1997 HST Calibration Workshop
Space Telescope Science Institute, 1997
S. Casertano, et al., eds.

A Package for the Reduction of Dithered Undersampled Images

A. S. Fruchter[1], R. N. Hook[2], I. C. Busko[1] and M. Mutchler[1]

[1]*Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21210*
[2]*Space Telescope European Coordinating Facility, D–85748 Garching, Germany*

Abstract. We present a set of tasks developed to process dithered undersampled images. These procedures allow one to easily determine the offsets between images and then combine the images using Variable-Pixel Linear Reconstruction, otherwise know as "**drizzle**". This algorithm, originally developed for the combination of the images in the Hubble Deep Field(Williams et al. 1996, Fruchter & Hook 1997), preserves photometry and resolution, can weight input images according to the statistical significance of each pixel, and removes the effects of geometric distortion both on image shape and photometry. In this paper, the method and its implementation are described, and measurements of the photometric accuracy and image fidelity are presented. In addition, we describe ancillary tasks developed for determining offsets between images, and discuss the use of drizzling to combine dithered images in the presence of cosmic rays.

1. Why Dither?

Although we have all observed the usefulness of dithering throughout our social lives, the attendees at this conference may find it surprising that a technique so important to politicians should also improve the quality of one's Space Telescope data. However, dithering or spatially offsetting the telescope detector between exposures has several benfits.

- Large scale dithering allows one to reduce the effects of flat field errors or spatially varying detector sensitivity on the final image.
- Shifts of a few pixels allow one to remove small scale defects such as hot pixels, bad columns, and charge traps from the image.
- Non-integral dithers allow one to recover some of the information lost to undersampling.

The latter reason is particularly important in the case of Hubble Space Telescope (HST) imaging, for while the telescope now provides the superb images for which it was designed, the detectors on HST are frequently unable to take full advantage of the resolving power of the optics. This is particularly true of the Wide Field Camera (WF) of WFPC2 and Camera 3 of NICMOS (NIC3). The width of a WF pixel equals the full-width at half maximimum of the optics in the the near-infrared, and greatly exceeds it in the blue. The NIC3 similarly undersample the image over much of its spectral range. The effect of undersampling on WF images is illustrated by the "Great Eye Chart in the Sky" in Figure 1.

Fortunately, much of the information lost to undersampling can be restored. In the lower right of Figure 1 we display a restoration made using one of the family of techniques we refer to as "linear reconstruction." The most commonly used of these techniques are shift-and-add and interlacing. The image in the lower right corner has been restored by interlacing dithered images. However, due to the occasional small positioning errors of

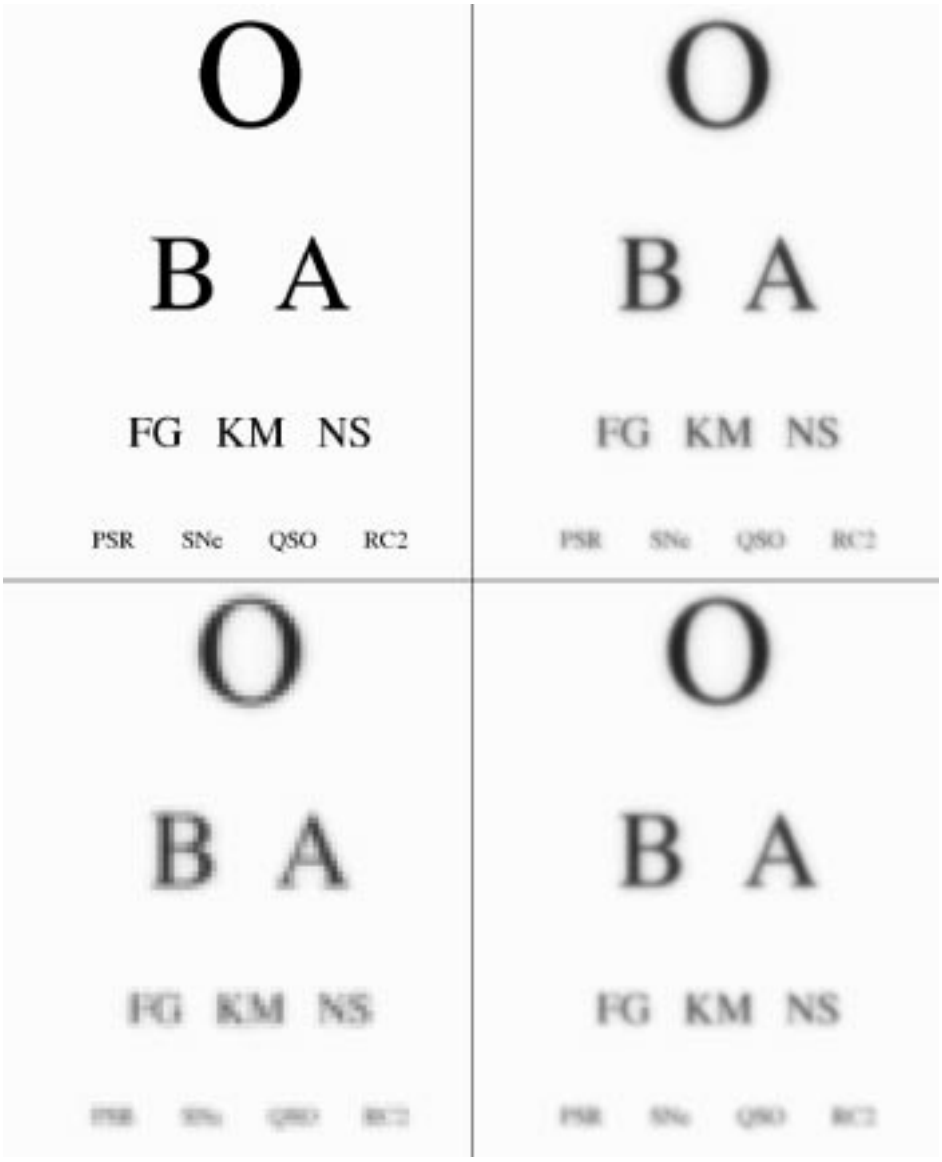

Figure 1. In the upper left corner of this Figure, we present the "true image", i.e., the image one would see with an infinitely large telescope. The upper right shows the image after convolution with the optics of the Hubble Space Telescope and the WFPC2 camera—the primary wide-field imaging instrument presently installed on the HST. The lower left shows the image after sampling by the WFPC2 CCD, and the lower right shows a linear reconstruction of dithered CCD images.

telescope and the non-uniform shifts in pixel space caused by the geometric distortion of the optics, true interlacing of HST images is often infeasible. The other standard linear reconstruction technique, shift-and-add, can easily handle arbitrary dither postions, but it convolves the image yet again with the orginal pixel, adding to the blurring of the image and to the correlation of the noise. The importance of avoiding unnecessary convolution of the image with the pixel is emphasized by comparing the upper and lower right hand images in Figure 1. The deterioration in image quality is due entirely to convolution of the image by the WF pixel. In the next Section we present a new method, **drizzle**, which has

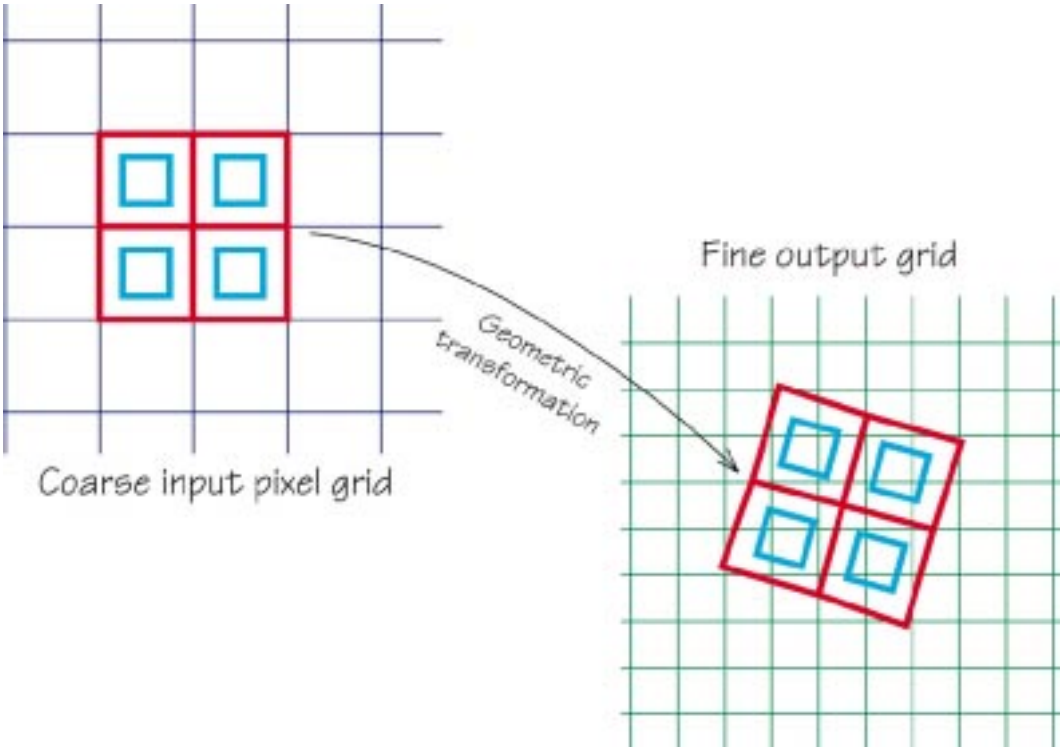

Figure 2. A schematic representation of drizzling. The input pixel grid (shown on the left) is mapped onto a finer output grid (shown on right), taking into account shift, rotation and geometric distortion. The user is allowed to "shrink" the input pixels. We refer to these shrunken pixels as drops (faint inner squares). A given input image only affects the output image pixels under drops. In this particular case, the central output pixel receives no information from the input image. Therefore, the dropsize (*pixfrac*) shown here would only be appropriate were many more images to be drizzled onto the output.

the versatility of shift-and-add yet largely maintains the resolution and independent noise statistics of interlacing.

2. The Method

The **drizzle** algorithm is conceptually straightforward. Pixels in the original input images are mapped into pixels in the subsampled output image, taking into account shifts and rotations between images and the optical distortion of the camera. However, in order to avoid convolving the image with the large pixel "footprint" of the camera, we allow the user to shrink the pixel before it is averaged into the output image, as shown in Figure 2.

The new shrunken pixels, or "drops", rain down upon the subsampled output. In the case of the Hubble Deep Field (HDF), the drops used had linear dimensions one-half that of the input pixel—slightly larger than the dimensions of the output subsampled pixels. The value of an input pixel is averaged into an output pixel with a weight proportional to the area of overlap between the "drop" and the output pixel. Note that, if the drop size is sufficiently small, not all output pixels have data added to them from each input image. One must therefore choose a drop size that is small enough to avoid degrading the image, but large enough so that after all images are "dripped" the coverage is fairly uniform.

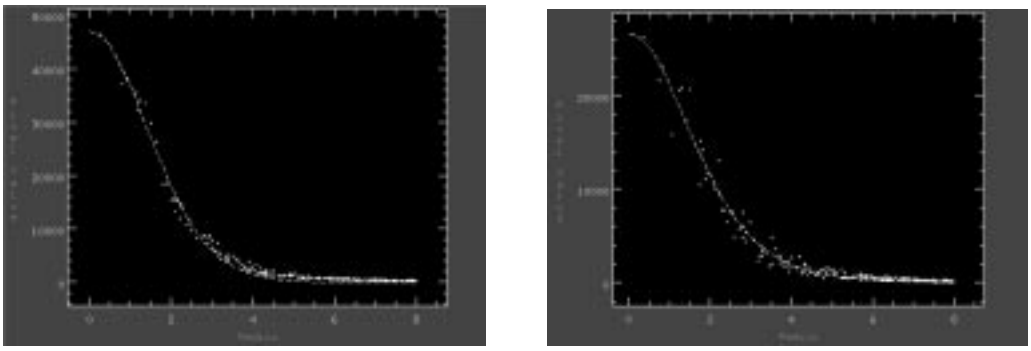

Figure 3. A comparison of two PSFs from the HDF. Note that one of these PSFs shows substantial noise about the Gaussian. Although some noise is expected as a result of the interpolation performed by **drizzle**, WFPC2 images occasionally show more noise than simulations would predict. This may be the result of defects in the detector, and could be related to the known charge transfer errors in WFPC2. A pixel in the processed HDF image has a linear size of $0\farcs04$.

The drop size is controlled by a user-adjustable parameter called *pixfrac*, which is simply the ratio of the linear size of the drop to the input pixel (before any adjustment due to the geometric distortion of the camera). Thus interlacing is equivalent to setting $pixfrac = 0.0$, while shift-and-add is equivalent to $pixfrac = 1.0$.

When a drop with value i_{xy} and user-defined weight w_{xy} is added to an image with pixel value I_{xy}, weight W_{xy}, and fractional pixel overlap $0 < a_{xy} < 1$, the resulting value of the image I'_{xy} and weight W'_{xy} is

$$W'_{xy} = a_{xy} w_{xy} + W_{xy} \quad (1)$$

$$I'_{xy} = \frac{a_{xy} i_{xy} w_{xy} + I_{xy} W_{xy}}{W'_{xy}} \quad (2)$$

This algorithm has a number of advantages over standard linear reconstruction methods presently used. Since the area of the pixels scales with the Jacobian of the geometric distortion, **drizzle** preserves both surface and absolute photometry. Therefore flux can be measured using an aperture whose size is independent of position on the chip. As the method anticipates that a given output pixel may receive no information from a given input pixel, missing data (due for instance to cosmic rays or detector defects) do not cause a substantial problem, so long as there are enough dithered images to fill in the gaps caused by these zero-weight input pixels. Finally, the linear weighting scheme is statistically optimum when inverse variance maps are used as weights.

3. Image Fidelity

Drizzle was designed to obtain optimal signal-to-noise on faint objects while preserving image resolution. These goals are unfortunately not fully compatible. For example, non-linear image restoration procedures, which attempt to remove the blurring caused by the point-spread function (PSF) and the large pixel by enhancing the high frequencies in the image (such as the Richardson-Lucy (Richardson 1972, Lucy 1974, Lucy & Hook 1991) and maximum entropy methods (Gull & Daniell 1978, Weir & Djorgovski 1990)), directly exchange signal-to-noise for resolution. In the drizzling algorithm no compromises on signal-to-noise have been made; the weight of an input pixel in the final output image is entirely

independent of its position on the chip. Therefore, if the dithered images do not uniformly sample the field, the "center of light" in an output pixel may be offset from the center of the pixel, and that offset may vary between adjacent pixels. The large dithering offsets which may be used in WFPC2 imaging, combined with geometric distortion, can produce a sampling pattern that varies across the field. The output PSFs produced by the combination of such irregularly dithered datasets may, on occasion, show significant variations about the best-fit Gaussian. Fortunately this does not noticeably affect aperture photometry performed with typical aperture sizes. In practice the variability about the Gaussian appears larger in WFPC2 data than our simulations would lead us to expect. Examination of recent dithered stellar fields leads us to suspect that this excess variability results from a problem with the original data, possibly caused by charge transfer errors in the CCD (Whitmore & Heyer 1997).

4. Photometry

The WFPC2 optics geometrically distort the images: pixels at the corner of each CCD subtend less area on the sky than those near the center. However, after application of the flat field, a source of uniform surface brightness on the sky produces uniform counts across the CCD. Therefore point sources near the corners of the chip are artificially brightened compared to those in the center.

In order to study the ability of **drizzle** to remove the photometric effects of geometric distortion, we have created a a sub-sampled grid of artificial stellar PSFs and adjusted their counts to reflect the effect of geometric distortion—the stars in the corners are up to $\sim 4\%$ brighter than those in the center of the chip. This image was then shifted and down-sampled onto four simulated WF frames and the results combined using drizzling with typical parameters. Aperture photometry on the grid of stars after drizzling reveals that the effect of geometric distortion on the photometry has been dramatically reduced: the RMS photometric variation in the drizzled image is 0.004 mags.

5. The Dither Package

In order to effectively use the **drizzle** algorithm, one must know the offset and rotation between input images. We have therefore developed a package of tasks to assist the user through the entire process of combining dithered images. This STSDAS package is know as "`dither`" and presently contains the following tasks:

- `precor`: determines regions of the image containing astrophysical objects and nulls the remainder of the image, thereby substantially reducing the effect of cosmic rays and chip defects on the offset measurement. The ouput from precor is only used for offset determination and not final image creation.

- `offset`: cross-correlates all four images in a WFPC image, creating output cross-correlation images with names that can be appropriately grouped by later tasks. Offset employs the task `crossdriz` to perform the cross-correlation.

- `crossdriz`: cross-correlates two images, after preprocessing which includes trimming, and, if requested, drizzling to remove geometric distortion or rotation. `crossdriz` will also perform a loop over a range of test rotation angles.

- `shiftfind`: locates the peak in a cross-correlation image and fits for sub-pixel shift information. The search region and details of the fitting can be adjusted by the user.

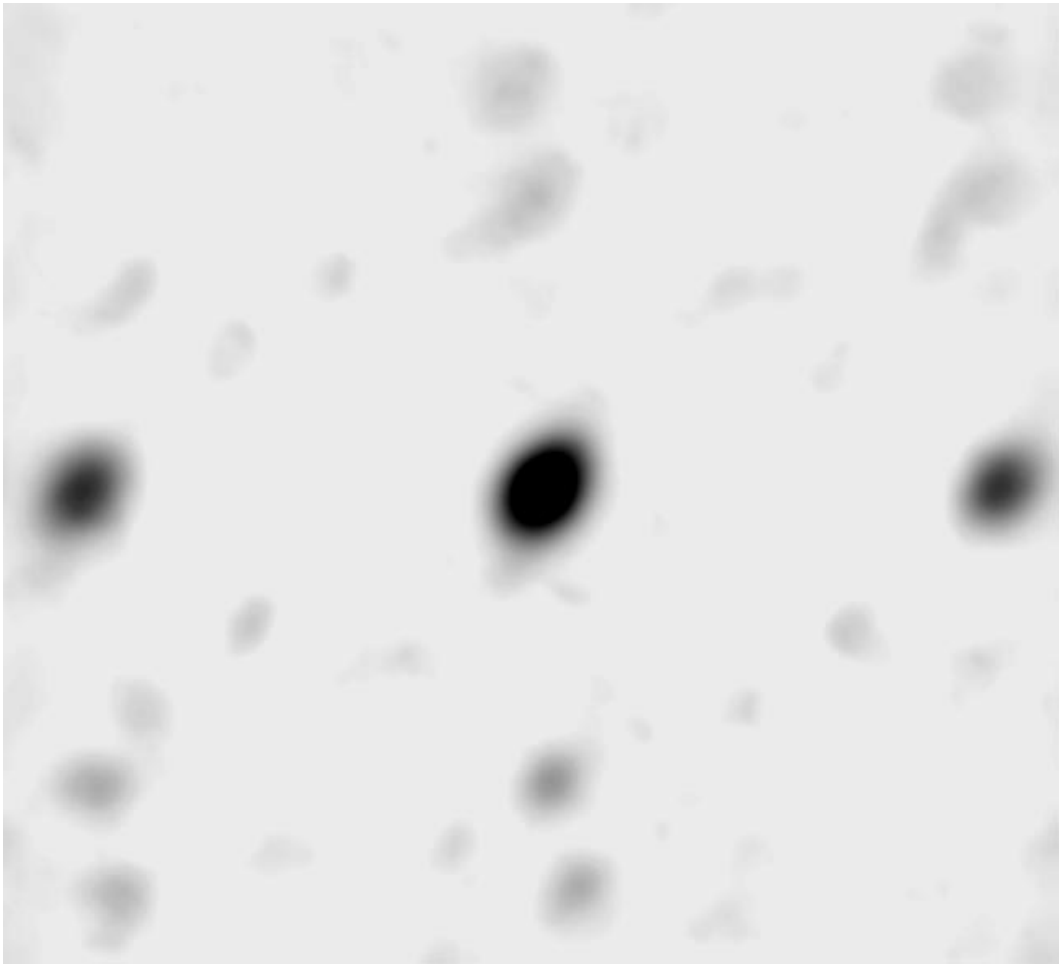

Figure 4. The cross-correlation image of one of the WF chips from two dither positions in the HDF. Note that the cross-correlation is symmetric about the maximum, and reflects the dominant objects in the image. In this case, a bright elliptical galaxy, rather than a star, dominates. Nonetheless, one can determine the shift between these two chips to about 1/20th of a pixel. The shift between the images is equal to the offset between the peak of the cross-correlation and the center of the image.

- `rotfind`: fits for the rotation angle between two images. `rotfind` is called when `crossdriz` has been used to loop over a range of test rotation angles between two images.

- `avshift`: determines the shifts between two WFPC2 images by averaging the results obtained on each of the groups after adjusting for the rotation angles between the four groups. `avshift` can also be used to estimate the rotation angle between two different WFPC2 images, when the rotation angle is a small fraction of a degree.

- `blot`: maps a drizzled image back onto an input image. This is an essential part of the tasks we are developing for removing cosmic rays from singly-dithered images.

Although a detailed exposition of all of these tasks is beyond the present review, here we will describe a number of the `dither` tasks and typical results from their use.

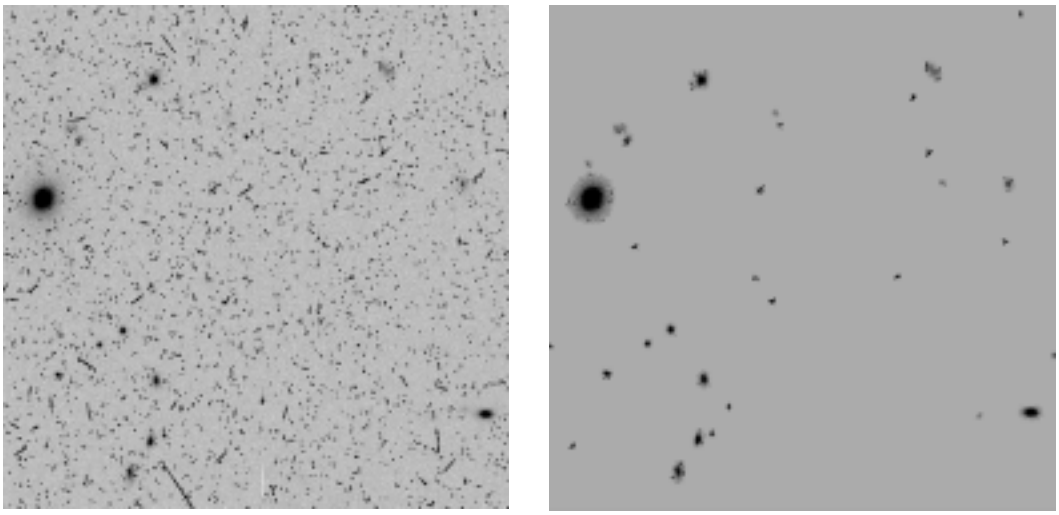

Figure 5. A comparison of a raw image and the output image from `precor`. The sky is set to zero, while true objects are preserved. As a result, essentially all cosmic rays, with the exception of those on top of detectable astronomical objects, are removed.

The major engine for determining the offset between two images in the `dither` package is cross-correlation. The following equations explain the reason for the choice of this algorithm. Assume that one has a known template image, $T(x_{i,j})$, and an observed image, $I(x_{i,j})$, and one wishes to know the offset between the observed and template images. Then

$$\sum_{i=1,j=1}^{N,N} (I(x_{i,j}) - T(x_{i,j}))^2 =$$

$$\underbrace{\sum_{i=1,j=1}^{N,N} I(x_{i,j})^2 + \sum_{i=1,j=1}^{N,N} T(x_{i,j})^2}_{\text{These terms invariant under shifts and rotations}}$$

$$-2 \sum_{i=1,j=1}^{N,N} I(x_{i,j}) T(x_{i,j}) \ .$$

Thus minimizing the difference of the sum of squares under shifts and rotations of $I(x_{i,j})$ is equivalent to maximizing the cross-correlation between the two images! If the errors in the images are Gaussian, then this is equivalent to minimizing χ^2. In practice, of course, images are a mixture of Gaussian and Poisson noise as well as, in many cases, very highly skewed noise additions such as cosmic rays and hot pixels. These can be largely removed by a task we call `precor`. This task separate objects from hot pixels and cosmic rays, by determining whether the fraction of pixels above the sky in a given box is above a user defined threshold. `precor` is not perfect, but is able to distinguish the vast majority of these defects from true objects. The regions of the sky-subtracted images containing objects are kept; other areas are set to zero. The image can also be converted from counts to signal-to-noise ratio, thus removing the distinction between Gaussian and Poisson noise in the minimization of χ^2 (in practice this step also appears to help reduce the errors

introduced by the spatial undersampling of images). Figure 5 shows an HST image before and after the application of `precor`.

After the cross-correlation image has been created, `shiftfind` determines the shift by locating the peak of the image. In the case of a single image (with no rotation) one is now done. However, in the case of WFPC, all four chips have been shifted (and rotated) by the same physical amount, and therefore one can average the shift values determined from each of the individual chips to obtain a better estimate of the shift. The script `avshift` performs this task, taking into account the (fixed) rotations between the four chips. Furthermore, in the case of a rotation between the template image and the observed image of significantly less than one degree, one can use `avshift` to estimate the rotation.

Below is an example of applying `avshift` to two dither positions of the HDF. In acquiring the second of these dither positions, the fine guidance sensors locked up on a secondary null, and as a result the pointing of the telescope was off by about an arc second, and was rotated with respect to the desired orientation by a few arc minutes. In the first use of `avshift` the user puts in a suggested rotation of 0.00 degrees, and a large error is found in the shift determination. However, in the second attempt the user (correctly) guesses the correct orientation and the error becomes a small fraction of a pixel. In practice, it only takes a few tries with `avshift` to home in on the rotation angle.

```
di> avshift pos2x11.out 0.00
Assumed angle of rotation = 0.  degrees
Weights: 1. 1. 1. 1.
#
# Image        xsh_in   ysh_in best_xsh best_ysh tot_sh_in delta_xsh delta_ysh
pos2x11_g1     45.068   31.945   45.709   31.410   55.461   -0.641     0.536
pos2x11_g2     14.095  -20.226   14.559  -20.780   25.372   -0.464     0.553
pos2x11_g3    -21.278  -14.065  -20.858  -14.446   25.372   -0.420     0.381
pos2x11_g4    -14.839   21.432  -14.245   20.996   25.372   -0.594     0.436
# RMS error (all chips on WF scale)):   0.6187
# RMS error (WF only):     0.6795
# weighted RMS error (all chips on WF scale)):   0.6187
# weighted RMS error (WF only):     0.6795
di> avshift pos2x11.out 0.08
Assumed angle of rotation = 0.08  degrees
Weights: 1. 1. 1. 1.
#
# Image        xsh_in   ysh_in best_xsh best_ysh tot_sh_in delta_xsh delta_ysh
pos2x11_g1     45.068   31.945   45.068   32.066   55.311    0.000    -0.121
pos2x11_g2     14.095  -20.226   14.105  -20.231   24.663   -0.011     0.005
pos2x11_g3    -21.278  -14.065  -21.289  -14.018   25.490    0.012    -0.047
pos2x11_g4    -14.839   21.432  -14.821   21.416   26.044   -0.018     0.016
# RMS error (all chips on WF scale)):   0.0392
# RMS error (WF only):     0.0322
# weighted RMS error (all chips on WF scale)):   0.0392
# weighted RMS error (WF only):     0.0322
di>
```

In cases where the rotation angle is large, or where one only has one image, we use a different approach. We apply test rotations to the template and cross-correlate after each test rotation. We then find the best rotation angle by finding the maximum of the cross-correlation in this three-dimensional space of shifts and rotations. This is done by fitting for the rotation angle that produces the highest cross-correlation, and then interpolating for the precise shift using the shifts determined by the cross-correlation of the two best template

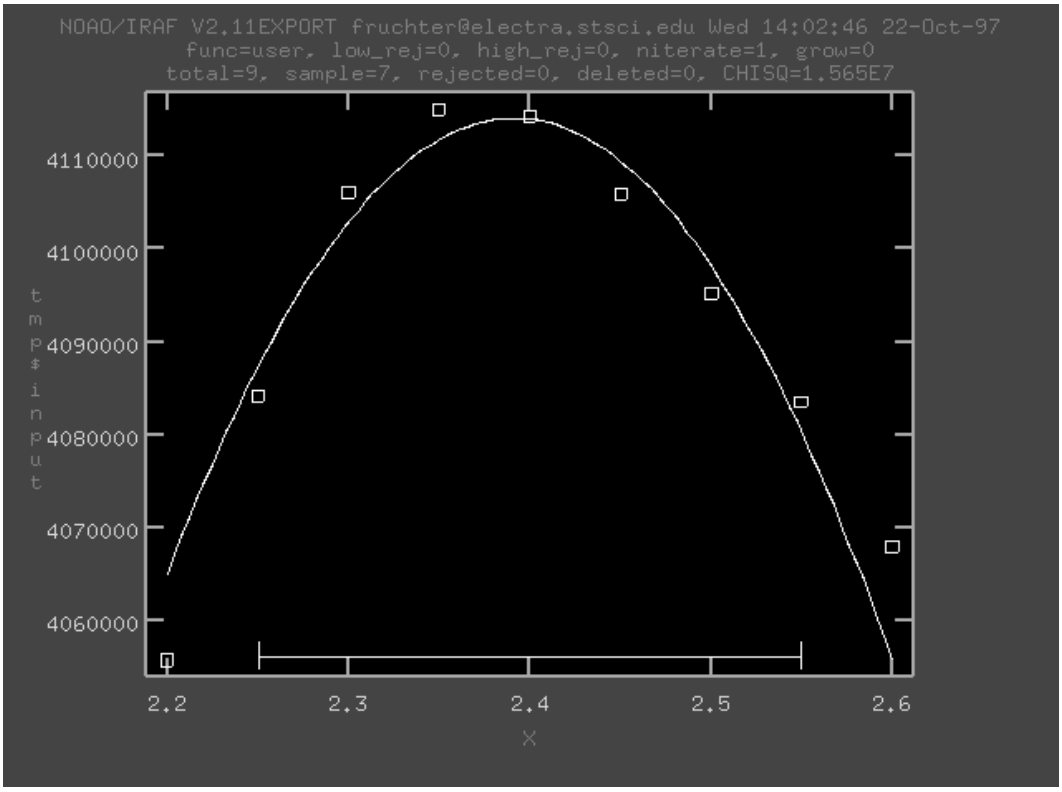

Figure 6. The graphical output of the task `rotfind`, used in this case two determine the rotation angle between two WFPC2 images of GRB 970228. According to the header, the change in orientation between these two images was 2.4 degrees. An error of 0.01 degrees is typical in measuring the rotation between two WFPC images.

rotations. The task `rotfind` is used to fit the rotation angle; an example of its graphical output is shown in Figure 6. Comparison of single group images produces a measurement which agrees with that in the header to 0.01 degrees. This level accuracy is typical for the comparison of two single group WFP2 images, and would result in an error of about 0.1 pixels at the very edge of the chip. Some of this error undoubtedly comes from the difficulty of attempting to interpolate an undersampled image. Higher accuracy should be obtainable (though is probably rarely required) by comparing the image with a first pass drizzled output of many dithered images of the same field.

6. Cosmic Ray Detection

Few HST observing proposals have sufficient time to take a number of exposures at each of several dither positions. Therefore, if dithering is to be of wide-spread use, one must be able to remove cosmic rays from data where few, if any, images are taken at the same position on the sky. We have therefore begun to adapt **drizzle** to the removal of cosmic rays.

A technique which appears quite promising is described below. We are presently working on creating a suite of robust scripts to perform these procedures. At present one of the authors (ASF) has a suite of very fragile scripts to do these operations. It is possible that some version of these scripts will be made publicly available in an unofficial, caveat-emptor release. Information will be posted on the web pages mentioned below.

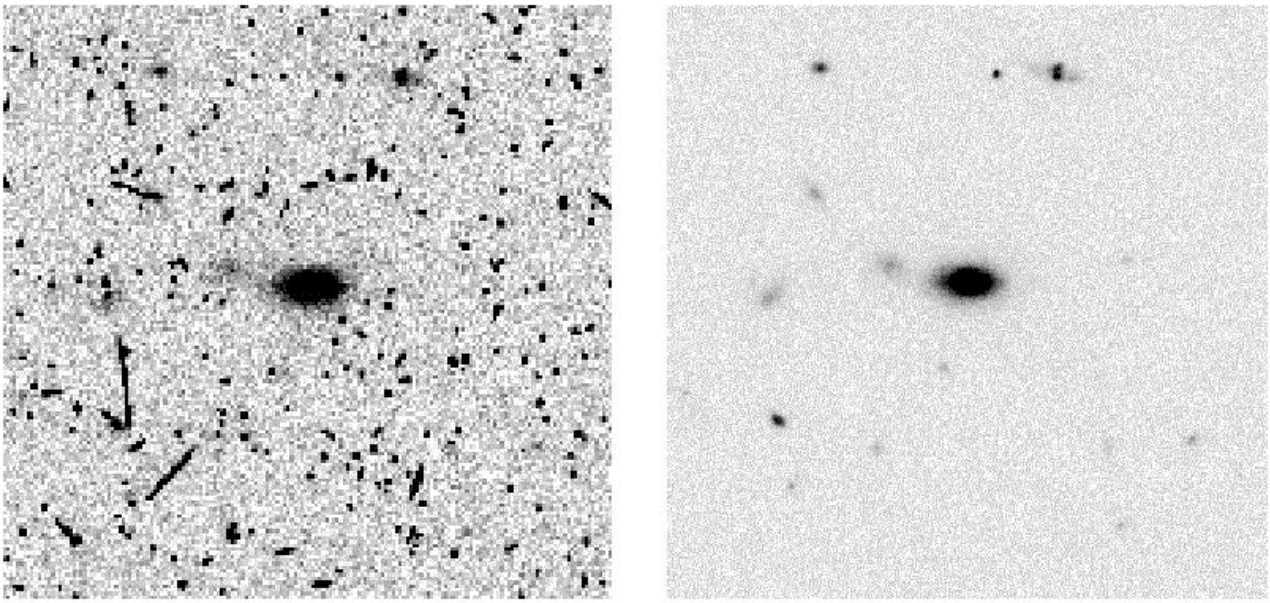

Figure 7. The image on the left shows a region of one of twelve 2400s archival images taken with a wide near-infrared filter on WFPC2. Numerous cosmic rays are visible. On the right is the drizzled combination of the twelve images, no two of which shared a dither position.

1. `drizzle` each image onto a separate sub-sampled output image using $pixfrac = 1.0$
2. Take the median of the output drizzled images.
3. Map the median image back to the input plane of each of the individual images, taking into accout the image shifts and geometric distortion. The is now done by interpolating the values of the median image using a program we call "`blot`".
4. Take the spatial derivative of each of the blotted output images.
5. Compare each original image with the corresponding blotted image. Where the difference is larger than can be explained by noise statistics, orthe flattening effect of taking the median, or perhaps an error in the shift (the magnitudes of the latter two effects are estimated using the image derivative), the suspect pixel is masked.
6. Repeat the previous step on pixels adjacent to pixels already masked, using a more stringent comparison.
7. Finally, `drizzle` the input images onto a single output image using the pixel masks created in the previous steps.

Figure 7 shows the result of applying this method to data originally taken by Cowie and colleagues (Cowie, Hu, & Songaila 1995), which we have reprocessed using `drizzle`.

7. Additional Information

Drizzle was originally developed for the Hubble Deep Field, a project whose purpose was to image an otherwise unexceptional region of the sky to depths far beyond those of previous

astronomical images. The resulting images are available in the published astronomical literature(Williams et al. 1996) as well as from the Space Telescope Science Institute via the World Wide Web at `http://www.stsci.edu/ftp/science/hdf/hdf.html`. This web page provides access both to the final drizzled image, and to the stacks of images at each of the 11 dither positions. These images are excellent material for those who would like to try out the tasks in the `dither` package.

The `dither` package is scheduled to be released shortly by STScI as an IFAF STSDAS package. However, additional, unofficial tasks may become available from time-to-time, before they are officially released. Interested IRAF users should occasionally examine `http://www.stsci.edu/~fruchter/dither` for more information.

References

Cowie, L. L., Hu, E. M., & Songaila, A., 1995, AJ, 110, 1576

Fruchter, A., & Hook, R. N., 1997, *A novel image reconstruction method applied to deep Hubble Space Telescope Images*, ed. A. Tescher, Vol. 3164 (S.P.I.E.), in press

Gull, S. F., & Daniell, G. J., 1978, Nature, 272, 686

Lucy, L. B., 1974, AJ, 79, 745

Lucy, L. B., & Hook, R. N., 1991, in *Astronomical Data Analysis Software and Systems I*, A.S.P. Conf. Ser., Vol. 25, eds. D. M. Worrall, C. Biemesderfer, & J. Barnes (San Francisco: ASP), 277

Richardson, B. H., 1972, J. Opt. Soc. Am., 62, 55

Weir, N., & Djorgovski, S., 1990, in *The Restoration of HST Images and Spectra*, eds. R. L. White & R. J. Allen (Baltimore: STScI), 31

Whitmore, B., & Heyer, I., 1997, Instrument Science Report WFPC2 97-08

Williams, R. E., et al., 1996, AJ, 112, 1335