

rgsfilter

January 27, 2025

Abstract


This task filters the combined event list for RGS. The filtering is limited to time intervals and specific event flags, and is applied to the frames, events and exposure maps. This is not a generalized selection engine, but should be used for any filtering operations that impact on the exposure maps.

1 Instruments/Modes

Instrument	Mode
RGS	Spectroscopy
RGS	High Time Resolution

2 Use

pipeline processing	yes
interactive analysis	yes



3 Description

rgsfilter removes unwanted frames and events from the combined event list for RGS. The internal node-specific exposure maps are also constructed (or reconstructed) so that they comply with the filtering of the frames and events. This is important for tasks such as **rgsspectrum** and **rgsrmfgen**, which use these maps to perform exposure corrections. Lastly, a single exposure map for the entire focal plane is constructed by combining the node-specific maps. An event list may be filtered repeatedly by **rgsfilter**, and typically the user will start with an event list that has already been filtered for bad events and corrupted frames, and apply additional filtering on hand-selected time-intervals. **rgsfilter** is not a generalized selection engine: it filters on frame timestamp and certain event flags only. All other filtering, for which the exposure maps do not need to be rebuilt (or for which this is either not feasible or not well defined), is left to other tasks. Throughout this document a † marks items that do not apply to High Time Resolution (HTR) mode data.


The filtering of frames is specified for each CCD separately through OGIP standard tables of Good Time Intervals (GTI). The time-filtering that has been applied to the event list is described by its **STDGTION** tables. Strictly speaking, these tables may be more prescriptive than descriptive on the input, but are always descriptive on the output. To further restrict the GTI, a list of auxiliary GTI tables may be specified (parameter **auxgtitables**) and are applied as a filter to each of the **STDGTION** tables. These tables are also filtered to remove any frames that lack attitude correction or lie partially outside the GTI. The output **STDGTION** tables, whether altered or not, are applied as a filter to the **EVENTS**, **EXPOSUON** and **REJPIXnn** tables. The filtering of events is specified through event flags (parameter **rejflags**). The event-filtering that has been applied to the event list is described by the **REJFLAGS** attribute of the **EVENTS** table. Events with these flags are removed from the **EVENTS** table, but not from the **REJPIXnn** tables. When the exposure maps are constructed the **REJFLAGS** attribute determines which entries in the **REJPIXnn** tables are selected. Event-filtering is restricted to the enumerated subset of the **EVENTS** table **FLAG** bits that are recommended for rejecting events. Note that all uplinked bad pixels are always left unexposed in the exposure maps regardless of what flags are set in **REJFLAGS**.

An exposure map provides the effective amount of active, photon-collecting time as a function of discrete (binned) position on the surface of the chip. The map values are in units of seconds, and the coordinate system is aspect-corrected RGS angles (dispersion and cross-dispersion). After **rgsfilter** version 1.18, the units for the dispersion axis in the exposure maps coordinate can be in:

- Beta space
- Lambda space

being the orthogonal axis always cross-dispersion. The **spectrumbinning** should be used to switch between these two dispersion axis units. Note that this definition of an exposure map differs somewhat from the common practice in X-ray Astronomy. Each node-specific map contains the contribution of active time from one node only, and in general there will be some overlap between the maps at the node boundary because of aspect drift. The total active time in the region of overlap is the sum of both contributions. The aspect-correction of the angular coordinates requires source-dependent parameters which are added to the event list by **rgsangles**. The validity of the exposure maps with respect to other possible targets is limited by this implicit choice of nominal target. The construction of the exposure map would be unreasonably slow if every frame were treated independently of the others. Instead, the frames are grouped together into bins according to their aspect correction and the average aspect correction in each bin is used to lay down the exposure for all of its frames. The bin size is configurable (parameter **driftbinsize**), but the default is both very fast and very accurate.

In principle one would not expect to find negative values in an exposure map, but because of certain algorithmic difficulties negative values do occur. The largest of these (in absolute value) are due to




the nearly negligible curvature of the angular coordinate system, and only manifest along the top and bottom edges of the chip. This effect should not be larger than roughly 10^{-5} of the total exposure (characterized by the `DATAMAX` attribute). Much smaller, and more widespread, are the errors due to round-off: inevitably, spurious infinitesimal values, both positive and negative, are occasionally found in unexposed pixels. To suppress this disconcerting effect any pixel with absolute value less than 3×2^{-24} of the total exposure is reset to zero in the final map. Any negative values that remain most likely represent a systematic error.

3.1 Node-specific Exposure Maps

The active time at each angular bin is the sum over frames of the active time contributions from all CCD pixels, within the node-specific telemetry window, which overlap the surface of the bin. Note that the chip coordinate system is related to the angular coordinate system through a fairly complicated transformation, which changes from frame to frame as the satellite drifts. A normal pixel contributes active time to a bin in proportion to its overlap, which naturally varies by frame. The active time, per frame, of a normal pixel is the full frame integration time multiplied by the node-specific fractional exposure; bad pixels contribute no active time. A pixel is normal if it can be used to detect x-rays. The on-board Data Preprocessor (DPP) should usually know the locations of all the hot pixels, but when a pixel it considers normal contains too much energy the DPP is configured to suppress it from the telemetry, effectively rendering the pixel bad. Since the active time lost this way cannot be assigned directly to the responsible pixels, it is simply deducted, proportionately, from every pixel (not including uplinked bad pixels) in the node via the fractional exposure calculated by `rgsframes`. For HTR mode data there are no such lost pixels, and accordingly no fractional exposure correction.

The `REJFLAGS` attribute selects which pixels are to be considered bad, either in particular frames or in every frame. The `REJPIXnn` tables record, by frame, the position of every pixel that might qualify as bad and is not redundant with the `BADPIXnn` tables. Filtering on `ON_BADPIX` sets all pixels listed in the `BADPIXnn` table bad in every frame, and filtering on `NEXT_TO_BADPIX` sets all pixels directly adjacent to these bad in every frame as well. Filtering on `ON_WINDOW_BORDER` sets all pixels at the edge of the telemetered window bad in every frame (in HTR mode only the vertical edges are marked bad). Filtering on `ON_NODE_INTERFACE` sets all pixels on either side of the trans-node boundary bad in every frame. These pixels tend to collect spurious events because the DPP does not perform event reconstruction across the trans-node boundary, and may suppress only part of a saturated event that crosses the boundary. Bad events (principally cosmic ray strikes) can render otherwise normal pixels bad on a frame-by-frame basis, in the sense that x-ray photons couldn't be detected there in that frame. Filtering on `BAD_SHAPE` sets all pixels associated with a bad event of this type bad in the frame where it occurs. But filtering on `BELOW_ACCEPTANCE`, which indicates a non-x-ray event on pixels that would have registered an x-ray if exposed to one, has no effect on the exposure map. If an event is rejected for this reason only, then its pixels remain normal in the exposure map. Note that due to the arbitrary alignment between the CCD grid and the angular grid, pixels marked bad in every frame tend to appear as dips rather than as holes in the exposure map.

The node-specific exposure maps are constructed on a sub-space of the two-dimensional channel space established by `rgsangles` for the event coordinates (`BETA_CHANNEL` or `MLAMBDA_CHANNEL` and `XDSP_CHANNEL`). Each pixel in the exposure map covers the same angular region as a corresponding pixel in the event channel space, when using `BETA_CHANNEL` calculations, but to minimize file size a node-specific exposure map covers only the exposed surface of its node. The FITS standard World Coordinate System (WCS) keywords are present in both the event coordinate columns and the exposure maps, and can be used to place events onto the node-specific sub-space. The following equations relate the various coordinate systems. b (dispersion) and c (cross-dispersion) are exposure map coordinates, with integer values referring to the midpoint of a channel. Following FITS convention, the first element of the map array is labeled (1,1). β (`BETA_CORR`) or λ (`LAMBDA`) and ξ (`XDSP_CORR`) are aspect-corrected RGS angular coordinates.



B (BETA_CHANNEL) and C (XDSP_CHANNEL) are event coordinates, or when used as subscripts on WCS attributes, refer to those respective columns; all other WCS attributes refer to the exposure map. Note also that $CDELTA1 \equiv TCDLT_B$ and $CDELTA2 \equiv TCDLT_C$.

- from event channels to map coordinates

$$\begin{aligned} b &= B + \frac{TCRVL_B - CRVAL1}{CDELTA1} + (CRPIX1 - TCRPX_B) \\ c &= C + \frac{TCRVL_C - CRVAL2}{CDELTA2} + (CRPIX2 - TCRPX_C) \end{aligned}$$

- from map coordinates to angular coordinates

$$\begin{aligned} \beta &= CDELTA1 \cdot (b - CRPIX1) + CRVAL1 \\ \xi &= CDELTA2 \cdot (c - CRPIX2) + CRVAL2 \end{aligned}$$

- from map coordinates to chip coordinates

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} T_BETA_X & T_XDSP_X \\ T_BETA_Y & T_XDSP_Y \end{pmatrix} \begin{pmatrix} b \\ c \end{pmatrix} + \begin{pmatrix} T_ORIG_X \\ T_ORIG_Y \end{pmatrix}$$

The parameters of the linear transformation from map coordinates to chip coordinates are attributes of the exposure maps. The actual transformation is not linear, but this approximation is good to a fraction of a CCD pixel within the limits of the map. Further, because these parameters describe the average aspect correction, coordinates at the edges of the map may transform to outside the limits of the node.


3.2 Combined Exposure Map

The combined exposure map covers the area defined by **rgsangles** as the minimum and maximum legal values (TLMIN and TLMAX) for BETA_CHANNEL or MLAMBDA_CHANNEL and XDSP_CHANNEL. This means that in general the combined map will extend well beyond the exposed surface of the focal plane. Even so it is nevertheless possible for the camera to drift outside these bounds, and in that case the combined exposure map is truncated with respect to the aggregate of node-specific maps (which are aligned to the same grid but not confined by it). The warning emitted in this unlikely circumstance will most likely indicate a serious problem with the aspect correction or GTI filtering. The combined map is output to its own dataset, separate from the event list.

Since **rgsangles** sets TLMIN=1, event channel coordinates can be located on the combined map without any change of origin. As in the node-specific exposure maps, the transformation from map coordinates to angular coordinates is:


$$\begin{aligned} \beta &= CDELTA1 \cdot (b - CRPIX1) + CRVAL1 \\ \xi &= CDELTA2 \cdot (c - CRPIX2) + CRVAL2 \end{aligned}$$

The combined map does not define a linear transformation from map coordinates to chip coordinates. Instead, the approximate mapping of the node boundaries is preserved as a linear transform on the dispersion coordinate:



$$\begin{aligned}n &= T_BETA_N \cdot b + T_ORIG_N \\ccd &= 1 + n/2 \\node &= \begin{cases} 0, & n \text{ is even} \\ 1, & n \text{ is odd} \end{cases}\end{aligned}$$

The two parameters of the transform are attributes of the output dataset.



4 Parameters

This section documents the parameters recognized by this task (if any).

Parameter	Mand	Type	Default	Constraints
-----------	------	------	---------	-------------

mergedset	yes	dataset	merged.ds	
------------------	-----	---------	-----------	--

The input combined event list. It must contain the `EVENTS` table, and for each `EXPOSUOn` table that it contains it must also contain the `REJPIXOn` and `REJPIX1n` pair of tables. For exposure map generation the `BADPIXOn` and `BADPIX1n` pair of tables is also required. This dataset is not modified.

evlist	no	dataset		
---------------	----	---------	--	--

The output event list, a filtered copy of the input `mergedset`. The default file name is inferred from the input `mergedset` name as: `*EVENLI????.*` if there are ten characters before the suffix, and otherwise just `evlist.*`.

auxgtitables	no	list of tables		
---------------------	----	----------------	--	--

A list of OGIP compliant auxiliary GTI tables for filtering the event list in conjunction with its internal `STDGTIOn` tables.

rejflags	no	list of strings	“BAD_SHAPE ON_NODE_INTERFACE ON_BADPIX NEXT_TO_BADPIX ON_WINDOW_BORDER BE- LOW_ACCEPTANCE”	BAD_SHAPE ON_NODE_INTERFACE ON_BADPIX NEXT_TO_BADPIX ON_WINDOW_BORDER BE- LOW_ACCEPTANCE
-----------------	----	-----------------	--	--

A list of rejection flags for filtering the events. See `rgsevents` for the meaning of each flag.

withexpmaps	no	boolean	yes	
--------------------	----	---------	-----	--


If true, an exposure map for each node is added to the output dataset, overwriting any previous versions of the maps. The maps are constructed to be consistent with the filtering of the events. Chips without an `EXPOSUOn` table in the event list are not considered. *Requires Spectroscopy mode data.*

withcombmap	no	boolean	yes	
--------------------	----	---------	-----	--

If true, parameter `expimageset` specifies a second output dataset in which the combined exposure map is placed as the primary image. *Requires withexpmaps=true.*

expimageset	no	dataset		
--------------------	----	---------	--	--

The output combined exposure map. The default file name is inferred from the input `mergedset` name



as: *EXPMAP?????.* if there are ten characters before the suffix, and otherwise just `expmap.*`. *Requires* `withcombmap=true`.

driftbinsize	no	real	1	0-
---------------------	----	------	---	----

The size in arcseconds of the square bins used to group frames according to their aspect drift in dispersion and cross-dispersion. *Requires* `withexpmaps=true`.

spectrumbin角度	no	choice	lambda	lambda beta
----------------------	----	--------	--------	-------------

dispersion direction binning type. Used to accumulate the exposure maps.

5 Errors

This section documents warnings and errors generated by this task (if any). Note that warnings and errors can also be generated in the SAS infrastructure libraries, in which case they would not be documented here. Refer to the index of all errors and warnings available in the HTML version of the SAS documentation.


badParameter (*error*)

One or more of the values of parameter `rejflags` is not recognized by the task implementation, even though permitted by the interface. This should not occur.

lostExposure (*warning*)

Some portion of a node does not lie entirely within the boundaries of the combined exposure map. This need only concern users of the combined exposure map. In particular, note that the tasks to generate spectra, lightcurves and response matrices *do not* use the combined exposure map. However, the default boundaries are ample enough that this situation should not occur, and may indicate some other problem.

corrective action: The combined exposure map is truncated at its boundaries, which are set by task `rgsangles`.



6 Input Files

The following input files are never modified by this task.

- **mergedset** — input event list

Attributes: TELESCOP, INSTRUME, DATAMODE, OBS_ID, EXP_ID, EXPIDSTR, DATE-OBS, DATE-END.

Extensions:

- **STDGTIO_n**

OGIP compliant table of Good Time Intervals. Originally constructed by **rgsframes** and possibly modified subsequently by **rgsfilter**.

- **EVENTS**

Attributes:

CCDOCB		rgsframes
SOURCEID	int	rgsangles
ASPFAC1	real	rgsangles
REJFLAGS	int	rgsfilter (optional, assumed zero if not present)

Columns:

TIME		real, s	rgsevents
FLAG		int	rgsevents
BETA_CHANNEL	XDSP_CHANNEL		rgsangles

The channel columns are required only for these attributes: TCDLT, TCRPX, TCRVL, TLMIN and TLMAX.

- **EXPOSUO_n**

Columns:

FRAME		int32	rgsframes
TIME		real, s	rgsframes
TIMEDEL		real32, s	rgsframes
FRACEXP0 [†]	FRACEXP1 [†]	real32	rgsframes
ASPCDSP	ASPCXDSP	real32, rad	rgsangles

- **REJPIX_{nn}**

Columns:

FRAME		int32	rgsevents
FLAG		int32	rgsevents
CHIPX	CHIPY [†]	int16	rgsevents

- **BADPIX_{nn}**

Attributes:


CCDNODE		int	rgsbadpix
ABUTTED		bool	rgsbadpix
WINDOWX0	WINDOWY0	int	rgsbadpix
WINDOWDX	WINDOWDY	int	rgsbadpix

Columns:

CHIPX	CHIPY [†]	int16	rgsbadpix
YEXTENT [†]		int16	rgsbadpix
BADFLAG		int16	rgsbadpix

- **auxgtitables** — auxiliary GTI tables

Optional OGIP compliant tables of Good Time Intervals.



7 Output Files

- `evlist` — output event list

The output event list is a clone of the input event list in which the tables `EVENTS`, `STDGTIOn`, `EXPOSUOn` and `REJPIXnn` have been destructively filtered and the node-specific exposure maps have been deleted or replaced. The `STDGTIOn` tables and the `REJFLAGS` attribute of the `EVENTS` table indicate what filtering has been performed. The specification of the node-specific exposure maps follows:

- `EXPMAPnn`

A node-specific exposure map is an image (FITS array) on a channel grid that is commensurate with the channel coordinate system defined in the `EVENTS` table, but covering only the exposed surface of a single node. Array values are seconds of exposure in single-precision, floating-point format. The two dimensions of the array are dispersion and cross-dispersion, arranged such that values adjacent in dispersion are adjacent in memory. The following attributes are provided:

<code>CCDID CCDNODE</code>	CCD chip and node identifiers
<code>CCDOCB SOURCEID</code>	copied from the <code>EVENTS</code> table
<code>CTYPE1 CUNIT1 CRPIX1 CRVAL1 CDELTA1</code>	dispersion axis WCS
<code>CTYPE2 CUNIT2 CRPIX2 CRVAL2 CDELTA2</code>	cross-dispersion axis WCS
<code>T_BETA_X T_XDSP_X T_ORIG_X</code>	average transform to CHIPX
<code>T_BETA_Y T_XDSP_Y T_ORIG_Y</code>	average transform to CHIPY
<code>DATAMIN DATAMAX</code>	minimum and maximum exposure

The two exposure maps for a chip are not constructed unless the `EXPOSUOn` table for that chip is present in the input event list. However, if this table does exist and contains no frames, two null exposure maps are added to the output. A null map has zero extent in both dimensions. In this case the WCS, the coordinate transform, and the data range attributes above cannot be defined and are omitted.

- `expimageset` — combined exposure map

If enabled, the combined exposure map is placed in this output dataset as the primary image (FITS array). It is formatted in the same way as a node-specific exposure map, except that the channel grid covers the full range specified in the `EVENTS` table. Coordinates are identical between this map and the event list. The following attributes are provided:

<code>CCDOCB SOURCEID</code>	copied from the <code>EVENTS</code> table
<code>CTYPE1 CUNIT1 CRPIX1 CRVAL1 CDELTA1</code>	dispersion axis WCS
<code>CTYPE2 CUNIT2 CRPIX2 CRVAL2 CDELTA2</code>	cross-dispersion axis WCS
<code>T_BETA_N T_ORIG_N</code>	average transform to CCD and node
<code>DATAMIN DATAMAX</code>	minimum and maximum exposure

Attributes `TELESCOP`, `INSTRUME`, `OBS_ID`, `EXP_ID`, `EXPIDSTR`, `DATE-OBS` and `DATE-END` are also copied from the primary header of the input event list. In the unusual case that there is no exposure on any of the chips (all of the node-specific maps are null) the combined map is zero everywhere and the attributes `T_BETA_N` and `T_ORIG_N`, which cannot be defined, are omitted.



8 Algorithm

```
Initialize the output event list:
  create a DAL clone of the input event list: *evlist
  delete all EXPMAPnn extensions

if *withexpmaps && *withcombmap:
  create the dataset for the combined exposure map: *expimageset
  set the primary image dimensions to the EVENTS table channel space
  initialize image to zero

Complete the STDGTI tables:
  auxiliaryGti = intersection(*auxgtitables)
  for each chip:
    for each valid frame contained in *auxiliaryGti:
      alignedGti |= gti(frame)
      table STDGTI0n &= *alignedGti

Filter the EVENTS table:
  Sn = "(CCDNR == n && GTI([STDGTI0n],TIME))"
  gtiExpr = "(*S1 || *S2 || *S3 || *S4 || *S5 || *S6 || *S7 || *S8 || *S9)"
  selectlib expression: "(FLAG & *rejflags == 0) && *gtiExpr"
  attribute REJFLAGS |= *rejflags


for each chip present:
  Filter the EXPOSU0n table:
    selectlib expression: "GTI([STDGTI0n],TIME)"


  for both nodes:
    Filter the REJPIXnn table:
      simultaneously traverse EXPOSU0n and REJPIXnn, deleting all
      rows in the latter that are not in correct FRAME order
      according to the former; unexpected frames are deleted

  if *withexpmaps:
    Construct the EXPMAPnn array:
      sort frames by group according to *driftbinsize
      for each group:
        adjust coordinates using group average aspect correction
        for each frame:
          subtractively overlay rejected pixels with frame timedel
          accumulate group timedel
          additively overlay exposure with group timedel
          subtractively overlay global bad pixels with group timedel
        for each pixel:
          pixel = 0 if |pixel| < 1.788e-7 * exposure

  if *withcombmap:
    additively overlay EXPMAPnn on combined map

if *withexpmaps && *withcombmap:
  compute dispersion-channel-to-node transform
  find min and max data values
```





9 Comments

The exposure maps suffer from a minor potential double-counting error. If a bad pixel is not known to the DPP, and is not present in the telemetry, and is of the variety that tends to be hot, then it is possible that it has been rejected on-board. Such bad pixels effectively subtract twice their active time from the map, once as an individual bad pixel, and again in the fractional exposure applied to the entire node.

References