Method of analysis

For each run, a RD27 Fortran program read the raw data offline, made a calculation (as described in Sec. 2) of the pedestal mean and rms (root mean square) for each of the 6×6 channels, and appended the result to a disk file. This *pedestal file* also contained the run parameters such as clock period, trigger thresholds, beam cell coordinates etc., as recorded in the paper logbook. These run parameters are summarized in an "electronic runlog", available along with the pedestal files on CERNVM. Appendix A.1 contains excerpts from the runlog and comments on the runs not included in the pedestal analysis.

The main subject of the first part of the Summer Student work was to look for patterns in the pedestal data, especially for sudden changes from one run to the next, and for various correlations. This was done by writing a Fortran program that read a merged file of pedestal files and wrote most of the information into a PAW/HBOOK Ntuple (see App. A.2 and the references [4, 5]). Then PAW was used on CERNVM to look at the data in the Ntuple. The analysis ended up in a memo E-mail on July 26, summarizing some characteristic features of the data [7]. Information from the paper logbook about cabling etc. which might explain these features was sent back in a reply E-mail [6] from the U.K., where the logbook resided. The two E-mails form the basis for sections 1 and 2, except 2.2.

The second part of the work was chosen to consist of a close look at almost the full raw data from 127 selected events of run 607, again collecting the data into a Ntuple (see App. A.3). The purpose was to find where the high pedestal rms's originated from, and also to gain an initial feeling of hit flags and energy sums. The results are described in section 2.2 and section 3. The reason for selecting run 607, of all runs, was twofold: it contains one of the dozen exceptionally high rms's (5.3) in one of the channels; and it belongs to the group of runs with the highest non-exceptional rms's (1.9–2.1) in certain channels, as described in Sec. 2.3.

During the Summer, many plots and other paperwork was produced. It is available for examination at the office of Nick Ellis at CERN Meyrin. Most of the plots required more than one PAW command to be made. All these plots were created by the KUIP MACros (KUMACs) listed in Appendix A.4.

Labeling & drawing channels

The 6×6 trigger channels covered by the Cluster-Finding Module (Fig. 1) are labeled (i, j) by RD27 and drawn such that i increases with η in the right-hand direction, and j increases with ϕ upwards, $1 \le i, j \le 6$.

Initially the CFM was connected to the left half of the RD3 calorimeter, except the lowest row ($j = 1$), which was unconnected because of missing RD3 electronics. After run 562, the cables were moved to the right half of the calorimeter, and two more rows were unconnected, as shown on Fig. 3. The letter "P" marks such a "Passive" channel, while "A" labels an "Active" (connected) one.

The number in the range 1–3 after the "A" or "P" (e.g. in "A1" or "P2") indicates to which of the three 12-channel ADC cards the channel belongs.

The 6 *emphasized* channels, (5,3)–(6,5), were analyzed with most care during the first part of the work, because their pedestal rms's after run 562 turned out to be particularly high and correlated. This result — which will be substantiated in the next section — was announced in [7] without knowledge of the connectivity. That knowledge was obtained by means of [8] and extended by [9].

6	P1	P1	P2	P2	P3	P3
5	A1	A1	A2	A2	*A3*	*A3*
4	A1	A1	A2	A2	*A3*	*A3*
3	A1	A1	A2	A2	*A3*	*A3*
2	P1	P1	P2	P2	P3	P3
1	P1	P1	P2	P2	P3	P3
$\uparrow$ j $i \rightarrow$	1	2	3	4	5	6

Figure 3: Definition of trigger channel groups A1–A3 & P1–P3

In summary, the emphasized channels are characterized as the "connected [or Active] channels of ADC card #3" after run 562, or simply the "A3-channels".

The runlog gives the beam spot for each run in calorimeter cell coordinates. The trigger cell (i, j) contains the $4{\times}4$ calorimeter cells given as $(4i-3 \ldots 4i, 4j-3 \ldots 4j)$ until run 562. After that, 24 must be added to the first calorimeter coordinate, because of the cabling movement "to the right" mentioned above. For example, the run 607 beam spot is given as "(43,11)". It thus falls into trigger cell (5,3).

2 Pedestal analysis

The pedestal analysis is based on data from 398 runs in the run number range 172–618. This is virtually all those runs for which a pedestal file exists, as detailed in Appendix A.1.

The pedestal means and rms's given in the pedestal files have been calculated from time slice 3+4+5 (1 slice = 12 to 25 ns) — before the rising edge of the pulse that triggered the read-out — from the ADC-values in the first n events of each run. n is either equal to 1000 or equal to the number of events in the run, whichever is smaller. (Some runs have only a few events).

A short general description of the pedestals was given in the memo of July 26 [7]:

> The means range from 18.2 to 34.2, but vary much less within each channel. The rms's are mostly stable around a value less than 1.0, typically 0.5, with a few high ones up to 7.03.

Only 9 of the 14328 rms's were above 4.0. (398 runs $\times$ 6 $\times$ 6 channels $\times$ 1 rms per channel = 14328 rms's). Most of these can probably be explained by double events (Sec. 2.2.1).

2.1 Characteristics of pedestal rms's vs run number

An example of a plot of the pedestal rms versus the run number is shown on Fig. 4. The run number ranges explicitly indicated on the figure will be discussed in Sec. 2.3.

The 36 such plots were described by the memo as follows:

> 1) The 36 plots of pedestal rms vs run can be divided into 6 groups characterized by their behaviour in the end of the plot, approximately (or perhaps precisely) after run 562. I was told that the ADCs were connected exactly in 6 such groups, but not

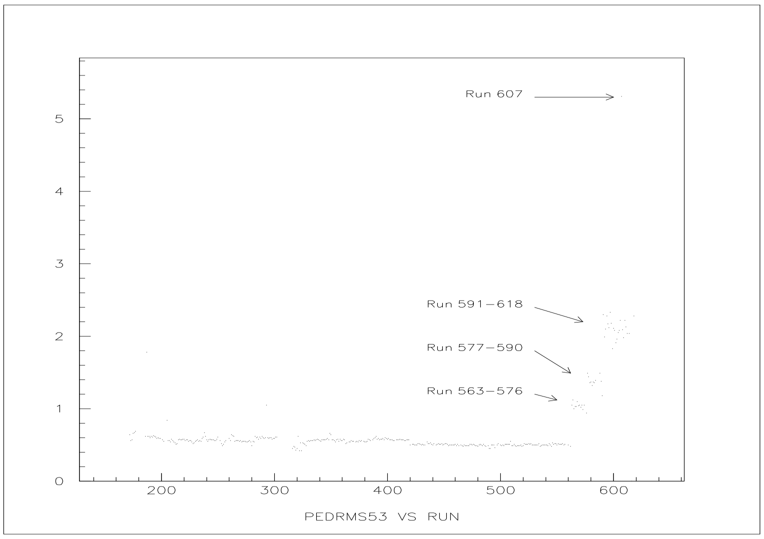

Figure 4: Pedestal rms vs run for trigger channel (5,3)

> which ADCs belonged to which groups. From my plots I can make a very qualified guess on the connections ...

The division into ADC boards and connected/unconnected channels after run 562, as given on Fig. 3 (Sec. 1), matches exactly the guess given in the memo. The memo's descriptions of the 6 channel groups are reproduced in table 1, using the channel group labels of Fig. 3.

Channel group	*Pedestal rms description (run 563–618)*
P1	Slight jump down to new plateau near run 600
A1	A slight rise in the end
P2	A slight "bump" up and back to normal
A2	A clear rise during the interval (563–end)
P3	As above, but slightly higher rise
A3	Very clear rise, ending 3–4 times over start

Table 1: Channel dependence of pedestal rms run-by-run development

Note that Fig. 4 shows one of the six A3-channels. So, after the movement of the trigger equipment to a new set of RD3 electronics, something started to go wrong in the channels handled by ADC card #3 and by the connected part of ADC card #2. (More on the hardware origin in Sec. 2.3). As further evidence of a common cause of the noise, a high correlation of the A3-channels was documented in the memo:

> To illustrate how closely these 6 rms's follow each other, we can note that for the interval $541 \leq$ run ≤ 618:
>
> rms(6,4) = 1.3 ± 0.6 (mean rms ± r.m.s. rms)
> rms(6,3)-rms(6,4) = 0.002 ± 0.07

2.2 Explanation of high rms's using a sample of raw data

Aiming at a deeper insight into the formation of the high pedestal rms's, the raw data for 127 events of run 607 was extracted for further analysis, as detailed in Appendix A.3. Three distinct effects which contribute to a high pedestal rms were found, having highly different characteristic time scales:

- Double events: Two electrons hit the cell within a few times 25ns
- Low frequency noise (occasional pedestal shifts lasting $\gg$ 25ns, but $\ll$ time between events)
- High frequency noise (background noise composed of periods $\ll$ 25ns)

The second effect in the list causes the channel–channel correlations.

2.2.1 Double events

Very rarely it happened that two particles hit the same trigger cell separated by only a few time slices. Event 575 in run 607 is an example of this (Fig. 5).

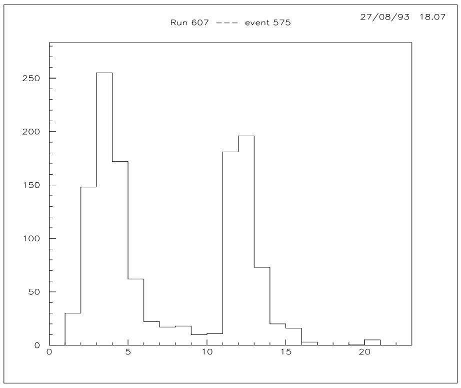

Figure 5: A double event

The peak around slice 12 is the normal one, expected from the RD3 scintillation trigger. The other covers the pedestal slices 3+4+5. The pedestal rms calculated from those slices over the first 1000 events was 5.3. Excluding this double event would give around 2.2, estimated from the rms of the first 100 events. The 2.2 is close to the rms's of the other 5 channels in the group (A3) in the runs 591–end, namely 1.9–2.1 (Fig. 4). So this event presumably explains all the discrepancy.

Knowing the mean time between the arrival of two particles to the calorimeter, it is possible to estimate how often they hit so closely in time that the effect on the calculated rms is observable. This will be attempted in the following, to achieve an order of magnitude estimate.

An upper limit for the mean time between events can be found from the number of events recorded on tape during 24 hours. The runlog entries for 07 May 1993, runs #563 to about #590 (ref. [9]), add up to a total of 1.3×10^5 events. 24 hours equals 0.9×10^5 seconds. These numbers give one event per second. The mean time between events was thus less than one second, by a factor that depends on e.g. the actual length of the data-taking day (sometimes only 8 hours or less), the dead time between runs and the amount of bunching/bursting used in the beam. An event rate of 10Hz will be tentatively assumed here. Additionally, the events are regarded as independent, with a uniform statistical time distribution (within each burst, possibly).

The first pulse in Fig. 4 would still have affected the calculated pedestal rms if it were shifted just a few time slices. The width of this critical time window can be taken to be about 5 time slices, or on the order of 100ns with a clock period of 25ns.

For a given pulse, the probability for another pulse to arrive within the critical time window defined by the first one is thus: 10Hz $\times$ 100ns = 10^{-6}.

The number of events used in the pedestal rms calculation is 1000 per run, in most of the 398 runs, as mentioned earlier. This gives a total of 4×10^5 events.

All in all, the expected number of significant double events becomes: $10^{-6} \times (4 \times 10^5) = 0.4$. This is somewhat lower than the observed number of unusually big pedestal rms's: 12 rms's above 3.0, including 9 rms's above 4.0 (up to 7.03). The expectation would match the observation better, if the event rate within a burst was assumed to be a factor of 20 higher, i.e. 200Hz.

The event on Fig. 5 seems to be extraordinarily unusual in another respect: The unexpected peak appears at an *earlier* time than the expected one. (Otherwise it wouldn't have been found, since it wouldn't affect the pedestal rms). The first electron should normally fire the RD3 scintillation trigger. Thus the second one should appear after the peak at slice 12 rather than before. Perhaps a third electron caused a trigger shortly before the first, at a time such that the read-out deadtime window happened to end *between* the two peaks of Fig. 5. Inefficiency of the RD3 scintillation trigger is another possibility.

If double events are the explanation for the other extremely high rms's too, then the runs involved should have the beam spot calo cell in exactly the high-rms trigger cell. A check of the 9 rms's above 4.0 gives the results summarized in table 2, where the beam cell coordinates

rms	*trig*	*run*	*calo*	*Containment?*	
7.03	(1,4)	455	(3,14)	yes	
5.49	(2,4)	492	(9,13)	almost,	calo cell touches that trigger cell
5.31	(5,3)	607	(43,11)	yes,	this is the well-analyzed one
5.18	(3,3)	519	(13,10)	almost,	calo cell touches trigger cell
4.60	(2,4)	462	(10,14)	no,	1 calo cell between them
4.56	(4,4)	422	(17,16)	almost	touches
4.15	(5,4)	435	(19,15)	yes	
4.07	(6,5)	427	(22,16)	almost,	touches
4.03	(2,4)	448	(7,15)	yes	

Table 2: Does the high-rms trigger cell contain the beam calorimeter cell?

come from the electronic runlog (App. A.1). It seems promising, but not perfect. Could stray particles cause the mismatch by missing the nominal target calorimeter cell by a little bit? Hardly. No stray hits were found in the 127 selected events of run 607; so the probability of *two* stray particles conspiring to make a double event is vanishing. The question remains open. It would probably be solved by a look at for example run 462.

2.2.2 Low frequency noise

Once in a while, an event is shifted up or down by as much as 10 ADC counts. The pedestal *within* the event varies much less. Figure 6 shows a normal event, a low pedestal event and a high pedestal event, superimposed on the same plot, for each of the 6 connected channels on ADC card #3.

The figure illustrates the fact that a given event is shifted nearly the same amount in all channels of a given 6-channel group. This explains why the 6 rms's follow each other so closely from run to run. The peaks in the channels other than (5,3) are interpreted as due to cross-talk, as further described in Sec. 3.1.

The size of the shift for the events of Fig. 6 is summarized in table 3. The magnitudes

Channel group	*Pedestal shift*	*(typical magnitude in ADC counts)*
#1	no observable shift	
P2	small shift	(high +0.5, low –3)
A2	a bit larger	(high –2, low –4)
P3	even larger	(high +9, low –5)
A3	maximal shift	(high +10, low –10)

Table 3: Channel dependence of pedestal shifts in run 607

listed here are different for other events. However, there is reason to believe that the ordering is general: It matches well the one according to rms's listed in table 1 (Sec. 2.1).

It may be questioned whether *low frequency noise* is an appropriate name for this pedestal shift effect. The plot of the event-by-event pedestal versus event number for the small 100-event sample is generally flat, with the shifted events scattered rather unsystematically in size and time. Of course, an entirely regular pattern cannot be expected, since the time from one event to the next varies randomly. Still, the time scale of the noise is much less than the mean time between events. Perhaps *medium frequency noise* would be a better name.

2.2.3 High frequency noise

High frequency noise is a suitable term for the variations within the pedestal slices of a single event. These variations are seen to be small on the connected part of ADC card #3 (Fig. 6), typically flipping back and forth between two neighbouring ADC values.

The unconnected channels on the three ADC cards are all even more silent within each event. However, the high frequency noise is considerable in the connected channels on ADC cards #2 and #1, covering 4 adjacent ADC values instead of only 2.

These conclusions are tentative, being based on 3 events only.

2.3 Correspondence of cabling with observed rms jumps

In the memo of July 26 [7], the three A3-wide pedestal rms jumps indicated on Fig. 4 were identified and described as in table 4. Similarly, the rise in the P3- and A2-channels (table 1) consists of smaller jumps at the same run numbers.

The paper logbook [8] revealed that changes of cabling took place exactly at the time of these jumps:

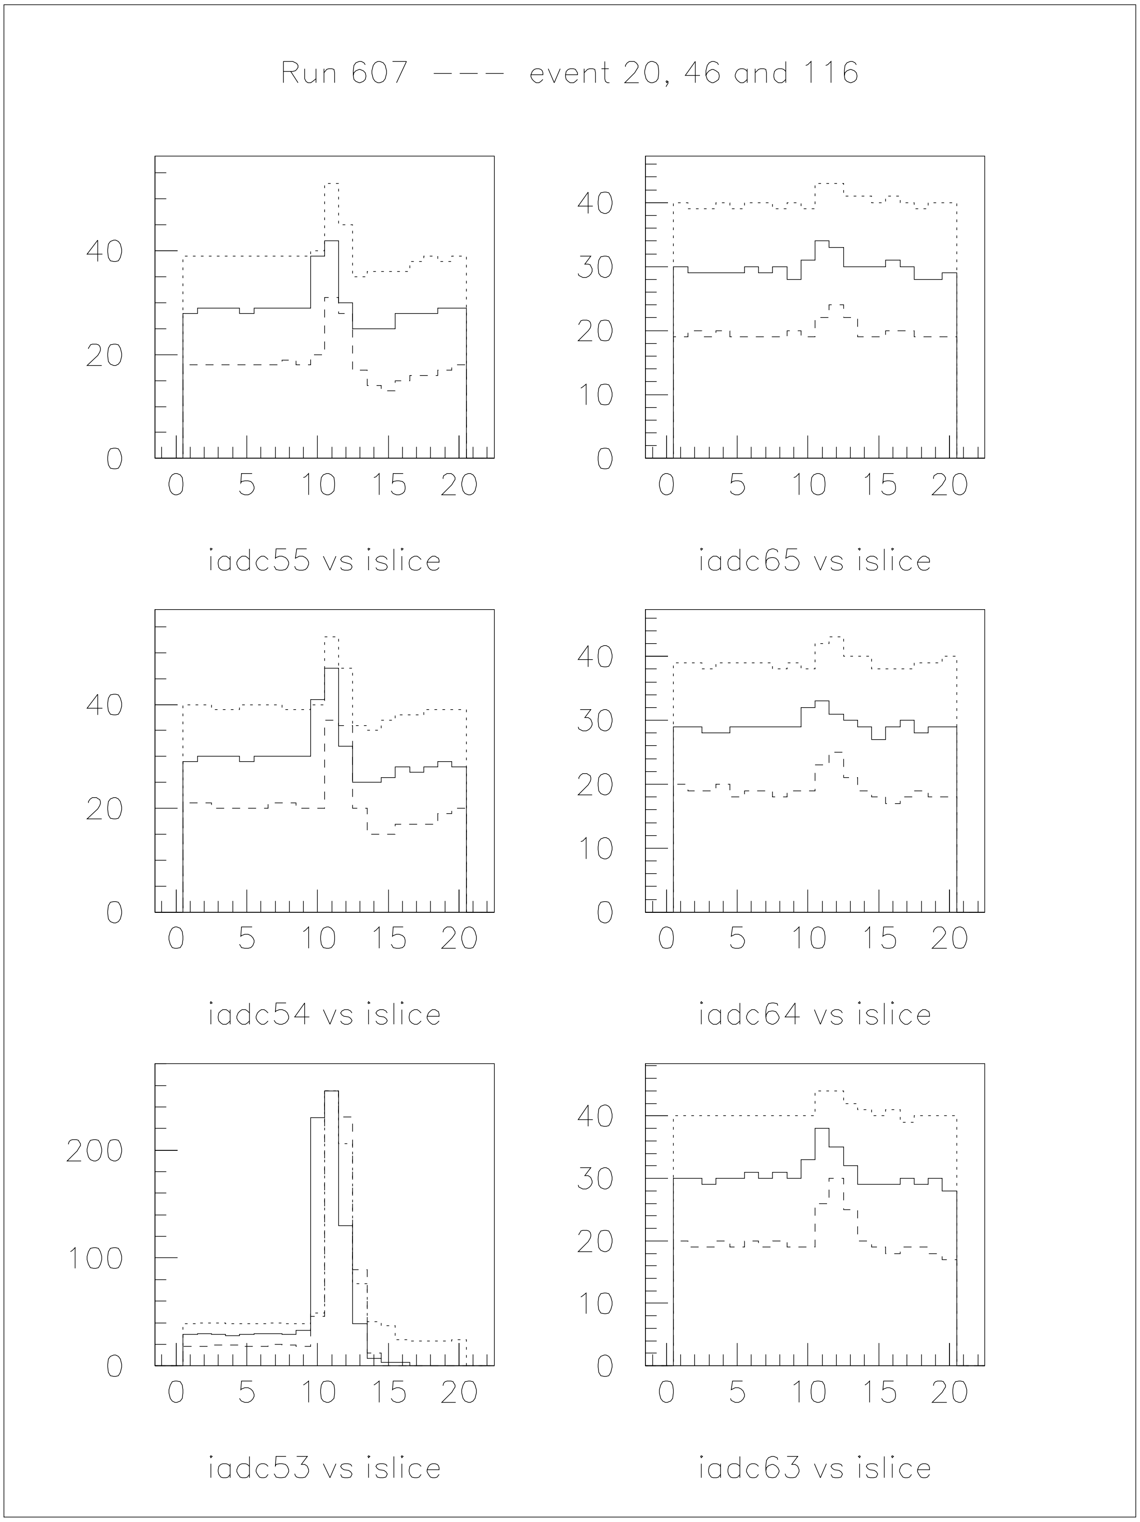

Figure 6: A3-channels of a normal, a low and a high pedestal event

run :	rms	(roughly, as seen from (run,rms)-plots)
–562:	0.5–0.65	(r.m.s. much less than 0.1 on this)
563–576:	1.0–1.1	(r.m.s. somewhat less than 0.1)
577–590:	1.4	(r.m.s. near 0.1)
591–618:	1.9–2.1	(r.m.s. 0.1–0.2)

Table 4: Pedestal rms jumps (magnitudes valid for the A3-channels)

- The jump from 0.6 to 1.0 coincides with the movement to different RD3 electronics (from Si based to GaAs based, in fact [9]), as already mentioned in Sec. 1.
- The jump from 1.0 to 1.4 happened when the analog summing of signals from the calorimeter was extended from 2 to all 3 depth samplings.
- The jump from 1.4 to 2.0 coincides with the change in cabling of the unused channels: They all became terminated in 50 Ohm at the inputs to the FADC modules (see Fig. 1). Prior to this, all FADC inputs were connected to the 40m BNC cables, via Lemo-to-BNC patch panels (not shown on Fig. 1) at each end of the 40m. For run 172–576, the 50 Ohm termination of all unused channels was done at the calorimeter end instead, i.e. in the end of the Lemo cables that would otherwise have been attached to the shaper modules. After run 576, some of these passive Lemos were moved to participate in the third depth sampling for the active channels. Thus, for run 577–590 some unused channels were terminated in the calorimeter end, while the rest of them were terminated at the fan-in inputs.

So apparently, each of these changes caused an increased low frequency noise on 2 of the 3 ADC cards. The noise must have originated in the FADC modules themselves, or at the entrance to them, since also the ADC-end-terminated channels in run 591–618 were affected (the P3-channels and to a small extent the P2-channels, see tables 1 and 3).

One feature from the memo, however, remains a mystery:

> The rms in channel (4,5) is considerably higher than in all other channels for run 201–302 & 387–562 (incl.), namely 1.43 ± 0.07 to be compared with the normal value of less than 1.0. (Note that the run after run 302 is run 316).

The two run number intervals are approximately defined by four changes of cabling (scaler connected; a NIM crate replaced and earthing improved; a faulty fan-in replaced; new RD3 electronics), but none of them single out channel (4,5). The third change almost does so: The faulty fan-in only did the summing for channel (4,3) and (4,5).

Perhaps some weakness in the channel (4,5)-electronics was exposed/suppressed by the other hardware changes.

2.4 Explanation of pedestal value characteristics

A typical plot of the run-by-run time development of the pedestal value in a single channel is shown on Fig. 7. Three characteristic features on the pedestal value plots were identified in the memo [7]:

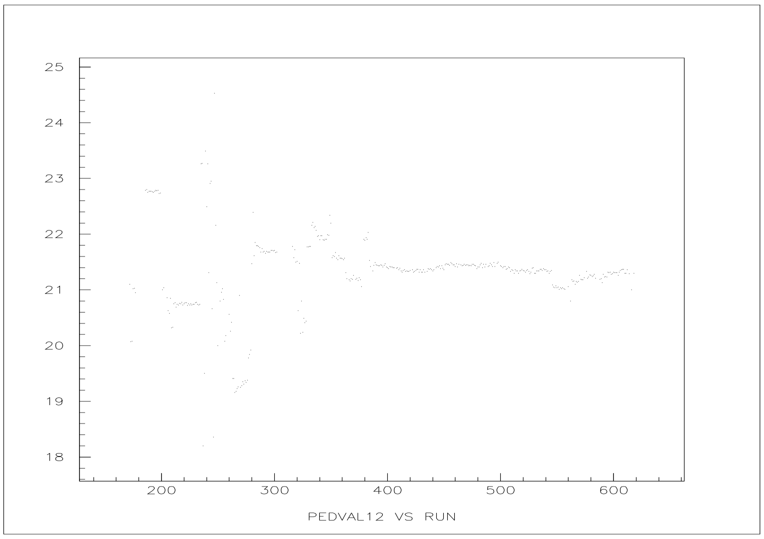

Figure 7: Pedestal value vs run for trigger channel (1,2)

1. Almost all channels: Value fluctuating until run 382 ± 4, then stable
2. Many channels: Value jumps from run 590 to 591
3. Some channels: Values form peak/dip (typically 1–2 counts) in runs 316–328

The first feature was rather elaborately described. To quote all of the description would be lengthy; but the numbers corresponding to Fig. 7 should not be left out:

Typical quantitative example (mean $\pm$ r.m.s.):

$$\begin{aligned} \text{val(1,2)} = \quad & 21.3 \pm 1.0, \quad \text{run} < 384 \text{ (182 runs, fluctuating)} \\ & 21.3 \pm 0.12, \quad \text{run} \geq 384 \text{ (216 runs, stable)} \end{aligned}$$

The run numbers defining all three features may coincide with changes in external conditions noted in the paper runlog [8, 9]:

1. (a) Run 382–: damaged cable at fan-in output for cell (5,5) replaced
   (b) Run 388–: Faulty fan-in replaced for cells (4,3) and (4,5)
2. Run 591–: 50 Ohm terminators of P-channels moved from calo/fan-in to ADC end
3. Run 3xx-3yy: RD3 high voltage scan (run range uncertain)

The first change does not fully explain the first feature, since

- The values stabilize shortly *before* the replacements, in most channels
- The replacements only directly involved 4 of the 36 channels

The second and perhaps the third features are well explained by the changes listed. The uncertainty on the third point is reflected by the following quote from [9]:

After run #329 RD3 start a high voltage scan but couldn't ascertain from the logbook when it finished. (this may conflict with earlier info. when the H.V. scan was said to be run #316–328 which could have been wrong).

3 Signal analysis

This section concerns the output of ADCs and ASICs in those time slices where the effect of an arriving electron is visible. All the conclusions build on run 607, which used 300 GeV electrons, 25ns time slices and all 3 depth samplings of 9 X_0 each.

3.1 Cross-talk

A rather large cross-talk has already been indicated by the three events in Fig. 6. The effect has been confirmed and further analyzed in three ways:

- by looking at these events in the other channels
- by summing all events
- by plotting one channel against another

This leads to the following description of how the cross-talk depends on the channel and on the signal in the beam cell.

3.1.1 Channel dependence

None of the 3×6 unconnected channels show any sign of cross-talk. The fact that they were terminated in 50 Ohm at the ADC end in the run considered leads to the conclusion that the cross-talk is generated before the signals arrive at the 3 ADC cards. Conversely, all the 3×6 connected channels show a detectable amount of cross-talk. It generally decreases with the distance to the beam cell. As an example, the magnitudes for the connected channels on ADC card #3 can be read off Fig. 6, with the results included in table 5. The maximum amount

Group	*Channel(s)*	*Peak size (counts)*	*Comment*
A3	(5,3)	(Overflow)	Beam cell, peak cut off at 255 counts (8 bit)
A3	(5,4)	13–17	Largest cross-talk here, closest to beam cell
A3	(5,5)	13–14	Next largest, perhaps electrically next closest
A3	(6,3)	4–10.5	Big difference between the 3 events
A3	(6,4)	4–6	
A3	(6,5)	3.5–4.5	
A2	(4,3)	1.5–6	Neighbour to the right hand side of beam cell
A2	other	0–3	Other five connected channels on ADC card #2
A1	all 6	0–2	
P#	all 18	0	Unconnected channels are silent

Table 5: Cross-talk peak heights for 3 events of run 607

of cross-talk here is seen to be 17 counts — corresponding to $\approx$ 17 GeV — for these 300 GeV electrons.

The few-count cross-talk on ADC cards #2 and #1 is almost invisible from single events, because it is on the limit of drowning in the background of high frequency noise, which covers 4-count intervals on those cards (see Sec. 2.2.3). However, when 100 events are summed, the cross-talk exposes itself as a defined, regular peak in every connected channel.

3.1.2 Beam signal dependence

The cross-talk interpretation of the non-beam-cell peaks is strengthened considerably by the study of how the ADC value in channel (5,5) (not even neighbour to beam cell) rises with the concurrent value in the beam cell, (5,3). A plot of one against the other, for all time slices of 100 minimum bias events, is shown in Fig. 8. First of all, it confirms that all events are affected by the cross-talk: while 27–29 is the usual pedestal value for channel (5,5), *none* of the event time slices with adc53$\geq$240 have adc55$<$30; most of those have adc55 in the range 40–44, corresponding to a cross-talk of approximately 42–28=14 counts or 14 GeV.

Next, a linear rise is clearly seen, with a slope around 7%. The ADC values in channel (5,5) mostly cover a 5–7 count interval for any given value in the beam cell. The only exception is the points (hollow crosses) coming from the events that are strongly affected by the Low Frequency Noise found in Sec. 2.2.2. They fall on lines with the same 7% slope, but shifted roughly 10 counts up or down in adc55. Note that the noise shifts *both* ADC values by 10 counts, since they come from the same ADC card; but the difference in scale tends to hide the shift in adc53.

3.1.3 Run dependence

The logbook [9] reports a first observation of cross-talk after run 331. The size is 4–10 ADC counts. It is seen in the connected trigger channels (4,2–6), as well as in the unconnected channel (4,1).

Another analysis [10] also sees cross-talk in unconnected channels. Both observations concern runs earlier than run 591, when the "unconnected" channels were electrically connected to a long cable chain: Lemos + 40m BNC cables + Lemos + NIM fan-ins + x, where x is partly terminators and partly terminated Lemos (Sec. 2.3 and Fig. 1).

Combining all information, the conclusion is that the cross-talk originated after the RD3 calorimeter+shapers, but before the RD27 FADC modules.

Whether or not the cross-talk was generated in the 16+16ns Lemo cables leading to the NIM fan-ins can be determined by comparing differently terminated channels in a run in the range 577–590:

- Cross-talk in *all* 18 P-channels, i.e. also the fan-in-terminated ones, will exclude the 16+16ns Lemos as sources of cross-talk
- Cross-talk absent from some P-channels (confirmed from logbook, if possible, that these are precisely the fan-in-terminated ones) will imply that the 16+16ns Lemos were the source

3.2 Hit flags and energy sums

Only a superficial study of the hit flag and energy sum outputs of the ASICs has been done in the present work. Still, a few general points can be made, starting from plots of each of them

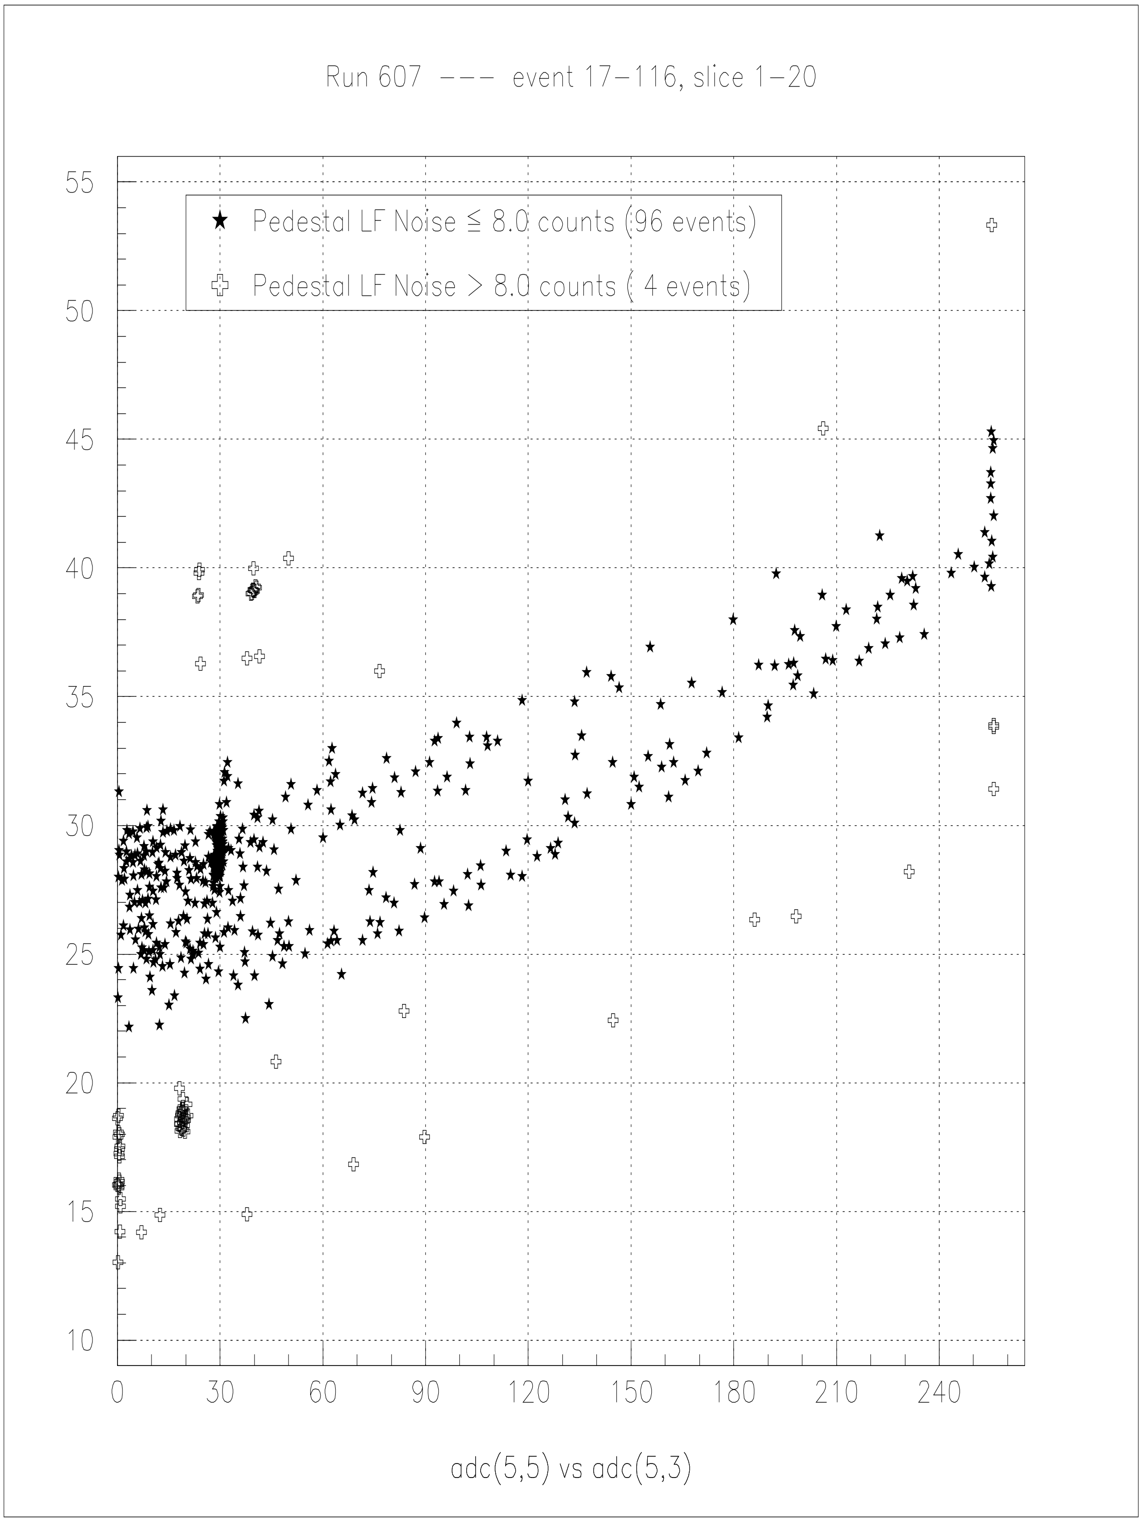

Figure 8: Correlation of ADC values in non-beam and beam channel (*cross-talk*)

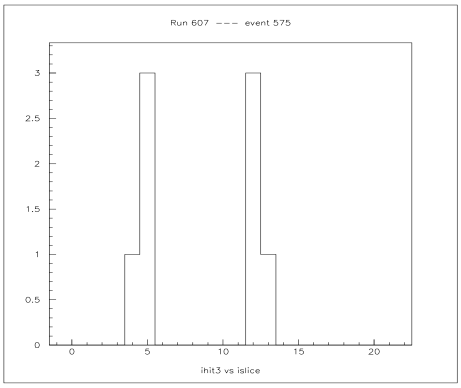

Figure 9: Hit values ($2\times$hitB+hitA) in ASIC #3 vs time, for a double event

versus the time slices in well selected events. Fig. 9 shows the 2-bit hit value for ASIC #3 in the double event of Fig. 5. As expected, no non-zero hits on any of the other 8 ASICs were found in this run.

Each slice number on the ADC plot is one less than the corresponding number on the hit plot, the reason being that the energy sums were output 6 clock cycles after the ADC input (except when *late* as explained below), and the hit flags 1 cycle after that. Both were then always shifted back 6 clock cycles relative to the ADC values, before all this was written to tape.

The cross-talk conspires with the Isolation Thresholds to give the "1-3" or "3-1" hit peak shapes characteristic of many events of run 607. The hit value "1" means that only hit flag "A" was set, while "3" means both "A" and "B". In both cases, the Cluster Sum (effectively here the sum of adc53 and the reference cell value adc42$\approx$28) had to be greater than the Cluster Threshold 200. This was the case for slices 3, 4, 11 and 12 on the ADC plot. Hit B only fired when the Isolation Sum (of 12 ADC values) was less than the Isolation Threshold 360, i.e. 30 per Isolation ADC on average, a few counts above the typical pedestal value. Hit B didn't fire in the first hit slice of the first peak and the last hit slice of the second peak on Fig. 9, because in those cases the beam cell signal was maximal (see Fig. 5), and so also the cross-talk in the Isolation cells was large enough to inhibit the hit. Conversely, the cross-talk became too low for inhibition in the other two hit slices, where the beam signal was below maximum.

Sometimes the hits appear 1 clock cycle too late. However, a check of a few of these events revealed that each one also contained late energy sums. Thus no contradiction with the Status Report [1] was found, since it only asserts that the hits are correct for those events where the energy sums are "in-time".

Fig. 10 shows the energy sum output of the 9 ASICs for one of the many events where part of the output is 1 clock cycle too late. Here, ASICs #3+6+9 are *late*, while the other 6 are in-time. However, ASIC #2 is immediately seen to have two time slices where the energy sum output is different from that expected in any of the nearby slices. Thus, this event also belongs to the class of rare *bad* events, a term used by Alan Watson (AW), who is doing a detailed study of them.

The first 1000 events of run 607 contain 14 *bad* events, according to AW's search algorithm, i.e. a rate of 1–2%. In each of these 14 events, only one of the 9 ASICs is affected. (The 14 events affect a total of 5 ASICs). 12 of them have exactly 2 *bad* slices in the affected ASIC,