WFC3 Instrument Science Report 2016-11

The Effect of Repeated Exposures on Measured Fluxes in the WFC3/IR Detector

Knox S. Long, Sylvia M. Baggett & Vera Kozhurina-Platais
August 1, 2016

ABSTRACT

Many observations carried out with the WFC3/IR detector consist of multiple exposures of a single field; almost all of these are dithered. Here we investigate the degree to which the measured flux depends on its previous history by examining a set of normal staring observations where multiple exposures with the same filter were taken at the same pointing position on the sky. We find that the measured flux does rise with the number of exposures, and that the increase can affect photometry at the level of 0.02-0.04 magnitudes. The phenomenon is clearly related to the slow rise in measured flux seen in grism observations of planetary transits, and is likely associated with the anomaly known as persistence, but is not estimated well using our current persistence model. Although most WFC3/IR programs will not be compromised scientifically by this behavior, observers should consider its effects when very high precision photometry is required.

1 Introduction

Observations of planetary transits generally involve a large number of short grism exposures. Originally most such observations were carried as staring mode observations see, e.g. Berta *et al.* 2012; Wakeford *et al.* 2013), but today most observers opt to scan sources across the detector to increase the duty cycle of the observations (see, e.g. Kreidberg *et al.* 2014). In both cases, observers notice that the flux observed from the

star appeared to rise at the beginning of an orbit and in the case of the scanned observations the rise time is sufficiently long that observers typically discard the first orbit of an observation, treating it as a "conditioning" orbit.

Persistence is an afterglow of earlier images that is generally seen in IR detectors, including WFC3. Persistence arises from traps in the diodes of the detector; as the bias voltage decreases during an exposure, these traps capture free charge that is released slowly in subsequent exposures (see Smith $et\,al.$ 2008ab). The amount of persistence is known to increase with the number of repeated exposures (Long $et\,al.$ 2013ab).

The fact that repeated exposures increase the amount of persistence can be understood if the traps in the diodes of the detector have trapping times which are comparable to or exceed the time of the original exposures. A greater percentage of the traps are filled at the end of a series of exposures than for a single exposure and these produce more persistence. Qualitatively, one can also understand why the measured flux would rise with time as well. The mostly empty traps in the first exposure capture charge that would otherwise discharge the diode. The mostly filled traps release additional charge in subsequent exposures that contribute to the discharge of the diode. Both effects tend to make the measured persistance signal greater in later exposures.

The model we currently use to predict persistence in WFC3 IR images has the following form:

$$P = A\left(\frac{t}{1000\ s}\right)^{-\gamma} \tag{1}$$

where A and γ are functions of fluence and t is the time since the earlier exposure. The parameters A and γ depend of the total number of electrons in a pixel (as measured in the flt file, hereafter the fluence) and the exposure time of the earlier image (see, Long $et\,al.$2015). The parameter A is of order 0.3 and γ is of order 1 for pixels with a fluence $>$ 70,000 e, which suggests persistence at 100 seconds is about 3 e s^{-1}. We do not however have a good calibration of persistence at delay times as short as this, and it is also clear that at very short times this power law approximation must eventually fail, as the formula predicts infinite persistence at zero delay time.

The purpose of this report is to investigate the magnitude of these persistence-related effects in normal staring observations, especially as they might affect stellar photometry. Our approach is to examine observations in which multiple exposures were obtained of a single field without dithering so that, in the absence of image motion, the photon flux in each exposure should be identical.

2 Observations

For this study, we make use of a series of observations from program 12694, which was one of the calibration programs used to show that persistence increased when a pixel is illuminated repeatedly (Long $et\,al.$ 2013b). Each visit in this program consisted

Table 1. Omega Cen images from proposal 12694.

Visit	Root	Num.Ims	Obs. Date	Target	RA	Dec.	PA V3
11	ibvd11	1	2012-02-27	OMEGACEN-1	201.6928	-47.4791	133.8990
13	ibvd13	3	2012-03-19	OMEGACEN-1	201.6928	-47.4791	151.0826
15	ibvd15	5	2012-04-28	OMEGACEN-1	201.6928	-47.4791	205.1857
22	ibvd22	2	2012-04-15	OMEGACEN-2	201.7025	-47.4847	194.8147
24	ibvd24	4	2012-03-17	OMEGACEN-2	201.7025	-47.4847	149.2986
26	ibvd26	6	2012-03-20	OMEGACEN-2	201.7025	-47.4847	151.8888
31	ibvd31	1	2012-04-27	OMEGACEN-3	201.6837	-47.4862	203.3965
33	ibvd33	3	2012-03-20	OMEGACEN-3	201.6837	-47.4862	152.1025
35	ibvd35	5	2012-03-03	OMEGACEN-3	201.6837	-47.4862	137.4953

of from one to six undithered exposures of a field in Omega Cen followed by a series of darks. For the analysis described here we are only interested in the Omega Cen exposures, which are listed in Table 1.[1] The sample sequence STEP50 with NSAMP of 13 (resulting in an exposure time of 349 s) and the filter F125W for all of the external exposures were selected to assure a wide range of saturations levels in the detector. For the purpose of this report, we have re-calibrated the data with the latest version of CALWF3 (version 3.3) and used reference files available in the spring of 2016. We use the flt products which are fit-up-the ramps of the multi-accum images and have cosmic-rays removed.

3 Pixel Comparisons

To see whether repeated observations had any effect at all on the level of persistence, we first compared pixel values in the flt files from multiple exposures. The results are shown in Fig. 1 for Visit 26, which had 6 repeat exposures, and in Fig. 2 for Visit 15, which had 5 repeated exposures. To prepare these figures, we eliminated pixels which had non-zero data quality flags. We then divided the pixels in the first image into bins on a logarithmic scale from 1000 e to 1,000,000 e, and took the median value of those pixels in the subsequent images. The results show that pixel values rise for each exposure at all fluence levels. The shape of the rise is relatively complex, however, with a pronounced change around the pixel saturation level of 70,000 e. At fluence levels below saturation, the ratio of the flux of an image to the first image rises in each exposure; above saturation most of the flux rise is in the second and third exposures. The maximum increase in fluence in both visits is about 3%.

There are two ways in which these figures can be interpreted, the simplest being that the flux increases are all due to changes in the response of pixels to light from ear-

[1] Visits 11 and 31 are included in the table for completeness but they were not part of this comparison study since they contain only a single Omega Cen exposure.

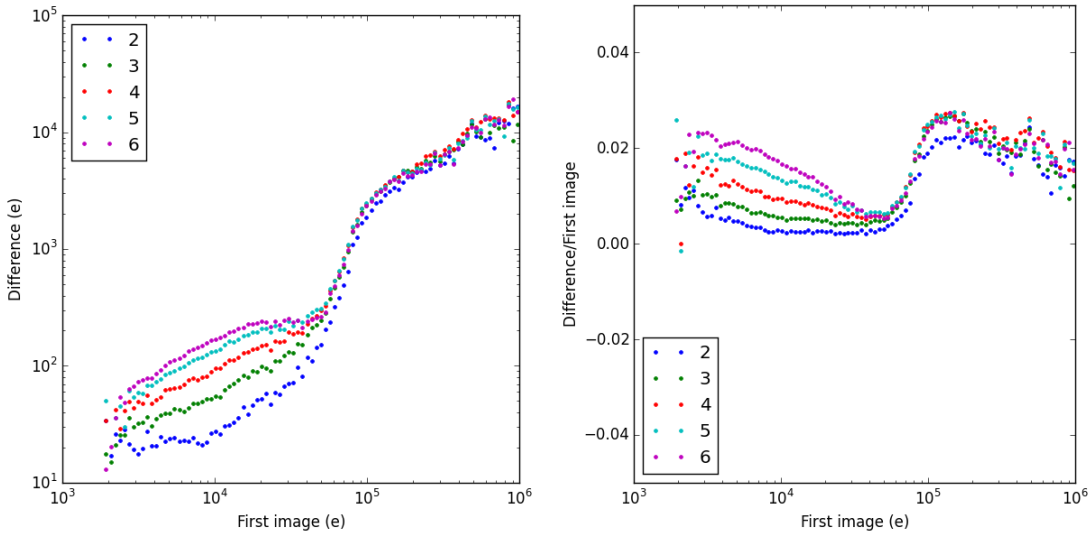

Figure 1. Both panels show how pixel levels change in exposures 2 through 6 of Visit 26 as a function of the fluence measured from exposure 1. Left: Increase in fluence measured in pixels for each of the exposures. Right: Ratio of the fluence in the exposures 2 through 6 to that in the first exposure.

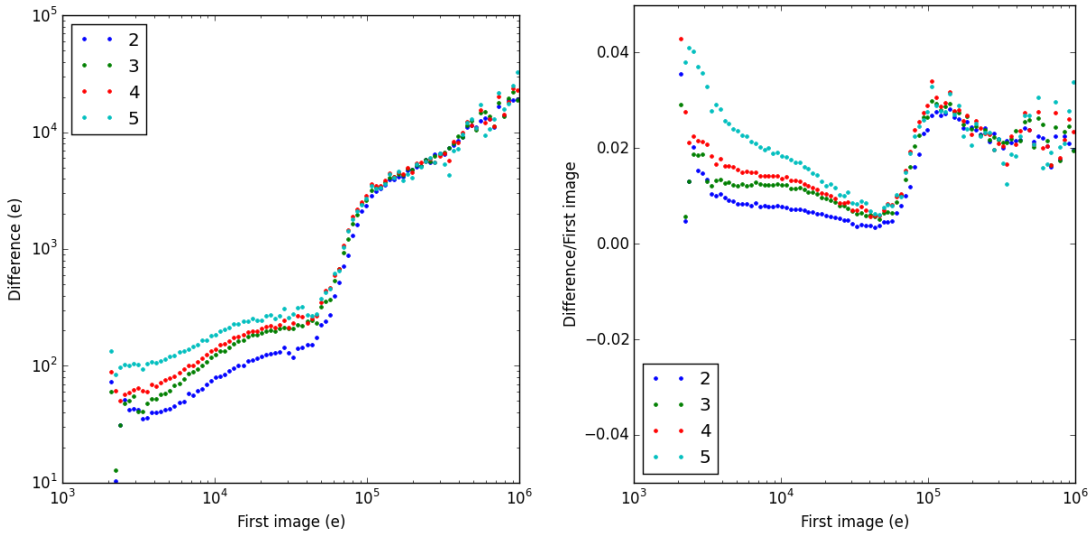

Figure 2. Same as Fig. 1 but for Visit 35, which had five repeated exposures.

lier exposures. However, there is another possibility. Although these images were not dithered, HST pointing does jitter and drift by very small amounts between pointings and this has the potential to change the fluence measured in a pixel in different exposures, especially since the WFC IR array undersamples the point spread function. Such pointing drifts can cause the measured flux to change, and for a pixel with low fluence adjacent to a pixel with a high fluence in the first image, there would be bias towards increased fluence in subsequent images.

One way to check whether image shifts are affecting the pixel comparisons shown in Fig. 1 and 2 is to use the last image in the sequence as the reference image. If drifts are causing the apparent flux increases, we should see them regardless of the "arrow of time". The results are presented in Fig. 3 and 4 for Visits 26 and 35, respectively. These figures show that regardless of whether one uses the first image or the last image as a reference, the measured fluence increases throughout the exposure sequence. If image motion was the root cause of the apparent changes in fluence, we should have seen the measured fluence increase in earlier exposures; instead we see it decrease. If image motion was the cause, then right hand panels of Fig. 3 and 4 should be qualitatively similar to that of Fig. 1 and 2; instead we see a decrease.

Another possible concern with the pixels comparisons as presented thus far is the possibility that variations in the sky background are affecting the results. The sky background for the F125W filter is about 1 e s^{-1}, and thus its contribution to the fluence is about 349 e. Since the effects we are measuring, particularly at low fluence levels, are of order 100 e, large changes in the sky background could affect our results. The sky background rates in the IR filters do vary from observation to observation (due primarily to changes in the zodiacal light), but short term changes have only been documented in the filters that pass He I λ 10830 airglow, such as the F110W filter (Brammer $etal.$2014). The F125W filter does not pass the He I λ 10830 airglow line, and there is no evidence that the sky background varies significantly within orbits.

We conclude that the primary effect we are measuring is due to changes in the behavior of pixels due to their exposure history.

4 Photometry

While a single-pixel based analysis is useful for trying to understand the "physics" of repeat observations, precision "astronomical" measurements generally involve measuring the flux of objects, typically stars or other point like objects, via photometry. We have therefore carried photometry experiments on the data using three separate tools: DAOPHOT (Stetson 1987), as implemented in IRAF; DOLPHOT, which is an updated version of HSTPHOT (Dolphin 2000); and an early version of a soon-to-be-released photometry package written by Jay Anderson.

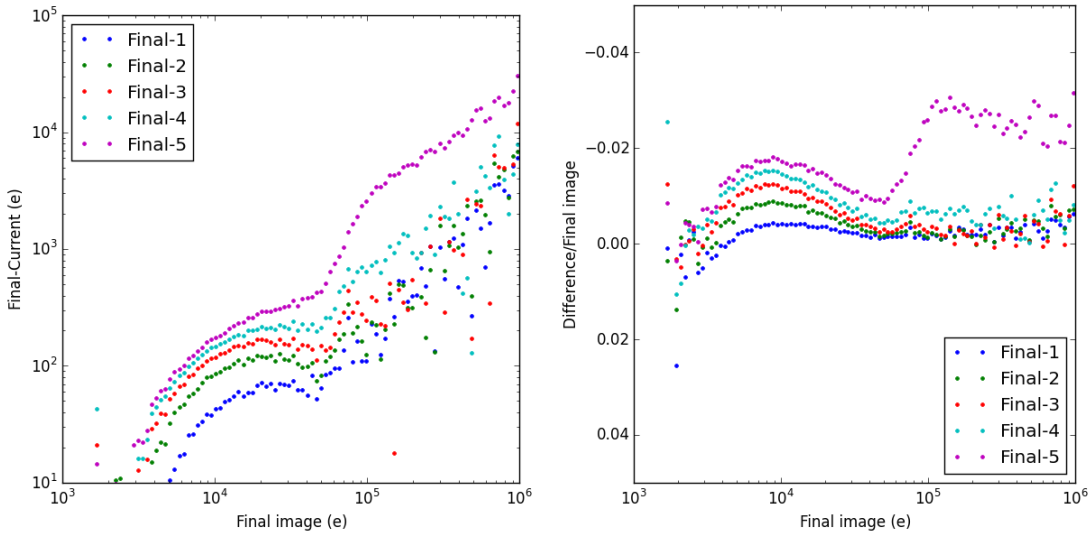

Figure 3. Similar to Fig. 1, but here the last (sixth) image has been used as the reference image. The curves labeled "Final-1" refer to the fifth exposure in the sequence, etc., so 'Final-5" is the first exposure. Left: Decrease in fluence measured in pixels for each of the exposures relative to the last image in the sequence. Note that the sense of the the comparison has been reversed compared to the earlier figures. Here we plot the difference between the last image and an earlier image, so positive numbers imply that the measured flux is less in earlier images. Right: Ratio of the fluence in the exposures to 1 though 5 to that in the last exposure. Note that the sense of the y axis has been reversed, for easier comparison with Fig. 1.

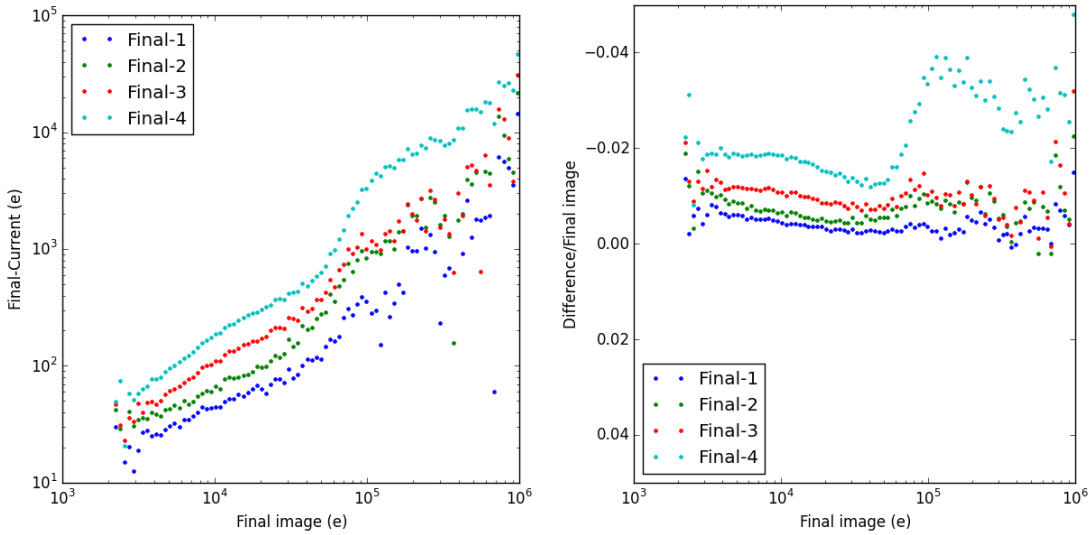

Figure 4. Same as Fig. 3 but for Visit 35, which had five repeated exposures.

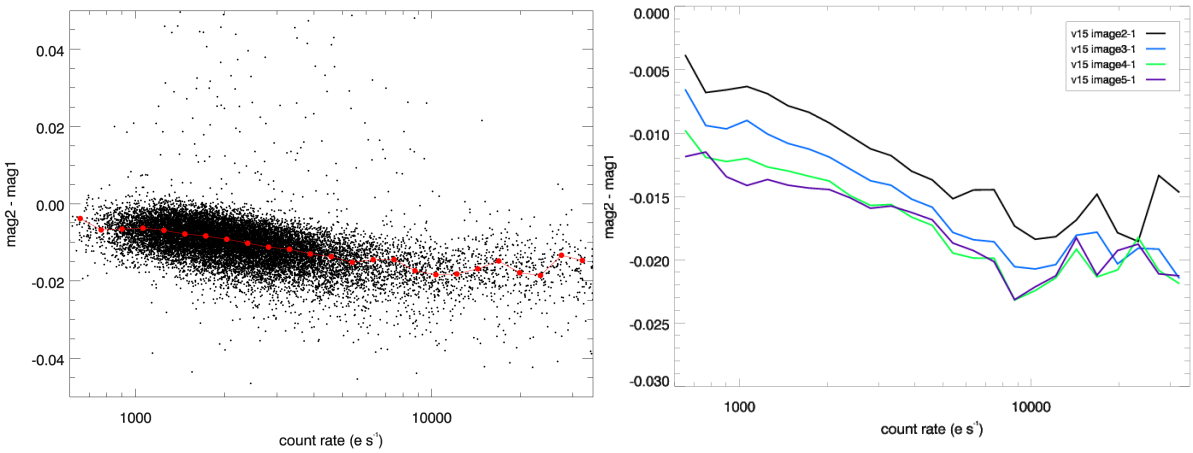

Figure 5. At left, difference in magnitudes of Omega Cen exposure 2 to exposure 1 aperture photometry as a function of log source flux (e s^{-1}) for Visit 15, Program 12694. Black points represent individual measurements, while red points show the binned averages. At right, magnitude differences of all Omega Cen images in Visit 15, each relative to exposure 1 from that visit. Although plotted as lines for clarity, these are the binned averages in each case. Primarily as a result of the varying numbers of sources as a function of count rate, errors in the binned averages are of order 1×10^{-4} and 1×10^{-3} for count rates less than and greater than about 10,000 e s^{-1}, respectively.

4.1 DAOPHOT

As a first step in estimating the photometric effects of repeated, undithered observations, we perform aperture photometry using the NOAO APPHOT routine within IRAF on each of the Omega Cen exposures in Program 12694. The procedure can be summarized as follows. We first used DAOFIND task to locate sources within the first flt image of the visit. About 18,000 sources were identified. The DAOFIND algorithm, which operates by identifying local maxima within the image based on the FWHM of the PSF, is less sophisticated than other algorithms (e.g. DOLPHOT, discussed in a later section) but adequate in terms of providing one way of evaluating the photometric behavior of undithered exposures. We kept the positions of the sources fixed in the remainder of the exposures in each visit. We then used APPHOT to carry out the aperture photometry, extracting the count rates using a 3-pixel radius, without recentering, background subtraction, or application of the bad pixel mask.

The results for Visit 15 are shown in Fig. 5. The difference between the exposure 2 and exposure 1 magnitudes is shown in the left panel as a function of the log source flux (e s^{-1}) in the first exposure. Results for smaller apertures (not pictured) are similar, but show more dispersion. The larger dispersion is likely partly due to lower signal-to-noise within the aperture, partly due to the effect of individual bad pixels, and partly

to slight image shifts that will have a greater impact on any signal measured within very small apertures. As Fig. 5 illustrates, magnitudes measured in the second exposure are about 1% larger than in exposure 1 (i.e. magnitude differences of about -0.01) for fainter sources, rising to about 1.5% difference for brighter sources. The right panel of Fig. 5 summarizes the magnitude differences for all Omega Cen exposures in Visit 15 relative to exposure 1. Note the relatively small Y scale in this figure, to allow for easier comparison of the differences. The fluxes measured in the subsequent exposures are all systematically higher than in the first exposure. The largest change is for the second exposure; fluxes in subsequent images continue to grow but the rate of the increase is less. That is, there is a steady progression e.g. at fluxes of a few thousand e s^{-1} from about 1% offset from image 1 (image 2-1) to what appears to be an asymptotic value of about 1.5% offset from image 1 (image 4-1 and image 5-1).

No obvious shifts are present between the images but since we have not re-centered the sources located with DAOFIND, we reverse the analysis order similar to the process described in Sec. 3. That is, we used DAOFIND to locate sources in the last exposure and then used these positions for carrying out aperture photometry to determine the magnitude difference of exposure1 photometry to 5, 2 to 5, etc. The trends in the results are equivalent to those shown in Fig. 5 except that, as expected if the trend is persistence-related and not due to image shifts, the magnitude differences were positive with levels at about 1% and asymptotically approached values of about 2%.

Fig. 6 presents the results for all visits from 12694 with more than one Omega Cen exposure per visit. Here we see first that for faint sources, all visits start with exposure 2 being about 0.5% brighter than exposure 1 (i.e. mag2-mag1 of about -0.005). For brighter sources, the exposure 2-1 magnitude differences range from about 1.5% to about 2% depending upon the visit. The exposure 3-1 magnitude differences are a few tenths of a percent higher than the exposure 2-1 differences. However, only visit 24 shows the asymptotic behavior seen in Visit 15; the other two visits with four or more Omega Cen images show more evenly spaced offsets between flux ratios. Visit 26, with the most repeat exposures, shows a reversal, where the magnitude differences 6-1 (red curve) are closer to the 2-1 differences than the 5-1 differences. In general, for sequences of 4 images or less, photometric effects top out at about 2%; when more exposures are added, those discrepancies can increase to more than 2.5%.

Finally, we evaluate the aperture photometry magnitude differences as a function of exposure in the sequence, combining all visits into single plots. The results are presented in Fig. 7. None of the behavior in the visits duplicates another but instead, there's a spread of about 0.5% between the magnitude difference values for any given exposure in the sequence.

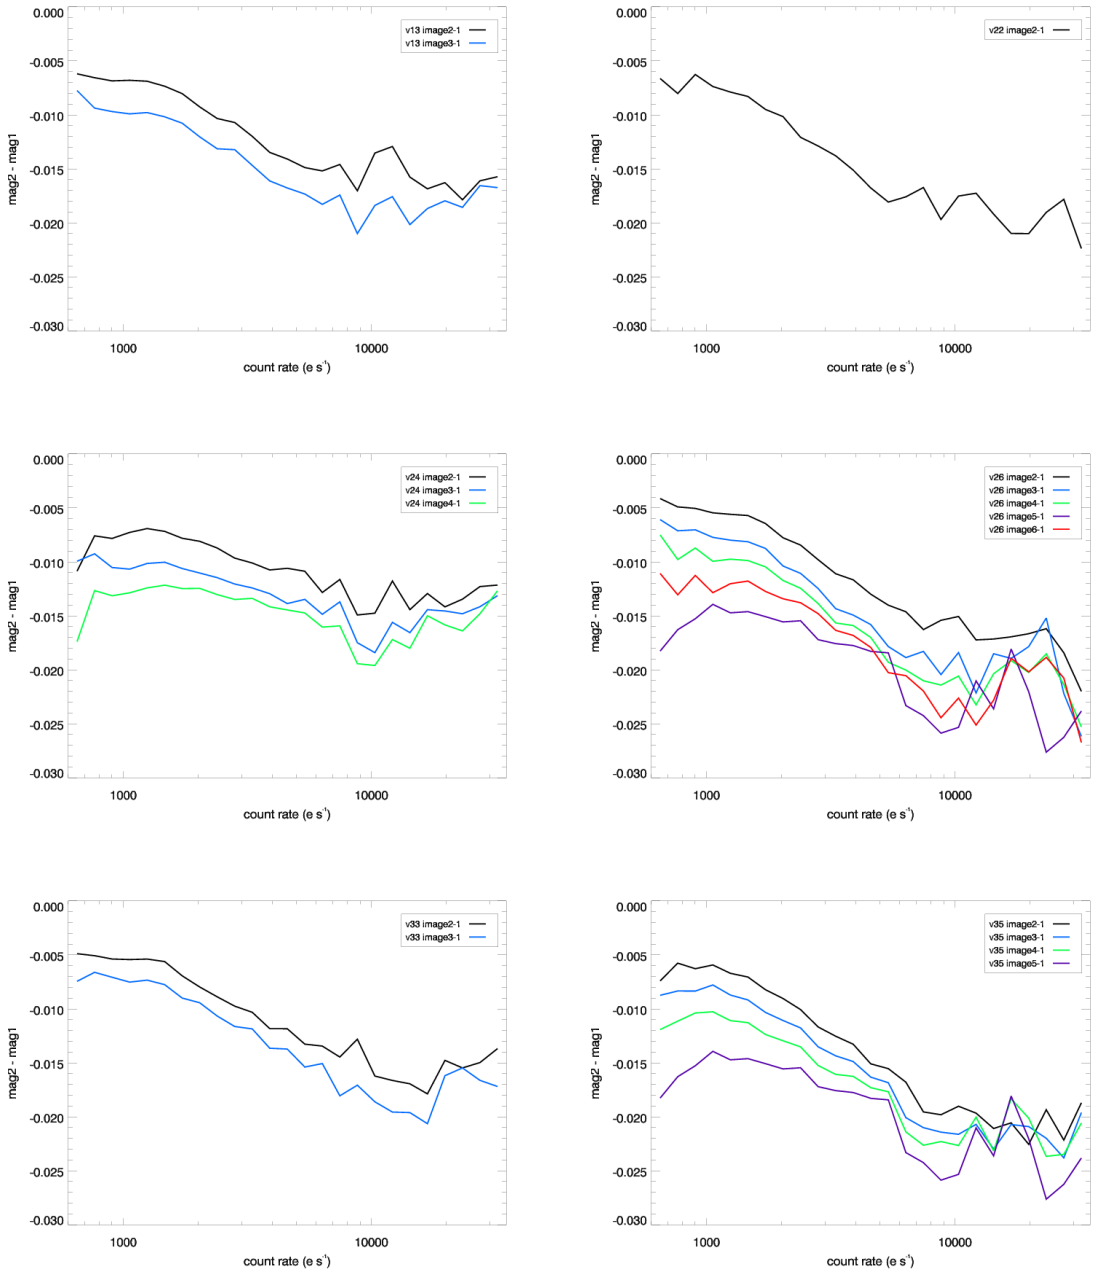

Figure 6. Omega Cen aperture photometry results relative to exposure 1 for each visit in 12694 (Visit 15 is in Fig. 5). Plotted as lines for clarity, these are binned averages; colors are consistent with those from Fig. 5. Errors of the binned averages are of order 1×10^{-4} and 1×10^{-3} for count rates less than and greater than about 10000 e s^{-1}, respectively.

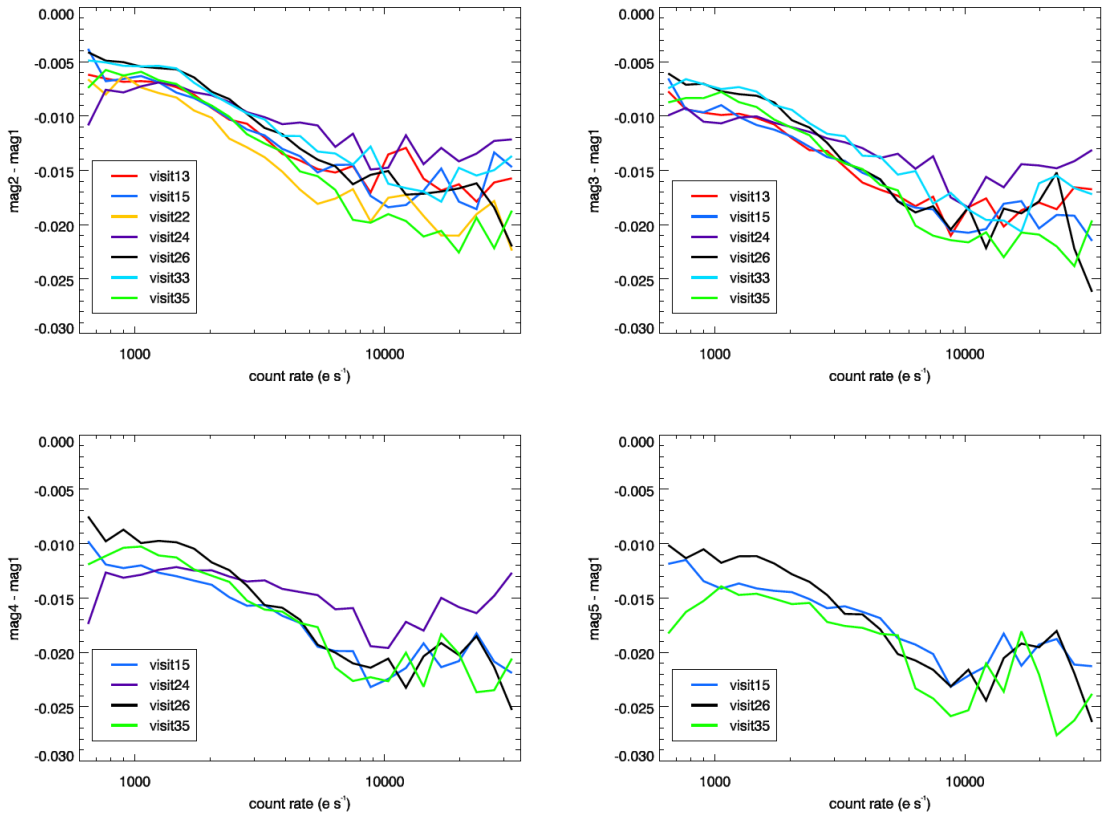

Figure 7. Omega Cen aperture photometry results for all exposures, relative to exposure 1 in each visit, as a function of log source flux ($e\ s^{-1}$). At upper left and right are the curves for all exposure 2-1 and exposure 3-1 magnitude differences, respectively, while at lower left and right at the curves for exposure 4-1 and exposure 5-1 differences. In all cases, binned averages are shown. As previously, errors in the binned average are of order 1×10^{-4} and 1×10^{-3} for count rates less than and greater than about 10000 e s^{-1}, respectively.

4.2 DOLPHOT

DOLPHOT is a publicly available photometry package[2] which was designed explicitly for the analysis of HST images (see Dolphin 2000 for a description of the algorithms involved). It has been used extensively for this purpose, and recently in the generation of the catalog for the HST survey of M31 (Williams $et\ al.$2014).

Although DOLPHOT can be used to measure fluxes via simple aperture photometry, here we have produced fluxes used DOLPHOT (Version 2.0) via psf fitting. Specifically, we produced catalogs by sequentially applying the routines WFC3MASK, CALCSKY, and DOLPHOT, on each of the external exposures, using the parameters recommended for use with WFC3/IR observations in the DOLPHOT User's Guide. Each im-

[2]http://americano.dolphinsim.com/dolphot

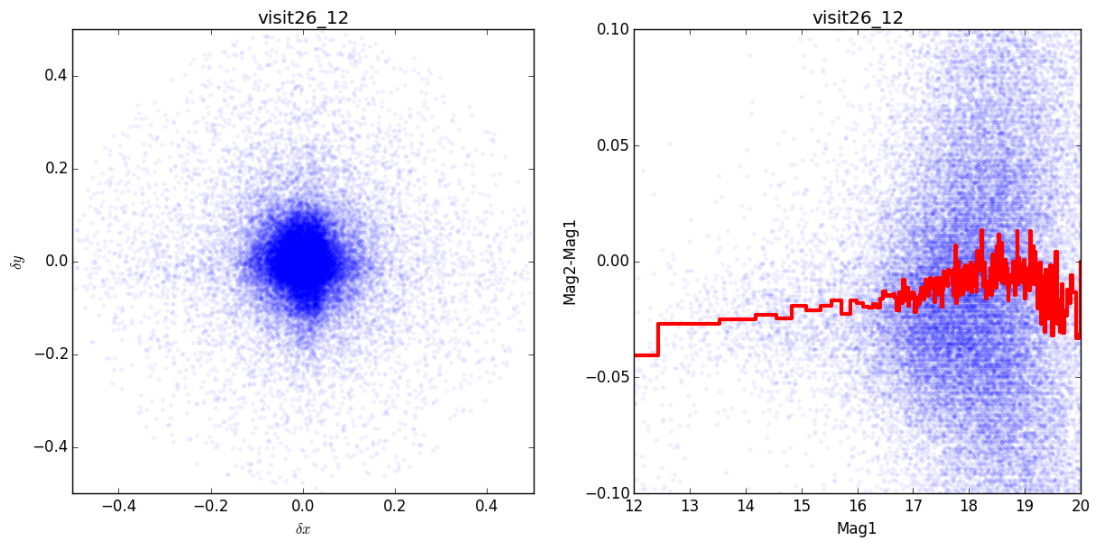

Figure 8. Left: The offsets in matched sources from a DOLPHOT comparison of sources in the first and second images of Visit 26. Right: The difference in Vega magnitude of sources in the first and second Omega Cen images of matched sources as a function of the magnitude in the first image. Individual objects are plotted in blue while average in various magnitude bands are plotted in red.

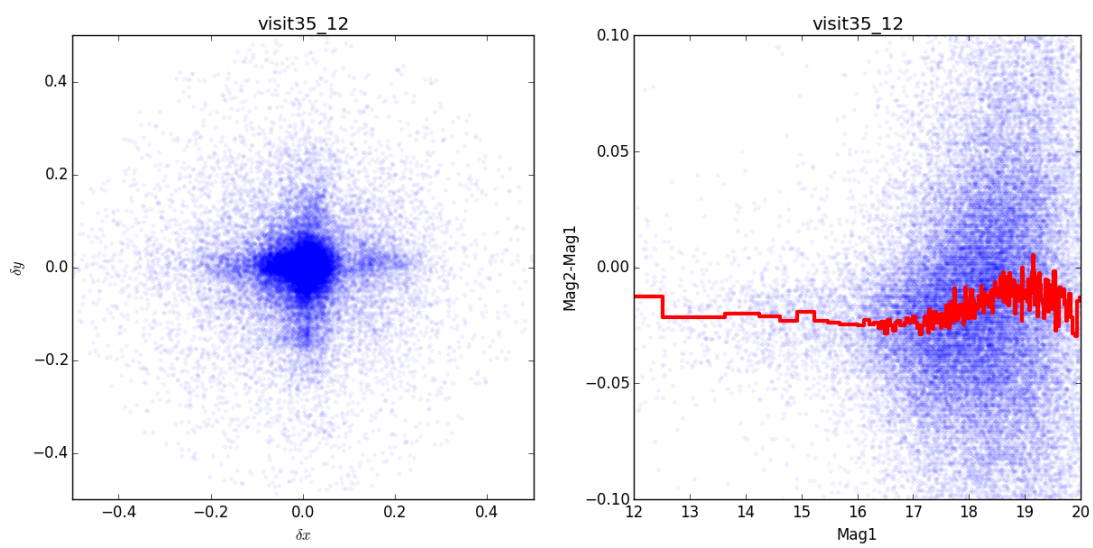

Figure 9. Same as Fig. 8, but for the first two images of Visit 35.

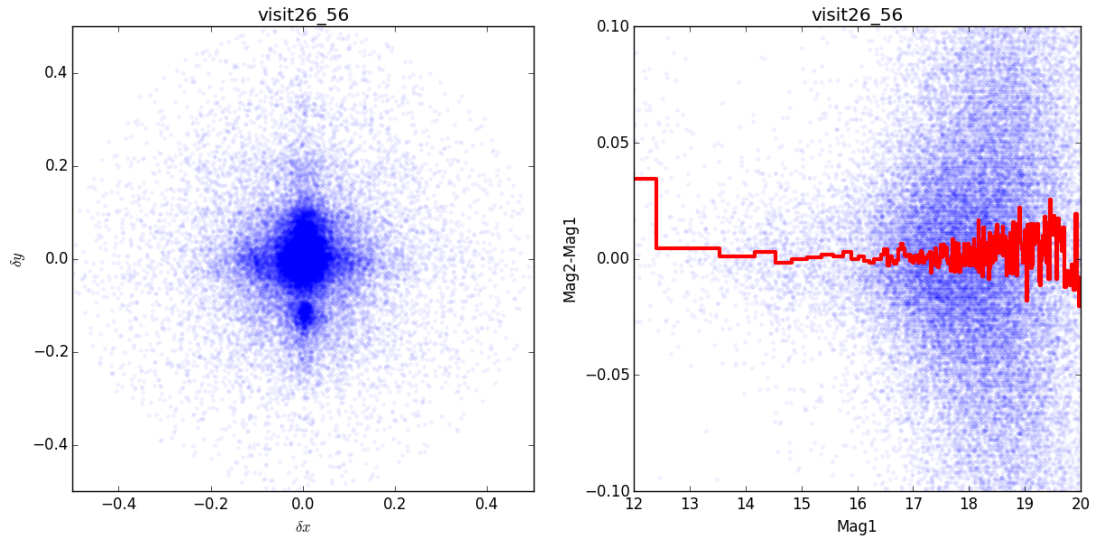

Figure 10. Same as Fig. 8, but for the last two images of Visit 26.

age was processed independently; we did not use source positions from one image in generating source fluxes in another image. Approximately 50,000 sources were identified in each exposure. To match sources between images we used our own custom-built PYTHON script. We eliminated all objects with degraded properties and matched objects in any two images based on the pixel positions. As part of the process of matching objects, we allowed for a small x-y shift in the position between the two images (which were typically of the order 0.01 pixels in each direction).

The results of the source matching and the difference in measured magnitudes of the first and second external exposure for Visit 35 are shown in Fig. 8. Bright sources are 0.03 magnitudes brighter in the second exposure than in the first image. The magnitude difference decreases at fainter magnitudes. The situation is similar for the first two images of Visit 26 as shown in Fig. 9. By contrast, as is shown in Fig. 10, the differences between the magnitudes measured between exposures 5 and 6 of Visit 26 are almost non-existent. Evidently, at least for this exposure time, by the time one has obtained 5 images, the detector is fully conditioned. The results shown here are typical of other comparisons which we made with DOLPHOT, and are broadly consistent with the results obtained with simple aperture photometry using DAOPHOT.

The results obtained from the pixel to pixel comparisons in Section 3 were presented as percentage changes; in that case we found typical brightness increases of 2-4%. The results here are presented as magnitude differences. A 2-4% change in counts corresponds to a magnitude change of 0.05-0.09 magnitudes, which should be compared to a typical magnitude change with DOLPHOT of order 0.03 magnitudes. Although we do see the same general trends in all three cases, the magnitude of the effect appears to be somewhat less with DOLPHOT.

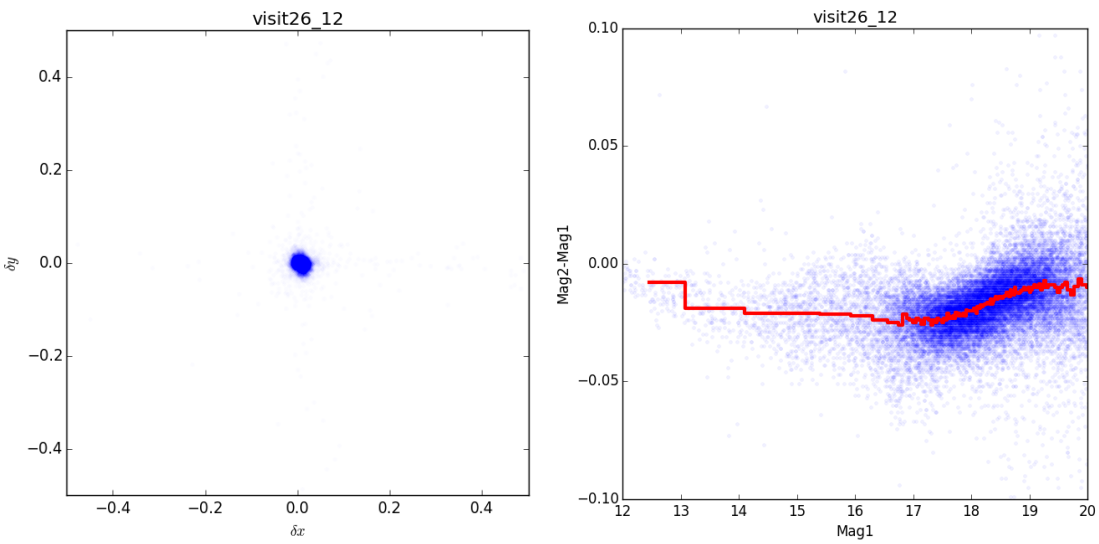

Figure 11. Left: The pixel offsets in matched sources in the first and second images of Visit 26 measured using the Anderson PSF photometry program. Right: The difference in Vega magnitude of sources in the first and second Omega Cen images as a function of the magnitude in the first image. Individual objects are plotted in blue while the average in various magnitude bands are plotted in red.

4.3 Anderson's photometry package

To check whether the results obtained with DAOPHOT and DOLPHOT were sensitive to the specific photometry package used, we have also analyzed the images with a photometry package being developed by Jay Anderson that is based on earlier work by Anderson & King (2000)[3]. The specific FORTRAN routine used here (IMG2XYM_WFC3IR) operates as an executable on individual images, fits a local sky to the data and carries out PSF fitting allowing for variations in the PSF across the face of the detector. It uses the IR image as well as a tabulated set of PSFs and returns a list of X-Y positions, fluxes for each star, and a quality factor for the measurement of the source flux.

As with DOLPHOT, we processed each of the images separately. Approximately 25,000 sources were identified in each image. We then excluded the small number of sources with uncertain flux measures, those with quality factors greater than 0.5, and matched the remaining sources on the basis of their pixel position.

The results of the source matching and the difference in measured magnitudes of the first and second external exposure for Visit35 are shown in Fig. 11. As with DOLPHOT, bright sources are brighter in the second exposure than they are in the first image though in this case the magnitude difference seems to increase slightly as one

[3]Anderson (private communication) is currently working on an updated version of this software in conjunction with developing more accurate PSFs for use with the routine. When this is completed, both the PSFs and the software will be released.

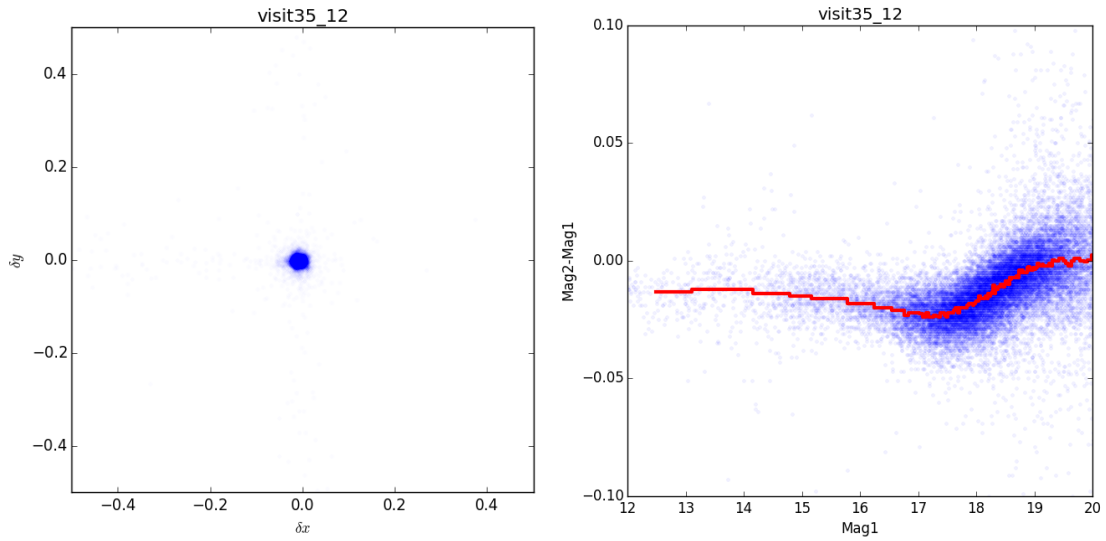

Figure 12. Same as Fig. 11, but for the first two images of visit 35.

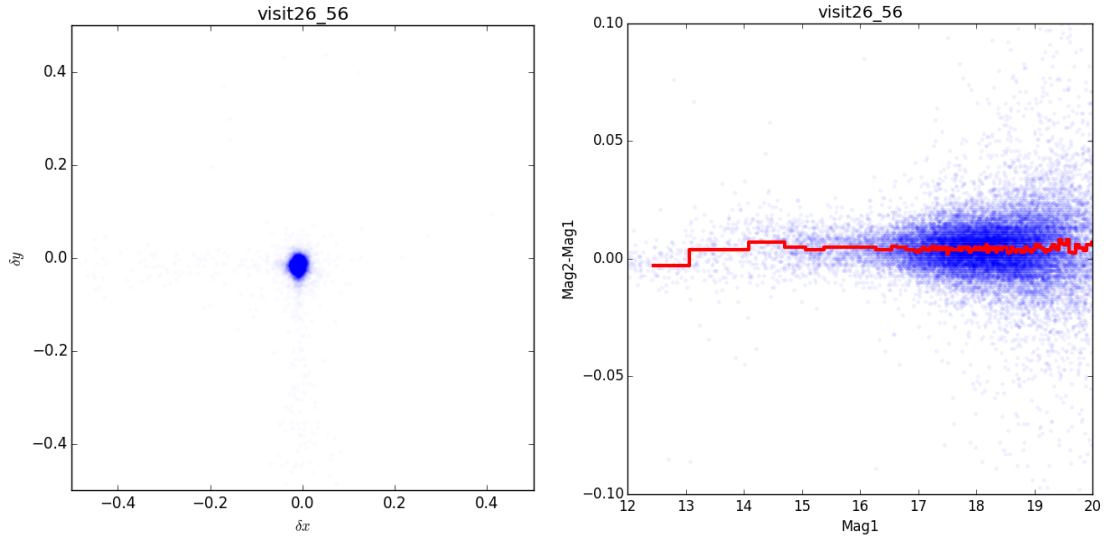

Figure 13. Same as Fig. 11, but for the last two images of visit 26.

goes from 12 to 17.5 mag, before decreasing at fainter magnitudes. The situation is similar for the first two images of visit 26 as shown in Fig. 12. By contrast, as is shown in Fig. 13, the difference between the magnitudes measured between exposures 5 and 6 of Visit 26 are almost non-existent. As was the case with DOLPHOT, the magnitude differences are of order 0.03 magnitudes, although the shapes of the curves obtained with Anderson's program differ somewhat from those obtained with DOLPHOT and the dispersion in magnitudes at any particular magnitude is clearly less.

5 Discussion

The analysis of of undithered observations in program 12694 shows (once again) that pixels in the NIR channel of WFC3 change their behavior depending on their previous exposure history. The effects in this case however are fairly small, but need to be considered if one's science goals are impacted by changes in photometry at the level of a few hundredths of a magnitude in situations where the source being measured has the same fluence as the source in the earlier exposure.

One reasonable question is whether the excess signal we see in these repeat observations can be accounted for simply as persistence within the context of the current persistence model. To find out, we have carried out the analysis described in Section 3 on the persistence-corrected flt files. Above saturation, we find that about 1/3 of the excess flux is removed; below saturation there is very little persistence predicted by our current persistence model. As a result, there is very little point in using the persistence-corrected flt files to mitigate changes in pixel behavior in repeated observations like the Omega Cen observations discussed here. There are at least two explanations for this. First, the current persistence model is derived from power law fits to persistence that span a time typically from about 300 s to more than 5,000 seconds; as a result, departures at short delay times from a simple power law power law are not captured well by the current persistence model. A Cycle 22 calibration program is intended to rectify this deficiency in our model. Second, our persistence model is derived from darks after bright exposures. We have observed that excess persistence can be "stimulated" by a bright (usually Tungsten) exposure (Ryan and Baggett, 2015). This is exactly the situation in the observations in this program.

The results here all involve 349 s exposures through the F125W filter. In the case, of persistence, we know that the amount of persistence varies with the exposure time. To the extent that the physical effects that cause the flux in repeat observations to rise involve traps, the details of the flux rise in repeat observations will also vary depending on the length and separation of exposures. Unfortunately (for this purpose) there are no other observations in the archive that are appropriate for analysis. Nearly all IR observations are dithered. To verify this, we conducted a search of the archive for observations that were repeated without dithering. We did find some examples of exposures taken at the same position, but generally these involved a short and a long exposure through the

same filter, or exposures with different filters. None proved suitable for detailed analysis. We cannot rule out, therefore, that the effect we see would be somewhat larger with different exposure sequences.

Characterizing this phenomenon generally would be quite difficult, especially if one wanted to go beyond establishing trends with, for example, exposure time. It would most likely require a substantial commitment of observing time, and it is not clear that this is warranted from a scientific perspective. We already know that persistence varies as a function of fluence and exposure time and position on the detector.

Since effectively all observations with the IR array on WFC3 are dithered, the real issue for observers is not whether there is a small offset in the apparent magnitude of a star as a result of repeated observations. The real issue is whether there is a pixel that had a bright star in an early image and a faint star in a later image, and one wishes to measure the magnitude of the faint star. There are least 3 ways to guard or check against this: (1) Use the persistence predictions to determine whether this has in fact happened, and then discard the measurement in the exposure that was impacted (which is the general advice we provide for observers on how to avoid persistence effects). (2) Examine the products of AstroDrizzle, and in particular compare the "single_sci" images to the original flt files to see whether there are specific stars that appear to have "varied" in the different exposures. (3) Carry out photometry on the individual (flt) rather than the combined drizzled images so measurements that are in doubt can be discarded.

6 Conclusions

Here we have examined changes in pixel fluxes and in stellar photometry in repeated undithered exposures of the IR detector on WFC3. We find that stars tend to brighten by several hundredths of a magnitude in repeated undithered exposures of duration 349 s over a fairly large range of fluence levels. The effect is sufficiently small that it is unlikely to be significant for most staring mode observations with HST. It does have the potential to affect the precision of both relative and absolute photometry, however, so observers need to be aware that the issue exists, and that it can affect stars in the same field differently depending on the dithering pattern. The phenomenon is larger, 0.02 - 0.04 mag, than predicted by our current model of persistence in WFC3/IR. Whether this is simply due to the fact that we do not have a good calibration of persistence at short time delays, or whether it is due to some other phenomenon is unclear.

Acknowledgments

We thank Jay Anderson for sharing with us the current version his PSF-fitting software, and Nor Pirzkal for a careful reading of this manuscript.

References

Anderson, J., & King, J., 2006, "PSFs, Photometry, and Astrometry for the ACS/WFC", ACS ISR 2006-01

Anderson, J., King, I., 2000, "Foward High-Precision Astrometry with WFPC2. I. Deriving an Accurate Point-Spread Function', PASP, 112, 1360

Berta, Z. K., Charbonneau, D., De sert, J.-M., et al. 2012, "The Flat Transmission Spectrum of the Super-Earth GJ1214b from Wide Field Camera 3 on the Hubble Space Telescope", ApJ, 747, 35

Brammer, G., Pirzkal, N., McCullough, & Mackenty, J. 2014, "Time-varying Excess Earth-glow Backgrounds in the WFC3/IR Channel", WFC3 ISR 2014-03

Dolphin, A 2000., "WFPC2 Stellar Photometry with HSTPHOT", PASP, 112, 1383

Hilbert, B & Petro L., 2012, "WFC3/IR Dark Current Stability," WFC3 ISR 2012-11

Hilbert, B. 2014, "Updated non-linearity calibration method for WFC3/IR," WFC3 ISR 2014-17

Kreidberg, L., et al. 2014,"Clouds in the atmosphere of the super-Earth exoplanet GJ1214b", Nature, 505, 69

Long, K. S., Baggett S. M., & MacKenty, J. W. , 2013a, "Characterizing Persistence in the WFC3 IR Channel: Finite Trapping Times," WFC3 ISR 2013-06

Long, K. S., Baggett S. M., & MacKenty, J. W. , 2013b, "Characterizing Persistence in the WFC3 IR Channel: Observations of Omega Cen", WFC3 ISR 2013-07

Long, K. S., Baggett S. M., & MacKenty, J. W. , 2015, "Persistence in the WFC3 IR Detetor: An Improved Model Incorporating the Effects of Exposure Time", WFC3 ISR 2015-15

Long, K.S., Baggett, S.M., MacKenty, J.W., and Riess, A.G., 2012, "Characterizing persistence in the IR detector within the Wide Field Camera 3 instrument on the Hubble Space Telescope," Proceedings of the SPIE, 8442, 84421W-9

Ryan, R., Baggett S. M., 2015, "Internal Flatfields for WFC3/IR", WFC3 ISR 2015-11

Smith, R.M., Zavodny, M., Rahmer, G. & Bonati, M., 2008a, "A theory for image persistence in HgCdTe photodiodes," Proceedings of the SPIE, 7021, 70210J-1

Smith, R. M., Zavodny, M., Rahmer, G., Bonati, M., 2008b, "Calibration of image persistence in HgCdTe photodiodes," Proceedings of the SPIE, 7021, 70210K-1

Stetson, P. B. 1987, "DAOPHOT - A computer program for crowded-field stellar photometry," PASP, 99, 191

Wakeford, H. R., et al. 2013, MNRAS, "HST hot Jupiter transmission spectral survey: detection of water in HAT-P-1b from WFC3 near-IR spatial scan observations", 435, 3481

Williams, B, *et al.*. 2014, "The Panchromatic Hubble Andromeda Treasury. X. Ultraviolet to Infrared Photometry of 117 Million Equidistant Stars", ApJS, 215, 9