

THE RHESSI SPECTROMETER

D. M. SMITH¹, R. P. LIN¹, P. TURIN¹, D. W. CURTIS¹, J. H. PRIMBSCH¹,
R. D. CAMPBELL¹, R. ABIAD¹, P. SCHROEDER¹, C. P. CORK², E. L. HULL²,
D. A. LANDIS², N. W. MADDEN², D. MALONE², R. H. PEHL², T. RAUDORF³,
P. SANGSINGKEOW³, R. BOYLE⁴, I. S. BANKS⁴, K. SHIREY⁴ and
RICHARD SCHWARTZ⁵

¹*Space Sciences Laboratory, University of California, Berkeley, U.S.A.*

²*Lawrence Berkeley National Laboratory, U.S.A.*

³*ORTEC, U.S.A.*

⁴*NASA Goddard Space Flight Center, U.S.A.*

⁵*NASA Goddard Space Flight Center/SSAI, U.S.A.*

(Received 12 September 2002; accepted 16 September 2002)

Abstract. RHESSI observes solar photons over three orders of magnitude in energy (3 keV to 17 MeV) with a single instrument: a set of nine cryogenically cooled coaxial germanium detectors. With their extremely high energy resolution, RHESSI can resolve the line shape of every known solar gamma-ray line except the neutron capture line at 2.223 MeV. High resolution also allows clean separation of thermal and non-thermal hard X-rays and the accurate measurement of even extremely steep power-law spectra. Detector segmentation, fast signal processing, and two sets of movable attenuators allow RHESSI to make high-quality spectra and images of flares across seven orders of magnitude in intensity. Here we describe the configuration and operation of the RHESSI spectrometer, show early results on in-flight performance, and discuss the principles of spectroscopic data analysis used by the RHESSI software.

1. Introduction

The Reuven Ramaty High-Energy Solar Spectroscopic Imager (RHESSI) was designed to study high-energy emission from flares over a broad energy range, from thermal X-rays through nuclear gamma-ray lines (Lin *et al.*, 2002). The design of its spectrometer was driven by the need to satisfy many requirements with a single instrument of modest cost and weight. The spectrometer had to be able to process up to millions of hard X-rays per second for imaging, while carefully collecting rare gamma-ray line photons with high efficiency, high energy resolution, and without interference from the X-rays. It was desirable for the energy range to extend down to 3 or 4 keV to clearly image the thermal components of flares and to be highly sensitive to microflares, while extending well above the 4.4 and 6.1 MeV gamma-ray lines to characterize the electron bremsstrahlung spectrum above the nuclear line region. We wanted to be sensitive to events covering at least seven orders of magnitude in intensity, from microflares to the largest X-class events.



Most importantly, RHESSI had to have extremely high energy resolution in order to make advances in spectroscopy comparable to the advances its high-angular-resolution Rotating Modulation Collimator (RMC) system would make in imaging (Hurford *et al.*, 2002). The Solar Maximum Mission Gamma-Ray Spectrometer (SMM/GRS), with its large collecting area and its long and successful mission operations, produced a wealth of data on gamma-ray lines from large flares (Vestrand *et al.*, 1999) and a significant contribution was also made by the instruments on the Compton Gamma-Ray Observatory (Share, Murphy, and Ryan, 1997) and the Gamma-Ray Spectrometer on *Yohkoh* (Yoshimori *et al.*, 1994). But all these instruments were scintillators with energy resolution $\geq 6\%$ full width at half maximum (FWHM) at gamma-ray line energies. In some cases (particularly around 1 MeV) this low resolution can result in the blending of several lines, and, more importantly, it leaves the issue of Doppler shifts and broadening of the nuclear deexcitation lines mostly unaddressed. Energy resolution an order of magnitude finer would allow these Doppler phenomena to be observed unambiguously, revealing information about the angular distribution of accelerated ions in large flares. In the hard X-ray band, high resolution (about 1 keV FWHM from 3–100 keV) would allow a clean separation between thermal and nonthermal emission never before achieved except on balloon payloads flying germanium spectrometers (Lin and Schwartz, 1987).

RHESSI has many new, unique features compared to other gamma-ray spectrometers flown in space, even beyond the obvious advance of using a large array of germanium detectors. These include an efficient, lightweight, and inexpensive mechanical cooler; movable attenuators that respond to the intensity of a flare, so that the detectors can record microflares without saturating even in the biggest flares; and extremely low-noise electronics and segmented detectors to cover from 3 keV to 17 MeV with a single crystal. These and other advances allow RHESSI to meet its goals of resolution, sensitivity, and dynamic range within the cost and weight constraints of a Small Explorer mission.

The sections below will address the composition and operation of the RHESSI spectrometer, the spectrometer data formats, in-flight performance, and the principles of the spectroscopic data analysis.

2. Components of the RHESSI Spectrometer

2.1. GERMANIUM DETECTORS

RHESSI's spectrometer is an array of segmented coaxial germanium detectors. When ultrapure 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 free pairs. If there is a high electric field (on the order of 1000 V cm^{-1}) across the crystal, the electrons and holes will be pulled to each electrode, creating a current




Figure 1. Cross-sections of a RHESSI detector. (A) A detector profile with field lines, with the field line marking the segment boundary in *bold dashes*. (B) 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.

pulse that can be amplified and digitized by suitable electronics. The total charge in the current pulse is proportional to the photon energy.

Figure 1 shows two cross-sections of the cylindrically-symmetrical RHESSI detector design. This design was a joint effort of the RHESSI co-investigators at U. C. Berkeley and Lawrence Berkeley National Laboratory and the manufacturer, ORTEC (currently a division of AMETEK). The shape is a variation of a 'closed-end coaxial' detector, the industry standard design for large volumes and high gamma-ray sensitivity. The ultrapure, slightly n-type germanium material is doped in a very thin outer layer with 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. When 2000–4000 V is applied between the inner and outer electrodes, the crystal is depleted of free charge carriers, with enough electric field in the crystal from the space charge and external voltage combined to cause the electron-hole pairs to reach terminal velocity.

For space applications, it is important to fly n-type material with this electrode configuration in order to minimize the effect of radiation damage on resolution.

Since radiation damage produces primarily hole traps rather than electron traps, the goal is to minimize the length of germanium traversed by the holes. For a coaxial detector, most of the volume is near the outer radius. Thus most interactions take place near the outside and the holes should therefore travel to the outside electrode. The effects of radiation damage can also be minimized by keeping the detectors very cold (72–76 K in the case of RHESSI) and making sure the detector high voltage is never turned off once there has been significant exposure to protons and neutrons (Hull, 1998; Koenen *et al.*, 1995).

The inner (lithium) contact is discontinuous at a point near the top, and signals are extracted separately from the two halves of this electrode. The line extending from this point to the outside edge of the detector in Figure 1(A) 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 crystal becomes a zero-gap 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 fine spectroscopy can be done without high deadtime from the X-rays.

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 that originates at the break in the inner contact 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. A ring of ‘graded-Z’ material keeps the shoulder of the rear segment from being swamped with flare hard X-rays (see below). 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.

Figure 2 shows a side view of the detector segmentation produced by scanning a 3 mm beam of 81-keV X-rays across the detector in discrete positions corresponding to each pixel. The brightness of each pixel is the sum of the front and rear segment count rates in the photopeak of the line. The segmentation boundary probably looks dark due to a combination of effects which cause charge to be split between the front and rear signal electrodes: Compton scattering and k-shell photon transport, which produce interactions on either side of the boundary, and charge-sharing between the segments for even single-site interactions that happen very close to the boundary.

Note that the segmentation boundary is not independent of azimuth angle, but rather has a wavy appearance, with four peaks and four valleys around the periphery of the detector. This is a natural result of the cubic crystal structure of germanium: electron mobility is highest along the three main crystal axes and lower between them. Thus electrons traveling to the central bore from different spots on the crystal surface at terminal velocity will not travel directly along field lines but instead their paths will bend toward the nearest axis. This distorts the segmentation boundary position as a function of azimuth about the cylindrical axis. The crystals




Figure 2. Result of a narrow-beam 81 keV X-ray scan of the side of a RHESSI detector showing the ripple in the segmentation boundary (see text). The front of the detector is to the left. The scan does not proceed all the way down the rear segment. Lighter shades represent a higher count rate in the 81 keV line.

are always grown so that one axis is parallel to the axis of the cylinder ('X' in Figure 2).

2.2. DETECTOR ELECTRONICS

The preamplifiers for the RHESSI detectors feature custom four-terminal, N channel junction field effect transistors (JFETS) with adjustable back-gate voltage for very-low-noise operation, one for each segment. These are immediately behind the detectors, and are thermally connected by wires to the outer thermal shield (see below) to achieve an ideal operating temperature between 130 and 170 K. Their signals are taken out to the rest of the preamplifier circuit, which resides in an individual box for each detector, clustered around the bottom of the spectrometer (see Figure 3), by a harness consisting of thin traces of manganin (for low thermal conductivity) layered in a flexible plastic film.

The preamplifiers use a pulsed reset circuit developed at Lawrence Berkeley National Laboratory (Landis, Cork, and Goulding, 1982). The integrated charge on the preamplifier's feedback capacitor is restored by this circuit when a total charge equivalent to interactions from 40 MeV worth of photons or cosmic rays has been deposited in each segment. The preamplifier's output-level-sensing and charge-restoration logic avoids transient signal reductions associated with high-energy events.




Figure 3. The RHESSI spectrometer, before being mounted on the spacecraft. The aluminum attenuator disks are visible on top (one set in the aperture, the other out), and the preamplifiers and high-voltage filter boxes can be seen around the skirt. The detectors are behind the thin, ribbed side wall.

Leakage current in the detectors is so low (a few picoamps) that cosmic rays and background photons dominate the total current through the detectors. Each reset event results in approximately $40 \mu\text{s}$ of deadtime in the resetting segment, but as this seldom happens at a rate of more than 100 Hz it represents a small contribution to the overall system deadtime. Leakage current is kept low by avoiding the introduction of any contamination on the back (insulating) surface of the crystals in two ways: (1) careful design of the cryostat (see below) and (2) a proprietary ORTEC passivation process used on the rear germanium surface.

The signals from the preamplifier are taken to a Detector Interface Board (DIB) in the primary Instrument Data Processing Unit (IDPU) (Curtis *et al.*, 2002) box on the spacecraft deck. Each board contains both the analog and digital signal-processing circuitry for both segments of one detector. The signal from each event is split, going to a fast-shaping channel (which produces a triangular pulse 800 ns wide at the base) and a slow-shaping channel ($8 \mu\text{s}$ peaking time) for spectroscopy.

The fast pulse is used for pileup rejection: two events very close together in time, which would be analyzed as a single higher-energy event by the analog-to-digital converter (ADC), can instead be recognized, and both rejected. If the second event comes shortly after the slow-shaped pulse from the first has already been

sampled at its peak and fed to the ADC, then only the second is rejected, since the first was read with its correct energy. The fast lower-level discriminator is set at about 7 keV in the front segments (just above the noise) and about 20 keV in the rear segments. Photons below 7 keV in the front segments may pile up with other counts without triggering the pileup veto, so careful data analysis is necessary for high count rates of soft photons.

The circuit also uses the height of the fast pulse for large events to choose quickly between one of two gain ranges available in the rear segments: a high gain setting that goes up to about 2.7 MeV and a low one that goes to about 17 MeV. The ADC samples the shaped pulse (see below) at its peak to produce a 13-bit (8192 channel) value for either range. Thus the ADC gain can be set before the slow-shaped pulse reaches its peak. Because the energy resolution of the fast pulse is poor and because its short integration time means that each pulse still has a different shape depending on the position of the interactions in the detector, the transition between the two rear gain ranges is gradual rather than abrupt, i.e., events from about 2.6 to 3.2 MeV may appear with either gain. Since the gain range that was used is recorded as part of the event, there is no loss of information. The front segments use only the first (high) gain value.

The bottoms of the usable energy ranges in each segment are set by both electronics and physics. The physical limits are absorption in the materials around the detectors (see below). The electronic thresholds, set by the noise level in the slow-shaping channel, are about 2.7 keV for most of the front segments and 20 keV in the rears. Both the physical and electronic cutoffs are somewhat gradual and have to be modeled precisely for low-energy spectral analysis. The front segment electronics use quasi-triangular shaping to optimize low-energy response, and the electronics for the rear segments use quasi-trapezoidal shaping to minimize ballistic deficit.

The ADC-busy and pileup rejection signals are ‘or-ed’ together and sampled at 1 MHz by a livetime counter. Because the pileup circuit has the ability to sometimes veto *both* piled-up events, no simple circuit can sample the true system livetime. There is, however, a simple mapping between the recorded and true livetimes as long as all events are above the fast discriminator threshold. When a large fraction of the counts are below this energy, deriving the true livetime is more complicated, but can still be done based on pre-flight and in-flight calibrations.

2.3. THE RHESSI CRYOSTAT

The RHESSI cryostat was designed to provide a lightweight but secure environment for the detectors. This involves maintaining a hard vacuum and very good thermal insulation, so that the detectors can remain at operating temperature (about 75 K) with a very low heat leak. Figures 3–5 depict the interior and exterior of the cryostat and the arrangement of the detectors.

The weight constraints of a Small Explorer did not allow a heavy active shield to veto solar photons that scatter out of the detectors and to keep out background




Figure 4. Exterior of the spectrometer shown with part of the outer housing cut away, revealing the detector modules inside. The numbering scheme of the detectors matches the grids.




Figure 5. The interior of the RHESSI spectrometer, showing the detectors mounted on the coldplate, the Sunpower cryocooler below it, and the fiberglass-strap 'cartwheels' that suspend this assembly in the interior of the cryostat.

photons. We therefore decided instead to design the side walls of the cryostat to be as thin as possible; there is about 4 mm of Al between the rear detector side surfaces and space. This allows hard X-rays (above about 25 keV) and gamma-rays from cosmic sources such as supernova remnants, pulsars, and gamma-ray bursters to enter the spectrometer, providing a wealth of secondary science, including astrophysical topics (Smith *et al.*, 2000, 2002) and the study of emissions from the Earth's atmosphere (Share *et al.*, 2002). Cosmic and terrestrial sources usually appear primarily in the rear segments rather than the front segments, so they can generally be distinguished from solar events.

The detectors are securely housed in independent modules that are not hermetically sealed but have a very small indirect gas conductance path to the outside to minimize contamination during handling. The modules are firmly attached to a common coldplate, and the detector/coldplate assembly is suspended by three sets of fiberglass straps for thermal isolation (see Figure 5). The coldplate assembly is surrounded by two thin aluminum thermal shields at intermediate temperatures, each of which is surrounded by multilayer aluminized-mylar insulation.

The cryocooler is a Sunpower model M77 Stirling-cycle refrigerator, a small, high-efficiency unit modified for long-term use in orbit by the cryogenics group at NASA's Goddard Space Flight Center. The cooler is visible in Figure 5 as the horizontal cylindrical object below the coldplate. It is equipped with an actively-driven balancer mass to cancel vibrations to first order. Its 'gas bearing' system means that there is no rubbing of solid parts during operation. With 55 W of input power, the cooler can provide the approximately 3.5 W of cooling needed to keep the entire coldplate/detector assembly at 75 K. The cooler is coupled to the coldplate by a flexible, S-shaped link made from 700 sheets of 0.025 mm aluminum foil for good heat conduction with minimum transmission of vibration. Also in the thermal path between the cooler's coldfinger and the coldplate is a rod of sapphire. This material has high thermal conductivity at low temperatures and vice-versa, and is important during cooldown and annealing (see below). As part of the thermal isolation, the outer thermal shield is actively cooled by its own separate connection partway down the coldfinger of the cooler, which keeps it at about 155 K. Waste heat from the cooler is radiated to space from the back surface of the spectrometer, which is also the rearmost surface of the spacecraft. It is covered with silver teflon tape. It alternately views deep space and the Earth during each orbit.

2.4. THE ATTENUATORS OR SHUTTERS

Attached to the top of the spectrometer are two lightweight, movable frames, each of which carries nine aluminum disks that can be moved in front of the detectors (see Figure 3). These disks, manufactured to our design by Tecomet, Inc., serve as attenuators, to keep the detectors from being saturated at high counting rates. One set of disks is thicker than the other, but they are not uniform: each has a small, thin spot in the center so that there is always some low-energy response. There is




Figure 6. Top: photographs of the attenuator disks. Left: thin disk. Right: thick disk. Bottom: cross-section of the thin and thick disks. The vertical scale has been exaggerated by a factor of four to make the features easier to see. The outer diameters are 61.5 mm to cover the detector front segments. Thicknesses vary from 0.05 mm to 1.53 mm.

also a slightly larger region that is thick on the otherwise thin disk and thin on the otherwise thick disk: thus, the full attenuation doesn't occur until both disks are in place (see Figure 6). The attenuator dimensions vary slightly from detector to detector based on the grid-slat spacing in order to avoid introducing biases into the imaging when the attenuators are put in.

The onboard computer monitors the front segment livetimes and puts the disks in (thin first) when deadtime gets too high. The state with the thick shutter in and the thin shutter out is not currently used in spacecraft operations. Because the system is particularly sensitive to pileup at low energies (see above), and because low-energy thermal emission dominates the detected flare photons when both shutters are out, the thin shutter is currently set to come in when deadtime exceeds 8%. The deadtime will then drop sharply. If the flare brightens further, the thick shutter will come in when the deadtime reaches 10%. Early in the mission the shutters were only brought in at higher deadtimes, until it was realized that pileup was too severe in the resulting data. Because bringing in each set of shutters reduces the deadtime so effectively, it is difficult for the onboard software to estimate when to pull them out again. As long as the deadtime is very low, the flight software will try briefly pulling out the last shutter it put in every 4 min until either (1) the deadtime is found to be low enough to leave it out or (2) it exceeds a set maximum number of shutter motions per spacecraft orbit.

Each shutter frame is moved by long, thin shape-memory alloy (SMA) wires that contract when a voltage is applied and current runs through them, producing

heat. There are four sets of six wires each: one set for moving each frame in each direction. Each motion takes less than one second, including the delay for heating, and about 20 s is required for cooldown before the same set of wires can be activated again. In addition, each frame is connected to a heavy SMA spring for an emergency backup that can pull it strongly into its default position (thin in, thick out) and lock it there. These are only to be used in the unlikely event of a mechanical jam or other failure of the wire SMA system. A second set of SMA springs was used to uncage each frame after launch.

With neither attenuator in place, the front segments view the Sun through four Be windows and four blankets of multilayer aluminized-mylar insulation (two inside the cryostat and one on each end of the imager tube). The blankets are thinned but not eliminated immediately above each detector to improve the transmission of solar photons. The net transmission is 10% at about 5.5 keV (see Figure 8 below). The attenuator system adds about four orders of magnitude to the dynamic range of microflares and flares detectable by RHESSI. In order to save onboard memory and downlink capacity, the thin shutter is often left in place for one or more days at a time.

2.5. OTHER SPECTROMETER SUBSYSTEMS

2.5.1. *Charged Particle Detector*

A solid-state silicon charged particle detector (CPD) is mounted on a strut that runs past the side of the spectrometer. The detector is a disk 1.0 mm thick and 5.64 mm in diameter, and is mounted so that it faces outward, its normal axis perpendicular to the spacecraft axis. It is encased in shielding consisting of at least 3.8 mm of Al in all directions except for an outward-facing pinhole of diameter 1.02 mm. This detector is not configured as a spectrometer, but count rates are read out every $\frac{1}{8}$ s above two energy thresholds: 50 keV and 620 keV. The higher threshold is commandable and the lower fixed. The commandable threshold may be set to any energy from about 20 keV (limited by noise) to 2 MeV.

Cosmic rays penetrate the shielding easily and usually trigger both thresholds. Thus the CPD can measure cosmic-ray flux, although the upper level discriminator (ULD) counter in the germanium detectors themselves provides a more accurate measure due to their much larger area. Electron precipitation events generally trigger the lower threshold only. During transit of the SAA, the upper threshold counts mostly trapped protons and the lower mostly trapped electrons, although the upper threshold is partially contaminated by SAA electrons, which are much more numerous, entering the pinhole. Because the detector rotates with the spacecraft it samples the pitch-angle distribution of the particles whenever the spacecraft axis is near perpendicular to the local magnetic field. At these times, the SAA particles are seen to have an extreme pancake distribution (circling at 90° to the field) as expected for mirroring particles.

2.5.2. *Beryllium Scatterer*

Since the rear segments see no direct flare photons below 100 keV, we can use them at these low energies as a hard X-ray polarimeter (McConnell *et al.*, 2002). There is a cylinder of beryllium 3 cm in diameter and 3.5 cm long attached to the coldplate, nestled among the rear segments near the center of the spectrometer. 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 Compton cross-section, differential in azimuth angle, is a function of the angle from the polarization axis. Thus, by watching the relative counting rates in these rear segments, we will measure the direction and degree of polarization for incoming photons of roughly 20–100 keV. Simulations (McConnell *et al.*, 2002) 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 rear detectors and which also vary with the spacecraft spin, albeit with one peak per spin instead of two. Since this component will not be a perfect sinusoid, any first-harmonic power it has could be mistaken for a polarization signal. Scattering in the coldplate and other passive materials into the rear detectors does not modulate with the spacecraft spin, and so it only interferes with polarization measurements by increasing the background count rate.

2.5.3. *Graded-Z Shielding*

In the absence of a heavy active shield, there were several places where a small amount of passive shielding to block hard X-rays was considered valuable. In each place, the material is a 'graded-Z' laminate of metals: from outside to inside, 0.5 mm of tantalum, 1.0 mm of tin, and 0.5 mm of stainless steel. The tantalum stops most incident photons up to 100 keV. Each successive metal absorbs the K-shell fluorescence photons of the one before it. The iron K-shell photons from the steel are absorbed in the aluminum structures backing the shielding, and the aluminum K-shell photons are both too low in energy to trigger the detectors and mostly stopped in the aluminum itself. The combination is mostly opaque below 100 keV, and fairly transparent above.

A ring of this material surrounds the side surface of the front segments to shield them from background photons from the Earth and space, so that they can be sensitive to the faintest possible microflares. A washer-shaped ring sits above the shoulder of each of the rear segments to protect them from direct solar photons. Both these rings are visible in Figure 1(B). Finally, a larger sheet of the material, with holes for each detector and the Be scatterer, is attached to the top cryostat wall, below the attenuator assemblies. Its purpose is to minimize the number of solar hard X-rays which scatter in the coldplate and other cryostat materials and enter the rear segments at low energies. This serves two functions: to keep the counting rates in the rear segments from saturating during the largest flares, and to

reduce the background against which polarization measurements will be made in the rear segments at low energies.

2.5.4. *Onboard Radioactive Source*

RHESSI carries a tiny onboard radioactive source (5 nanocuries of ^{137}Cs) that makes a line at 662 keV, far from any line expected to occur in flares or in RHESSI's variable background. This source is so weak that it only gives a useful number of counts in spectra accumulated over many hours. Having a line of known intensity lets us monitor loss of narrow-line efficiency in the detectors due to radiation damage, which happens on a time scale of months to years. 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 from before launch since there was always some of the isotope in the laboratory.

2.5.5. *Onboard Pulser*

The IDPU can put regular, small pulses onto the detectors' high voltage (HV) line. The electronics see this 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.

2.6. CRYOSTAT THERMAL OPERATIONS

Cooldown of the spectrometer began within a few hours of launch and lasted about six days. Although the cryostat was kept on a vacuum pump until being sealed off shortly before launch, continued outgassing from the interior, particularly water from the mylar thermal blankets, remained a concern. The thermal impedance of the sapphire rod between the coldplate and the cryocooler was included to force the first stage of cooldown of the detectors to proceed extremely slowly. The secondary connection between the cryocooler and the intermediate thermal shield has no such impedance, so the shield cools much more quickly. This allows all the water vapor and other volatiles in the cryostat vacuum to freeze out onto the shield first instead of the detectors, so that we do not need to use detectors that are individually hermetically sealed. At about 200 K the cooling curves cross and the detectors become colder than the shield from then until they reached their operating temperature of around 75 K.

Once the radiation damage that has built up in the detectors is significantly affecting RHESSI's energy resolution (which we predict will occur 2–4 years after launch), we can anneal the detectors at high temperature to remove most of the effects of the damage. Although some damage is due to cosmic rays and the secondary neutrons they generate, most of it will be caused by trapped protons encountered when the spacecraft passes through the SAA. When the anneal cycle

begins, the detectors will be powered down and then a one-time-use SMA actuator will be fired to open a valve in the cryostat's back surface to the vacuum of space so that gases released during the warmup can escape. Then, diode heaters on the coldplate will be powered to begin heating it up. The cryocooler will be run at a low power level during the early phases of the warmup to keep the intermediate shield cold, so that it doesn't release any deposited water onto the detector surfaces. The detectors will be held near 100 °C for several days before being cooled down again to operating temperature.

3. Data Formats

3.1. EVENT DATA

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 is 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) gain setting of the event,
- 13 bits 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 that are a partial measurement of the livetime counter (Curtis *et al.*, 2002).

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 the remaining memory and the position of the attenuators. 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. Operating modes that have been used in flight have decimated the rear segments (with $N = 3$ or 4 and $E = 150$ keV to 450 keV) either during spacecraft night or at times when RHESSI is at high magnetic latitudes and electron precipitation is common (see below). For most of the mission, front segment events have been turned off during spacecraft night, and all events are always turned off during passage through the SAA.

3.2. MONITOR RATES

Monitor rates are counters that are read out every second for each of the 18 segments. They include: number of resets, number of events triggering the slow-channel lower-level discriminator (LLD), number of events triggering the fast-channel LLD, number of high-energy (ULD) events beyond the scale of the ADC

(usually cosmic rays), and livetime. This information is normally used to check the health of the detectors and is not necessary for spectroscopy.

3.3. FAST RATES

Fast rate data are produced only when the count rates are very high – they are the last line of defense against detector saturation for flares big enough to overwhelm the normal electronics chain with both attenuators in place. The data are count rates in four broad energy bands, roughly 7–15 keV, 15–25 keV, 25–70 keV, and 70 keV and higher, with significant variation from detector to detector. The pulses are sampled from the fast electronics chain. The rates for the three detectors under the finest grids (and therefore the fastest imaging modulations) are sampled at 16 kHz; the next three at 4 kHz; and the three coarsest grids 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.

4. RHESSI's Performance in Flight

4.1. RESOLUTION

At low energies (below about 200 keV in the RHESSI front segments), the width of a spectral line in a germanium detector is dominated by noise in the electronics, and is roughly constant. At higher energies, one of two noise components will dominate: the counting statistics of the electron-hole pairs, which increases as the square root of energy, and broadening due to trapping of electrons or holes in the crystal, which goes linearly with energy. Electron trapping is intrinsic to the detectors, due to impurities when they are grown. Hole trapping is the result of radiation damage in space.

Eight of the nine RHESSI detectors have been performing well since the start of the mission, with energy resolution of about 1 keV FWHM and a low-energy cutoff of about 2.7 keV in the front segments and with resolution of better than 3 keV FWHM in the rear segments to over 1 MeV (see Table I). Figure 7 shows the 2.223 MeV line from neutron capture on hydrogen seen by RHESSI from the X4.8 flare of 23 July 2002. This is a sum of gain-corrected spectra from the rear segments of all the detectors except the one below grid No. 2 ('G2'; see below). The FWHM is 4.4 keV. The line is expected to be intrinsically \ll 1 keV because the neutrons at the Sun thermalize before being captured, so this is a good measurement of the current resolution capability of the instrument at this energy. The slight asymmetry of the line is probably due to a small amount of trapping. If interpreted as due to radiation damage, it is similar to our prediction for this phase of the mission (six months after launch). The resolution of the summed rear segments at low energies (i.e., the 93.3 keV background line) is 2.5 keV FWHM, which is the contribution of the electronics plus any blurring due to imperfect gain calibration.

TABLE I

RHESSI in-flight resolution by detector segment. Resolution is measured at 93.9 keV in the front segments and 1117 keV in the rears.

Detector	Front FWHM (keV)	Rear FWHM (keV)
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




Figure 7. The 2.223 MeV neutron-capture line from the X4.8 solar flare of 23 July 2002, as observed by the RHESSI rear segments.

The gain calibration of each segment is currently re-calculated about once a week, using a full orbit's worth of background data. The entire channel-to-energy conversion is represented as a purely linear fit over each energy scale (0.0–2.7 MeV for the front segments and 0.0–2.7 and 2.7–17 MeV in the rear segments). The gain and offset of the linear relation are currently determined by finding the centroids of the two brightest background lines, a germanium activation line at 93.3 keV and the positron annihilation line at 511 keV. The high degree of linearity of the electronics




Figure 8. RHESSI front and rear segment effective areas for photopeak absorption, summed over all 9 detectors. The front-segment traces peak at lower energy, and from top to bottom represent the case with no attenuator in, with the thin attenuator only in, and with both in. Below the rear-segment trace is the trace representing events which are split between the front and rear segment of one detector. The total of the front (no attenuator), rear, and coincidence modes is also shown.

allows this fit to produce results as good as Figure 7 even though the energy is extrapolated by a factor of four from the higher line used in the gain solution. Gain variations within one orbit and within one day are negligible thanks to the thermal isolation of the detectors and the careful design of the electronics to be temperature-independent. Gain solutions using more background lines and longer integrations are currently being produced and may yield a slight improvement on the already excellent results.

4.2. DYNAMIC RANGE AND EFFECTIVE AREA

With both sets of attenuators out, RHESSI can see microflares in the 3–10 keV range that are too small to be seen by GOES (Krucker *et al.*, 2002). With both sets in, the X4.8 flare of 23 July 2002, at its peak, caused no more than 30% downtime in the front segments. This still requires a pileup correction for good hard X-ray spectroscopy (see below), but not a particularly difficult or nonlinear one. The count rate at the peak of the X4.8 flare was between 20 000 and 25 000 counts s^{-1} in each front segment, close to the maximum throughput. The maximum throughput is about 20% higher when both shutters are out and the flare is dominated by soft photons that don't trigger the pileup suppression. A count rate several times higher can be processed by the fast rate counters.

Figure 8 shows the effective area of the RHESSI array for the front segments alone and the rear segments alone, excluding the effects of the grids. The effect of the attenuators on the front segment effective area is also shown; their effect on the rear segment is negligible. The curve labeled ‘F/R COINC.’ represents events that leave part of their energy in each segment, adding up to the full energy. These events are not yet used by the analysis software, but the adaptation, which requires new ways of dealing with background and gain, is in progress. Such ‘coincidence’ events between segments of *different* detectors present even more complications and add less added effective area, so they will be addressed only after front/rear coincidences have been incorporated.

The grids reduce the effective areas shown by factors strongly dependent on energy and the offset of the flare from the imaging axis, ranging from roughly 4 to 6 below 100 keV. The response of the spectrometer is generated by GEANT3, a high-energy photon and particle transport code that is used to generate the response matrix of the instrument (see below). Before flight, the GEANT3 performance was calibrated by taking spectrometer data in the laboratory using radioisotopes with lines from 3.7 keV to 6.1 MeV, placed in many positions above and around the spectrometer. The laboratory data were compared with GEANT3 simulations of the same configurations, and the detectors’ internal segmentation boundaries in GEANT3 were adjusted until the GEANT3 results matched the data.

Although the effective areas of the front and rear segments in Figure 8 cross at about 250 keV, this is not where the *sensitivity* of the two segments crosses. Because the background in the rear segments is a factor of 10 higher (see below), the front segments are more sensitive in a background-dominated situation up to about 800 keV.

4.3. BACKGROUND

RHESSI is not designed to be a low-background instrument. Heavy shielding was not possible within the Small Explorer weight constraints. A launch into an equatorial orbit was not available, which would have reduced every background component except cosmic diffuse radiation by minimizing cosmic-ray exposure and eliminating exposure to trapped protons in the South Atlantic Anomaly (SAA). Fortunately, solar flares are bright, and most RHESSI flare science will still be done using data that are source-count limited rather than limited by statistical fluctuations in the background.

Figure 9 shows the background in the RHESSI front and rear segments. The front spectrum is shown both with the segment operated on its own and in anticoincidence with the rear segment. The latter mode is similar to ‘phoswich’ detectors made of two layers of scintillator, i.e. the rear segment serves as a closely-placed anticoincidence shield to reject events that Compton scatter into or out of the front segment and charged-particle events. At most energies continuum is dominant, which is typical for an unshielded instrument. The continuum is comparable to




Figure 9. Typical RHESSI background spectra, accumulated over 2.5 hr. The three *solid traces*, top to bottom, are the backgrounds in the rear segments, the front segments, and the front segments using front/rear anticoincidence ('phoswich mode'). The *dotted lines* are pre-flight predictions of the background that included continuum from the cosmic diffuse background and from cosmic rays interacting in the Earth's atmosphere and in the spacecraft.

preflight estimates (dotted lines in Figure 9). In the simulations that produced the estimates, the dominant background component is continuum from the Earth's atmosphere. Below about 100 keV in the rear segments, cosmic diffuse emission dominates. Prompt cosmic-ray interactions with the spacecraft, induced radioactivity of spacecraft materials and β -decays in the detectors also contribute.

Most of the prominent lines in the background spectra are due to the activation of the germanium detectors themselves (e.g., 10, 54, 67, 93, and 198 keV). The brightest higher-energy line, however, is the positron-annihilation line at 511 keV. Much of this comes from the atmosphere, but there is a significant contribution due to the activation of spacecraft materials: immediately after RHESSI exits the South Atlantic Anomaly (SAA), the 511 keV line flux is several times its average value for orbits that don't include the SAA, and the excess decays over the next hour or so.

The three primary sources of variation in the RHESSI background are passes through the SAA (several a day on successive orbits), smooth modulations due to changes in geomagnetic latitude (and therefore cosmic-ray flux) over each orbit, and occasional periods of electron precipitation from the outer radiation belt when the spacecraft is at its highest geomagnetic latitudes (about 40–50°). Brems-

strahlung produced by these electrons in the spacecraft and the Earth's atmosphere can temporarily increase the background rate by more than an order of magnitude at times. These events can be distinguished from solar flares because they appear more strongly in the larger rear segments, whereas flares always produce more counts in the front segments. In addition, the CPD counts electrons at a significant rate in all but the smallest precipitation events.

4.4. INSTRUMENTAL ANOMALIES

4.4.1. *Detector Dropouts*

On orbit, we find that each detector segment can spontaneously shut off for a brief period ranging from tens to hundreds of milliseconds. These dropouts occur up to once every few seconds in the front segments and generally much less often in the rear segments, although this is reversed in detector G5. The cause is still under investigation, but we expect it is a response of the electronics to heavy cosmic-ray nuclei, which leave more energy in the detector than anything available on the ground during pre-flight testing. This phenomenon does not have a significant effect on spectroscopy, since it is energy independent, nor on imaging, since the dropouts do not correlate with the count-rate modulations used to reconstruct images. It does have an effect on photometry and lightcurves, and this is addressed by the data-analysis software, which can recognize the dropouts and compensate for them by adjusting the livetime accordingly. During periods when the count rate is relatively low, dropouts are distinguished from merely quiet moments because they always begin with a preamplifier-reset event.

4.4.2. *Individual Detector Anomalies*

Shortly after its high voltage was turned on, detector G2 suffered from a breakdown either on the rear surface of the crystal or somewhere along its high voltage circuit, such that it can no longer be operated above about 2400 V. At this voltage, the innermost part of the crystal is not depleted, and therefore the contact break on the inner bore (see Figure 1(A)) is not seen: G2 operates as a monolithic detector, with all events anywhere in the crystal going through the front segment electronics. In addition to losing the various advantages of segmentation, this also results in extremely noisy operation: energy resolution of about 10 keV FWHM and a low-energy cutoff closer to 20 keV than the 2.7 keV of most of the front segments. Therefore this detector should not be used for spectroscopy, although it can still be used for imaging from about 20 keV to the energy at which its grids, which are the second-finest of the nine, become transparent.

Detector G8, which is located just a few centimeters from the aft spacecraft antenna, becomes noisy when that antenna is switched on. This can occur several times a day for a few minutes at a time during contacts with the spacecraft. A database of times when this antenna is active is on line, but has not yet been integrated with the data analysis software.

Finally, detector G7 has a slight problem with its front segment signal contact, which causes its energy resolution to be degraded from 1 keV to 3 keV FWHM and its energy threshold to be raised from about the normal 2.7 keV to 7 keV. This behavior has been understood and unchanged since before launch.

4.4.3. *Spectral Artifacts*

A very small fraction of events near or above 3 MeV that should be analyzed with the high-energy (low-gain) range of the rear segments do not trigger the gain shift in the fast channel. These pile up in the top 64 channels (8128–8191) of the high-gain scale, and produce a spectral artifact there. A similar artifact appears in the 64 channels preceding the halfway point of the spectrum (i.e., channels 4032–4095) (Share *et al.*, 2002). Routines to automatically eliminate these channel ranges for each detector will be implemented but are not yet in place.

4.4.4. *Image Events*

When the clouds of electrons and holes liberated by a gamma-ray event move through the detector, they create induced charges on the electrodes of the segment in which they are moving. The change in time of this induced charge is the current pulse that is amplified and integrated by the electronics to become the detected event. However, they also induce charges on the electrodes of the empty segment. The difference is that this image charge reverses sign in the empty segment as the clouds approach the electrode of the segment in which they are actually moving. The result is that the image signals (current versus time) in the empty segment are bipolar in shape and integrate to zero charge.

The RHESSI electronics, however, do not integrate the signals for an infinite time. Therefore a small amount at the very end of each pulse is not counted. Since it is the negative part of the bipolar signal in an empty segment that comes last, the result is that there is a very small, positive residual from the bipolar signal: in other words, a 1 MeV event detected in one segment will create a simultaneous (but false) event in the other segment of a few keV. Fortunately, the timing of an image event is a couple of microseconds different from that of a front/rear coincidence due to Compton scattering, so the image events can be weeded out of the data on the ground without eliminating Compton scattered events, if that is desired.

5. Spectral Data Analysis

Spectral analysis is an inverse problem that begins with a spectrum of counts per spectrometer channel and seeks to recover the spectrum of photons per energy interval initially incident on the spacecraft. Figure 10 is a block diagram of the data analysis process for RHESSI spectra. The RHESSI data analysis software is written to be part of the SolarSoft system (Freeland and Handy, 1996), and, like all the data, is publicly available without a proprietary period. Spectral analysis in par-




Figure 10. Flow chart of a spectroscopy analysis. ‘Advanced’ features are needed only for a few flares, mostly the highest-energy ones.

ticular is accomplished with the SPEX spectral-inversion code (Schwartz, 1996), which has heritage from the Burst and Transient Source Experiment (BATSE) on the Compton Gamma-Ray Observatory. It would be beyond the scope of this paper to begin a tutorial on the use of the software, but it is worthwhile to summarize the steps that are needed for an accurate spectral result. As the software evolves, more and more of this will become automated and transparent to the user.

5.1. GAIN, LIVETIME, AND BACKGROUND

The first stage of analysis is to correct for gain drift and deadtime. Even our preliminary gain analysis is yielding excellent results (see Figure 7) due to the high linearity of the electronics and the narrow, easily identified background lines. The deadtime correction routines start with the livetime counter in the electronics, correct for the double rejection of piled-up events, and compensate for data dropouts.

The next task is to identify and subtract background. At present this is most easily done by selecting data intervals just before and just after the flare and sub-

tracting the spectra in those intervals from the spectrum at the time of interest (whether this is a small part of the flare or its entire duration). When the flare is long, this can be inaccurate, since the background varies significantly over tens of minutes. The results are improved by SPEX's ability to take several such intervals before and after the flare and fit a low-order polynomial to the background variation in each energy channel. At low energies, most flares are bright enough that the background is small compared to the flare flux, and this method of subtraction is suitable.

For a very few flares, more precise methods of background estimation are desirable. These include very long flares and flares with faint hard emission or faint gamma-ray lines. We are developing several additional background-subtraction options, including: (1) using as background the periods exactly 15 orbits (approximately one day) before and after the flare interval, since at these times most of the parameters that control the background had the same value as during the flare; (2) for flares that do not extend much above 50 keV, using the simultaneous background in the rear segments to predict the background in the front segments; and (3) training a neural network to produce a predictive model of the background at any time given a few parameters as input (such as geomagnetic latitude, time since last passage through the SAA, etc.).

5.2. RESPONSE MATRIX

Once gain, livetime and background are accounted for, a series of instrumental effects that together constitute the 'response matrix' of the instrument to flare photons must be removed. There are many effects that modify the input spectrum, including:

- absorption in the mylar blankets, cryostat windows, and grids;
- Compton scattering into and out of the detectors;
- Compton scattering off the Earth's atmosphere, which can dominate the flare count rate in the rear segments below 100 keV;
- noise in the electronics;
- resolution degradation due to radiation damage; and
- the low-energy cutoff imposed by the electronics.

All these effects are accounted for when the analysis software creates a response matrix. Each of these effects contributes either to the diagonal elements of the matrix (i.e. the efficiency of the instrument for detecting photons at their proper energy) or to the off-diagonal elements (moving photons from their true energy to another energy, usually lower), or to both.

When the user is only interested in isolated gamma-ray lines, the response is just the efficiency for photopeak detection, and the conversion from counts to photons is done immediately by dividing the count spectrum by the diagonal terms of the response matrix. This is also adequate for hard X-ray flares with no significant component above 100 keV, since the response of the front segments below this




Figure 11. RHESSI response matrices: photopeak to Compton-continuum ratio versus energy. The continuum contains every simulated photon which leaves less than its full energy but more than 3 keV in the detectors.

point is dominated by complete absorption, not scattering. Figure 11 shows the ratio of photopeak to partial-energy (Compton-scattered) detections in RHESSI versus energy. Since most flare spectra are steeply falling, the measured count rate in any given energy band is usually dominated by photopeak counts up to an energy well beyond 100 keV. One exception is at the lowest energies (below 15 keV) when one or both of the shutters are in. In this case, the true low-energy counts are strongly reduced and the spectrum can be dominated by non-photopeak counts created when a 15–25 keV photon is photoelectrically absorbed in the germanium but the K-shell fluorescence photon (about 10 keV) escapes.

For studying low energies with the shutters in, or for studying continuum and broad-band emission at high energies, the non-diagonal response must be accounted for. This can be done automatically by SPEX using forward-folding. The user specifies a model form for the flare spectrum, which can be a combination of simple functions (power laws, Gaussians, etc.) and physics-based spectral forms (e.g., a set of known nuclear lines from a particular element bombarded by energetic protons, or a thin-target bremsstrahlung spectrum from a monoenergetic electron beam). The software will then multiply this spectrum by the full response matrix, check the goodness of fit to the observed count spectrum, and repeat the process, varying the parameters of the input model until the best fit is found. The output of this process is either the best-fit parameters themselves, or else a spectrum created by multiplying the observed count spectrum by the ratio of the model photon spectrum to the model count spectrum.

Figure 12 shows the modeled RHESSI front-segment response to input photons at three energies: 50, 350, and 2500 keV. At 50 keV, the photopeak dominates but




Figure 12. RHESSI response matrices: sample responses at 50, 350, and 2500 keV (top to bottom).

the K-shell escape peak is visible. At 350 keV, a strong Compton continuum is seen, along with small K-shell fluorescence peaks from nearby passive material, including the tungsten RMC grids. At 2500 keV, pair-production can occur and there are three narrow lines due to positron annihilation (511 keV) and to the escape of one or two positron-annihilation photons from the detector.

In addition to SPEX, we are developing tools to export the RHESSI count spectra and response matrices to the XSPEC package (Arnaud, 1996). XSPEC has a wider variety of built-in spectral models than SPEX and is often used to combine simultaneous spectra from different instruments, which may be useful in the future if RHESSI and INTEGRAL or another instrument observe the same flare. We also plan to develop an algorithm for model-independent inversion of the spectra (Johns and Lin, 1992; Smith *et al.*, 1995). This algorithm will probably be most useful for




Figure 13. 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.

spectra that are dominated by the continuum, not lines, but that extend to high enough energies that simply dividing by the efficiency is insufficiently accurate.

5.3. PILEUP

The pileup correction is the most difficult to implement. It is an off-diagonal effect in that it redistributes counts to different energies, but, unlike Compton scattering, its effect is also non-linear: the contribution of 100 keV photons to the piled-up spectrum depends not only on the number of 100 keV photons coming in, but also on the shape of the rest of the spectrum as well. For this reason it cannot be made part of the response matrix. A full, correct treatment of pileup requires it to be treated as a *separate* stage of forward folding, occurring after the model spectrum has been folded through the response matrix. The (real) background must be added to the forward-folded model flux before pileup is simulated. Furthermore, since the deadtime and degree of pileup will vary rapidly as the grids modulate a bright flare, working with the *average* count rate to estimate pileup is only an approximation. For imaging in particular, future livetime and pileup corrections will have to be fully time-dependent down to the millisecond scale.

These refinements are still far in the future for the RHESSI analysis software, but in the meantime we are implementing an approximate pileup correction that

can be applied early in the analysis flow (see Figure 10) without the need of a spectral model. Figure 13 shows the effect of that pileup correction on a spectrum taken with the thin shutter in and an average of 12.4% deadtime. As we noted above, under the current flight software version the thick shutter would normally be brought in to cut down on pileup at 10% deadtime.

Acknowledgements

The authors would like to thank Gerald Share of the Naval Research Laboratory for providing a valuable external perspective and advice throughout the development process and since launch. We would like to recognize our colleagues at the Space Sciences Laboratory who participated in calibration and testing of the spectrometer: engineers Peter Harvey (also the RHESSI Project Manager) and Jane Hoberman, technician Michelle Galloway, and students Daniel Main, Albert Shih, Lewis Hyatt, Eunyong Cho, Scott Pollack, Mark Chu, and Carmel Levitan. Finally, we would like to thank the technicians, machinists, and others who helped build the RHESSI spectrometer.

The instrument could not have been what it is without the contributions of our colleagues who passed away before RHESSI was launched: Reuven Ramaty and Natalie Mandzhavidze, whose work on gamma-ray line emission in flares inspired its high-resolution design, and Robert D. Campbell, who designed its particle detector and worked tirelessly to characterize and calibrate its germanium detectors. We remember them with gratitude and respect.

References

- Arnaud, K.: 1996, in G. Jacoby and J. Barnes (eds.), *Astronomical Data Analysis Software and Systems V*, Astronomical Society of the Pacific, San Francisco, U.S.A., p. 17.
- Curtis, D. W. *et al.*: 2002, *Solar Phys.*, this volume.
- Freeland, S. L. and Handy, B. N.: 1998, *Solar Phys.* **182**, 497.
- Hull, E. L.: 1998, Ph. D. dissertation, Indiana University, U.S.A.
- Hurford, G. *et al.*: 2002, *Solar Phys.*, this volume.
- Johns, C. M. and Lin, R. P.: 1992, *Solar Phys.* **137**, 121 (Erratum: *Solar Phys.* **142**, 219).
- Koenen, M., Bruckner, J., Korfer, M., Taylor, I. and Wanke, H.: 1995, *IEEE Trans. Nucl. Sci.* **42**, 653.
- Krucker, S., Christe, S., Lin, R.P., Hurford, G.J., and Schwartz, R.A.: 2002, *Solar Phys.*, this volume.
- Landis, D. A., Cork, C. P., and Goulding, F. S.: 1982, *IEEE Trans. Nucl. Sci.* **29**, 1125.
- Lin, R. P. and Schwartz, R.: 1987, *Astrophys. J.* **312**, 462.
- Lin, R. P. *et al.*: 2002, *Solar Phys.*, this volume.
- McConnell, M.L., Ryan, J.M., Smith, D.M., Lin, R.P., and Emslie, A.G.: 2002, *Solar Phys.*, this volume.
- Schwartz, R. A. 1996, 'Compton Gamma Ray Observatory Phase 4 Guest Investigator Program: Solar Flare Hard X-ray Spectroscopy,' Technical Report, NASA Goddard Space Flight Center.

- Share, G. H., Murphy, R. J., and Ryan, J.: 1997, in C. D. Dermer, M. S. Strickman, and J. D. Kurfess (eds.), *Proceedings of the Fourth Compton Symposium*, American Institute of Physics, Woodbury, New York, U.S.A., p. 17.
- Share, G. H. *et al.*: 2002, *Solar Phys.*, this volume.
- Smith, D. M. *et al.*: 1995, *J. Geophys. Res.* **100**, 19675.
- Smith, D. M. *et al.*: 2000, 'Extra-Solar Astrophysics with the High-Energy Solar Spectroscopic Imager (HESSI),' *AIP Conf. Proc., The 5th Compton Symposium*, AIP New York, AIP), pp. 510, 671.
- Smith, D. M. *et al.*: 2002, 'Non-solar Astronomy with the Reuven Ramaty High-Energy Solar Spectroscopic Imager (RHESSI),' *SPIE Conf. Proc.* **4851**, in press.
- Vestrand, W. T. *et al.*: 1999, *Astrophys. J. Suppl.* **120**, 409.
- Yoshimori, M. *et al.*: 1994, *Astrophys. J. Suppl.* **90**, 639.