# CHAPTER 7: Observation Strategies

## In This Chapter...



## 7.1 Observing Faint Targets

For broad band filters the sky background will limit the detection of faint targets. For example, an 8 orbit observation in F555W gives a 5$\sigma$ detection limit at V=28.6 for an average sky level 23 mag arcsec$^{-2}$ in V. Note that the sky background is a strong function of position, especially for targets near the ecliptic; the sky level can vary from V=23.3 mag arcsec$^{-2}$ at the ecliptic pole to about V=20.9 mag arcsec$^{-2}$ on the ecliptic near the solar avoidance limit. (See Table 6.3 on page 113 for sky level as function of ecliptic coordinates.)

If these higher sky levels would severely impact the science data, observers should consider specifying the special requirement LOW-SKY on the Phase II proposal. This parameter forces the observation to be made when the sky background is within 30% of the minimum value for the target. Note, however, that this will also reduce the number of HST calendar windows available to the observation, and so could result in scheduling delays or may even make the observation infeasible if there are other constraints such as ORIENTs. A minor decrease in the per-orbit visibility period also results from LOW-SKY, but for background limited programs this is a minor price to pay for the guarantee of a much lower background. In summary, LOW-SKY will reduce the sky background, but should only be used if the science goals *require* it.

*Note that LOW-SKY cannot be used for CVZ targets, as they imply mutually exclusive pointing constraints.*

Scattering of bright Earth light in the OTA can produce non-uniformities in the background which may hamper analysis of faint target images. Most often these take the form of diagonal bars of suppressed background light in several of the CCDs. These effects tend to occur for broad band filters when the OTA axis is about 25° from the bright Earth. This effect is most often seen in observations of targets in the CVZ (continuous viewing zone), since the Earth limb is never very far from the OTA axis when observing in the CVZ. Figure 7.1 shows a typical case. LOW-SKY will eliminate this effect for non-CVZ targets, as it places the OTA axis more than 40° from the bright Earth. Alternatively, one can place the target away from the CCD center to avoid these artifacts.

Another option for reducing the sky brightness, is the special requirement SHADOW, which forces the observation to be made when HST is in the Earth's shadow. This usually has a large negative impact on the observing efficiency, and is recommended only when observing far-UV emission lines (e.g. Ly $\alpha$ and OI 1304Å). Its primary goal is only to reduce geocoronal emission lines. Moreover, it does not attempt to minimize zodiacal emission, which dominates at visible wavelengths.

## 7.2 Observing Bright Targets

Saturation is the primary concern when observing bright targets. The analog-to-digital converter will run out of bit codes at ~28,000 $e^-$ $pixel^{-1}$ in the ATD-GAIN=7 (default) setting, and at ~53,000 $e^-$ $pixel^{-1}$ in the ATD-GAIN=15 setting. Count levels above these are merely reported as 4095 DN in the data. Hence ATD-GAIN=15 is recommended for targets approaching 28,000 $e^-$ $pixel^{-1}$. The disadvantage of this setting is that the read noise is poorly sampled by this coarse digitization, and hence the read noise is slightly increased.

At count levels above ~90,000 $e^-$ $pixel^{-1}$ charge will overflow the potential well of each pixel, and begin to bloom up and down the CCD columns. For example, this occurs in the F555W filter at about V=13.5 for a 10s exposure on the WFC, and at about V=13.0 on PC1.

At very high count levels, above ~$10^8$ $e^-$ per CCD column, the charge bloom will reach the top and bottom of the CCD and flow into the serial registers. CLOCKS=YES will dispose of this charge as it reaches the ends of the CCD, and thus prevent it from leaking back into adjacent CCD columns. This exposure level corresponds roughly to a 10s exposure of a V=7 star in F555W. Note that CLOCKS=YES offers no benefit unless the bloom reaches the ends of the CCD, and that it may slightly compromise the bias and dark calibration. Moreover, CLOCKS=YES will result in anomalous exposure times; exposure times are rounded to the nearest integral second, minus a delay time of up to 0.25s for the

**Figure 7.1:** Example of Scattered Earth Light. Scattered light contributes ~100 $e^-$ of background throughout this image. The camera spiders block some of this scattered light along CCD diagonals, hence forming “X” patterns and bars where the background is reduced by ~40% in this image.

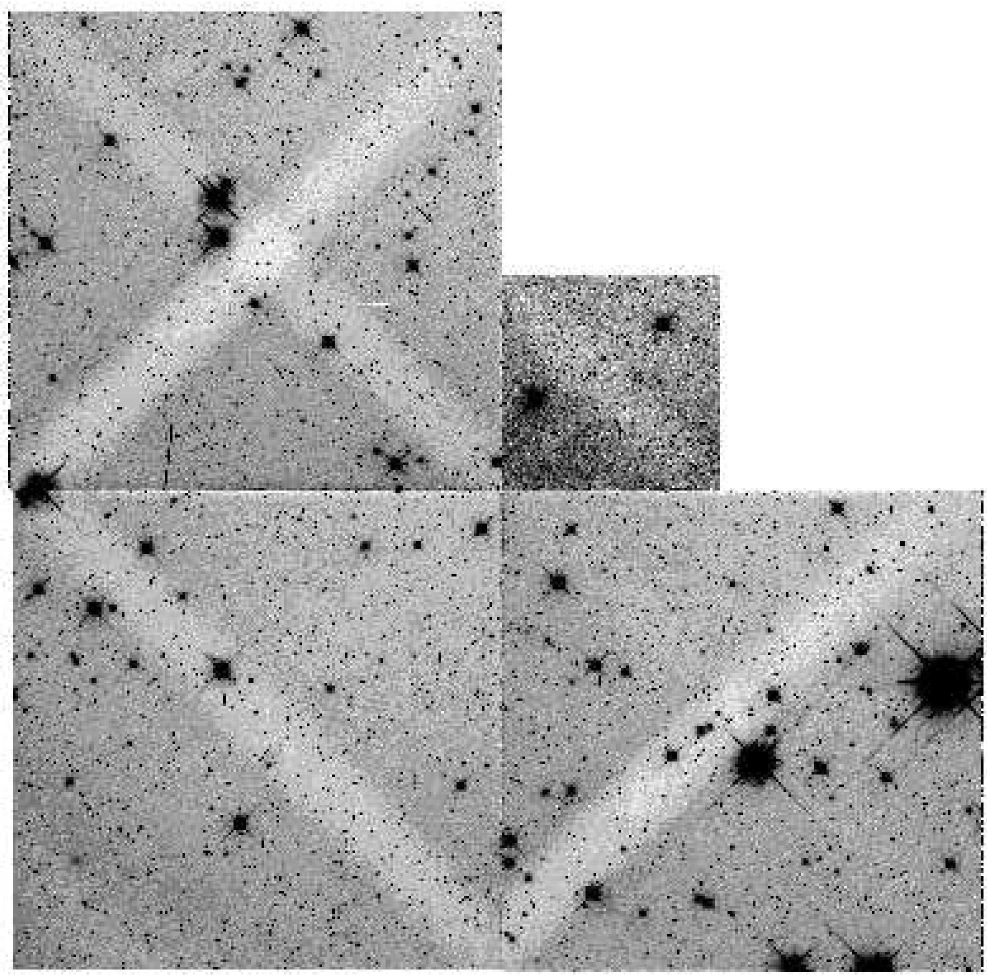

shutter to open. (See “Serial Clocks” on page 27 for further discussion of exposure time anomalies caused by CLOCKS=YES.)

Besides setting ATD-GAIN=15, the PC CCD can be used to reduce saturation effects for stellar objects. The peak of the PSF will be spread over more pixels on the PC (vs. WFC), so stars can be exposed about 50% longer on the PC before saturation sets in.

Note that the narrow band filters may be useful when observing very bright targets. For example, stars as bright as V~4.4 can be observed without saturation in F502N using the PC at ATD-GAIN=15 with a 0.11s exposure time.

# 7.3 Observing Faint Targets Near Bright Objects

The concerns here are similar to those for observing bright targets; saturation and blooming of the bright companion PSF must not impact the faint target. Also, one may need to consider subtracting the PSF of the bright object, and effects which limit the accuracy of that subtraction.

If the bright companion will saturate and bloom, it will be necessary to rotate the CCD so that blooming along the CCD columns does not obliterate the faint target. See Figure 7.8 on page 161 for an illustration of the bloom directions. It may also be useful to orient the field so that the OTA diffraction spikes from the bright companion (along diagonal lines on the CCDs) avoid the faint target. Table 7.1 summarizes ORIENTs which can be used to avoid CCD blooming tracks and OTA diffraction spikes caused by bright objects. For example, if a faint companion is at PA 60˚ on the sky relative to a bright companion, it would be advantageous to observe on PC1 with ORIENT= PA + 45˚ = 105˚. Ideally, some range in ORIENT would be specified to ease scheduling, hence “ORIENT=90D TO 120D” might be specified on the Phase II proposal. Note that “ORIENT=270D TO 300D” is also feasible, and should be indicated in the visit level comments.

**Table 7.1:** ORIENTs for Avoiding Bloom Tracks and Diffraction Spikes. “PA” is the position angle of the faint target relative to the bright object. Note that ORIENT should be between 0D and 360D, so subtract 360˚, if necessary. On proposal these are specified as, e.g., “ORIENT=231D TO 261D”.

| CCD | ORIENT |
|---|---|
| PC1 | PA+30˚ to PA+60˚, PA+210˚ to PA+240˚ |
| WF2 | PA+120˚ to PA+150˚, PA+300˚ to PA+330˚ |
| WF3 | PA+30˚ to PA+60˚, PA+210˚ to PA+240˚ |
| WF4 | PA+120˚ to PA+150˚, PA+300˚ to PA+330˚ |

If instead of observing a known companion, one is *searching* for companions, it is advisable to observe at several ORIENTs so that the CCD bloom track and OTA diffraction spikes will not hide possible companions. For example, three ORIENTs, each separated by 60˚, would give good data at all possible companion position angles.

If PSF subtraction will be needed during data analysis, then the PC CCD may have some advantage, since it provides better sampling of fine undulations in the PSF. It may also be useful to obtain observations of a second bright star for PSF calibration, though these may be of limited utility since thermal effects and OTA “breathing” can modify the telescope focus, and hence the PSF, on timescales of less than one hour. Any such PSF star should be similar in color to the target, and should be observed at the same CCD position (within 1″) and with the same filter. Sub-pixel dithering may also be useful, so as to improve sampling of the PSF (see “Dither Strategies” on page 155).

Figure 7.2 illustrates the effect of OTA breathing, and periodic focus adjustments, on PSF subtraction. It shows the difference between an “in focus” PSF and one where the OTA secondary mirror has been moved by 5μm. This amount of focus change is comparable to the range of OTA “breathing” effects (timescale <1 hour), and the periodic (semi-annual) focus adjustments of the OTA. Each panel shows a different contrast setting; the percentages indicate the energy per pixel which is plotted as white, expressed as a fraction of the total (un-subtracted) PSF energy. For example, features which are just white in the “0.003%” panel contain 0.003% of the total PSF energy in each pixel. In other words, the feature labeled “a” is, in effect, ~10 magnitudes fainter than the PSF of the bright object, so that it may be very difficult to detect a “real” companion object ~10 magnitudes fainter than the bright object, at this distance from the bright object. In a real PSF subtraction situation, other effects including PSF sampling, noise, and pointing instability would further degrade the subtraction. (The elongated appearance of the residuals in the PSF core is due to astigmatism in PC1.)

Table 7.2 gives some quantitative indication of the performance expected for PSF subtractions in the high signal-to-noise limit. It gives the magnitude of “star-like” artifacts remaining in the subtracted image, as a function of distance from the bright object, and magnitude $m_{bright}$ for the bright object. The right-most column gives an effective magnitude limit imposed by artifacts from the PSF subtraction. These results are derived for the 5μm focus shift described above, and are for PC1 and filter F555W. It may be possible to do somewhat better than these limits by subtracting accurate model PSFs, or by finding an observed PSF with matching focus.

**Table 7.2:** Approx. PSF Subtraction Artifact Magnitudes and Magnitude Limits.

| Distance from Bright Object | Effective Magnitude of Subtraction Artifacts | Effective Faint Object Detection Limit (3σ) |
|---|---|---|
| 0.1″ | $m_{bright}$+4.7 | $m_{bright}$+3.5 |
| 0.3″ | $m_{bright}$+8.6 | $m_{bright}$+7.4 |
| 1″ | $m_{bright}$+11.4 | $m_{bright}$+10.2 |
| 3″ | $m_{bright}$+13.2 | $m_{bright}$+12.0 |

Large angle scattering may also impact identification of very faint objects near very bright ones. This scattering appears to occur primarily in the camera relay optics, or in the CCD. Hence, if a faint target is more than ~10″ from a bright object (i.e. very highly saturated object), it would be advisable to place the bright object on a different CCD, so as to minimize large angle scattering in the camera

**Figure 7.2:** Impact of OTA Focus Shift on PSF Subtraction. Each image shows the difference between an “in focus” and a 5 micron defocused PSF at different contrast settings. Numbers indicate the energy per pixel which is plotted as white, as a percentage of total energy in the un-subtracted PSF. Based on Tiny-TIM models for PC1 in F555W filter.

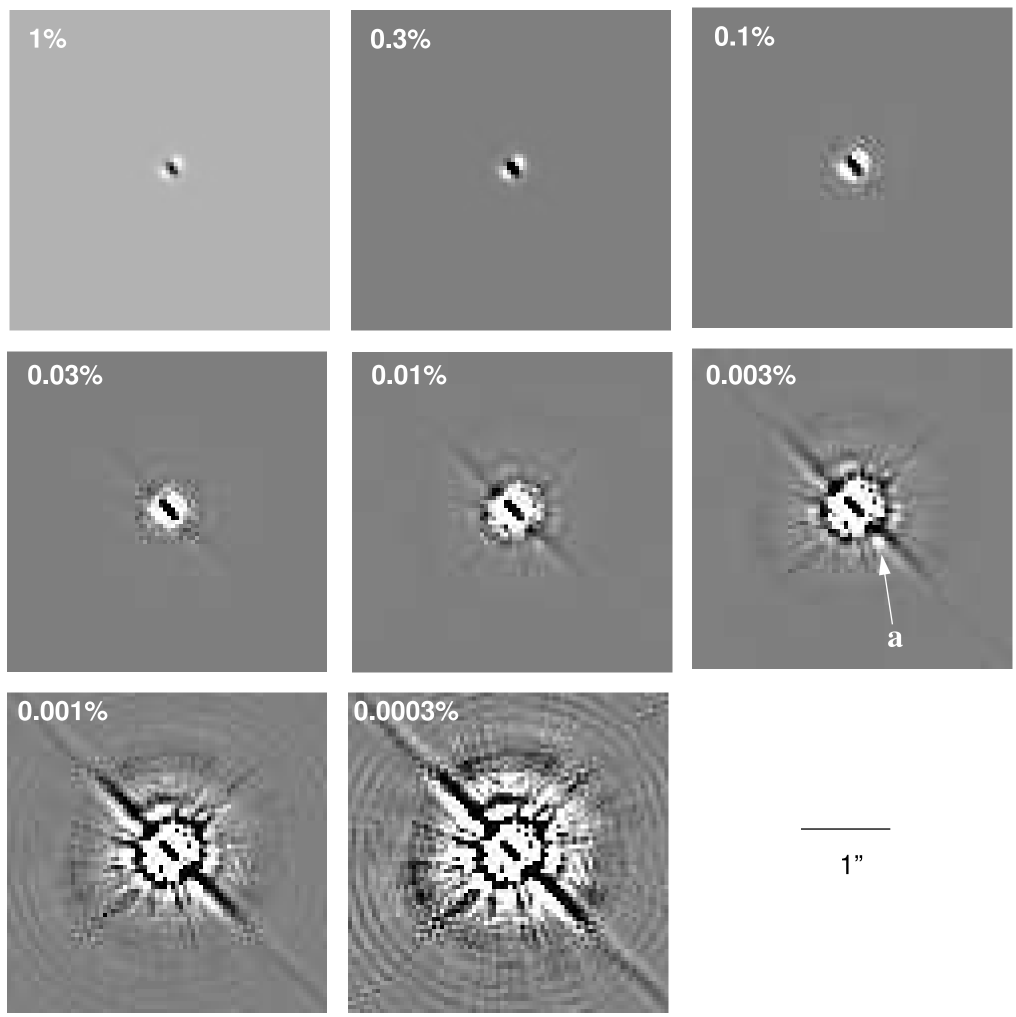


containing the faint target. See section on “Large Angle Scattering” on page 99. Note also that highly saturated PSFs exist for PC1 in filters F439W, F555W, F675W, and F814W, and for F606W on WF3; these may be useful when attempting to subtract the large-angle scattered light. As of this writing TinyTIM does not accurately model the large angle scattering, and should be used with caution when analyzing highly saturated images (Krist 1996).

It is generally unwise to place bright companions or other bright objects just outside the area imaged by the CCDs. The region of the focal plane just outside the CCDs (within about 6″ of the CCDs) contains a number of surfaces which can reflect light back onto the CCDs, hence placing bright targets there can have undesired results. Also, the un-imaged "L" shaped region surrounding PC1 should be avoided, since incomplete baffling of the relay optics allows out-of-focus images of objects in this region to fall on the CCDs. Figure 7.3 illustrates various bright object avoidance regions near the WFPC2 field-of-view; the indicated avoidance magnitudes will produce 0.0016 $e^{-}$ $s^{-1}$ $pixel^{-1}$ in the stray light pattern for F555W. Figures 7.4 and 7.5 show examples of artifacts which can result from bright stars near the PC1 CCD. The report "A Field Guide to WFPC2 Image Anomalies" (WFPC2 ISR 95-06, available on the WFPC2 WWW pages and from `help@stsci.edu`) gives more detailed discussions of artifacts associated with bright objects, and their avoidance.

**Figure 7.3:** Bright Object Avoidance Regions Near WFPC2 FOV.

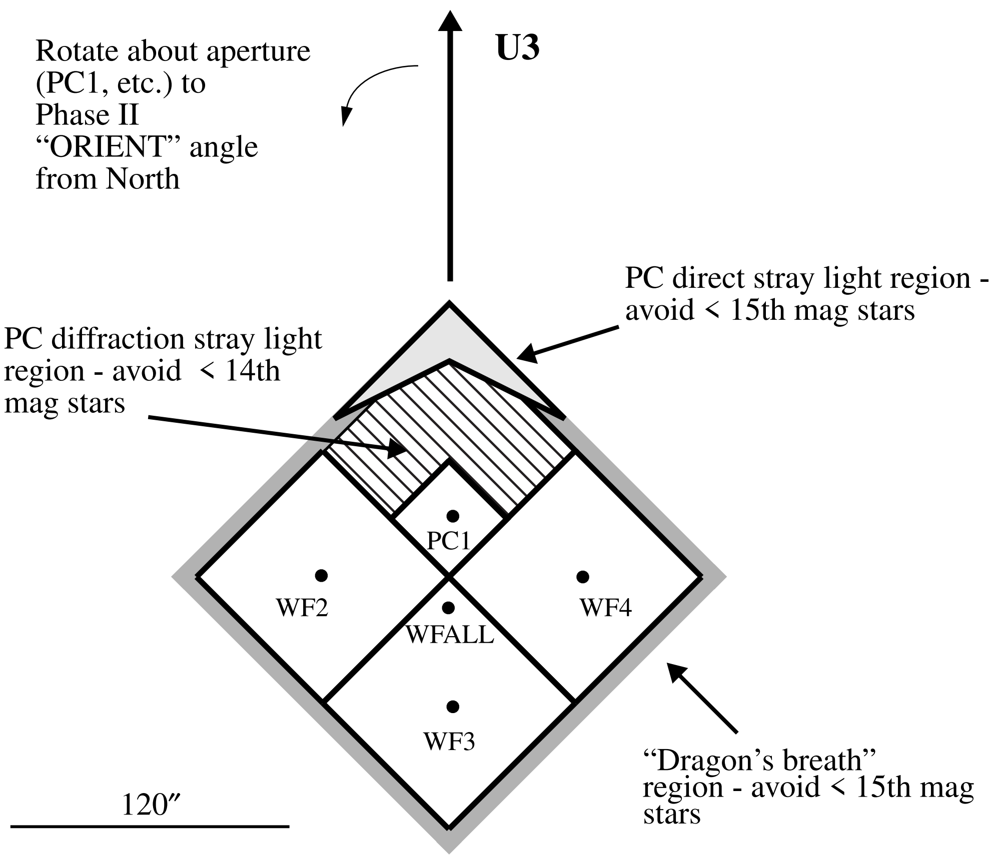

**Figure 7.4:** Example of PC1 “Direct” Stray Light Ghost.

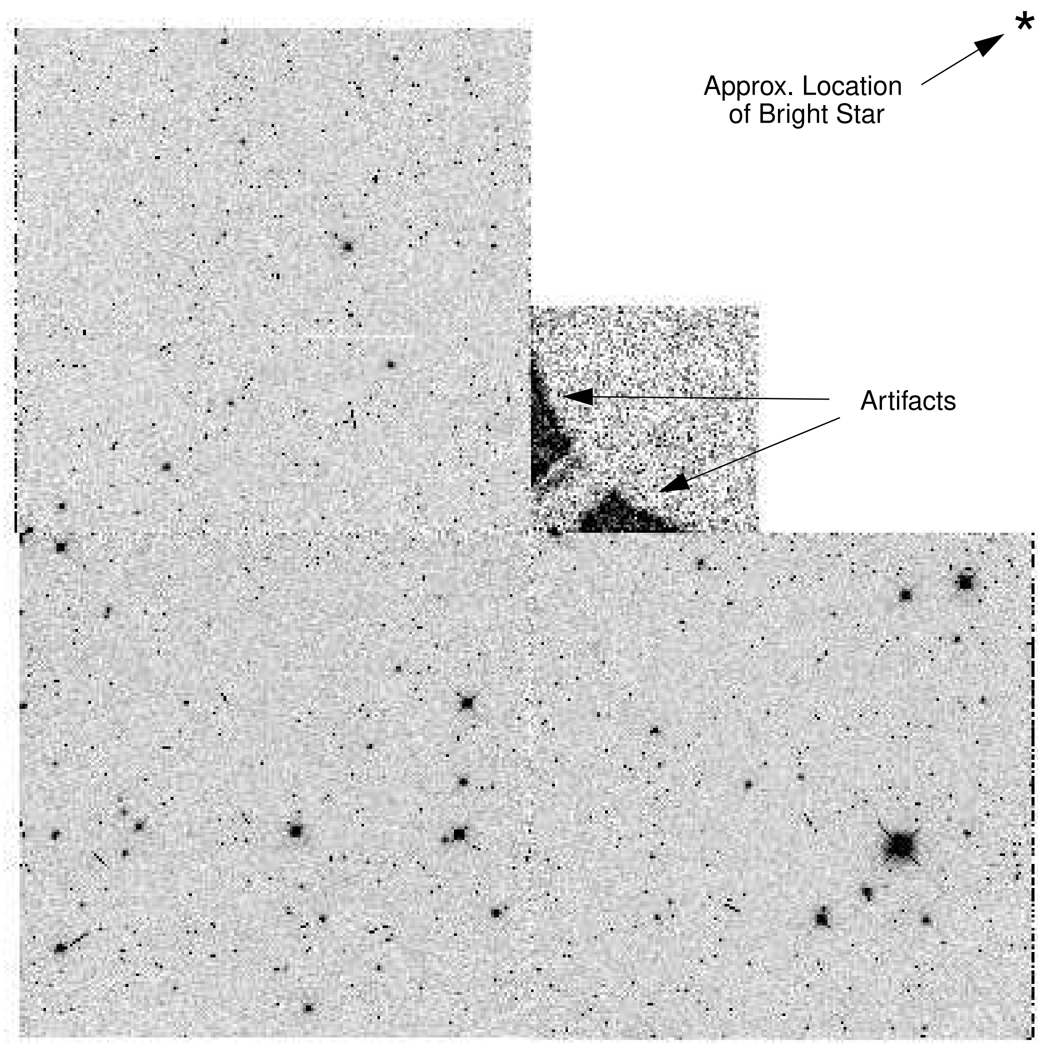

**Figure 7.5:** Example of PC1 "Diffraction" Stray Light Ghost.

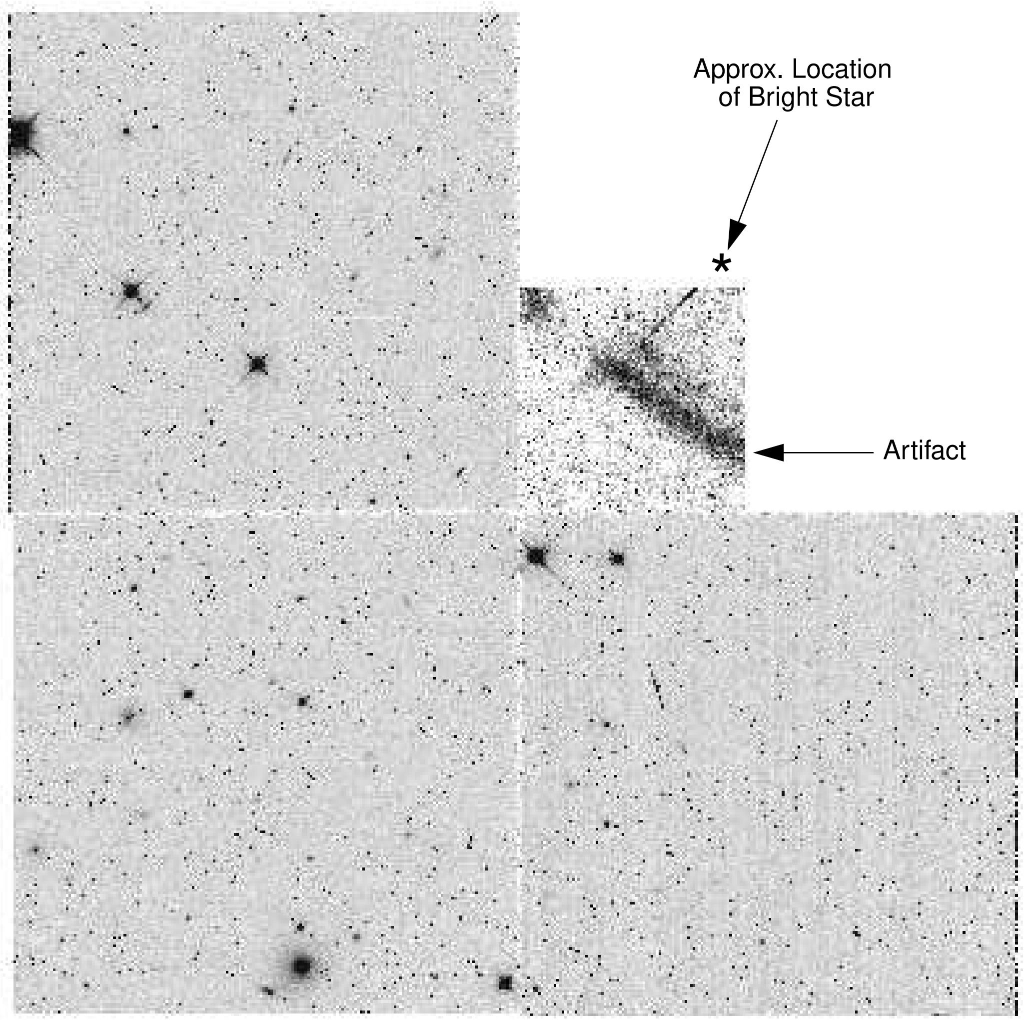


## 7.4 Cosmic Rays

Cosmic rays will obliterate ~20 pixels per second per CCD. It is imperative that two or more images be obtained at each pointing position, if these artifacts are to be removed from the data. The default action by the Phase II proposal processing software is to split exposures longer than 600s into two near equal parts, so as to allow removal of the cosmic ray tracks. The CR-SPLIT and CR-TOLERANCE optional parameters on the Phase II proposal allow observers to adjust this behavior. CR-SPLIT can be set to either DEF (default), NO, or a numeric value (0.0 to 1.0) giving the fraction of the total exposure allotted to the first sub-exposure of the pair. CR-TOLERANCE indicates the spread allowed in dividing the exposure, as a fraction of the total exposure time. For example, the default CR-TOLERANCE=0.2 allows the first sub-exposure to range from 0.3 to 0.7 of the total exposure. Setting CR-TOLERANCE=0 will force equal-length sub-exposures.

The required degree of cosmic-ray avoidance will depend on the science goals of the proposal; observations of a single small target will usually suffer much less impact from cosmic rays than programs needing very "clean" data over a large area. Table 7.3 gives very rough recommendations for the number of sub-exposures for a given total exposure time. Note that splitting into *many* sub-exposures introduces additional overhead time and will increase the noise for "read noise" limited exposures (usually exposures in UV or narrow band filters), and hence one should not use more sub-exposures than are truly required by the science goals.

**Table 7.3:** Recommended Exposure Splittings.

| Total Exposure Time (s) | Rough Recommended Number of Sub-exposures | |
|---|---|---|
| | Programs with Single Small Target | Wide-area Search Programs |
| <300 | 1 | 3 |
| 300 - 600 | 1 or 2 | 4 |
| 600 - 1600 | 2 or 3 | 4 |
| 1600 - 5000 | 3 | 5 |
| 5000 - 10000 | 4 | 6 |
| >10000 | One exposure per orbit (2400s each) | |

## 7.5 Choosing Exposure Times

The choice of exposure time generally depends on the signal-to-noise ratio required to meet the science goals. This can be assessed using information in Chapter 6 or Appendix B herein, or by using the on-line WWW Exposure Time Calculator tool.

However, when packing orbits, one must often compromise somewhat and decide which exposures to lengthen or shorten. Table 7.4 may be helpful in this regard. It shows the total time required to execute a single CR-SPLIT=NO exposure, *excluding* any time needed to change filters.

Note that the most efficient exposure times are those whose length approaches or equals, but does not exceed, an integral number of minutes *plus* 40s. Figure 7.6 illustrates event timings during a typical 60s WFPC2 exposure, similarly, Figure 7.7 illustrates events during a (more efficient) 100s exposure. (See "Overhead Times" on page 29 for more information about exposure timings.)

**Figure 7.6:** Event Timings During a 60s WFPC2 Exposure. All events, except shutter opening, start on 1 minute spacecraft clock pulses. Both the CCD clear and readout of each CCD require 13.6s. This 60s exposure, including the filter change, requires 4 minutes.

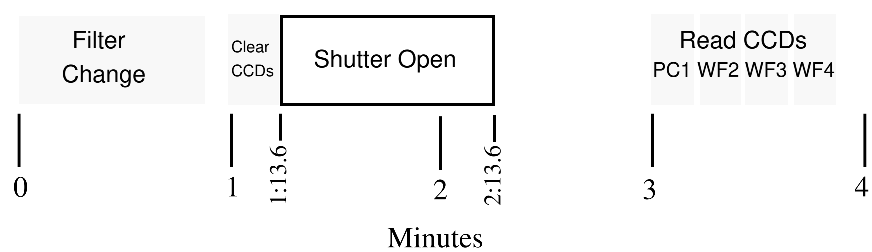


**Figure 7.7:** Event Timings During a 100s WFPC2 Exposure. This exposure, including the filter change, requires 4 minutes.

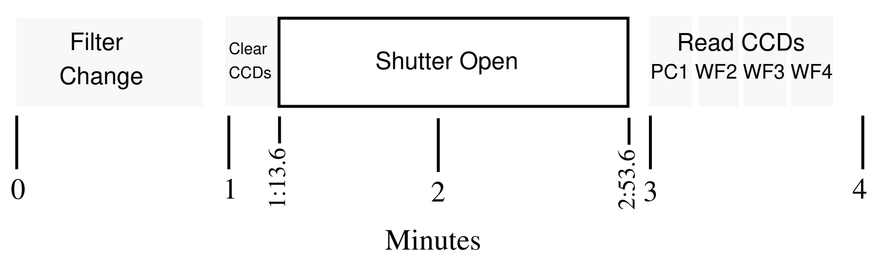


Due to the various overheads, shortening or lengthening an exposure can have unexpected effects on the orbit packing. For example, shortening an exposure from 400s to 350s has no effect on orbit packing; they both require 9 minutes to execute (CLOCKS=NO, the default setting). On the other hand, shortening an exposure from 180s to 160s trims the execution time by 2 minutes (again CLOCKS=NO, the default setting).

CLOCKS=YES may have some advantage in a long series of exposures whose lengths are 180s or somewhat greater. Each savings of 1 minute can add up to a few more exposures per orbit. The down side is that most calibrations are derived for exposures with CLOCKS=NO, so the calibration maybe slightly compromised. The largest calibration error is expected to occur on the dark current, where there may be a slight increase near the top and bottom of each CCD. In many situations this error may be acceptable, such as a small target near a CCD center, or broad band filter images where the sky completely dominates the dark current. CLOCKS=YES will have more impact on calibration of narrow filters, or situations requiring an extremely flat background. (Also, see “Serial Clocks” on page 27 for discussion of exposure time anomalies associated with CLOCKS=YES, though these are most important for exposures <30s.)

An exposure with CR-SPLIT=YES would simply require the total time for each sub-exposure as given by Table 7.4, again, plus any time needed to change filter before the first exposure. However, the default CR-SPLITting allows schedulers some latitude in dividing the exposures (CR-TOLERANCE=0.2 is the default) so the exact overheads are unpredictable. For example, a 700s exposure with CR-SPLIT=0.5 (the default) could be split into a pair of 350s exposures totaling 18 minutes, or a 300s and 400s exposure totaling 17 minutes.

**Table 7.4:** Basic Time to Execute Single Non-CR-SPLIT Exposure. This includes time to prep the CCD, execute the exposure, and readout the CCDs. Times needed to change filter (1 minute), or insert a second filter (1 minute), are *excluded*. See “Overhead Times” on page 29 for more discussion and other overheads.

| Exposure Time (s) | Total Execution Time (min.) | |
|---|---|---|
| | CLOCKS=NO (default) | CLOCKS=YES |
| 0.11 to 30 | 2 | (not recommended) |
| 35, 40 | 2 | 2 |
| 50,60,70,80,100 | 3 | 3 |
| 120,140,160 | 4 | 4 |
| 180,200 | 6 | 5 |
| 230,260 | 7 | 6 |
| 300 | 8 | 7 |
| 350,400 | 9 | 8 |
| 500 | 11 | 10 |
| 600 | 13 | 12 |
| 700 | 14 | 13 |
| 800 | 16 | 15 |
| 900 | 18 | 17 |
| 1000 | 19 | 18 |
| 1100 | 21 | 20 |
| 1200 | 23 | 22 |

# 7.6 Dither Strategies

There is no single observing strategy that is entirely satisfactory in all circumstances for WFPC2. One must consider cosmic rays, hot pixels (i.e. pixels with high, time variable dark count), spatial undersampling of the image, and large-scale irregularities such as the few arcsecond wide region where the CCDs adjoin. One strategy that can be used to minimize the effects of undersampling and to reduce the effects of hot pixels and imperfect flat fields is to dither, that is, to offset the telescope by either integer or sub-pixel steps. The best choice for the number and size of the dithers depends on the amount of time available and the goals of the project. In the following we will address a few issues related to dithering:

1. **Undersampling:** Individual images taken with ***sub-pixel*** offsets can be combined to form an image with higher spatial resolution than that of the original images. A single dither from the original pixel position -- call it (0,0) -- to one offset by half a pixel in both x and y, (0.5,0.5) will produce a substantial gain in spatial information. On the other hand very little extra information is gained from obtaining more than four positions, if the standard four point dither is used, and if the telescope has successfully executed the dither. Therefore the recommended number of sub-pixel dither positions is between 2 and 4.

2. **Hot Pixels:** There are three ways to deal with hot pixels: correct using "dark frames" that bracket the observation, dither by an ***integer*** amount of pixels, or use a task such as "WARMPIX" within STSDAS to filter out the known hot pixels. Note that the integer dither strategy would ideally use six images, i.e. two CR-SPLIT images at each of three different dither positions. This is because in addition to hot pixels, low or "cold" pixels[1] can be present and simple strategies selecting the minimum of two pixel values can fail. However, even four images (two each at two dither positions) will greatly aid in eliminating hot pixel artifacts.

3. **Cosmic Rays:** Although dithering naturally provides many images of the same field it is better to take several images at each single pointing in order to remove cosmic rays. In principle, it should be possible to remove cosmic rays using only sub-pixel dithered data. At present, however, no publicly released software is available for this task. *Hence we recommend obtaining two or more images (i.e. CR-SPLITing)* ***at each position*** *in the dithered sequence.* For very long integrations it is convenient to split the exposure into *more than two separate images.* As an example, for two 2000s exposures, about 1000 pixels per chip will be hit in both images and will therefore be unrecoverable. Moreover, since CR events typically affect 7 pixels per event, these pixels will not be independently placed, but rather will frequently be adjacent to other unrecoverable pixels.

1. Cold pixels usually result from hot pixels in the dark calibration file which do not actually appear in the science data.

4. **Accuracy of dithering**: We do not yet have good statistics on the accuracy of HST dither offsets. During the Hubble Deep Field, nearly all dithers were placed to within 10 mas (during ±1.3″ offsets and returns separated by multiple days), although in a few cases the dither was off by more than 25 mas, and on one occasion (out of 107 reacquisitions) the telescope locked on a secondary FGS peak causing the pointing to be off by approximately 1″ as well as a field rotation of about 4 arcminutes. The software which was developed for the Hubble Deep Field is able to reconstruct images even for these non-optimal dithers, still gaining in resolution over non-dithered data. This software will be available in STSDAS in mid 1996 and is based on the variable pixel linear reconstruction technique developed by Fruchter and Hook (this procedure is also known as “dripping and drizzling”). The Richardson-Lucy-Hook non-linear deconvolution technique (STSDAS task ACOADD) can also use non-optimal dithers.

The simplest way to schedule dithers is to use the options DITHER-TYPE=LINE (e.g. for two-point diagonal dithers) or DITHER-TYPE=BOX (for four-point dithers). Alternative approaches are to specify a spatial scan or to use POS TARGs. Note that when the WF3 is specified as an aperture, the POS TARG axes run exactly along the WF3 rows and columns. For the other chips, they only run approximately along the rows and columns due to the small amount of rotation between CCDs. For small dithers these rotations are not important.

Some specific offsets allow one to shift by convenient amounts both the PC and the WFC chips. For instance an offset of 0.5″ is equivalent to 5 WFC pixels and 11 PC pixels. The default DITHER-LINE spacing of 0.3535″ along the diagonal is equivalent to shifts of (2.5,2.5) pixels for the WFC and (5.5,5.5) pixels for the PC.

Note that large dithers will incur errors due to the camera geometric distortion which increases toward the CCD corners and alters the image scale by about 2% at the corners. Hence a 1.993″ offset will be 20.3 WF pixels at the field center, but suffer a 0.4 pixel error at the CCD corners. Large dithers may also occasionally require a different set of guide stars for each pointing, thus greatly reducing the expected pointing accuracy (accuracy only ~1″ due to guide star catalogue). One set of software presently available to handle the reduction of large dithers (which show the effects of geometric distortion) is the Fruchter-Hook code developed for reduction of the Hubble Deep Field images.

For related articles on dither strategies, see the January, 1995 issue of the WFPC2 Space Telescope Analysis Newsletter and the February, 1995 issue of the ST-ECF Newsletter. For more information on dithering, please see “Questions about Dithering WFPC2 Observations” on the “Frequently Asked Questions” page of the WFPC2 WWW area. For detailed information on POS TARGs, please see the report “Dithering: Relationship Between POS TARGs and CCD Rows/Columns” obtainable from the WFPC2 WWW pages or `help@stsci.edu`.

# 7.7 Pointing Accuracy

Some WFPC2 programs have critical target positioning constraints (i.e. the target must be as close as possible to a specified aperture). A sure way to meet such requirements is to include an interactive acquisition. However, INT ACQs are costly in terms of allotted orbits. The only alternative is to provide accurate target coordinates. This section provides some guidance in doing that, and indicates the level of residual error that can be expected.

## 7.7.1 Absolute Pointing Accuracy

We have looked carefully at a sequence of images to assess the absolute pointing performance that HST delivers to WFPC2. The apertures used in the observations studied were either PC1, PC1-FIX, or WF2. The observed positions of stars on WFPC2 images were measured and compared with the proposed coordinates and apertures. Where necessary, coordinate and proper motion errors were accounted for (with the assumption that SAO catalog coordinates are exact - they form the astrometric basis for the guide star coordinate system). The typical residual pointing error is 0.86″, with 1.84″ being the largest error seen. This study did bring out several easy-to-make target coordinate errors (which we corrected in the analysis, but which frequently dominated the pointing error), so we discuss these first.

In a number of cases studied, the proposal coordinates were from the printed version of the Yale Bright Star Catalog. One problem is that the equinox 2000 positions in the BSC are given in the FK4 (Besselian) reference system. The proposal system assumes that equinox 2000 and later coordinates are in the FK5 (Julian) reference frame, and that earlier ones are in the FK4 frame. This can be overridden by specifying B2000 instead of J2000 for the equinox in the proposal. The latest digital version of the BSC (BSC5) is in J2000. The 1950 edition of the SAO catalog is in B1950 (FK4), and a digital version is available for J2000 (FK5). An error of up to 1.5″ can result from assuming BSC positions are J2000 instead of B2000 in the proposal.

Another common problem with target coordinates is that they lack precision. For example, in the BSC, RA is given to the nearest $0.1^s$ and DEC to the nearest arcsecond. This can cause an error of up to 0.75″ in RA and 0.5″ in DEC. The SAO coordinates have higher precision, $0.001^s$ in RA and 0.01″ in DEC, and should be used when possible.

A common error source is not specifying proper motion or specifying it in the wrong units. It is critical to follow the latest version of the proposal instructions on this because the units have changed. Even small proper motions are significant at the resolution of HST images.

Residual pointing errors (after coordinate errors and aperture location changes) range from 0.26″ to 1.84″. The average is 0.93″ and the median is 0.86″. There are no obvious trends in any coordinate system. These are errors which cannot be accounted for by a proposer, being due to guide star position errors, FGS alignment uncertainties, and residual aperture location errors. Using Guide Star Catalog positions may help reduce the error between target and guide stars. Most of the targets used in this study were too bright to have true Guide Star Catalog positions.

In summary, a target with good coordinates (and proper motion) referenced to the SAO catalog can typically be placed within 0.9″ of a specified aperture. However, errors of around 1.5″ occasionally happen.

## 7.7.2 Updates to Aperture / Coordinate Systems

We note that on 11 April, 1994 an update was made to the spacecraft database which tells HST where to place targets relative to the FGSs. This impacted both the location of targets in the WFPC2 field-of-view, and the position reference frame in the image headers. The nominal (or intended) pixel locations of the apertures in the WFPC2 focal plane did not change. Only the (V2,V3) coordinates of the apertures changed, as their locations relative to the FGSs became better known. For example, PC1 and PC1-FIX are designated to be at pixel (420,424.5). Before April 1994, this aperture was thought to be at (V2,V3) =(4.95″,-30.77″), which using the most current information was actually located at pixel (X,Y) = (459.8,377.3). Since April 1994, the aperture in the spacecraft's database has been at (V2,V3) = (1.87″,-30.96″) or, assuming the current best estimate is exactly correct, at (X,Y) = (414.9,428.1). Thus, for the same coordinates and aperture, the pixel position of a target in an image taken before April 1994 could be nearly 3″ different from its position in later images, due to aperture updates. Similar corrections apply to all WFPC2 data taken before this date.

This update also affects the position information placed in the image headers, which maps sky coordinates onto each individual CCD. Observations taken before April 11, 1994 have preliminary plate scales, rotations, and reference pixel locations in their image headers. Thus, the sky coordinates associated with a given pixel will be different for otherwise identical images taken before and after April 11, 1994, due to improvements in the aperture locations. The change is primarily an approximate 3″ shift, as well as a small rotation. There is a 0.8° rotation for WF2, and smaller rotations for the other chips (0.28° in PC1, 0.46° in WF3, and 0.06° in WF4).

The STSDAS task METRIC, which is used to assemble the 4 CCDs into a single image, uses this header information. Hence images taken before 11 April 1994 will not properly assemble into a single image. A new STSDAS task UCHCOORD will be available in the next STSDAS release (early 1996) to update the header group parameters to reflect the improved plate scale, shift, and rotations. This task will supersede the task UCHSCALE in the current STSDAS release, which can update only the scale, *not* the rotation or shift.

We also note, that in April and May 1996 two updates were made to the (V2,V3) coordinate system. This update should be transparent to most observers. The purpose was to remove a slow drift in the position of WFPC2 in the HST focal plane; the largest change was 0.6″. (See Table 3.15 on page 61 for details.)

All the apertures are now thought to be correct to within 0.3″, and updates in the near future should be small.

### 7.7.3 Pointing Repeatability

The Hubble Deep Field (HDF) afforded an opportunity to study the repeatability of pointing over many images and acquisitions of the same field. The pointing appears to have been stable to better than 5 mas accuracy while taking many images of the same field without interruption over several orbits. The accuracy for full-up acquisition of the same field after slewing to other targets appeared to be ~10 mas typically, with occasional 20 mas errors seen. However, a few large errors were seen; in about 1 in 100 acquisitions the FGSs locked-up incorrectly resulting in a ~1″ error.

Other programs report similar 3 mas pointing accuracy if simple re-acquisitions are done between orbits. Approximately once per day a “full-up” acquisition is usually required (for engineering reasons) where the dominant FGS is fixed in position, but the sub-dominant FGS performs a spiral search for the guidestar and tracks wherever the star is found. On rare occasion these full-up acquisitions produce position errors of several hundred mas, and field rotations of up to ~0.1°, relative to previous images of the same field. This may impact long sequences of exposures requiring half a day or more to execute.

### 7.7.4 Tracking Modes

Three guiding modes are available: (1) Gyro Hold, (2) Coarse Track, and (3) Fine Lock. Fine Lock (PCS MODE FINE) is used by default, since use of Coarse Track may be harmful to the Fine Guidance Sensors. Use of Gyro Hold (PCS MODE GYRO) is not generally recommended, even for snapshot (SNAP) observations, since the pointing accuracy is only 14″. Also the drift rate is 0.0014″ $s^{-1}$ so exposures >100s can result in smeared images. However, if the reduced pointing accuracy can be tolerated, and the exposures are only a few seconds or less, Gyro Hold can give a significant savings in the target acquisition overhead time.

## 7.8 CCD Position and Orientation on Sky

During observation the target is placed at the aperture (PC1, WF2, WFALL, etc.) specified on the Phase II proposal. Locations of the principal apertures are shown in Figure 7.8 (Table 3.14 on page 61 gives a complete list of apertures; the (V2,V3) system here is post 1996 day 127).

The POS TARG special requirement can be used when a position offset is needed. The target is positioned with offset “POS TARG x,y”, measured in arcseconds, from the specified aperture. The approximate directions (within 1°) of the POS TARG offsets are shown in Figure 7.8. The exact directions of the offsets are parallel to the rows and columns of the CCD on which the aperture is specified. There are small rotations (few tenths of a degree) between the CCDs. (For detailed information see “Dithering: Relationship Between POS TARGs and CCD Rows/Columns” obtainable from the WFPC2 WWW pages or help@stsci.edu.)

It is often useful to explicitly specify the desired rotation of the WFPC2 field-of-view on the sky. This is specified in the Phase II proposal using the ORIENT special requirement. It is defined as the PA (measured from North through East) of the +U3 axis on the sky. Figure 7.8 shows the CCD orientation and aperture locations relative to the U3 axis. In effect, the sequence of events is to first move the target to the desired aperture, then offset by any specified POS TARG from the aperture, and finally to rotate the target “in place” on the CCDs to the desired ORIENT.

Observers should try to specify all possible ORIENTs which would give the desired data, since having a range of values, or several ranges, will make the observation much easier to schedule. Often two ORIENTs separated by 180° will both give useful data. Sometimes ORIENTs separated by 90° will also give similar data.

The ORIENT for any observation can be computed as follows:

1. Obtain the Position Angle (PA) of the source axis on the sky, measured in the standard way, North through East.
2. Look at Figure 7.8 and decide what angle you want, measured clockwise, from the +U3 axis to the source axis.
3. Sum the angles in steps 1 and 2.
4. ORIENT must be between 0° and 360°, so subtract 360°, if necessary. The result is the ORIENT you should specify on the proposal.

Another way to select the ORIENT, is to place Figure 7.8 on an image of the target, shift and rotate to get the desired alignment, and then simply measure the position angle of the +U3 axis relative to North.

Note that the +V3 axis is quite different from the +U3 axis. They are exactly parallel, but oppositely directed. The +U3 axis is used for specifying orientation (ORIENT) in the proposal, while the +V3 axis is used in the data headers to indicate field orientation. Data header keyword PA_V3 gives the position angle of the +V3 axis on the sky.

We now give two examples of how the POS TARG and ORIENT special requirements might be used. The first example (Figure 7.9) shows placement of a 100″ long jet along the CCD diagonals in PC1 and WF3 (i.e. along the -U3 direction). The coordinates of the nucleus are given on the proposal. Aperture PC1 together with POS TARG +10, +10 are used to place the nucleus near the outer corner of PC1. We want to rotate the WFPC2 field-of-view about the nucleus so the jet is diagonal on PC1 and WF3. We can thus compute the desired orientation as

**Figure 7.8:** ORIENT Definition, Aperture Positions, and CCD Alignments. “FIX” apertures are in same locations, unless otherwise indicated. Dashed lines show vignetted regions along CCD boundaries. Short lines and “X”s in outer CCD corners indicate directions of bloom and OTA diffraction spikes, respectively. Origin of the (V2, V3) system is at the origin of the plot axes, with V2 and V3 exactly along diagonal lines as marked. POS TARGs are offsets measured from the aperture specified on the proposal (PC1, WF2, WFALL, etc.); their directions are as indicated. CCDs have pixel (1,1) where the four CCDs overlap.

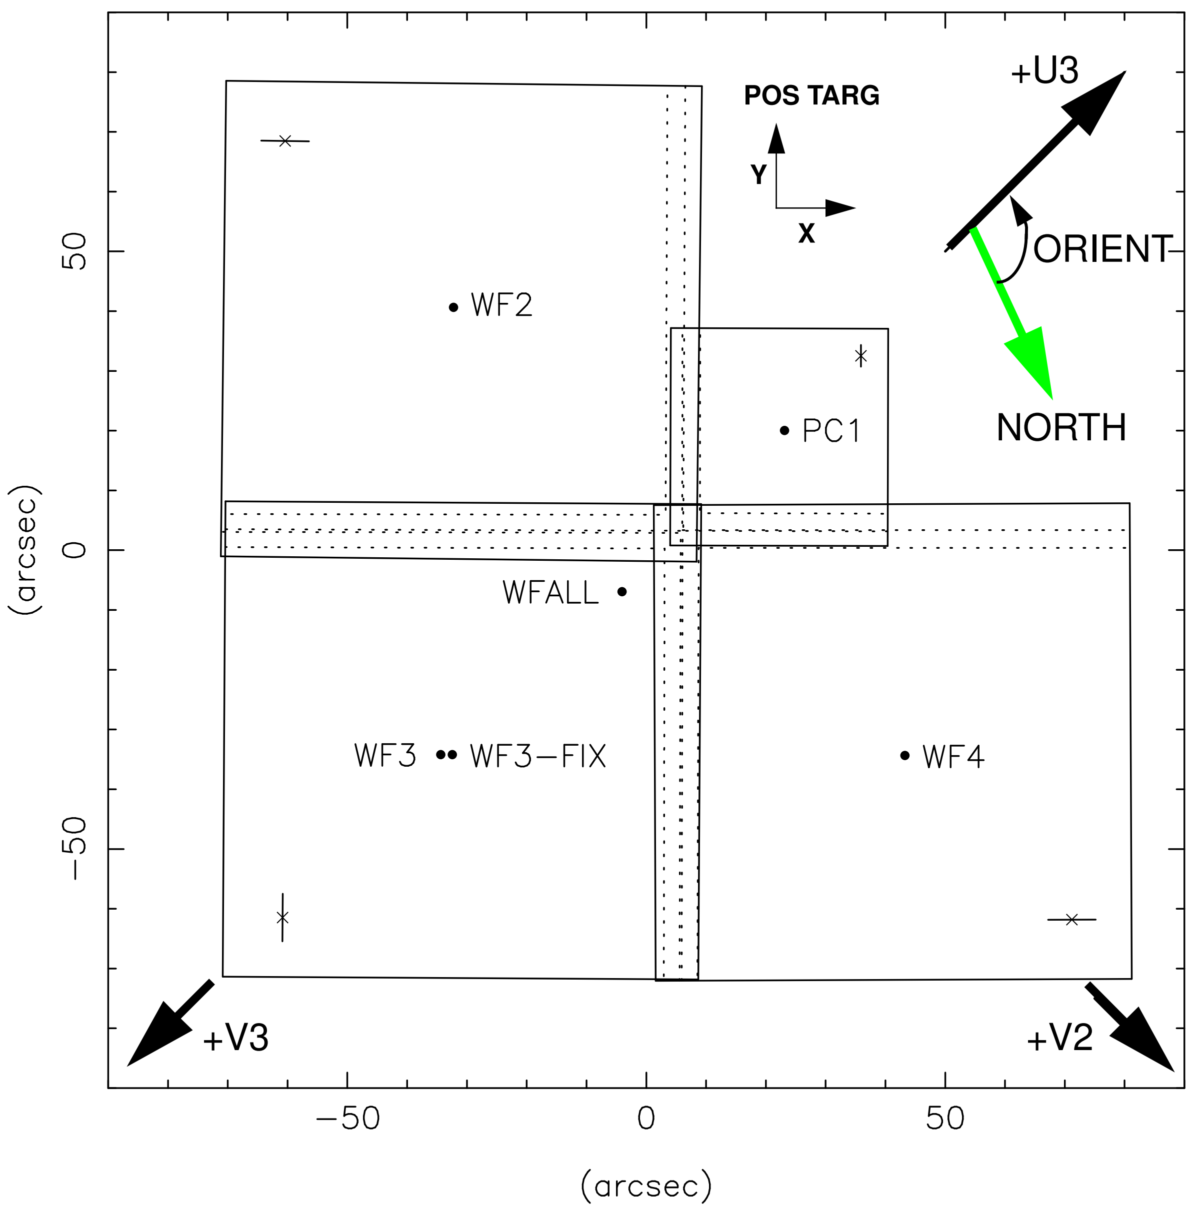


ORIENT is defined as the Position Angle of the +U3 Axis on the Sky.

ORIENT = (source PA on sky)

+ (desired source angle in field-of-view measured CW from +U3 axis)

$= 290^\circ + 180^\circ = 470^\circ - 360^\circ = 110^\circ$

On the Phase II proposal we would allow some range in the angle (to ease scheduling), hence "ORIENT 105D TO 115D" might be specified.

The second example (Figure 7.10) shows placement of a galaxy across WF2 and WF3, with the nucleus on WF3 safely away from the vignetted region. Aperture WF3 together with POS TARG +20, 0 is used to place the nucleus near the outer edge of WF3. We want to rotate the WFPC2 field-of-view about the nucleus so the galaxy major axis is across WF2 and WF3. We can thus compute the desired orientation as

ORIENT = (source PA on sky)

+ (desired source angle in field-of-view measured CW from +U3 axis)

$= 60^\circ + 315^\circ = 375^\circ - 360^\circ = 15^\circ$

On the Phase II proposal we would again allow some range in the angle (to ease scheduling), hence "ORIENT 5D TO 25D" would be specified. Note that "ORIENT 185D TO 205D" is also feasible, and should be indicated in the visit level comments. Note also, that WF3 and WF4 could be used with either "ORIENT 95D TO 115D" or "ORIENT 275D TO 295D"

## 7.9 Polarization Observations

Polarization observations require three or more images with the polarizing filter spanning a large range of position angles on the sky. For WFPC2, this may be achieved by using different quads of the polarizing filter (each quad being oriented 45° to the others), by rotating the spacecraft though different angles, or by a combination of these methods. Rotating the spacecraft through use of ORIENTs provides the simplest calibration, as a single polaroid can be used for all images. However, in practice, it will be the most difficult method to schedule. For further information see Biretta and Sparks (1995, WFPC2 ISR 95-01).

We note that WFPC2 has significant instrumental polarizations which will make measurements on targets with less than 3% polarization difficult. The pick-off mirror introduces about 6% instrumental polarization. Furthermore, the pair of mirrors in the calibration channel, which is used to generate the polarizer flat fields, introduces ~12% polarization. In principle these effects can be calibrated out, but this has yet to be demonstrated.

The polarizers are most effective in the range from 3000Å to 6500Å; this corresponds roughly to filters in the range F255W to F675W. At shorter wavelengths the transmission decreases sharply, and at longer wavelengths they cease to polarize the incoming light.

**Figure 7.9:** Example of ORIENT and POS TARG Selection. (A) A jet at PA=290° is observed using PC1 and WF3; the position of the nucleus is used for the target position. (B) The aperture is specified as "PC1" and the nucleus is placed near the outer corner of PC1 using "POS TARG +10,+10." To place the jet across PC1 and WF3 "ORIENT 105D TO 115D" is specified.

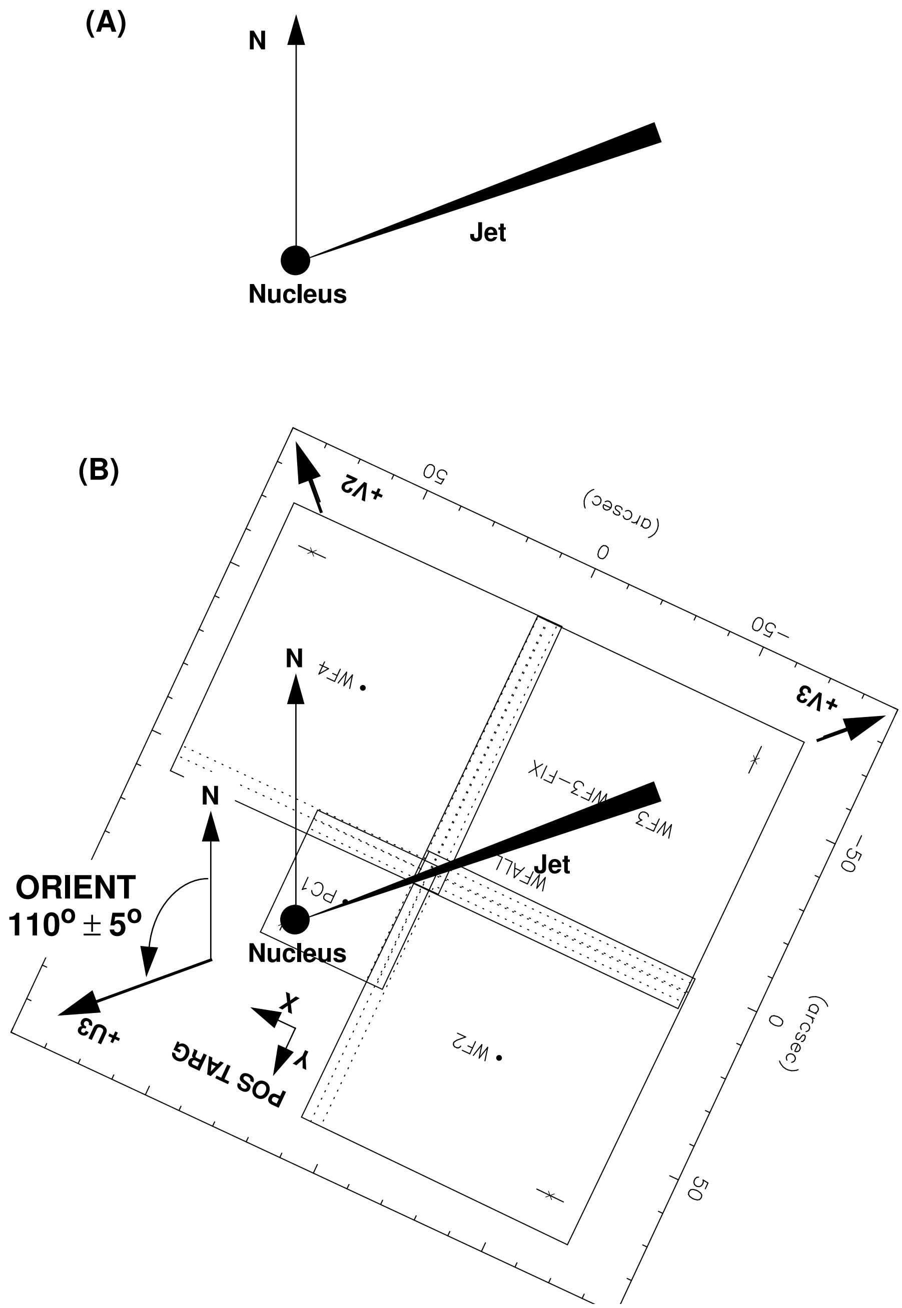

**Figure 7.10:** Example of ORIENT and POS TARG Selection. (A) A galaxy with major axis at PA=60˚ is to be placed across WF2 and WF3. (B) The aperture is specified as “WF3” and the nucleus is placed near the outer edge of WF3 using “POS TARG +20,0.” To place major axis across WF2 and WF3 “ORIENT 5D TO 25D” is specified. Note that “ORIENT 185D TO 205D” is also feasible.

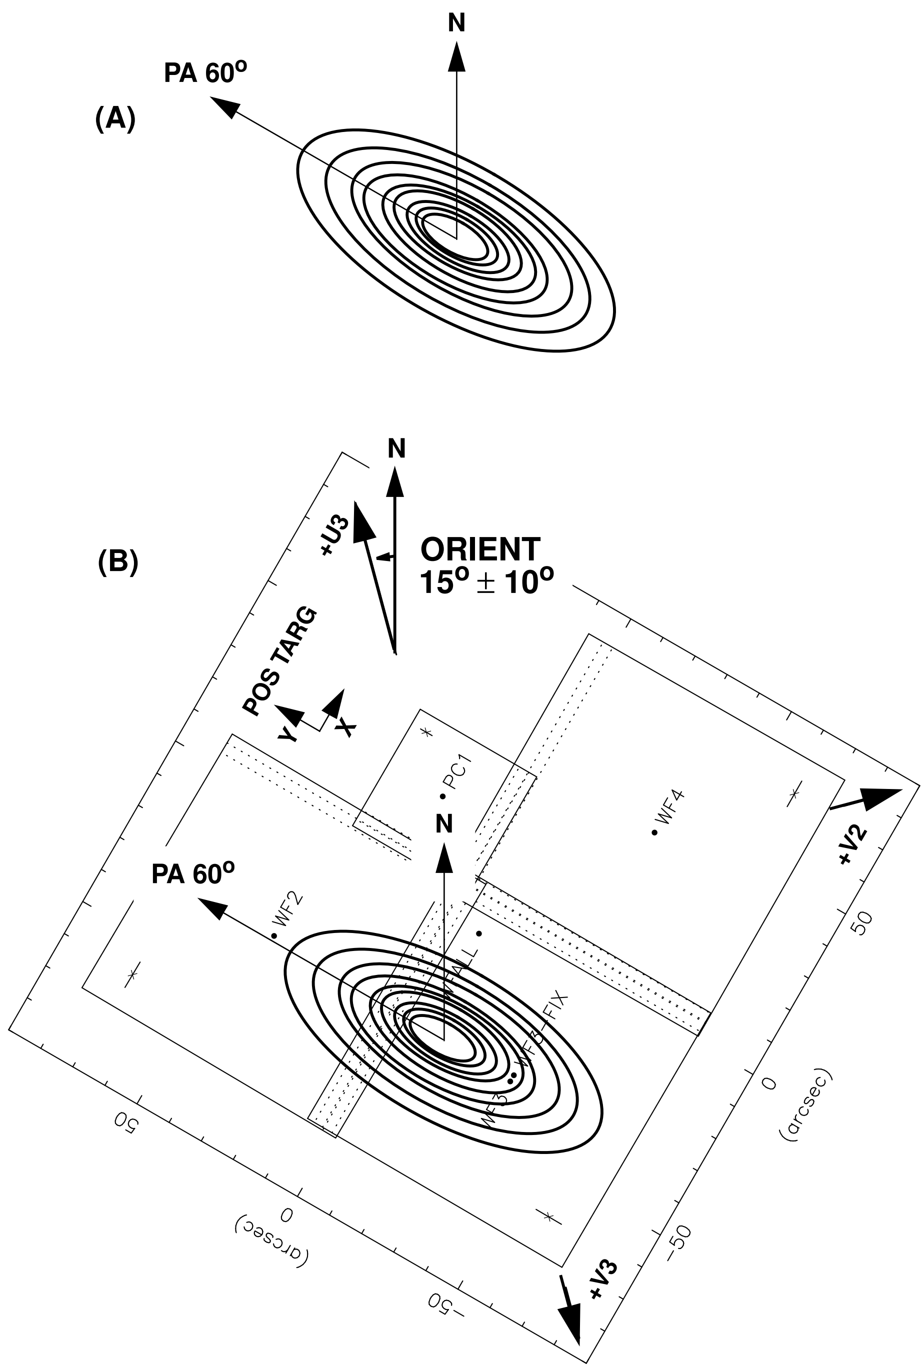

# 7.10 Observing with Linear Ramp Filters

The Linear Ramp Filters provide a narrow band ( $\Delta\lambda/\lambda = 0.013$ ) imaging capability which is continuously tunable from 3710Å to 9762Å. These are essentially a collection of narrow band interference filters whose central wavelength varies with position on the filter glass. The filter and aperture should be specified as LRF on the Phase II proposal, and the desired central wavelength should also be specified. The HST scheduling software will then select the target position so as to provide the desired wavelength.

Note that it is not possible to choose between PC1 and WFC for the LRFs; one must use whatever CCD is automatically assigned by the scheduling software. If it is necessary to know which CCD will be used, observers can consult Table 3.7 on page 48 or Table 3.8 on page 49, or use the on-line LRF calculator tool on the WFPC2 WWW pages.

It is possible to use POS TARGs with LRF observations; the offsets are made from the default pointing for the specified wavelength. Observers should be mindful that the unvignetted field-of-view has a minimum size of ~10″ in diameter, so that only small POS TARGs (<4″) should be used.

While it is recommended that observers assume a 10″ diameter field-of-view when using the LRFs, larger elongated (e.g. 15″ x 10″) targets can be sometimes be accommodated by placing the target major axis along the direction of the wavelength variation on the filter. This will result in a small reduction in throughput (i.e. small central wavelength offset) at the outer edges of the target. However, placing targets outside the central 10″ of each ramp is strongly discouraged; outside the central 10″ width the light will pass through more than one ramp segment, hence mixing light from different wavelengths, and making the data very difficult to calibrate. (See page 40 for further details on LRFs.)

A common situation is one in which observers desire to make observations through an LRF filter, and then repeat the observation in a standard broad or narrow band filter at the same position on the CCD. The LRF Calculator Tool, available on the WFPC2 WWW pages, will tell observers the aperture (PC1-FIX, WF2-FIX, etc.) and POS TARG for any wavelength setting of the LRFs. Observers merely need to use this same aperture and POS TARG for the exposure through the other filter. If it is necessary to calculate the POS TARG manually, one can do this using the information in Table 3.7, Table 3.8, Table 3.14, and Figure 7.8. For example an LRF observation at 5034Å would be made on WF2 at pixel (673.4, 235.7) (from interpolation by wavelength between X1 and X2, and between Y1 and Y2 in Table 3.7). These offsets are referred to the WF2-FIX aperture which is located (Table 3.14) at pixel (423.5,414). From Figure 7.8 we can deduce that pixel X direction is parallel to POS TARG "+Y" on WF2, and that pixel Y direction runs in the POS TARG "-X" direction. Using the pixel scale in section 5.10, "Optical Distortion", on page 103, we have

POS TARG "X" = -0.09961 (235.7-414) = 17.76″, and

POS TARG "Y" = 0.09961 (673.4-423.5) = 24.89″, hence POS TARG=+17.76,+24.89 would be requested for the non-LRF exposure.

## 7.11 Emission Line Observations of Galaxy Nuclei

Saturation is a common problem for narrow band filter images of galaxy nuclei. Often the surface brightness of the emission line is estimated from ground-based images with 1″ resolution; sometimes line fluxes are quoted for apertures several arcseconds in radius. However, at HST resolution, much of this flux may occur in a single unresolved spot at the galaxy nucleus, thus leading to saturated images. Observers should be cautious and consider this possibility when estimating exposure times.