Instrument Science Report WFC3 2015-08

A Study of the Time Variability of the PSF in F606W Images taken with WFC3/UVIS

Jay Anderson, Matthew Bourque, Kailash Sahu,
Elena Sabbi, Alex Viana
May 28, 2015

ABSTRACT

This document is the first in a series designed to study how the F606W WFC3/UVIS Point-Spread Function (PSF) changes with focus. We have extracted star images from a large number of WFC3/UVIS exposures and use these images to explore the variation of the WFC3/UVIS PSF over time — both the long-term changes due to focus drifts and resets and the short-term changes related to breathing. We find that the PSF tends to vary in regular ways such that we are able to develop an empirical model to describe this variation. Future documents will aim to construct tools that can be used to provide accurate PSFs for individual exposures.

1. INTRODUCTION

The point-spread function (PSF) of the telescope modulates the scene that the telescope is able to deliver to the observer. For objects that are larger than this fundamental resolution element, an intimate knowledge of the PSF is not necessary to do high-precision science. However, many astronomical studies can be pursued only when we have an accurate understanding of a detector's point-spread function (PSF). For example, astrometry and photometry of point sources, bright and faint, cannot be done with high precision without PSF-fitting. In a similar vein, weak lensing and other studies of objects that are close to the resolution limit can be very dependent on the fidelity of the PSF model.

Unfortunately, even though accurate PSF models are critical to many astronomical studies, there are many reasons that very few published papers make use of good PSF models. For one thing, it is hard to construct good PSF models. The WFC3/UVIS detectors are mildly undersampled, which means that an accurate PSF can only be constructed from a dithered set of data, and one must take exquisite care to

accurately represent the sub-sampled nature of the PSF. In addition, the PSF changes with position across the field, both due to variations in optical distortion and variations in the thickness of the detector (related to charge diffusion). Furthermore, the PSF also changes over time due to secular and breathing-related changes in instrument focus. All of these issues make it difficult to have perfect knowledge of the PSF in an image *a priori*.

Even when accurate PSF models are available, it is hard to use them to do science. PSF models are most accurate in the individual flat-fielded frames (the `_flt` images), since the pixel values in these images are the only true and direct constraints[1] that we have on the astronomical scene. Even so, because of the undersampling and detector artifacts, a single exposure is not able to contain all the information the telescope can collect about the scene. It is necessary to dither the scene by whole pixels and fractional pixels in order to fully constrain the astronomical scene that has been delivered to the detector. Unfortunately, the large distortion that is present in HST's detectors makes it difficult to interrelate the `_flt` pixels in different dithers. For this reason, many users make use of the `Drizzle` software suite, which is designed to combine the individual distorted and undersampled exposures into a single composite image that has better sampling and no distortion. This resampling process can be done in a way that preserves flux, but it is very hard to perform this operation without introducing irregularities in the sampling or introducing correlations among the output pixels. All this means that it is hard to do high-accuracy PSF analysis on the `drizzle` product.

For all of the above reasons, the astronomical community has been slow to make use of accurate PSFs. In some ways, the quest for using PSFs to do better science is at an impasse. Even if people had access to better PSFs, they could not use them, since there is no easy way to relate `_flt` pixels with each other and with the scene to use them as well-posed constraints. At the same time, even if there existed a tool to properly place each `_flt`-pixel constraint in context, there do not yet exist PSF models to make use of them.

This document is an initial effort to short-circuit this chicken-and-egg problem by characterizing some static "library" PSF models and exploring how we might adapt them to account for breathing and other focus-related variations. An initial attempt to address the `flt`-pixel-mapping aspects of the challenge can be found in the recent ISR by Anderson (WFC3/ISR 2014-23). The ISR describes a software routine, named `hst2bundle`, that assembles together into a single bundle all of the pixels in the vicinity of a source and provides an easy way to relate them to each other and to the scene so that users can properly treat each pixel as a single constraint on the scene and can make use of all of the constraints from all the exposures in a single simultaneous fit.

This document is organized as follows. We first describe the static "library" PSF models that serve as the starting point for this exploration. We then describe the database that has been extracted from the entire WFC3/UVIS archive of external F606W observations. We use this database to study the variation of the PSF relative to the static "library" models in the upper-left corner of the detector. This corner has previously been identified as being most sensitive to focus and breathing. We identify eight different

[1] By "constraint" here we simply mean that a pixel in a raw value image means that a certain number of electrons were recorded within the boundaries of a particular pixel over a certain amount of time.

focus levels from the observed star images in the upper-left corner. This allows us to measure a focus-dependent PSF model for the entire detector. Finally, we take a peek at how focus variations may affect the detector as a whole.

2. THE STATIC "LIBRARY" PSF MODEL

The static "library" PSFs that we start with here were constructed from the dithered set of images taken in program GO-11911 of the core of Omega Centauri during the commissioning period for the WFC3/UVIS detector. A recent report by Kozhurina-Platais & Anderson (ISR WFC3 2015-02) describes the field at the center of this great cluster and shows why it is the perfect place (in terms of density and flatness of the star distribution) to explore issues of distortion and PSFs.

The procedure used to extract PSFs from the images is very similar to that described in Anderson & King 2006 (ISR ACS 2006-01, AK06) for the WFC camera of ACS. In that document, the spatial variability of the PSF was treated by solving for fiducial PSFs at a 9×5 array of locations on the corners, edges and across each chip. PSFs in between these fiducial locations were constructed by means of simple linear interpolation of the neighboring fiducial PSFs. The 9×5 array for the ACS's WFC was not perfectly evenly spaced across the detector, since because of charge diffusion, there were some locations where the PSF shape was found to be at an extremum (on account of a thin spot in the CCD, which resulted in less charge diffusion and a sharper PSF, see ACS ISR 03-06 by Krist), so we shifted some of the fiducial points to better represent the variation. We explored whether the WFC3/UVIS PSF showed any similar behavior. The "happy bunny" (described in WFC3 ISR 2013-011 by Sabbi and Bellini) is one such location on the WFC3/UVIS chip and we found that if we represent the WFC3/UVIS PSF with a 7×4 array for each chip, we just happened to get a fiducial point at the right location to deal with this feature.

The reader should consult Anderson & King (2000 PASP 112 1360, AK00) and AK06 for a more detailed description of how to construct an undersampled PSF. We will give only a brief summary here. The basic purpose of the PSF is to tell us what fraction of a star's light will land in a pixel that is offset by $(\Delta x,\Delta y)$ from the center of a star. This can be easily represented as a two-dimensional scalar function $\psi(\Delta x,\Delta y)$. The value of $\psi(0,0)$ tells us what fraction of the star's light lands in its central pixel if the star is centered on that pixel. It is an interesting number to have in hand. For detectors that are considerably undersampled, such as WFPC2, WFC3/IR, and Kepler, this can be anywhere from 0.30 to 0.65 (or larger). For moderately undersampled ACS/WFC this tends to be between 0.18 and 0.22 and for WFC3/UVIS it tends to be 0.16 to 0.19.

The challenge of modeling the PSF is two-fold. We must find a way to accurately parameterize this two-dimensional function, and we must determine optimal values for the parameters. It turns out that the representation developed for WFPC2 by AK00 has proven to be sufficient for WFPC2, ACS's WFC and HRC, and is also adequate for WFC3/UVIS and WFC3/IR. To represent the PSF, we specify its value at a 4×-supersampled array of grid-points across its two-dimensional domain. Unsaturated stars in HST images have measurable signal out to about 10 pixels, so we represent the PSF with Ψ_{IJ}, an array of 101×101 grid points that cover the PSF out to a radius of about 12.5 pixels. The grid-points are spaced every quarter pixel and tell us directly the fraction of a star's light that will land in a pixel that is offset by the corresponding distance from a star's center.

Since stars can land anywhere within a pixel, we of course need to determine a value for the PSF at any location within its domain, not just the points on the fiducial array. To do this we use bi-cubic interpolation, since bi-linear interpolation is not able to preserve the structure inherent in these sharp PSFs. The light in each of pixel [i,j] in the vicinity of a star with a center at (x_{STAR},y_{STAR}) of flux F_{STAR} on background S_{STAR} can be modeled by:

$$P_{ij} = F_{STAR} \times \psi(\Delta x,\Delta y) + S_{STAR} = F_{STAR} \times \psi(i - x_{STAR},j - y_{STAR}) + S_{STAR}.$$

To solve for the PSFs, we must determine the best possible value for each grid point, Ψ_{IJ}. Since the value of a PSF grid point represents the fraction of a star's light that will land in a pixel at the corresponding offset, then to determine an accurate value for the array Ψ_{IJ}, we simply need to identify observed pixels at the right offset for which we know the flux of the star and the background. In this case,

$$\psi(\Delta x,\Delta y) = (P_{ij} - S_{STAR}) / F_{STAR}$$

is an estimate of the value of the PSF at the offset corresponding to the location of the pixel. As with the evaluation of the PSF where we did not have grid-points to sample every possible location of a star, here we do not have observations of the PSF at the exact location of the gridpoints. As such, the construction of the PSF must be done by iteration.

We start by grouping together all the observations near a particular (I,J) grid point and assign the PSF a robust-average value from these observed estimates. We then assemble all these estimates into a model and impose some simple smoothness constraints (see AK00) in order to ensure that the PSF will not vary chaotically from gridpoint to gridpoint. Next, we iterate, adjusting the PSF according to the difference between the observation and the current model. This process converges relatively quickly.

It is easy to generalize this procedure to account for the fact that the PSF varies spatially across the detector. We end up solving for the PSF at each of the 7×8 fiducial points (α,β), $\Psi_{IJ;\alpha\beta}$, using the stars observed in its vicinity and do a four-dimensional interpolation to arrive at a model for each pixel in each star, using $\psi(\Delta x,\Delta y;x,y)$. The interpolation in $(\Delta x,\Delta y)$ is bi-cubic and the interpolation in (x,y) is bi-linear.

The PSFs are stored in simple fits images. The PSF at each of the 7×8 fiducial locations (here, we are working in a two-chip system) is a 101×101 image, and these are all stored in a single fits image with 7×8 = 56 extensions. The header information in the image tells where each fiducial PSF is located.

Figure 1.1 shows the PSFs arrayed roughly at their fiducial locations, though the edge and corner PSFs are located at the edges and corners of the detector. The left half of the PSF shows the PSFs themselves. It is hard to find a grayscale stretch where they do not all look nearly identical. The right half shows the difference between each PSF and the average. Here, it is easy to see the spatial variability. In general, the PSF becomes duller as we go towards the outskirts of the detector, with the exception being the region near the "happy" bunny in the lower right, where the PSF is sharper.

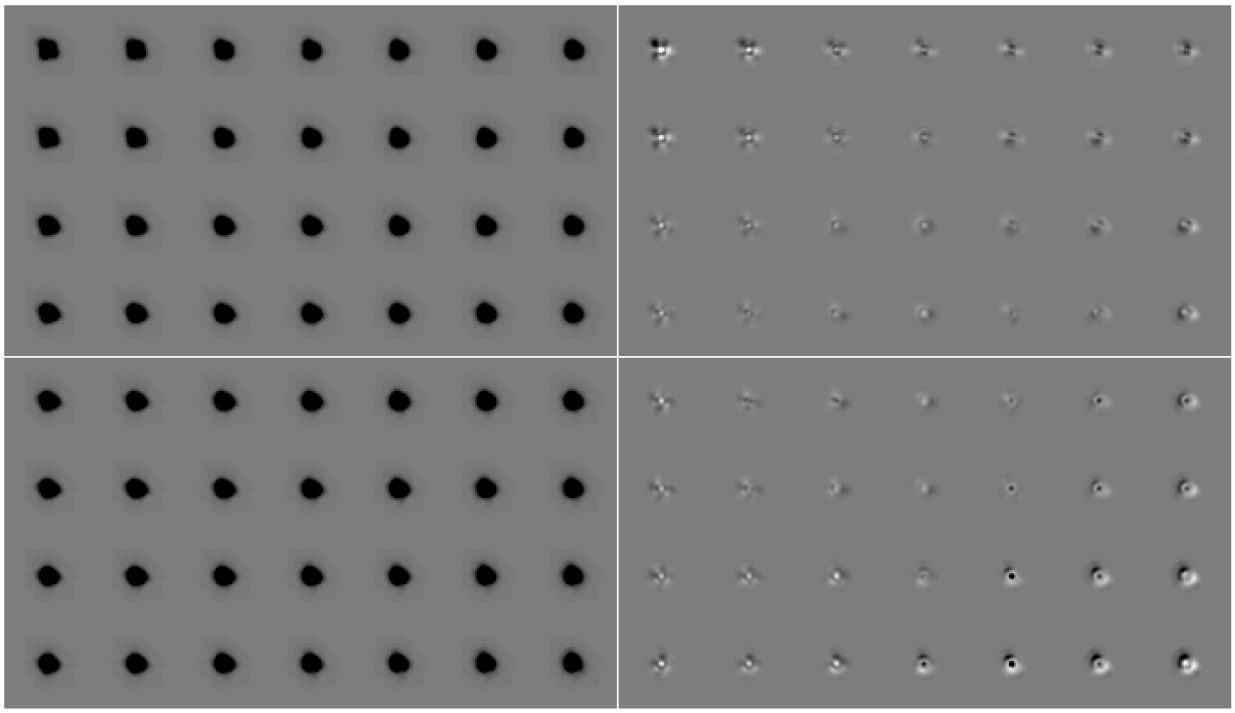

Figure 1.1: The 7×4 array of fiducial F606W PSFs for the two WFC3/UVIS chips. The left half of the figure shows the actual PSF and the right the residual with respect to the average. The array of PSFs covers each of the 4096×2048 chips with the fiducial PSFs at the corners and edges and a spacing of 682.67 pixels. (Although each WFC3/UVIS chip reads out 2051 pixels in the vertical direction, the top three pixels have different sizes and are not well calibrated.) The placement of fiducial points at the corners and edges ensures that we can construct all PSFs by interpolation of existing models; no extrapolation is necessary. In the grayscale shown, dark corresponds to positive flux and white to negative.

3. A DATABASE OF STAR IMAGES

The PSFs constructed above were extracted from images taken at a single time and focus. In reality, although the HST PSF is much more stable than ground-based PSFs, breathing and secular drifts do cause the focus to change over time. Therefore, it is good to determine how well the above "library" PSF can be used to represent the PSF at any given time, and perhaps whether one can add time or "focus" dimensionality to the model.

To this end, we decided to go through the entire WFC3/UVIS archive of F606W observations in order to collect as many images of stars as possible. We constructed a software program, based on the publicly available routine `img2xym_wfc3uv`[2], to go through each exposure in the archive and identify every object that met the following finding criteria: its brightest pixel is [1] below the saturation threshold (65,000 counts) and is [2] more than five pixels away from a brighter pixel, and [3] it has more than 10,000 counts in its central 3×3 pixels, so that it will have a S/N greater than 100. For each of these

[2] http://www.stsci.edu/~jayander/WFC3/WFC3UV/PSFs

objects, the routine measured a sky value (s) from the pixels between a radius of 8 and 12 pixels and fit the "library" PSF described above to its central 5×5 pixels and determined a position (x, y), a flux (z), and a quality-of-fit metric (q), which corresponds to the absolute value of the residual between the observed pixel values and the best-fit model divided by the total flux of the star.

In addition to fitting for these basic parameters, we also recorded in a database a raster of 21×21 pixels centered on the star's brightest pixel, along with the image name, the central pixel location, and the MJD time of the observation. In all, our database contains 5,013,607 star images taken across the lifetime of the WFC3/UVIS detector (as of 99 October 2014). This database has been converted into a simple 3-dimensional fits file with dimensions 21×21×5013607, to make it easy to access the data.

There are of course many ways to break down and examine this database. To give a sense of what it contains, **Figure 2** below shows the observed brightness of stars as a function of MJD. Each dot is a star. It is clear that WFC3/UVIS has been observing stars throughout its lifetime.

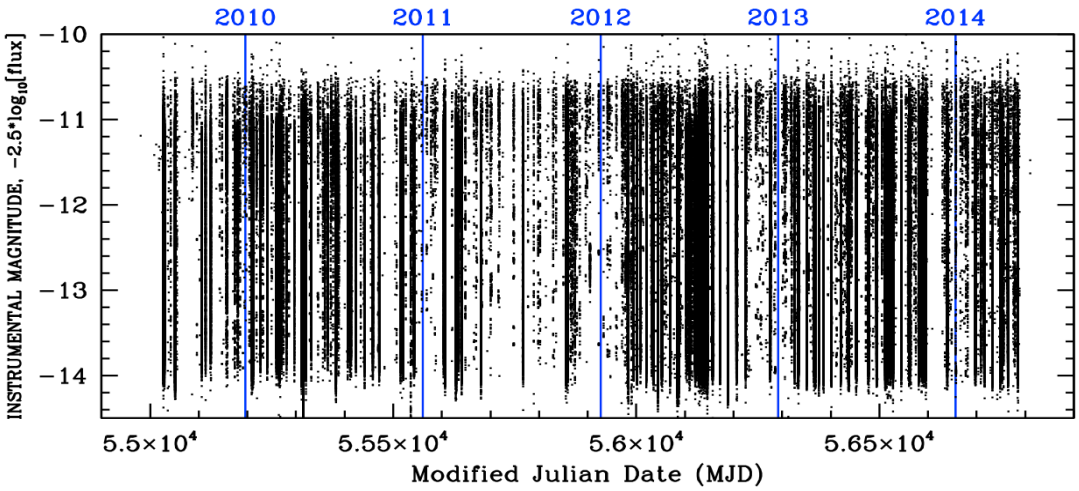

Figure 2: This shows the distribution of star images we have as a function of observation date and instrumental magnitude. Each dot is a bright star observed at a particular place in a particular image.

Figure 3 gives the distribution of quality of fit as a function of instrumental magnitude. Because we selected sources for the database based simply on their brightness and isolation, the list contains some real stars, some bright galaxies, and some cosmic rays. The dotted line separates the stars from the non-stars. We find that roughly 97.5% of the sources identified are true point-like objects based on their resemblance to the PSF. It is clear that all the stars we have selected are bright enough to have a clear PSF shape. If we had allowed fainter objects, then the quality-of-fit metric would suffer as the star got fainter.

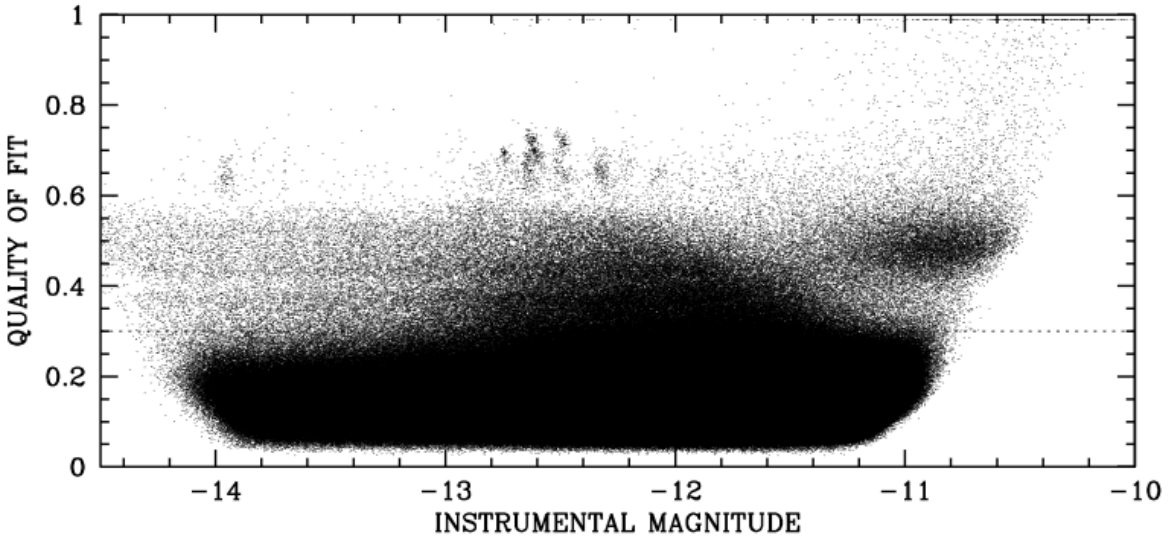

Figure 3: The quality of fit metric as a function of instrumental magnitude for all extracted star images. Images with a quality-of-fit metric greater than 0.25 are likely cosmic rays or otherwise not point sources.

Figure 4 shows the distribution of sources (and qualified stars in solid) in terms of instrumental magnitude. An instrumental magnitude of −10 corresponds to 10,000 total electrons under the PSF and therefore to a signal-to-noise of about 100. There are certainly more faint stars than bright ones, but there are a good number of stars at all brightness levels.

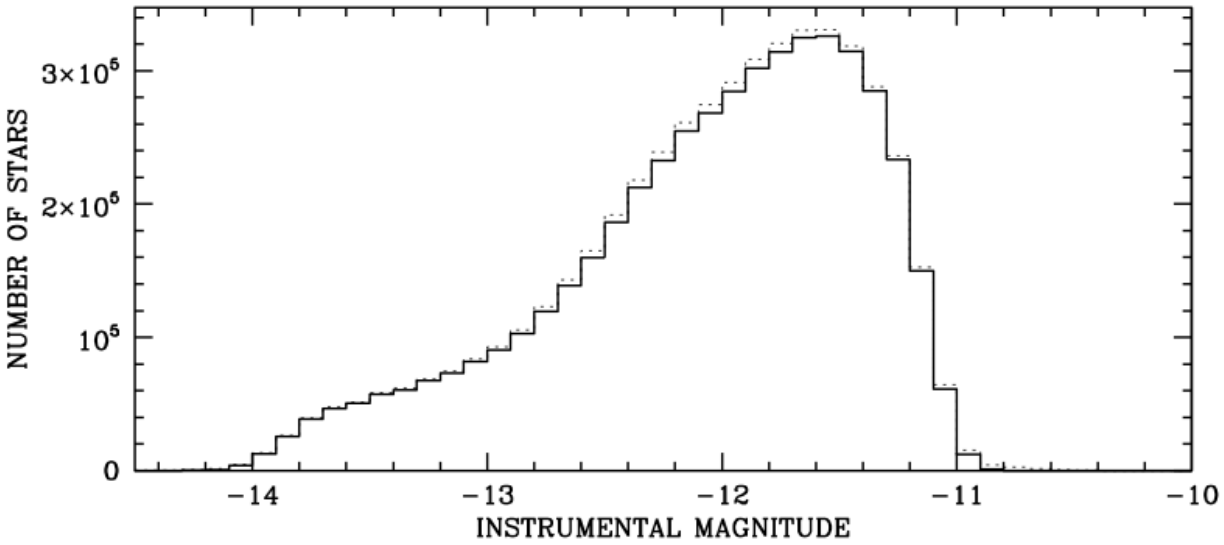

Figure 4: Histogram showing the distribution of instrumental magnitude for the extracted star images.

Finally, **Figure 5** shows the distribution of stars per image. In the top plot, each exposure is a dot and it is plotted according to its observed MJD on the vertical axis and the number of stars it contains on the horizontal axis. The bottom plot distills the top plot into a histogram. There is a bimodal distribution. Typical exposures have between 5 and 20 stars, but about 200 exposures have about 10,000 stars in them. These are likely observations of the bulge or star clusters. Several of these are observations of the Omega Centauri calibration field.

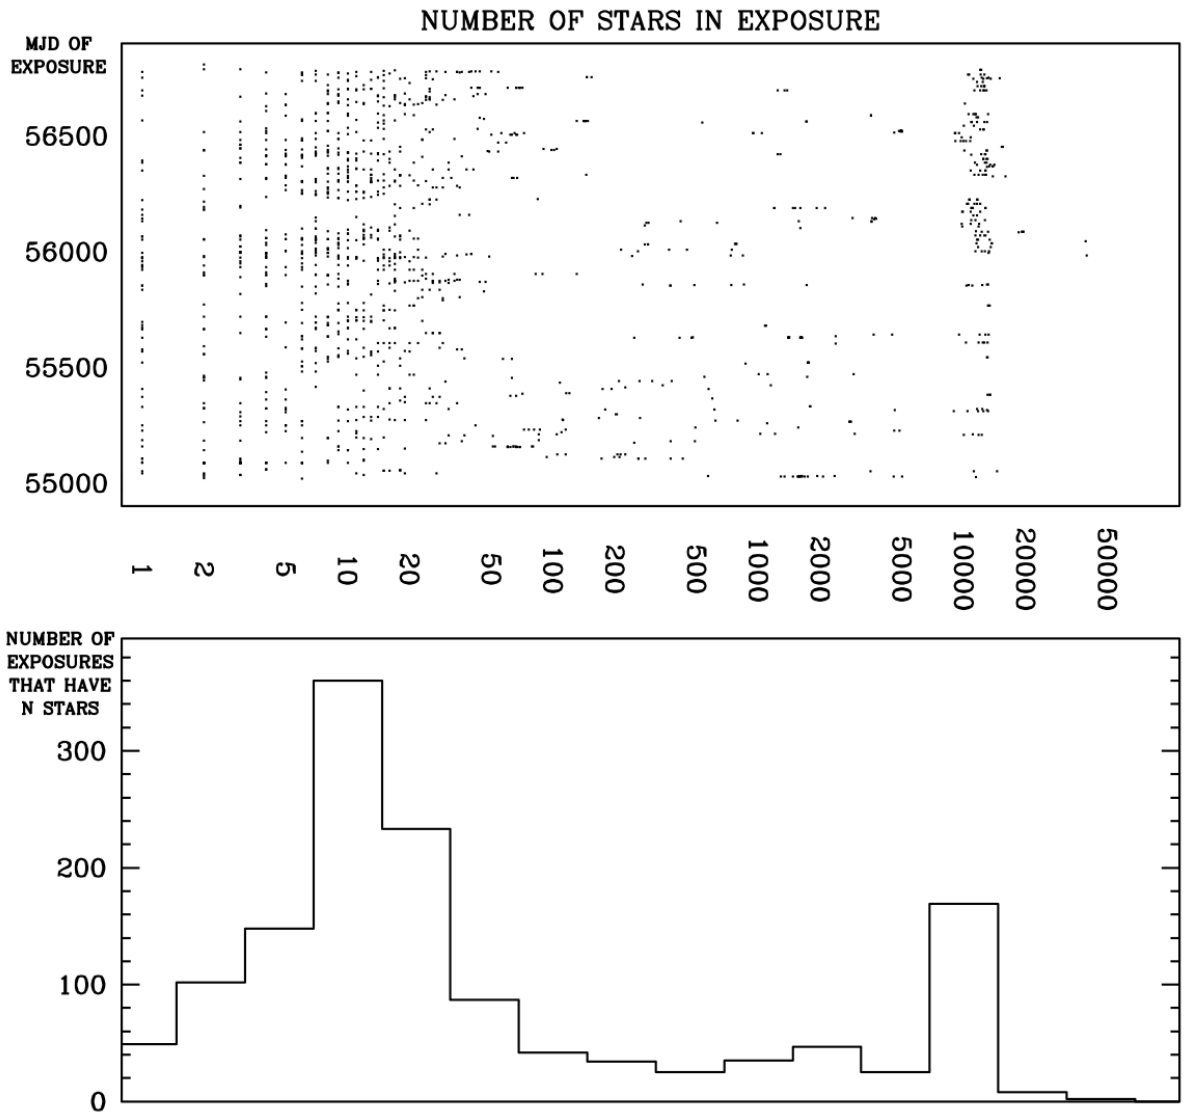

Figure 5: The number of star images observed in each exposure: (top) Shown against observation date on the vertical axis and (bottom) in histogram form.

4. THE PSF IN THE MOST SENSITIVE CORNER

Studies by Sabbi & Bellini (2013) have shown that the upper left corner of the UVIS detector (the far corner of Amplifier A in the top chip, UVIS1) is the most sensitive to breathing, as it is nominally out of focus relative to the rest of the detector. Out of the 5,013,607 stars in the database, 134,887 are in the 682×682 pixels of the sensitive corner (which we will refer to as 'A-corner'). In this section, we will focus on these stars.

One thing that happens when HST goes out of focus is that the fraction of light in the core goes down. But it is well known (see Chapter 6 of WFC3 Instrument Handbook; WFC3 ISR 2013-11 by Sabbi and Bellini) that the PSF also becomes astigmatic. On one side of focus there is an elongation along the position angle (PA[3]) of 45°(corresponding to the even quadrants, shown by the A's in the illustration below), and on the other side of focus the elongation is along the 135° axis (corresponding to the odd quadrants, shown by B's in the illustration below)[4]. We will exploit this phenomenon in an attempt to characterize the focus in each exposure by the images of the stars taken in it.

We read in the raster for each of these 134,887 corner stars from the database and fit each with the PSF as described above. The PSF fit provides an estimate of the fraction of flux that should land in each of the star's central pixels. In order to measure the asymmetry, we first characterize each pixel in terms of whether it is larger on one side of focus or the other. In the following table, pixels labeled "A" should receive more flux on one side of focus, and pixels labeled "B" should receive more flux on the other side of focus. The star is centered somewhere within pixel "C", which is also the brightest pixel.

A	A		B	B
A	A		B	B
		C		
B	B		A	A
B	B		A	A

For each star image, we determined the quantity f_A, the fraction of the star's observed flux in the "A" pixels and also determined the quantity ψ_A, the fraction of flux we would expect to find in these pixels based on the PSF model and the star's fitted position. The difference $f_A - \psi_A$ is a measure of the excess or deficit of flux in the odd-quadrant moment. Similarly, we determine $f_B - \psi_B$ as a proxy for the even-quadrant moment. (Note: these quantities may not be the best way to measure focus in some optimal sense, but they are easy to access and should roughly track focus. Since Tiny Tim models are only

[3] PA is the angle as measured in degrees from the vertical in an anticlockwise direction.
[4] It is unclear what property of the HST optics leads to this phenomenon. It is likely some kind of misalignment (George Hartig, personal communication).

qualitatively accurate, it is hard to come up with a better metric, so we will start with this. These rough proxies may eventually help us identify better proxies.)

We construct these moments for each of the 134,887 A-corner stars in the database. **Figure 6** shows the correlation between the two moments. There is a clear negative correlation present, but there is also a lot of noise.[5]

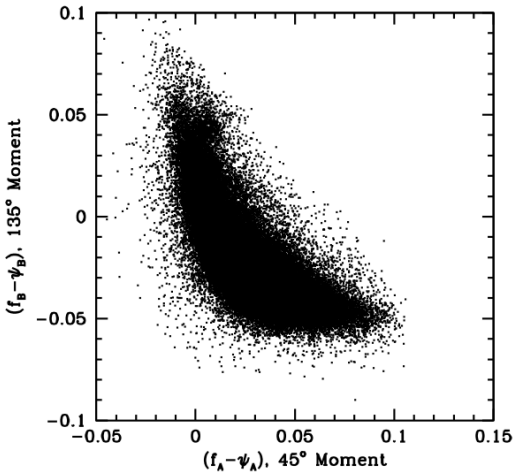

Figure 6: Correlation between the two extracted asymmetric moments for the stars in the upper-left quadrant.

Since these two focus-proxies are negatively correlated with each other, we can construct a single proxy for focus that is simply the difference between the two: $F_{A-B} = (f_A - \psi_A) - (f_B - \psi_B)$. We can then plot this against the other proxy we have for focus, namely the fraction of light in the central pixel relative to the model expectation: $F_C = (f_C - \psi_A)$. The left panel of **Figure 7** shows this correlation. Since the center of the star can be anywhere within the central pixel and since sometimes this places it very close to some of the A and B flagged pixels, we are likely to get a better estimate of the focus in an exposure if we concentrate on the stars that happen to be centered on pixels, so that the moments will be measuring similar things for all the stars. The panel on the right shows this correlation for the restricted sample.

[5] Due to an error in the data-reading algorithm, the y-values in the images were initially read in reverse order, which caused some error in the proxy values used in the banana plots. This error was corrected by effectively interchanging the A's and B's shown in the illustrative diagram on page 9. However, such a correction works perfectly only if the reference PSF is fully symmetric. The fact that the reference PSF is slightly asymmetric introduces a slight residual offset (~0.01) in the y-axis of the banana plots. The main conclusions presented here remain unaffected by this small error, but we plan to fully correct for this effect in a future version of the ISR.

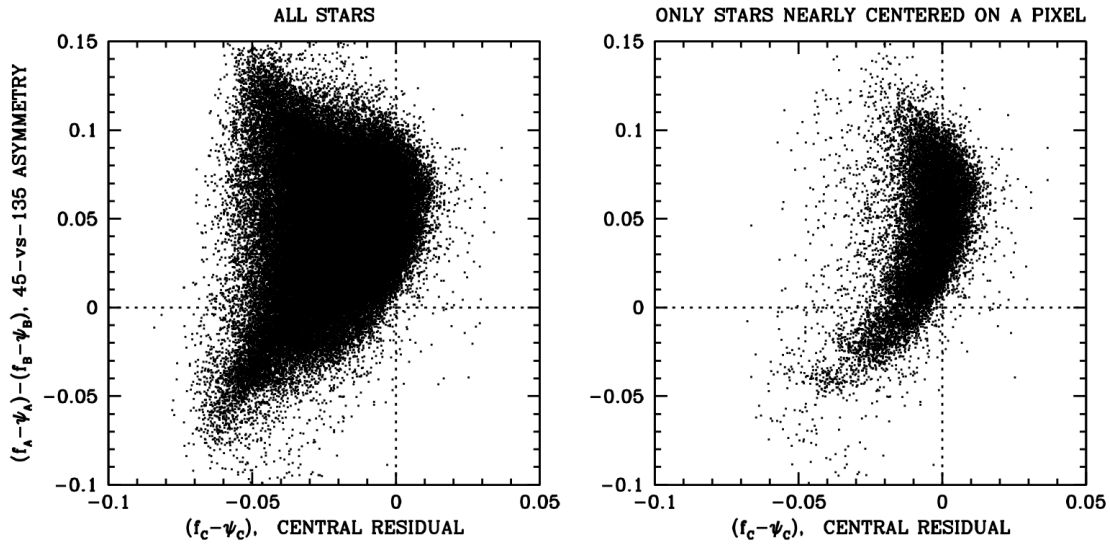

Figure 7: (Left) The "banana" plot for all stars in the upper left quadrant. (Right) Same, but for only those stars that are nearly centered on a pixel.

We will call this a "banana" plot, for obvious reasons. It shows that the focus appears to vary along a curve. On one side of focus, the PSF is less sharp and asymmetric in one direction, but as it goes through focus, the asymmetry decreases and the PSF becomes sharper. As we go beyond the optimal focus, the asymmetry continues through zero into the opposite sense and the PSF becomes less sharp again.

It is worth noting that the best focus is not at (0,0) on this plot, but rather is at (0,+0.06). This is an indication that the "library" PSF model that we are using may not be perfect. To explore this possibility, we subtracted the observed PSF model from each of the 134,887 rasters and distilled them into a residual PSF, in a procedure similar to the residual-plus-smoothing iterative procedure we used constructed the PSF in the first place above in §2.

Figure 8 shows an image of the residual PSF, with dark corresponding to excess flux and white to a flux deficit. There is some clear even-odd-quadrant asymmetry in this image, even though it appears that the central value of the PSF is well-represented by the model. The largest positive and negative pixels below have values of about 0.005, meaning that the worst pixels get 0.5% more of the total flux than the model predicts. Each of the pixels below is a quarter the size of a real image pixel. The dark blue box is marked around the inner 7×7 image pixels.

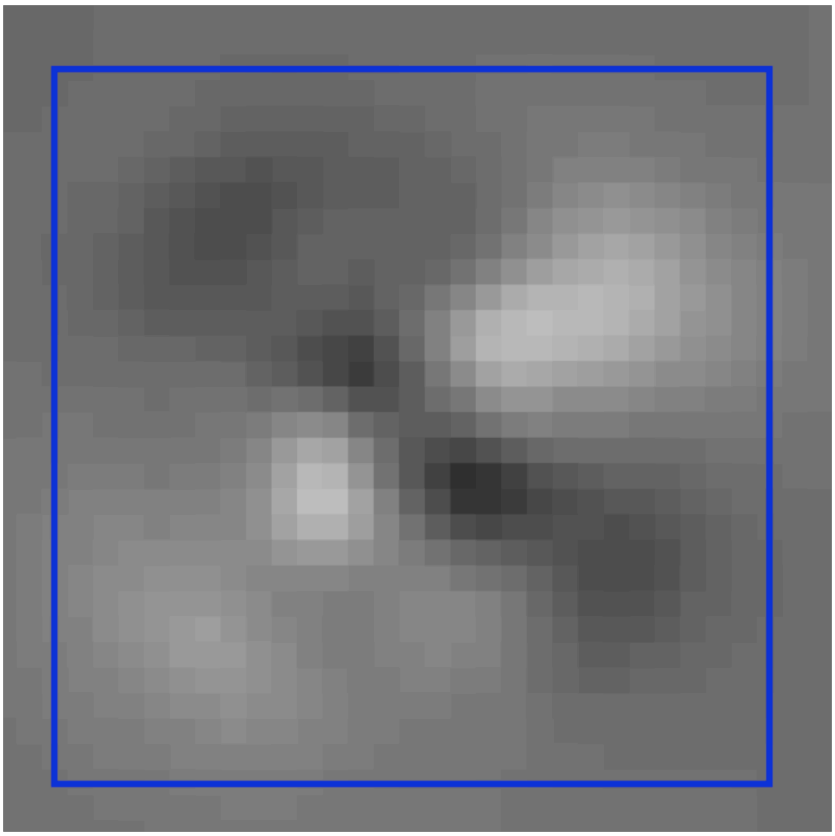

Figure 8: Residual of the PSF fits to all the stars in the upper-left quadrant. Dark is positive flux and white is negative flux.

To verify that this PSF works, we regenerated the banana plots and show them in Figure 9. As expected, the new points are centered on (0,0). The residual images in Figure 10 show that most of the residual has been removed. The largest peaks/troughs are down by a factor of about four from the first iteration.

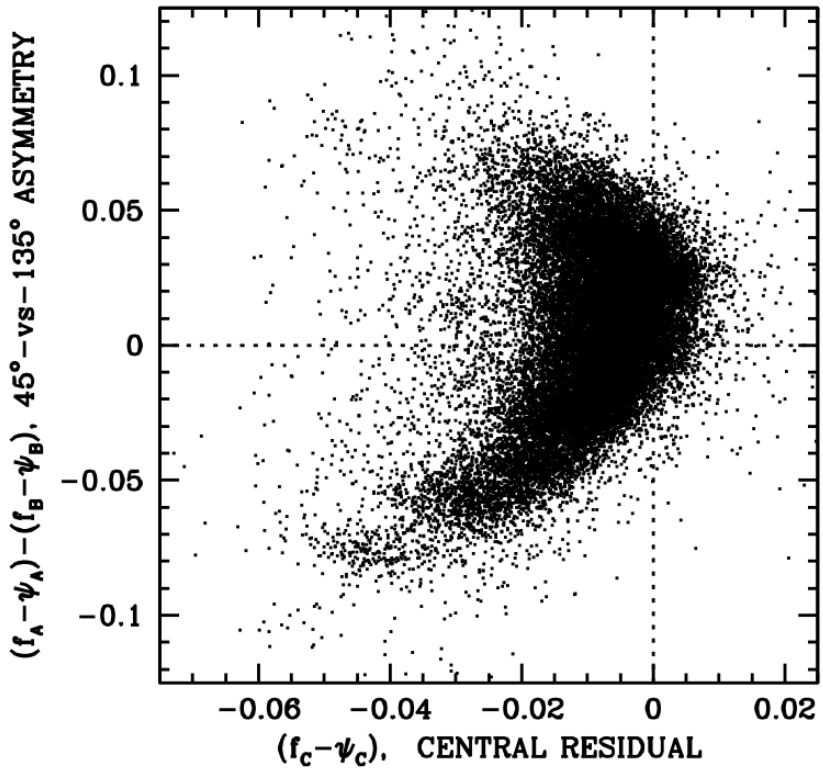

Figure 9: New "banana" plot, making use of the original PSF with the adjustment from Figure 8.

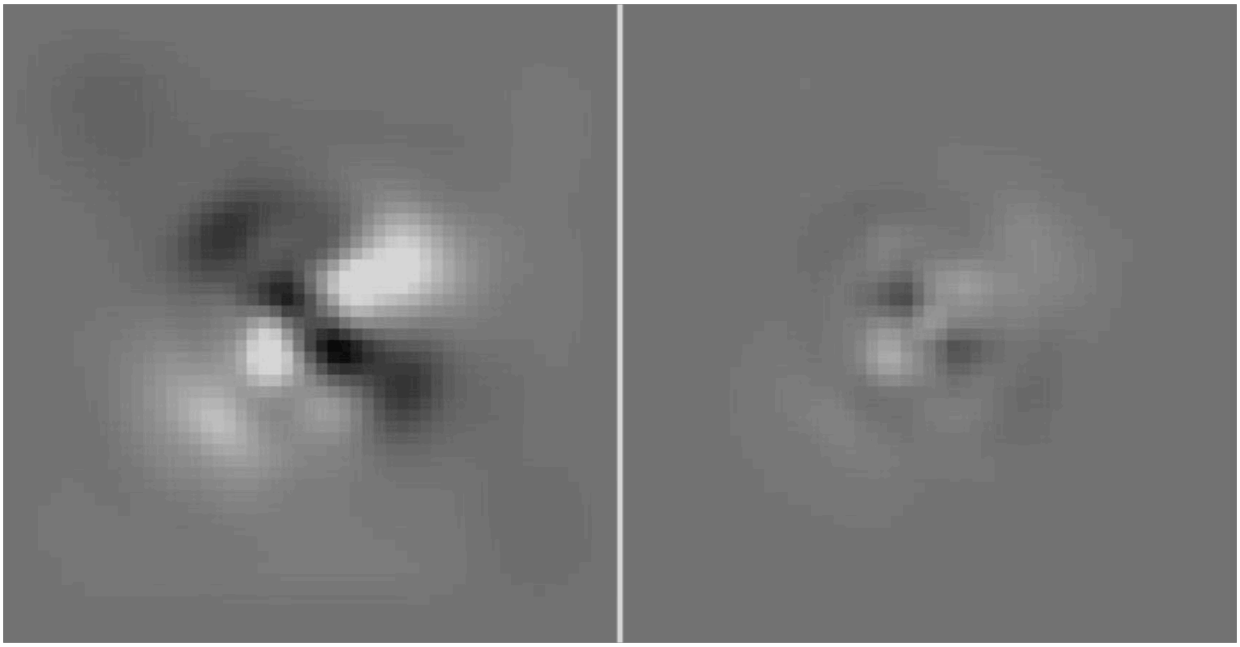

Figure 10: (Left) Residual of the star images with respect to the library PSF (same as in Figure 8). (Right) Residual of the star images with respect to the library PSF adjusted by the residuals on the left.

5. CHARACTERIZING THE IMAGES BY FOCUS

The individual points in the previous "banana" plots each came from different stars observed in a large number of different exposures over the lifetime of WFC3/UVIS. If the trend in the banana plots represents variations of focus, then all the stars from a given exposure *should* cluster in a particular location on this diagram.

To test this, we identified a few images that had a large number of stars in them, so that we would be able to define a coherent cluster. The left panel of **Figure 11** below is similar to that in **Figure 5**, except that here we are counting only the number of stars in the A-corner.

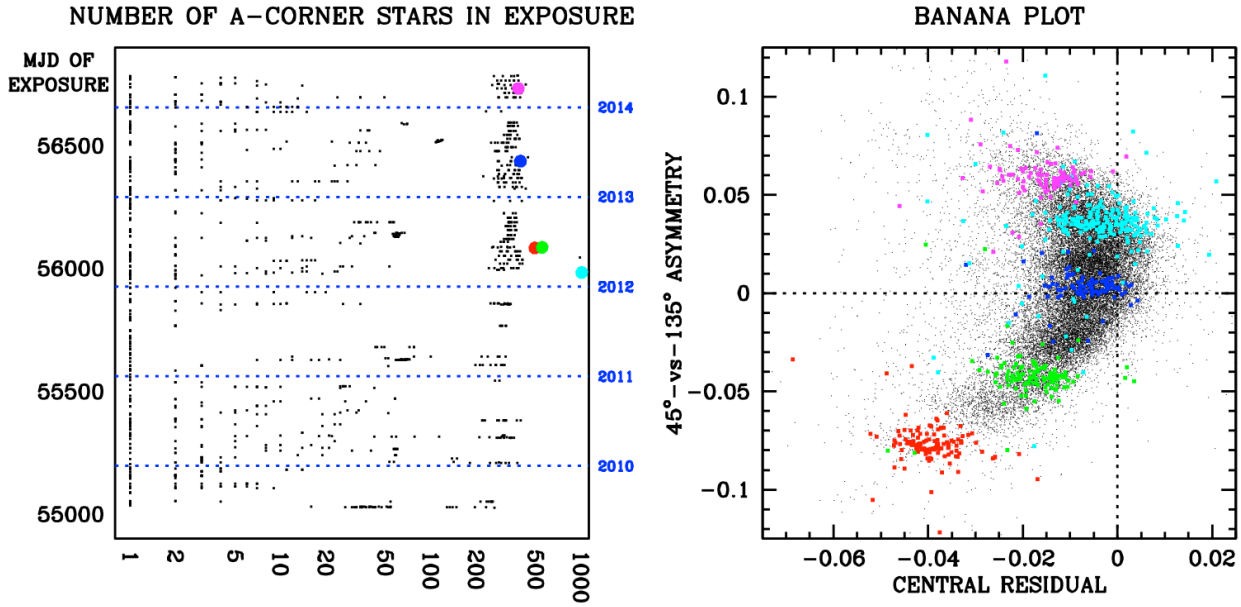

Figure 11: (Left) The number of stars identified in each exposure as a function of the date of the exposure. A few particular exposures with large number of stars are identified with large dots of various colors. (Right) The small dots are all the points in the previous banana plot; the colored dots represent those stars in each of the similarly colored exposures coded on the left.

We have identified a few exposures that have a good number of stars and color-coded them red (ibqs11cv), green (ibqs13lf), blue (ic3f28co), cyan (1bvd11a6), and magenta (ic8x12m9). We then color-coded all the pixel-centered stars coming from these exposures on the banana plot on the right. It is clear that stars from a given exposure do indeed all occupy a similar location in the focus space that we have defined.

If we can now identify specific images that have stars in various locations in the focus diagram, then perhaps we can characterize the PSF that is typical of that focus level. **Figure 12** below shows eight designated focus zones in the banana plot, as labeled to the right. The solid blue dots represent robust averages of all the well-centered stars in the exposures that have at least 25 well-centered stars. The number in parentheses next to the focus-zone label reports the number of exposures found to be in that focus zone.

Of course, if we could predict the focus for a given exposure from engineering data, then this exercise would not be necessary. There does exist a focus model for HST, but it is not accurate enough to use for this purpose. The eventual hope for this project is that we can use this study of correlated variations in the PSF to help us understand the focus better, and perhaps do a better job of either predicting the focus for a given observation, or at the very least determining it from the observation itself.

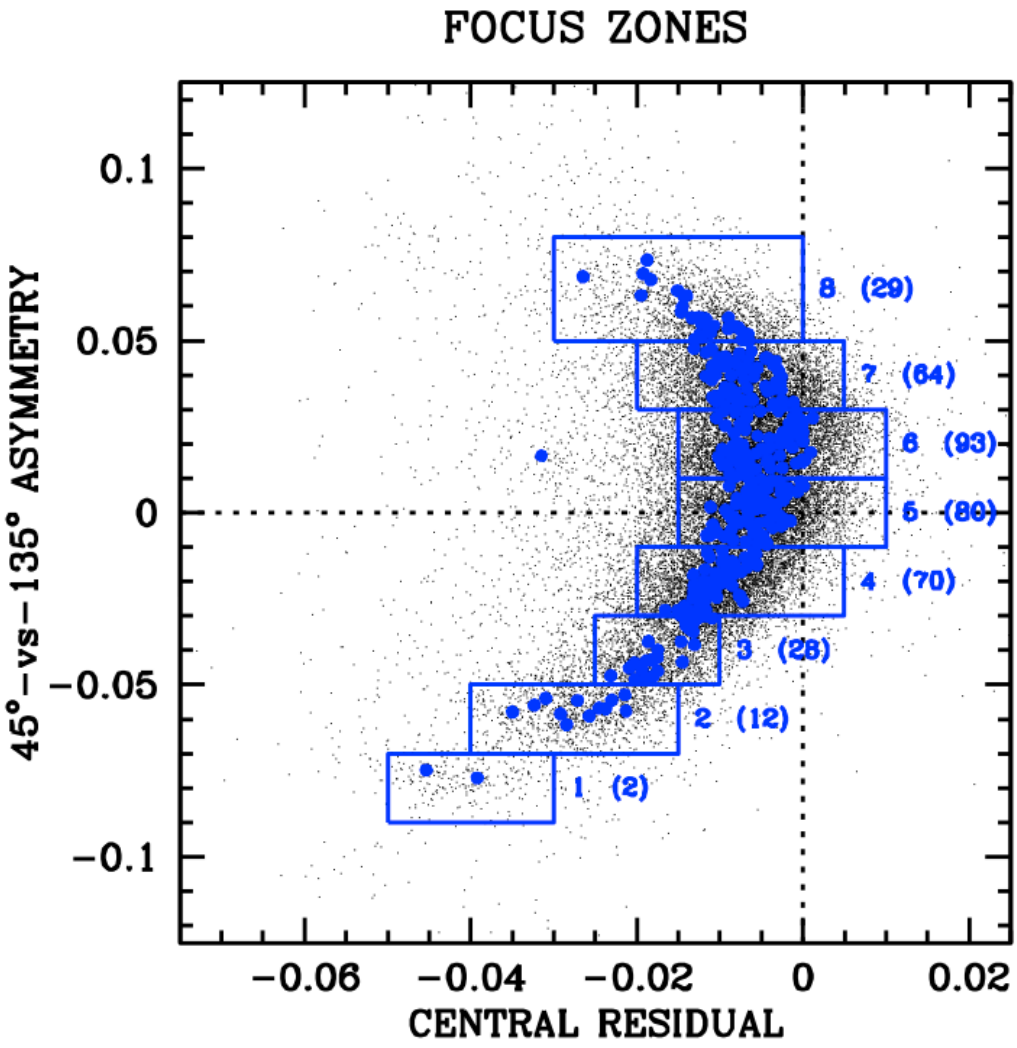

Figure 12: The same banana plot shown in previous figures, but with the zones identified.

6. THE PSF BY FOCUS

Now that we have managed to characterize each exposure in terms of some kind of monotonic focus parameter, we can take the next step. We can explore the residual PSF to see how it changes as the focus level goes from "1" to "8".

So, we ran the PSF-residual-finding routine from §4 on the exposures that had been traced to the eight different focus zones. This resulted in eight different residual PSFs, which we show in **Figure 13** below. As before, dark means excess flux and white means a deficit of flux.

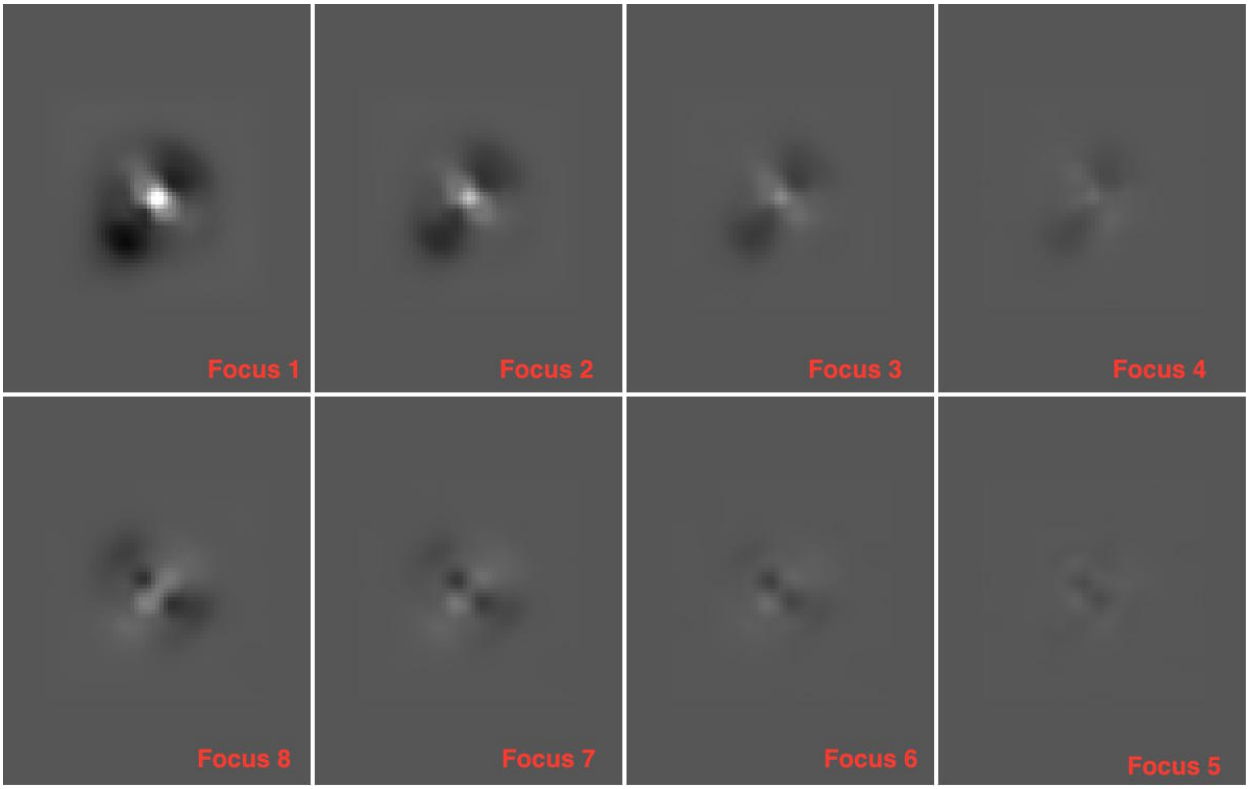

Figure 13: The residual PSF as extracted from star images taken at various empirically identified focus levels. Black means more flux, white less.

The PSFs go from the bottom of the banana to the top of the banana, from upper left clockwise to the lower left. The PSF goes in a nice, smooth way from having asymmetry in the 45° direction in the upper left, to having it in the 135° direction −− though the first asymmetry is much more severe than the second. This is likely related to the fact that this "A" corner is generally out of focus relative to the rest of the detector, so it has much more aberrated excursions on one side of focus than on the other.

The largest pixel values here are in the Focus-1 PSF. At the center of the PSF, the library-plus-static-residual-PSF predicts that about 4% more of the star's total flux will land in the central pixel than is observed. Conversely, it predicts that 1.5% more of the star's total flux will land in a pixel (−2,+2) relative to the center. The fact that the PSF transitions smoothly between successive focus levels tells us that perhaps we might be able to try to fit these eight different PSFs to the stars in an exposure in hopes of determining the focus level for that exposure.

7. PSF-FITTING FOR FOCUS

Rather than use moments, ideally it would be better to involve actual PSF-fitting in the determination of the focus level. We can use the residual PSFs determined above to construct a PSF for each of the focus levels. We can fit these eight PSFs to each star and determine a quality of fit for each. This will tell us which stars prefer which focus levels, but we shouldn't expect to be able to measure the focus level from an individual star. We will likely need to find which PSF provides the best fit "overall" to all the stars in an image. We remain focused for now on the most sensitive corner, but it may be that even though the PSF is not as sensitive to focus in the rest of the detectors, the large increase in the number of stars in the rest of the chip will compensate for the lack of sensitivity.

So, in order to identify the empirical focus level that best describes each exposure, we measure each star in the exposure with the eight different PSFs, where each PSF is a combination of: (1) the library PSF we started with (which is spatially variable), (2) the residual PSF constructed in §4 (which is constant across the chip, but was constructed only from stars in the upper-left corner, and (3) the residual PSF constructed for each focus level (again constant across the detector). For each star we determine eight different values for `qfit` (see Fig 3) and will attempt a statistical analysis to infer the overall best fitting PSF.

The left panel of **Figure 14** below shows the variation of the `qfit` parameter as we change the PSF from that of focus-level 1 through that of focus-level 8. We find that this particular star is best fit with a focus level between 4 and 8, but it is hard to tell the exact best focus. In the middle panel, we fit all 26 stars in the sensitive corner for this particular exposure. Most of these curves have a minimum near focus-level 5, but since they all have different minimum qualities, it's hard to define a single best focus. The left panel shows the percentiles for the plots in the middle panel. At every focus level, we took the distribution of qfit values and connect the lines for the 10%, 20%, 30%, … 50%, .. and 90% locations within this distribution. Pretty much all of these have a minimum at focus-level 5. In fact, if we were to fit some kind of quadratic to each one, we would arrive at a focus-level of about 4.75.

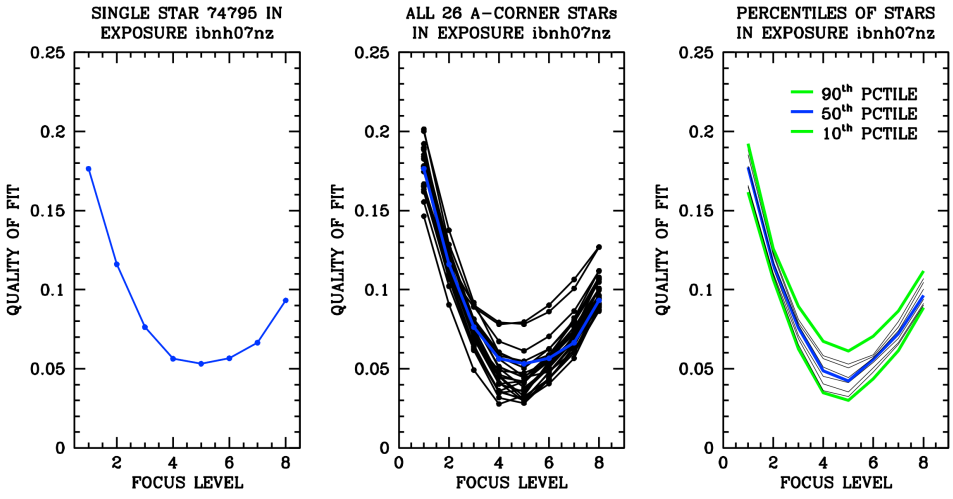

Figure 14: (Left) The qfit parameter for a given star as determined from PSFs at different focus levels. (Middle) Same for all 26 A-corner stars in the same exposure. (Right) Percentiles of the curves in the middle panel.

8. NEXT STEPS

We began this study by going through the entire archive of WFC3/UVIS observations taken through the F606W filter. We searched each exposure for high-signal-to-noise stars (S/N > 100) and constructed a database of over 5 million star images, each of which was observed at a particular place on the detector in a particular exposure.

We next took a spatially variable "library" PSF and investigated how, in a crude, moment-based sense, the pixels in a given star image differ from the best-fit model PSF. We then focused on stars that were in the upper-left corner of the top chip and were also centered on their central pixels, since the "moment-based" signal should be the clearest for these stars. We examined two moments in particular: one moment was related to the star's asymmetry and the other to the star's central intensity. We found a clear correlation showing a progression from out-of-focus on one side, through focus, to out-of-focus on the other side.

We next plotted on this focus diagram all the stars that came from a given exposure and found that these stars tended to cluster in a specific place along the focus curve. Clearly this represents a way to measure the focus in an exposure. We then identified 8 fiducial focus "zones" and associated with each exposure the corresponding zone. This allowed us to derive a "residual" PSF for each focus level. This corresponds to the difference between the model PSF and the PSF that corresponds to that focus level.

It was trivial, then, to construct a new model PSF for each focus level. We did this and fit every star in every exposure with the PSF for the eight different focus levels to determine how its quality of fit depended on the assumed focus level. We found that by combining the quality of fit metrics for as few as 25 stars in an image, it is possible to pin-point the focus level to better than one zone. This seems to be a very powerful way to determine the focus for a given exposure. If we can generalize the analysis we have done here on the upper left corner to the entire chip, then it might be possible to use all the stars in an exposure to help us determine the focus level, which will in turn give us a PSF that can be used anywhere on the detector.

To this end, it will be interesting to see how many stars of what signal-to-noise level we need to constrain the focus adequately. It will also be interesting to see what "adequately" means in terms of astrometric, photometric, and modeling errors. However, these questions are beyond the scope of this exploratory document.

Finally, once we have an estimate of the focus level for each exposure we can see how the focus varies within an orbit and over longer periods of time. We can then investigate how this empirical measure of focus may correlate with the actual focus (as measured using phase-retrieval technique) and the focus model. Future documents may explore these issues, and more.

Acknowledgements: We would like to thank Mike Fall for helpful comments on the manuscript. We thank George Hartig, Matt Lallo and Linda Dressel for useful discussions.

REFERENCES

Anderson, J., 2014, "Local Pixel Bundles: Bringing the Pixels to the People," Instrument Science Report WFC3 2014-24

Kozhurina-Platais,V. & Anderson, J. 2015, "Standard Astrometric Catalog and Stability of WFC3/UVIS Geometric Distortion," Instrument Science Report WFC3 2015-02

Anderson, J., & King, I. 2006, "PSFs, Photometry, and Astronomy for the ACS/WFC," Instrument Science Report ACS 2006-01 (AK06)

Krist, J., 2003, "ACS WFC & HRC Field Dependent PSF Variations Due to Optical and Charge Diffusion Effects," Instrument Science Report ACS 2003-06

Sabbi, E. & Bellini, A., 2013, "UVIS PSF Spatial & Temporal Variations," Instrument Science Report WFC3 2013-011