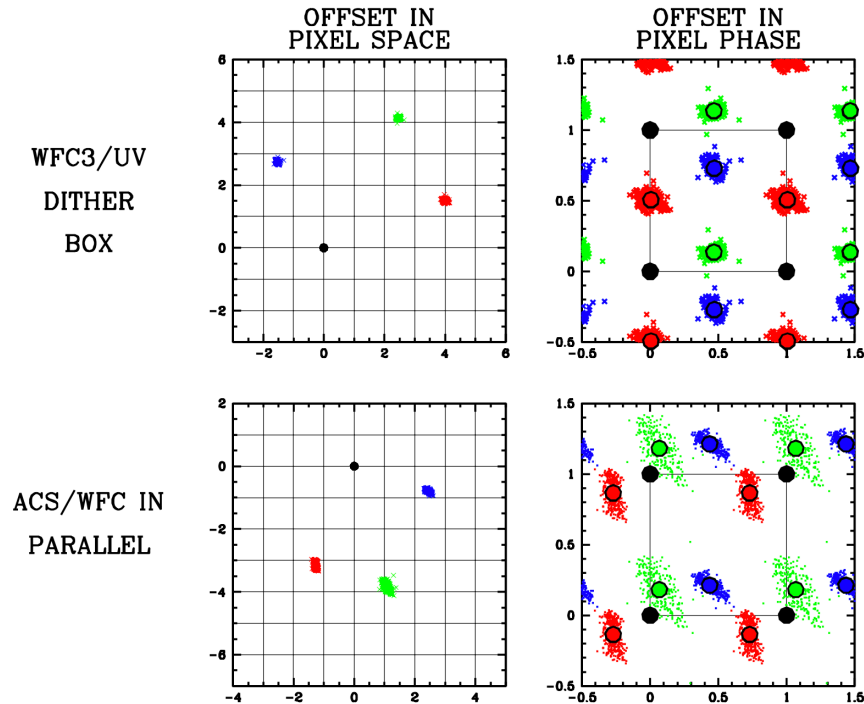


Figure 5: (Top) the actual achieved full-pixel shifts for the WFC3-UVIS-BOX dither on the left and the corresponding pixel-phase shifts on the right. (Bottom) The corresponding shifts observed for the ACS/WFC, observed in parallel.

2.4 WFC3/IR Box dithers

The top panels of [Figure 6](#) show the WFC3/IR box dither pattern for four WFC3/IR images `ib6k03y[d,h,l,o]q`, taken by program GO-11731 using the set of POS-TARGs: : (0.0000, 0.0000), (0.54235,0.18176), (0.33891, 0.48481), (-0.20345, 0.30304). The sub-pixel coverage achieved is almost perfect.

The bottom panels show the pattern achieved in the parallel instrument, ACS/WFC. The pixel-phase coverage in the parallel instrument is clearly sub optimal. Furthermore, even though the dither is quite tight for WFC3/IR (a span of 5.5 pixels in x and 3 pixels in y), the pattern in ACS spans more than 12 pixels in both x and y , which introduces considerable spread in pixel phase for a given dither. Thus, even if the average locations of the red, blue, and green points were better spaced out, the dither would still be increasingly compromised the farther we go from the center of the detector.

For convenience, [Table 1](#) lists all the dithers discussed in this section.

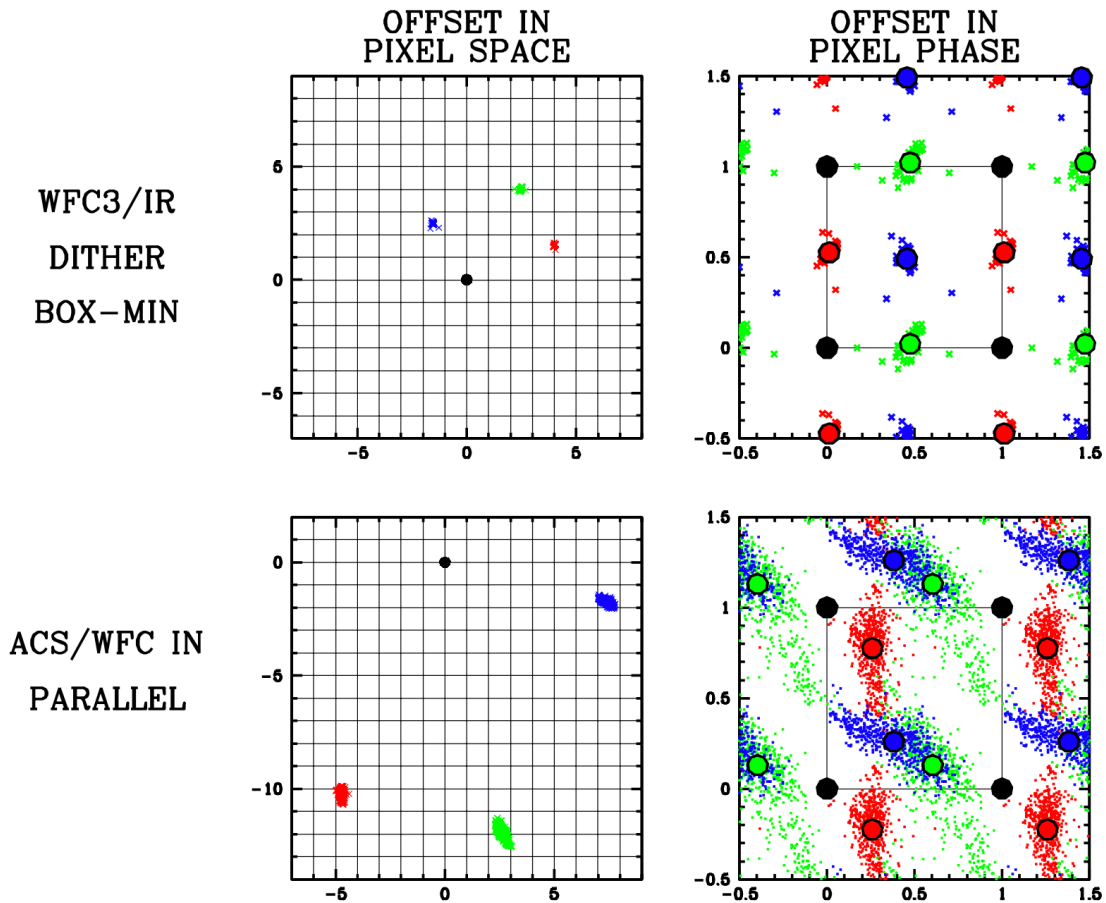


Figure 6: (Top) The whole-pixel and sub-pixel dither achieved in the WFC3/IR dither box pattern. (Bottom) The whole-pixel and sub-pixel dither seen in the ACS/WFC channel when it is observed in parallel.

Table 1: The box dithers in the various instruments. The upper two are no longer valid. The first is for before SM4 and the second was recently found to be erroneous and has been replaced.

Prime Instrument	Dither #1	Dither #2	Dither #3	Dither #4
ACS/WFC Pre-SM4	0, 0	0.24794, 0.09354	0.12374, 0.23334	− 0.12420, 0.13980
ACS/WFC Post-SM4	0, 0	0.24862, 0.08266	0.12456, 0.22920	− 0.12406, 0.14654
ACS/WFC Fixed	0, 0	0.24668, 0.09321	0.12326, 0.23183	− 0.12342, 0.13862
WFC3/UV	0, 0	0.15819, 0.07005	0.09867, 0.16492	− 0.05951, 0.09488
WFC3/IR	0, 0	0.54235, 0.18176	0.33891, 0.48481	− 0.20345, 0.30304

3. Issues to consider when dithering

Dithering accomplishes several beneficial objectives. One major goal of dithering is to mitigate the effect of bad pixels, bad columns and artifacts (such as WFC3/IR blobs) on the dataset as a whole. To mitigate bad pixels, one must typically dither by a few pixels such that the core pixels of a point source (which typically has a FWHM of 2-3 pixels) will not be impacted by the same bad pixel in more than one exposure. In order to mitigate bad columns in CCDs⁵, then the x -spacing between all dithers must be 2 to 3 pixels, irrespective of the y -spacing. To mitigate image artifacts (such as the IR blobs), the dither must be larger than the artifact size. For WFC3/IR blobs, this means ± 15 pixels. Unfortunately, as we will see in the next section, such large steps cannot be made while maintaining pixel-phase coherence. Another issue to consider for WFC3/IR is persistence: if the dithers are less than a few pixels, then the central pixels of bright stars will dither onto pixels that were bright in previous exposures, thus making the correction for persistence less certain.

An additional benefit of dithering for undersampled detectors (such as the three HST detectors that we consider here) is that it can also improve the sampling of the scene. This is done this by placing objects that may have been positioned at the center of a pixel in one exposure at the corner or edge of a pixel in a different exposure. This allows image-stacking routines such as Astrodrizzle⁶ to reconstruct a higher resolution image of the scene than can be realized in a single exposure.

This document is focused on the scene-sampling aspect of dithering. This means that we will select dithers that accomplish the other above benefits, when possible, but in this effort we will sacrifice all other benefits of dithering in order to achieve optimal sampling.

The main dithering goal that is often incompatible with sub-pixel sampling has to do with mitigating artifacts or detector gaps, which often require dithering of tens of pixels (or more). The reason for this incompatibility is that when a detector has significant non-linear geometric distortion, dithers that are large enough to mitigate artifacts can suffer from pixel-phase drift and

⁵ CMOS detectors, such as WFC3/IR, do not have bad-column issues.

⁶ <https://www.stsci.edu/scientific-community/software/drizzlepac>

incoherence. A non-linear distortion means that the local pixel scale (and possibly pixel orientation) changes with location on the detector⁷. In the present context, a non-linear distortion means that if we execute a dither of 10.50 pixels along the x direction in ACS, then a star at the center of the detector will be moved from (say) (2048.00, 2000.00) to (2058.50, 2000.00). However, on account of non-linear distortion, a star on the left edge of the detector will be shifted from (4000.00, 2000.00) to (4010.75, 2000.0) — a dither of 10.75 pixels rather than 10.50 pixels. For this reason, a dither that is meant to provide good pixel-phase coverage at one point of the detector can end up producing bad pixel-phase coverage at another point. The only way to avoid this is to make the dithers smaller.

On account of its off-axis location in the focal plane (see [Figure 1](#)), the ACS/WFC detector has the largest non-linear distortion of all the detectors. [Figure 7](#) shows the pixel-phase decoherence that results from whole-pixel offsets from zero up to 10 pixels in Δx and Δy . In each panel, the dotted square represents dither shifts of 0.25 pixel. Staying within this box is a good rule of thumb for preventing 0.5-pixel-phase offsets from encroaching on one another. From this, we can see that, for ACS/WFC, we really need to keep the total sub-pixel throw below 6 pixels in either Δx or Δy .

[Appendix A](#) shows the same figures for WFC3/UV and WFC3/IR. From them, we see that WFC3/IR can dither up to 8 pixels in either Δx or Δy , and WFC3/UVIS can dither more than 10 pixels without encountering significant pixel-phase creep.

⁷ This is why HST’s detectors must have “pixel area maps” to account for the fact that pixels can subtend different amounts of sky at different locations in the field.

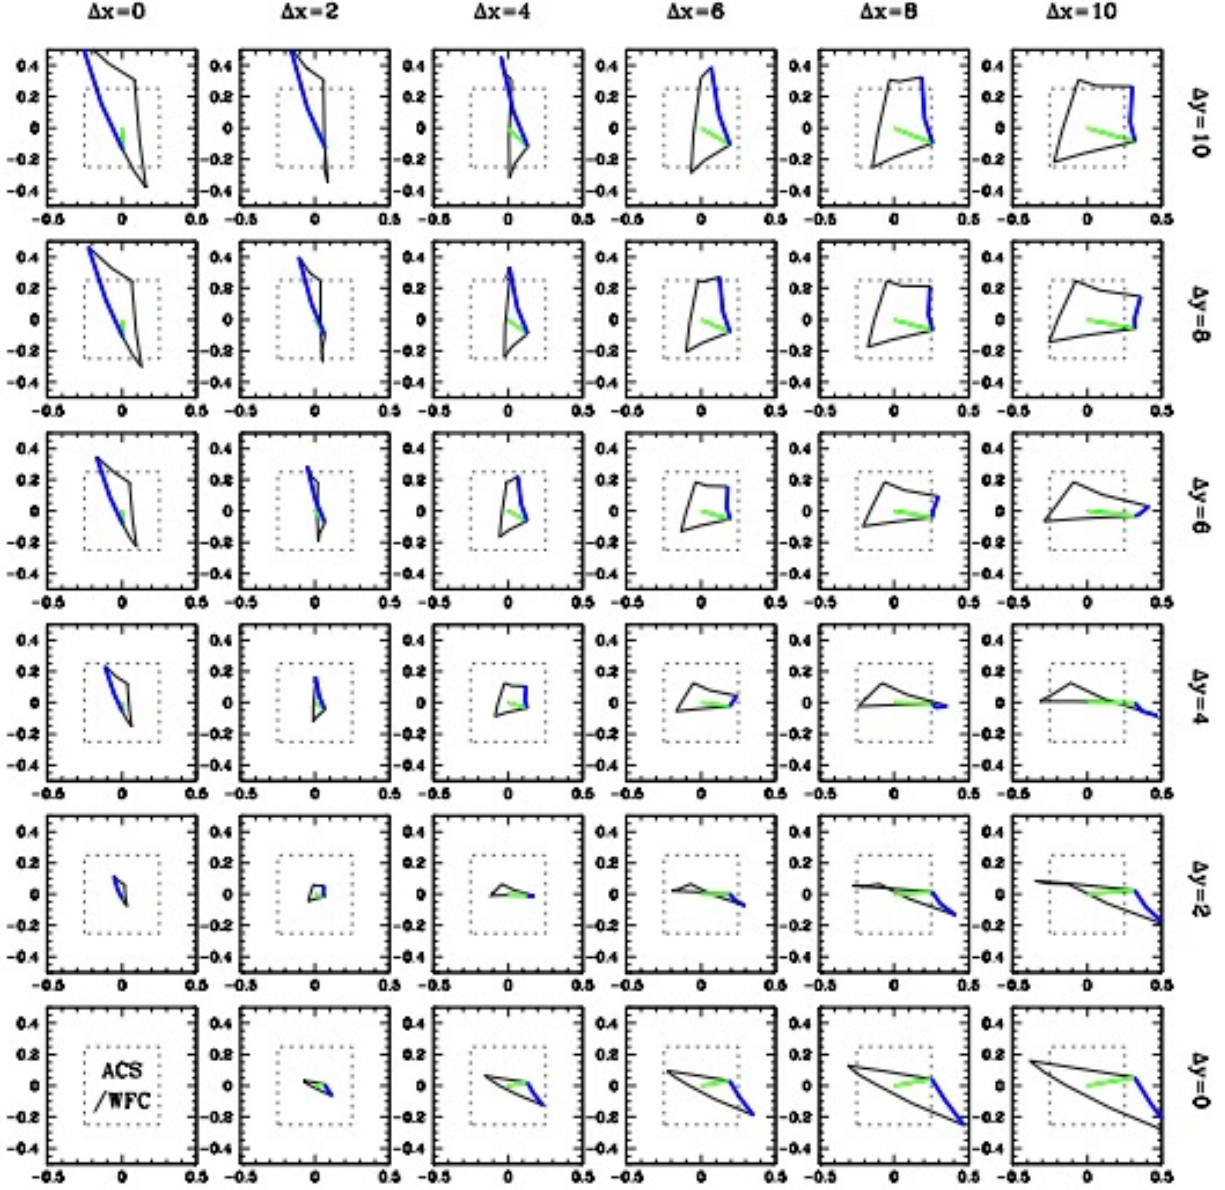


Figure 7: These panels show the dither decoherence achieved at the outer edges of the ACS/WFC detector for an array of whole-integer dithers executed at the center of the detector. Dither decoherence is the difference in achieved dither at a particular location on the detector $[I, J]$ relative to the dither at the center $[2048, 2000]$: $(\Delta\Delta x, \Delta\Delta y) = (\Delta x^{[I, J]} - \Delta x^{[2048, 2000]}, \Delta y^{[I, J]} - \Delta y^{[2048, 2000]})$. Dither decoherence is a result of non-linear distortion in the detector. The panels in the bottom row show the decoherence achieved for dithers made along the x-axis ($\Delta y = 0$). The green line connects the dither decoherence $(\Delta\Delta x, \Delta\Delta y)$ achieved at the bottom left corner of the detector (4096, 0) with that at the detector center (which is zero). The blue line connects the bottom left with the bottom right of the detector, and the black lines follow the other three edges of the detector. The outer boundary drawn in each panel shows the worse possible case of the dither decoherence for a given integer shift $(\Delta x, \Delta y)$. The dotted line shows a decoherence of 0.25 pixel, which is a proxy for a badly decoherent dither.

4. Mapping POS-TARGs to prime and parallel dithers

Tailor-made dithers can be achieved through APT by providing a list of POS-TARG offsets⁸ for the prime instrument. POS-TARGs can be specified either implicitly through a line or box pattern with particular spacing and angles or explicitly by simply specifying the desired shift in Δx and Δy for each exposure relative to the nominal aperture location. POS-TARGs are specified in arcseconds with the POS-TARG y axis aligned closely (though not exactly) with the detector y axis. The POS-TARG x axis is not defined along the detector x axis; rather it is defined as being perpendicular to the POS-TARG y axis. Because the ACS/WFC detector axes are significantly skewed from perpendicular, the ACS/WFC POS-TARG x axis thus has a significant y -detector component.

In order to design good dithers, we must evaluate the exact mapping between the POS-TARG frame and the detector frame, both for the prime instrument and for the parallel instrument. To this end, we examined the entire set of full-frame exposures in the archive taken between 2009 and 2022 with all possible prime/parallel combinations of ACS/WFC and WFC3/IR and WFC/UVIS. Full-frame HST images typically have at least 10 moderately bright stars that can be used to align images and measure exposure-to-exposure shifts empirically. Observation pairs that did not have an adequate number of stars in each of the four exposures were rejected.

We evaluated pairs of prime/parallel exposures taken together within an HST orbit through the same filter⁹ and extracted from each exposure the POS-TARGs (specified in the prime-instrument frame). We can thus compare from each pair of exposures ΔPT_X and ΔPT_Y , from the headers, and Δx_{prime} and Δy_{prime} taken from the difference in shifts between the pair of prime exposures and $\Delta x_{\text{parallel}}$ and $\Delta y_{\text{parallel}}$ taken from the parallel exposures. Recall that [Figure 7](#) and [Appendix A](#) show that distortion can significantly affect the measured shifts for all but the smallest dithers. For this reason, we measured the image-to-image shifts in a distortion-corrected system, namely the distortion corrected system described in the `hst1pass` documentation (Anderson 2022)¹⁰. These distortion corrections were constructed to be as close as possible to the detector frame and, as such, they are the most relevant ones to use for this task.

We used the following procedure to relate the POS TARGs to the offsets in each detector. We took the list of all full-frame ACS/WFC or WFC3 exposures taken in prime and parallel and sifted them into the three directories, one for each camera. We then ran `hst1pass` on each image, specifying the closest PSF and distortion solution. We had the routine output the `uvw`

⁸ POS TARG offsets are the way to manually command a shift of the image aperture relative to the nominal commanded center.

⁹ The reason that our comparison observations for the same instruments must be through the same filter is that a filter can introduce a zeropoint shift in the astrometry, and we do not want that to complicate our ability to associate shifts with commanded POS-TARGs. The prime observations must all be taken with the same filter, and the parallel observations have to all be taken with the same filter, but the prime and parallel observations can use different filters.

¹⁰ Note that there is no single distortion corrected frame for a given detector. A distortion-corrected frame can be offset, rotated, or uniformly rescaled and still be free of distortion. We are therefore at liberty to choose a frame that is convenient for our aims.

coordinates, which correspond to positions in the distortion-corrected frame. We then listed all pairs of exposures taken in the same orbit with the same nominal target, along with the POS-TARGs listed in the header. The first several lines of the ACS/WFC prime and WFC3/UVIS parallel file appear as the following:

#	WFC3/UVIS PRIME		ACS/WFC PARALLEL			POS-TARGs	
#	-----		-----			-----	
#FILTER	IMAGE#1	IMAGE#2	FILTER	IMAGE#1	IMAGE#2	ΔPT_X	ΔPT_Y
F435W	ib6w61uaq	ib6w61uzq	F435W	jb6w61ubq	jb6w61v0q	1.446	2.926
F435W	ib6w61uzq	ib6w61uaq	F435W	jb6w61v0q	jb6w61ubq	-1.446	-2.926
F814W	ib6wd41lq	ib6wd41jq	F814W	jb6wd4kzq	jb6wd4lhq	0.300	1.700
F814W	ib6wd41jq	ib6wd41lq	F814W	jb6wd4lhq	jb6wd4kzq	-0.300	-1.700
F814W	ib6wd2bzq	ib6wd2cpq	F814W	jb6wd2bxq	jb6wd2cnq	0.300	1.700
F814W	ib6wd2cpq	ib6wd2bzq	F814W	jb6wd2cnq	jb6wd2bxq	-0.300	-1.700

There were over five thousand pairs of exposures taken in the same orbit through the same filter with WFC3/UVIS prime and ACS/WFC in parallel.

For each pair, we next extracted the four .uvw files for the four full-frame images and determined Δu_{prime} and Δv_{prime} and $\Delta u_{\text{parallel}}$ and $\Delta v_{\text{parallel}}$ for each pair in the file above by cross-identifying stars and comparing the positions in the first and second files for each detector. We rejected any case for which there were not enough moderately bright stars to enable the empirical shift to be measured accurately.

This procedure then allowed us to associate a shift in each of the detectors with ΔPT_X and ΔPT_Y , the prime-instrument-based POS-TARG difference between the frames. Finally, we used least squares to fit these two linear systems for the A, B, C, and D linear terms:

$$\begin{pmatrix} \Delta u_{\text{prime}} \\ \Delta v_{\text{prime}} \end{pmatrix} = \begin{bmatrix} A_{\text{prime}} & B_{\text{prime}} \\ C_{\text{prime}} & D_{\text{prime}} \end{bmatrix} \begin{pmatrix} \Delta PT_x \\ \Delta PT_y \end{pmatrix}$$

and

$$\begin{pmatrix} \Delta u_{\text{par}} \\ \Delta v_{\text{par}} \end{pmatrix} = \begin{bmatrix} A_{\text{par}} & B_{\text{par}} \\ C_{\text{par}} & D_{\text{par}} \end{bmatrix} \begin{pmatrix} \Delta PT_x \\ \Delta PT_y \end{pmatrix}$$

In general, not all dithers are executed perfectly, so we iteratively rejected PT/dither pairs that were inconsistent with the trend seen in the others until we arrived at a normal distribution.

The final parameters for these linear fits are given in [Table 2](#). Note that the pixel scales (column 6) for the different instruments are quite consistent from comparison to comparison. The orientation for the prime instrument (column 7) is generally very close to zero. This is by design since the distortion solution was constructed to have its y axis aligned with the detector y axis, and the POS-TARG system was similarly designed to align its y axis with that of the detector.

Table 2: The linear relationship between POS-TARGs in the prime instrument and distortion-corrected pixel offsets in the prime and parallel instruments.

INSTRUMENT	A	B	C	D	SCL (pix/")	θ (°)
ACS/WFC PRIME	20.11280	0.07812	-0.06475	20.12717	20.120	0.203
WFC3/IR PARALLEL	-5.61565	-6.08584	6.06511	-6.61779	8.274	-132.753
ACS/WFC PRIME	20.13796	0.08805	-0.09412	20.08902	20.114	0.259
WFC3/UV PARALLEL	-17.10117	-18.43432	18.47071	-17.11431	25.163	-132.834
WFC3/UV PRIME	25.12966	-0.01044	0.01057	25.14709	25.138	-0.024
ACS/WFC PARALLEL	-13.72998	14.75823	-14.58244	-13.74834	20.136	133.026
WFC3/IR PRIME	8.26621	0.02781	-0.02821	8.26823	8.267	0.194
ACS/WFC PARALLEL	-13.73303	14.71498	-14.69005	-13.78217	20.135	133.098

We can now combine the two equations to estimate the output parallel dithers in terms of the input prime dithers:

$$\begin{pmatrix} \Delta u_{par} \\ \Delta v_{par} \end{pmatrix} = \begin{bmatrix} A_{par} & B_{par} \\ C_{par} & D_{par} \end{bmatrix} \begin{bmatrix} A_{prime} & B_{prime} \\ C_{prime} & D_{prime} \end{bmatrix}^{-1} \begin{pmatrix} \Delta u_{prime} \\ \Delta v_{prime} \end{pmatrix}.$$

It is worth remembering that these dithers are in distortion-corrected coordinates, so both the prime shifts and the parallel shifts will need to be run through the inverse distortion solutions to get actual shifts on the pixel grid. This is clearly a non-trivial operation, but we do have all the tools in hand to compute the pixel-grid shifts.

5. How to optimize the parallel dithers?

Now that we can determine exactly how HST will dither the parallel instrument for a given dither in the prime instrument, it is time to set up the problem for optimization. We already have optimal prime dithers for ACS/WFC, WFC3/UV, and WFC3/IR. However, [Section 2](#) showed us that, at least for the case of ACS/WFC in parallel with the WFC3 instruments, the dither achieved in the parallel instrument tends to have quite bad pixel-phase coverage. This should not be surprising, since until now the impact on the parallel instrument has not been assessed.

5.1 The art of the possible

It is not, in general, possible to achieve a perfect dither in both the prime and the parallel instruments. Therefore, our strategy here will be to start with a perfect sub-pixel dither in the prime instrument and then find the whole-pixel offsets in the prime detector that produce the best

dither in the parallel instrument. For a given set of sub-pixel offsets in the prime instrument, we have the flexibility of choosing whole-pixel offsets to meet other needs without affecting the sub-pixel dither, provided the whole-pixel offsets are small enough to prevent distortion spreading in pixel phase (see the discussion in [Section 3](#)). These whole-pixel offsets are usually chosen to accomplish other goals, but here we will make use of the flexibility of the whole-pixel offsets to optimize the sub-pixel dither in the parallel instrument. The whole-pixel offsets give us a six-dimensional parameter space that we can explore to determine which combination of $(\Delta i_2, \Delta j_2)$, $(\Delta i_3, \Delta j_3)$, and $(\Delta i_4, \Delta j_4)$ provides the best sub-pixel dither in the parallel instrument.

While we mentioned above that we will prioritize pixel-phase above all else, we will also impose a minimum offset of 2.0 pixels between dithers and will also insist that for the CCDs (WFC3/UVIS and ACS/WFC), no dithers will be within 1.25 pixels of each other in x coordinate, so as to minimize the impact of bad columns.

5.2 Dither-Quality Metrics

How, then, do we evaluate the quality of sub-pixel dithers? The “conventional wisdom” interlace dither pattern discussed in [Section 2](#) is shown in the left panel of [Figure 8](#). There are some other dithers that are nearly this good. For instance, the two dithers shown in the middle and right panels have similarly good spacing and provide similarly good image reconstructions.




Figure 8: Three good fiducial dither patterns. The open points show the wrap-around nature of dithers.

As we try to design “optimal” dithers in the parallel instrument, it is worth asking the question: what is the most important aspect of a good dither in terms of pixel phase? If we were to come up with a trial pattern by taking one of the above sub-pixel pattern and adding whole-pixel dithers for the prime dither and evaluating the corresponding dithers and pixel phases in the parallel instrument, how can we say how good the parallel dither is or whether one dither is better than another?

We can imagine two different metrics for evaluating sub-pixel quality. The first metric has to do with how far apart the dithers are in pixel phase space, the second has to do with how well they cover the face of the unit pixel. We construct the first metric M_1 by determining for each pair of

dithers the minimum distance between the two dithers (being aware that the unit pixel wraps around):

$$M_1 = \sum_{ij} \sqrt{(\|x_i - x_j\|^2 + \|y_i - y_j\|^2)},$$

where the “ $\| \cdot \|$ ” operator here represents the minimum wrap-around distance between points i and j on the unit square. For the “perfect interlace” dither above, the number is 6.83 (each of the four points has neighbors with distances of $1/2$, $1/2$, and $\sqrt{2}/2$). The numbers for the alternative patterns in panels (b) and (c) is: $6.47 = 4 \times (1/2 + \sqrt{5}/4 + \sqrt{5}/4)$, which is lower, and therefore *less* good in this metric. Note that if all four dithers are on top of each other, this number is 0.

The second metric has to do with the average distance squared from the typical point on the unit square to the closest dither point:

$$M_2 = \int_{-0.5}^{+0.5} \int_{-0.5}^{+0.5} \min_i (\|x - x_i\|^2 + \|y - y_i\|^2) dx dy.$$

This involves integrating over the two dimensions of the unit square the distance squared from each point (x,y) to the nearest dither (x_i, y_i) . For the interlace dither, this integral is: 0.191. If all four dithers are on top of each other, this number is 0.382. This worst possible case is a factor of two worse than the interlaced case. It turns out that if we take one of the alternative “good” dither patterns from the middle or right panel, we get 0.189, a slightly better result. The reason for this is that the maximum distance from a point in the square to the closest dither point in the “interlace” pattern is larger than that in the alternate patterns (see [Figure 9](#)). Since image reconstruction is more constrained by the location of scene-samplings than the distance between them, it could be that the second metric is more relevant to the goal of reconstructing a scene. However, in a practical sense, the two metrics track each other quite well, and we will evaluate them both for all patterns.




Figure 9: Contour plot showing distance from a point on the unit square to the closest dither point (shown in blue/green). The left plot corresponds to case (a) in Figure 8, and the right plot to case (b). The red contours show the maximum distances and the left plot has more deep red (high distance) contours.

6. Optimizing the parallel dithers

Summarizing all of the above considerations, our strategy will be to determine from the limits in [Section 3](#) how far a dither pattern can extend in native pixels for each instrument so as to avoid excessive pixel-phase creep, and then adopt that as a hard constraint on the maximum size of the dither pattern.

For ACS, we cannot dither more than six 50-mas pixels (0.3 arcsec) before the pattern loses coherence by the edge of the detector. For WFC3/IR, the pattern can extend up to eight 110-mas pixels (0.88 arcsec) before losing coherence. For WFC3/UVIS, the pattern can extend more than ten 40-mas pixels (0.4 arcsec). In all cases, ACS provides the tightest constraints of 0.3 arcsec, so this maximum throw restricts the 4-point dithers we consider here.

Our strategy is as follows. We constrain dither #1 to have a POS-TARG of (0, 0). We then have a 7-dimensional integer parameter space to consider $(\Delta i_2, \Delta j_2, \Delta i_3, \Delta j_3, \Delta i_4, \Delta j_4, F)$, where the shifts are in the primary-instrument pixel space and where F goes from 1 to 3, corresponding to the three fiducial cases shown in [Figure 8](#).

For each possibility in this 7-dimensional space, we will construct the set of dither shifts in the prime instrument ($\Delta x_2 = \Delta i_2 + \phi_{x,2}^{[F]}$, etc). We evaluate these offsets at the center of the detector. For WFC3/UVIS and ACS/WFC, we use $(x_1, y_1) = (2048, 2000)$ and for WFC3/IR we use (507, 507). We next convert these raw pixel coordinates (x_i, y_i) into distortion-corrected positions using the STDGDC files provided along with `hst1pass` (Anderson 2022). We will refer to these distortion-corrected coordinates as: $(u_1, v_1, u_2, v_2, \text{etc.})$. Next, we use the linear relationships described in [Section 4](#) (distilled into [Table 2](#)) to convert these offsets (i.e., $\Delta u_{2-1} = u_2 - u_1$, etc.) into POS TARGs in the frame of the primary instrument. With these POS-TARGs, we can use the second row in the same tables to determine the corresponding offsets in the distortion-corrected parallel frame ($\Delta u_{2-1}^{\parallel}, \Delta v_{2-1}^{\parallel}, \text{etc.}$). Finally, we can convert these offsets into offsets in the actual distorted pixel frame: $(\Delta x_{2-1}^{\parallel}, \Delta y_{2-1}^{\parallel}, \text{etc.})$. This long procedure results in eight (!) steps:

$$\Delta x \rightarrow x \rightarrow u \rightarrow \Delta u \rightarrow \text{PT} \rightarrow \Delta u^{\parallel} \rightarrow u^{\parallel} \rightarrow x^{\parallel} \rightarrow \Delta x^{\parallel}.$$

These steps allow us to go from pixel offsets in the prime detector frame to pixel offsets in the parallel detector frame.

Now that this procedure has been laid out, we can now execute the original plan. We convert each set of seven parameters $(\Delta i_2^{\parallel}, \Delta j_2^{\parallel}, \Delta i_3^{\parallel}, \Delta j_3^{\parallel}, \Delta i_4^{\parallel}, \Delta j_4^{\parallel}, F)$ into a set of three offsets in the parallel frame $(\Delta x_{2-1}^{\parallel}, \Delta y_{2-1}^{\parallel})$ $(\Delta x_{3-1}^{\parallel}, \Delta y_{3-1}^{\parallel})$, and $(\Delta x_{4-1}^{\parallel}, \Delta y_{4-1}^{\parallel})$. These offsets can then be trivially converted into three pixel phases. Before we evaluate these pixel phases, though, we impose the constraints that the total span of the offsets in POS TARG space cannot be more than 0.3 arcsec. In addition to this, we required that no dithers be within 2 pixels of another dither, to provide some modicum of protection from a single bad pixel affecting the same source twice, and finally we insist on at least a 1.25 pixel spacing in x for the CCD detectors in order to mitigate the impact of bad columns.

At this point, we have a list of possible dithers in the primary frame and the consequent dithers in the parallel frame. We determine metrics M_1 and M_2 from [Section 5.2](#) for each set of parallel-

frame pixel phases and determine which of the dithers provides the optimal fit for each of the four prime+parallel instrument combinations. We go through these four combinations one by one in the following subsections.

6.1 ACS/WFC Prime, evaluating the revised box dithers in APT

In our discussion in [Section 2](#) of the ACS box dithers available in APT, we focused on the need to improve them, but we did not have any WFC3/UVIS or WFC3/IR images taken in parallel to evaluate the quality of the parallel dither for the new pattern in APT. With the linear transformations and dither-evaluation tools we have developed here, we can now evaluate the implied parallel dithers.

Figure 10 shows the new ACS pattern along with the dither achieved in parallel for WFC3/IR on the left and WFC3/UVIS on the right. The resulting dither in WFC3/IR in parallel provides quite bad pixel-phase coverage, but the dither in WFC3/UVIS turns out to be satisfactory. If all the parallel dithers had turned out to be as good as this ACS-vs-UVIS case, we probably would not have conducted this detailed investigation.

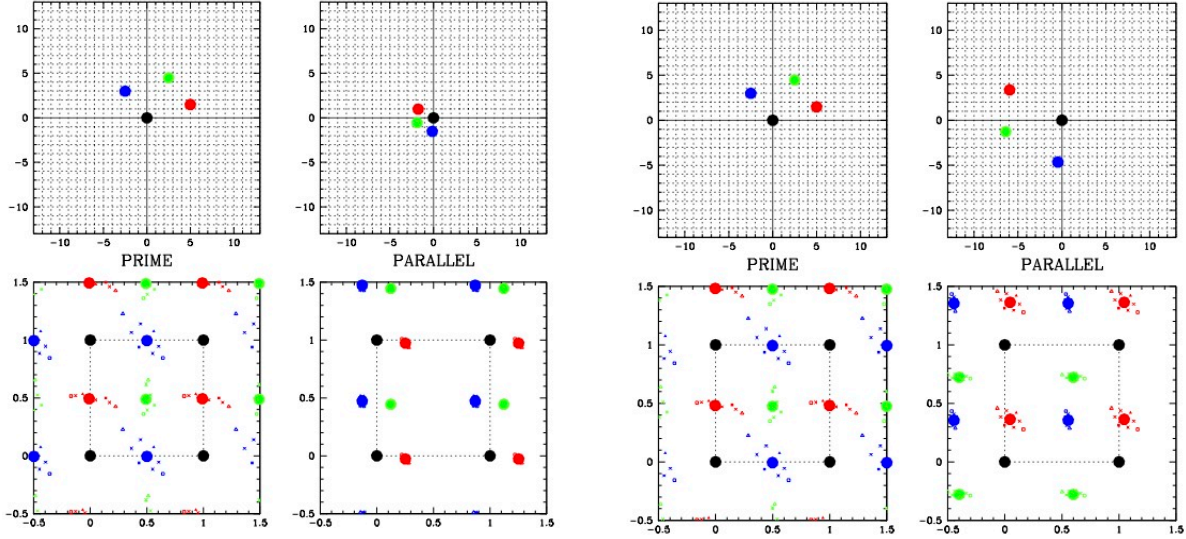


Figure 10: (Left) The new APT pattern for ACS and the consequent dither in WFC3/IR when it is in parallel. (Right) The same pattern for ACS, along with the consequent dither in WFC3/UVIS. The filled circle of each color corresponds to the dither at the center of the detector. The small points surrounding this correspond to the dither achieved at the corners/edges of the detector.

6.2 ACS/WFC Prime, WFC3/IR parallel

The uneven coverage in the lower-right plot in the left half of [Figure 10](#) shows that we really do need to devise a better box dither pattern for situations where pixel-phase sampling is important and ACS/WFC is prime and WFC3/IR is in parallel. We saw in [Figure 7](#) that we cannot dither more than 0.3 arcseconds without the sub-pixel pattern in ACS/WFC losing coherence across the detector. Unfortunately, this “maximal” dither in ACS/WFC corresponds to less than 3 pixels in WFC3/IR — hardly enough to dither over bad pixels. Clearly some compromises are required: it is not possible to achieve good pixel phase coverage and the other aspects of dithering that we are usually focused on. In this case, pixel-phase coverage can be achieved only by sacrificing other dithering goals. See the [Section 7](#) for more on ways to deal with this compromise.

We evaluated the sub-pixel dither in the parallel instrument for all of the possible integer-pixel offsets consistent with the 6-pixel limit for ACS. We then ordered them according to their M1 metric. The best 16 of these patterns are shown in [Figure 11](#), with the best dither shown in the upper left. [Table 3](#) provides the pixel shifts and the POS-TARGs for this dither and [Figure 12](#) shows the total and phase shifts for the prime and parallel instruments.

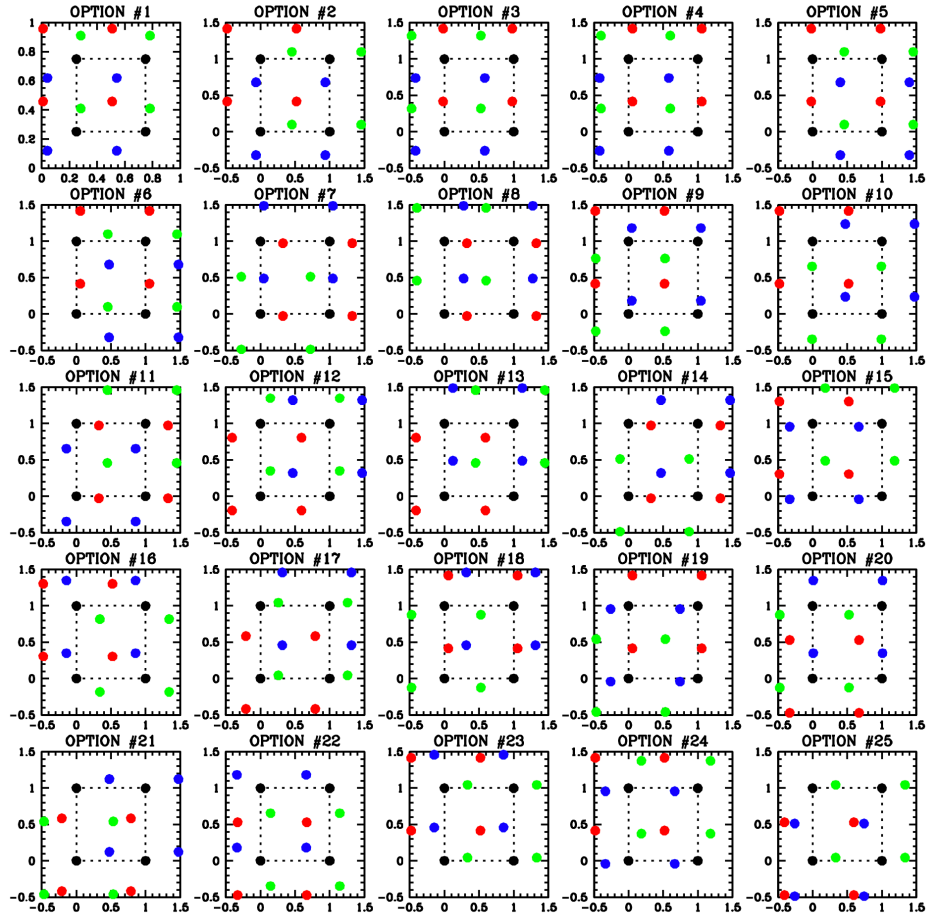


Figure 11: This shows the top 16 options for a pixel-phase pattern in WFC3/IR when ACS/WFC is prime. The two metrics for each are shown above each plot. They are ordered with the best in the upper left. The pattern in WFC/ACS is perfect. The four different colors show the four different dithers in the pattern in the order black (nominal), red, green, then blue.

Table 3: The set of POS-TARGs that provides a perfect dither for ACS/WFC (prime) and the best possible dither for WFC3/IR (in parallel).

DITHER #	POS TARG		DITHER IN ACS/WFC		DITHER IN WFC3/IR	
	PTX (")	PTY (")	Δx (pix)	Δy (pix)	Δx (pix)	Δy (pix)
1	0.00000	0.00000	0.000	0.000	0.000	0.000
2	0.17302	0.11255	3.500	2.000	-1.481	0.417
3	0.11114	-0.11507	2.250	-2.500	0.064	1.320
4	0.28420	-0.00253	5.750	-0.500	-1.417	1.738

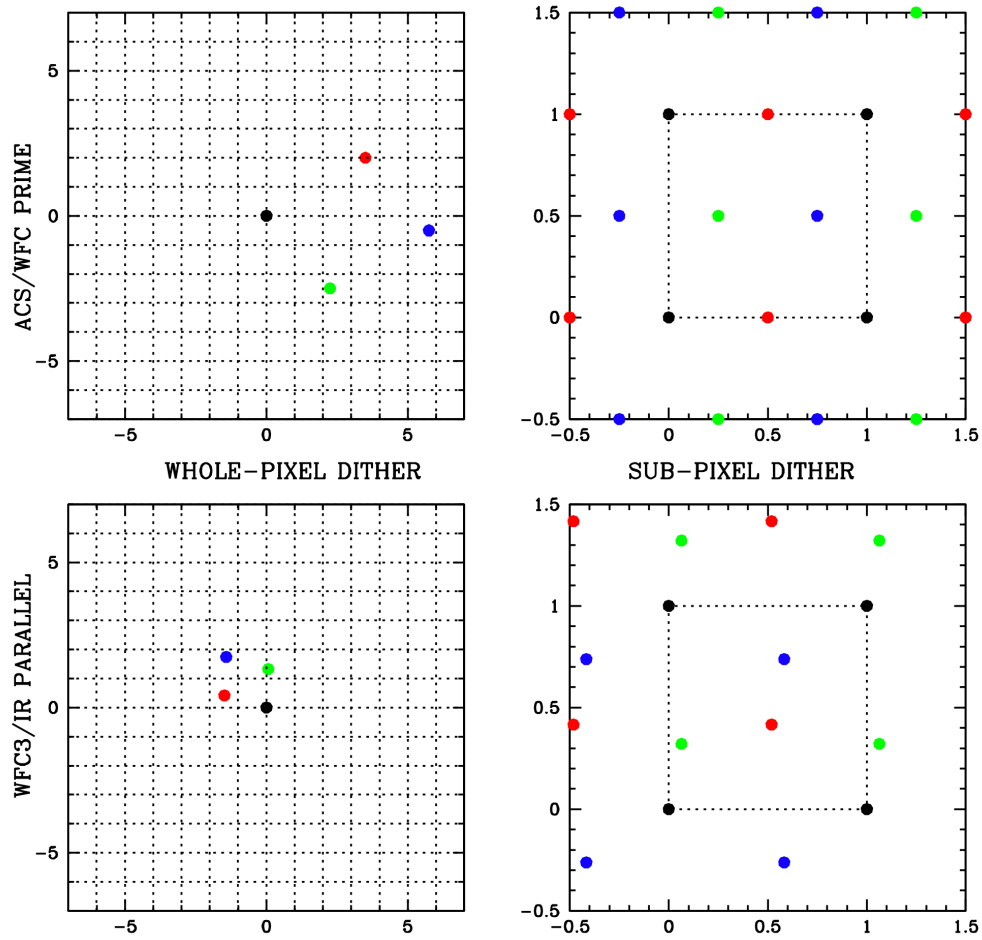


Figure 12: This shows the optimized dither for ACS/WFC prime and WFC3/IR parallel. See the caption of Figure 10 for a detailed description.

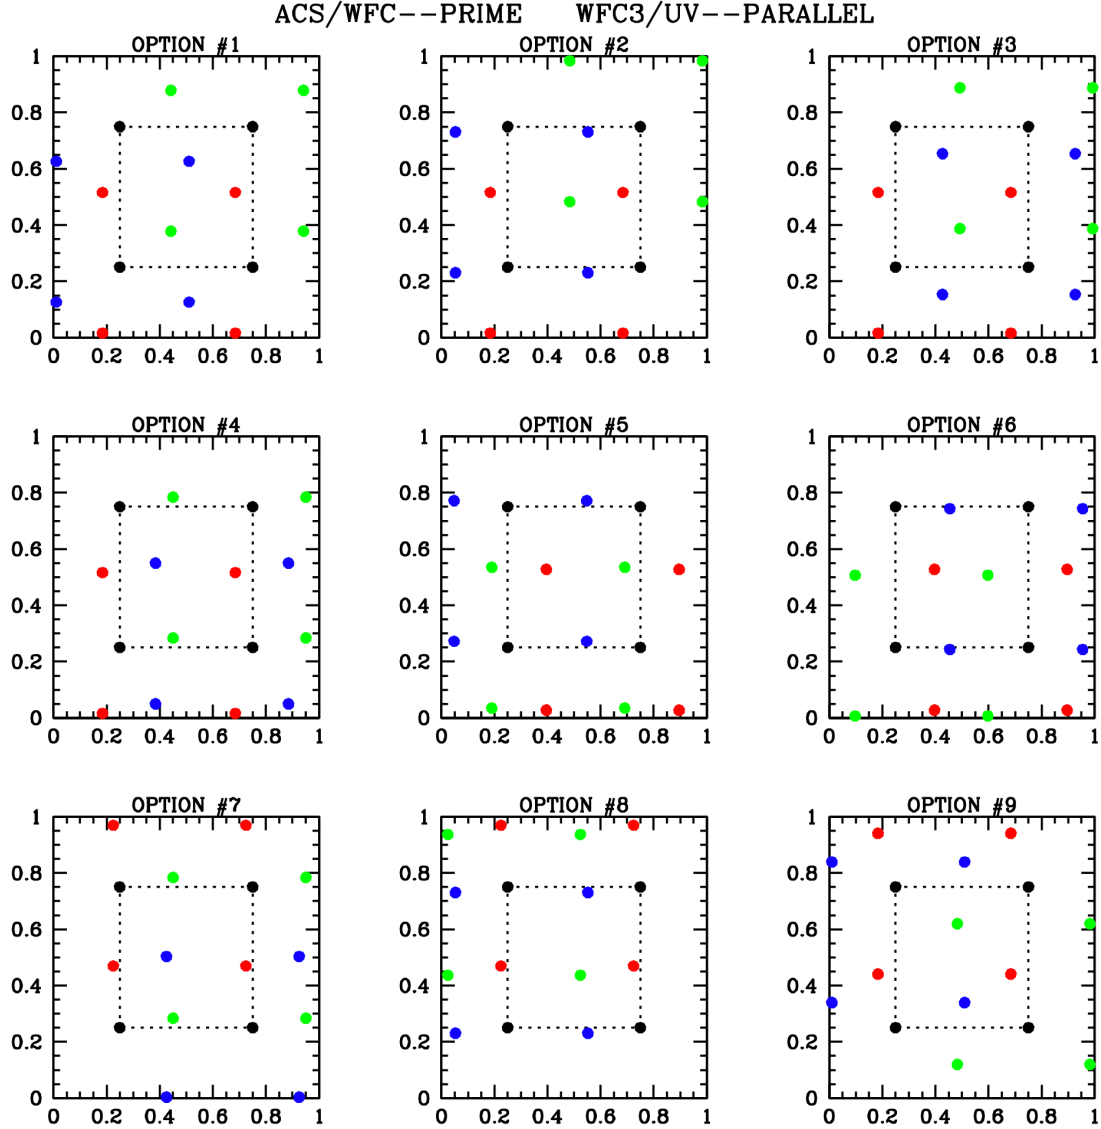


Figure 13: This shows the top 9 options for a pixel-phase pattern in WFC3/UVIS when ACS/WFC is prime. The two metrics for each are shown above each plot. They are ordered with the best in the upper left. The pattern in WFC/ACS is perfect. The four different colors show the four different dithers in the pattern.

6.3 ACS/WFC Prime, WFC3/UVIS parallel

From the left plot in Figure 10, we saw that the new ACS/WFC box dither pattern provides a reasonably good dither in WFC3/UVIS in parallel, but since we have the optimization machinery set up, we may as well try to construct one that might be even better. Figure 13 above shows the top 9 options and Table 3 gives the winning dither in terms of POS-TARGs and pixel offset. Figure 14 shows the achieved full-pixel and pixel-phase dithers for prime and parallel.

Table 4: The set of POS-TARGs that provides a perfect dither for ACS/WFC (prime) and the best possible dither for WFC3/UVIS (in parallel).

DITHER #	POS TARG		DITHER IN ACS/WFC		DITHER IN WFC3/UVIS	
	PTX (")	PTY (")	Δx (pix)	Δy (pix)	Δx (pix)	Δy (pix)
1	0.00000	0.00000	0.00	0.00	0.00	0.00
2	0.14807	0.08624	3.00	1.50	-4.131	1.533
3	0.22215	-0.11866	-4.50	-2.75	-1.616	6.258
4	0.07405	-0.20489	1.50	-4.25	2.521	4.754

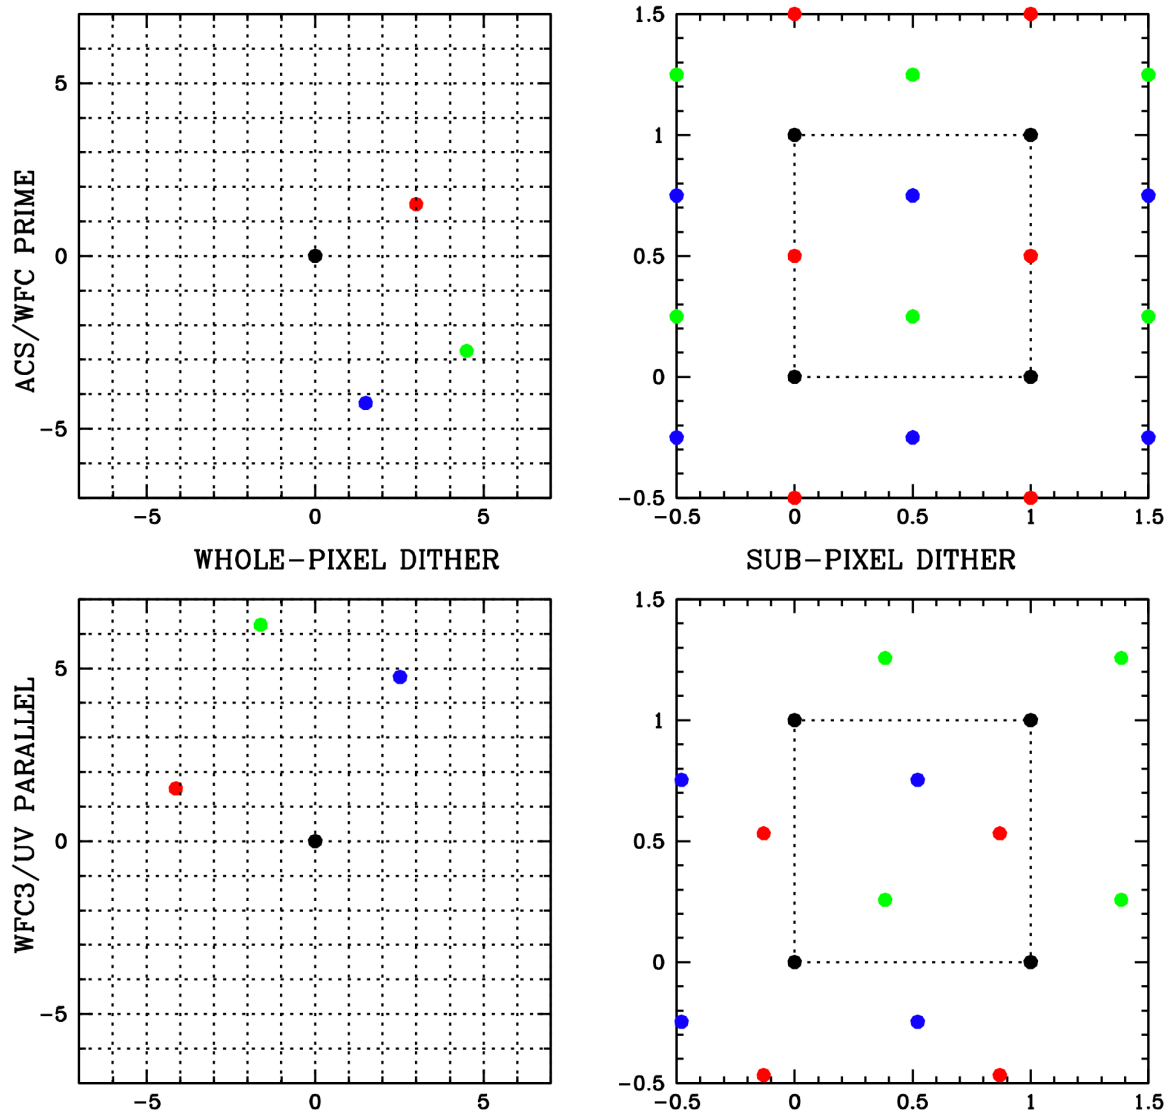


Figure 14: This shows all aspects of the prime+parallel dither for the new ACSWFC/WFC3UV solution: the actual offset in real pixels and the sub-pixel offsets for the four dithers (black, red, green, blue).

6.4 WFC3/IR in prime and ACS/WFC in parallel

The WFC3/IR prime with ACS/WFC in parallel situation is analogous to that in [Section 6.2](#): a dither span of much more than 2 pixels in WFC3/IR will end up dithering more than 6 pixels in ACS/WFC, which will in turn result in significant pixel-phase decoherence at the edges of the field. [Figure 6](#) showed that the current small-box dither in APT for WFC3/IR creates a very poor dither in ACS.

We went through the same parameter-space search as in the previous sections and identified a dither that works perfectly for WFC3/IR and works well for ACS/WFC. The specifications of this dither are given in [Table 4](#) and shown in [Figure 15](#).

Table 5: The set of POS-TARGs that provides a perfect dither for WFC3/IR (prime) and the best possible dither for ACS/WFC (in parallel).

DITHER #	POS TARG		DITHER IN WFC3/IR		DITHER IN ACS/WFC	
	PTX (")	PTY (")	Δx (pix)	Δy (pix)	Δx (pix)	Δy (pix)
1	0.00000	0.00000	0.00	0.00	0.00	0.00
2	0.06773	0.24217	0.50	2.00	2.668	-4.547
3	0.27082	0.06142	2.00	0.50	2.812	-4.631
4	0.20309	-0.18073	1.50	-1.50	-5.480	-0.084

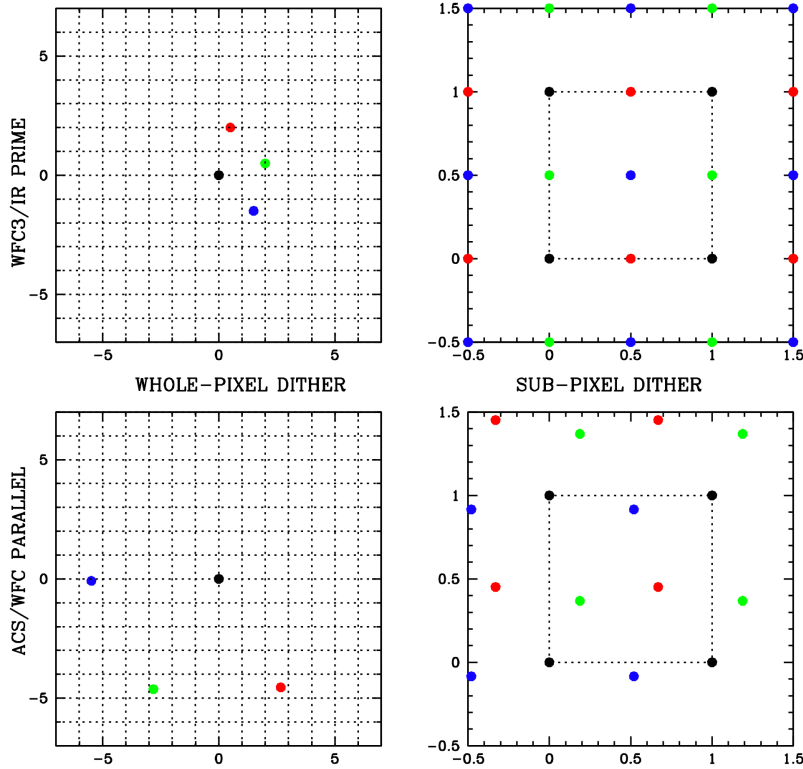


Figure 15: The full-pixel and sub-pixel dithers for the pattern WFC3/IR prime and ACS/WFC parallel pattern shown in [Table 4](#).

6.6 WFC3/UVIS in prime and ACS/WFC in parallel

Figure 5 showed that the standard WFC3/UVIS dither box produces quite an unsatisfactory dither in ACS/WFC when it is in parallel. We used the same procedure discussed above to optimize the whole-pixel offsets in combination with optimal sub-pixel offsets to construct a dither that is perfect for the prime instrument and as good as possible for the parallel instrument. **Table 6** provides the POS-TARGs and whole-pixel offsets, and **Figure 16** shows them graphically.

Table 6: The set of POS-TARGs that provide a perfect dither for WFC3/UVIS (prime) and the best possible dither for ACS/WFC (in parallel).

DITHER #	POS TARG		DITHER IN WFC3/UVIS		DITHER IN ACS/WFC	
	PTX (")	PTY (")	Δx (pix)	Δy (pix)	Δx (pix)	Δy (pix)
1	0.00000	0.00000	0.00	0.00	0.00	0.00
2	0.15868	-0.05024	4.00	-1.50	-2.931	-1.425
3	-0.05961	-0.15241	-1.50	-3.75	-1.452	-3.090
4	0.09907	-0.20266	2.50	-5.25	-4.384	1.665

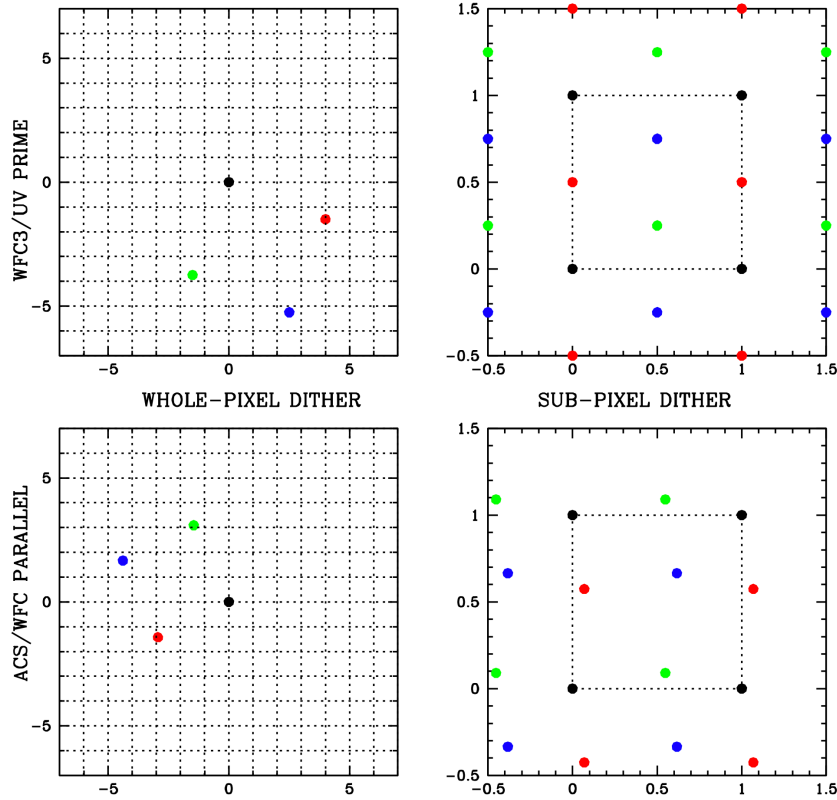


Figure 16: The full-pixel and sub-pixel dithers for the pattern WFC3/UVIS prime and ACS/WFC parallel pattern shown in Table 6.

7. Conclusions

In this document we have shown that the box dithers that were devised to provide good pixel-phase sampling in the ACS/WFC and WFC3/UVIS and WFC3/IR instruments when they are prime do not in general provide good pixel-phase coverage in complementary instruments that are observing in parallel. Since parallel observations are quite common and good pixel-phase coverage in both prime and parallel instruments is often desirable, we have constructed a set of POS-TARGs for the four combinations of prime/parallel cameras that provide optimal pixel-phase coverage in the prime instrument and the best possible coverage in the parallel instrument.

The optimized dithers are provided in [Tables 3, 4, 5](#), and [6](#) of [Section 6](#). These dither patterns are not yet available as selectable options in APT, so interested users must input the POS-TARG specifications directly into the “special requirements” tab for each individual exposure.

It is worth reiterating some of the discussion from [Section 3](#). Pixel-phase coverage is only one of the goals of dithering. Other common goals are: bad-pixel mitigation, bad-column mitigation, artifact mitigation, and the averaging-down of flat-field errors. It is not possible to achieve all of these goals for all detectors simultaneously. The ACS/WFC and WFC3/IR suffer from significant non-linear distortion, which introduces decoherence in pixel phases at the edges of the detectors when dithers are larger than about 6 pixels. This means that a 10-pixel dither set that has good pixel-phase coverage at the center of a detector will end up having bad coverage at the edge of the field. Our patterns were designed to be as compact as possible in order to mitigate this issue.

Most observations taken with parallels tend to be full frame, since the goal is to maximize the field covered. Many of these observations (such as the Hubble Ultra Deep Field) take many sets of four-point dithers in the course of getting many orbits of depth for the field. It is possible to take multiple sets of the four-point dithers recommended here, and even though a single set might not do a complete mitigation of bad pixels or artifacts, multiple sets of four-point dithers (spaced by POS-TARGs of $\sim 1''$) can accomplish *all* the goals of dithering: achieving even pixel-phase coverage as well as bad pixel/column/blob/artifact mitigation.

It is worth noting that it can take quite a large number of “random-pixel-phase” dithers to provide good pixel-phase coverage. The strategy of having multiple sets of 4 coherent points can ensure even coverage of all pixel phases, particularly as it is rarely necessary (or beneficial) to increase sampling by more than $\times 2$ in ACS or WFC3 cameras.

This study has encouraged us to investigate how sensitive Drizzle reconstructions are to the dither pattern. In our early explorations, we have found that with a good dither, Drizzle can indeed improve the resolution of the image, but with a bad dither, the drizzle product suffers clear sampling irregularities. This analysis is in-progress and will be published as a separate ISR in the near future.

Acknowledgements

The authors gratefully acknowledge Joel Green and Karla Peterson for useful discussions, and for assistance with implementation of these dither patterns into the STScI Astronomer's Proposal Tool (APT).

References

ACS Instrument Handbook: <https://hst-docs.stsci.edu/acsihb>

Anderson, J. & King, I. R. 2000 PASP 112 1360. *Toward High-Precision Astrometry with WFPC2. I. Deriving an Accurate Point-Spread Function*

Anderson, J. ACS ISR 2022-02. *One-Pass HST Photometry with `hst1pass`*

Hoffmann, S. L. & Kozhurina-Platais, V. ACS ISR 2020-09: *Validation of New ACS/WFC Geometric Distortion Reference Files Derived Using Gaia Data Release 2*

Lauer, T. 1999 PASP 111 227: *Combining Undersampled, Dithered Images*

APPENDIX A

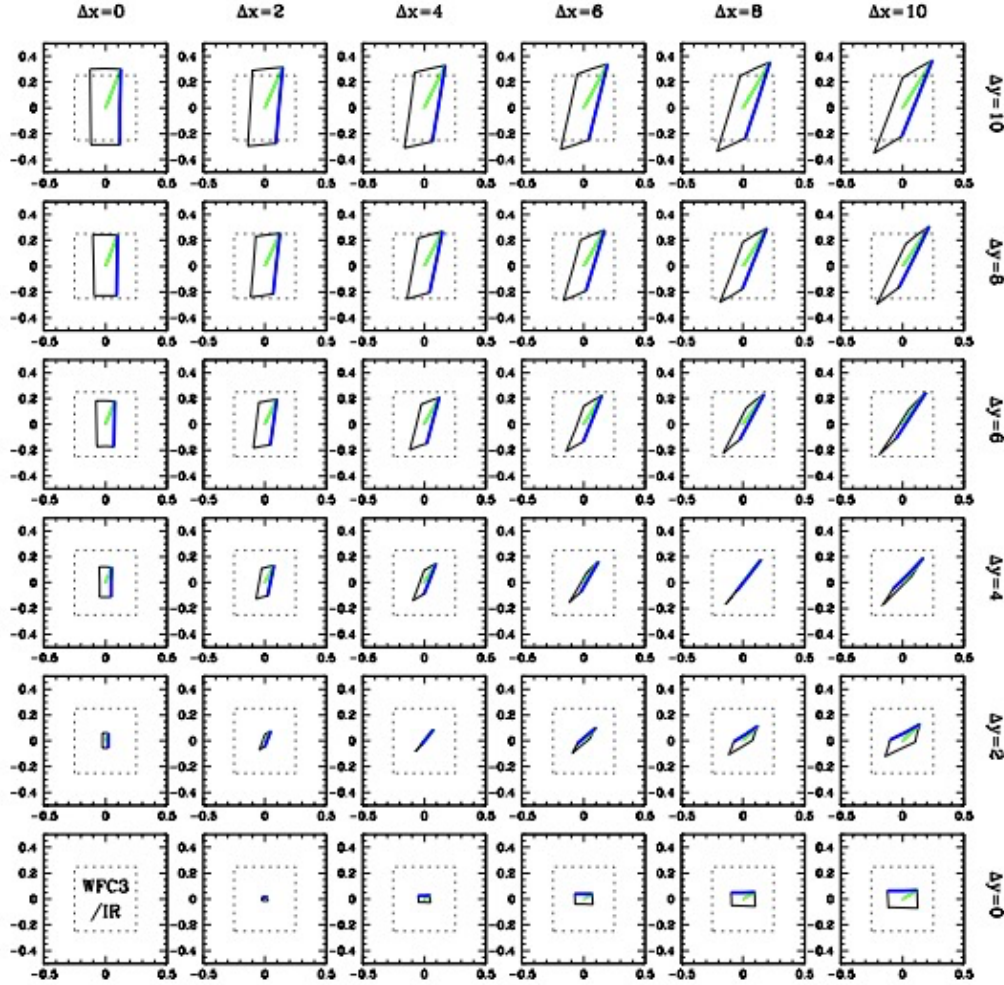


Figure A1: Similar to Figure 7, but for WFC3/IR. This shows how a whole pixel dither with offset $(\Delta x, \Delta y)$ can lose pixel-phase coherence at the edges and corners of the WFC3/IR detector. A dither of ~ 6 pixels ends up shifting pixel phase by 0.25 pixel, which is enough to cause problems with a 4-point dither.

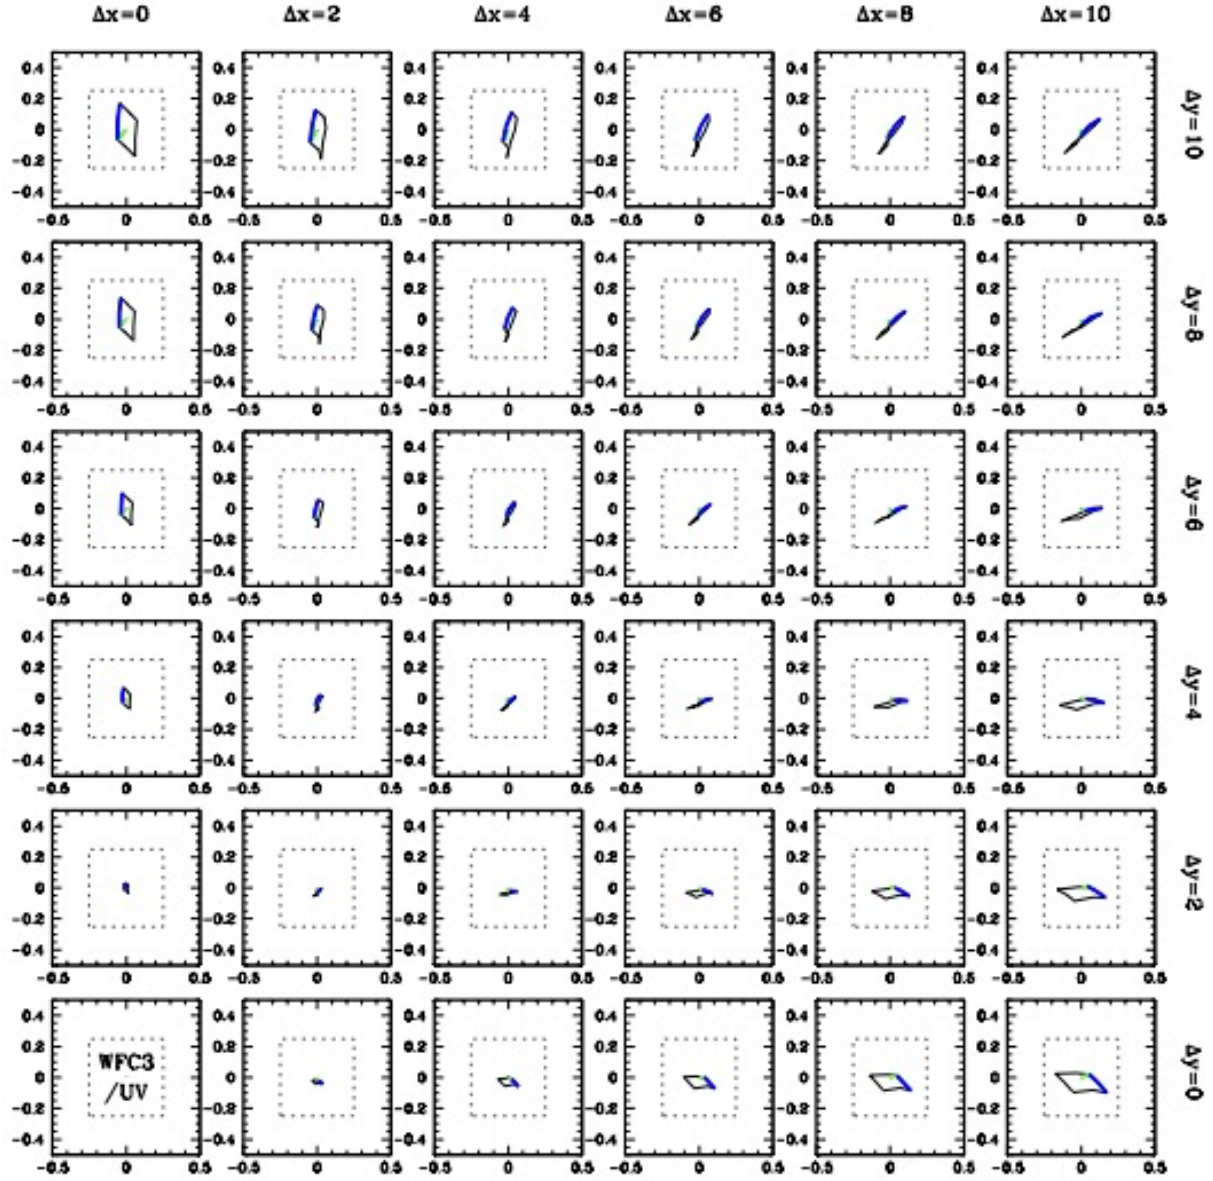


Figure A2: Similar to Figure 7, but for WF3/UVIS. This shows how a whole pixel dither with offset ($\Delta x, \Delta y$) can lose pixel-phase coherence at the edges and corners of the WFC3/UVIS detector. A dither of ~ 10 pixels ends up shifting pixel phase by 0.2 pixel, which is still not enough to cause problems with a 4-point dither.

APPENDIX B

Since the initial publication of this ISR, we received some requests from users to provide POS-TARGs to enable good sub-pixel dithers in prime and parallel instruments for 2-point and 3-point observation sets. We provide them here. With these 2-, 3-, and 4-point dithers it is possible to construct a set of dithers for N exposures that accomplishes good dithering in both prime and parallel by dividing the N exposures up into groups of 2-, 3-, or 4-pt dithers with ~ 1 arcsecond offsets between the groups. This will provide a much more consistent dither in the parallel instrument than will a set of N essentially random offsets.

For the 2-point dither, we adopt for the prime instrument (0.000, 0.000) and (0.500, 0.500) and use the flexibility of the integer offsets to get a sub-pixel dither as close as possible to this in the parallel instruments. **Table B1** provides the POS-TARG specifications for the optimized dithers, and **Figure B1** shows the achieved dithers graphically.

Table B1: List of the 2-point-dither POS TARGs (in arcseconds) that achieve a “perfect” sub-pixel dither in the prime instrument and an “optimized” sub-pixel dither in the parallel instrument.

PRIME/PARALLEL	ACSWFC/WFC3UV	ACSWFC/WFC3IR	WFC3UV/ACSWFC	WFC3IR/ACSWFC
DITHER#1	0.00000, 0.00000	0.00000, 0.00000	0.00000, 0.00000	0.00000, 0.00000
DITHER#2	0.12337, 0.18347	-0.07405, 0.16747	0.17859, 0.06987	0.06773, 0.30266

For the 3-point dither, we adopted two possible sub-pixel dithers for the prime instrument, one was (0.000, 0.000), (0.666, 0.333), and (0.333, 0.666) and the other was (0.000, 0.000) (0.333, 0.333) and (0.666, 0.333). They are both lines with spacing of 0.471 pixel. **Table B2** provides the POS-TARG specifications for the optimized dithers and **Figure B2** shows the achieved dithers graphically.

Table B2: List of the 3-point-dither POS TARGs (in arcseconds) that achieve a “perfect” sub-pixel dither in the prime instrument and an “optimized” sub-pixel dither in the parallel instrument.

PRIME/PARALLEL	ACSWFC/WFC3UV	ACSWFC/WFC3IR	WFC3UV/ACSWFC	WFC3IR/ACSWFC
DITHER#1	0.00000, 0.00000	0.00000, 0.00000	0.00000, 0.00000	0.00000, 0.00000
DITHER#2	0.16449, 0.24459	0.16475, 0.07888	-0.14548, +0.04326	-0.36113, -0.04164
DITHER#3	0.08221, 0.12230	0.08234, 0.03942	-0.21160, 0.09172	-0.18062, 0.28160

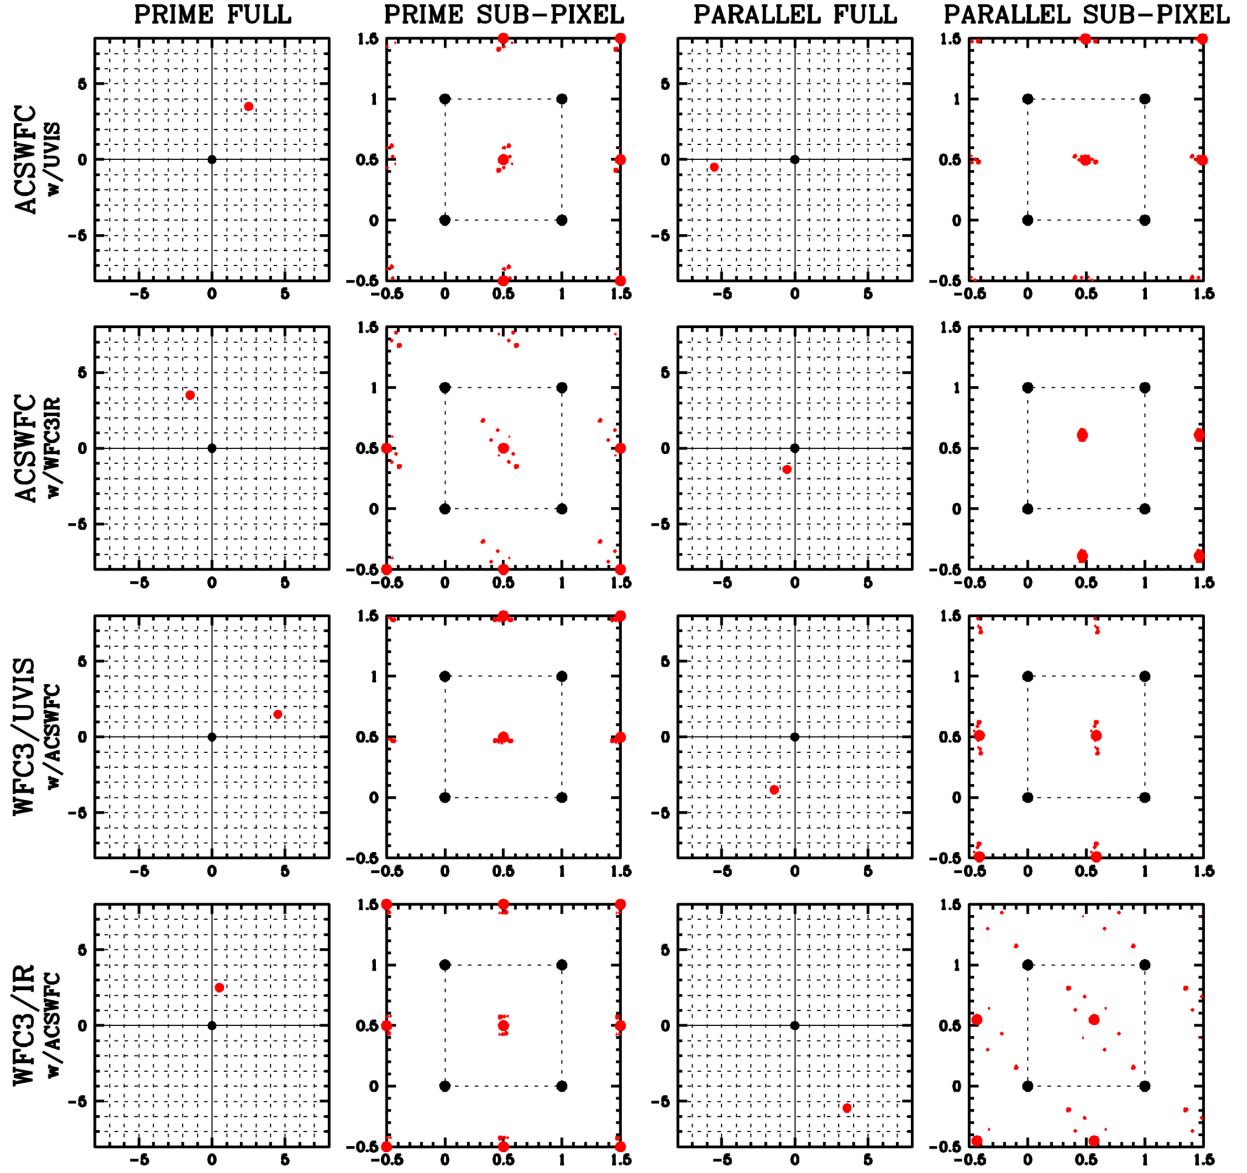


Figure B1: For each prime/parallel combination (labeled on the left), we show the dithers achieved for the optimized POS-TARGs for the 2-point dithers. The left panels show the achieved dither in the prime instrument, the second panel shows the sub-pixel dither in the prime instrument, and the right panels show the same for the parallel instruments. The large symbols show the dither at the center of the detector and the smaller ones at the edges/corners. The lower right panel shows the dither in ACS/WFC when WFC3/IR is prime; there is clearly a large spread in pixel phase at the edges/corners, but there were no smaller dithers with good pixel-phase distributions in ACS/WFC.

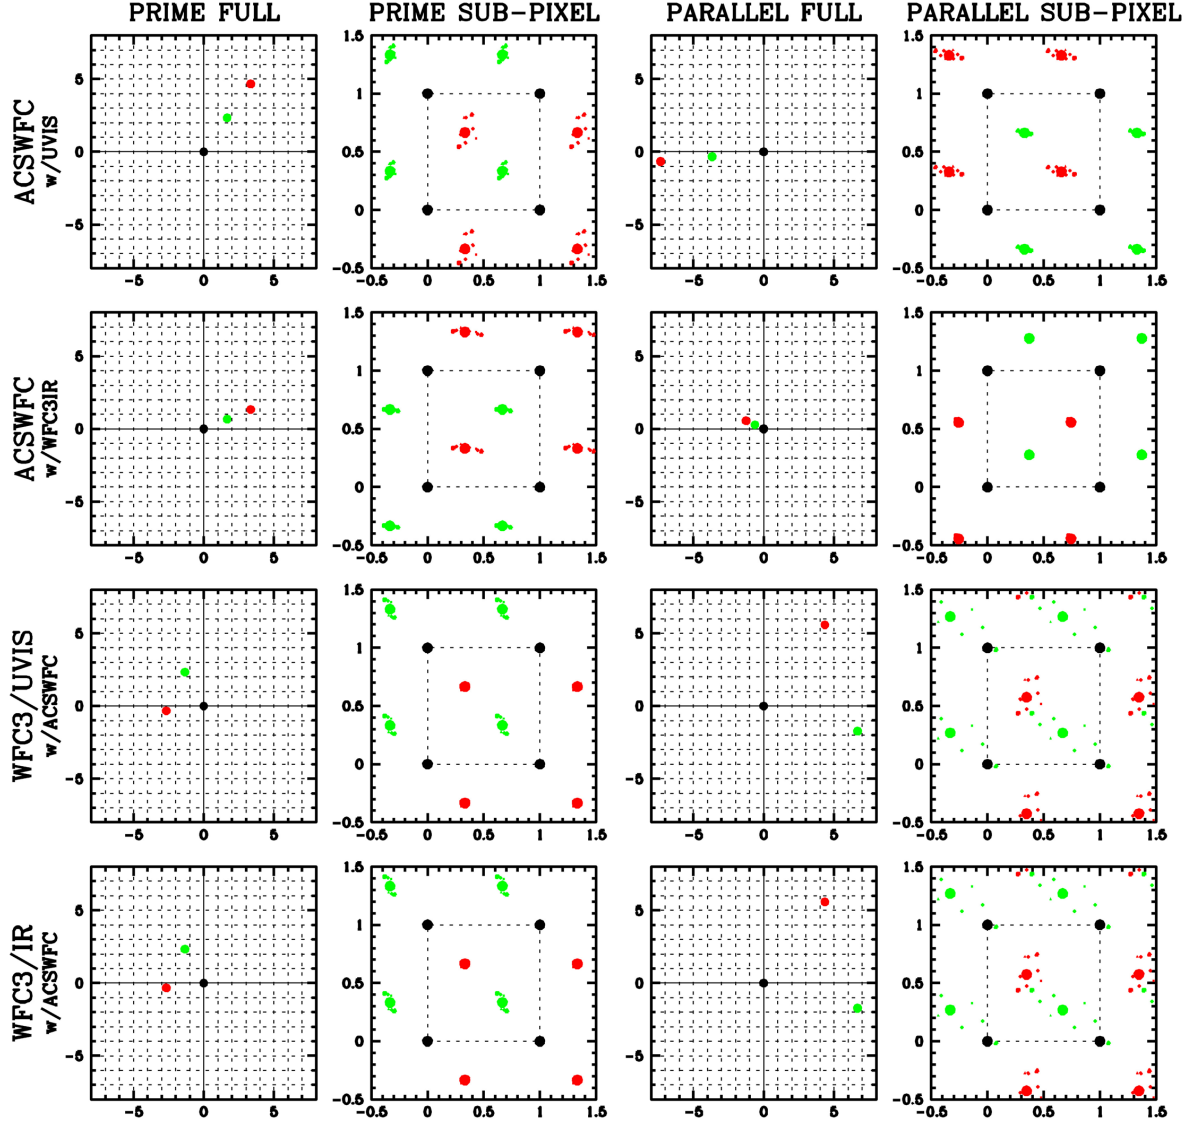


Figure B2: For each prime/parallel combination (labeled on the left), we show the dithers achieved for the optimized POS-TARGs for the 3-point dithers. The left panels show the achieved dither in the prime instrument, the second panel shows the sub-pixel dither in the prime instrument, and the right panels show the same for the parallel instruments.