DS–RAP–RP–0021 Issue: 1
Date: Oct 11, 2004 Rev. : 0

RAPID Recalibration

The Wherefore and the Why

Patrick W. Daly
Max-Planck-Insitut für Sonnensystemforschung
Katlenburg-Lindau, Germany

Following RAPID Team Meeting
at Boston University
Sept. 25–26, 2004

Document Status Sheet			
1. Document Title: **RAPID Recalibration Plan**			
2. Document Reference Number: **DS–RAP–RP–0021**			
3. Issue	4. Revision	5. Date	6. Reason for Change
Draft	0	2004 Oct 4	Draft
1	0	2004 Oct 11	First issue

Contents

Summary

RAPID is an imaging spectrometer for ions and electrons of energies over ~ 30 keV, on board the 4 Cluster spacecraft.

Since the start of the operations phase of Cluster in early 2001, a number of problems involving RAPID have been recognized, problems which were given only provisional solutions at the time. These problems have now been analyzed to the extent that solutions are now possible. A complete overhaul of the calibration factors, their determinations, and their usage by the analysis software will have to be undertaken.

This document outlines those problems, their causes, and shows how they are to be solved. It is intended as an explanation for those investigators, mainly the RAPID Co-Is, already using the data set, as well as a blueprint for the recalibration activities.

As such, it is primarily an “insider’s” document, assuming the reader is well familiar with the instrument’s functioning and design.

1 The Current Situation

1.1 The RAPID analysis software package

The RAPID raw data are obtained on CDs or are downloaded from ESOC. They are then combined into *merged science files*, or MSFs, which is simply a repackaging of the original raw data into a more convenient, synchronized form.

The basic analysis program actually combines several different functions, and could in fact be split off into several different programs. However, since they share much common coding, they have been retained as one big program. By means of aliases, these different functions can be invoked by various calling names, as though they were different programs.

To be more specific, the one big program is called `gcdc_pmgr` and its first argument specifies the functionality to be used. This first argument is also normally the alias for invoking it.

For example,

`gcdc_pmgr msf2iff /D=`*yyyymmdd* `/S=`*n*

reads the MSF data for a specified day and spacecraft, and creates an IFF file (*Intermediate Format File*) which is the basis for generating the Prime Parameter CDF files for CSDS. There is an alias `msf2iff` equated to this command. Similarly,

`gcdc_pmgr msf2sci /D=`*yyyymmdd* `/S=`*n* `/K=`*classes* ...

creates a so-called SCI file (an ascii file with scientific parameters) used for all subsequent data analysis, plotting, and submission to the Cluster Active Archive. The alias for this is `msf2sci`.

Both of these functions read in raw data from the MSF files, and generate physical parameters, as count rates, differential or integral flux, making use of *calibration factors* that need to known for the processing. These factors include geometry factors, efficiencies, energy levels. They need to be properly determined in order that the conversion from raw counts to fluxes is carried out correctly.

1.2 The need for revision

The original calibration factors were established before the launch and then revised after the commissioning phase. However, it was long recognized that a major overhaul would be needed at some point, simply because there were puzzling aspects in the data sets.

Our understanding of the operation of the RAPID measurements has reached the point where a redefinition of the calibration factors can be undertaken.

The following points need to be fixed up:

1. The counts in the 2nd hydrogen energy channel are definitely too low compared to the neighbouring channels. Allowing for the narrow energy bandwidth compensates partially, but not entirely, for this discrepancy.
2. The 3D ion distributions (`I3DD`) when summed over sectors and polar directions, is a factor of 2–3 lower than the corresponding omnidirectional values (`HSPCT`/`ISPCT`), except for the first energy channel, where the agreement is very good.
3. There would appear to be some hydrogen contamination in the first helium channel. This should be corrected for.
4. The actual energy thresholds vary slightly among the different sensors. A spectral correction should be made to reduce all the sensors to a common set of thresholds. This must be done for both ions and electrons.

5. The pedestal shifting in the electron data must be compensated for. This is a additional, varying, spectrum shift. (The RAL software already does this—its algorithms must be incorporated into the Lindau software package.)

In addition, cross-correlations among the spacecraft and among the sensor heads on each spacecraft must be routinely checked and corrected. An attempt to do this is a rough manner already exists in the current "expert" calibrations.

1.3 The present calibration files

The calibration factors for the IIMS and IES parts of RAPID are stored in *calibration sets* which are valid for particular time periods. These sets are integrated into *calibration files*. At the moment, there are two sets of calibration files in use:

Normal sets contain constant ideal geometry factors determined before the launch; no allowance for the loss of the central ion heads;

Expert sets (ions only) try to compensate for relative differences among the spacecraft and their detector heads.

In both sets, the changes in MCP voltages and IES look-up-tables are incorporated.

The normal sets should be abandoned. The idea was that at least you knew what you had, even if it was wrong.

The contents of the current calibration sets are listed in Sections A.1 and A.2.

2 Direct Event Analysis

In September 2004, we undertook an analysis of the ion behaviour by examining the direct events in the RAPID data. These contain the most detailed information about what IIMS is measuring at the most basic level. Each event consists of 4 numbers:

- the time-of-flight channel (0–255, for 0–80 ns)
- the measured energy channel (0-255, for 0–1500 keV)
- the sector number (0–15)
- the polar direction number, i.e., the sensor head and its subdivision into 4 fine directions (per sensor).

  This must be elaborated somewhat. The direction number is 0–15, and is encoded as follows:

0–11	for the 12 fine directions, sensor head plus subdivision
12–14	for sensor head only, without fine subdivision
15	for no direction determination

In nominal mode (NM) there are at most 20 events per data record (spin), and in burst mode (BM) at most 106. A prioritization scheme is used to ensure that heavier ions are read out before protons.

Every point in the 256×256 Energy–TOF space is assigned to a set of bins, 8 (energy) bins for each of several mass ranges, as shown in Figure 1. (Figure 2 is the same but with real direct events added as a scatter plot.)

An important aspect to appreciate here is that the measured energy must exceed 30 keV before an energy signal is generated (energy bin 5). If a TOF signal is obtained without a corresponding energy signal, the event is registered as being in the *underrange* region. Particle classification is then carried out solely by TOF analysis. The question remains whether this is good enough. See Section 2.5.

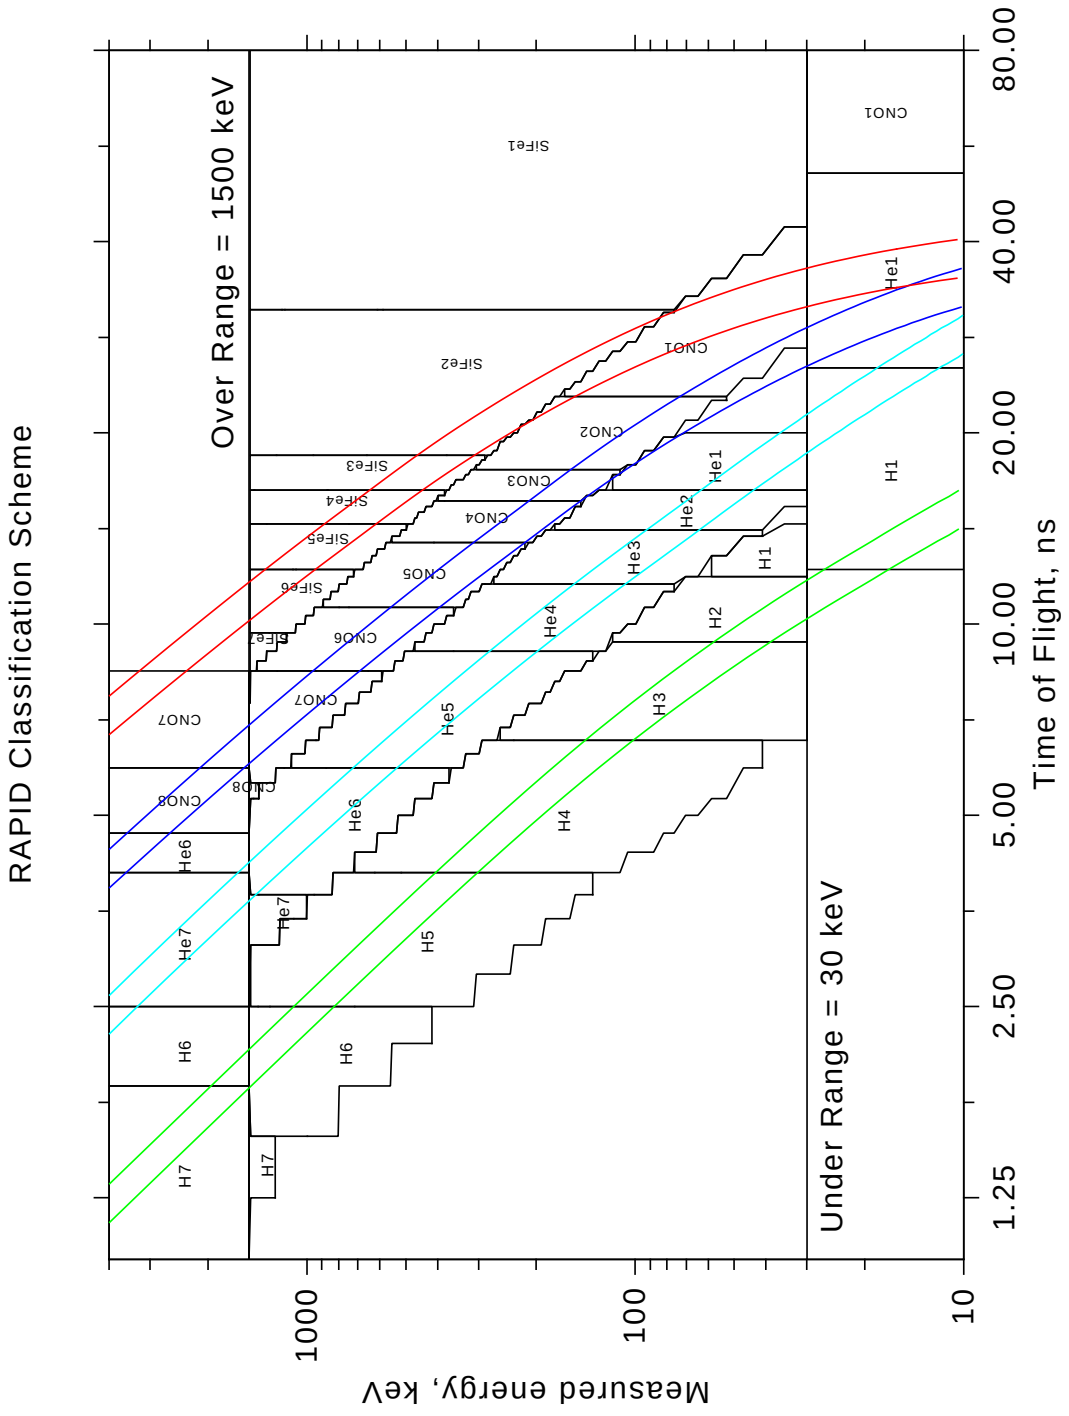

Figure 1: The bin assignments in the energy–TOF space, with theoretical loci for various species (coloured lines), allowing for energy loss (measured energy < time-of-flight energy). The left curve of each pair is for direct angle of incidence (mininum flight path), the right curve for incidence at 30° (maximum flight path).

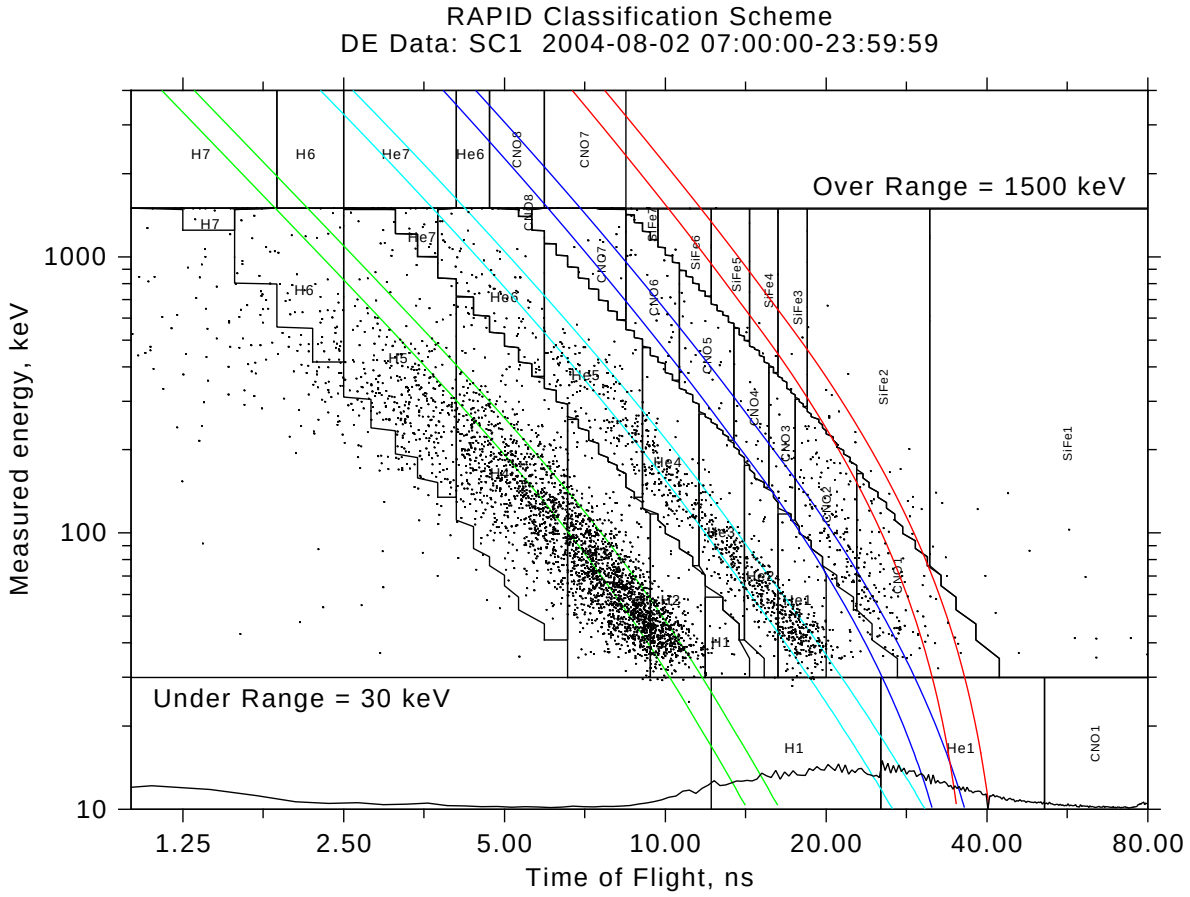

Figure 2: Same as Figure 1 but with actual direct events plotted. In the underrange region, there is no energy measurement, so only relative counts are plotted as a curve.

Note: the sharp drop in the underrange counts between H1 and He1 in Figure 2 is a result of the prioritization system that enhances helium over hydrogen, rather, it enhances everything in the He bins over that in the H bins.

The direct events offer an insight into the raw input before classification. In this sense, *classification* means determining the mass and energy bin for each event. The main output is actually the counts in each of the bins. The direct events are used here to investigate what is really going on at a low level.

2.1 The directional efficiency

When events are classified, they are added to various counters. These are:

HSPCT/ISPCT omnidirectional hydrogen, helium, CNO counts each in 8 energy bins; these are sorted regardless of the directional information.

I3DD sorted by species, sector, polar direction, and energy bin; these clearly require a full directional signal.

A full directional signal is produced by the start time-of-flight pulse, which can be located to within a quarter of the sensor head's 60° range. However, this determination can be ambiguous or otherwise missing, in which case only the head information is given to the event. Such events contribute to HSPCT and ISPCT but not to I3DD. Thus the summed I3DD counts can be significantly lower than

the omnidirectional ones. The difference depends on the *directional efficiency*, the probability that a TOF signal is accompanied by a full directional signal.

This has long been recognized, but the puzzling feature was that, although the summed I3DD counts were 2–3 times lower than H/ISPCT in energy channels 2–8, in channel 1 they were very nearly the same. Why is the directional efficiency apparently near 100% only in channel 1?

Direct event analysis can give us the ratios of events with and without a full directional signal, to give us the directional efficiency. However, it turns out that for events in the underrange region, without an energy signal, there is either a full directional value or no value (0–11 or 15), never a head-only value (12–14). Otherwise, the directional efficiency for these events is similar to with energy signal, when one compares the number with full information to the total. (For example, on August 11, 2002, SC 1, H and He show 61% with full direction in the underrange, and 49% in the middle range; CNO has 55% and 35%, respectively.)

The head-only information can in principle be determined several ways, even from the ambiguous start signal. Especially in serial mode, only one head at a time is being processed, so the time phasing is sufficient. However, it would appear that for event classification, it is the energy signal itself that is used to specify the head when the full direction is missing. If this is so, then it is no wonder that the underrange region is lacking head information—there is no energy signal here. And direct event analysis confirms this.

> **Consequence:** the directional efficiency can be routinely found from direct events, which can be used to correct the I3DD counts against the omnidirectional H/ISPCT ones, which are being registered at 100% for energy bins 2 and above. (At the moment, this factor is fixed at 2.5.) *However, contrary to present practice, the channel 1 values for both I3DD and the omnidirectionals must be similarly enhanced.* That means, we have been underestimating the ion fluxes in the first channel by a similar factor.

2.2 The true range of bin H2

We have long been aware that the counts in the second hydrogen channels are surprisingly lower than those in the third, something that is unexpected from the normal spectrum. That the width of this channel is narrower than most (28 keV compared to 73 keV for the 3rd) offered some explanation. However, when converted to differential flux, this channel was still lower than the 3rd at many times.

Examining the theoretical (green) hydrogen curves in Figure 1, one sees that they cut the 30 keV limit in the middle of the H2 bin. This suggests that this bin is not really seeing the full width for which it is intended.

Figure 3 shows the distribution of direct events by TOF channel for the first four hydrogen channels. Indeed, the H2 bin (upper right) exhibits a distribution sharply skewed to the lower channels (higher energies). The counts in channel 34 are only 1/3 of those in 30–32, and the others drop off even further. (Note that the vertical scale is logarithmic.) This means, bin H2 is effectively measuring only TOF channels 33–30, corresponding to incoming energies 80–95 keV, rather than the presumed range of 68–95 keV. This halves the actual energy range, which would double the differential flux.

> **Consequence:** the energy ranges must be specified with a lower and upper limit to allow for the gap in the hydrogen spectrum. So far, only lower thresholds are applied, with the assumption that the energy ranges are contiguous. The handling of energy ranges must be redone completely.

2.3 Varying time-of-flight distances

There is a second aspect to the above analysis. To date, the energy levels have been determined once for each ion head, based on ions flying through a mean TOF distance of 34 mm. However, the

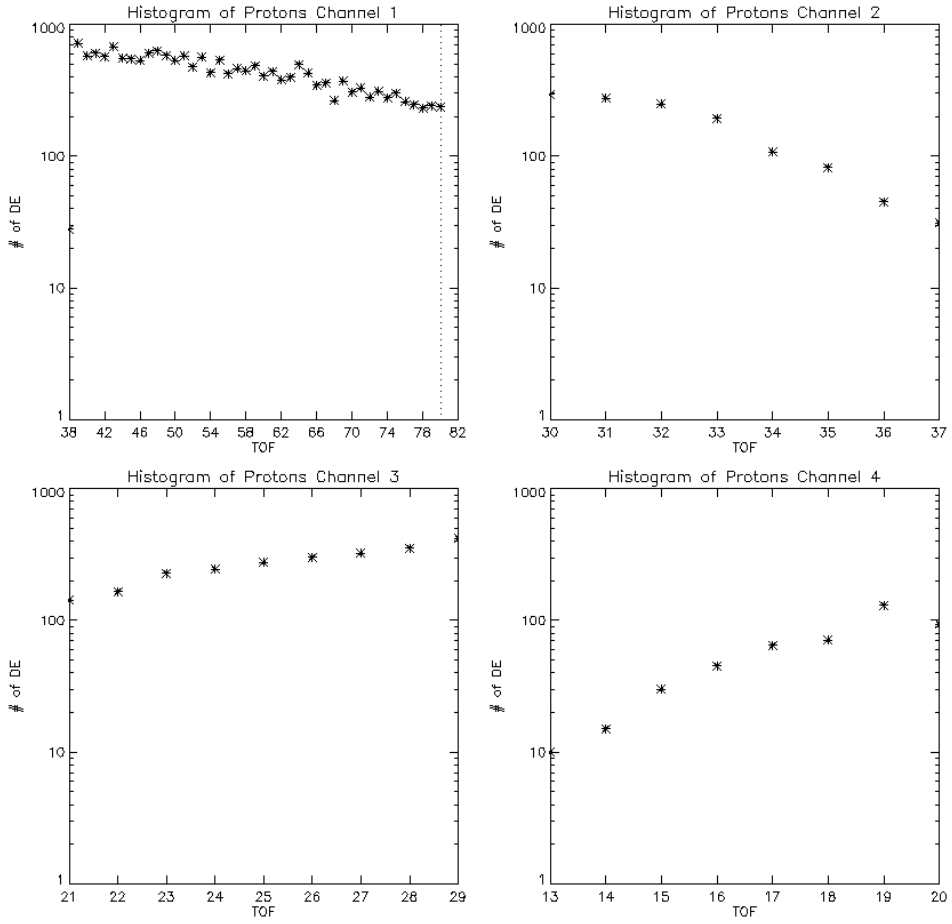

Figure 3: Direct events from SC 1, on August 11, 2002, UT 14:45–15:45, showing the distribution by TOF channel for the first four hydrogen bins.

IIMS heads can detect ions entering at angles from 0° to ±30°. In fact, the counts are sorted with the full directional information into four 15° segments. At 30°, the effective flight path is increased by 15%, meaning the TOF measurement is correspondingly higher for the same particle. (This is the reason for the two theoretical curves per species in Figure 1, the right one being shifted 15% from the left one.)

Evidence for this effect is given in Figure 4, where the direct events are sorted according to directional signals for the inner and outer segments. If the inner segments have a mean incoming direction of 7.5° and the outer ones 22.5°, then their TOF spectra should be shifted by 7%, the outer ones more to the left. Bin H2 provides a means to test this, because of its skewed distribution. Indeed, the outer distribution (brown diamonds) is about two channels to the right of the inner one in the upper right panel of Figure 4 (7% of 34 is ~2). This effect is also evident in the lower right panel, bin H4. In this case, the uneven distribution is caused by the ion spectrum, not by detector efficiency.

> **Consequence:** the energy ranges must be calculated for all 12 directions on each spacecraft, and not just for each head, as has been done so far.

2.4 Split bins

One sees from Figure 1 that the first energy bins are split between the underrange and middle regions. How does this complicate the geometry factor calculations and energy ranges?

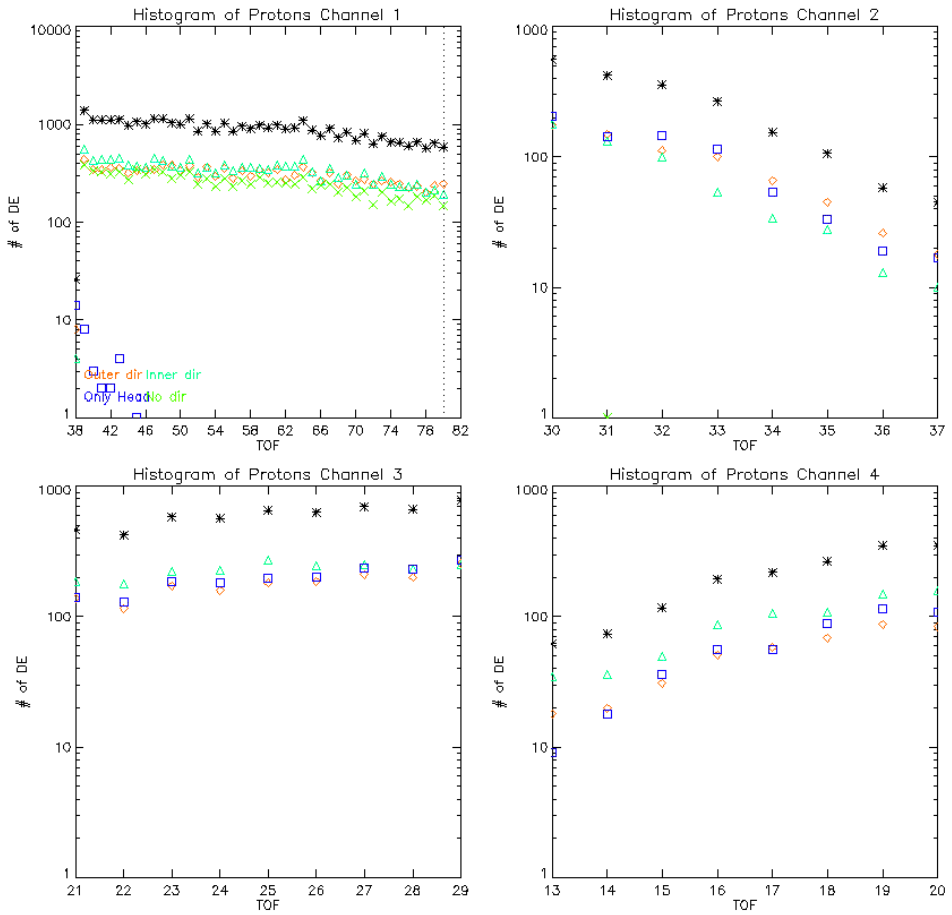

Figure 4: Direct events from SC 1, on November 18, 2003, UT 00:00–12:00, showing the distribution by TOF channel for the first four hydrogen bins, but sorted by segments within each head: brown diamonds for the outer segments, green triangles for inner segments, blue squares for head-only, cyan crosses for no directions. The black crosses show the total counts.

Figure 1 also indicates that the upper part of H1 rather misses the theoretical curve almost completely. In fact, there is only one TOF channel difference between the upper and lower parts, channel 38. In the upper left of Figure 3, we see that single channel at the far left, more than an order of magnitude below the others. So, yes, this part could really be neglected.

For He1, the case is different. There is even a gap between the upper and lower parts, seen very clearly in the direct event distributions for helium in Figure 5, upper left panel. The TOF channels 52 to 63 are the middle region, with energy, while 81 to 163 are the underrange. The large difference between the two parts of the TOF spectrum is a strong indication that there is considerable hydrogen contamination in He1, the topic of Section 2.5.

2.5 Hydrogen and helium separation

For the events with energy, the separation between the species seems fairly good, as seen in Figure 2. But is this always so?

For example, the underrange border between H1 and He1 is at 25 ns. The theory is that any hydrogen ions with longer flight times has such low energy that it does not trigger the TOF mechanism, while any helium ion with shorter times has energy over the 30 keV limit and thus appears with an energy signal (not underrange). However, both these assumptions appear to be false. It seems that

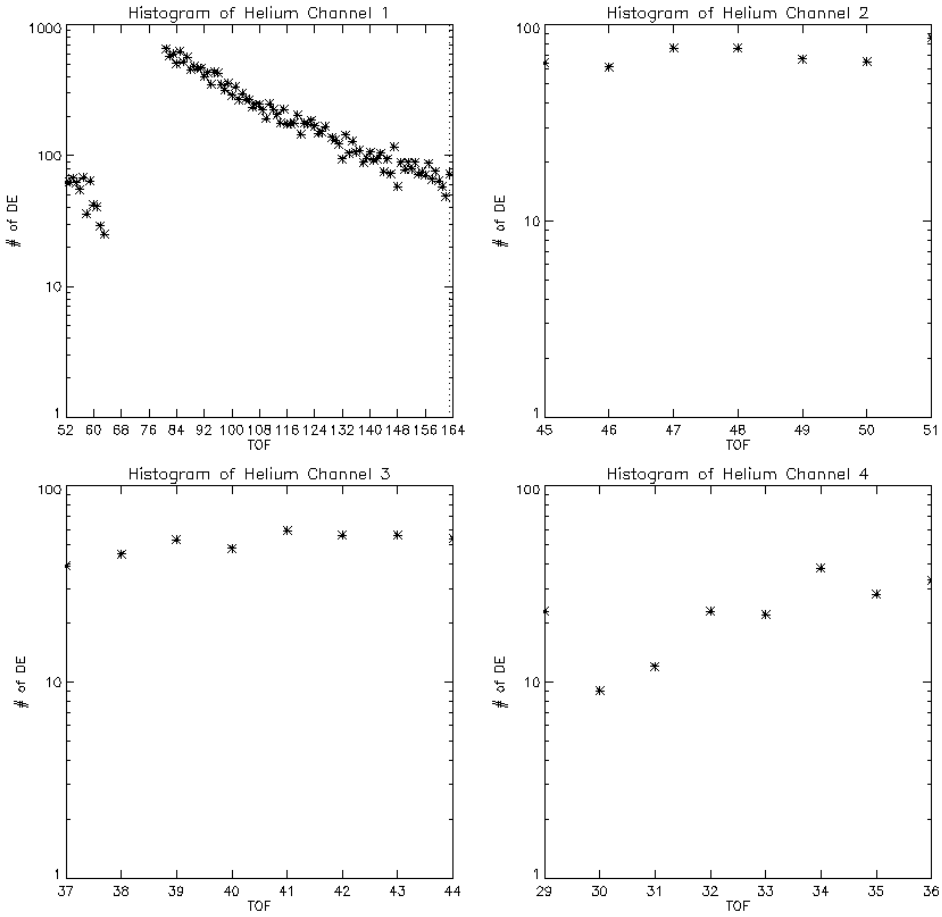

Figure 5: Direct events from SC 1, on August 11, 2002, UT 14:45–15:45, showing the distribution by TOF channel for the first four helium bins.

considerable numbers of hydrogen do enter the He1 bin, while Figure 1 indicates that the 30 keV limit is reached more at 20 ns for helium, well inside the H1 bin.

As mentioned in Section 2.4, the upper left panel of Figure 5 indicates a strong hydrogen contamination in He1. In fact, probably most of the events in the underrange He1 are hydrogen. The effect of helium in H1 can probably be neglected since one expects considerably less helium ions than H anyway.

In order to test how much hydrogen is entering He1, we have looked for times when helium is hardly present, as indicated by the He2 and higher bins. Such a case occurs on January 3, 2002, from 12:00 to 12:10. Here we count a total of 223 counts in H1 and 160 in He1. If these were all hydrogen erroneously in He1, then that represents 71% of the H1 value.

Another check: the efficiency for the time-of-flight measurement decreases rapidly with decreasing energy. For hydrogen, this efficiency is near zero for 19 keV, corresponding to TOF channel 120, well inside the He1 bin near 37 ns. At channel 80, the division between H1 and He1, that efficiency is 20% of its maximum. (At the left side of H1, the efficiency is about 70% maximum.) The ratio of the integrated hydrogen TOF efficiencies within the He1 and H1 bins is 19%. This would suggest that hydrogen ions in the He1 range should have low probability of being registered. However, the spectrum also rises strongly at the lower energies. Combining the TOF efficiency curve with a power law spectrum of $\gamma = 4$ and integrating over the H1 and He1 energy ranges, yields a ratio of He1 to H1 of 45%. (However, the H1 distributions in Figures 3 and 4 show a flatter spectrum, reflecting the TOF efficiency itself.)

The above argument is only suggestive, since it depends on the spectral shape and on the precise efficiency curve, but it does indicate that indeed an appreciable fraction of the H1 counts can also be inside the He1 bin, and since the helium contribution is expected to be small anyway, most of what is seen in He1 is likely to be hydrogen.

> **Consequence:** one could try to estimate a He1 to H1 ratio for hydrogen to be used to correct the helium content of He1. However, the error bars would be greater than the remaining count. Perhaps the best thing is to give up on He1, and start the helium spectra with He2. This is certainly what will be done with the CSDS Prime Parameters.

3 Spectral Corrections

Each of the different sensors and their subsections actually have different energy limits. This statement applies to both the ion IIMS and the electron IES parts of RAPID. It therefore makes sense to try to correct the measurements to bring them to a uniform set of energy thresholds.

This is already available in the SCI file data, as an option, at least for HSPCT and ESPCT. It must also be applied to the 3D products. Furthermore, with IES, there is the additional complication that the entire spectrum shifts with count rate as the pedestal moves (*pedestal shift*).

Martin Carter, formerly of RAL, has routines to determine and correct for this shift. These must be incoporated into the `msf2sci` software.

3.1 Effective energy of a bin

The correction system that I have devised to date fits the integral flux in two bins to a power law, and then integrates that over the energy range to be shifted. The tricky thing is to determine the power law.

For a single bin, all we have measured is the integrated flux between two energies, E_1 and E_2; dividing by $(E_2 - E_1)$ produces a differential flux, which is the mean differential flux between E_1 and E_2. The question then is, to what energy does this mean flux apply? This really depends on the power law, on the spectral index.

Let us represent the differential flux as a function of energy ε as

$$j(\varepsilon) = \mathcal{A}\varepsilon^{-\gamma} \tag{1}$$

where γ is the spectral index and $\mathcal{A}$ is a constant. The integral flux between two energy thresholds is then

$$\begin{aligned} J(E_1, E_2) &= \int_{E_1}^{E_2} \mathcal{A}\varepsilon^{-\gamma}\, \mathrm{d}\varepsilon \\ &= \frac{\mathcal{A}}{\gamma - 1}\left[{E_1}^{-\gamma+1} - {E_2}^{-\gamma+1}\right] \end{aligned} \tag{2}$$

Let $E_m = (E_2 + E_1)/2$ and $\Delta = (E_2 - E_1)/2$ and $\delta = \Delta/E_m$. Then

$$\frac{E_1}{E_m} = 1 - \delta \qquad \text{and} \qquad \frac{E_2}{E_m} = 1 + \delta$$

We then use equation 2 to get the mean differential flux as

$$\begin{aligned} \frac{J}{E_2 - E_1} &= \frac{\mathcal{A}}{\gamma - 1}\,\frac{(1-\delta)^{-\gamma+1} - (1+\delta)^{-\gamma+1}}{2\delta}\,{E_m}^{-\gamma} \\ &= \mathcal{A}\,{E_F}^{-\gamma} \end{aligned} \tag{3}$$

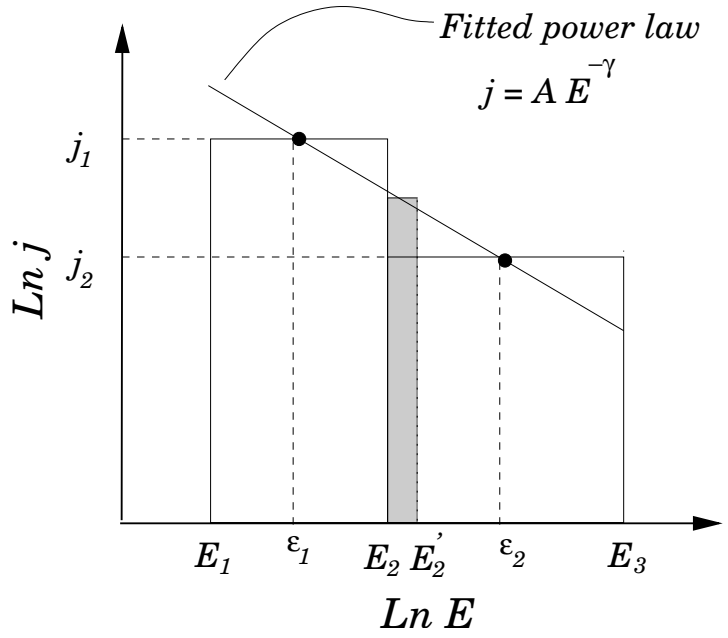

Figure 6: Demonstration of bin correction: a power law is fitted between the two bins with energy limits $E_1 \rightarrow E_2$ and $E_2 \rightarrow E_3$, with integral fluxes $J_1 = j_1 \cdot (E_2 - E_1)$ and $J_2 = j_2 \cdot (E_3 - E_2)$ respectively. The effective energies are ε_1 and ε_2. The new energy limit between the two bins is to be E_2'. The fitted power law is integrated between E_2 and E_2' (gray area), and this is added to J_1 and subtracted from J_2.

Equation 3 is the definition of the effective energy E_F. It is the energy for which equation 1 has the value equal to the mean diffential energy. The only problem is, it depends on the spectral index γ (but not on $\mathcal{A}$).

We can expand equation 3 as follows:

$$\begin{aligned} \left(\frac{E_F}{E_m}\right)^{-\gamma} &= \frac{(1-\delta)^{-\gamma+1} - (1+\delta)^{-\gamma+1}}{2\,(\gamma-1)\,\delta} \\ &\approx 1 + \frac{\gamma(\gamma+1)\,\delta^2}{6} + \mathcal{O}(\delta^4)\cdots \\ \text{or,} \quad \frac{E_F}{E_m} &\approx 1 - \frac{(\gamma+1)\,\delta^2}{6} + \mathcal{O}(\delta^4)\cdots \end{aligned} \tag{4}$$

As an example, even when the energy bins are very broad (say $E_2 = 2 \cdot E_1$) and the spectrum steep (say $\gamma = 5$), the right hand side of equation 4 yields 0.9, meaning E_F can be well approximated by the mean energy E_m.

3.2 Power law fit

The spectral correction requires a power law fit between two adjacent energy bins. The procedure is straight forward:

1. get the mean differential flux for each bin, j_1 and j_2, and assign them to effective energies ε_1 and ε_2, the mean energies of each bin;
2. find the $\mathcal{A}$ and γ for these two points that fit equation 1

$$\gamma = \frac{\ln(j_1/j_2)}{\ln(\varepsilon_2/\varepsilon_1)} \qquad \mathcal{A} = j_1 \cdot \varepsilon_1^{\gamma}$$

3. use γ to refine ε_1 and ε_2 with equation 4;
4. reiterate point 2 until required convergence is met.

3.3 Bin correction

Having found the power law connecting two adjacent bins, they may be corrected for the ideal energy boundary between them by integrating the power law with equation 2 over the energy shift. The integral is then added to the one bin and subtracted from the other (Figure 6). Note that this is a correction to the *integral* flux; and new differential flux must then be calculated with the corrected energy limits.

If the energy bins are not contiguous (like H1 and H2), then the correction is still possible, but each boundary must be shifted on its own.

The procedure is to handle the bins in pairs, fitting a separate power law to each pair and carrying out the correction between them. Only the first energy bin is a problem. Here, the power law must be found between bins 1 and 2 and then extrapolated to the lower energy end. (The upper bin can be neglected, as it rarely has non-zero counts and its upper threshold is not well defined to begin with.)

4 Conclusion

Based on the knowledge obtained about the workings of RAPID, in particular IIMS, new procedures will have to be created for handling the energy thresholds and converting to differential flux.

A new set of calibration factors must be employed. Some of these will require constant checking to obtain the current values for any time. I would envisage a set of values for every month.

I would hope to have a new functioning system working by the beginning of 2005.

A The Current Calibration Factors

I list the current calibration parameters here for reference. New sets will have to be worked out.

A.1 The IIMS calibration set

The current set of IIMS calibration parameters consist of:

- conversion factors for 8 energy channels, 12 polar segments, and 3 species, for a total of 288 values;
- background count rates, for 12 polar segments, 16 sectors, 3 species, and 2 modes (serial/parallel), for 1152 values;
- omnidirectional factors for 8 energy channels and 3 species (a multiplying factor for `HSPCT` and `ISPCT`), for 24 values;
- energy thresholds for 8 channels (plus a 9th as the upper limit), 3 IIMS sensors, and 3 species, for 81 values;
- energy masks (values 0 or 1) for forming the CSDS omnidirectional fluxes, for 8 channels, 2 energy ranges, and 3 species, for a total of 48 values (used only for IFF files, not SCI files);

- triggering modes allowed, of which there are 6 (numbered 0–5); 6 parameters are given, one per mode, = 0 for those that are not allowed for this calibration set;
- high voltage information: for each of the three high voltages (`STA`, `STO`, `DEF`) 3 parameters are given; the first is the default step number (0–15) which is only used by the main routines to determine the configuration mode; the other two are the minimum and maximum values of the voltage that are compatible with the calibration set, stored as coded numbers between 0–255;
- timing values: the live times for the 3 heads in serial mode, the single live time in parallel mode, and the dead time, for 5 values.

This makes a grand total of 749 parameters for each IIMS calibration set.

Summarizing the calibration parameters in terms of Fortran `dimension` statements:

			Number
Conversion factors	:	`CF(` N_{chan} , N_{polar} , N_{spec} , −)	288
Background rates	:	`BG(` N_{polar} , N_{sect} , N_{spec} , N_{mode})	1152
Omnidirectional factors:		`OF(` N_{chan} , − , N_{spec} , −)	24
Energy thresholds	:	`EL(` $N_{\mathrm{chan}}+1$, N_{sensor} , N_{spec} , −)	81
Spectral masks	:	`EM(` N_{chan} , N_{range} , N_{spec} , −)	48
Trigger modes	:	`TR(` 6 , − , − , −)	6
High voltage	:	`HV(` 3 , 3 , − , −)	9
Timing values	:	`TM(` 5 , − , − , −)	5
		Total:	1613

Note that $N_{\mathrm{range}} = 2$, the number of energy ranges for the CSDS output omnidirectional fluxes, and $N_{\mathrm{sensor}} = 3$, the number of IIMS sensor heads. For IIMS, $N_{\mathrm{chan}} = 8$, $N_{\mathrm{polar}} = 12$, $N_{\mathrm{species}} = 3$, and $N_{\mathrm{mode}} = 2$.

Note also that it is the leftmost index that varies the most rapidly, as in standard Fortran.

A.2 The IES calibration set

The current set of IES calibration parameters consist of:

- conversion factors for 8 energy channels, 9 polar segments, and 4 integration times, for a total of 288 values;
- background count rates, with the same dependencies as the conversion factots, for another 288 values;
- energy thresholds for 8 channels (plus a 9th as the upper limit), 9 polar segments and 4 integration times, for a total of 324 values;
- energy masks (values 0 or 1) for forming the CSDS omnidirectional fluxes, for 8 channels, 2 energy ranges, and 4 integration times, for 64 values (IFF only);
- pedestal widths in keV for 9 polar segments and 4 integration times, for 36 values; these widths are not currently used but were provided for future developments that might want to track pedestal shifting;
- PAD masks (values 0 or 1) specifying which of the 8 energy channels are included in the 2 `EPAD` channels, for 16 values (IFF only).

This makes a grand total of 1016 parameters for each IES calibration set.

Summarizing the calibration parameters in terms of Fortran `dimension` statements:

		Number
Conversion factors:	CF(N_{chan} , N_{polar} , N_{lut})	288
Background rates :	BG(N_{chan} , N_{polar} , N_{lut})	288
Energy thresholds :	EL($N_{\text{chan}} + 1$, N_{polar} , N_{lut})	324
Spectral masks :	EM(N_{chan} , N_{range} , N_{lut})	64
Pedestal widths :	PW(N_{polar} , N_{lut} , $-$)	36
EPAD masks :	PM(N_{chan} , N_{epad} , $-$)	16
	Total:	1016

Again $N_{\text{range}} = 2$, the number of energy ranges for the output omnidirectional fluxes. For IES, $N_{\text{chan}} = 8$, $N_{\text{polar}} = 9$, and $N_{\text{lut}} = 4$, the number of integration times, or *look-up tables*. Finally, $N_{\text{epad}} = 2$, the number EPAD energy channels.