2002 HST Calibration Workshop
Space Telescope Science Institute, 2002
S. Arribas, A. Koekemoer, and B. Whitmore, eds.

Growth of Hot Pixels and Degradation of CTE for ACS

A. Riess

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218

Abstract. The anneal rate of hot pixels on the ACS WFC is ∼60%–65%, significantly lower than the characteristic anneal rate of 80–85% seen for other CCDs flown on *HST* (i.e., WFPC2, STIS, and ACS HRC). The ACS WFC is annealed in the same way as the other *HST* CCDs and there is no firm understanding at this time of the source of the difference. After ∼7–8 successive anneals, the cumulative fraction of annealed pixels reaches an apparent plateau at ∼70%. The fitted, successive annealing function is used to project forward in time the expected fractional coverage of the CCD by hot pixels. Approximately 2 years after launch the coverage by hot pixels is expected to exceed that by cosmic rays in a ∼1000 sec exposure. At the nominal end of the *HST* mission (2010) the coverage by hot pixels would be ∼6%, i.e., one out of every 16 pixels. Because hot pixels are readily flagged and corrected or discarded they do not pose a serious threat to science observations, but their growing presence require careful dithering and consideration. For CTE, internal tests such as the cosmic ray tail measurements show the degradation of CTE on ACS which is most pronounced for WFC. Simple scaling from WFPC2 provides some quantitative estimates for photometry.

1. Introduction

The anneal rate of new hot pixels (dark current > 0.04 e/s) on the ACS WFC has been a disappointingly-low ∼60% in the first 8 monthly anneal cycles of the instrument (Riess, Mutchler, Van Orsow 2002, *Instrument Science Report* 02-06). This rate is significantly lower than the observed and characteristic value for other *HST* CCDs; 80–85% for WFPC2, STIS, and ACS HRC. The anneal rate is significantly less than 60% for pixels which are much hotter than 0.04 e/s. The likely consequence of poor annealing is a greater fractional coverage of the camera by pixels with elevated dark current than the experience of other *HST* CCDs.

The reason for the low anneal rate for the ACS WFC is not clear at this time. The WFC is annealed at approximately the same temperature (∼+20 C), for no shorter a time interval (∼24 hours) and with the same frequency (monthly) as the other instruments. The answer may lie in the details of the manufacturing process of the CCD or in the way in which the CCD is read-out. Expert CCD consultants have been unable to provide a reliable explanation for the poor annealing (private communication, Blouke and Janesick). One possibility may involve a difference in the way the chip is operated during integration and read-out. During integration the chip is used in MPP mode, but when it is read-out, it is switched to non-MPP mode. This results in "turning-on" the Si-SiO interface states that are normally passivated when integrating in MPP mode. (Mark Clampin, private communication). Unfortunately the chip electronics dictate this switching which cannot be disabled, even for a test.

Here we seek to quantify, empirically model and project forward in time the likely future population of hot pixels on the ACS WFC. Using our predictive model we then test the utility of increasing the frequency of anneals to bi-monthly.

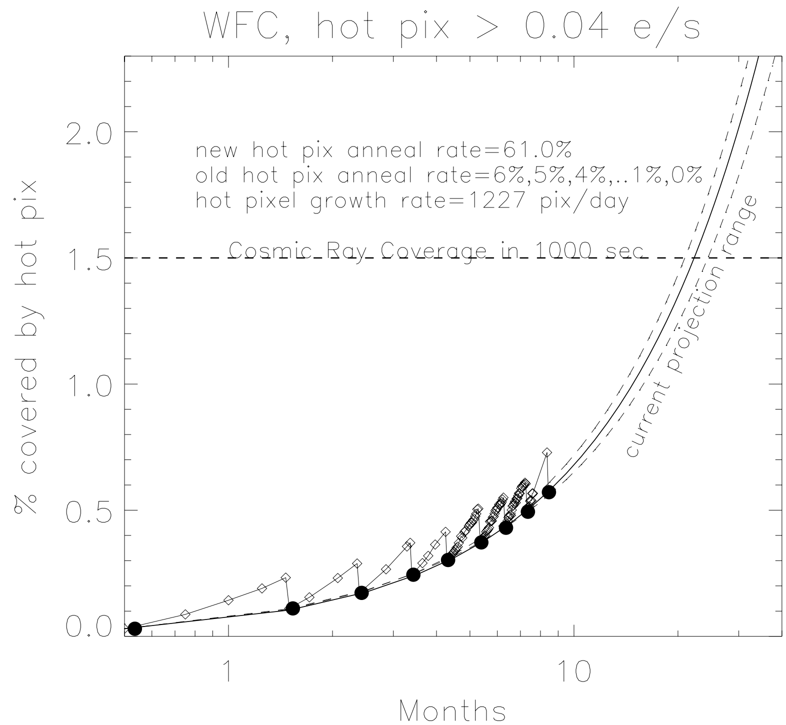

Figure 1. The growth of WFC hot pixels (dark current <0.04 e per s) with time (log scale). The "saw-tooth" pattern reflects the continual production and monthly annealing of hot pixels. Anneals heal ∼60% of new hot pixels.

2. Observations

The characteristic "saw-tooth" pattern of the growth and annealing of hot-pixels seen for ACS WFC during the first two anneal cycles (and all other CCDs flown on *HST*) has continued through the ensuing 6 anneal cycles. New hot pixels with dark current > 0.04 e/s (a number which is ∼18 times and ∼13 standard deviations above the mean dark current) steadily develop at a rate of ∼1200 pixels per day (see Figure 1). Once a month, the CCD is annealed and between 60% and 65% of the hot pixels created during the preceding month return to normal dark current production. Without any annealing, ∼10% of the CCD would be covered by hot pixels by early 2006. For the 40% of "persistent" hot pixels which do not anneal at their first opportunity, the likelihood that they anneal in any future anneal drops precipitously, a result which is consistent with WFPC2 and the ACS HRC.

Current projections for the hot pixel coverage of WFC are highly dependent on both the anneal rate of new hot pixels as well as the anneal rate of the persistent hot pixels. New hot pixels which fail to "heal" become increasingly unlikely to heal in each successive anneal. After approximately 7 or 8 anneals (the total number of anneal cycles we have observed to date) the cumulative anneal rate reaches an apparent plateau for pixels which became hot during the first or second anneal cycle (i.e., in April or May or 2002). The "persistent-pixel-annealing function," as seen in Figure 3, can be approximated by a simple series of monthly anneal rates of 62%, 6%, 5%, 4%, 3%, 2%, 1%, 0%. After 7 or 8 anneal cycles a cumulative fraction of ∼70% of hot pixels will have been annealed leaving ∼30% to remain hot (presumably for the long term).

Using this function we can project forward in time the expected hot pixel population. The result are Figures 1 and 2. Using as a reference the fractional WFC coverage by cosmic rays in a 1000 sec exposure (∼1.5%), we expect to reach this coverage by hot pixels less

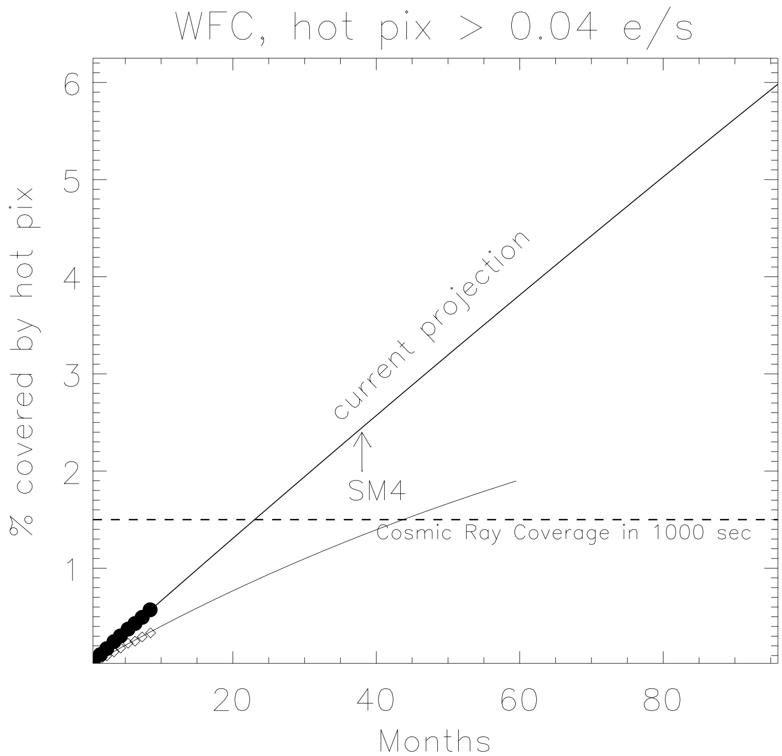

Figure 2. As Figure 1 on a linear scale.

than 2 years after launch or by the beginning of 2004. We would project a ∼3% coverage by ∼4 years after launch or by early 2006. We stress that accurate projections are difficult and grow more uncertain with the time interval of the projection.

For comparison, 2.5% of WFPC2 is currently covered by hot pixels (dark current > 0.02 e/s). However, a comparison between WFPC2 and ACS WFC coverage by hot pixels must account for the difference in the dark current limit used to define a hot pixel. For WFPC2 the limit is 0.02 e/s or half of the ACS WFC value. A simple correction derived from ACS WFC is that there are 1.4 times as many hot pixels at the WFPC2 threshold as for the ACS threshold. Using this conversion we conclude that currently ACS WFC has 1/3 the fractional coverage by hot pixels as currently exhibited by WFPC2.

2.1. Would Bi-Monthly Annealing Help?

Would bi-monthly annealing help reduce the long term growth of hot pixels? The evidence in hand suggests it would not. If the persistent-pixel-annealing function really does reach a plateau (as indicated by the data; see Figure 3) then simply increasing the number of anneals (or their frequency) would have negligible impact in the long term.

Another way in which bi-monthly annealing would differ from the current monthly annealing (and perhaps yield better results) is by reducing the mean time interval between the formation and attempted annealing of a hot pixel (from ∼14 days to ∼7 days). However, comparisons of the anneal rate for hot pixels formed at the beginning and at the end of an anneal cycle indicate that the anneal rate does not depend on this time interval. For both the ∼30 day old and the ∼1 day old hot pixels, the anneal rate was the same ∼60% and the cumulative anneal rate after 8 cycles was ∼70%. Therefore, simply decreasing the interval between annealings would not appear to have an impact.

Given our generally poor understanding of the anneal process we cannot rule out the possibility that bi-monthly annealing could impact the anneal rate in some more subtle way.

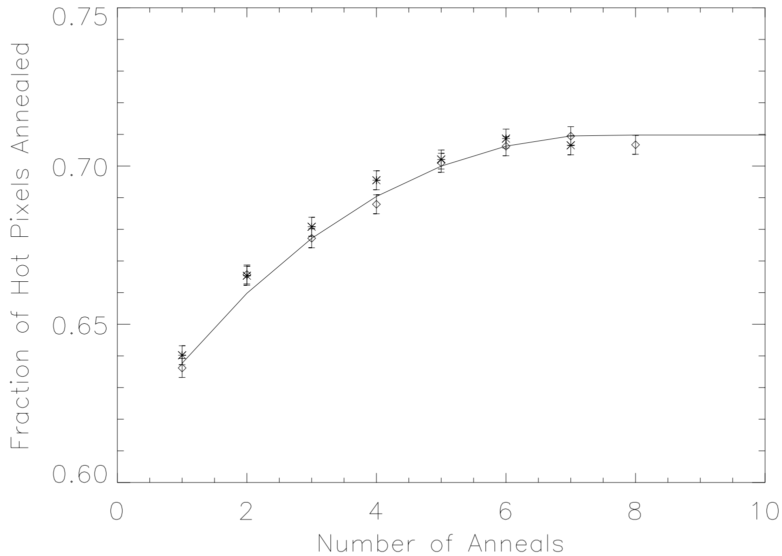

Figure 3. The cumulative fraction of hot pixels which are healed after successive anneals. Hot pixels formed during the first anneal cycle (April 2002) are shown as open symbols while the asterisks show those formed during the second anneal cycle (May 2002). After the first anneal which is ∼63% successful, succeeding anneals are less successful and after ∼7–8 anneals the cumulative anneal rate reaches an apparent plateau.

We can only say at this time the data indicates it will not help and that bi-monthly annealing is probably not a solution to the rapid hot pixel growth on WFC. We do recommend further study of this problem which might culminate in additional experiments. From past experience with other *HST* instruments, it has been shown that the length of time of the anneal is not an important parameter. However, other possibilities might include (but are not limited to) warming the CCD by pointing at the bright Earth during the anneal (or other ways of increasing the temperature of the CCD during the anneal), reading out the CCD during the anneal to increase the energy available to the lattice to break the bonds of the damaged site of the crystal, annealing during warm attitudes, or running the CCD colder to reduce the dark current in the already hot pixels.

2.2. Science Impact

Because the location of hot pixels is known from dark frames, they are readily flagged and discarded. In principle they can be corrected without discarding, but because the noise in hot pixels is greater than Poisson, corrections are of only limited value. The best strategy for mitigation is dithering. For a well-dithered image, a given fractional coverage by hot pixels of the CCD represents an equal fractional reduction in the effective exposure time. Over the next few cycles this will result in an effective reduction of exposure time of 2%–3% which will have little to no science impact. For searches for rare and faint transients (e.g., high-redshift supernovae), an additional exposure (e.g., 5 instead of 4 in an orbit) may be required in future cycles to insure each pixel is clean from contamination. Alternatively, contemporaneous dark frames can be used to reject transients found in the position of hot pixels. It is possible that a reduction in the operating temperature of the ACS WFC due to the aft-shroud cooling system could further mitigate the hot pixel problem by reducing the dark current of hot pixels. Upcoming tests include raising the temperature of the camera.

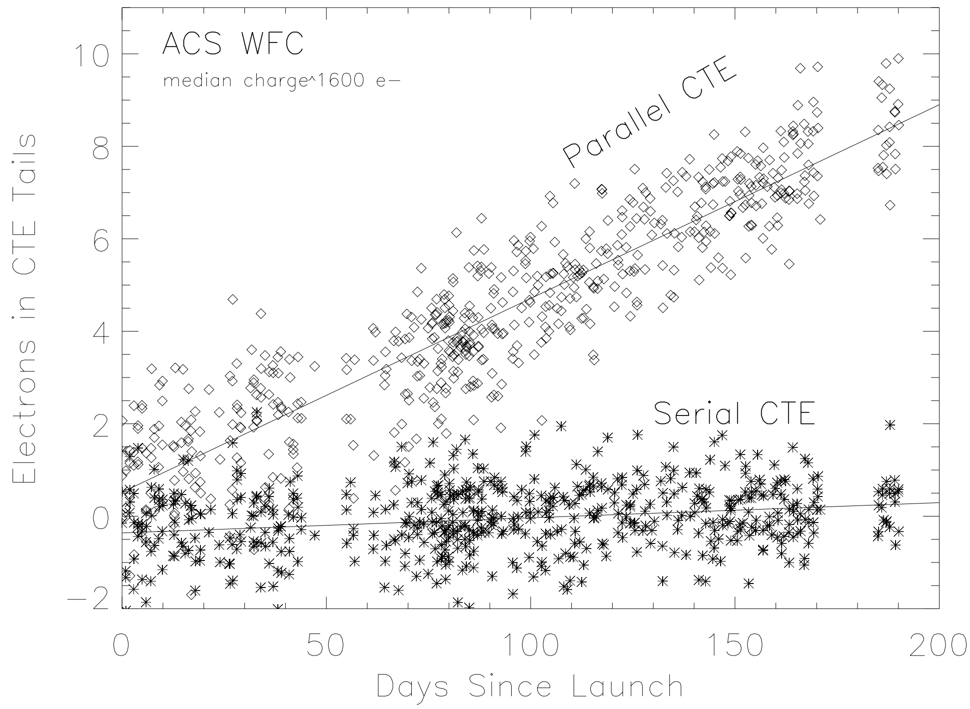

Figure 4. Degradation of ACS WFC CTE from cosmic ray tails.

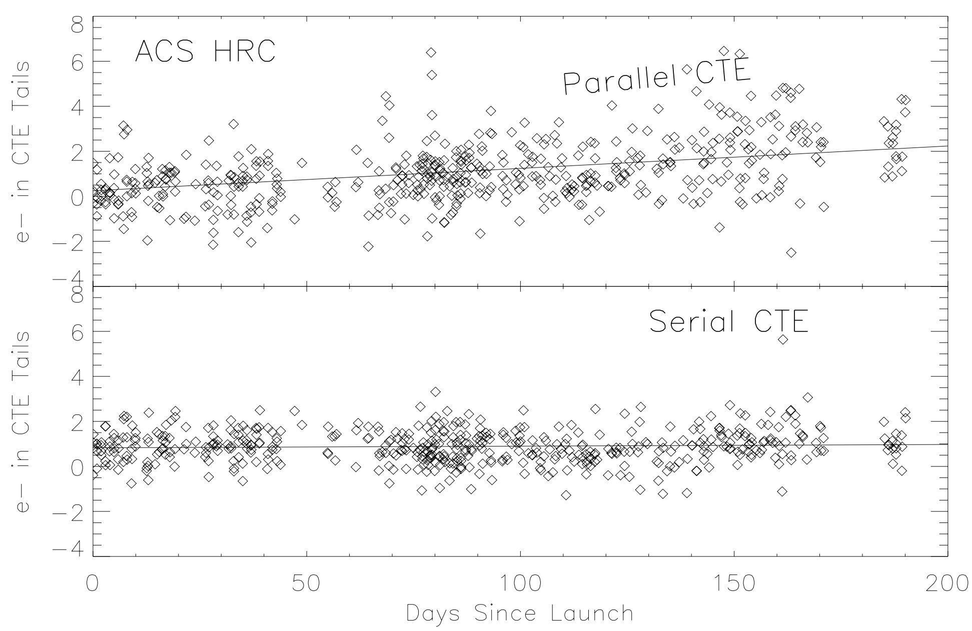

Figure 5. Degradation of ACS HRC CTE from cosmic ray tails.

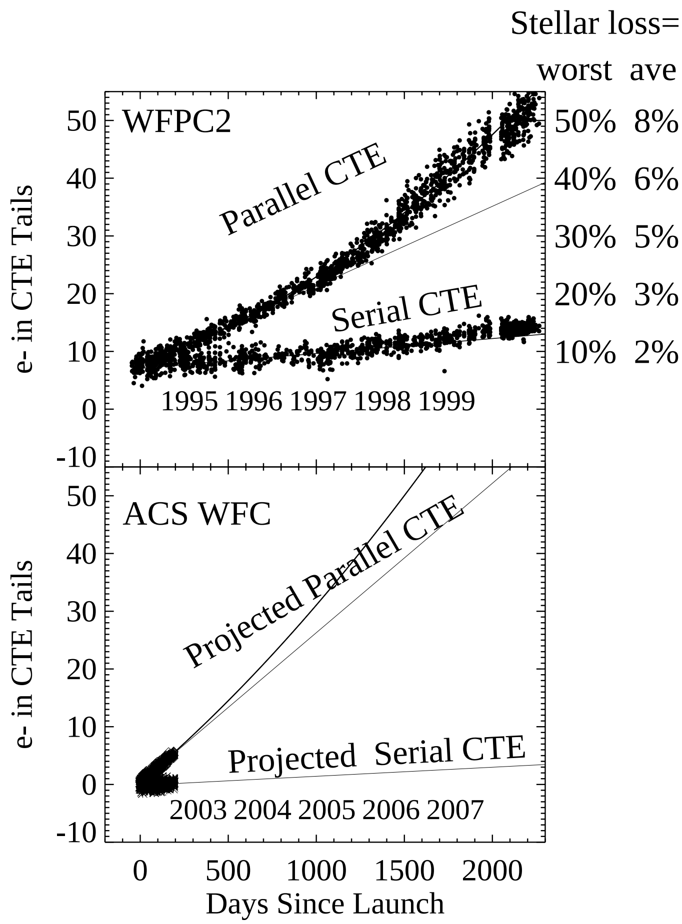

Figure 6. Projected ACS WFC CTE degradation versus WFPC2.

2.3. CTE

It's too early and the data is not yet available to determine the impact of imperfect CTE on photometry. However, we have used internal diagnostics to determine the relative degradation of CTE on WFC and HRC. As seen in Figures 4, 5, 6 of this paper, only the WFC parallel is getting markedly worse. Even this level is still not bad. Using a simple scaling from WFPC2 we would expect typical sources in the middle of the chip with average background to have only ∼1% to 2% losses to CTE (but as much as 5% to 10% in the worst cases such as very faint sources on little background at the edge of the chip). However, until an external measurement is available (in early 2003) its too soon to provide a calibration of CTE for ACS.

Acknowledgments. We wish to thank Mark Clampin, Doug van Orsow, Max Mutchler, Roeland van der Marel, Marco Sirianni and members of the ACS and ID Teams for helpful discussions.