

Instrument Science Report WFC3 2015-01

IR “Snowballs”: Long-Term Characterization

M. J. Durbin, M. Bourque, S. Baggett

March 10, 2015

ABSTRACT

We present a comprehensive analysis of the properties of the WFC3/IR anomalies known as “snowballs”, using all in-flight WFC3/IR data through July 2014. The source of snowballs is unclear at present, but is hypothesized to be radionuclides in the detector or bonding material. Nearly 7400 unique snowball events have been identified from the search of ~ 5 years of WFC3/IR images (over 6200 hours of exposure time), translating to a mean rate of about 1.2 snowballs per hour of IR exposure time. Typical snowballs affect about 10 pixels and saturate 2-5 of those, and deposit about 200,000 to 500,000 electrons on the detector. We find no trend over 5 years in the rates of snowball occurrences, consistent with the hypothesis that the uranium-238 decay chain is the source of snowballs.

We also publish a table of all hitherto identified snowballs, available at <http://www.stsci.edu/hst/wfc3/ins-performance/anomalies/ir-snowball-table.txt>.

keywords: HgCdTe, Hubble Space Telescope, HST - instrumentation: detectors - Wide Field Camera 3, WFC3/IR

Introduction

Snowballs are transient events observed in some HgCdTe detectors that occur instantaneously and deposit at least 200,000 electrons over an area of up to ~ 40 pixels. They have

been hypothesized by McCullough (2009, hereafter M09) to be the result of alpha decays from either thorium-232 or uranium-238 at concentrations of ~ 1 ppm within the detector or its surrounding material. We may be able to put constraints on this hypothesis with the data presented here, although on-orbit observations alone are not sufficient to fully determine the physical source of snowballs.

Snowballs do not currently and are not expected to pose a significant problem for observers, as the calwf3 pipeline does not differentiate between snowballs and cosmic rays (Rajan et al. 2011), and successfully rejects the vast majority of both. Nevertheless, an understanding of snowball properties is a useful aspect of detector characterization, particularly with regards to the question of whether their rate of occurrence changes over time. These results will have important implications for future ground- and space-based missions using HgCdTe detectors; detector manufacturers may be able to use some of the results presented here in future to reduce the number of snowballs an instrument produces. The NIRISS/FGS (Near-Infrared Imager and Slitless Spectrograph/Fine Guidance Sensor) team, for example, during their evaluation of ground test data of detector candidates have ruled out the use of two of them for flight due to unacceptably high rates of snowballs (Volk 2014).

Analysis

Identification Procedure & Criteria

We use the same methodology applied by Hilbert (2009, hereafter H09) to search for snowballs. For each IR multiaccum ramp, a difference image is generated i.e., read $N + 1$ – read N , for all N in the ramp. Each difference image is examined for any pixel more than 50σ above the background; if it is high in all difference images, we discard it as a likely hot pixel. To avoid the weaker cosmic rays, the 4 nearest neighbor pixels of a potential snowball central pixel are required to have flux level more than 25σ above the background. This set of candidates included many false positives due to relatively energetic cosmic rays, trailed sources in failed guide star observations, and other anomalous events, so a manual inspection of each potential snowball was deemed necessary to evaluate whether it should remain in the snowball catalog or be discarded.

Visual snowball assessment was performed with the aid of a script that displays three views of each detection: the read in which the snowball appears in the `_ima` file¹, a saturation map, and a flux map. After inspection, the user has the option to either keep the detection in the snowball candidate list, reject it, or view the previous or following read. When manually vetting snowballs, an effort was made to be conservative with rejections; if there were any doubt about whether a given candidate qualified as a snowball, we erred on the side of keeping it. One can see that several of the snowballs presented in Figure 1 do not meet all of the criteria listed below, but as we wished to limit our preconceptions of what snowball

¹The `_ima` extension indicates “a file that has had all calibrations applied (dark subtraction, linearity correction, flat fielding, etc.) to all of the individual readouts of the IR exposure” (Rajan et al 2011).

properties are universal, ambiguous candidates were kept more often than rejected.

For manual vetting, candidates were kept if they met at least one of the following criteria, which were established based on the snowball properties found in H09:

1. At least $\sim 200,000$ electrons deposited in ~ 4 -30 pixel area
2. At least one saturated pixel
3. Fairly well-defined edges in flux map

We find a total of 7392 unique snowball events. We have performed the vetting routine twice over on 5 months' worth of data, and found that out of 764 total snowballs identified over the first and second trials, 648 or 85% appeared in both vetting runs; thus, we can say that 85% of snowballs are unambiguous. Of the snowball candidates that were not selected consistently, 71% have flux values under $300,000 e^-$, whereas only 32% of the final catalog of all snowball detections do (see Figures 5 and 6).

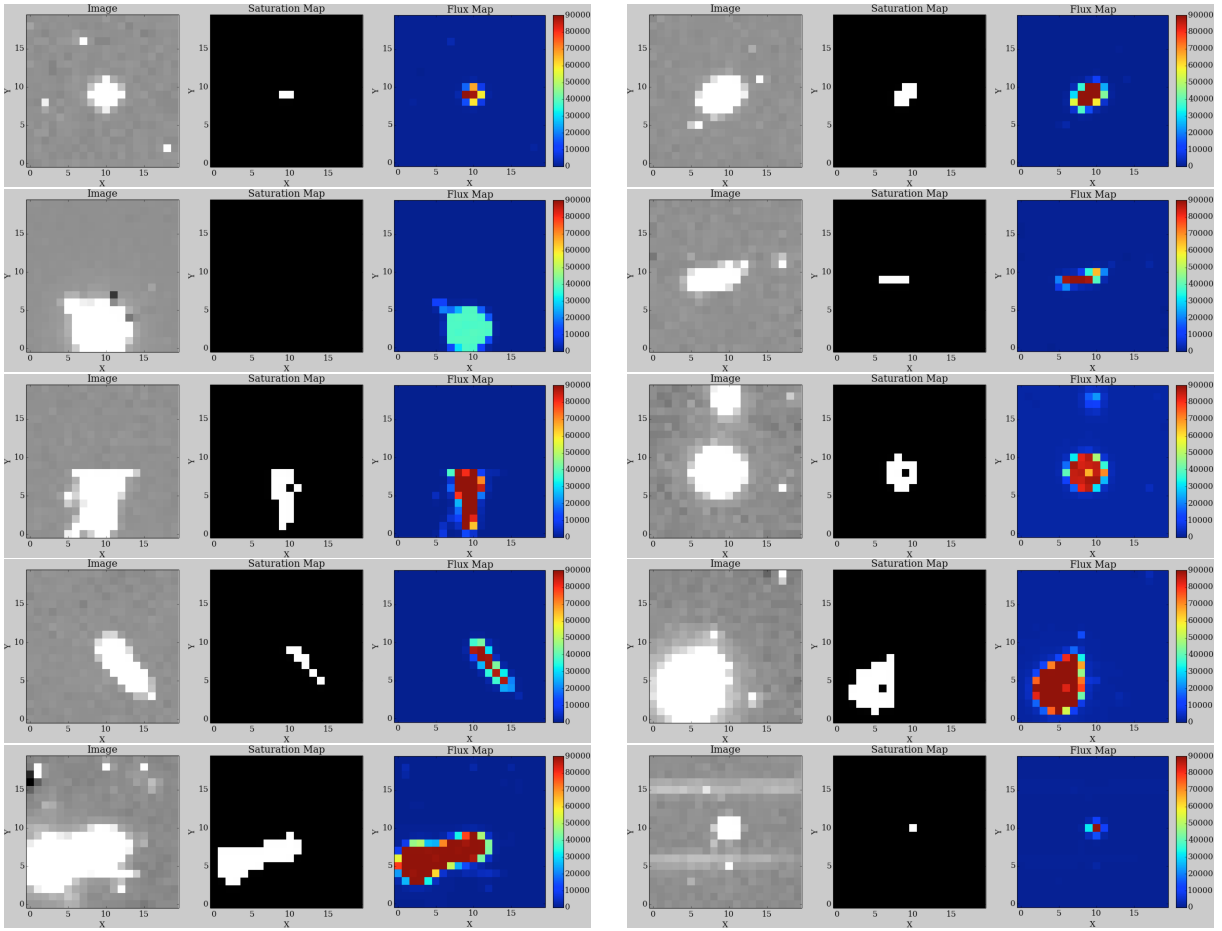


Figure 1: A sample of observed snowballs, as shown in the snowball vetting routine. Here the left image in each set is the region of the `_ima` file containing the snowball, center is the saturated pixel map, and right is the flux map, where red represents saturation.

Snowballs Over Time

M09 predicts an increase in snowball occurrences by a factor of 4 over 10 years if the source of alpha decays is solely Th-232, and a factor of 2 increase if the source is equal parts Th-232 and U-238. A histogram of snowball rate over time is thus a useful diagnostic for constraining the source of snowballs. If the source is at least 50% Th-232, we should be able to see an increase in snowball rate over 5 years by a minimum factor of approximately 1.5; however, we find no evidence of any increase in snowball rate over time. Assuming snowballs are due to radioactive decay, this is consistent with the hypothesis that U-238 (not Th-232) is the dominant radionuclide in the detector or bonding material (M09).

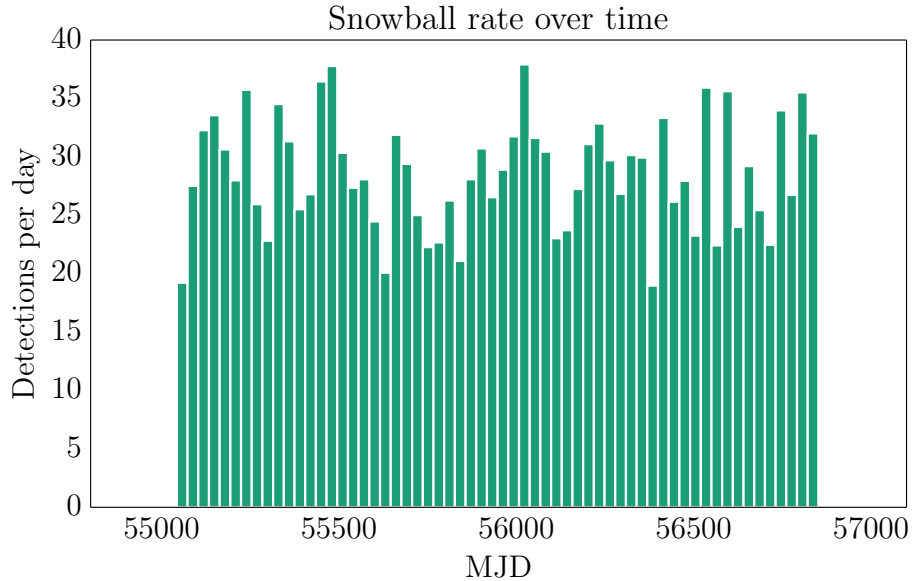


Figure 2: Snowball rate over time (snowball count normalized to total IR exposure time), binned approximately by month. We note here that the variance is of order 30%, which is higher than expected based on Poisson noise; this will be investigated in a future ISR.

Rauscher et al. (2014) observe a decreasing rate of snowballs in JWST’s NIRSpec instrument over time, and assert that the decay rate is consistent with known radioisotopes used in the detector fabrication process. They predict that NIRSpec’s snowball rate will be negligible by the time of JWST’s launch in 2018; however, our results appear to contradict their findings.

Furthermore, we find no trends over time in snowball fluxes, saturated pixels, total area, or positions on the detector, as illustrated in Figures 3 and 4. Snowball fluxes were calculated in the original snowball-finding script using simple aperture photometry with an aperture radius of 8 pixels and background annulus of radius 10 pixels starting at 10 pixels from the snowball center. Total snowball area, also calculated in the original script, is the number of pixels with signal over 9σ higher than the mean background signal.

Snowball positions on detector

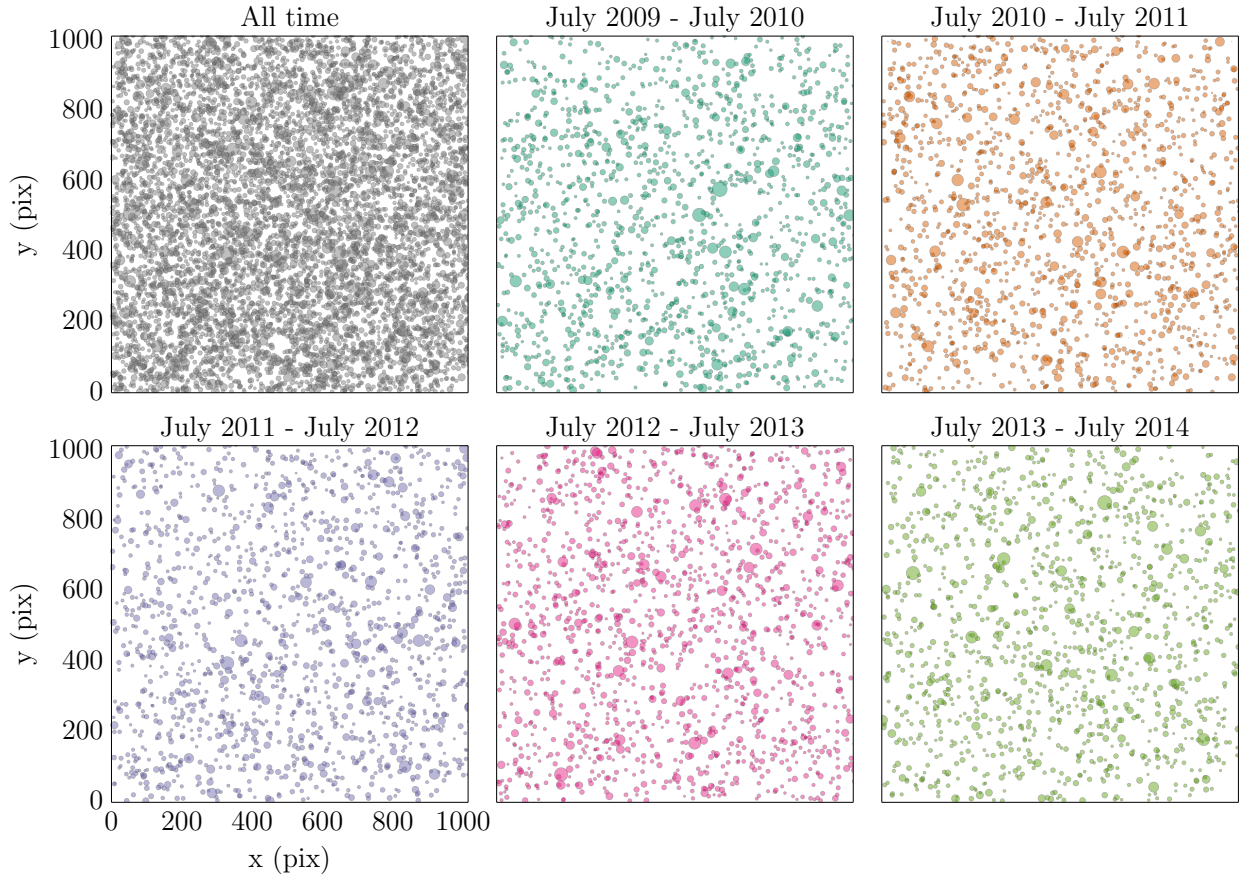


Figure 3: Snowball positions by epoch, with marker size scaled by total snowball flux. We observe no noticeable patterns in snowball locations over time, and we will analyze the spatial distribution of snowballs in detail later in this report.

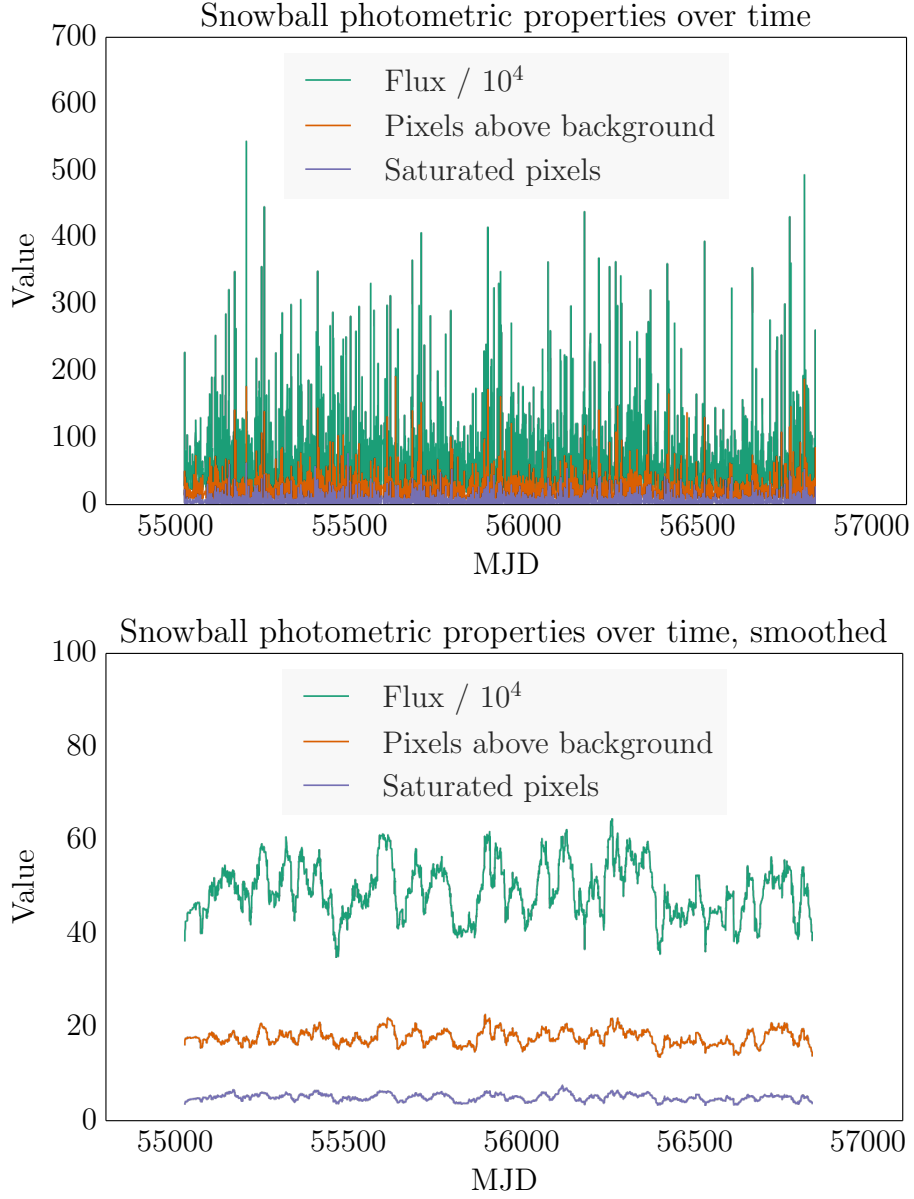


Figure 4: Snowball fluxes, total affected pixels, and saturated pixels by date. The datasets in the bottom plot have been convolved with a boxcar kernel of size 100 days to improve readability of the plot and demonstrate that there are no trends hidden by scatter.

Snowball Properties

We find that snowball fluxes follow an approximately log-normal distribution with a long tail (Figure 5), with the caveat that all flux measurements are lower bounds due to saturation. The best-fit flux distribution has a mode of $3.26 \times 10^5 \text{ e}^-$, a median of $4.36 \times 10^5 \text{ e}^-$, and a mean of $5.04 \times 10^5 \text{ e}^-$. The number of saturated pixels follows a similar roughly log-normal distribution with a peak at 2, while the total number of affected pixels appears closer to exponential decay, with almost no snowballs affecting fewer than 8 pixels (Figure 7).

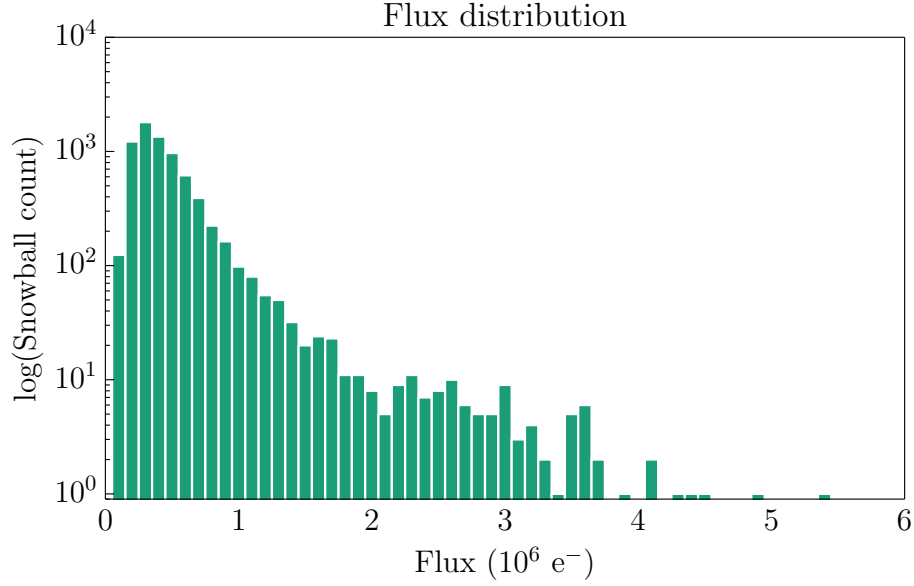


Figure 5: Histogram of measured snowball fluxes. There are several points at the high end of the distribution with flux measurements of up to $5.4 \times 10^6 \text{ e}^-$; however, the vast majority of snowballs deposit fewer than $2 \times 10^6 \text{ e}^-$.

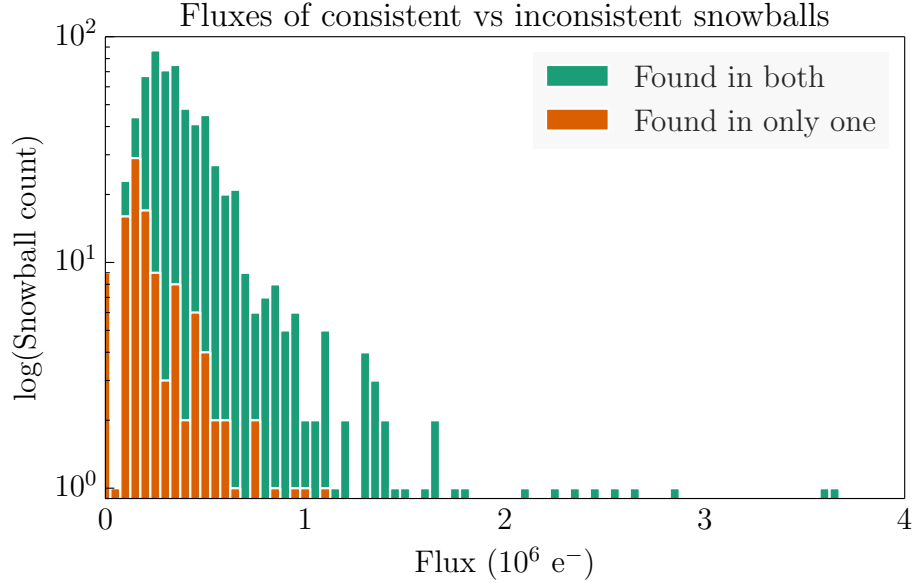


Figure 6: A comparison of the flux distributions of snowballs that were identified consistently between two runs of the snowball vetting routine on the same data set, and those that were identified only once. We can see from this that manual identification is least reliable for snowball candidates that deposit relatively few ($< 3 \times 10^6$) electrons.

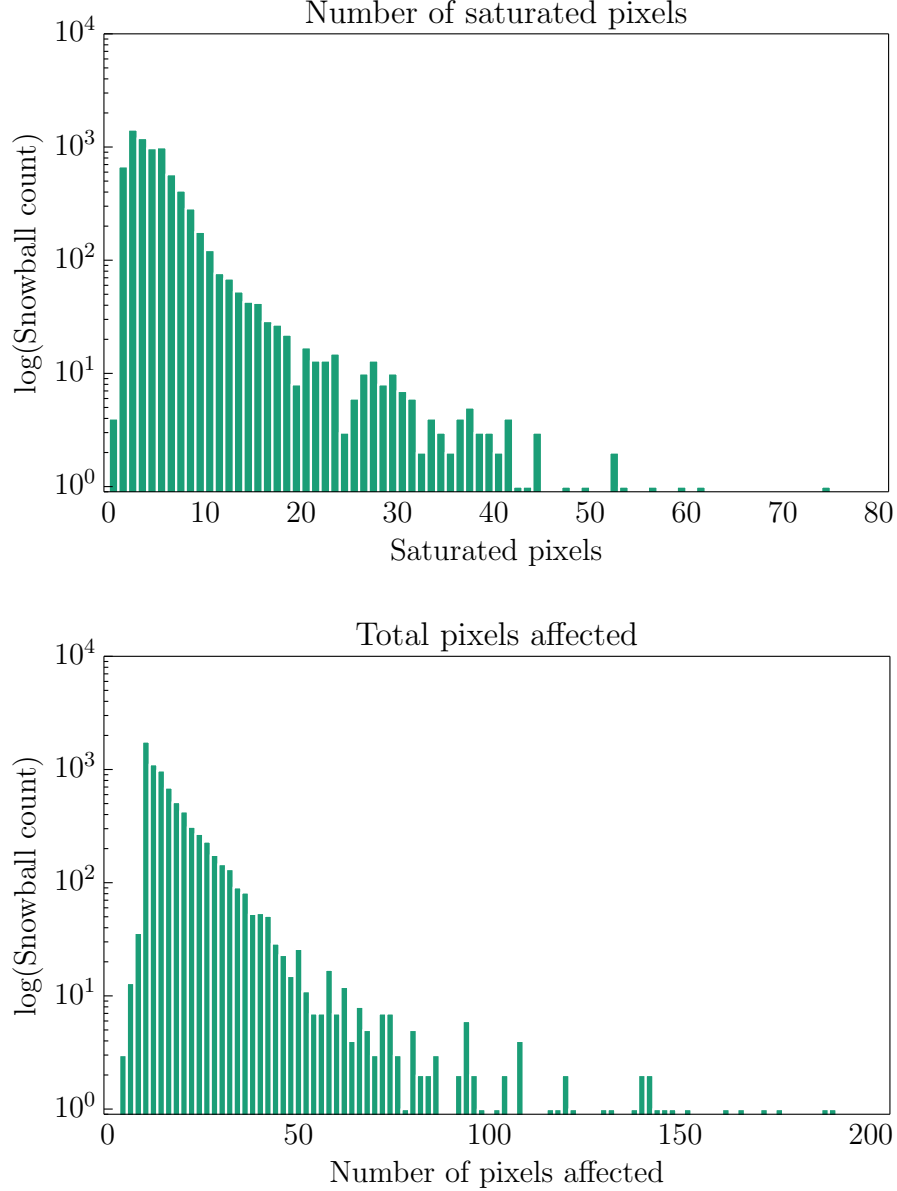


Figure 7: Histograms of number of saturated pixels (top) and total snowball area (bottom). As in the flux histogram, there are several points at the high end of each distribution with up to 75 saturated pixels and 191 total affected pixels; however, the bulk of snowballs saturate no more than 30 pixels and affect no more than 50.

Spatial Correlation & Repeat Decays

Although we find that the snowball rate is approximately stable over time, we wish to examine whether there is any evidence for daughter decay production as hypothesized by M09. To do this, we first examine the evidence for spatial autocorrelation among snowballs; that is, whether snowballs tend to appear closer to each other than expected if snowball positions are random. We calculate the number of snowball pairs that have distances between

r and $r + 3$ for $0 \leq r \leq 20$ pixels, scale that number by the area of the search annulus $\pi[(r + 3)^2 - r^2]$ and compare this to a set of 500 numerical simulations with randomly generated snowball positions, as seen in Figure 8.

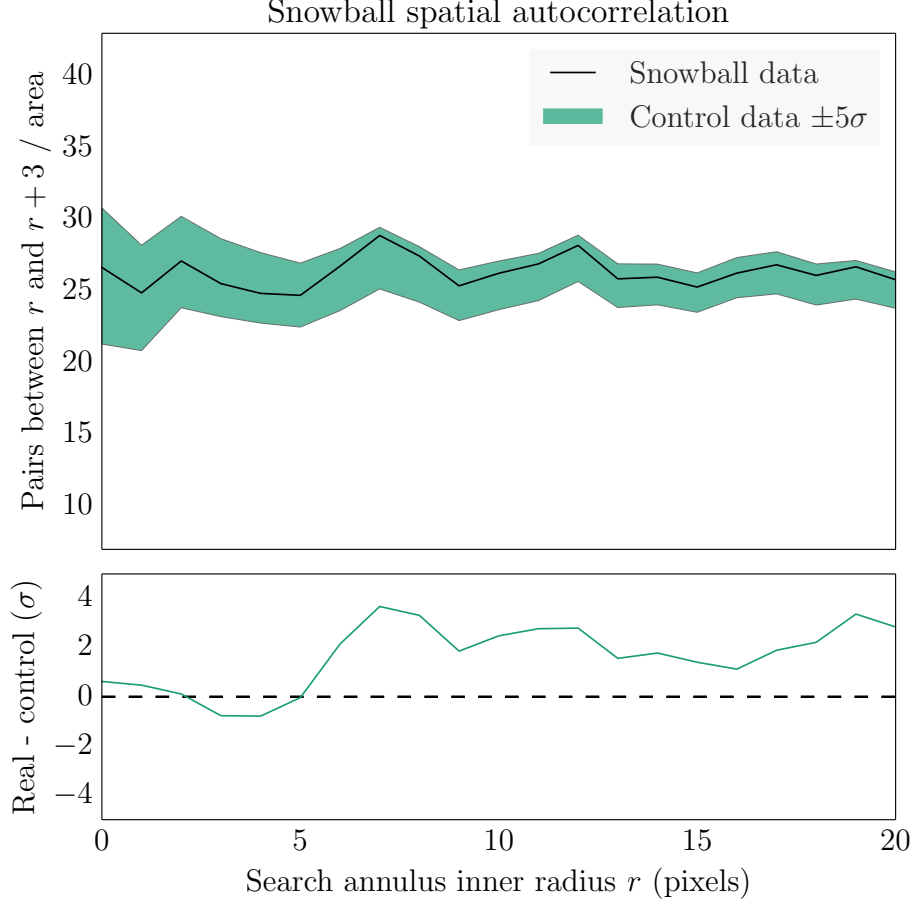


Figure 8: Number of possible snowball pairs with distances between r and $r + 3$ as a function of r , scaled by annulus area $\pi[(r + 3)^2 - r^2]$.

We find that for $r \leq 5$ the number of snowball pairs is approximately in line with predictions, whereas at larger distances the number of snowball pairs grows up to 4σ larger than the predicted number. This indicates that there may be regions larger than 5 pixels in radius that have relative overdensities of snowballs, but snowballs are not more likely to appear within 5 pixels of each other than random, as would be expected if there are a significant number of daughter decays.

Another way to check for spatial correlation is finding the distance to the nearest neighbor of each snowball and creating a histogram (Figure 9). If snowball positions are random, we expect an approximately Poissonian distribution of nearest neighbor distances.

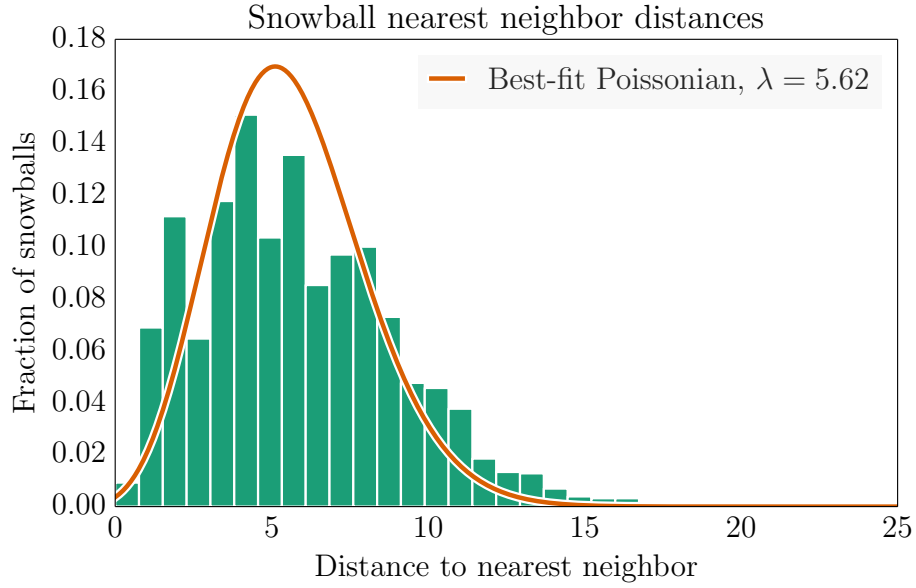


Figure 9: Snowball nearest neighbor distance histogram and its best-fit Poisson distribution.

We also perform the same nearest neighbor search and Poisson fitting on 500 sets of randomly generated coordinates, and find a mean λ of 5.71 ± 0.10 . While the true snowball nearest-neighbor distance distribution does not fit the Poisson curve perfectly, its λ is within 1σ of the predicted λ for randomly distributed snowballs, and although there are more snowballs with low nearest-neighbor distances than expected, there are also more with high nearest-neighbor distances.

Although there is little spatial evidence for repeating snowballs, they may also be detected using temporal correlation. Th-232 daughter decays are expected to appear within less than a day of the original snowball (see Table 1 in M09 for details on the predicted decay chain). We can test whether there are an overabundance of such pairs by finding the fraction of pairs of snowballs within r of each other that have a time difference Δt less than one day, as shown in Figure 10.

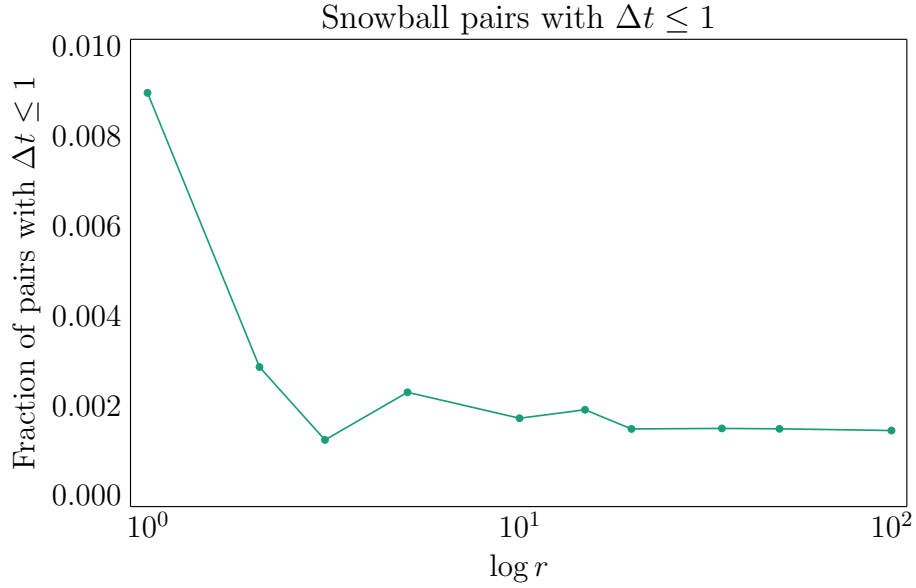


Figure 10: Fraction of all snowball pairs with distance less than r that have a delay time Δt less than or equal to one day as a function of r . While it appears as though there are substantially more snowball pairs with $\Delta t \leq 1$ day at $r = 1$, the fraction represents only 1 pair out of 111.

The fraction of snowball pairs that appear within 1 day of each other is below 1% for all search radii, and below 0.25% for $r \geq 3$, the latter of which is in particular small enough to be considered negligible. However, it is possible that the fraction of snowballs with $\Delta t \leq 1$ day may be an underestimate due to gaps between IR exposures. If so, there is little to be done about that, and given the stable snowball rate over time it is improbable that we are missing a significant amount of potential daughter decays.

While we cannot rule out a weak spatial correlation between snowballs at large radii, we find no evidence that said spatial correlation is caused by the daughter decays that M09 predicts. It is possible that the distribution of radionuclides over the detector is non-isotropic, or there may be another explanation altogether. We will continue to monitor the positions of observed snowballs and search for trends, but it is not presently a major cause for concern.

Unstable Pixels

H09 found evidence for a moderate spatial correlation between snowballs and unstable pixels (see Hilbert 2012 for details on processes used to identify unstable pixels), which we are now equipped to investigate more thoroughly. We select pixels flagged as unstable (DQ value = 32) or unstable and affected by blobs (DQ value = 512 + 32 = 544) from the bad pixel table w681807ii.bpx.fits (described in Hilbert 2012). As with snowball autocorrelation, we compare the number of pairs between unstable pixels and known snowballs with distance between r and $r + 3$ to the mean number of pairs between unstable pixels and 500 sets of randomly generated snowball positions in the same annulus for $0 \leq r \leq 20$.

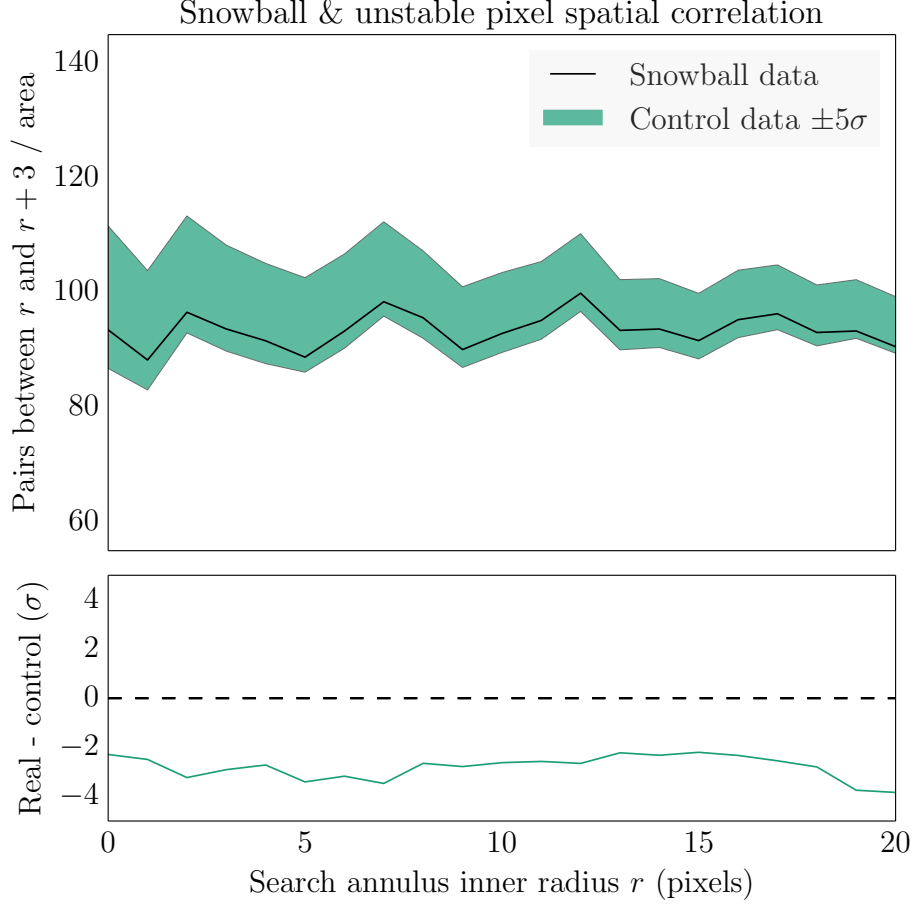


Figure 11: Number of possible pairs of snowballs and unstable pixels with distances between r and $r + 3$ as a function of r , scaled by annulus area $\pi[(r + 3)^2 - r^2]$.

Contrary to the findings of H09, we appear to find that snowballs and unstable pixels are somewhat anticorrelated; at most radii, the number of pairs between snowballs and unstable pixels is $\sim 3\sigma$ below the expected number.

However, there is a known overdensity of unstable pixels in the bottom right corner of the detector in the anomalous region known as the “wagon wheel.” While the mere fact that there is not a comparable overdensity of snowballs in that region points to snowballs and unstable pixels being unrelated, we run the pair routine again excluding the region $(x \geq 960, y \leq 80)$ for both snowballs and unstable pixels, as shown in Figure 12.

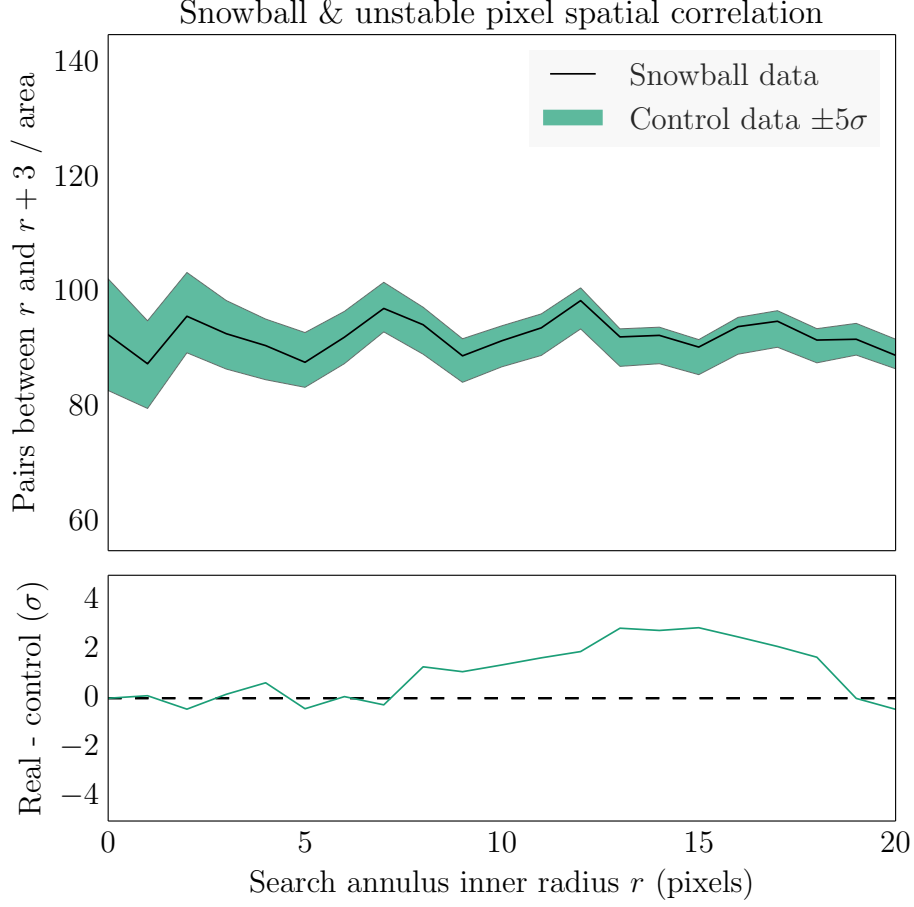


Figure 12: Number of possible pairs of snowballs and unstable pixels with distances between r and $r + 3$ as a function of r , scaled by annulus area $\pi[(r + 3)^2 - r^2]$ and excluding all snowballs and unstable pixels with coordinates $(x \geq 960, y \leq 80)$.

We find that excluding the densest part of the “wagon wheel” region effectively reverses the anticorrelation for radii greater than 7 pixels, although not sufficiently to confirm that snowballs and unstable pixels are linked. However, it is possible that the lack of observed snowballs in regions of high unstable pixel density can be explained by the fact that unstable pixels tend to be found near other types of bad pixels, such as dead pixels, which may prevent snowballs from being observed in the first place, and that these regions may effectively disguise a relationship between them.

As an additional test, we follow the method used in H09 to evaluate the relationship between snowballs and unstable pixels (see Figure 9 in H09). For $0 \leq r \leq 25$ pixels we calculate the fraction of snowballs that have at least n unstable pixels within r . As a control, we also calculate the fraction of all pixels that have at least n unstable pixels within r , which gives us an overall probability for how often we can expect to find at least n unstable pixels within r of any pixel.

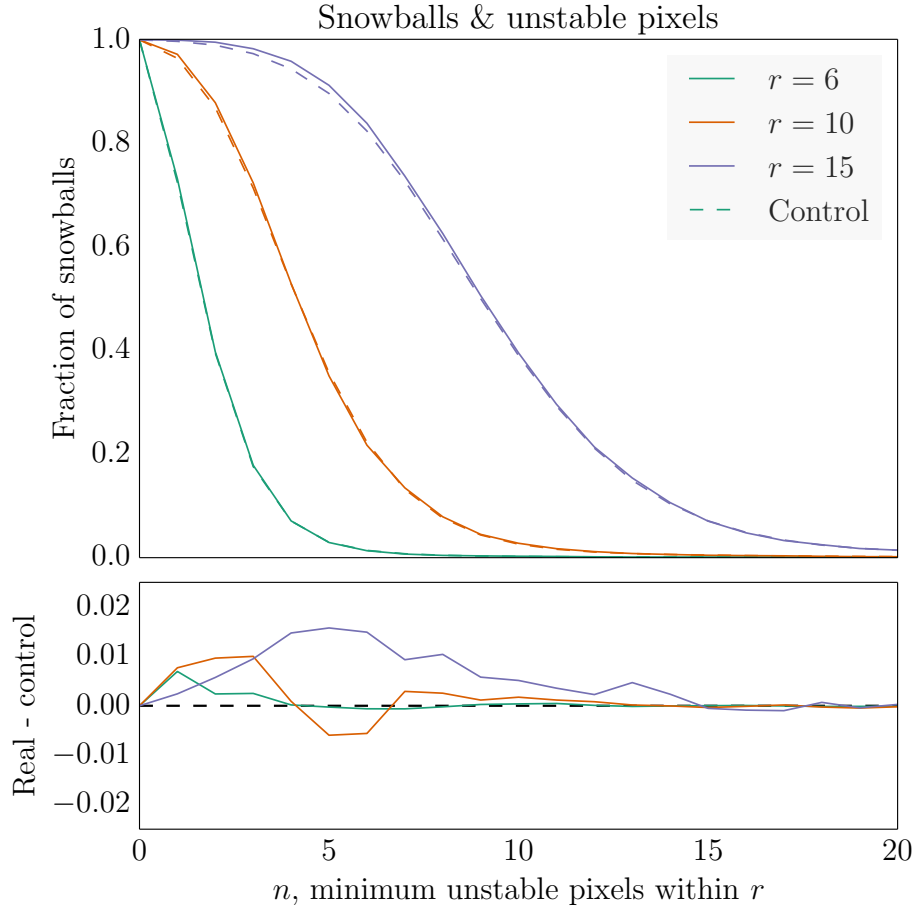


Figure 13: Fraction of snowballs with at least n unstable pixels within radius r (solid lines), and the corresponding control values (dashed lines).

In the most extreme case in Figure 13, it is 1.5% more likely to find 5 unstable pixels within 15 pixels of a snowball than any random pixel. Based on the available data, we conclude that there is no significant spatial correlation between snowballs and unstable pixels. This result runs counter to expectation if snowballs and unstable pixels are both caused by radioactive impurities in the detector, as previously hypothesized in H09.

Another caveat to all of this is that the most recent unstable pixel mask in the pipeline (and the one used in the analysis here) is based on data from 2009-2010. Preliminary results indicate that the unstable pixel population has not grown significantly over time, but an updated bad pixel table is required to confirm. A repeat of this analysis will be performed once a new bad pixel table is available.

Conclusions

We produce a catalog of 7392 snowball detections based on all in-flight WFC3-IR data (~ 6200 hours of exposure time). The “typical” snowball is approximately circular, has a flux of between 2×10^6 and 5×10^6 e^- , saturates between 2 and 5 pixels, and affects between 8 and 25 pixels. We find that the rate at which snowballs are observed is approximately constant over time, suggesting that if the source of snowballs is in fact radioactivity in the detector or bonding material, the primary contaminant is U-238. We find no evidence for snowballs being spatially autocorrelated at radii under 5 pixels, although they may be moderately autocorrelated on larger scales. We also do not find sufficient evidence to establish a spatial correlation between snowballs and unstable pixels, although this may change once the analysis is repeated with an updated bad pixel table.

We will continue to monitor the snowball population for as long as WFC3 continues to operate, and report any notable changes to any of their properties should they occur, although we do not anticipate any based on the current data.

Acknowledgements

We would like to thank Michael Fall for his helpful insights during the review process, and Peter McCullough and Bryan Hilbert for their contributions.

References

- Hilbert, B. 2009, “Snowballs” in the WFC3-IR Channel: Characterization, ISR WFC3 2009-43
- Hilbert, B. 2012, WFC3/IR Cycle 19 Bad Pixel Table Update, ISR WFC3 2012-10
- Hilbert, B. 2015, priv. comm.
- McCullough, P. R. 2009, Radioactivity in HgCdTe devices: potential source of “snowballs”, ISR WFC3 2009-44
- McCullough, P. R. 2015, priv. comm.
- Rajan, A. et al. 2011, “WFC3 Data Handbook”, Version 2.1, (Baltimore:STScI)
- Rauscher, B. J. et al. 2015, “New and Better Detectors for the JWST Near-Infrared Spectrograph,” PASP 126, 942
- Volk, K., 2014, NIRISS/FGS Detector Selection: “Snowball” Analysis, Doc #JWST-STScI-003929, SM-12