



Pre-Flashing WFC3/IR Time-Series, Spatial Scan Observations

K. B. Stevenson & W. Eck
September 27, 2019

ABSTRACT

Spatial scan observations using WFC3's IR channel exhibit time-dependent systematics (in the form of a ramp or hook) that have been attributed to the effects of persistence. The amplitude of these systematics is often two orders of magnitude larger than the signal sizes of interest and, therefore, must be carefully modelled and removed. The goal of this calibration program (CAL-15400) is to mitigate these systematics by continuously illuminating the detector while repeatedly reading it out during Earth occultation (termed pre-flashing). Compared to standard observations, we are able to reduce the amplitude of the systematic effect by a factor of ~ 7 (from 1.30% to -0.19%), thus confirming our hypothesis that the detector more quickly reaches an equilibrium state when subjected to higher flux levels. Compared to the latest modeling techniques (Zhou et al., 2017), we achieve a marginal improvement in the white light curve precision ($\Delta_{rms} = -8 \pm 9$ ppm); therefore, pre-flashing is an equally effective means to mitigate WFC3's instrument systematics. We conclude that pre-flashing does not warrant future consideration due to the increase in the number of channel select mechanism (CSM) motions, effort required to implement, and equivalent ability to model instrument systematics with current techniques.

Introduction

In 2011, the *Hubble Space Telescope* (HST) performed its first Wide Field Camera 3 (WFC3) time-series, spatial scan observations. This includes a test observation of GJ 1214

from program 12325 (PI: MacKenty) and two transiting exoplanet observations from a general observer program (GO-12181, PI: Deming). In the eight years since, WFC3 has made over 200 time-series, spatial scan observations, mostly targeting exoplanets as they pass in front of or behind their parent stars and fast-rotating brown dwarfs.

During this time, the standard observing strategy has been to acquire an additional orbit of data at the start of each time-series observation as a means to stabilize the detector. By excluding the first orbit, the systematics from all subsequent orbits within that observation exhibit a repeatable behavior. The repeatability of *HST*'s orbit-dependent systematics (often referred to as the “ramp”) is key to fitting time-series observations and achieving < 100 ppm precision. This ramp is typically attributed to charge trapping in the detector pixels.

Recent work by [Zhou et al. \(2017\)](#) has introduced a new, physically-motivated model based on charge-trapping theories that fits the ramp in all *HST* orbits. They find that the scatter in the residuals from the first *HST* orbit are consistent with those from subsequent orbits (within a given observation) and conclude that data from the first orbit no longer need be discarded. Out of precaution and to improve the baseline flux constraint, the WFC3 team continues to support the acquisition of a stabilizing orbit at the start of each visit.

[Long et al. \(2014\)](#) attempt to mitigate charge trapping by pre-conditioning the detector during Earth occultation (CAL-13573, PI: Long). They found that by filling detector pixels with light from the tungsten lamp and holding the charge for ~ 42 minutes before reading out the electronics, [Long et al. \(2014\)](#) were able to reduce the amplitude of the effect by a factor of ~ 2 . They also note that science observations with a larger number of scans per orbit have a shorter exponential timescale and, thus, stabilize faster. They conclude that each time a pixel is illuminated, a fraction of the available traps are filled irrespective of how long the pixel maintained its fluence level.

In reviewing their work, we adopt an alternative interpretation, namely that the detector stabilizes faster due to higher flux levels. We note that brighter targets require faster scan rates and, as a result, contain more scans per orbit; therefore, the two parameters are closely linked. We use this interpretation for the basis of our calibration program. The hypothesis is that by pre-flashing the detector with a bright source (i.e. high flux) while repeatedly reading out the electronics, the detector will more quickly reach its equilibrium state. Performing this strategy within each Earth occultation would then mitigate the ramp seen during science observations.

In the sections below, we provide the details of our observations, data reduction, and analyses; present our findings from this calibration program; and consider the impact of implementing the pre-flashing strategy on future time-series, spatial scan observations.

Label	Program ID	Visit #	Obs. Date	Median Fluence (e^-/pixel)
Pre-Flashed	15400	01-04	2018-08-10	$19,230 \pm 50$
Standard	13021	14	2013-07-06	$21,860 \pm 1300$

Table 1: *Observation summary.*

Observation Strategy

In this calibration program (CAL-15400; PI: Stevenson), we perform time-series observations of a point source in spatial scan mode in coordination with Tungsten lamp pre-flash observations during Earth occultation. We select GJ 1214 as our point source because we need a reference system that was previously observed without pre-flashing, thus providing a baseline for comparison. GJ 1214 is a relatively quiet, slowly-rotating M dwarf and the transiting planet in this system has no detectable thermal emission at these wavelengths, which would have made our analysis more challenging.

We apply the same observing strategy as the final 10 visits from GO program 13021 (PI: Bean). Namely, we obtain 103 second exposures with the G141 grism (NSAMP = 15; SAMP-SEQ = SPARS10) in the 256×256 subarray mode. Exposures scan perpendicular to the dispersion direction at a rate of $0.12''/\text{sec}$ and alternate directions (i.e., round trip mode). During test building for the calendar we find that Visits 02 – 04 are too long (as implemented by Program 13021) because they did not allow enough time for gyro bias updates. Unlike the original program, bias updates are now necessary to mitigate pointing failures. To recover ~ 5 minutes at the end of each orbit for gyro bias updates, we scheduled three fewer scans per *HST* orbit (16 vs 19). Similar to GO-13021, we obtain a direct image of GJ 1214 at the start of each *HST* orbit using the F130N filter (NSAMP = 3; SAMP-SEQ = RAPID). Table 1 summarizes the observations from both programs. Visit 14 represents a typical observation from GO-13021 and is representative of time-series, spatial scan observations in general.

Unlike GO-13021, we precede each science orbit with our programmed persistence mitigation strategy. This involved exposing the WFC3/IR detector to light emitted by the on-board Tungsten lamp during Earth occultation. Each of 55 exposures lasts 22.3 seconds (NSAMP = 4; SAMP-SEQ = SPARS10). The goal is to achieve a median fluence comparable to that of GO-13021 ($21,860 e^-/\text{pixel}$). Using the G141 grism, we reach a median fluence of $19,230 e^-/\text{pixel}$. This corresponds to a median flux of $862 e^-/\text{s}$. See Appendix A for additional information on how this observation was scheduled.

Data Reduction and Analysis

To reduce the GJ 1214 data, we used software that is part of the WFC3 Quicklook project (Bourque et al., 2017) and is intended to monitor time-series observations that use the spatial scan mode (Stevenson, 2019). Briefly, the Spatial Scan Monitor is used to track data quality over time and quantify how it varies with certain parameters (e.g., pointing drift, fluence, etc.). The data reduction software makes use of the IMA files rather than the FLT files because the former often yield more robust results. The reduced data consist of band-integrated light curves (flux vs. time) with auxiliary information. No reduction is required for the tungsten lamp data. We use the FLT frames and compute median flux values over the same subarray region used to analyze the GJ 1214 data. We fit the white light curves using standard techniques (e.g., Stevenson et al., 2014). For both datasets, we adopt a quadratic correction in time and a custom-written version of the ramp model described by Zhou et al. (2017).

Results

Pre-Flashed Data

We measure the flux during Earth occultation (when the Tungsten lamp is turned on) to assess its stability. Figure 1 shows a sharp increase in flux at the beginning of each orbit followed by a steady decline. Tests indicate that the observed trend is consistent across the entire detector, not just the science region exposed by GJ 1214. We consider whether the observed trend is the result of variations in the brightness of the tungsten lamp, charge trapping in the detector, or both. The tungsten lamp has a warm-up period of at least 40 seconds before subsequent commanding can start; however, despite this delay, data from the UVIS Bowtie Monitor program (CAL 11808, 11908, 12344, 12688, 13104, and 13072; Bourque & Bagget, 2013) shows that the first tungsten lamp exposure is consistently lower than the third exposure (> 200 seconds later) by a few percent. We note that the two channels (UVIS and IR) use physically different, but identically-produced lamps. This information suggests that the tungsten lamp (at least partially) contributes to the ramp shown in Figure 1. A future calibration program is necessary to fully quantify the lamp’s contribution to this ramp.

Persistence Mitigation

Figure 2 depicts the normalized flux as a function of time (the white light curves) from both programs. Data acquired using the standard method (GO-13021) exhibit significant ramps, with flux increases of 1.26%, 0.22%, and 0.14% over orbits 1, 2, and 4, respectively. The total increase from the first to fourth orbit is 1.30%. By comparison, data acquired with

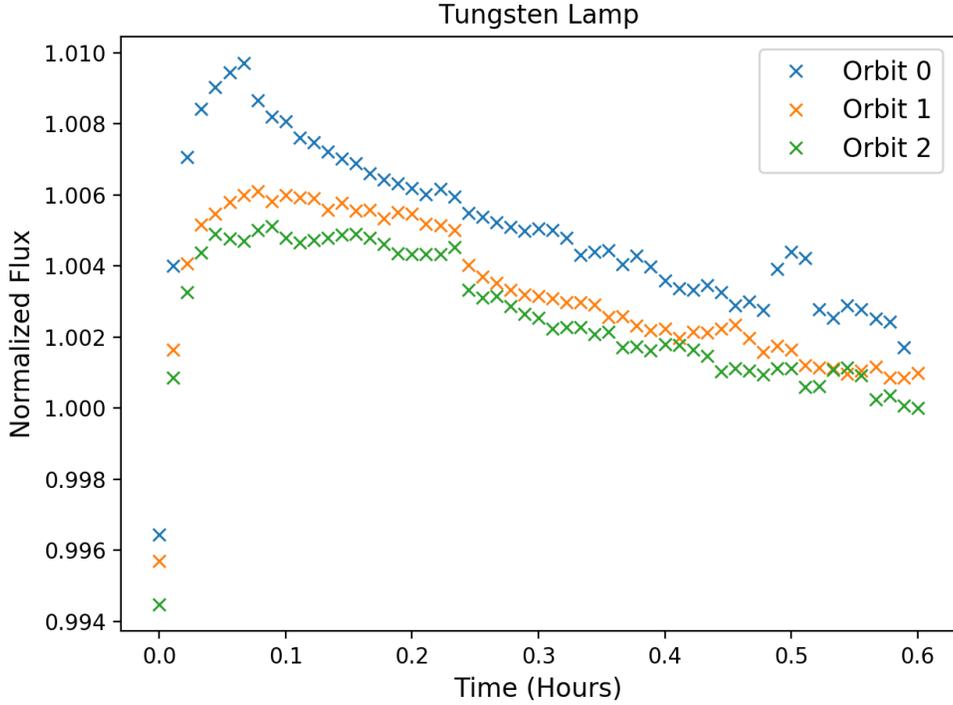


Fig. 1.— Measured flux from the tungsten lamp over three orbits. The flux is normalized to the last measurement in Orbit 2. The shape of the measured flux variation is consistent between orbits, but its cause is unclear.

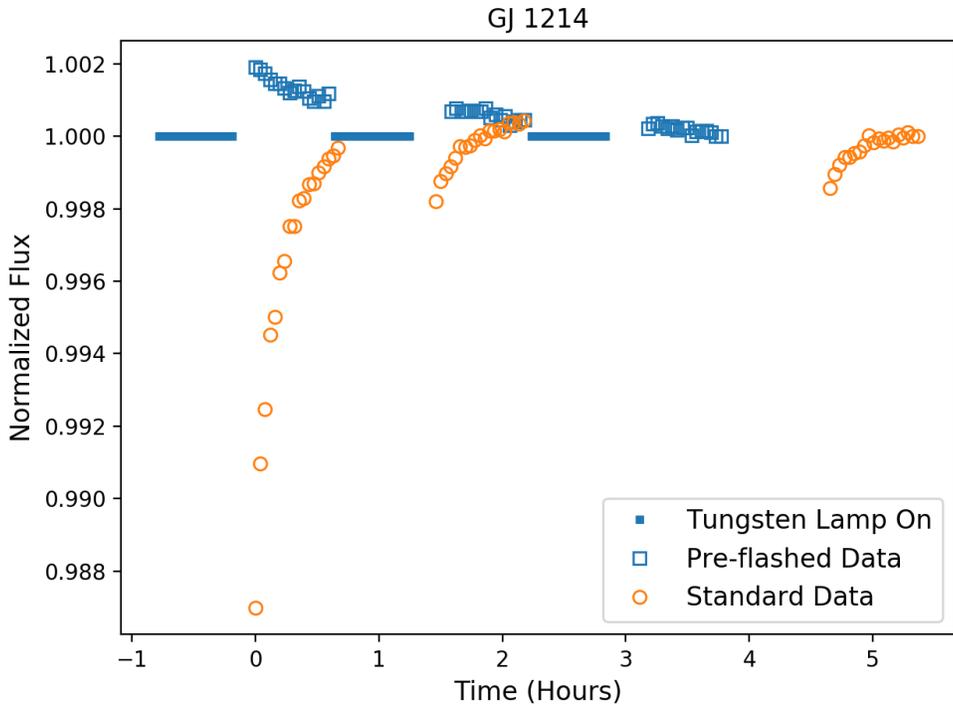


Fig. 2.— Normalized white light curves with and without pre-flashing the detector. The pre-flashed data (blue squares) yield a less-pronounced ramp (HST orbit-dependent systematic) relative to the standard dataset (orange circles). The improvement is the result of illuminating and cycling the detector during Earth occultation (blue lines). [Note: we have removed the third orbit from the standard dataset because it contains the transit of GJ 1214b.]

Label	Orbit 1 (ppm)	Orbit 2 (ppm)	Orbit 3 (ppm)	Orbit 4 (ppm)	Total (ppm)
Pre-Flashed	73	62	56	N/A	64
Standard	61	73	N/A	78	72

Table 2: Measured rms per HST orbit and combined over all orbits.

the pre-flashing technique exhibit smaller decrements in flux (-0.09%, -0.02%, and -0.02% over orbits 1, 2, and 3, respectively). The total change over all three orbits is -0.19%. The factor of 13 reduction in the persistence systematic within the first orbit suggests that, with pre-flashing, it may be possible to retain this orbit for science purposes. Additionally, after the first orbit the pre-flashed data exhibit minimal curvature in the persistence systematic.

Light Curve Residuals

After fitting each white light curve with a quadratic function in time and a ramp model similar to that of [Zhou et al. \(2017\)](#), we compute the residual rms values. Figure 3 plots a histogram of the residuals. We find that the pre-flashed data exhibit less scatter than the standard data (64 vs 72 ppm rms); however, the standard error on the residuals is ± 9 ppm, so the difference is not statistically significant. For comparison, [Kreidberg et al. \(2014\)](#) report an rms of 70 ppm for their example observation. The expected photon-limited rms is ~ 65 ppm. Table 2 breaks down the residual rms values by orbit number. The values given for Orbit 1 do not trend in a consistent direction (larger or smaller) compared to the other orbits. This suggests that random noise dominates the orbit-to-orbit variations. As confirmation, the standard error on the residuals in each orbit is ~ 15 ppm.

We note that the decrease in rms for the pre-flashed data cannot be attributed to improved centroid stability. Figure 4 demonstrates the increased scatter experienced by the pre-flashed data (relative to the standard data). Observations with 1σ scatter of $\gtrsim 0.1$ pixels along the dispersion direction (x axis) experience residual noise due to imprecise flat fielding ([Stevenson, 2019](#)). The pre-flashed data exhibit an rms of 0.09 pixels, whereas the standard data exhibit an rms of only 0.04 pixels. [Stevenson \(2019\)](#) describes how the centroid is computed.

When considering whether or not to implement pre-flashing, an important factor is the additional burden on WFC3's channel select mechanism (CSM). When fully inserted, the CSM directs the light beam into the IR channel. When the CSM is completely removed, light continues to the UVIS channel. An intermediate position is necessary to insert the diffuser that directs light from the tungsten lamp into the IR channel. Thus, the current strategy would typically require two CSM moves per *HST* orbit, which translates into an additional ~ 220 moves per year. This non-negligible increase in wear on the CSM, combined

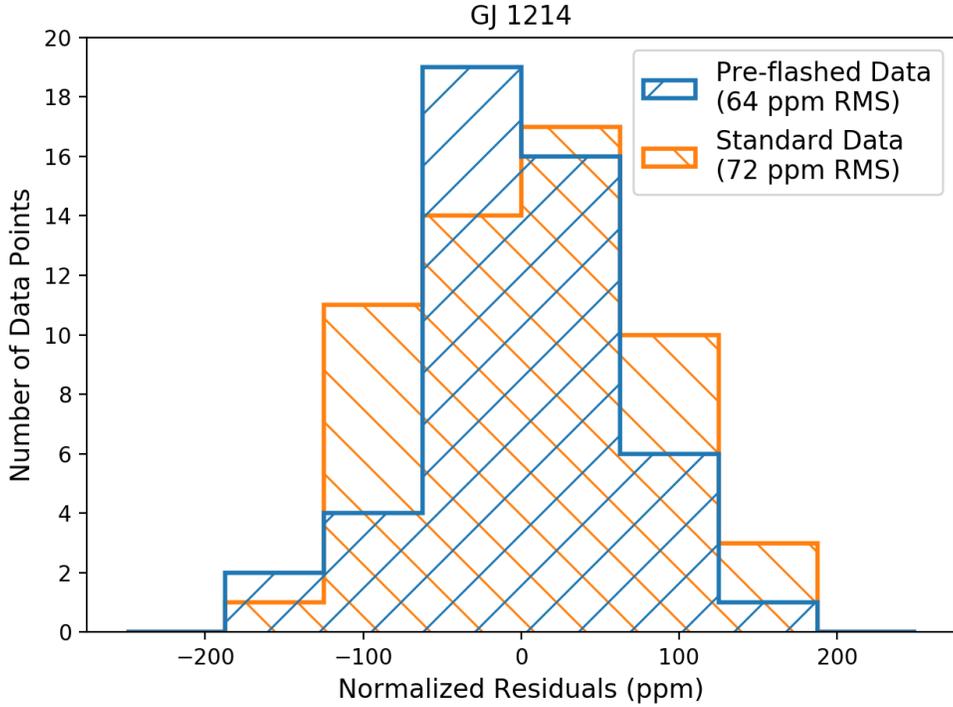


Fig. 3.— Histogram of normalized residuals. The pre-flashed data achieve slightly smaller residual scatter (i.e., narrower histogram) relative to the standard data (64 ppm vs. 72 ppm). The expected photon-limited rms is 65 ppm.

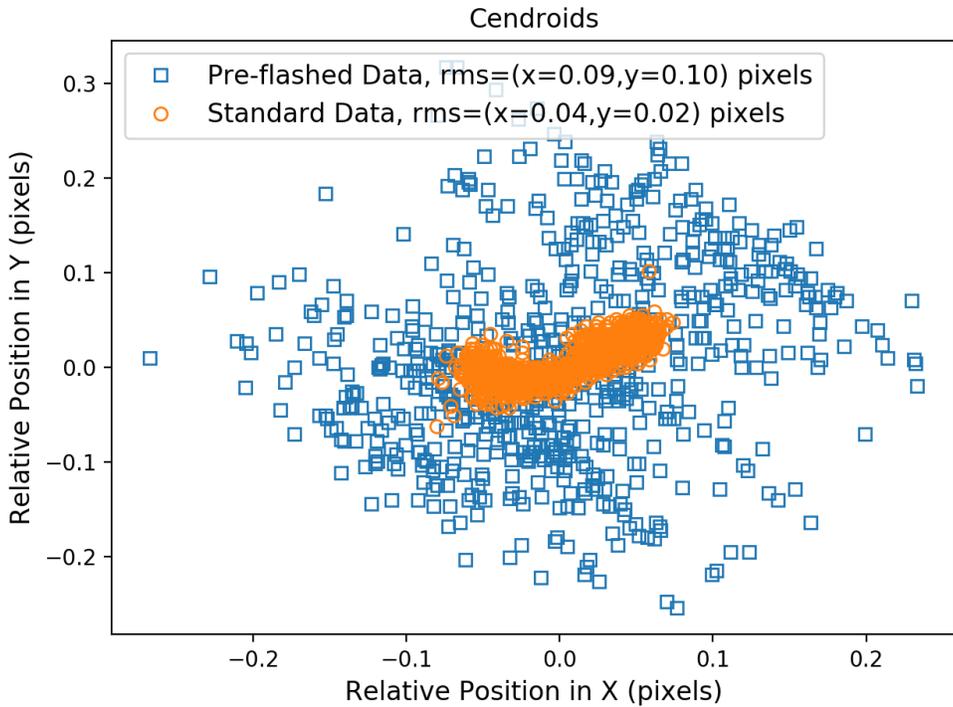


Fig. 4.— Measured centroid positions of the pre-flashed and standard data. The increased scatter in the pre-flashed data confirms that the smaller residual scatter shown in Figure 3 is not due to improved centroid stability.

with the effectiveness of current modeling techniques, does not warrant the implementation of pre-flashing on a regular basis.

Summary

We have demonstrated that by pre-flashing the detector during Earth occultation, the magnitude of the ramp is reduced by a factor of ~ 7 (from 1.30% to -0.19%) and, after correcting for instrument systematics using the ramp model described by Zhou et al. (2017), the improvement in precision is not statistically significant ($\Delta_{\text{rms}} = -8 \pm 9$ ppm). Our strategy of continuously illuminating and reading out the detector (as opposed to holding the charge) validates our hypothesis that the detector more quickly reaches an equilibrium state when subjected to higher levels of flux. The increase in wear on the CSM and lack of improved precision confirms that the use of pre-flashing is unjustified with WFC3.

Appendix A

Successfully implementing this program required several iterations with the scheduling team that we record here for posterity. When initially submitted, the program contained a number of timing and orient constraints that were incompatible with the scheduling software. Specifically, Visit 01 (the first pre-flash orbit) had separate orient ranges, despite being internal, as well as a SAME ORIENT restriction linking it to Visits 02 – 04. It also had both GROUP WITHIN and PHASE requirements. The combination of all these restrictions crashes the scheduling software. After testing with the Spike team (Spike is a general framework for planning and scheduling *HST* observations), it became clear that visits that do not have an external target cannot have Orient constraints placed upon them. In order to fix this, all Orient ranges need to be placed on one of the “pointed” visits (Visit 02). The unpointed visit (Visit 01) can then use the SAME ORIENT AS Visit 02 restriction. For the same reason, the phase constraint must also be placed on the pointed visit.

It was a requirement of this program that the first lamp exposure take place during occultation immediately preceding the first pointed visit. To implement this requirement, we added a timing link to Visit 02 to be After Visit 01 by 30-45 minutes. Since Visit 02 has to execute at the start of its orbit, this requirement forces Visit 01 to take place half an hour beforehand, which is during Earth occultation.

We briefly examined if the pre-flashing strategy could be implemented as a calibration program that runs parallel to time-series, spatial scan observations and determined that it would likely be feasible. Having the pre-flashing program run in parallel would minimize any additional work performed by observers when planning their Phase II programs. A similar strategy was used to schedule darks after South Atlantic Anomaly crossings with NICMOS

(Bergeron & Dickinson, 2003).

Acknowledgements

We would like to thank Jules Fowler, Peter McCullough, Sylvia Baggett, and John Mackenty for their thorough reviews of this ISR. We would also like to thank Mario Gennaro and Knox Long for their useful insight and discussions on this program.

References

- Bergeron, L. E., & Dickinson, M. E. 2003, Removal of Cosmic Ray Persistence From Science Data using the Post- SAA Darks, Tech. rep.
- Bourque, M., & Bagget, S. 2013, WFC3/UVIS Bowtie Monitor, Tech. rep.
- Bourque, M., et al. 2017, in IAU Symposium, Vol. 325, Astroinformatics, ed. M. Brescia, S. G. Djorgovski, E. D. Feigelson, G. Longo, & S. Cavuoti, 397–400
- Kreidberg, L., et al. 2014, *Nature*, 505, 69
- Long, K. S., et al. 2014, Attempts to Mitigate Trapping Effects in Scanned Grism Observations of Exoplanet Transits with WFC3/IR, Tech. rep.
- Stevenson, K. 2019, Analyzing Eight Years of Transiting Exoplanet Observations Using WFC3's Spatial Scan Monitor, Tech. rep.
- Stevenson, K. B., et al. 2014, *AJ*, 147, 161
- Zhou, Y., et al. 2017, *AJ*, 153, 243