

Calstis1: Basic Two-Dimensional Image Reduction

Phil Hodge, Stefi Baum, and Paul Goudfrooij

and the Spectrographs Group pipeline block: Ivo Busko (SSG), J.C. Hsu (SSG), Steve Hulbert, Rocio Katsanis, Melissa McGrath, and Dick Shaw

October 16, 1998

ABSTRACT

This ISR describes calstis1, the two dimensional image reduction portion of calstis, which outputs a flat fielded image. This ISR contains a description of the need for and method of converting between image and detector pixels, the algorithmic steps in calstis1 and the reference files used by calstis1. This document supercedes STIS ISR 95-07.

1. Introduction

In this ISR we describe the steps in *calstis1*, the two dimensional image reduction portion of *calstis*. The output of *calstis1* is a flat-fielded image, in which the photometric header keywords have been populated for imaging data. The flow of data through the *calstis* pipeline and the output products from it are described more generally in STIS ISR 95-06.

The remainder of this ISR is described as follows:

In Section 2, “Image to Detector Pixel Mapping” on page 2 we describe the work *calstis1* has to do in mapping from image to detector pixels, in order to perform two-dimensional image reduction.

In Section 3, “Algorithmic Steps in *calstis1*” on page 5 we describe each calibration step in the 2-D image reduction.

In Section 4, “Integrating Timetag Data” on page 13 we describe how to integrate TIMETAG data to an ACCUM image.

In Section 5, “Calibration Reference Files and Tables Used by *calstis1*” on page 14, we describe the calibration reference files and tables used in 2-D image reduction, briefly describing the structure of the files and giving a brief overview of how they are expected to be produced.

In Section 6, “Propagation of Errors” on page 19 we describe the propagation of errors.

In the Appendix we describe how to run *calstis1*.

2. Image to Detector Pixel Mapping

Calstis must map between image pixels and detector pixels in order to perform two-dimensional image reduction. There are two important components to this for STIS data:

- formats other than full frame (use of subarrays and on-board pixel binning) and
- Doppler compensation.

Each of these is discussed in turn below.

Subarrays and Binning

An image may have been taken as a “subarray” which covers only a subset of the detector. In addition, each image pixel may correspond to a bin of more than one pixel on the detector. The subarray size and offset and the binning must be taken into consideration when applying a reference file during image reduction. The subarray offset and binning are given by keywords LTV1, LTV2, LTM1_1, and LTM2_2 in the extension headers. See the Appendix to STIS ISR 98-10 for detailed information about these keywords for STIS data.

The binning may be two pixels (‘lowres’) in either or both directions for the MAMAs. For the CCD, the pipeline will only calibrate data taken with binning of 1, 2, or 4 in each axis and for which the subarray includes the entire serial overscan region. CCD data taken with other subarrays can be calibrated, but the BLEVCORR step will use a default bias level from the CCDTAB.

The primary header keywords describing the image size, binning, and location are SIZAXIS1, SIZAXIS2, BINAXIS1, BINAXIS2, CENTERA1, and CENTERA2. These keywords are derived from the proposal, and they are not very useful for determining the actual size and location of the image. First of all, they remain fixed, even as the image size or binning is changed (e.g. by BLEVCORR or LORSCORR). Secondly, the values of SIZAXIS1 and SIZAXIS2 are not, in general, the actual image axis lengths. The actual image size is given by NAXIS1 and NAXIS2 in the SCI extension header. The LTV and

LTM keywords give the mapping from detector coordinates Xref, Yref to image pixels X, Y as follows:

$$X = X_{\text{ref}} * LTM1_1 + LTV1$$

$$Y = Y_{\text{ref}} * LTM2_2 + LTV2$$

The IRAF convention (which is used by *calstis1*) is that the center of a pixel has integer pixel number. Pixel number *i* therefore runs from *i*-0.5 to *i*+0.5. By mapping the corners of a pixel back to the detector, we find that pixel number [*i*,*j*] in the science image corresponds to the following ranges of pixels on the detector:

$$\text{in X: } (i - 0.5 - LTV1) / LTM1_1 + 0.5 \text{ to } (i + 0.5 - LTV1) / LTM1_1 - 0.5$$

$$\text{in Y: } (j - 0.5 - LTV2) / LTM2_2 + 0.5 \text{ to } (j + 0.5 - LTV2) / LTM2_2 - 0.5$$

Subarray offset and bin size must be considered when interpreting the values in reference files. For initializing the data quality array, the value to assign to a pixel is the logical OR of all the values in the data quality initialization file that overlap the given pixel. When applying bias, dark, and flat field reference files, a logical OR is again used to combine the data quality array value in the file being calibrated with all corresponding pixels in the data quality array for the reference image. The limits to use when checking for MAMA local non-linearity will be affected by bin size. When combining CRSPLIT CCD images, the bin size affects the noise model (since the CCD binning is on-chip and so read noise and dark current will be different for unbinned and binned data). On the other hand, the PHOTFLAM keyword value and correction of extracted spectra to absolute flux units are not affected by binning, since the counts within a bin were summed rather than averaged. Binning of the reference files for MAMA low-res data will not be necessary if the reference files are themselves low-res; however, for those modes where a Doppler correction is applied on board (medium resolution and echelle spectroscopy with the MAMAs), the full high-res (2048 x 2048) reference files will need to be used, in order to correct for the Doppler smoothing, which is done in high-res pixels.

Doppler correction

For MAMA ACCUMULATION (“ACCUM”) mode observations using the medium or high resolution echelles, the flight software corrects the location of each photon event on-board for the Doppler shift. In TIMETAG mode this correction is *not* applied on-board, but it is done by Generic Conversion when making an ACCUM image from the TIME-TAG table. The Doppler correction must therefore be computed in the following situations:

- When processing TIMETAG data taken with the medium- or high-resolution echelles, the timetag stream must be corrected for Doppler shifts in the pipeline. Generic Conversion does this and writes the Doppler-corrected X coordinates in a new column `AXIS1` in the TIMETAG table (`_tag.fits`). Then when the TIMETAG data are integrated in time (in Generic Conversion or by using **inttag**) to produce an uncalibrated accumulated science image (`_raw.fits`), it is the Doppler-corrected positions that are used.
- When processing ACCUMULATION mode data for which the Doppler correction was applied by the flight software. In this case the flat field and dark reference files should be convolved with the Doppler smearing function, since the counts in a single image pixel were actually detected at different (Doppler-shifted) detector pixel locations. If there is a data quality initialization file, bad pixels should be extended by the Doppler smearing function as well, assigning to each image pixel all bad pixel values from detector pixels which contributed to it.

The Doppler shift in high-res pixels at time T (MJD) will be computed by

$$\text{shift} = - \text{DOPPMAG} * \sin ((2 * \pi / \text{ORBITPER}) * (T - \text{DOPPZERO}))$$

where DOPPMAG is the Doppler shift amplitude in high-res pixels, and DOPPZERO is the time (MJD) when the Doppler shift was zero and was increasing (i.e. near when HST was closest to the target). ORBITPER is the orbital period of HST in seconds, but it may be an assumed value (the one used on-board) rather than a true value. DOPPMAG, DOPPZERO and ORBITPER are all SCI extension header keywords populated by Generic Conversion. If a photon actually hit detector location (X,Y), then pixel (X+shift,Y) in the image would have been incremented; so a positive shift means that the reference files (e.g. dark, flat) should be shifted to the right before being applied.

Allowing Users to Convert from Image Pixels to Detector Pixels

The use of standard IRAF keywords for tracking subimages (LTV1, LTV2, LTM1 and LTM2) allows users to easily convert from image to detector coordinates when examining their data in IRAF/STSDAS. A number of IRAF tasks (e.g. **rimcursor**, **listpix**) that display coordinate information allow the user to specify which coordinate system they want. The choice is usually between “logical”, “physical”, and “world”. World coordinates would be wavelength and distance along the slit, for example, or right ascension and declination. Logical coordinates are image pixel coordinates. Physical coordinates are also pixels, but they are in the “original” image in case the image at hand is a section of another image. The LTV (vector) and LTM (matrix) keywords give the linear transformation between logical and physical coordinates, while the FITS coordinate parameters CRVAL, CRPIX, etc., give the transformation between logical and world coordinates. In the case of STIS, the user’s image may be only a portion of the full detector, and the pixels may also have been binned, so we adopt the coordinates of the full detector (without overscan, for the CCD) as the physical system. If an IRAF task such as **imcopy** is used to further extract

a subset (or to subsample, or to flip an axis), the LTV and LTM keywords are updated by that task so they still transform back to the original detector coordinates. Thus when a user displays an image and reads coordinates interactively using **rimcursor**, for example, the user can choose to see either image pixel coordinates, original detector coordinates, or wavelengths, simply by setting the value of the **rimcursor** wcs parameter to “logical” “physical”, or “world” respectively.

3. Algorithmic Steps in *calstis1*

In Table 1 and Table 2 we summarize the steps in *calstis1* for CCD and MAMA data, respectively. The processing steps (column two) are listed in order of occurrence within *calstis1*.

Table 1. Two-dimensional Image Reduction (*calstis1*) for the CCD

Switch	Processing Step	Reference File
	initialize error array	CCDTAB
DQICORR	initialize data quality	BPIXTAB, CCD-TAB
ATODCORR	correct for A to D conversion errors	ATODTAB
BLEVCORR	subtract bias level computed from overscan	
BIASCORR	subtract bias image	BIASFILE
DARKCORR	subtract dark image	DARKFILE, CCD-TAB
FLATCORR	divide by flat field image	PFLTFILE, DFLTFILE, LFLTFILE
SHADCORR	apply shutter shading correction	SHADFILE
PHOTCORR	populate photometric header keywords	PHOTTAB, APERTAB

Table 2. Two-Dimensional Image Reduction (*calstis1*) for the MAMAs

Switch	Processing Step	Reference File
	initialize error array	
DOPPCORR	apply Doppler convolution to reference files	
LORSCORR	bin data to LTM1_1 = LTM2_2 = 1 before processing	

Switch	Processing Step	Reference File
DQICORR	initialize data quality	BPIXTAB
GLINCORR	correct for global non-linearities	MLINTAB
LFLGCORR	flag local and global non-linearities	MLINTAB
DARKCORR	subtract dark image	DARKFILE
FLATCORR	divide by flat field image	PFLTFILE, DFLTFILE, LFLTFILE
PHOTCORR	populate photometric header keywords	PHOTTAB, APERTAB

Here are the algorithmic details of each calibrations step. The steps are presented in alphabetical order.

ATODCORR

The analog-to-digital correction will need to be applied if the CCD electronic circuitry which perform the analog-to-digital conversion is biased toward the assignment of certain DN (data number) values. The ATODCORR is a statistical correction which reassigns DN numbers based on a lookup in a reference file, the ATODTAB. To apply the correction, the raw data value is used as an index into the array read from the table. That is, if a raw data value is i , then the value in the reference table array at element i (zero indexed) is the corrected value. Currently, ATODCORR is set to OMIT, and no reference table has been prepared. There is evidence of a small bias, but the effect may be too small to be significant.

BIASCORR

The BIASCORR step removes any two-dimensional additive stationary pattern in the electronic zeropoint of each CCD readout. To remove this pattern a bias image is subtracted. A bias image is the pattern left when an exposure is taken with zero exposure time, following the removal of the overscan bias level (see BLEVCORR, below).

The bias reference file (BIASFILE) is a full-format image. Different bias files are available for different values of CCD gain and binning. If the science image is a subarray, a section of the bias image must be extracted to match the science image. It may be that two or more images were added together (e.g. for cosmic ray rejection) to form the input science image. The number of such images is given by the extension header keyword NCOMBINE, and the bias image will be multiplied by this value prior to being subtracted. The bias image will have accompanying error and data quality arrays, which are included with the science data during bias subtraction.

BLEVCORR

The BLEVCORR step removes the electronic bias level for each line of the CCD image. The electronic bias level may vary with time and temperature; its value must be determined in the particular exposure being processed from the overscan regions. Virtual overscan is created by continuing the readout of each line of the CCD past its real physical extent, while physical overscan is created by unilluminated pixels in the detector. For the STIS CCD in full-frame unbinned mode, there are 20 rows of virtual overscan in the parallel direction, and 19 leading and trailing columns of physical overscan in the serial direction. Physical overscan pixels can be affected by cosmic rays and dark current.

When on-chip pixel binning is used, since the binning factor cannot divide evenly into 19 and 1024, the raw image produced will contain both pure overscan pixels, mixed overscan plus science pixels, and science pixels. A figure illustrating this can be found in the Appendix to STIS ISR 98-10. The pipeline will only calibrate pixel binnings of 1, 2, or 4 in either X or Y. The bias level will then be determined from a portion of the trailing overscan columns of each line of the science image (see Table 3). For each line of the image, the median of the values in the overscan regions will be taken, in order to reject discrepant pixels. This gives a value of overscan for each image line, but the uncertainty in each value is large because it is based on relatively few pixels. A straight line is therefore fit to these values as a function of image line number, and the fitted line is evaluated at each image line number. This value will be subtracted from the science data for that line. Actually, there's one more detail. There's a slope within each image line, amounting to of order one DN over the entire line. The parallel overscan region is used to determine this slope, and the value subtracted from each pixel in the line is adjusted by a small amount depending on this slope. Effectively, then, a two dimensional bias level image is created by performing linear fits to the parallel and serial overscan regions, and this image is subtracted from the science data.

Table 3. Region in Overscan Used to Determine Bias Level

Binning	Columns in Raw Image
unbinned	2 through 16
subarray	1 through 14
1, 2, or 4	2 through 8

The output science image from this step will be smaller than the input image by the width of the overscan regions. Both the serial and parallel overscan regions will be trimmed away in this step. Table 4 shows the sizes of the raw and calibrated (trimmed) images. This table only shows explicitly the cases of equal binning in the X and Y axes. For unequal binning, the values can be taken from two rows. For example, for binning of 2

x 4, the size of the raw file would be 532 x 266 (the X size is from the 2 x 2 row, and the Y size is from the 4 x 4 row). This explains why there is a 1 x 1 row; while “binning” of 1 x 1 is not allowed, it is permitted to use binning of, say, 1 x 2 or 2 x 1.

Table 4. Sizes of Raw and Calibrated CCD Images

Binning	Size of Raw Image	Size of Calibrated Image
unbinned	1062 x 1044	1024 x 1024
subarray	1060 x N	1024 x N
1 x 1	1054 x 1034	1024 x 1024
2 x 2	532 x 522	511 x 512
4 x 4	271 x 266	255 x 256

DARKCORR

The DARKCORR step removes the dark signal (count rate created in the detector in the absence of photons from the sky) from the science image by subtracting an integration time- and gain-scaled dark reference file image (DARKFILE). If the science image is a subarray or was binned, a section of the dark image must be extracted and binned (summed) to match the science image. If Doppler correction was applied on-board for the science data (i.e., if DOPPCORR is set to PERFORM), the Doppler smearing function should be computed and convolved with the dark image to account for the contributions of various detector pixels to a particular image pixel. The Doppler convolution must be done before binning the dark image. To apply the correction, the dark image is multiplied by the exposure time, divided by the gain (if CCD), and subtracted from the science image. (The exposure time is taken from the SCI extension header keyword EXPTIME, as opposed to the total exposure time given by the primary header keyword TEXPTIME.) The error and data quality arrays in the calibrated file are updated to include error and data quality information from the dark reference file. The average of the dark values that were subtracted is written to the output SCI extension header with the keyword MEANDARK; the average is restricted to those pixels in the darkfile that were not flagged as bad in the DQ extension in the darkfile.

For the CCD detector, it takes approximately half a minute to read out a full-frame image. Dark counts accumulate during this time, but this is *not* corrected for in the DARKCORR step. Instead, the dark time beyond the exposure time is accounted for in the BLEVCORR and BIASCORR steps.

DOPPCORR

This is not an independent calibration step. Instead, it's a modifier for DQICORR, DARKCORR, and FLATCORR, meaning that the reference file should be convolved with the Doppler shift integrated throughout the exposure, in order to compensate for the on-board correction for Doppler shift. This is only relevant for MAMA medium or high resolution spectroscopic data. DOPPCORR will be silently set to OMIT if DOPPMAG is zero. Since DOPPCORR is a modifier for more than one calibration step, DOPPCORR is not reset to COMPLETE in the output file, because *calstis1* could be run multiple times to perform those individual steps.

DQICORR

The DQICORR step takes the data quality file output from Generic Conversion for the science data and bitwise ORs it with the values in the bad pixel reference file table (BPIXTAB) to initialize the science data quality array for propagation through subsequent steps in *calstis1*. (The data quality array is a bit mask.) If DOPPCORR is PERFORM, *calstis1* will Doppler smear the BPIXTAB prior to performing the OR operation with the (unsmearred) science input data quality image. See Table 6 on page 17 for a description of the data quality flags. Since new data quality values are combined with existing ones by ORing them, DQICORR can be repeated if *calstis1* is rerun. If the same bad pixel table is used, DQICORR will have no additional effect, while using different bad pixel tables results in combining all the flags.

For CCD data, this step also includes a check on saturation, comparing the science data values with the saturation level read from the CCD parameters table (CCDTAB).

Another part of DQICORR that is only performed for CCD data is to flag regions that are beyond the aperture, when an aperture between six and 28 arcseconds in height was used for the observation. The flag value is 4. This was found to be necessary to prevent problems with cosmic ray rejection. When a large fraction of the image is actually not illuminated, the computation of sky level can be seriously affected. The aperture size is read from the APER_FOV keyword, which is a string of the form "28X50", giving the height by the width in arcseconds, and that height is then converted to image pixels using the CD matrix. The aperture is assumed to be centered on the detector. To allow for errors in centering and aperture size and to avoid flagging illuminated pixels, the computed aperture size will be expanded by 10 pixels in each direction. For imaging type data, pixels beyond the aperture in both X and Y directions can be flagged. For spectroscopic data, however, only pixels beyond the aperture in the cross dispersion (Y) direction will be flagged. The flagging does not extend into the overscan regions, although mixed illuminated/overscan pixels can be flagged.

Error Array Initialization

If the error array for the input data is entirely zero, then values will be computed based on a simple noise model and assigned to the output error array. For MAMA data the error is simply the square root of the uncalibrated science array value I (with a minimum value of zero). For CCD data the noise model is:

$$\sigma \approx \sqrt{\frac{(I - \text{bias})}{\text{gain}} + \left(\frac{\text{readnoise}}{\text{gain}}\right)^2}$$

where bias is the bias level in DN, gain is electrons/DN, and readnoise is the readout noise in electrons, as read from the CCD parameters table (CCDTAB).

FLATCORR

The FLATCORR step corrects for pixel-to-pixel and large-scale sensitivity gradients across the detector by dividing the data by a flat field image. The flat field image used to correct the data is created from up to three flat field reference files¹:

- PFLTFILE - this flat is a configuration (grating, central wavelength and detector) dependent ‘pixel-to-pixel’ flat field image, from which any largescale sensitivity variations have been removed (i.e., it will have a local mean value of unity across its entirety). Such configuration dependent flats are expected to be produced only once per year.
- DFLTFILE - this flat is a ‘delta flat’ which gives the changes in the small scale flat field response relative to the pixel to pixel flat (PFLTFILE). Delta flats are currently not used, as they are not needed; the idea was that there would be a single delta flat for each detector, and they would be taken as often as necessary, perhaps monthly.
- LFLTFILE - this flat is a subsampled image containing the largescale sensitivity variation across the detector.

To create the single combined flat field file, *calstis* will use bilinear interpolation to expand the largescale sensitivity flat (LFLTFILE) to the subarray format of the science data and multiply it by the appropriate (subarray) portion of the pixel-to-pixel flat (PFLTFILE) and delta flat (DFLTFILE). After the pixel-to-pixel flat, delta flat, and low order flats have been multiplied together, the Doppler convolution will be applied to the product, if appropriate, and the product binned (averaged), if necessary, to match the binning of the science image. The science data are then divided by the combined flat field. The error and data quality arrays in the calibrated file will be updated to reflect the error and data quality from

1. The rationale for maintaining 3 types of flat field reference files rather than a single integrated reference file is described in detail in STIS ISR 96-015, “On-Orbit Flat Fields and Absolute Calibration of STIS”, Bohlin et al., Jan 1996

the flat field. If Doppler convolution is applied, a correction is also applied for the loss of counts at the image edges due to the flight software’s “effective subarray”; see STIS ISR 98-05 for further information.

GLINCORR and LFLGCORR

The MAMAs are photon counting detectors. At high photon (pulse) rates, the MAMA response becomes nonlinear due to three effects.

1. Pore paralysis in the Micro Channel Plates arises when charge cannot flow rapidly enough to replenish channels that have been depleted of electrons due to high local photon rates. This produces a *local* non-linearity. The local count rate is roughly linear up to counts rates of ~100 counts/second/pixel, and it turns over at high count rates, showing an inverted V shape. Thus it is not possible to reliably correct for or flag in the pipeline pixels which have exceeded the local linearity limit (since the relation is bi-valued).
2. The electronic processing circuitry has a “dead-time” of roughly 350 nanoseconds between pulses; thus at global count rates (across the detector) of 300,000 counts (pulses) per second, the electronic circuitry counts roughly 90% of the pulses. This electronically induced *global non-linearity* can be corrected for in the pipeline.
3. The MIE electronics/flight software is specified to process 300,000 pulses per second (i.e., is matched to the expected global count rate performance of the electronic circuitry). At count rates higher than this, the MIE will still count only 300,000 pulses per second - thus this represents a hard cutoff beyond which no information is available to allow correction to true count rate. In addition, because of an unexpected feature of the MIE flight software at global count rates above 285,000 counts/sec an additional software-induced non-linearity sets in.

When LFLGCORR = PERFORM, the pipeline will flag pixels in the accompanying data quality array when they either themselves have count rates greater than the local linearity limit or when they are within 4 pixels of a pixel which has a count rate exceeding that limit. The local linearity limit is read from the LOCAL_LIMIT column in the MLINTAB. When the global count rate (across the detector, from science header keyword GLOBRATE) is greater than the global linearity limit all pixels in the accompanying data quality file will be flagged, and the primary header keyword GLOBLIM will be set equal to “EXCEEDED”. The global linearity limit is read from the GLOBAL_LIMIT column in the MLINTAB.

When GLINCORR=PERFORM, the pipeline will correct all the pixels in the input image for global non-linearity using expressions developed and tested on the ground. The expression below seems to fit the existing test data well (Ebbets 94, STIS SER CAL-005)

$$y = xe^{(-tx)}$$

where

- x = true input count rate (counts/s)
- y = observed count rate (counts/s)
- t = time constant ~ 290 nanoseconds.

For subarrays, the hard cutoff limit of the MIE electronics/software will differ from that for full frame processing. The first step in the MIE processing is to determine the detector pixel origin of the incoming pulse; if it is outside of the specified subarray it will be thrown out, if it falls within the subarray it will be processed further. The performance of the sorting for subarray data, which is dependent on the global count rate across the full detector, slows down the MIE (i.e., lowers the cutoff); however, the fact that the next processing step is performed on only those pulses within the subarray speeds up the MIE (i.e., raises the limit). Ground tests will have to be performed to determine the relevant time constants for the two parts (sorting and processing) in order to allow correction for global non-linearity when subarrays are used. The global count rate (across the entire detector) is determined as part of the bright object protection sequence and is passed down with the exposure as the SCI extension header keyword GLOBRATE.

LORSCORR

MAMA data are taken in high-resolution mode (2048 x 2048 pixels), but they may be binned on-board to 1024 pixels in either or both axes. If they are not binned on-board, the binning may be done by *calstis*. This is a simple step, just adding counts in pairs of pixels. If binning was not done on-board, BINAXIS1 and BINAXIS2 in the primary header will be one, and LTM1_1 and LTM2_2 in the extension headers will be two. After binning is complete, the LTM keywords will be reset to one (indicating low-res data), but the BINAXIS keywords will remain unchanged, because they reflect what was specified in the proposal.

PHOTCORR

For imaging type data, the total system throughput will be read in from a reference table (PHOTTAB), and the photometric keywords PHOTFLAM, PHOTBW, and PHOTPLAM will be computed and saved in the primary header of the output image. No change is made to the data numbers themselves. Thus, this step can be repeated if *calstis1* is rerun on its own output. For spectroscopic type observations, this step is not performed.

There is a related step that is actually performed regardless of the PHOTCORR keyword. Names describing the instrument configuration (e.g. detector, optical element, etc.) are concatenated together and written to the output primary header with keyword PHOTMODE. This string is what one would use for obsmode when using **synphot** tasks.

SHADCORR

This step corrects for shading by the shutter in very short integration time exposures. The STIS CCD shutter is specified to produce exposure non-uniformity less than or equal to 5 milliseconds for any integration time: the shortest possible STIS CCD exposure time is 100 milliseconds. At the present time this step is not performed, as it appears to be unnecessary.

The shutter shading correction reference file is a full-frame image containing the extra time that each pixel is exposed, compared to the nominal exposure time. That is, if the nominal exposure time is EXPTIME, pixel [x,y] was exposed for a time EXPTIME + shadfile[x,y], where shadfile is the shutter shading correction image. The algorithm for correcting the data is:

$$\begin{aligned} \text{corrected}[x,y] &= \text{raw}[x,y] * \text{EXPTIME} / (\text{EXPTIME} + \text{shadcorr}[x,y]) \\ &= \text{raw}[x,y] / (1 + \text{shadcorr}[x,y] / \text{EXPTIME}) \end{aligned}$$

STATFLAG

The STATFLAG keyword has a boolean value, T or F, rather than a string “PERFORM”, “OMIT”, etc. If STATFLAG = T, simple statistical values will be computed for the calibrated data and saved in the output SCI and ERR extension headers. STATFLAG will not be reset to F in the output file, as it would be very appropriate to rerun this step if *calstis1* is rerun on its own output. Values are computed for pixels which have non-negative error array value and which are not flagged as bad in the data quality array. Being flagged as bad means that the bitwise logical AND of the DQ extension value with the SDQFLAGS keyword value is non-zero. The statistics that are computed for both the SCI and ERR arrays are the number of good pixels NGOODPIX, the minimum good pixel value GOODMIN, the maximum good pixel value GOODMAX, and the mean of good pixels GOODMEAN. In addition, the ratio of SCI to ERR value (if the latter is non-zero) is used to compute minimum signal-to-noise (S/N) ratio SNRMIN, the maximum S/N ratio SNRMAX, and the mean S/N ratio SNRMEAN; these keywords are saved in the SCI extension header.

4. Integrating Timetag Data

Raw time-tagged data will be stored in one or more FITS binary tables which contain at least three columns, the X and Y pixel coordinates and the time of each photon event, with column names AXIS1, AXIS2, and TIME respectively. The AXIS1 coordinates will already have been corrected for the Doppler shift, if appropriate. There may be an associated table containing “good time intervals.” In order to process time-tagged data through *calstis1* as a pseudo ACCUM mode image, the events should be integrated over time and

within pixels. Generic Conversion does this integration, producing an `_raw.fits` file, but it can also be done off-line using the **inttag** IRAF task. See also the **odelaytime** task, for correcting the times from spacecraft to heliocentric. For each event, the time will be compared with the set of good time intervals, and if the event falls within a good interval, then the X and Y coordinates will be scaled to the output pixel location, and the output pixel closest to that location will be incremented by one.

5. Calibration Reference Files and Tables Used by *calstis1*

Table 5 lists all the calibration reference files used by *calstis1* and provides a brief description of them. The calibration reference images will be stored in FITS image extension files using the same conventions as for science data, as triplets of science, error, and data quality images (in extensions 1, 2, and 3, respectively; the primary data unit will be left empty). The calibration reference tables will be stored in FITS binary tables (see STIS ISR 95-06). For the CCD, reference images will be 1024 x 1024 format, except for binned bias images and the low-order flat. For the MAMA, for the echelle and medium-resolution modes, reference files must be maintained as 2048 x 2048 images (since the on-board Doppler correction is performed on high-res pixels). However, for the low-resolution gratings, we may choose to maintain both 1024 x 1024 and 2048 x 2048 format reference images so as not to burden low-res users with enormous reference files. A brief description of each reference file is provided below.

There are some cases where a reference file is (intentionally) not used, although it could have been, for calibrating a particular file. For example, for flat fielding, any one, two, or all three of PFLTFILE, DFLTFILE, and LFLTFILE may be used. The string “N/A” for a reference file name means that the file is not available and need not be applied.

Table 5. Reference Files for *calstis1*

reference file	format	calibration step	selection parameters
apertab	binary table	PHOTCORR	configuration ^a
atodtab	binary table	ATODCORR	detector (CCD), gain, amplifier ^b
biasfile	image	BIASCORR	detector (CCD), amplifier
bpixtab	binary table	DQICORR	detector
ccdtab	binary table	DARKCORR DQICORR, PHOTCORR	detector (CCD), amplifier
darkfile	image	DARKCORR	detector
dfltfile	image	FLATCORR	detector,
lftfile	image	FLATCORR	configuration
mlintab	binary table	GLINCORR LFLGCORR	detector (MAMA)
pfltfile	image	FLATCORR	configuration
phottab	binary table	PHOTCORR	configuration
shadfile	image	SHADCORR	detector (CCD)

- a. A STIS configuration is defined as a filter, optical element, central wavelength, detector combination
- b. Only a single amplifier will be used during any period - however, it is possible that over time the amp in use will change.

APERTAB

The aperture throughput table is used for imaging type data in order to include the filter throughput when computing photometry keywords for the PHOTCORR step.

ATODTAB

The ATODTAB is a binary table with one or more rows, each containing an n-element array, where n is the number of possible values of DN in the input image. Each row of the table gives the correction for a different value of some parameter, such as temperature, given by a header keyword in the input image. The row is selected that has the value of this parameter closest to the value in the image header, and the analog-to-digital correction array is read from that row. The value of element i (zero indexed) in this array gives the corrected value (value to be assigned in output image) corresponding to input DN values of i.

BIASFILE

The BIASFILE is a two-dimensional image. The process of creating bias reference files will be described in a separate ISR. There is a separate bias reference file for each supported binning setting, in contrast with other reference files which are binned by *calstis* to match the science data. For bias files that are binned in the first axis direction, the length of the first axis is smaller than 1024/BINAXIS1 because the mixed illuminated/overscan pixels are excluded.

BPIXTAB

The bad pixel table is a tabular listing of detector pixels (x and y detector coordinates) with known conditions (non-zero values) and their data quality conditions. It is used, with the input data quality array if one exists, in the data quality initialization step, to determine the data quality flags appropriate for the uncalibrated data. The possible error conditions which the data quality conditions flag and their values are given in Table 6 below; they are also described in the HST Data Handbook, Volume 1, October 1997, table 20.8. The data quality values are set as specific bits in a 16 bit word; thus multiple conditions can be flagged for a given pixel via the bitwise OR operation.

Table 6. Data Quality Flags

value	macro name	meaning
0	GOODPIXEL	OK
1	SOFTERR	Reed-Solomon decoding error
2	DATALOST	data replaced by fill value
4	DETECTORPROB	bad detector pixel or beyond aperture
8	DATAMASKED	masked by occulting bar
16	HOTPIX	hot pixel
32	LARGEBLEM	large blemish
64		not currently assigned
128	OVERSCAN	can be used for finding bias level
256	SATPIXEL	saturated pixel
512	CALIBDEFECT	bad pixel in reference file
1024	SMALLBLEM	small blemish
2048	X1D_BAD_BACKGROUND	bad pixels in background region
4096	X1D_DISCARDED	pixel discarded from extraction region
8192	DATAREJECT	rejected during image combination
16384		not currently assigned

CCDTAB

The CCD parameters table gives the gain (electrons per DN), bias (DN), readnoise (electrons), and saturation limit (DN). Different rows of the table give different values for the parameters, and the appropriate row is selected on amplifier, nominal gain, CCD offset, and binning.

DARKFILE

The dark file is a two-dimensional image of the dark counts scaled to a one-second integration.

DLFTFILE

The delta flat file is a two-dimensional image with a value indicating the change in the flat field response at that pixel relative to the response in the pixel-to-pixel flat field file (PFLTFILE). We don't currently provide delta flats, as they do not appear to be needed.

LFLTFILE

The low order flat field file is a two-dimensional image of the largescale (low order) sensitivity variations across the detector. Not all modes will require low order flats. In principal, they could be produced by median filtering (in a boxcar region) the values in a lamp-illuminated flat field from which the lamp illumination pattern has been removed. In practice, to date they have been constructed from dithered external source observations. Since only the low-order response is being determined, the calibration exposures for the low-order flats need not be long (i.e., can be of relatively low signal-to-noise). *Calstis* uses bilinear interpolation to expand the image to the pixel format of the science image in order to correct for the low-order flat field response.

MLINTAB

The MAMA linearity table contains the count rates at which to flag the data quality array for local and global nonlinearities, and it contains parameters needed to correct for global non-linearity.

PFLTFILE

The pixel-to-pixel flat field is a configuration-dependent pixel-to-pixel flat field from which the low-order sensitivity variations have been removed. It is intended that pixel-to-pixel flats have a signal to noise of at least 100:1 per resolution element (2 x 2 low-res pixels) for the MAMAs, and appreciably higher for the CCD.

PHOTTAB

Each row of the photometry table contains an array of wavelengths and an array of associated total throughputs (quantum efficiency), with a different row for each configuration. The throughput values actually do not include the aperture throughput, which in this case means filter throughput, since the filters are in the aperture wheel, and the phottab is only used for imaging mode in *calstis1*. The filter throughputs are taken from the APERTAB.

SHADFILE

The shutter shading correction reference file is a full frame image containing the extra time that each pixel is exposed, compared to the nominal exposure time. That is, if the nominal exposure time is EXPTIME, pixel [x,y] was exposed for a time EXPTIME + SHADFILE[x,y], where SHADFILE is the shutter shading correction image.

6. Propagation of Errors

The calibration steps in *calstis1* are simple subtraction, multiplication, or division. The values in the associated error (ERR) array will be updated to reflect the combination of data and errors from the science and reference files. This section describes how this is done.

When the data are multiplied by a constant, the error array is multiplied by the same constant. Adding a constant does not affect the error. When operating on two arrays, let *a* and *b* be the values of the two arrays at a pixel, and let *da* and *db* be their respective errors. Then the error in the sum or difference is given by:

$$\sqrt{da^2 + db^2}$$

The error in the product is given by:

$$\sqrt{(a \times db)^2 + (b \times da)^2}$$

The error in *a* / *b* is given by:

$$\sqrt{\left(\left(\frac{da}{b}\right)^2 + \left(a \times \frac{db}{b^2}\right)^2\right)}$$

7. Appendix

This section gives some details regarding how *calstis1* can be executed. When running the pipeline task *calstis0*, either in the pipeline or off-line, *calstis1* will be invoked as a subroutine if the appropriate header switches are set to PERFORM. *Calstis1* can be run by itself either as a host-level executable **cs1.e** or as the IRAF task **basic2d**. (The executable is located in the *stdas* bin directory.) As with *calstis0*, the names of reference files are taken from the primary header of the input file. In contrast to *calstis0*, however, the steps to be performed are specified by the user as command-line arguments (or *cl* task parameters). The command-line arguments that are recognized can be found by running **cs1.e** without any arguments, in which case the following message will be printed:

```
syntax:  cs1.e [-t] [-v] input output [outblev]
```

```
command-line switches:
```

```
-dqi  -atod -blev
```

```
-dopp -lors -glin -lflg
```

```
-bias -dark -flat -shad -phot -stat
```

The `input` and `output` arguments are the names of the input and output files. If a wildcard is used it should be enclosed in quotes (e.g. `'*_raw.fits'`) to prevent Unix from expanding the list to multiple names. That is, `calstis1` expects to take the input name or names from a single command-line argument, and similarly for the output name and optional `outblev` (described below). The output name is actually optional as well; if no name is given, the output name will be constructed from the input by replacing the suffix (see below). The `outblev` argument, if given, is the name for a text file to contain the bias overscan level that was subtracted from each image line during the BLEVCORR processing step. Of course, this is relevant only if BLEVCORR is performed. If `output` and/or `outblev` is given, the number of names in the list must match the number in the `input` list. See STIS ISR 98-10 for further information and examples.

The command-line switches `-t` and `-v` serve the same function as in `calstis0`, to print timing information and verbose output, although in `calstis1` at the present time there is no additional information printed in verbose mode.

The remaining switches are for specifying which calibration steps to perform. If no switch is specified on the command line, all relevant steps will be performed. If one or more switches are specified, however, only those steps will be performed. All of the command-line calibration switches are named from the corresponding header keywords with the suffix “CORR” removed, except `-stat` which is from STATFLAG. The term “relevant steps” used above needs some clarification. Some steps (e.g. BIASCORR) are only done for the CCD detector, while others (e.g. LORSCORR) are only done for the MAMAs. So if an image to be calibrated was taken with one of the MAMA detectors, BIASCORR will not be performed even if it was specified. Furthermore, for a CCD image (for example), if BIASCORR is set to COMPLETE in the input header, that step will not be performed again even if it was specified. Most calibration steps in `calstis1` should not or cannot be repeated. DQICORR, PHOTCORR, DOPPCORR, and STATFLAG can be repeated in principle, however.

For CCD data, `calstis1` is normally run twice, first to subtract bias prior to cosmic ray rejection (`calstis2`), and then to complete the other steps such as dark subtraction and flat fielding. Here is an example to illustrate the command-line switches.

```
cs1.e o47s01020_raw.fits o47s_tmp.fits -dqi -blev -bias
```

```
cs2.e o47s_tmp.fits o47s01020_crj.fits
```

```
cs1.e o47s01020_crj.fits o47s01020flt.fits
```

Note that for the second execution of *calstis1*, no command-line switches were specified. This means that we want all relevant steps to be done, which no longer includes BLEVCORR and BIASCORR since they’ve already been done.

Calstis1 can also be run from IRAF. The **basic2d** task in the **stis** package is a cl script that constructs a Unix command line based on the task parameters. For most users this is the preferable way to execute the task, because the parameters are visible in **eparam**, and the user doesn’t have to remember the directory location of the executable.

It was mentioned above that the output name need not be specified, in which case it will be constructed from the input name. If the input root name ends in a recognized suffix, that suffix will be replaced by a matching suffix; otherwise, “_flt” will be appended to the input root name to make the output name. The list of input suffixes and their replacements for output is given in the following table.

Table 7. Input and output suffixes

Input Suffix	Matching Output Suffix
_raw	_flt
_blv_tmp	_flt
_crj_tmp	_crj
_wav	_fwv

ACKNOWLEDGEMENTS

We gratefully acknowledge many many useful discussions with Don Lindler, and the use of his excellent ‘Science Data Management and Analysis Software Requirements Document’, Version 1.0. We also benefited greatly from discussions with Ralph Bohlin, Mark Clampin, George Hartig, and Mary Beth Kaiser.