

Behind, Beneath, and Before RHESSI Spectroscopy

David Smith, dsmith@ssl.berkeley.edu

Version 3, 04/12/02

This document is intended to introduce all the issues surrounding RHESSI spectroscopy, i.e. the process of converting raw RHESSI data to a spectrum in photons/cm²/s/keV originally incident on the spacecraft. **Most of this conversion will be done automatically by the data analysis software with few decisions required of the user.** Nonetheless, it is good for anyone intending to interpret RHESSI spectra to know what the spacecraft and software are doing behind the scenes, beneath the surface, and before you see your spectrum. This document contains sections on

- The RHESSI detectors
- The detector electronics
- The RHESSI data formats
- Other devices on the instrument related to spectroscopy
- The phases of a spectroscopy analysis.

It is not intended to be a comprehensive reference on any of these topics, rather a broad introduction to all of them for the scientist intending to do RHESSI spectroscopy. For an introduction to doing spectroscopy with the RHESSI Graphical User Interface and the SPEX fitting package, see Brian Dennis's Spectroscopy First Steps document at http://hesperia.gsfc.nasa.gov/~dennis/spectroscopy/first_steps.htm ("First Steps").

For a listing of quirks of individual detectors and the system as a whole which can trap the unwary, see the companion document to this one, called **Problems, Peculiarities and Phenomena in RHESSI Spectroscopy** ("PPP").

DETECTORS

The arrangement of the RHESSI detectors in the spectrometer (one beneath each grid) is shown in Figure 1. Figure 2 is a photograph of the assembled spectrometer before it was inserted into the bottom of the spacecraft. We refer to the detectors by the name of their associated grid pair: G1, G2, ... G9.

Each detector is a single crystal of hyperpure germanium in the general shape of a cylinder with a bore partway up the middle (see Figure 3). When hyperpure germanium is at cryogenic temperatures, no electron-hole pairs are in the conduction band, but a hard x-ray or gamma ray interacting in the crystal will release one or more energetic electrons, which lose energy by creating many free electron/hole pairs. If there is a high electric field (on the order of 1000 V/cm) across the crystal, the electrons and holes will be pulled to the electrodes, creating a current pulse. The total charge in the current pulse is proportionate to the photon energy.

Two conductive layers are implanted on the crystal surfaces to serve as electrodes: a thin, p-type layer of




Figure 1: Arrangement of the RHESSI detectors inside the cryostat. The small cylinder nestled among the detectors is the beryllium scatterer for polarization measurements (see below). The detectors are numbered with the number of the grid they sit beneath.




Figure 2: The assembled spectrometer, showing the aluminum attenuator disks (see below) on top (one set in the aperture, the other out).

implanted boron on the front and side surfaces, and a thicker, n-type layer of diffused lithium ions on the inner bore. The rear surface is left as an insulator. The material overall is very slightly n-type, and when 2000-4000V is applied between the inner and outer electrodes the crystal is depleted of these charge carriers, with enough electric field in the crystal from the space charge and external voltage to cause the electron-hole pairs to reach terminal velocity. Most of the RHESSI detectors will be operated at a high voltage (HV) of 4000V.

The step seen near the top of the inner bore in Figure 3 (top) is a break in the lithium contact. The signals are extracted separately from the two halves of this electrode. The line extending from this step to the outside edge of the detector represents a boundary electric field line: photons stopping above this line are detected in the front channel, and those stopping below it in the rear channel. **Thus a single crys-**




Figure 3: Cross-sections of a RHESSI detector. Top: a detector profile with field lines, with the field line marking the segment boundary in bold dashes. Bottom: a detector in the cryostat, showing Ta/Sn/Fe/Al shielding around the side of the front segment and above the shoulder of the rear segment.

tal becomes a stacked pair of detectors. The front segment will absorb all the hard x-rays up to about 100 keV, letting most gamma-ray line photons through. The rear segment will stop many of the latter, so that gamma spectroscopy can be done without high deadtime from the x-rays.

You may also hear the front segment referred to as the “hemisphere” or “cap” segment, and the rear segment referred to as the “coax” segment. We try to use the words “segment” and “detector” consistently: RHESSI has 9 detectors and 18 segments.

The notch on the outer edge of the detector serves two purposes. First, it concentrates the electric field lines at the corner of the notch, so that the field line which originates at the inner step always hits the proper place on the outside of the detector. In addition, it removes some mass from in front of the rear segment, so that fewer high-energy gamma rays Compton scatter before entering the rear. To keep the “shoulder” part of the rear segment from being swamped with flare hard x-rays, a ring of thin “graded-Z” passive shielding is placed above the shoulder (Figure 3, bottom). This shield is just as effective as the front segment in photoelectrically absorbing hard x-rays, but with much less Compton scattering of gamma-rays. The graded-Z shielding consists of 0.5mm of Ta, 1mm of Sn, 0.5mm of steel, and the aluminum which is part of the cryostat structure. Each element stops the K-shell fluorescence photon from the higher-Z element before it. The whole is opaque up to about 100 keV. Additional graded-Z shielding is put in a ring around the edge of the front segment to stop background photons from the Earth and cosmic diffuse background, and on top of the spectrometer in the area between detectors to stop x-rays from scattering around in the spectrometer and entering the rear segments from the side.

ELECTRONICS

The RHESSI electronics produce a series of photon events (4 bytes each) containing the energy, time, and detector segment of the photon plus some encoded livetime information. There are two other data formats produced as well: “monitor rates”, which are 1-second accumulations of various counters related to activity in the detectors (see below), and, at times of very high count rate, “fast rates”, which record the count rate in each segment at a very fast cadence (also see below).

Front-End Amplification

The first stage of amplification comes in an FET immediately behind the detector. The signal from the FET is taken by a special low-thermal-conductivity harness out to the rest of the Charge Sensitive Amplifier (CSA). One CSA for each detector sits along a ring around the outside of the spectrometer, alongside the HV filter which conditions the HV for the same detector. The CSA amplifies the signal further and also drives a reset transistor for each segment (which is located next to the signal FET inside). The reset transistor drains the accumulated charge whenever the total accumulated energy from photon interactions, particle interactions, and leakage current for its segment is about 40 MeV. Figure 4 shows, schematically, the raw output of the CSA.

You may also hear the CSA called the “preamp”.

Fast & Slow Channels

The CSA signals for front and rear segments are sent to a Detector Interface Board (DIB). There is one DIB for each detector, and they are located in the Instrument Data Processing Unit (IDPU) box. The signal chain for each of the 18 segments is split into two: a fast chain which converts each event (step in the raw CSA signal) to a pulse of about 800 ns total

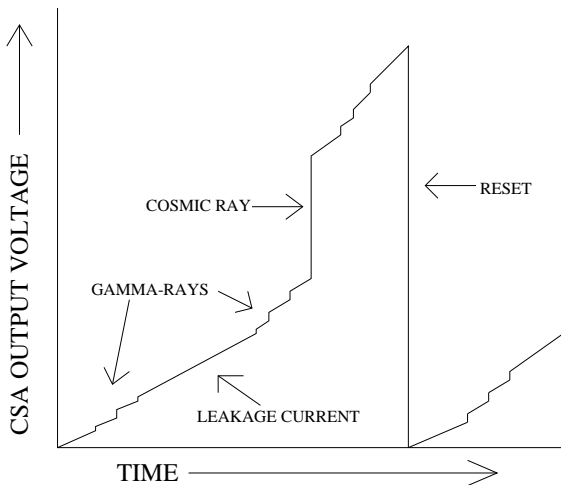


Figure 4: Output voltage of the CSA vs. time as seen by an oscilloscope. On top of a constant leakage current, photons and cosmic rays represent sudden increments of charge. Periodic resets drain the accumulated charge and restore the voltage baseline.

width, and a slow chain which converts each event into a pulse of about 12 μ s total width. The slow signal, which has very good energy resolution because of the long integration time, is sent to the analog-to-digital converter (A2D) to produce the pulse heights written into each event. The fast pulse has poorer energy resolution but is used for decisions which have to be made quickly, i.e. pileup rejection and gain switching (see below), and for generating fast rate data when counting rates are too high to do event-by-event spectroscopy (also see below).

Both the fast and slow channels have lower-level discriminators (LLDs). For the front segment, these are set as low as possible without triggering on random noise: about 3 keV for the slow channel and 6 keV for the fast channel. One front segment (G7) has these thresholds set a little more than a factor of 2 higher due to an extra white noise source in that channel. For the front segment, the slow LLD must trigger to signal the A2D to begin digitizing and produce an event. For the rear, both thresholds must be exceeded for the A2D to produce an event, and both are set near 20 keV. Below 20 keV, photons approaching the rear segment from any direction will be stopped by passive material or by the front segment.

A2Ds & Gain

The A2D produces a 13-bit (8192 channel) pulse height value for each event. Gain is extremely linear, with deviations of < 1 keV from linearity across the whole scale. The approximate gain function for the front segments is $\text{Energy(keV)} = \text{Channel} * 0.35 - 12 \text{ keV}$, giving an energy range up to about 2.8 MeV. Note that this places zero energy around channel 34. There are significant variations from segment to segment, but variations with time, temperature, etc. for any one segment should be small ($\ll 1\%$). Nonetheless, because the resolution is so high, we anticipate keeping a database of gain variations with time in order to get the best possible performance from the detectors. These gains will be calculated orbit-by-orbit from fitting the positions of background lines. The spectroscopy tools will select the appropriate gain function in a way that should be invisible to the user.

The rear segments have a gain function near $\text{Energy(keV)} = \text{Channel} * 0.39 - 12 \text{ keV}$, or up to about 3.2 MeV. To analyze gamma-rays at even higher energies, there is a second gain range added to the rear segments. The pulse from the fast chain is sent out ahead of the slow spectroscopy signal. If its amplitude is higher than a certain threshold (about 2.9

MeV), the gain of the A2D is changed so that the 8192 channels represent an energy scale about a factor of 5.2 higher (to 17 MeV). There is a bit in the 4-byte rear-segment event to tell which gain level the 13-bit channel value was taken with. Because the resolution of the fast signal is poor, the transition from one gain level to another is not abrupt: i.e, photons in the range 2.6-3.2 MeV can appear at either gain. By using events from both gain ranges, however, an undistorted spectrum can be reproduced. Figure 5 shows how data taken with a high-energy source (^{56}Co) are split between the two gain ranges. Most of the time, the spectroscopy tools will automatically search both gain ranges for events if the user has asked for data in the 2.6-3.2 MeV range.

Pileup Rejection & Livetime

If there are two events less than 800ns apart, their energies will be added together (piled up) - this is unavoidable, but should happen very seldom except in the largest flares. For event separations of 800ns to $6\mu\text{s}$, the DIB will veto both events based on the two fast signals, since the slow signals would overlap and contaminate each other, distorting the energy value. For event separations from 6 to $9\mu\text{s}$, only the second event would have its energy distorted, so only the second event is vetoed. In the front segment, events below about 6 keV (which don't trigger the fast LLD) can pile up unrestrictedly on other events. The pileup correction algorithm will eventually be incorporated into the routine analysis, but **as of now**,




Figure 5: Transition of counts from one gain range to another. Note the minor pileup of events in the top 64 channels of the lower energy range (see “Spectral Artifacts” in PPP).

pileup correction is NOT a regular part of the spectroscopy process outlined in First Steps.

Pileup correction is an iterative process which is based on the shape of the raw spectrum and knowledge of the count rate and measured livetime. After the spectral shape is corrected once, the corrected spectrum can be fed back into the algorithm for successively finer corrections. Second-order pileup corrections (for three photons adding together) and higher will eventually be needed for the brightest flares, and for times when both attenuators are out (see below).

Livetime, on the other hand, is independent of the energy, and the livetime correction changes only the overall magnitude of the spectrum (i.e. photometry) rather than the spectral shape. The livetime of a detector is simply the probability that one additional photon, put in at a random time, would in fact be recorded. Deadtime is $(1. - \text{livetime})$, the percent time that the circuit is “tied up” by other events. The livetime signal for each segment is a logic line which is “OR”-ed together from several lines in the circuit, representing several reasons the circuit could be unable to process an event (A2D in use, pileup rejection active, CSA reset in progress, etc.) This line is sampled at 1 MHz and the livetime reported in the monitor rates is the fraction of these sample strobes which read the line as clear.

Because the pileup circuit can sometimes choose to veto BOTH events rather than just one, there is a loss of livetime due to pileup rejection which cannot be measured just by sampling the pileup-rejection logic line. Therefore the livetime correction routines do not use the “raw” livetime measurement as reported in the monitor rates, but a corrected livetime which is somewhat lower. This correction is automatically taken into account by the spectroscopy format.

The maximum event throughput is about 25,000-30,000 counts per segment per second and it is reached at about 50% livetime.

DATA FORMATS

Event Packets

Most of the time, whether a flare is active or not, every photon event in the RHESSI detectors is stored in onboard memory as 4 bytes of data and telemetered to the Berkeley ground station (or a backup station) within a day or two. Each event contains:

- 5 bits identifying the segment and (if a rear segment) energy range of the event,
- 13 for the energy channel,
- the last 10 bits of the time counter, giving the event time in units of binary microseconds (i.e. 2^{-20} seconds), and
- 4 bits which are a partial measurement of the livetime counter. The livetime counter is read every $1/2048$ second and its 9-bit value is broken up into 3-bit chunks and placed in this data slot in three consecutive events from that segment. The 4th bit is set to 1 to indicate where this sequence of three events begins. If the events are not coming quickly enough to produce a continuous livetime readout this way, they still produce a sampling of it.

If the count rate is low enough that the last 10 bits of the clock roll over between events, dummy “time stamp” events are inserted so that the time on each event is not ambiguous.

If the onboard memory starts to fill up, a decimation algorithm automatically throws out all but one out of every N events in the front segments below a certain energy E , with N from 2-16. E and N are functions of both the remaining memory and the position of the attenuators (see below). Decimation in the rear segments can be commanded as a routine way of keeping background (mostly photons from the Earth’s atmosphere or the cosmic diffuse background) from filling up the memory. For example, a high level of decimation can be commanded for times when RHESSI is in shadow.

Monitor Rates

Monitor rates are a separate kind of data packet (“ApPID” in telemetry parlance) produced once every ten seconds. Each packet consists of 10 1-second accumulations of a set of counters for each of the 18 segments. The counters are: number of resets, number of slow LLD triggers, number of fast LLD triggers, number of upper-level discriminator (ULD) triggers on the slow channel (i.e. events above the energy of channel 8192 - usually cosmic rays), and fraction of “live” results from the livetime strobe. This information is normally used to check the health of the detectors and is not necessary for spectroscopy.

Due to an oscillation in the CSA that occurs only during a reset, the reset counters in the rear segments

increment multiple times on each actual reset; the multiplicity factor for each segment is known and will be available to any software that uses these data.

Fast Rates

Fast rates are yet another data format, produced only when the count rates are very high. The data are count rates in four broad energy bands (all in the hard x-ray range). The pulses are sampled from the fast electronics chain. The rates for the three detectors (1,2,3) with the finest grids (and therefore fastest imaging modulations) are sampled at 16 kHz; the next three (4,5,6) at 4 kHz; and the three coarsest grids (7,8,9) at 1 kHz. Events are not shut off when the fast rate data turn on; however, at these very high count rates the event data will naturally taper off due to very high deadtime.

RELATED DEVICES

Attenuators

There are two attenuators designed to cut down the photon flux during bright flares so that the detector deadtime stays below 50%. Each attenuator is a set of 9 aluminum disks held in a lightweight frame. Each frame has only two positions: one with the disks covering the detectors and one with the detector field of view completely clear. They are moved automatically in response to increases and decreases in the monitor rate counters. The frames are moved by putting




Figure 6: Attenuation factor of the attenuator configurations as a function of energy.

a high current through shape-memory alloy (SMA) wires which then contract. The move takes less than 1 second, so there will be only a short time when the instrument response is in transition between well-calibrated states.

Each attenuator disk consists of three concentric regions of differing thickness. The “heavy” disks have a thick outer region covering most of the detector area, a thin middle region of smaller area, and a tiny, very thin spot in the center. The “light” disks have only thin aluminum in the large outer area, thick aluminum in the middle region, and a very thin spot in the middle somewhat larger than that of the heavy disk. Together, these disks provide four levels of attenuation which allow RHESSI to have an enormous dynamic range of sensitivity and make useful measurements over 6 orders of magnitude in incident flare (or microflare) flux. When the attenuators are both in place, there is thick aluminum over both the outer and middle regions, and only the smallest spot in the very center admits the soft photons which are so prolific in the thermal (post-impulsive) phase of large flares.

You may also hear the attenuators referred to as “filters” or “shutters”.

In the early parts of the mission, we have been experimenting with several modes of shutter operation. Sometimes the thin shutter has been fixed “in” and “out” for long periods, sometimes we have enabled it to move in automatically if the livetime drops below a certain level, sometimes we have had the thin shutter fixed in and the thick shutter free to move in if a truly enormous flare happened, and sometimes in the fully automated mode which enables transitions between states “0,1, and 3” (no shutter, thin in, and both in). **Although the analysis software automatically knows the shutter state at the time of your observation and accounts for its opacity,** you may want to choose your flares based on the shutter status near the time of observation. This information is summarized for the whole mission on this web page: <http://sprg.ssl.berkeley.edu/~dsmith/hessi/shutterrecord.html>.

Figure 6 shows the attenuation caused by the thermal blankets and beryllium windows above the detectors (top trace) and the attenuation with the thin, thick, and combined attenuators. There are currently no plans to use state “2” (thick in, thin out), the second to last trace on the plot.

Onboard Radioactive Source

RHESSI carries a tiny onboard radioactive source (5 nanocuries of ^{137}Cs) which makes a line at 662 keV, far from any line expected to occur in flares or in RHESSI’s variable background. The count rate from this source is so small that it can only be detected in spectra accumulated over many hours. The line of known intensity lets us monitor any changes in the efficiency of the detectors such as might occur via radiation damage (see below). The same function can be served by a line at 1460 keV from naturally occurring ^{40}K in the spacecraft, but the absolute flux of this line is not known before launch since there is always some of the isotope in the laboratory. Other sources of background will be far more important in determining the sensitivity of RHESSI.

Scatterer for Polarization

Since the rear segments see no direct flare photons below 100 keV, we can use them at these low energies as a crude hard x-ray polarimeter. There is a cylinder of beryllium 3 cm in diameter and 3.5 cm long nestled among the rear segments (see Figure 1). Above this cylinder is a thin spot in the spectrometer shell and a hole in the grid trays, so that solar photons > 20 keV can reach the cylinder and scatter into the adjacent rear segments. The Klein-Nishina cross section, differential in azimuth angle, is a function of the angle from the polarization axis. Thus, by watching the relative rates of these rear segments, we will measure the direction and degree of polarization for incoming photons of roughly 20-100 keV. Simulations suggest that we will be able to detect polarization fractions as low as a few percent for the largest flares. The key difficulty in the analysis will be photons scattered from the Earth’s atmosphere, which also produce low-energy counts in the rears and which also vary with the spacecraft spin.

Onboard Pulsar

The IDPU can put regular, small pulses into the detectors’ HV supply, which the electronics see as equivalent to photon events. The pulse energy can be tuned across the detectors’ full range, but the front/rear ratio is fixed at roughly 1:3 because the HV is shared by both segments and their response is proportional to their capacitance. Pulse rates can be commanded separately for each detector at 11 discrete frequencies, spaced by a factor of two, up to 1024 Hz. This pulser is not required in flight, but may be used for diagnostics.

PERFORMANCE

FWHM Resolution (keV) in flight: front segments measured at 93.3 keV, rear segments measured at 1117 keV.

Det	Front	Rear
G1	1.13	2.90
G2	7.94	—
G3	0.98	2.77
G4	0.98	2.82
G5	1.47	2.73
G6	1.01	3.05
G7	3.15	2.98
G8	1.26	3.36
G9	1.19	2.27

The effective areas of single RHESSI front and rear segments are shown in Figure 7, from a simulation with an accuracy of a few percent. The nominal resolution performance is approximately 1 keV FWHM in the front segments and 2.5 keV FWHM in the rears; the values for individual detectors are tabulated above and discussed in PPP. By comparison with the expected solar line widths (Figure 9) it is

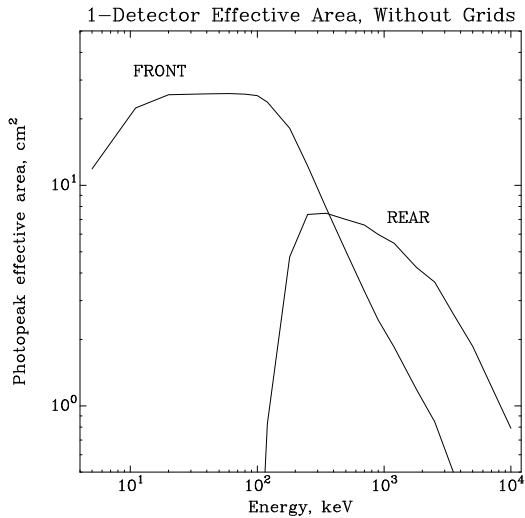


Figure 7: Effective area of the front and rear segments of a single RHESSI detector for photopeak detection of a downward-incident photon, without the shadowing of the grids.

apparent that all of the solar lines except the very narrow 2.2 MeV line will be resolved, and the very steeply falling continuum spectrum of the superhot component will not be distorted. This is a striking contrast to previous missions based on scintillators.

The effective area, resolution, and predicted background level allow us to calculate RHESSI's sensitivity to the gamma-ray lines in flares. Even though RHESSI carries no shielding to block background photons from the Earth or the cosmic diffuse sky, the flares which will show gamma-ray lines are so intense that the bremsstrahlung emission from the flare itself will overwhelm these sources of background, and will therefore define the continuum above which the lines must be detected. As of this writing (4/12/02), no X-class flares have yet been seen by RHESSI.

Figure 8 shows the expected RHESSI response to a intense gamma-ray line flare, with the origin of many of the lines labeled. Figure 9 shows the expected sensitivity to these and other lines, defined in two ways: the statistical significance of each line for the largest expected gamma-ray line flare, and the line flux required for a 3σ detection.




Figure 8: Simulated RHESSI count spectrum from a large gamma-ray line flare.

SPECTROSCOPY ANALYSIS

Figure 10 gives a schematic representation of the various steps involved in a spectral analysis. In this document, we discuss only the philosophy of a spectral analysis in general; the First Steps document will show you how to use the RHESSI software. One of the most important things in a spectral analysis is that the steps are done **in order** as shown in Figure 10. We will make sure the RHESSI software makes it difficult or impossible to do otherwise. The default analysis described in First Steps is a relatively simple one, optimized for hard x-ray flares of modest intensity.

Imaging spectroscopy is beyond the scope of this document. The modulation introduced by the grids does have some energy dependence. Therefore, to make sure that the modulation has no distorting effect on your spectrum, **you should take spectra whose duration is roughly an integral number of half spins whenever possible.**

A general note: spectroscopy analyses will be more accurate when the whole procedure down to the last

Line Energy (MeV)	Excited Nucleus	Width (keV)	3 σ Line Fluence (ph/cm ²)	Large Flare			
				Line Fluence (ph/cm ²)	HESSI Line Counts	HESSI Continuum Counts *	Number of σ
Prompt lines							
0.339	⁵⁹ Ni	4	1.6	9.2	1512	37475	7.7
0.429†	⁷ Li	5	2.4	9.5	838	15839	6.5
0.478†	⁷ Be	10	3.5	9.5	854	25342	5.3
0.452†	⁷ Li- ⁷ Be	30	2.3	19.0	3195	89004	10.5
0.847	⁵⁶ Fe	5	1.4	17.3	1074	4471	14.4
0.932	⁵⁵ Fe	5	1.2	2.5	274	3193	4.7
1.369	²⁴ Mg	16	1.7	25.5	1567	6995	16.9
1.634	²⁰ Ne	20	1.6	75.3	4297	6863	40.7
1.778	²⁸ Si	20	1.9	30.4	1446	5245	17.7
2.618	²⁰ Ne	60	3.5	11.4	445	4187	6.5
4.439	¹² C	145	6.9	73.3	1294	5383	15.8
6.129	¹⁶ O	145	11.4	57.6	549	2716	9.6
Delayed lines							
0.511	e ⁻ /e ⁺	5	2.4	196.9	21060	15432	110.2
2.223	² H	2.5	0.6	298.7	13345	346	114.1

* For most lines, about 90% flare continuum and 10% instrumental background for this bright flare.

† The narrow lines are for a downward beam or a fan beam; the single broad line at 0.45 MeV is for an isotropic distribution. This line has an intrinsic width of 0.1 keV, so we used the instrument FWHM resolution (2.5 keV).

Figure 9: Expected sensitivities to the principal gamma-ray lines in a large flare.

step is done separately for every segment. For example, imagine that one segment has a much higher livetime than another, and is also picking up a slightly different-shaped spectrum (more background, perhaps). If you do an averaged livetime correction to a spectrum which is the sum of these two segments, their relative contributions will be incorrect. Similarly, if the count rate varies dramatically during the period of your analysis, the time intervals of high and low livetime in EACH segment should also be analyzed separately. If you want a summed spectrum over a long time, you can analyze it in pieces and average together the results at the end. We will try to make this easy for you to do, but it will sometimes be easier, for a quick look, to sum events from many segments and time intervals together and do a single pass through the steps listed below considering the entire spectrometer (or a set of segments you've chosen) as the detector.

Event Selection

A spectrum can be produced from one or more segments, and with and without coincidence in various modes. In general, for work below 200 keV, you should choose just the front segments. See the PPP document for hints on when you might want to leave

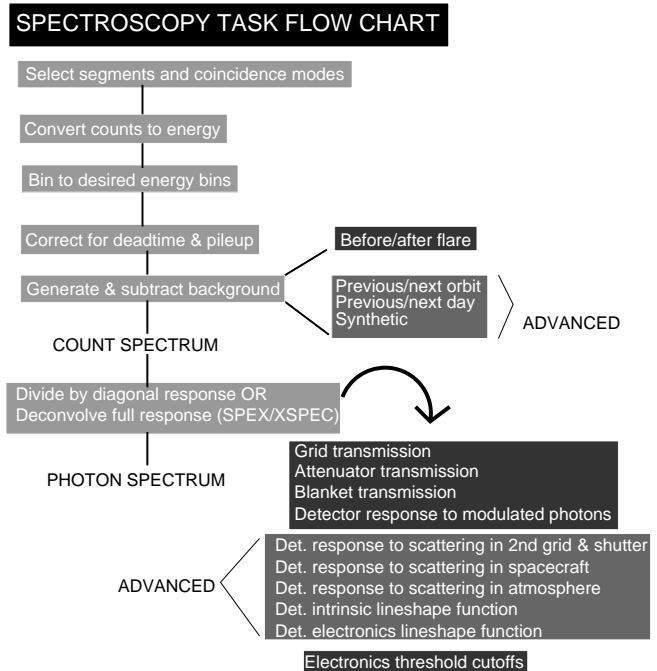


Figure 10: Flowchart of a spectroscopy analysis. "Advanced" features are discussed in the text.

one or more detectors out because of their particular characteristics. For higher energies, include the rear segments. For even higher energies (above the annihilation line at 511 keV), it will probably be useful to start including (true) coincident events between the front and rear segments and between adjacent rear segments. The spectroscopy software treats each possible combination of two segments as a “virtual segment”, with its own gain correction, efficiency, etc. **Any use of coincident events at all is considered an advanced function, and the implementation is not yet complete.**

Gain Correction & Binning

The software will use the best possible gain data for each segment; currently a single calibration for each segment, and, when it becomes available, one which includes the small variations with time.

The software can bin the data from all the segments (except the high-energy range of the rear segments) up to a standard finest binning of exactly 1/3 keV per bin, starting exactly at 0 keV, and going to 3 MeV in 9000 bins. This preserves just about all the information in the original spectra. The standard binning for the high-energy range is exactly 2 keV per bin.

It would be prohibitively difficult in terms of computer memory and processing time to try to do the spectroscopy analysis in this full set of bins, however, particularly if using the full response matrix (see “Full Response Correction” below). The spectrum would look pretty bad, too, since there might be few counts in each bin in most cases, and there is little scientific benefit to choosing bin sizes significantly less than the energy resolution. Therefore you must choose your own binning, and you should choose a binning fine enough to accomplish the scientific goal but otherwise giving as small an array as possible.

On the other hand, if the bins are TOO broad, the response, background, etc. will change too much across each of your bins for the average value to be a good representation - this is analogous to the perils, described above, of combining segments or time periods which are very different. As a rule of thumb, you shouldn’t expect good results from bins with $\Delta E > E/10$ (at the very lowest energies, from 3-10 keV, this rule will be violated).

You will often want to specify binnings with different widths at different energies. One example would

be a logarithmic binning, where $\Delta E/E$ is nearly constant; another would be a binning that gets fine only around narrow gamma-ray lines and is coarser everywhere else. We provide a selection of convenient binnings that you can choose just by referring to a number, and these will probably be enough for most uses, although you can always make up your own. To view the “canned” energy binnings available, you can call up the “SPECTRA” window on the Graphic User Interface (GUI) in the RHESSI software, and it the button labelled “Show Binning Codes”. This prints out a file called “energy_binning.txt” which lives in the SolarSoft tree under /hessi/dbase/spec_resp.

Pileup & Livetime Correction

Livetime correction, as mentioned above, affects only the intensity and not the shape of a spectrum, although it can affect the shape indirectly if different livetime corrections are applied to different segments which recorded different spectra.

Although the pileup correction is not yet available as part of the regular spectroscopy process, you can run it by hand to estimate the magnitude of its effect on your spectra. This is done with the routine “hsi_correct_pileup.pro”, which takes as its input the count spectrum and energies of the binned channels (these can be written out of the GUI into a .FITS file then read back into IDL) and the livetime (which

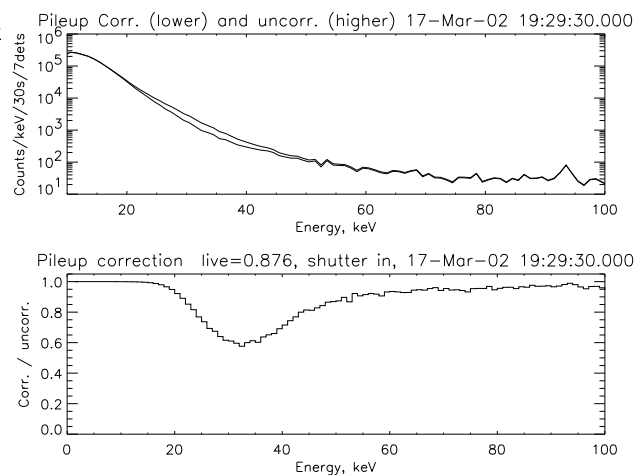


Figure 11: Pileup correction at modest deadtime with the thin shutter in place. Top: count spectra before and after pileup correction. Bottom: Ratio of corrected to uncorrected spectrum as a function of energy.

can be read from the monitor rate data if it is not available otherwise). This is a good idea for any flare with more than about 6000 counts per second in the front segments, or more than 2000 counts/second if both shutters are out. This lower threshold is because such flares will generally be dominated by low-energy photons, many of which will not pass the $\sim 6\text{-}8$ keV threshold for the pileup detection circuit, and thus will be much more likely to appear as piled-up events (artificially high energies) in the data.

An example of a currently uncorrectable, false high-energy feature caused by extreme pileup is shown in PPP. Figure 11 in this document illustrates a correctable level of pileup. Here we have a spectrum from an M-class flare on 2002-Mar-17 taken from 19:29:00 to 19:30:00 UT. The average recorded livetime in the 7 front segments used was 0.876, and so the true livetime was about 0.793. The thin shutters were in place, and count rates were roughly 12,000 c/s/front segment. Note that even at this modest deadtime, at around 30 keV the result is changed almost by a factor of two. This is due to the softness of the initial spectrum: since it's falling rapidly, even a very small fraction of the 15 keV photons adding together will create a relatively important number of 30 keV counts.

Background Subtraction

RHESSI is an unshielded, high-background instrument like BATSE on *CGRO* or TGRS on *Wind*, not a




Figure 12: Typical RHESSI background spectra in the front and rear segments.

shielded instrument like the Gamma-Ray Spectrometer on *SMM*. Therefore it will be necessary to subtract background even in the case of bright flares. A typical background spectrum in each segment is shown in Figure 12. Details of how the background varies with time are discussed in PPP. The simplest way to estimate the background is to average spectra from just before and just after your flare; this can be done within SPEX (see First Steps), which also allows you to fit the background at many different non-flare times to a polynomial vs. time to get the background underneath the flare. SPEX also allows you to pick different time periods for different energies. This is useful because the flares often have a long-lived thermal component. In these cases you may want to use only data **before** the flare to generate the background at the lowest energies, while using high-energy data during the late, soft, gradual part of the flare for generating background for use with the early, hard, impulsive part of the flare.

Sometimes, however, a flare will last for the better part of an orbit, or occur just after sunrise or just before sunset, so that this technique won't give a good result. Eventually, there will be other choices, e.g.:

- Background data taken from one orbit before and one orbit after the flare, at the same orbital phase as the flare data.
- Background data taken from 15 orbits (1 day) before and 15 orbits after the flare, at the same orbital phase as the flare data.
- Background generated by interpolation from a large database of many background spectra, sorted by various parameters such as the Sun-Earth-RHESSI angle, magnetic latitude, time since last SAA passage, etc.

Efficiency Correction or Full Response Correction

There remains only to correct for the efficiency of the instrument (i.e. the loss of counts to absorption in passive material), and possibly for scattering (i.e. for the movement of counts from their original energy to some other energy). The former correction is much easier, since it involves only multiplying each bin of the spectrum by a correction factor. Once this is done, we no longer give the spectrum units of "counts", but rather of "photons" (or photons/cm²/s/keV, etc.).

In some cases, it will only be necessary to do the efficiency correction. Below 100 keV, more photons

leave all their energy in the detector than scatter and leave less energy. This is shown in Figure 13, which shows the ratio between the number of counts in the photopeak and the number at lower energy (in the “Compton tail”). Below about 40 keV the Compton tail is truly negligible; but between 40 keV and 100 keV virtually all flare spectra fall, more or less steeply, so that even if the number of Compton tail counts from 100 keV photons is equal to the number of photopeak counts, it is dwarfed by the photopeak events from the lower energy photons. Thus, for hard x-ray work, the simple correction based only on the photopeak efficiency will often give a spectral shape very close to the correct one. One exception to this is at very low energies if at least one shutter has been put in. Then it is likely that many of the counts in the 3-10 keV band are actually due to photons about 10 keV higher, which have an easier time getting through the shutter, but which then suffer from the escape of a Ge-K fluorescence photon from the detector, moving their detected energy downward.

For the study of narrow gamma-ray lines, only the photopeak efficiency is needed, since the user is not concerned with the shape or intensity of the continuum. This is certainly true for the very narrow lines at 2.2 MeV and 511 keV, but will also be a good approximation for most of the nuclear excitation lines as well.




Figure 13: Peak to Compton ratio for a single RHESSI detector, combining front and rear segments (but vetoing coincident events). This can be improved at high energies by including coincident events and by including the narrow lines corresponding to the escape of one or two 511 keV annihilation photons (“escape peaks”).

The full response matrix will have to be used when looking at the overall shape of the gamma-ray spectrum above 100 keV, and looking at the continuum x-rays below 10 keV when the shutter is in. Figure 14 shows the response of the instrument to incident photons of three energies: 50, 350, and 2500 keV. They are each scaled so that the photopeak energy

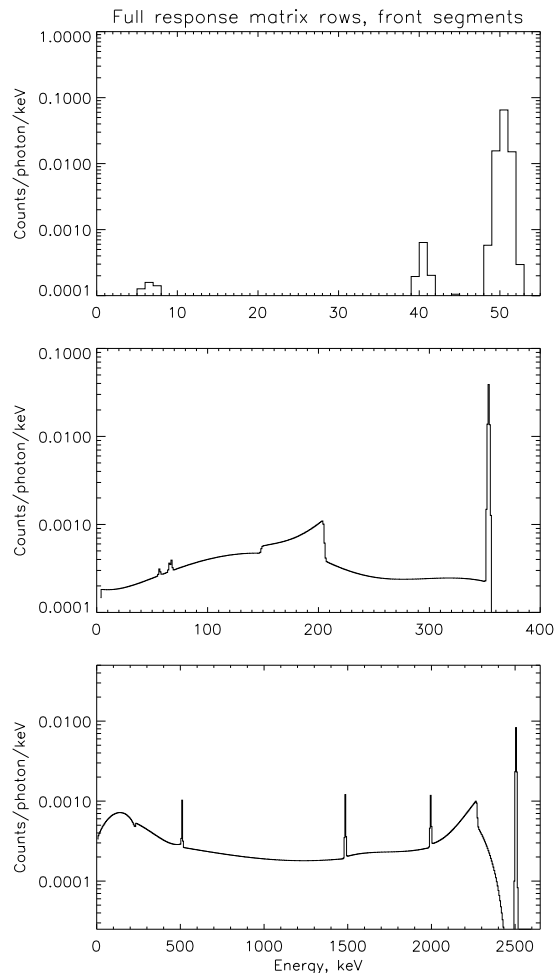


Figure 14: Full instrument response (including scattering) for three energies: 50, 350, and 2500 keV. Top: 50 keV input photons. Note the Ge K-shell escape peak at 40 keV. Middle: 350 keV input photons. Note the weak K-shell fluorescence lines between 40 and 60 keV from nearby passive material, including the tungsten grids. Bottom: 2500 keV input photons. Note the positron annihilation line at 511 keV and the lines corresponding to single and double escape of 511 keV photons.

is near the right edge of the plot. At 2500 keV, the photons can undergo pair production. Thus three narrow lines are visible in the spectrum: the 511 keV line from photons that pair produce outside the detector and then send an annihilation photon in, and lines at $2000-511$ and $2000-2*511$ keV from photons which pair-produce in the detector and have one or more 511 keV photons escape (“escape peaks”). The varied shapes of the scattered continuum are a function of the behavior of the Klein-Nishina formula and the distribution of passive material around the detectors.

These curves are produced by fitting analytic functions to Monte Carlo simulations of photon transport in the spacecraft. The Monte Carlo simulations are done at a few discrete energies, and then the analytic functions are used to interpolate to arbitrary energies. The Monte Carlo simulations are calibrated to measurements of the response of the instrument to radioactive sources in the laboratory.

Figure 15 shows a full response matrix in greyscale. A vertical slice through this plot gives the spectral response to photons of a single energy (i.e. one of the traces in Figure 14).

The various energy-dependent contributors to instru-

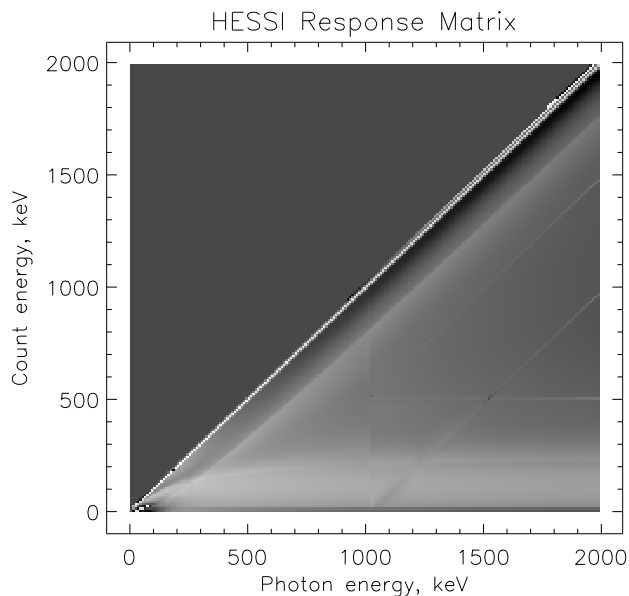


Figure 15: Full response matrix, showing the photopeak (diagonal line) Compton edge (the ridge below the diagonal line), and backscatter peak (nearly horizontal ridge near 200 keV).

mental response are listed in Figure 10. The response matrix software calculates the full response (Figure 15) by calculating each of these “submatrices” separately and then either multiplying them together (for the diagonal or photopeak parts) or adding them in (for the off-diagonal or scattering parts). This is done by the master response-matrix building routine, `hessi_build_srm.pro` (q.v.). When some of the submatrices are time-consuming to build and can be reused (for instance attenuator transmission will never change throughout the mission), they can be separately stored to disk and `hessi_build_srm` can be instructed to read them from disk instead of regenerating them. The submatrices which affect photopeak efficiency are:

- **Grid transmission:** this cannot be simulated, because the transmission of a grid pair is a strong function of the exact shape, regularity and alignment of the slats. The grid transmission submatrix is derived directly from laboratory measurements at the x-ray and optical grid calibration facilities at GSFC. It uses transmission averaged over a half rotation (which is equivalent to a full rotation by symmetry).
- **Attenuator transmission:** The attenuator state (none, one, or both in place) is of course a necessary input to the response matrix software. The data analysis software should automatically find this for you in the data when it is building a response matrix for you. You should not attempt to take spectra in a long time interval containing more than one attenuator state unless your only interest is in high energies (well above 100 keV) where the attenuator has little effect.
- **Blanket transmission:** There are layers of thermal blanketing above and below the imager and inside the spectrometer. They are virtually transparent above 10 keV but are important for the lowest energies.
- **Detector efficiency:** Photons which enter the detector with their full energy can, at high energies, go right through, or else leave less than their full energy (for a photopeak efficiency correction these are equivalent). Most frequently the photon will Compton-scatter out, but annihilation photons, germanium K-shell fluorescence photons, or even Compton- or photo-electrons can also escape. Absorption in the Be windows in front of the detectors (a total of 40 mils) is included here.

- **Electronic threshold cutoffs:** The slow LLD must trigger in the front segment in order for an event to be produced, and in the rear segment both the slow and fast LLDs must trigger (see above). These requirements produce a low-energy cutoff in the spectrum just as the blankets and attenuator do. Because of the finite resolution of the electronics, these cutoffs are not sharp and the loss of events they produce must be compensated for just like any other absorption in order to get the proper spectral shape below 10 keV (fronts) or 30 keV (rears).

The effects that produce off-diagonal (non-photopeak counts) are:

- **Scattering in the detector:** This is actually part of the same submatrix as **Detector efficiency** above, and consists of those photons which enter the detector with their full energy but don't leave it all.
- **Scattering in the spacecraft:** For this component we simulate photons on all paths *except* those which could possibly interact with the detectors first. Thus there can be no photopeak events, since any photon which hits the detectors must have Compton scattered first. Many of these photons are stopped by the graded-Z shielding on top of the cryostat, but they are still significant, particularly in the rear segments at relatively low energies (below 200 keV). There is relatively little contribution above a few hundred keV due to the nature of the Compton formula: since most of the photons have undergone a large-angle scatter before hitting the detector, they have usually lost the majority of their energy.
- **Scattering in the grids and attenuator:** Photons which scatter in the passive material above the detectors *can* contribute to energies below but near the photopeak, since they can hit the detector after a small-angle scatter. In addition, this submatrix includes the fluorescence photons from the grids. We are only simulating scattering of photons in the lower grid, since a photon which scatters in the upper grid has a relatively low chance of hitting a detector due to the large distance.
- **Scattering in the Earth's atmosphere:** Since the rear detectors have only a little aluminum between them and space on the sides (and only somewhat thicker aluminum behind them), many

flare photons Compton-scattering off the atmosphere will be counted (these are sometimes called albedo photons). In fact, we expect these photons to dominate the rear segment count rate in X-class flares (for weaker flares, it is dominated by background). They will be even more likely than spacecraft-scatter photons to have gone through large-angle scatters and multiple scatters, so they will be soft (mostly < 100 keV). We simulate this effect by a two-part Monte Carlo. We first bounce simulated flare photons off a simulated spherical atmosphere, recording what spectrum comes from it as a function of position around the Earth and angle from the zenith. We then convolve this matrix with another, the output of a simulation of the spectrometer's total response to photons at various energies and various (large) angles from the solar direction. In order to use the resulting submatrix properly, you must provide the Sun-Earth-RHESSI angle (or a distribution of angles).

The finite energy resolution (line shape) is also an effect that moves counts away from the energy of the input photon. For continuum work its effect is negligible, since it only moves energies a small amount. For work with narrow lines, it will be useful to know the best current guess at the instrumental line shape, so as to see whether a solar line is genuinely broadened. The line shape is a convolution of two submatrices: the intrinsic lineshape from the detector and the noise introduced by the electronics. The electronic noise is independent of line energy, usually symmetrical, and often Gaussian. It can be independently measured by turning on the pulser occasionally (see above). The intrinsic noise has a Gaussian component proportional to the square root of energy (due to counting statistics of the electron-hole pairs) and a trapping component, linearly proportional to energy. Both the intrinsic and electronic lineshapes may change so fast that the database for the spectroscopy software cannot always keep up. To be sure you understand the instrumental lineshape, you can look at the shape of background lines during non-flare times. Most background lines are narrow, although the 511 keV line from the Earth's atmosphere has an intrinsic width of about 3 keV.