

Cosmic Ray and Hot Pixel Removal from STIS CCD Images

Robert S. Hill and Wayne B. Landsman

Hughes STX Corp., NASA/GSFC/LASP

Don Lindler

Advanced Computer Concepts, NASA/GSFC/LASP

Richard Shaw

ST Science Institute

Abstract. The problem of cosmic ray (CR) removal is a general one plaguing spaceborne CCDs, as is the gradual accumulation of single high-dark-rate pixels between CCD annealings. The STIS team at Goddard has developed IDL implementations of standard techniques for dealing with these problems as part of the STIS GTO software. This report summarizes the methods and discusses the pitfalls.

1. Why Worry About Cosmic Rays?

Table 1 shows data obtained from long dark images on the rate of accumulation of CR pixels in a given exposure. The typical rate of accumulation is $\sim 25 - 30$ pixels/s at a detection threshold of 4σ . How long does it take to fill the entire detector with CR pixels? The situation is described by exponential decay, because only the first CR to hit a given pixel counts. Thus, the instantaneous probability of another pixel being affected by a CR is proportional to the number of remaining non-CR pixels. The “half-life” of the image in this sense is $\sim 27,000$ s.

Be that as it may, in a typical STIS CCD exposure of ~ 1000 s, $\sim 2.5\%$ of the pixels will be affected by CRs. In most cases, omitting the CR correction would interfere with object detection, spectral extraction, and photometry. Nor should the cosmetic problem be discounted, because pattern recognition by the scientist is a necessary part of data analysis.

2. Cosmic Ray Removal Methods

There are two main kinds of CR removal technique. One uses the image being cleaned to determine empirically the statistical outlier pixels. Pixels differing from some computed background by a specified threshold are repaired. Although the photons that should have been detected can never be recovered, at least the CR pixels can be flagged and cosmetically improved.

The most rigorous kind of CR removal uses multiple images together with the known, calibrated readout noise and gain of the CCD. Outlier pixel values are defined in relation to the expected distribution of differences between the input images. Usually, the STIS/CCD or WFPC2 observer plans ahead to divide long exposure times into two or more actual exposures (CR-splits) in order to use this method. However, for faint, extended sources, there is a potential trade-off, since with N exposures, the read noise of the final co-added image increases by a factor of $N^{0.5}$ over a single exposure of the same duration.

Table 1. Accumulation Rates of Pixels Affected by Cosmic Rays

Date	Number of Pixels	Exp. Time (s)	Rate (pixels s ⁻¹)	STISLOG Entries
50527.71	228032	5*1800	25.3	716,720,722,724,726
50531.08	221603	5*1800	24.6	737,739,741,745,747
50546.67	296034	6*1800	27.4	792,794,796,827,829,831
50563.83	218052	4*1800	30.3	941,943,945,955
50573.17	235042	5*1800	26.1	1026,1030,1032,1034,1036
50580.09	224395	5*1800	24.9	1104,1106,1108,1110,1120
50587.11	256936	5*1800	28.5	1269,1271,1273,1275,1277
50594.41	270249	5*1800	30.0	1518,1520,1666,1668,1670
50601.22	204008	4*1800	28.3	1973,1975,1977,2006
50605.19	239604	5*1800	26.6	2119,2123,2125,2127,2138
50612.53	178210	5*1350	26.4	2421,2423,2425,2427,2433
50618.08	152793	5*1350	22.6	2701,2707,2709,2711,2713
50618.52	157647	5*1200	26.3	2732,2734,2736,2738,2740
50619.70	191219	5*1200	31.9	2750,2752,2754,2756,2758
50625.10	145831	5*1350	21.6	3035,3037,3039,3041,3043
50631.74	173537	5*1350	25.7	3191,3193,3195,3197,3199
50639.02	131846	4*1350	24.4	3359,3365,3367,3369
50646.49	155334	5*1350	23.0	3691,3693,3695,3697,3701
50653.04	215986	5*1350	32.0	3865,3867,3869,3871,3873
50660.12	214942	5*1350	31.8	4033,4035,4039,4041,4043
50667.10	154941	5*1350	22.9	4168,4170,4172,4174,4176
50674.09	155016	5*1350	23.0	4324 4326 4328 4330 4332
50681.03	148671	5*1350	22.0	4540,4542,4544,4548,4550
50685.13	184481	5*1350	27.3	4613,4615,4617,4619,4621

3. What are Hot Pixels?

Hot pixels are individual pixels with a high dark current, persistent and occurring at fixed positions on the detector. If left alone, they continually increase in number. However, many of them can be repaired physically by turning off the thermoelectric cooler (TEC) and letting the CCD warm from its operating temperature of -83°C to -5°C , where it remains for 12 hours. WFPC2 experience indicates that it may be possible to reach a near-steady state in which the net growth rate in the number of hot pixels is $\sim 8\%$ of the instantaneous growth rate (Kimble 1997).

Another poster in this conference (Beck et al.) discusses hot pixels.

4. Hot Pixel Repair Methods

There are two ways to remove hot pixels, analogous to the two ways of removing CRs. One way is to subtract the excess signal from each hot pixel, leaving behind the legitimate astronomical source flux. The other way is to estimate the correct pixel value from surrounding pixels on the same image.

Users of the GTO IDL software often invoke a hybrid method. Many pixels perceived as hot in an image display are actually removed quite well by subtracting a standard weekly dark, which is a high S/N image made by combining five long-duration dark frames. After this step, the software uses a pixel list that is generated from the same darks. The rates given in the list are only used to decide which pixels to correct, and the corrected values are estimated from the surrounding pixels. Thus, two ways of fixing hot pixels are combined. The user must decide at what count rate to put the transition between pure dark subtraction, and dark subtraction followed by interpolation.

5. Statistical CR-Removal Programs

A classification of some commonly used CR-removal programs is as follows (Massey 1997, Wells & Bell 1994, and on-line documentation for the various programs):

1. Several images
 - (a) Empirical noise
 - i. Iterative: None apparently in common use; such an algorithm was used in processing ground-based STIS/CCD flats
 - ii. Non-iterative
 - A. IRAF tasks `combine` and `imcombine` with some options, e.g., `avsigclip`
 - B. Similarly for STSDAS task `gcombine` (works on multi-group GEIS data)
 - (b) Calibrated noise model
 - i. Iterative
 - A. IDL program `cr_reject`, as called by `stis_cr` in the GTO IDL system
 - B. STSDAS task `hst_calib.wfpc.crrej`, or `calstis-2` in the STIS pipeline
 - ii. Non-iterative
 - A. IRAF tasks `combine` and `imcombine`, with the `crreject` option
 - B. Similarly for STSDAS task `gcombine` (works on multi-group GEIS data)
2. One image: IRAF task `cosmicrays`

6. Hot Pixel Removal Programs

A classification of some commonly used hot-pixel repair programs is as follows (Massey 1997, Wells & Bell 1997, as well as on-line documentation for the various programs):

1. History-based
 - (a) Dark frame subtraction
 - i. GTO IDL program `calstis` (`darkfile` option)
 - ii. IRAF task `ccdproc` (`darkcor` option)
 - iii. STSDAS pipeline tasks `calwp2`, `calstis-1` (`darkcorr` option)
 - (b) Hot pixel lists or masks
 - i. GTO IDL program `calstis` (`hrepair` option)
 - ii. IRAF task `ccdmask`, followed by `fixpix`
 - iii. STSDAS tasks `warpix` or `wfixup`
2. Single image
 - (a) GTO IDL program `hotterp`
 - (b) IRAF program `cosmicrays`

7. What is the Standard CR Rejection Method for STIS?

The Institute pipeline and the GTO IDL software use very similar programs to remove CRs. Both are based on the STSDAS task `crrej` in the package `hst_calib.wfpc`. The algorithm is iterative, and it uses a calibrated noise model. Typically, the first iteration clips at some high number of σ , such as 6 or 8, then ramps down to 3 or 4σ . The STScI routine is called `calstis-2`, and it can be invoked from STSDAS. The GTO IDL program, which is called `cr_reject`, is usually invoked as part of the higher-level routine `stis_cr`.

No one set of parameters for either of these routines can handle all cases. STScI estimates that the `calstis-2` product will be quantitatively usable for $\sim 50\%$ of the observations (Baum et al. 1996). The STScI pipeline tunes the parameters of `calstis-2` depending on the observation, as shown in Table 2.

8. Some Subtleties of the Algorithms

Besides the iterative $N\sigma$ clip, both programs mentioned above have additional features:

- The CR flags can be propagated in the neighborhood of the initially detected CR pixels, to take into account a failure to detect fainter pixels around the CR edges
- The initial guess at a CR-free image can be either the pixel-by-pixel minimum of the input images, or the pixel-by-pixel median; for only 2 input images, there is really no other good way than starting with the minimum
- The STScI algorithm clears the CR flags between iterations, so that a pixel that has once been flagged can regain its good standing if the average creeps back up toward it; the GTO IDL algorithm by default says, “once a CR, always a CR,” but allows re-initialization as an option; this difference probably only affects tight clips, say, 2 iterations at 2.5 (N.B.: the GTO IDL default may change)

Table 2. CR Rejection Parameters in STScI STIS Pipeline

Inputs		Outputs ^a	
CR-Split Images	Mean Exp. Time	Initial Guess	Clipping Sigmas
2	3.	minimum	4.0
2	21.	minimum	3.5
2	100.	minimum	3.0
2	3.0E5	minimum	3.0
3	2.	minimum	4.0
3	17.	minimum	3.5
3	82.	minimum	3.0
3	3.0E5	minimum	3.0
4	2.	median	4.0
4	15.	median	3.5
4	71.	median	4,3,0
4	319.	median	5,4,3
4	3.0E5	median	6,5,4,3
5	2.	median	4.0
5	13.	median	3.5
5	63.	median	4,3,0
5	1479.	median	5,4,3
5	3.0E5	median	6,5,4,3
6	2.	median	4.0
6	12.	median	3.5
6	58.	median	4,3,0
6	1479.	median	5,4,3
6	3.0E5	median	6,5,4,3
7	2.	median	4.0
7	11.	median	3.5
7	53.	median	4,3,0
7	3187.	median	5,4,3
7	3.0E5	median	6,5,4,3
8	1.	median	4.0
8	10.	median	3.5
8	50.	median	4,3,0
8	3187.	median	5,4,3
8	3.0E5	median	6,5,4,3

^aThe following outputs are constant: no sky adjustment; no scale noise; CR propagation never done; no input quality flags avoided; mask always produce

- The input images can be sky-adjusted before CR rejection; the GTO IDL program uses a DAOPHOT-like sky computation, whereas `crrej` finds the histogram mode directly
- Scale noise, such that $\sigma = k \times \text{flux}$, can be added into the noise model; this is used either for images that are already flat-fielded (not the pipeline default) or in case of a slight PSF mismatch

9. What Are the Potential Problems?

- Any sufficiently stringent CR removal program will remove some real data because statistical fluctuations can be mistaken for CR pixels; be cautious, and do not go for better than a 3 or 4 σ clip in the final iteration
- PSF mismatches or image shifts, even at the sub-pixel level, cause real flux to be removed
- Some residual CR pixels will be added into the final image, resulting in noise that affects faint source detection; one may need to model this, depending on the goal of the analysis
- Some hot pixels will usually be left over, regardless of the method used, and may need fixing by hand

To assure good results from CR rejection, the user should do the following:

- Check both the co-alignment of the CR-split images and the consistency of the PSFs
- Verify that flux is conserved in the result
- Examine the mask images that most programs produce in order to show which pixels are rejected; if a mask shows the obvious shape of a real source or the contour of a steep flux gradient, then the work probably needs to be redone with new parameters or a new procedure

The user should not take any CR-rejection result on faith. Depending on the situation, even a simple algorithm can sometimes do the job, and with the wrong parameters, even the most sophisticated algorithm can either reject good pixels or retain too many bad ones. Checking the general performance of a program using darks or other test data is not sufficient, as the presence of sources affects the outcome.

References

- Baum, S., Hsu, J. C., Hodge, P., & Ferguson, H. 1996, The STScI Pipeline V: Cosmic Ray Rejection, STIS ISR 96-018
- Kimble, R. et al. 1997, ApJL submitted
- Massey, P., 1997, A User's Guide to CCD Reductions with IRAF (NOAO IRAF Web site)
- Wells, L. A. & Bell, D. J. 1994, Cleaning Images of Bad Pixels and Cosmic Rays using IRAF (NOAO IRAF Web site)