



STScI | SPACE TELESCOPE
SCIENCE INSTITUTE

WFC3 Instrument Science Report 2018-03

Persistence in the WFC3 IR Detector: Intrinsic Variability

Knox S. Long, & Sylvia M. Baggett
March 29, 2018

ABSTRACT

When the WFC3 IR detector is exposed to a bright source or sources, the sources can appear as afterimages in subsequent exposures, a phenomenon known as persistence. This can affect the science obtained with the IR channel. We have been involved in an effort to predict the brightness of the afterimages so that users can (at a minimum) flag the affected pixels and remove them from their analysis or (even better) subtract the afterimages from their science images to salvage the data. The ability of any model to remove afterimages depends on the degree to which persistence is the same for identical sets of exposures. We investigate possible time variability of persistence in the WFC3 detector using sets of (almost) identical visits comprised of single exposures of Omega Cen followed by a series of darks in which persistence is measured. We analyze 8 data sets, each consisting of two or three identical visits, with stimulus exposures between 49 and 1199 s, and find clear evidence of variability in several of the datasets in darks taken within 1000 s of the stimulus exposure. In most of the datasets, the difference in persistence for saturated pixels in the stimulus exposure is $< 0.01 e s^{-1}$ for darks taken 1000 s after the initial exposure. One of three 274-second visits has significantly more persistence than its two identical visits. Persistence in this visit was higher in all 4 detector quadrants. The persistence in all three visits is well modeled as a power law decay; the visit with higher persistence has a higher power law amplitude. There was nothing unusual about the observing conditions preceding and during each of these visits that can explain the discrepancy in persistence levels. Variation in persistence implies that: (1) Unless and until the source of the variability is understood, any persistence model for the WFC3 array will be limited in its ability to predict persistence in a single observation, and, (2) as a consequence, users should always carefully inspect the results of any attempt to subtract persistence from WFC3 IR data based on a model prediction.

1 Introduction

Persistence is the afterglow produced in an IR detector of sources from earlier exposures. The amount of persistence in the WFC3 IR array depends on the degree of saturation in the earlier exposure, the time since the earlier exposure, and the exposure time of the earlier exposure. Persistence is higher when 1) the saturation in the earlier exposure, hereafter the stimulus image, is higher 2) the exposure time in the stimulus image is longer and 3) the time delay between the stimulus image and the observer’s science image is shorter.

As discussed by Long, Baggett, & MacKenty (2015a), the area-averaged persistence of the WFC3 IR detector can be described in terms of a model of the form:

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

where A and γ are functions of fluence (the effective number of electrons produced in the exposure) and exposure time. This is the model currently being used to produce the persistence data products for all WFC3/IR observations (available through MAST). Model parameters were determined from the 8 visits of Program 13572. Each visit started with a single external image of a globular cluster (either 47 Tuc or Omega Cen), the so-called “stimulus image”, followed by a long series of darks in which persistence was measured. The exposure times in the stimulus images ranged from 49 to 1403 s. Interpolation is used to estimate persistence for exposure times not used in the persistence calibration programs. The full persistence model includes a “correction flat”, described by Long, Baggett, & MacKenty (2015b), which provides a zeroth-order correction for changes in persistence amplitude across the face of the detector.

We are in the process of developing a more sophisticated model of WFC3 IR persistence which involves measuring A and γ in subsections of the detector. As part of this effort we have acquired additional data for calibrating the model, including a number of datasets that are nearly identical to one another. Here we investigate the degree to which the persistence in these identical visits is the same, or whether persistence itself is time-variable. This is a key question as any variations in persistence place fundamental limitations on the accuracy of any persistence model.

2 Data

The data used in this analysis consist of a selection of the observations from Cycles 18 through Cycle 23 obtained explicitly for studying the WFC3 IR persistence. These visits begin with a single external exposure of a globular cluster followed by a long series of darks. We only consider identical visits taken with the same filter and the same

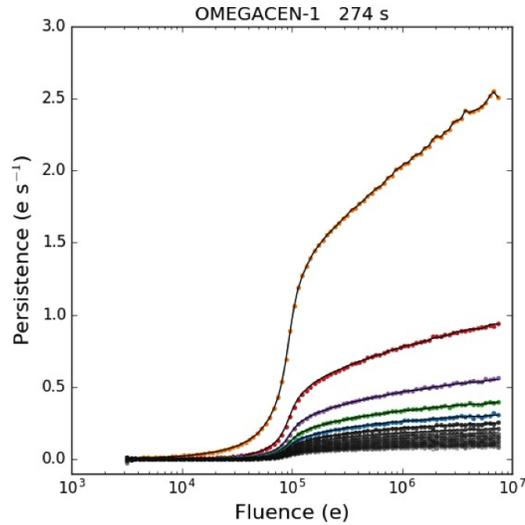


Figure 1. The persistence following visit 1 of program 12351, which consisted of a single 274 second exposure of Omega Cen followed by a series of darks. Each curve represents the persistence measured in one of the darks. Persistence is highest in the first dark (top curve) after the stimulus image and decays to progressively lower levels in each successive dark (lower curves).

exposure time, and whose dark sequences are obtained with the same exposure times and time delays from the stimulus exposure. A significant caveat is that even for these identical visits some differences will inevitably remain. The most obvious disparity is different illumination levels in the individual pixels of the stimulus images due to offsets in pointing position and field of view orientations between the visits. Since persistence varies somewhat across the detector array, this could cause some apparent variability between different visits even if there is no intrinsic variability in detector persistence.

Table 1. Observation Log

ProgID	Visit	Obs Date	Dataset	Target	Filter	Exp. (s)	Saturated (%)
14381	71	2016-03-23	id1s71o5q	OMEGACEN-1	F140W	149	4.6
14381	81	2016-03-31	id1s81whq	OMEGACEN-2	F140W	149	4.4
14381	91	2016-04-05	id1s91krq	OMEGACEN-3	F140W	149	4.3
12351	01	2011-01-27	ibmf01vfq	OMEGACEN-1	F110W	274	13.5
12351	03	2011-02-01	ibmf03nbq	OMEGACEN-3	F110W	274	12.7
12351	AC	2011-05-25	ibmfaenyq	OMEGACEN-2	F110W	274	13.2
12694	11	2012-02-27	ibvd11aaq	OMEGACEN-1	F125W	349	9.6
12694	31	2012-04-27	ibvd31vuq	OMEGACEN-3	F125W	349	9.2
14381	D1	2016-05-09	id1sd1afq	OMEGACEN-1	F125W	499	14.9
14381	E1	2016-06-03	id1se1brq	OMEGACEN-2	F125W	499	14.7
14381	F1	2016-06-10	id1sf1xyq	OMEGACEN-3	F125W	499	14.8
14381	11	2016-01-12	id1s11jiq	OMEGACEN-1	F125W	599	18.5
14381	21	2016-01-19	id1s21j0q	OMEGACEN-2	F125W	599	18.4
14381	31	2016-02-02	id1s31ppq	OMEGACEN-3	F125W	599	17.1
14381	A1	2016-04-15	id1sa1haq	OMEGACEN-1	F127M	899	5.1
14381	B1	2016-04-29	id1sb1w6q	OMEGACEN-2	F127M	899	5.1
14381	C1	2016-05-20	id1sc1dfq	OMEGACEN-3	F127M	899	5.0
14381	41	2016-02-12	id1s41dkq	OMEGACEN-1	F127M	1199	7.3
14381	51	2016-03-07	id1s51csq	OMEGACEN-2	F127M	1199	7.2
14381	61	2016-03-16	id1s61ohq	OMEGACEN-3	F127M	1199	6.8

The observations are listed in Table 1 ordered by exposure time. The table lists the (1) program number, (2) visit number¹, (3) observation date, (4) dataset name, (5) target name (where -1, -2, -3 denote different pointings), (6) filter, (7) exposure time and (8) percentage of pixels saturated in the external exposure. There are 7 different stimulus exposure times ranging from 149 s to 1199 s distributed across 20 different visits, about three visits per exposure time.

As a preliminary step for the data analysis, we reprocessed all of the data with version 3.1.6 of CALWF3. To have the calibration software apply the gain correction to the dark exposures, we activate the flatfield correction and use a “unity flat” (containing only 1’s). After subtracting our best estimate of the dark level in each dark exposure, we construct persistence curves for all darks within each visit using a standard set of stimulus levels. An example of these persistence curves for one of the visits is shown in Fig. 1. Based on inspection of the curves for a number of other visits with different stimulus exposures, those shown in Fig. 1 are fairly typical.

3 Analysis

Our goal in this study is to determine whether persistence varies between identical sets of exposures. Variations in persistence levels could occur in the observations we are using for a number of reasons including 1) changes in the distribution of bright stars across the detector due to pointing offsets, 2) uncertainties in background subtraction due to variations in the dark rate², and 3) intrinsic variations i.e. non-repeatability in persistence.

Because the visits in our sample are identical it is possible to perform tests that are largely model-independent to evaluate whether, under nominally similar observing conditions, we measure the same persistence levels. The outcome should provide insight into the reliability of our model of the average persistence.

The most straightforward comparison of two identical visits is to simply subtract their persistence curves (like those shown in Fig. 1) from one another. An example is shown in Fig. 2 which compares 3 visits of 899-second exposures of Omega Cen followed by identical sets of darks. The dark current in the WFC3 IR channel varies from exposure to exposure (see, e.g. Sunnquist, Baggett & Long 2017); we assume that there is no persistence in pixels exposed to a stimulus of less than 10,000 e and thus take the median value of those pixels in the dark image to be a representative estimate of the dark current. As is evident in Fig. 2, we measure almost no variation in persistence below about 40,000 e. This is not surprising since there is very little persistence due

¹For scheduling reasons, in program 14381 each visit was actually broken into two visits scheduled immediately after one another (e.g. if the original full visit was 81, it was split into two smaller visits: 81 and 82). For simplicity, in the table we list only the first visit, i.e. the one in which the external exposure was acquired.

²The mean rate is about 0.05 e s^{-1} and the standard deviation about the mean is about $\pm 0.03 \text{ e s}^{-1}$ (Hilbert & Petro 2012, Sunnquist, Baggett, & Long 2017).

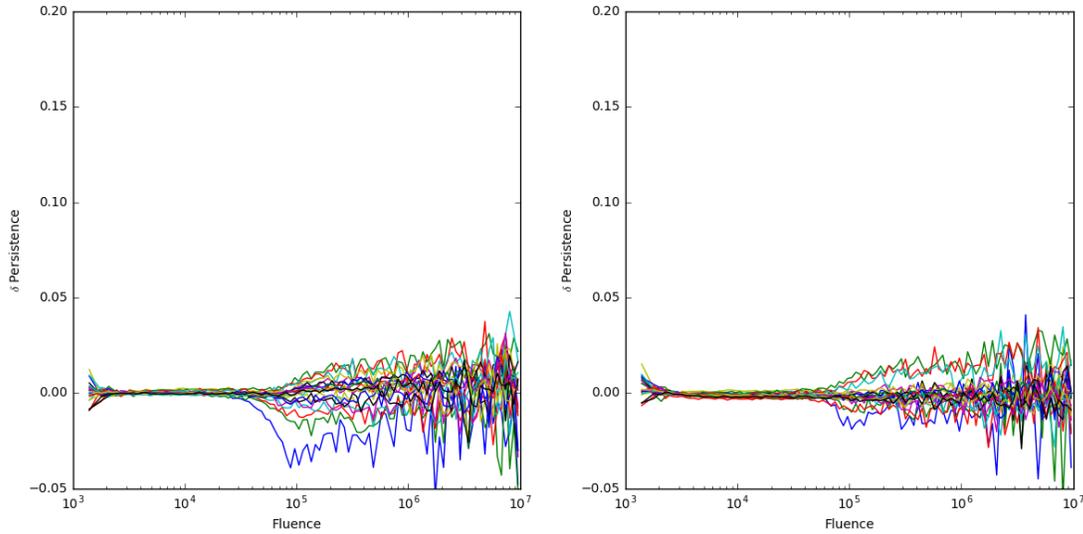


Figure 2. Left: The difference in persistence as measured in between visits A1 and B1 of Program 14381, both of which began with a stimulus exposure of 899 s and had the same time delays for the follow-on darks. Right: Identical to the left panel, except for visits A1 and C1 of the same program. The differences between the first few darks after the stimulus are plotted in lines of various colors, the first one being blue. Differences between later darks are all plotted as black.

to stimuli in this range (as can be seen by inspecting Fig. 1). At higher fluence levels, we do see some variations, with the largest difference occurring in the first few darks after the stimulus exposure, where the persistence is largest. In the remaining darks, the differences are typically less than 0.01 e s^{-1} , as we will quantify later. Similar figures for the remaining visits with other stimulus exposure times appear in Appendix 1 or in one case with more evidence of persistence variability later in the main body of this report.

We would not necessarily expect a comparison between the visits to produce curves that are identically zero. The specific pixels that are illuminated in the various visits differ, and the amplitude of persistence varies by of order 10% across the face of the detector. Read noise for individual pixels is significant; for a dark exposure time of 350 s and a read noise of 16 e, the equivalent current is 0.046 e s^{-1} , so we rely on the averaging of typically 1000 (a few 1000 in some cases) pixels to reduce the measurement error. One way to estimate the error would be to compare trends in the data to the scatter in the individual plots. From inspection of Fig. 2 and those in the Appendix, it is fairly clear that there are small differences in persistence at early delay times (the colored lines) and that these differences decay later (the black lines).

Although one can make “by-eye” estimates of how much persistence varies from plots like those presented in Fig. 2, one would prefer something a little more systematic. For this purpose, we have calculated the mean and standard deviations in persistence

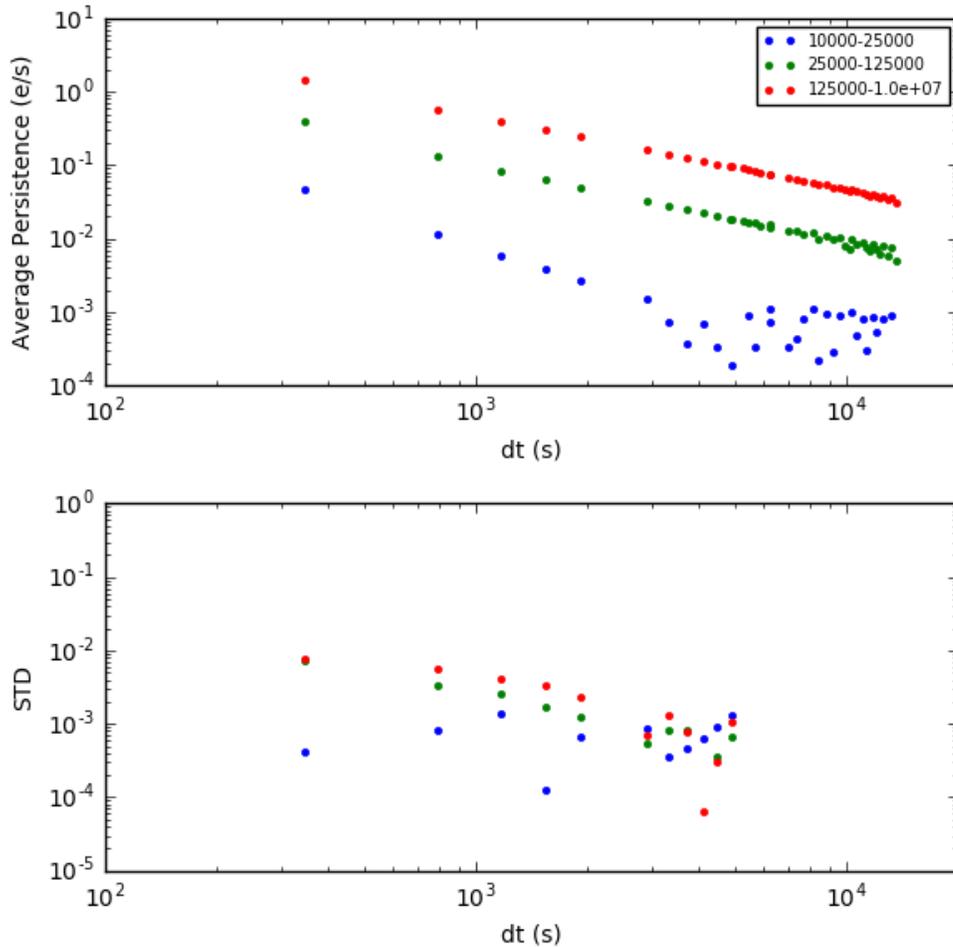


Figure 3. Top: The average persistence for visits A1, B1, and C1 of program 14381, which all began with a 899 s exposures of Omega Cen through the F125W filter. The three curves represent three different fluence ranges in the stimulus images, 10,000-25,000 e, 25,000 - 125, 000 e, and 125,000 to 10,000,000 e. Bottom: The difference between the measured values, represented as the standard deviation (STD) at each fluence level.

over three broad ranges in of stimulus, 10,000 - 25,000 e, 25,000 - 125, 000 e, and 125,000 - 10,000,000 e for each of the various exposure times as a function of the time delay. The results of such a calculation are shown in Fig. 3 for the visits that began with an 899 exposure of Omega Cen. Specifically, the upper panel summarizes the average value of the persistence in the three ranges as constructed from curves like those shown in Fig. 1. The persistence in all three ranges decays as a power law. The power law index is steeper in the 10,000 - 25,000 e band, and after about 4,000 s it is clear that the persistence signal is becoming difficult to detect at all. For the two higher flux bands, the power law continues throughout the observations. The lower panel shows the standard deviation of the average persistence for the 3 visits in the same 3 bands. For the lower band the standard deviations are about 10^{-3} e s^{-1} throughout the observations, whereas they are about $5 - 7 \times 10^{-3} \text{ e s}^{-1}$ in the two higher bands at times less than 1000s then drop to $2 - 3 \times 10^{-3} \text{ e s}^{-1}$.

To allow easy inter-comparison of the amount of apparent variation we have tabulated the results in Tables 2 and 3 for delay times less than 1000 s and between 1,000 and 3,000 s, respectively. The tables give the mean persistence, the standard deviation of the mean persistence, and the ratio of the two for each of the 3 ranges of stimulus: 10,000 to 25,000 e, 25,000 to 125,000 e, and 125,000 to 10,000,000 e. While the standard deviation is a well-defined way to estimate variations in a dataset, it is not the same as the maximum variation (which for two datasets would be twice the standard deviation.) At low stimulus levels the ratio of the standard deviation to the mean persistence is high, especially after 1000 s. Any model we construct at low stimulus levels will have significant errors. Fortunately, persistence is quite low in this stimulus range, so this does not have much practical effect on models to subtract persistence. At the highest stimulus levels, the standard deviation in the persistence is in most cases only a few percent of the mean, which suggests that variations in persistence are not too large, as a fraction of the total persistence. Unfortunately, if one is trying to subtract persistence from an image, the ratio may be less important than the absolute variation. Finally, there is one case that appears to be out of family with the rest of the of the datasets, and that is the series of visits with a stimulus exposure time of 274 s.

Table 2. Variations less than 1000 s after Stimulus

Exp	10,000 - 25,000 e			25,000 - 125,000 e			125,000 - 10,000,000 e		
	Mean	STD	Ratio	Mean	STD	Ratio	Mean	STD	Ratio
149	0.0058	0.0006	0.469	0.1171	0.0049	0.042	0.5870	0.0119	0.021
274	0.0102	0.0002	0.074	0.2110	0.0110	0.053	0.9325	0.0376	0.040
499	0.0114	0.0030	0.367	0.1613	0.0028	0.021	0.7410	0.0070	0.010
599	0.0152	0.0012	0.172	0.1798	0.0052	0.027	0.7633	0.0118	0.015
899	0.0293	0.0006	0.042	0.2613	0.0053	0.022	1.0008	0.0066	0.008
1199	0.0315	0.0008	0.057	0.2791	0.0032	0.013	0.9947	0.0059	0.007

Table 3. Variations 1000 to 3000 s after Stimulus

Exp	10,000 - 25,000 e			25,000 - 125,000 e			125,000 - 10,000,000 e		
	Mean	STD	Ratio	Mean	STD	Ratio	Mean	STD	Ratio
149	-0.0005	0.0005	0.976	0.0321	0.0017	0.052	0.1938	0.0046	0.024
274	0.0004	0.0002	3.089	0.0365	0.0023	0.064	0.1955	0.0083	0.043
499	0.0002	0.0017	1.083	0.0462	0.0009	0.021	0.2620	0.0032	0.012
599	0.0015	0.0014	0.930	0.0529	0.0013	0.023	0.2733	0.0035	0.012
899	0.0035	0.0008	0.272	0.0574	0.0015	0.025	0.2739	0.0026	0.009
1199	0.0039	0.0005	0.125	0.0648	0.0012	0.020	0.2832	0.0022	0.008

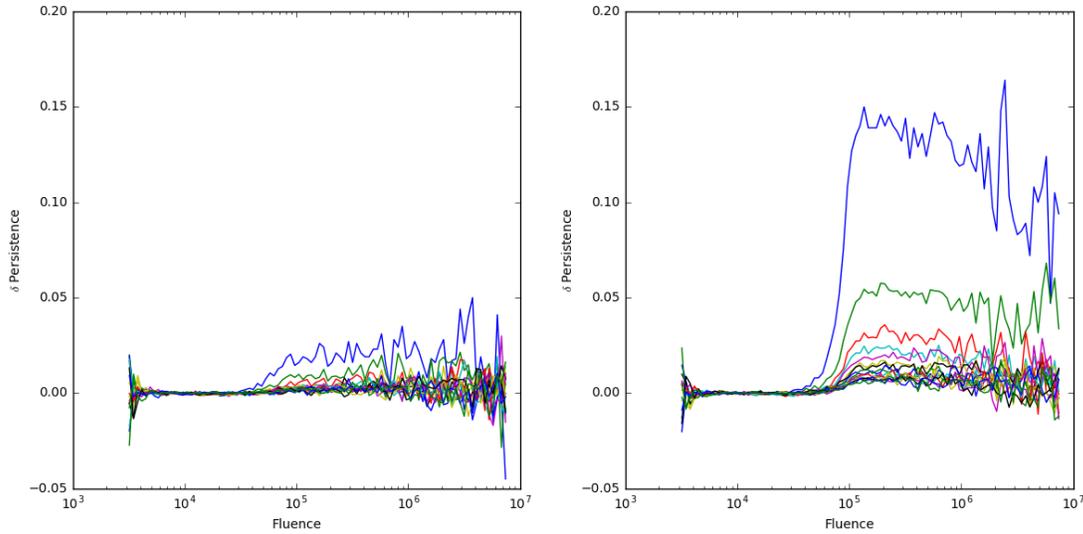


Figure 4. Left: The difference in persistence as measured between visits 01 and 03 of Program 12351 in each of the follow-on darks. Right: Identical to the left panel, except for visits 01 and AC. Persistence is considerably larger in visit AC, especially in the first few darks after the stimulus exposure.

4 An Outlier Among the Datasets

As noted above, the variation in the standard deviation between the 3 visits with 274 s external exposures is significantly more than for the other exposure times. This was also apparent from the direct comparison of the persistence measured in the 3 visits as is shown in Fig. 4. The persistence between two of the visits (01 and 03) is nearly the same: the difference is less than 0.02 e s^{-1} in the first dark after the stimulus exposures and less than 0.01 e s^{-1} in the remaining darks that followed. The story changes, however, when visit 01 is compared to visit AC, as is shown in the right panel. The persistence is considerably higher in visit AC than it was in either visit 01 or 03, and it remains higher in all of the subsequent darks.

We have been unable to find any reason that would explain why visit AC is anomalous. The stimulus image of all 3 visits consisted of a 274 s external exposure acquired with the F110W filter. The history of exposures preceding visit AC, available in the MAST persistence logs, looks completely innocuous in terms of anything that could have caused significant persistence. Although the exposures were taken at different positions within Omega Cen, all three exposures saturated a similar fraction of the detector, about 13% of the pixels. In addition, as shown in Fig. 5, the distribution of bright stars across the field of view appears relatively similar in each of the three stimulus images.

Persistence in the WFC3 IR is known to vary somewhat across the face of the detector; in particular, persistence in quadrant 1 of the array (upper left) is higher than in quadrant 3 (lower right), as discussed in Long, Baggett and MacKenty (2015b). Fig.

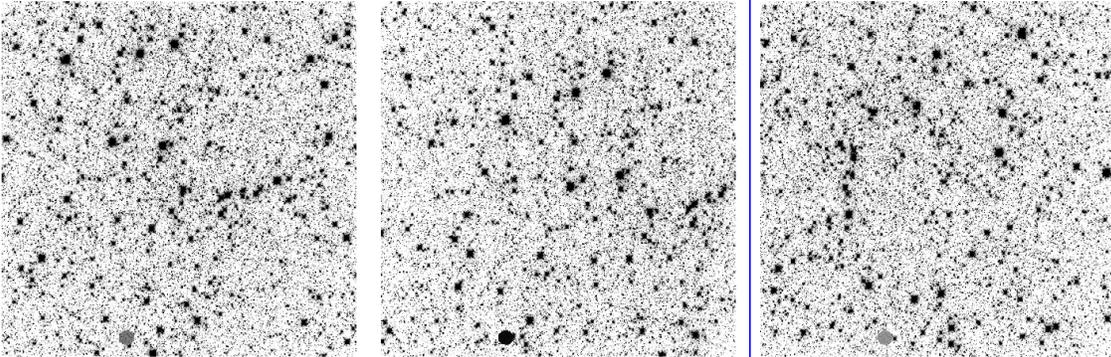


Figure 5. The first stimulus exposures for visits 01, 03 and AC. There are perhaps more bright stars in quadrant 4 for visit AC than the other two images, but there is nothing to suggest a cause for the anomalous persistence in visit AC.

6 shows the average persistence in visit 01 and visit AC as a function of time for pixels responding to a fluence in the stimulus image between 10^5 and 10^6 e. It is clear that persistence in visit AC is higher than in visit 01 not just in one but in all four quadrants. The difference in the measured persistence is not due a difference in sources across the field.

Given that we cannot find any reason for visit AC to be anomalous, we have attempted to explore exactly what is different by fitting each of the visits in this collection with the power law model shown in Equation 1 and described by Long, Baggett, & MacKenty (2015a). The results for A and γ are shown in Fig. 7. The models are very similar for visits 01 and 03. For visit AC, the power law exponents γ are almost identical to those of visits 01 and 03 which implies that persistence decays at the same rate, but the amplitude of the persistence in visit AC is simply greater than in the other two cases.

5 Discussion

Having concluded that evidence for variations in persistence is strong in the case of the visits with a stimulus exposure of 274 s, the obvious question is whether this is an outlier or just an extreme example of variations that are generally present in the persistence. Many of the the comparison curves in Appendix 1 show differences particularly at early times that are difficult to explain as measurement error or noise. Our general conclusion is that there is a continuum of variability, though quantifying the amount of variability more than has been accomplished with Tables 2 and 3 is difficult.

There have been earlier indications that persistence is a variable phenomenon. In particular, Long et al. (2010) described a set of calibrations involving Tungsten flat field exposures followed by a series of darks. In that case, three identical visits were taken several months apart with the stimulus flat at a mean fluence level of about 2x saturation

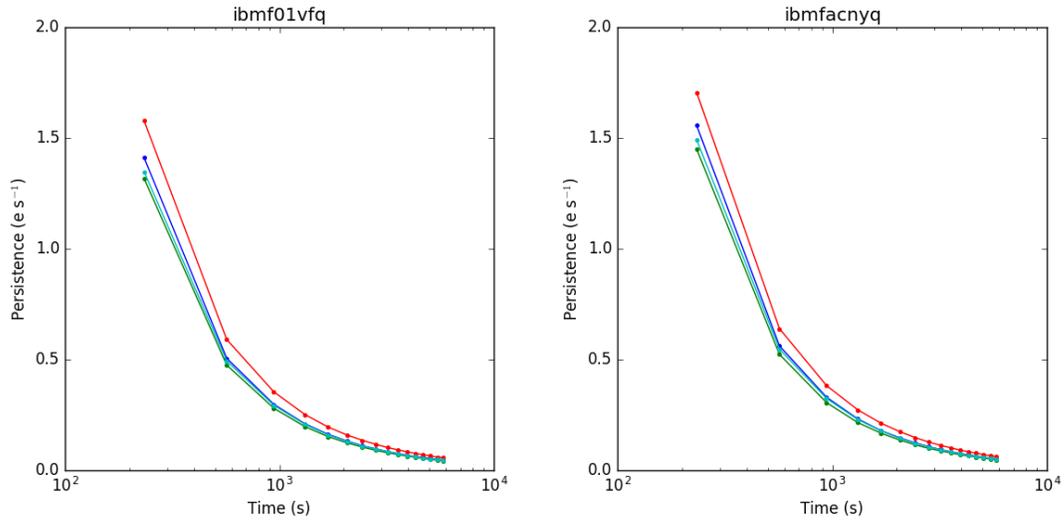


Figure 6. Left: Visit 01 average persistence in pixels responding to a stimulus fluence between 10⁵ and 10⁶ e for each detector quadrant. Right: Identical to the left panel except for visit AC. All four quadrants show more persistence in visit AC than in visit 01.

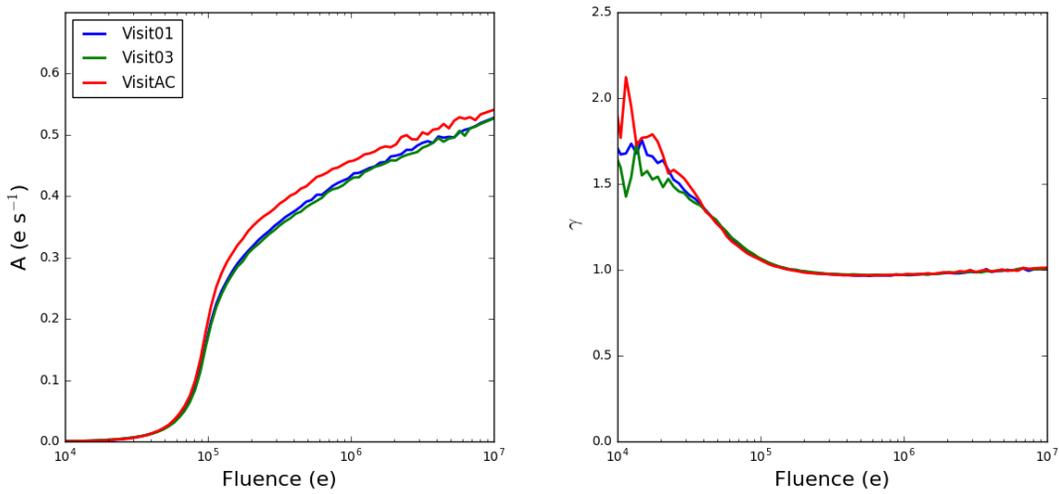


Figure 7. The derived amplitude (left) and power law exponents (right) from a power law fits to each of the 3 visits with a stimulus exposure time of 274 s.

(140000 e). One visit inexplicably showed more than twice the level of persistence than the other two visits (0.65 e s^{-1} versus 0.32 e s^{-1}). The disparity in persistence levels then and now cannot be explained by the known differences in persistence as a function of position in the detector. Investigation into the scheduling and commanding of the visits revealed nothing obviously unusual about the visits that show higher persistence. The Tungsten flat field outlier visit was closer to an SAA passage than the other two identical visits but WFC3 IR does not experience elevated dark rates following SAA passages. The IR array was used more in the day preceding visit AC (the extreme outlier in this report) than in the days preceding the other two identical visits, but none of those earlier exposures should have produced the anomalous persistence seen in AC. In addition, there are counterexamples in other visits (e.g. 71/91 and D1/E1) where the IR usage was higher for one of the visits and not the others yet the persistence was at normal levels.

Both dark current and persistence are generally associated with imperfections in the material that comprises the diodes of the detector. We know that dark current varies and that some of this variation is associated with the position of HST in its orbit (Sunquist, Baggett, and Long, 2017). The presumption is that these dark current variations arise either from small changes in the detector temperature, or more probably, from changes in some of the voltage levels related to times when the spacecraft is in the sunlight. However, the observations described here span an entire orbit, and so if there were orbital variations we should see departures from a power law in all of the visits, which we do not.

In considering the implications of this study, one should be aware that despite the fact that in this report we have focussed on searching for, finding and characterizing variability, the amount of variation we have discovered is fairly small. Even for the case of the three visits with a 274 second stimulus exposure, the variation is of order 10% of the average persistence (where here we have taken the variation to be twice the standard deviation). In the other cases, the variation is a 2-4% by the same standard. If one wishes to use these data as a basis for a model intended to predict when one should be concerned about persistence in science observations then the variations we have seen here are not a problem. However, if one wishes to subtract persistence from observations then the variations, as long as they are unexplained, limit the accuracy of the subtraction.

6 Future Work

A limitation of the current study is that we typically only have three visits at any particular stimulus exposure time. As a result, it is difficult to isolate the factors that might cause persistence in the detector to vary. In principle, one could address this problem by obtaining 10 or 20 identical visits at one or two different stimulus exposure times. This would allow an assessment of whether visit AC was truly out of family with the other

two visits taken with a stimulus exposure time of 274 seconds, or whether it was merely part of a quasi-gaussian distribution. Such a study might be able to shed some light on the underlying cause of the variability in persistence.

If persistence were a limiting factor in the science that could be achieved with the WFC3/IR channel, acquiring significantly more calibration data would be the logical course of action to take. Fortunately, persistence is not a limiting factor. There are operational procedures currently in place for restricting IR usage after observations that are expected to cause significant persistence. These procedures have successfully limited the number of cases where persistence is easily detectable in the images. The existing model of persistence is good enough to identify which pixels in an image are likely to be affected by persistence, even though it may not be sufficiently accurate to be used to remove persistence completely. These predictions, in the form of the persistence products, are available to all observers through MAST. In addition, in situations where persistence affects a small number of pixels scattered over the face of the detector and observers follow the recommended practice of dithering their observations, the pixels affected by persistence can be treated in the same manner as hot or dead pixels.

A higher priority, from a science perspective, is to encourage observers to make full use of the available persistence products. To facilitate this, one approach would be to flag the affected pixels in the science image data quality extensions. The advantage is that the effects of persistence would be taken into account in the standard products delivered by STScI, including the fit files and drizzled images which many observers use for their analysis.

7 Summary

In order to accurately remove persistence from observational data, persistence must be a stable phenomenon. Here we have analyzed a series of nominally identical visits that were to serve as a basis for a new, improved model of persistence in WFC3.

We find that persistence is a relatively stable phenomenon. Variations in persistence are typically small, usually 2-4 % of the persistence signal itself resulting from saturated pixels as measured by the standard deviation, about twice this value if one considers the maximum variation between a visit with less persistence and one with the most persistence.

That said, our analysis also implies that persistence is variable. In the worst case, the variations appears to be as large as 10% of the average persistence in a set of visits with a stimulus exposure time of 274 s. There is no obvious reason why this visit should have been an outlier.

These ostensibly intrinsic variations constitute one of the primary limiting factors to improving the persistence model. As discussed earlier, the existing model is well-suited for masking pixels affected by persistence. However, to improve the model's reliability so that the results can be used to accurately correct science data and recover

persistence-impacted pixels will require an understanding of the underlying cause of the non-repeatability in the persistence behavior.

Acknowledgements

We thank George Chapman, Merle Reinhart, Alan Welty, and Bill Januszewski for their help in crafting and executing this program and Susana Deustua for a careful reading of this manuscript.

References

- Hilbert, B & Petro L., 2012, “WFC3/IR Dark Current Stability,” WFC3 ISR 2012-11
- Long, K. S., Baggett S. M., Deustua, S., & Riess, A., 2010, “WFC3/IR Persistence as Measured in Cycle 17 using Tungsten Lamp Exposures,” WFC3 ISR 2010-17
- Long, K. S., Baggett S. M., & MacKenty, J. W. , 2015a, “An Improved Model of Persistence for the IR detector on WFC3”, WFC3 ISR 2015-15
- Long, K. S., Baggett S. M., & MacKenty, J. W. , 2015b, “Persistence in the WFC3 IR Detector: Spatial Variations”, WFC3 ISR 2015-16
- Sunnquist, B., Baggett S., & Long, K. S. 2017, “An Exploration of WFC3/IR Dark Current Variation”, WFC3 ISR 2017-04

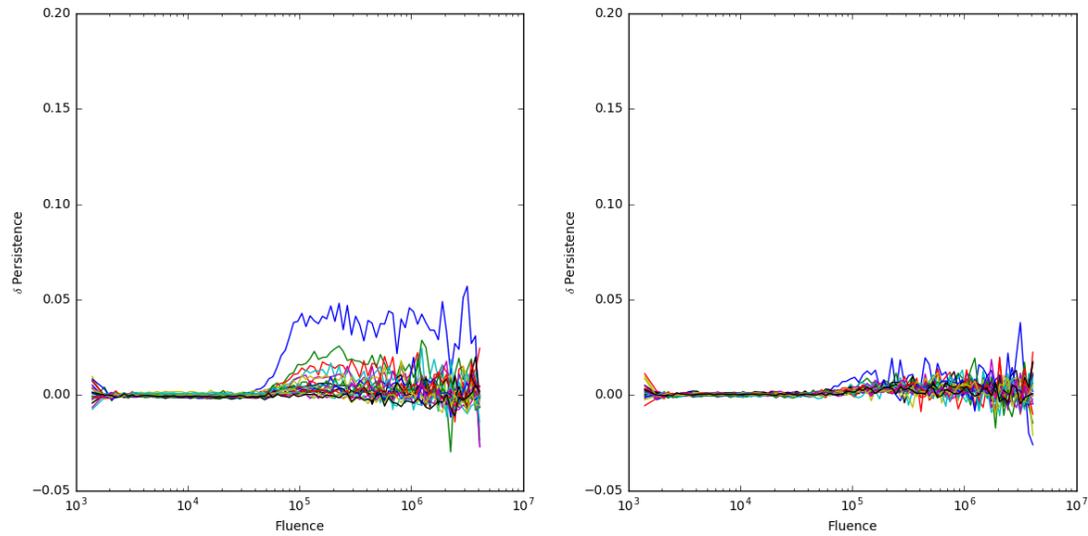


Figure A-1. Left: The difference in persistence as measured between visits 71 and 81 of Program 14381 in each of the follow-on darks. The stimulus exposure was 149 s in duration. Right: Identical to the left panel, except for visits 71 and 91.

1 Appendix

The complete set of difference comparisons for all visits in Table 1 (except those already presented in the main body of the report) are provided in this Appendix.

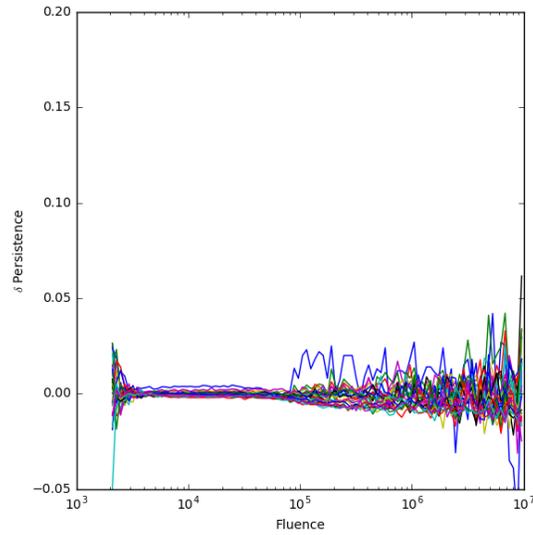


Figure A-2. Identical to one of the panels in A-1 except for visits 11 and 31 of Program 12694, in which the stimulus exposure lasted 349 s.

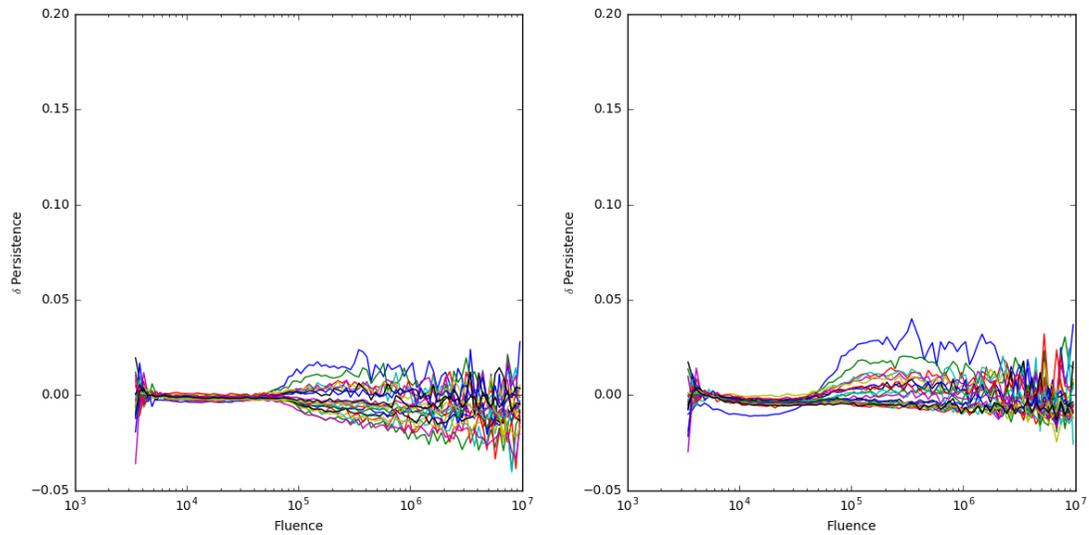


Figure A-3. Identical to one of the panels in A-1 except for visits D1 and E1 of Program 14381, in which the stimulus exposure lasted 499 s.

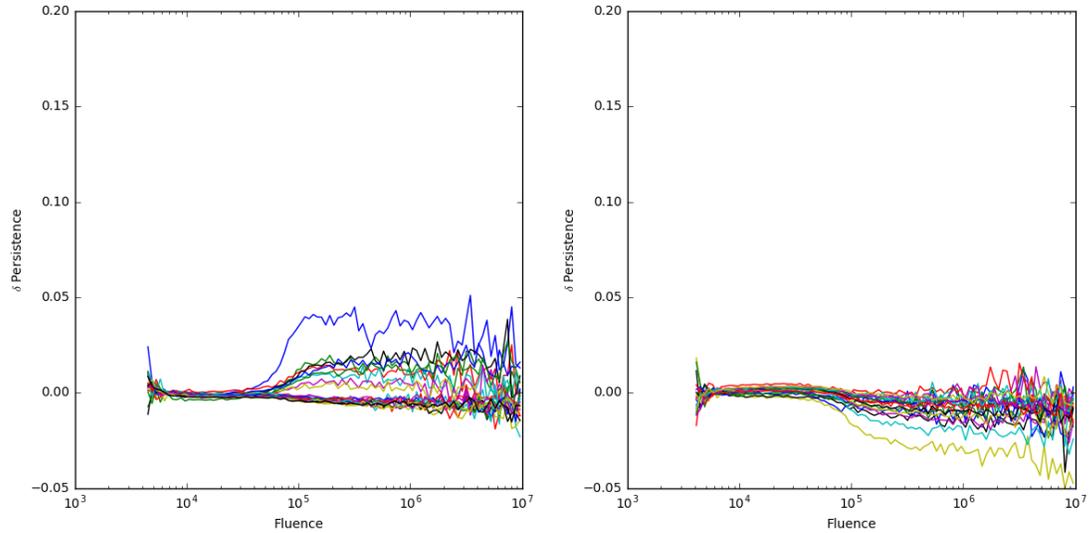


Figure A-4. Identical to Fig. A-1 except with a stimulus exposure of 599 s: left panel visits 11 and 21 of Program 14381, and right panel visits 11 and 31 of the same program.

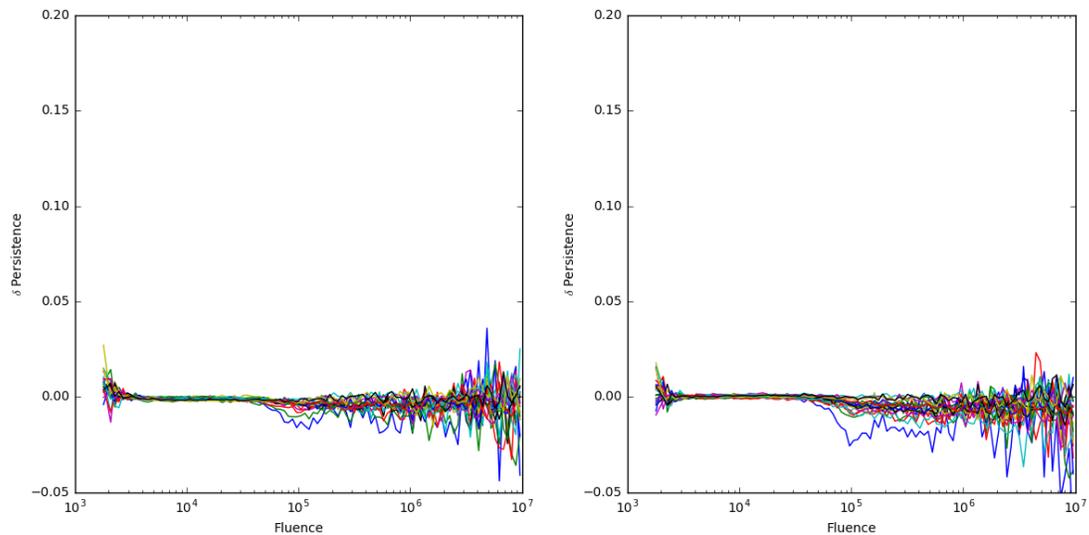


Figure A-5. Identical to Fig. A-1 except with a stimulus exposure of 1199 s: left panel visits 41 and 51 of Program 14381, and right panel visits 41 and 61 of the same program.