Signal Processing in the Radiation Assessment Detector for MSL

Stephan Böttcher

Revision 1.5, Date 2011/07/13

Contents

List of Tables

List of Figures

1 Requirements

The RAD instrument shall measure the flux of all known and unknown elementary and composite particles that contribute to the radiation dose exposure on the surface of Mars. In particular:

PLD-150 The RAD instrument shall have the ability to measure energetic ions with energies in the range of 10 to 100 MeV/n for protons, He; 20 to 100 MeV/n for Li−Na ($Z = 3-11$); and 30 to 200 MeV/n for Mg−Fe ($Z = 12-26$).

PLD-151 The RAD instrument shall have the ability to measure neutral particles (neutrons and gamma-rays) with energies up to 100 MeV.

PLD-153 The RAD instrument shall have the ability to measure dose and LET spectra in the range of 1 to 1000 keV/μm.

PLD-154 The RAD instrument shall measure energetic particles with an energy resolution sufficient to distingush between major particle species (i.e., electrons, ions), low-Z ions (i.e., protons, He, Li ions), medium-Z ions (i.e., C, N, O ions), and high-Z ions (i.e., heavier nuclei up to Fe).

PLD-155 The RAD instrument shall measure energetic particles with an observing cadence sufficient to identify the onset of solar particle events (SPEs), and resolve the time profiles associated with such events.

2 RAD Sensorhead

The RAD Sensorhead employs three kinds of sensitve detectors to measure the radiation energy deposited by particles.

- Silicon diodes, PIN.
- Thallium doped Cesiumiodide, CsI(Tl)
- Organic scintillator, BC432m.

Both the CsI(Tl) and BC432m scintillators emit orange light with wavelength that can be detected with silicon photodiodes. The photodiodes are identical in structure as the silicon diodes that directly measure the deposited charge.

Photodiode readout is less demanding on bias supply voltage than photomultiplier tubes when operating in the Mars athmosphere.

Charged particles are measured in a view-cone of 60° full opening angle, traversing a sequence of sensitive detectors until the are stopped or leave the detector at the bottom.

SSD-A Silicon detector, 300 μm thick, segmented in an inner hexagon and an outer hexagonal ring.

SSD-B The second silicon detector, 300 μm thick, is a smaller hexagon, about 50 mm below SSD-A. A coincidence between SSD-A and SSD-B defines the view cone of the instrument.

SSD-C The third silicon detector, 300 μm thick, is right below SSD-B, hexagonal and slightly larger than SSD-B to cover the full projection of the view cone.

CsI(Tl)-D The CsI crystal is a 29 mm deep hexagonal prisma with sides that follow the projected view code. Three photodiodes are glued to sides of the crystal to detect the scintillation light.

NC-E Below the CsI crystal is a hexagonal prisma made of organic scintillator BC432m. Three photodiodes are glued to sides of the scintillator to detect the scintillation light.

AC-F The CsI crystal and the organic scintillator are surrounded by a round, 12 mm thick anticoincidence detector made of BC432m. Three silicon photodiodes are glues to the sides of the anticoincidence near the bottom, viewing into the bottom plate. Three outer segments of SSD-C are glued to the top of the anticoincidence.

2.1 Solid State Detectors

The three silicon detectors SSD-A, SSD-B, and SSD-C are part of three identical segmented silicon detectors. Fig. 2 shows the arrangement of the silicon detectors in ralation to the CsI crystal. Fig. 1 shows hoe the segmentation of the silicon detectors are used

SSD-A used all segments, the inner two together being one part, and the outer four segments being the other part. This radial split of SSD-A limits the path length variation for particle tracks, and keeps the detector capacitances small, to achieve better resolution to distinguish ion charge (PLD-154).

SSD-B is using only the innermost segment. Two further segments are part of the silicon anticoincidence channel C2. The outer segments are not used at all.

SSD-C is using the two innermost segments, since it needs to be slightly larger than SSD-B. The third segment is part of the silicon anticoincidence channel C2. The outer segments are detecting scintillation light from the anticoincidence.

SSD-C provides a closure of the anticoincidence around the calorimeters for neutral particle detection, improves the resolution of charged particle detection, allows for positron detection, and provides some redundancy.

2.2 Calorimeters

The depth of the CsI(Tl) crystal of 29 mm is sufficient to stop protons with energy up to 100 MeV (PLD-150), and Fe ions up to 350 MeV/n.

The high-Z material CsI is efficiently interacting with gammas, with charged secondaries (e^-) that produce scintillation light. The volume is sufficient to contain secondaries of gammas up to 100 MeV with some probability.

The BC432m organic scintillator has a high content of H atoms, which interact with neutrons. The recoil protons from these interactions produce scintillation light. The NC-E detector inside the AC-F has sufficient mass (30 g) to count neutron interactions.

The AC-F scintillator is thick enough to detect minimum ionizing particles (MIP) and allow to distinguish energy deposits on the CsI-D and NC-E detectors of neutral particles, from energy deposits by charged particles (PLD-151)

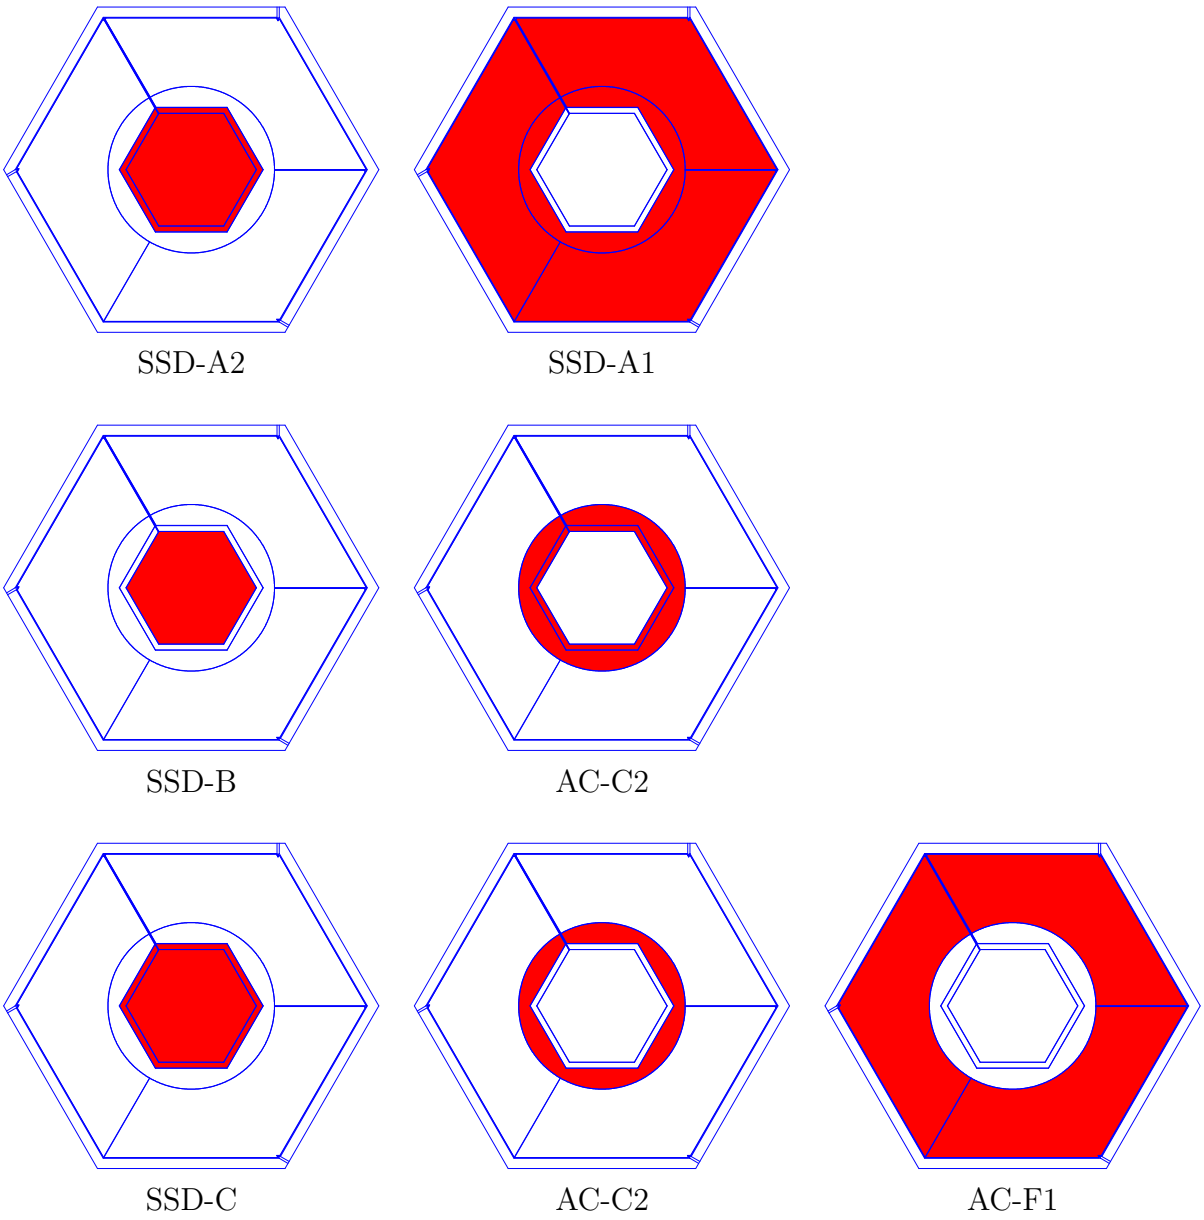

Figure 1: SSD segments. The figures in each row show the segments used on the solid-state detectors A, B, and C.
Detector A is split into an outer (A1) and an inner (A2) SSD segment. The large area of A1 results in poor resolution for single charge MIPs, but provides a large geometry factor for high-Z partices.
Detector B is split between SSD-channel B and the anticoincidence channel C2.
Detector C is split into the SSD-channel C, the anticoincidence channel C2, and the plastic anticoincidence readout F1. SSD-C is slightly larger than SSD-B, so that all particle tracs projecting through SSD-A and SSD-B also hit SSD-C. The AC-C2 segment is combined with those on detector B. The three segments of the AC-F1 channel are connected to separate CSA, so that each CSA sees only a small input capacitance for low noise. The three CSA may differ in gain, which limits the energy resolution of the analog sum.

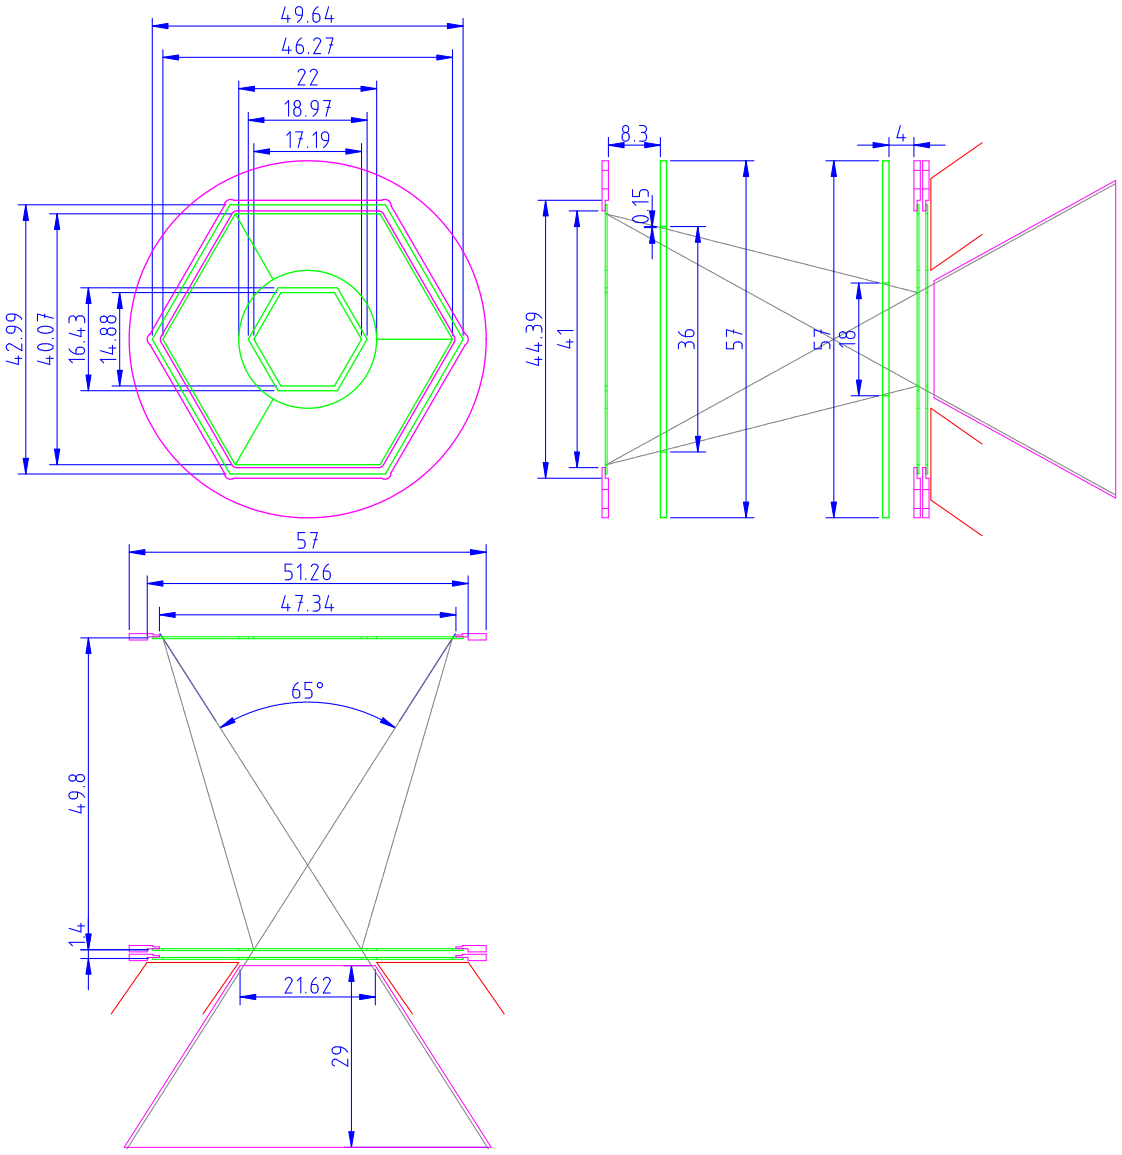

Figure 2: RAD field of view.

3 RSH Front-End-Electronics

All detector signals of RAD end up as electric charge pulses emitted from silicon diodes. The charge is either set free by ionizing radiation directly in the silicon, or by scintillation light absorbed by the sensitive side of the diode, where it is glued to the scintillator.

The front-end-electronic is distributed inside the RAD sensor head, close to the silicon diodes that each module is reading. The purpose of the FEE is

- Charge Sensitive Amplifiers (CSA) pick up the charge from the silicon diodes and convert the charge pulses to voltage steps.
- Shapers (FSH) convert the voltage step outputs from the CSA to voltage pulses.
- The detector bias voltage is distributed to the silicon detectors.

3.1 Electronics constraints

The CSA are characterized by their gain and time constant. The FSH are characterized by their gain and pole-zero compensation. The pole-zero compensation of each shaper must match the time constant of the driving CSA. The gain of the CSAs is limited to the range

$$g_{\mathrm{CSA}} = \frac{1}{8.2\,\mathrm{pF}} \cdots \frac{1}{1\,\mathrm{pF}}$$

by stability concerns and parasitic capacitances. The voltage gain of the FSH is selected between 1 and 16. The maximum output amplitude of both CSA and FSH is about 3.5 V.

3.2 Charge scale

A charged particle penetrating silicon deposits 3.6 eV energy for every electron-hole pair generated, that is 44 fC/MeV. A particle traversing the CsI(Tl) crystal needs about 15 times more energy to release a photo electron in a pin diode. A particle going through the organic scintillator looses about 100 times more energy for each photo electron.

3.3 Detector signals

The RSH has 17 CSA to read detector signals:

(2) Inner and Outer SSD-A partitions.

(1) SSD-B.

(1) SSD-C.

(1) Silicon anticoincidence, parts of SSD-B and SSD-C chips.

(3) AC-F1 top readout, three segmants of the SSD-C chip.

(3) CsI-D, three diodes.

(3) NC-E, three diodes.

(3) AC-F2, three diodes at the bottom of AC scintillator.

3.4 Dynamic Range

The detectors that are part of the charged particle view cone need to be read out with a sufficient dynamic range to cover MIPs up to heavy ions.

3.4.1 SSD

A MIP going through 300 μm of silicon deposits 60 keV/115 keV (minimum/average) energy. An iron ion that stops in tha same depth of silicon can deposit up to 1.2 GeV. It is not possible a achive this range within the constraints of section 3.1. The maximal measurable energy deposit is

$$E_{\max} = \frac{3.5\,\mathrm{V} \cdot 8.2\,\mathrm{pF}}{44\,\mathrm{fC/MeV}} = 650\,\mathrm{MeV}. \tag{1}$$

The output amplitude of a MIP can be as small as

$$V_{\min} = \frac{60\,\mathrm{keV} \cdot 44\,\mathrm{fC/MeV}}{8.2\,\mathrm{pF}} = 0.3\,\mathrm{mV}. \tag{2}$$

650 MeV is sufficient to distinguish stopping medium-Z ions from stopping high-Z ions (PLD-150). The LET spectrum of relativistic ions up to Fe requires no more than 200 MeV (PLD-153). Stopping ions will be rather infrequent on the Mars surface.

A voltage pulse of $0.3mV$ amplitude may be rather difficult to pass to the back-end electronics without adding noise. The SSD detector signals are therefore each processed by two shapers driven by a common CSA. The low-gain shaper with a voltage gain of 1 covers the full dynamic range to the maximum, while the high-gain shaper amplifies small signals so we can reasonably expect the backend to receive them with sufficiently low noise. The back end is required to measure RSH output pulse amplitudes as low as 4.5 mV

SSD-A and SSD-B are read out with the largest dynamic range. SSD-C is part of the anticoincidence for neutral particle detection, and thus it is read with more gain in the CSA, to make sure that no MIP are missed. Tab. 1 lists the exact gains of each channel.

3.4.2 Calorimeters

The CsI crystal can stop Fe ions with up to 20 GeV energy. A MIP deposits about 25 MeV. γ-Lines are located in the sub-MeV range. The 512 keV positron anihilation line is worth being measurable, as well as the ^{40}K line at 1.4 MeV. Except that the RTG will blast that energy range, so we may not see a lot of those in the background.

The organic scintillator neutron channel receives about 4 MeV energy from a penetrating MIP.

The CsI crystal (D-detector) and the organic scintillator (E-detector) each have three PIN diodes glued to side surfaces that collect the scintillation light. To cover the large dynamic range, each diode is read out with a different gain, so that they cover a part of the required range. Tab. 2 lists the gain configurations.

CSA	gain	FSH	gain	$E_{\min}$	$E_{\max}$
A1	1/8.2 pF	A1L	1	840 keV	650 MeV
		A1H	16	50 keV	40 MeV
A2	1/8.2 pF	A2L	1	840 keV	650 MeV
		A2H	16	50 keV	40 MeV
B	1/8.2 pF	BL	1	840 keV	650 MeV
		BH	16	50 keV	40 MeV
C	1/3.3 pF	CL	1	330 keV	260 MeV
		CH	8	40 keV	33 MeV

Table 1: SSD channel ranges.

The largest measurable energy in the CsI crystal is listed as 10 GeV, which is significantly less than the 25 GeV stopping Fe-ion deposit. And the scintillation light conversion efficiencies from section 3.2 are rather uncertain. But CsI is know to show significant light quenching effects at high LET, so that the real range is expected to exceed the stopping iron deposit by a large margin.

CSA	gain	FSH	gain	$E_{\min}$	$E_{\max}$
DL	1/8.2 pF	DL	1	12 MeV	10 GeV
DM	1/1.0 pF	DM	1	1.5 MeV	1 GeV
DH	1/1.0 pF	DH	8	200 keV	150 MeV
EL	1/1.0 pF	EL	1	10 MeV	8 GeV
EM	1/1.0 pF	EM	4	2.5 MeV	2 GeV
EH	1/1.0 pF	EH	16	700 keV	500 MeV

Table 2: Calorimeter channel ranges.

3.4.3 Anticoincidence

To properly count neutral particles, the anticoncidence around the calorimeters must very efficiently reject charged particles. Both the surrounding organic scintillator and the SSDs must detect MIP that enter the calorimeter. The organic scintillator (BC432m) anticoncidence light is read out by three photodiodes at the bottom (F2) and three segments of SSD-C at the top (F1). The gap between these segments and the proper SSD-C segment is closed by another diode segment on the same chip, which is called the C2 channel.

The three segments on the top and the three diodes at the bottom are each read out with a separate CSA. This keeps the detector capacitance on each CSA below 100 pF and the noise down. The outputs of three CSA are summed into one shaper. The gain mismatch of the three CSA will degrade the energy resolution, but since the anticoncidence is not used directly for spectroscopy this does not matter.

The BC432m anticoncidence is read out at the highest possible gain for maximal efficiency. Tab. 3 lists the anticoncidence CSA and shaper gain configuration.

CSA	gain	FSH	gain	$E_{\min}$	$E_{\max}$
C2	1/1.0 pF	C2	8	10 keV	10 MeV
F1a	1/1.0 pF				
F1b	1/1.0 pF				
F1c	1/1.0 pF	F1	16	700 keV	500 MeV
F2a	1/1.0 pF				
F2b	1/1.0 pF				
F2c	1/1.0 pF	F2	16	700 keV	500 MeV

Table 3: Anticoincidence channel ranges.

Pin	**Name**	**Det**	**CSA**	**FSH**	**Board**
1	F1	AC-F1	1 pF	×16	BC
2	A1H	SSD-A1	8.2 pF	×16	A
3	A1L	SSD-A1	8.2 pF	×1	A
4	A2H	SSD-A2	8.2 pF	×16	A
22	A2L	SSD-A2	8.2 pF	×1	A
23	BH	SSD-B	8.2 pF	×16	BC
6	BL	SSD-B	8.2 pF	×1	BC
8	C2	AC-C2	1 pF	×8	BC
25	CH	SSD-C	3.3 pF	×8	BC
9	CL	SSD-C	3.3 pF	×1	BC
11	DH	CsI-DH	1 pF	×8	DH/F
13	DM	CsI-DM	1 pF	×1	DM/F
31	DL	CsI-DL	8.2 pF	×1	DL/F
32	EH	n-EH	1 pF	×16	E
15	EM	n-EM	1 pF	×4	E
16	EL	n-EL	1 pF	×1	E
18	F2	AC-F2	1 pF	×16	F

Table 4: RSH output signals. The columns are: Pin number on the MDM 51 connector, signal name, detector name, CSA feedback capacitance, fast shaper gain, and the FEE circuit board where the CSA/FSH reside.

3.5 Signal channels

The electrical interface to the Rad Sensor Head is a single MDM 51 connector, which carries redundant pins for power, three pairs of pins for temperature sensores, and 17 non-redundant output signal pins. For the complete pinout description please refer to Fig. 3 and the RSH-REB electrical interface control document. The signal pins are listed in table 4.

4 VIRENA, Level 1 Trigger

The VIRENA is an ASIC that processes the shaped detector signals provided by the RSH. Fig. 4 is a block diagram of a VIRENA signal channel. The ASIC has 36 such signal channels. Two channels on the edge of the chip are not suitable for physics signal processing. Since there are twice as many VIRENA channels as RSH output signals, each RSH output signal is connected to two VIRENA inputs.

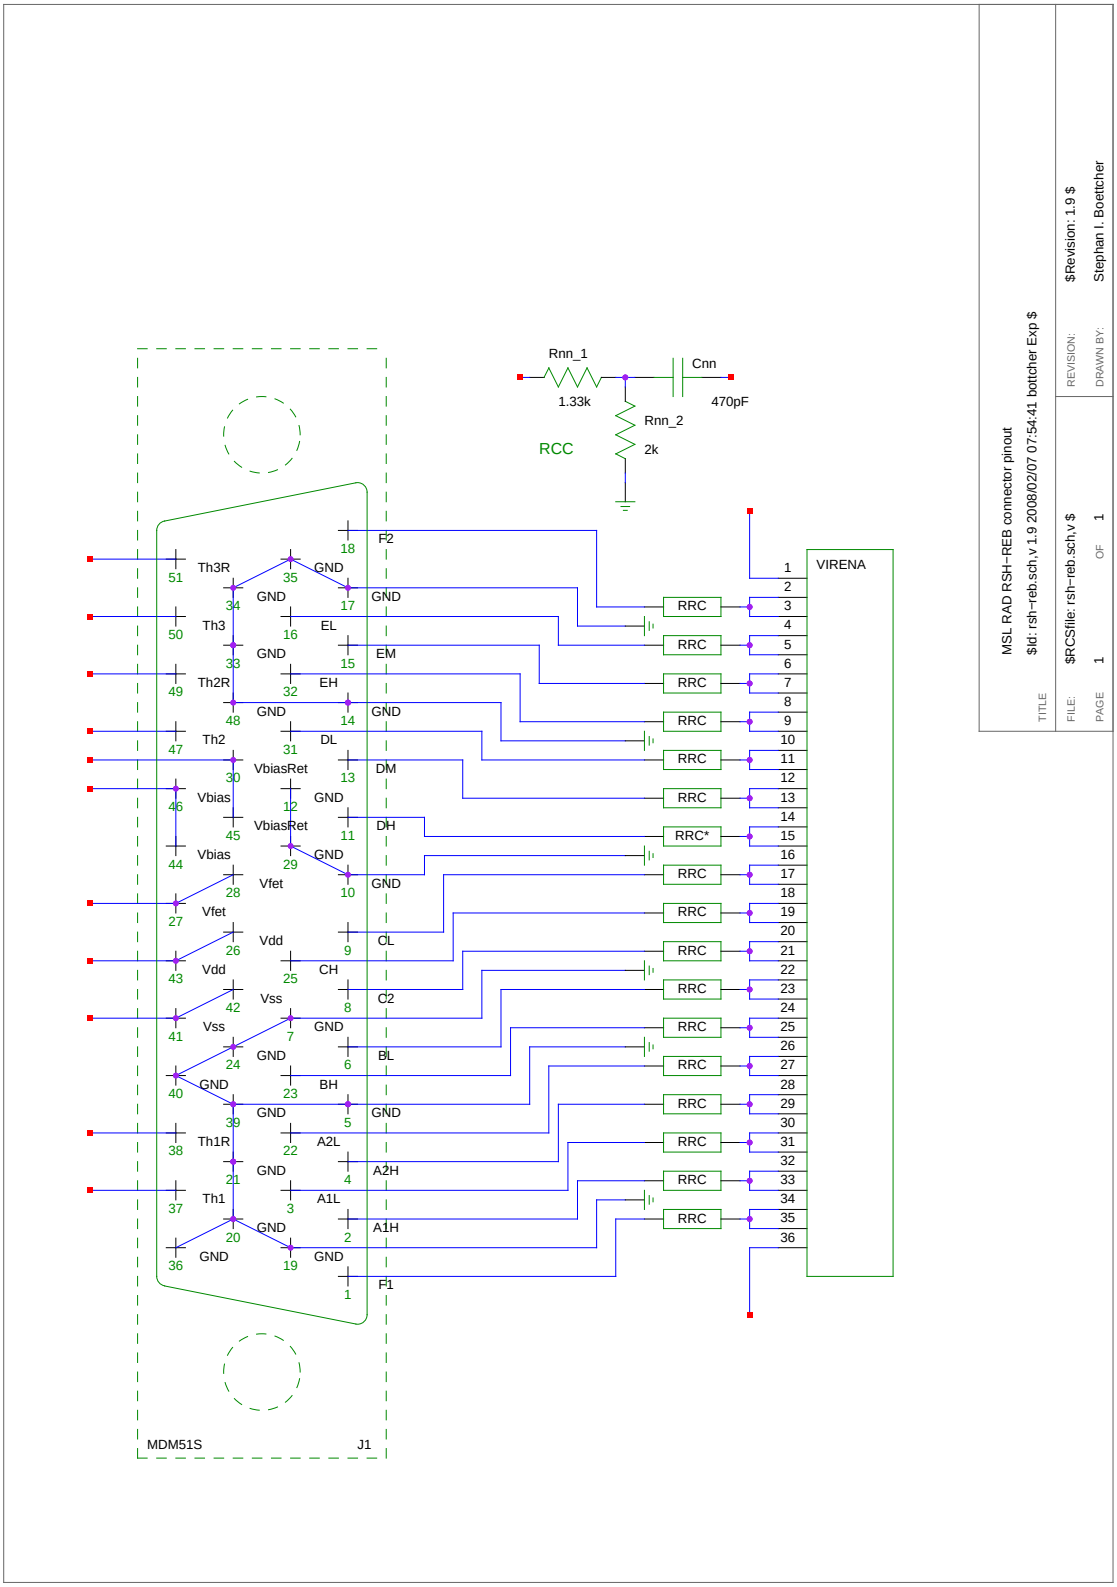

Figure 3: RSH interface connector and VIRENA input configuration

4.1 RSH-VIRENA signal matching

The VIRENA accepts positive going signal pulses in the range from 1.5 V to 3.5 V. The RSH delivers voltages between 0 V and 3.5 V. To match the RSH outputs to the VIRENA inputs, the signal is attenuated by a resistive devider by a factor 0.6, and AC-coupled with a capacitor to the VIRENA input.

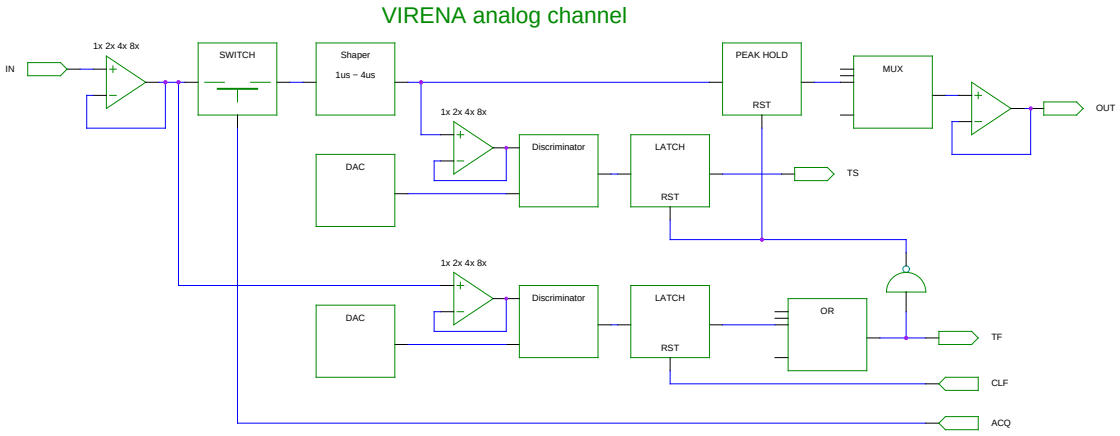

Figure 4: VIRENA channel.

The DC bias level on the VIRENA inputs is provided by 100 kΩ resisters internal to the VIRENA. There are two VIRENA inputs connected to the each capacitor, so the effective bias resistor is 50 kΩ. Internally the resistors are connected to $V_{\text{ref}} = 1.5\,\text{V}$.

The AC-coupling capacitance value is chosen for a time constant of 22 μs. The baseline shift caused by a full scale input pulse has decayed to the ‰-level after 1 ms.

4.2 Input gain

At the input each channels has linear amplifier with an individyally configurable gain of ×1, ×2, ×4, or ×8. For most RSH signals one of the two channels is opperated with gain ×1 to benefit from the full dynamic range of the signal. The other channel is operated at higher gain for additional resolution.

4.3 Fast discriminators

Each signal channel has two discriminators with an additional linear gain at the input, a DAC to set the threshold, and a flip-flop to latch pulse events that exceeded the discriminator threshold.

The fast discriminators operate on the (linear amplified) input pulses. Each channel's fast discriminator can be enabled individually, and all enabled fast discriminators contribute to the *level 1 trigger*. When any input pulses on an enabled fast discriminator exceeds the threshold a fast trigger (TF) is emmited.

The values stored in the flip-flops at the output are called the fast tokens.

4.4 Slow shapers and discriminators

Parallel to the fast discriminator, the input signal is processed by the *slow shaper*, which is a four pole integrator with configurable time constant. The time constant will be set to about 4 μs.

The output of the shaper is compared to the second, *slow* discriminator. The output latch of the slow discriminator is forced to reset as long as no fast trigger occured.

The flip-flops at the output are called the slow tokens.

4.5 Peak detecotor

The peak detector is looking at the output of the slow shaper. Normally the peak-hold capacitor is discharged by a constant current, such that the output of the peak detector follows the input voltage. When a fast trigger occured the discharge current is turned off, and when the input pulse drops the peak detector output stays at the peak voltage.

4.6 Tokens, Output multiplexer

Some time after a fast trigger was emmitted the VIRENA input will be disabled by deasserting the ACQ signal, to prevent further signal pulses to pile up on the captured event. At that point, the event is stored inside the VIRENA in three sets:

1. The fast tokens tell which channel's fast discriminator thresholds were exceeded.

2. The slow tokens tell which channel's slow discriminator threshols were exceeded.

3. The pulse height information is stored in the peak-hold caps.

The tokens are accessible via serial shift interfaces. The token patterns can be shifted in and out of the VIRENA. The slow token pattern controlls the analog output multiplexer. Only channels with s set slow token will be offered for digitization on the analog output of the VIRENA.

4.7 Channel assignment

The second level trigger processing (see below) relies on the shifting order of the tokens and the analog multiplexer output.

The RSH signal channels are grouped by detector. The RSH detectors are read out via multiple signal channels are different gain. Channels belonging to the same detector are connected to adjacent VIRENA channels, with hight gain signals connected to higher VIRENA channel numbers (see Fig. 3).

The tokens shift in and out with the highest VIRENA channel number first, channel 35 downto 0. The L2 trigger relies on seeing the highest gain signal tokens first.

While digitizing, the VIRENA offers the peak voltage of the lowest channel number first. The gain selection processing in L2 relies on seeing the lowest gain signal first.

5 EVIL, Level 2 Trigger

The VIRENA ASIC is controlled by an FPGA called EVIL. The EVIL logic also includes the second level trigger. This trigger decides which signal channels shall be digitized based on the value of the slow tokens. The L2 logic also processes the digitized peak height information, by applying calibration, and performing gain selection.

VIRENA	**Level 2**	**Name**	**Gain**	**L1**	**Readout**	h-E
0		T2	×1			
1	31	F2L	×1		First	✓
2	30	F2H	×8		Last	
3	29	EL	×1		First	
4	28	EN	×8		Yes	✓
5	27	EM	×4		-	
6	26	EI	×4		Yes	
7	25	EH	×2		Last	
8	24	EU	×8	✓	-	
9	23	DL	×1		First	
10	22	DN	×8		Yes	✓
11	21	DM	×8		-	
12	20	DI	×4		Yes	
13	19	DH	×1		Last	
14		DU	×8	✓		
15	18	CL	×1		First	
16	17	CM	×8		Yes	✓
17		CH	×1			
18	16	CU	×8		Last	
19	15	C2L	×1		First	✓
20	14	C2H	×8		Last	
21	13	BL	×1		First	
22	12	BM	×8		Yes	✓
23	11	BH	×8		-	
24	10	BU	×8	✓	Last	
25	9	A2L	×1		First	
26	8	A2M	×8		Yes	✓
27	7	A2H	×1		-	
28	6	A2U	×8		Last	
29	5	A1L	×1		First	
30	4	A1M	×8		Yes	✓
31	3	A1H	×1		-	
32	2	A1U	×8		Last	
33	1	F1L	×1		First	✓
34	0	F1H	×8		Last	
35		T1	×1			

Table 5: VIRENA channels. The gain menmonics are somewhat arbitrary: *Ultra, High, Intermediate, Medium, Next, Low.*

5.1 Level 1 Trigger

The EVIL controls the processing of the event in response to a fast trigger on the TF signal from the VIRENA. A configurable time after the arrival of TF, the EVIL deasserts the ACQ signal to the VIRENA. The delay must cover the peaking

time of the slow shapers, so that slow discriminators and the peak detectors can capture the peak of the shaped input pulse. The delay shall be 10 μs, typically.

After ACQ was deasserted the fast and slow tokens are shifted out of the VIRENA into the L2 trigger processer. The immediate result from the L2 trigger processor is a priority bit. The EVIL decides if the event shall be digitized based on the priority bit and the amount of free space in the EVIL output FIFO. High priority events will be digitized when at least one full-sized event fits into the FIFO. Low priority event are digitized only when the FIFO is at most half full.

When the EVIL decides that the event shall be digitized, it commands the L2 trigger to produce readout tokens, which are shifted into the VIRENA, to tell which channels shall be digitized.

Subsequently, the EVIL controlls the ADC to digitized the VIRENA output, switching the output multiplexer through all requested channels in turn. The digitized values are forwarded to the L2 processor for calibration, and gain selection.

When the L1 event processing is done, the EVIL resets and reenables the VIRENA to wait for the next event. The reenable sequence may be delayed by a configurable amount of time, starting from the arrival of the TF signal, when any one from a configurable subset of fast tokens is set. The purpose of this delay is to allow the analog signal chain to go basck to baseline after a large energy deposit. The delay will be enabled for low-gain channels, including the redundant F1/F2 channels, and last for about 1 ms.

5.2 L2 trigger decision

Level 2 processing is applied to a configurable subset of the 36 VIRENA channels. Up to 32 channels can be selected. Channels which are not selected are completely ignored by the L2 trigger and any downstream processing. Typically, channels 1 to 34 are selected, apart from two channels that are not essential. Table. 5 shows the assignment of signals to VIRENA channel numbers and L2 trigger channel numbers for this typical case, with DU and CH ommitted. The extra gain in DU is just amplifying noise. The Si-H channels are not required to cover the dynamic range, and CH is not even required for L2 triggers.

32-bit bitmasks are just impossible to handle efficiently on the level three processor. So we just pretend that the VIRENA has only 32 channels. There is one 36-bit configuration register in the L2 processor which tells what channels to use. Fewer than 32 channels may be specified. When more than 32 bits are set in the register, the excess bits are ignored at the LSB end.

5.2.1 Slow token in

The level 2 trigger consists of 16 configurable coincidence conditions. It's a *trigger menu*, with 16 items.

Each coincidence condition requires a given channel's slow token to be set, reset, or don't-care. Each coincidence condition is associated with a priority, a set of channels that shall be digitized, and a bit that may require all channels to be digitized even when that detector has no slow token hit (see below).

The complete set of configuration registers for each trigger menu item is:

Mask: A 32-bit register, one bit for each channel, telling if that channel is contributing to this coincidence condition.

Value: A 32-bit register, one bit per channel, telling if the slow token must be set or reset. If any bit in the Value register is set while the corresponding Mask bit is not set, then this coincidence condition will never be true.

Read: A 32-bit register telling which channels shall be digitized if this coincidence condition is true.

Pri: A single bit telling if this coincidence condition is a high priority trigger.

Too: A single bit that, when set, tells to digitized only channels that have a slow token set on any channel of the same detector, sort of.

The **Too** and **Pri** bits are combined into one 32-bit configuration register.

An event is high priority if any coincidence condition is met that is marked high priority.

5.2.2 L2 Counters

While the tokens are shifted out of the VIRENA they get counted. There is a counter for each slow and fast token for each of the 32 selected channels. As soon as the last token arrived, when the final trigger decision is available, one of two global event counter is incremented, either the low priority or the high priority counter. Afterwards each L2 coincidence trigger gets countered. There is a counter for each coincidence condition.

5.3 Readout tokens

When the EVIL decides not to digitize, because the event does not fit into the allocated FIFO space, event processing is finished at this point. Otherwise, the EVIL requests readout tokens from the L2 trigger module.

Each coincidence condition is assotioated with a set of channels that shall be digitized when the condition is met. The overall set of digitized channels is the union of sets requested by the L2 triggers.

5.3.1 The TOO logic

There is one more global configuration register for the L2 trigger, a 32-bit register that tells which channels belong to the same detector. Each bits corresponds to a channel, when the bit is set it tells that the channel belongs to the same detector as the previous channel. The highest-gain channel for each detector shall have this bit reset. All other channels, which follow the highest-gain channel, have their *same-detector* bit set. This is one reason why the channels of each detector must be adjacent and ordered on the VIRENA inputs. We assume that the highest gain channel has the lowest overall slow threshold and that its slow token will be set too whenever any lower-gain token is set.

When the **Too** bit of a coincidence condition is *not* set, and the coincidence condition is true, then all channels specified in the corresponding **Read** register will be digitized. This is typically used for high-priority triggers, where a complete set of peak heights is desirable for ground processing.

When the **Too** bit is set, for a channels to get digitized (in addition to the corresponding bit in the **Read** register) the corresponing slow token is required

to be set, or the previous channel was selected for digitization and the channel belongs to the same detector:

$$Too[i] = \text{Slow}[i] \textbf{ OR } (\text{Too}[i-1] \textbf{ AND } \text{Same}[i]) \quad (3)$$

$$\text{Readout}[i] = \text{Too}[i] \textbf{ AND } \text{Read}[i] \quad (4)$$

for each L2 channel number i.

5.3.2 Readout counters

When the EVIL requests readout tokens for digitization, the L2 modules increments a second set of global event counters and L2 trigger counters, but no token counters. Dividing the number of triggered events by the number of digitized events of a given class gives the prescale factor for that class of events. In practice only two event classes will be relevant: low priority and high priority events.

5.3.3 Event header

The request for readout tokens causes yet another action from the L2 trigger module: The event data packet header will be formatted and sent to the FIFO.

The event data packet consists of three parts, the EVIL packet header, the event data header, and the channel data. The EVIL packet header are five bytes defined by the RAE to REB data stream protocol. Event data are sent in EVIL packets of type 0.

The event data header is five words (twenty bytes) long:

Slow Token: The slow discriminator tokens. Only the 32 selected channels.

Readout token: The readout token. This bitmask tells which channels the appended data records belong to, 32 bit.

Misc/Trigger The LSB are the matched trigger conditions. 6 bits for the number of channels being digitized. 8 bits L1 event number. The priority bit, and an attention bit ???.

Fast Token: The fast discriminator tokens, 32 bits.

EVIL Timestamp: A 24-bit timestamp with 1 μs resolution, and an 8-bit delta-timestamp, which measures the time since the previous trigger in units of 4 μs.

5.4 Calibration

After the readout tokens have been shifted out, and the triggers were counted, the L2 modules waits for the digitized peak heights from the EVIL.

The ADC result is provided to the L2 module as a 14-bit unsigned number, with a pedestal around 1000 and peaks going positive up to 16383.

Each digitized peak height is processed and formatted into an eight-byte record which is sent to the FIFO. The correct number of digitizations is required to complete a well formed event data packet. The channel record contents is:

PHA data: 8-bit floating point representation of the ADC reading.

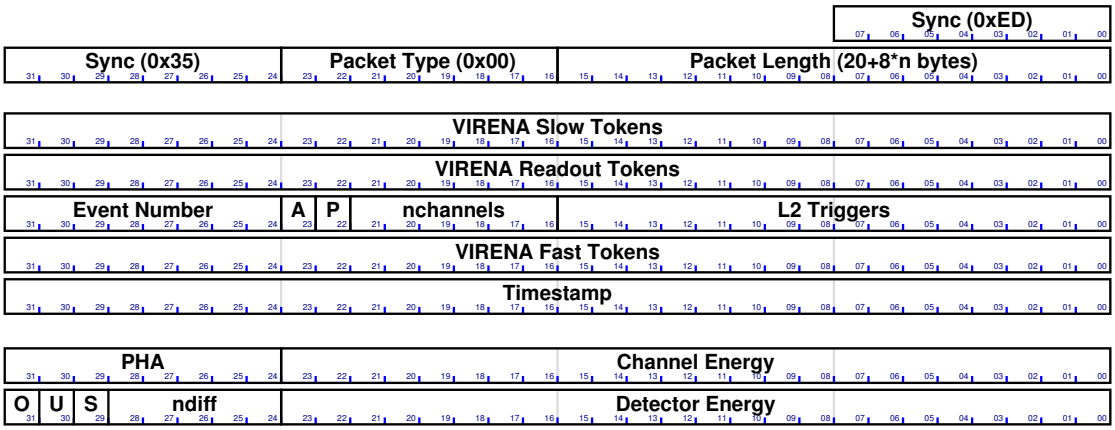

Figure 5: Event packet format. The EVIL packet header is five bytes long, including sync bytes, packet type, and packet payload length. The event data header is 5 words long, with 4 bytes per word. The event packet header is followed by two words (eight bytes) for each digitized channel.

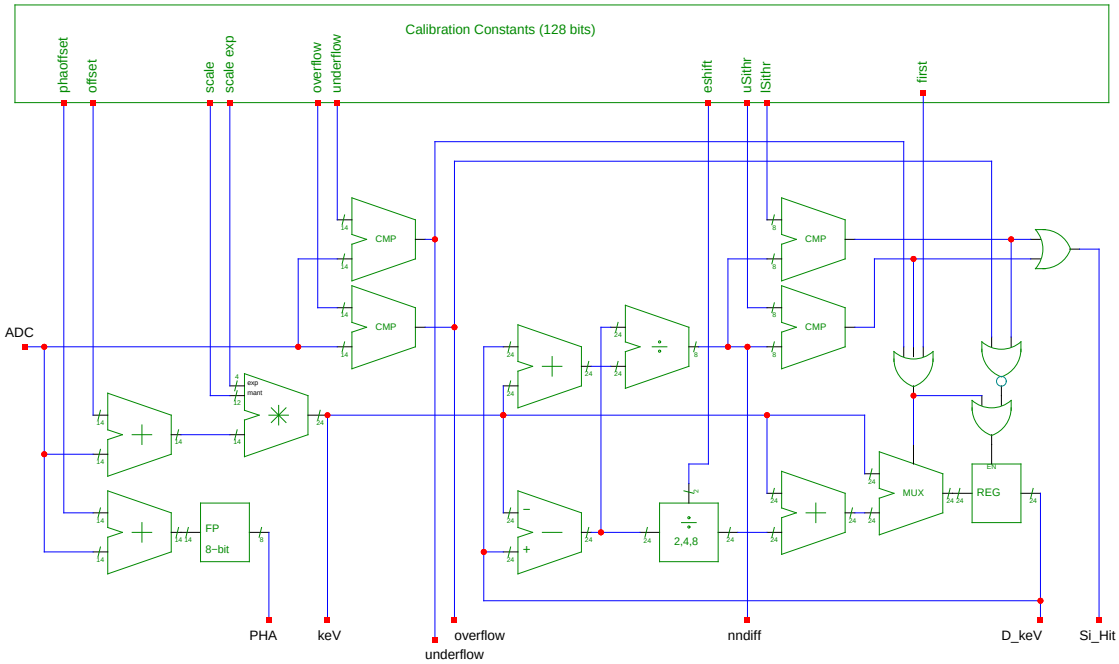

Figure 6: L2 calibration block diagram.

Calibrated Energy: 24-bit unsigned integer representation of the ADC reading converted to deposited energy in keV.

Detector Energy: Result of the gain selection for this detector up to this channel, 24 bits.

Flags: Miscellaneous flags from the gain selection process, 8 bits.

sym	bits	name	description
p	14	offset	pedestal of the ADC reading
m	12	mantissa	energy scale factor
e	4	exponent	energy scale factor
p'	14	PHA offset	a little below pedestal

Table 6: Configuration parameters for the energy calibration. There is one set of these parameters for each L2 channel.

5.4.1 Calibrated energy

The ADC delivers the pulse height as a 14-bit unsigned number, with some pedestal that varies from channel to channel, and a scale that depends on the detector sensitivity, preamp sensitivity, shaper gain, and the configurable VIRENA input gain.

To obtain the calibrated energy, at first a 14-bit signed offset p is added to the ADC result D. This offset is negative in most cases, effectively subtracting the pedestal.

The addition is performed with overflow/underflow detection and clipping, so that the unsigned result does not wrap around in case of overflow. Afterwards a calibration factor is multiplied to the result, yielding a 24-bit unsigned result. The factor is a 12-bit unsigned mantissa m and a 4-bit exponent e:

$$E = \mathrm{floor}((D + p) \cdot m \cdot 2^{-2-e} + \frac{1}{2}) \tag{5}$$

5.4.2 PHA data

The level 3 trigger will select a subset of events for analysis on the ground. For that purpose, the ADC results of each channel shall be converted to an 8-bit floating point number to fit more events into the telemetry allocation.

Before the number is converted, an offset p' is added to the 14-bit ADC result, just like for calibration. A different offset is used because the PHA data shall cover the baseline peak, so that a little positive pedetal shall remain. But that remaining pedestal shall be small enough that the 5-bit mantissa can represent the baseline peak and small energy deposits with full ADC resolution.

The floating point value is unsigned, 3-bit exponent e, 5-bit normalized mantissa m with implicit leading 1. An exponent value of zero means a non-normalized mantissa. To recover an approximation of the original ADC result, use:

$$D = \begin{cases} 4m - p' & \text{for} \quad e = 0 \\ (m + 32) \cdot 2^{e+1} - p' & \text{for} \quad e > 0 \end{cases} \tag{6}$$

name	bits	description
first	1	First (lowest gain) channel of the detector
last	1	Last (highest gain) channel of the detector
overflow	14	ADC reading threshold for overflow
underflow	14	ADC reading threshold for underflow
usihit	8	upper silicon hit threshold
lsihit	8	lower silicon hit threshold
eshift	2	linear combination index

Table 7: Gain selection configuration parameters

5.4.3 Gain Selection

The Level 2 module attempts to estimate the energy deposited in a given detector based on the ADC results of all its channels. Each channel is processed in turn, starting with the lowest gain. The result is a current best estimate of the detector energy, called E_{sum}. This value is sent to the FIFO for every channel. E_{sum} is also carried over for the next channel processing. The last, highest gain channel's E_{sum} can be used by level 3 as the overall result for the detector.

The lowest gain channel shall be marked by the *First* bit in the calibration memory. When this bit is set for a channel, the calibrated energy is copied to the E_{sum}, and no further conditions considered, since the previous channels belong to a different detector.

Otherwise, the ADC result is compared to the *overflow threshold.* When the ADC result exceeds that threshold, the channel result is ignored, and the previous E_{sum} value is kept. The threshold shall be configured so that ADC values below are in the linear response range of the channel. Values beyond overflow threshold are discared in preference of the results of lower gain channels. An overflow is flagged in the channel data record.

Otherwise, the ADC result is compared to the *underflow threshold.* When the ADC result is below the underflow threshold, the previous E_{sum} is ignored, and the calibrated energy from this channel is copied to E_{sum}. The underflow threshold shall be configured so that ADC values below indicate that the previous, lower gain channels cannot usefully contribute anything but noise to the result. An underflow is flagged in the channel data record.

5.4.4 Si-hit detection

When the channel is not first, not in overflow nor in underflow, then a silicon hit detection is attempted. The scintillation detectors are read out with multiple photo diodes. When a particle deposits energy in the scintillator, all diodes should see about the same charge, and the calibrated energy of all channels should match. When a particle hits the pin diode itself, that channel will show a much large signal compared to the others. The larger value shall be discared.

The L2 module computes a normalised difference of the calibrated energy E and the previous E_{sum}:

$$N = 64 \cdot \frac{E_{\text{sum}} - E}{E_{\text{sum}} + E} \tag{7}$$

The result N is compared to an *upper* or *lower silicon hit threshold* depending on the sign of N. The five most significat bits of N are sent as flags to the FIFO.

Index	Condition	Readout	TOO	Pri	Description
0	A1U·BU·!C2H	all	✓		low-LET
1	A2U·BU·!C2H	all	✓		low-LET
2	A1M·BM	all		✓	high-LET
3	A2M·BM	all		✓	high-LET
4	BH	A1·A2·B	✓		Si dosimetrie
5	EU·EI	E	✓		plastic dos.
6	EH·EI·!AC·!BU·!CU	F1·F2·E			neutrons
7	DH·DI·!AC·!BU·!CU	F1·F2·D			gammas
8	DH·DI·EU·EI·!AC·!BU·!CU	F1·F2·D·E		✓	neutrals
9	DM·DN	all		✓	heavy ions
...					
13	BU·!F1·!F2·!DU·!EU	none			stopped in Si
14	NONE	none			count empty
15	true	none			count all

Table 8: Typical Level 2 trigger menu. AC = (F1H+F2H+C2H). NONE requires all slow tokens to be zero.

If N is positive, and exceeds the *upper silicon hit threshold,* then the previous $E_{\rm sum}$ is much larger than what this channel says is should be. In this case the previous $E_{\rm sum}$ is ignored, and the current channel calibrated energy is copied to $E_{\rm sum}$.

If N is negative and below the *lower silicon hit threshold,* this channel has a much larger energy signal than the previous channels. The new result is discared and the old value for $E_{\rm sum}$ is kept. A silicon hit is flagged in the channel data record.

5.4.5 Linear combination of energy readings

When none of the above conditions are satisfied, we assume that the previous $E_{\rm sum}$ may contribute signal above noise and this channel even more, since it is higher gain. The new $E'_{\rm sum}$ is computed as a linear combination of the previous $E_{\rm sum}$ and this channels calibrated energy E:

$$E'_{\rm sum} = E + k \cdot (E_{\rm sum} - E) \tag{8}$$

where k can be be configured to be 0, 1/2, 1/4, or 1/8. The selection is done via the *eshift* bits:

eshift	k	Formula
0	0	$E'_{\rm sum} = E$
1	$\frac{1}{2}$	$E'_{\rm sum} = \frac{1}{2}E + \frac{1}{2}E_{\rm sum}$
2	$\frac{1}{4}$	$E'_{\rm sum} = \frac{3}{4}E + \frac{1}{4}E_{\rm sum}$
3	$\frac{1}{8}$	$E'_{\rm sum} = \frac{7}{8}E + \frac{1}{8}E_{\rm sum}$

The new value $E'_{\rm sum}$ is sent to the FIFO as the current best estimate of the detectro energy, and becomes the saved $E_{\rm sum}$ for the next channel gain selection, if any.

5.5 L2 trigger menu

Table 8 is a typical trigger menu for science observation on the Mars surface.

Triggers 0 to 3 are charged particle triggers. The frequent low-LET particles (protons, muons, He, e^-) require an anticoincidence condition in C2. All detector will be digitized if they trigger the slow token. For high-LET charged particles (heavy ions) we do not require an anticoincidence condition because stray secondaries may trigger the AC detector without invalidating the event. High-LET charged particles are high-priority, and we always digitize all channels. The slow discriminator thresholds in the Ultra-high-gain SSD channels will be set for efficient detection of single charge minimum ionizing particles. The thresholds for the Medium-gain channels will be set to discriminate heavy ions from light ions.

Triggers 4 and 5 are dosimetry triggers. Both are single detector triggers, therefore we use slightly higher thresholds. The Si dosimetry trigger covers both total dose and LET, the SSD-A detectors are digitized too when they show activity. The EU channel shall have a higher threshold than the EH channel, but with better resolution. Since EU is not read out, the no harm is done when we find an EH token without an EU token.

Triggers 6 to 8 are cover neutral particles. All anti-coincidence channels must be quiet, and either D or E or both must be triggered. We require a coincidence between DH and DI, or EH and EI, to reject silicon hits at L2 level.

SSD-B and SSD-C are part of the anti-coincidence. The F1/F2 anti-coincidence is also read out independently of any slow discriminator tokens, so that the L2 trigger can fine-tune the trigger cuts. Neutral particles that hit both D and E are high priority, because they go into a separate L3 histogram to study the leakage of secondaries from one into the other calorimeter scintillator.

Trigger 9 causes a high-priority readout of the complete instrument for any really heavy ion hit from any direction into the CsI crystal.

Trigger 13 tests if a charged particle stopped in the SSD, or if it made it into the calorimeter. This is provided to aid the L3 trigger processor to quickly see which class of events this is.

Triggers 14 and 15 are used for counting only. Trigger 14 counts all events with not slow tokens set at all. Trigger 15 counts all events.

Three triggers are unused, and the two counter triggeres are not required. That leaves five spare triggers for unforeseen needs.

6 Level 3 Trigger

The third level trigger is implemeted in software executed on the embedded 8051 processor in the RDE board FPGA. L3 counts events in histograms based on the energy reading digitized by the ADC and processed by L2. The compressed ADC readings are saved in PHA buffers for telemetry, until the buffers are full. At the end of an acquisition period, the histograms, PHA bufferes, and houskeeping records are packaged into observation packets and stored in non-volatile memory until the rover retrieves them.

L3 receives the particle event packets from L2 through a FIFO. Each event contains bitfields for the VIRENA discriminator tokens and the L2 trigger hits, and the digitization results. Each digitization record contains the energy observed in the physical detector derived from all channels up to the current one by the L2

gain selection module. The last channel of each physical detector (highest gain) contains its best reconstructed energy.

There are nine physical detectors: A1, A2, B, C, D, E, C2, F1, F2.

L3 maintains a 4-bit timestamp with divides the observation period into 16 segments.

The L2 trigger assigns a binary priority value to each event. A different fraction of high and low priority event may be digitized. Each L3 event counter must be maintained in two copies, one for high and one for low priority events.

The L3 processor can efficiently calulate the logarithm of an enrgy value:

$$L_2(E) = 8 * log_2(E) \tag{9}$$

given a unsigned 24-bit E, yieling an unsigned 8-bit results. These L_2 values are used for most arithmetic, cuts, and histogram indices.

6.1 Total Dose

Two detectors are used for dosimetry: the SSD-B detector and the organic scintillator (E-detector). The L2 trigger provides for each of these detectors a trigger. L3 tests these bits to decide if total dose processing is required.

The total dose is the sum of all energy deposits observed in these detectors. There are individual counters for each timestamp period.

The total number of total dose counters is (detectors, timestamp, priority):

$$N_{\mathrm{TD}} = 2 \times 16 \times 2 = 64. \tag{10}$$

The total dose must only be counted when the corresponding slow token (BH, EH) is set. Otherwise we get a biased noise contribution.

The light output of the plastic scintillator is not linear, but depends on the LET of the particle track (quenching). Early calibration tests show that this effect is significant. Therefore it is proposed to count the total dose from the NC-E detector in 16 LET bins, instead of 16 time slots.

6.2 LET

Any particle event that penetrated SSD-A and SSD-B will be processed for LET spectra.

Condition: The L2 trigger tests a set of L2 to decide if a potential AB coincidence happened. This set may not actually require an actual SSD-A hit, if the trigger menu has no space for entries. L3 must further test for non-zero SSD-A energy.

Which A: The LET and charged particle spectra are separately counter for particles going through the inner or outer SSD-A segments.

$$A = A_1 > A_2 \ ? \ A1 : A2 \tag{11}$$

When attempting to calculate $L_2(A)$, the coprocessor may provide a convenient bit to tell that A is zero, and LET processing can be aborted.

Cuts: Penetrating particles will deposit a similar amount of energy in SSD-A and SSD-B. L3 tests if this is the case:

$$K_1 < \frac{A}{B} < K_2, \tag{12}$$

or in log:

$$K_1 < L_2(A) - L_2(B) < K_2, \tag{13}$$

or practically, with unsigned byte math

$$L_2(A) < L_2(B) \quad ? \quad L_2(B) - L_2(A) < K_1 \tag{14}$$

$$: \quad L_2(A) - L_2(B) < K_2. \tag{15}$$

Histogram: The LET events are counted in a $2 \times 2 \times 256$ histogram, binned by

$$\text{prio} \times A_1{>}A_2 \times L_2(B). \tag{16}$$

At the end of the observation, a configurable $L_2(B)$ range of the histogram is rebinned for telemetry. The range shall include energy deposits from 60 keV to 120 MeV, in bins of $L_2(B)/2$.

6.3 Charged Particles

All charges particle events were necesarily already processed for LET spectroscopy. The values of $A_1{>}A_2$, $L_2(A)$, and $L_2(B)$ are carried over.

Condition: When the event was considered for LET, an additional set of L2 trigger bits is tested to see if it also qualifies for charged particle processing.

The L2 trigger menu must include a trigger on the condition $\bar{D} \cdot \bar{E} \cdot \bar{F}_1 \cdot \bar{F}_2$. The L3 code tests this trigger bit to distinguish charged particles that stopped in an SSD, from particles that made it into or through the calorimeters.

6.3.1 Stopping in Silicon

When the L2 trigger bit $\bar{D} \cdot \bar{E} \cdot \bar{F}_1 \cdot \bar{F}_2$ is true, the charged particle was stopped in either SSD-B or SSD-C.

Cuts: Four cuts are required to pass for an event to be counted in the stopping charged particles histogram:

The silicon anticoincidence C_2 must show significantly less energy deposit than B:

$$\frac{C_2}{B} < K_1, \tag{17}$$

with unsigned byte math:

$$L_2(C_2) + K_1 < L_2(B). \tag{18}$$

And C_2 must show less energy than C, but only if there is a token T_C:

$$\bar{T_C} \quad \text{OR} \quad L_2(C_2) + K_2 < L_2(C). \tag{19}$$

Finally, the ratio of energy deposits in A and B must be in range:

$$K_3 < \frac{A}{B} < K_4, \tag{20}$$

or practically, with unsigned byte math

$$L_2(A) < L_2(B) \quad ? \quad L_2(B) - L_2(A) < K_3 \tag{21}$$
$$: \quad L_2(A) - L_2(B) < K_4. \tag{22}$$

Histogram: The event will be counted as a stopping charged particle with total energy:

$$E_{\text{tot}} = A + B + C. \tag{23}$$

and LET:

$$\frac{\mathrm{d}E}{\mathrm{d}x} = A + B + C. \tag{24}$$

6.3.2 Stopping in the Calorimeter

When the L2 trigger bit $\bar{D} \cdot \bar{E} \cdot \bar{F}_1 \cdot \bar{F}_2$ is not true, the charged particle reached the calorimeters.

Condition: Processing will be aborted at this point when one on the slow tokens T_C or T_D is not set. The particle must have deposited energy in C and the CsI D.

Cuts: The energy scale for all cuts is calculated as the total energy in the SSDs:

$$E_{\text{SSD}} = A + B + C. \tag{25}$$

6.3.3 Penetrating Particles

6.4 Neutral Particles

7 Ground processing