CMS-IN 2001/???
March 15, 2001

CMS Internal Note

On Calibration, Zero Suppression Algorithms and Data Format for the Silicon Tracker FEDs

I.R. Tomalin
R.A.L.

Abstract

This note proposes clustering algorithms which the Tracker FED's could use to reduce the data volume which transmitted to the DAQ. It also considers how common-mode noise should be subtracted. Using Monte Carlo simulation, it examines the characteristics of the clusters, and compares the performance of the clustering and common-mode noise subtraction algorithms, both in terms of cluster finding efficiency and their effect on the data volume (tracker strip occupancy). The question of how the output data should be formatted to minimise its size is also addressed.

1 Introduction

The main aim of this document is to help define how the FED will perform zero suppression on the tracker data. It also considers how the FED should obtain the necessary calibration constants. Furthermore it addresses the issues of tracker occupancy and how the FED output data should be formatted.

Section 2 describes the calibration procedure which must be performed before the FEDs can do cluster finding, and considers how the calibration constants should be monitored to ensure data qualtity. In Sect. 3, the simulation tools which are used for the studies presented in this note are described. Section 4 shows some typical characteristics of clusters in minimum bias events, to give some idea of what the FED is looking for. Section 5 compares the performance of possible clustering algorithms, beginning with a comparison of algorithms used by existing experiments and continuing with studies of cluster finding efficiency and ghost rates. Section 6 studies the related problem of how to estimate the common-mode offset. Finally, Sect. 7 shows the expected occupancy using the best algorithm found, and also examines how the output data should be formatted to minimise its size. Section 8 concludes by proposing the best algorithms on the basis of these studies.

2 Calibration and Monitoring

2.1 Calibration

A calibration run could be taken at the beginning of each fill to update the constants needed by the FED to do cluster finding. This might require 1000 events to be written without zero suppression to a computer which would calculate the following quantities:

- The pedestal of each strip. (i.e., The mean pulse height excluding signals.)
- The rms noise on each strip. (i.e., The average variation of the pulse heights about the pedestal, excluding signals.)
- The rms variation from event to event of the common-mode offset per APV, (where the common-mode offset is defined as the mean pulse height on all strips in an event, excluding those with signals.)
- Flagging of bad strips, defined as those which are either dead or very noisy.

Using 1000 events ensures that the statistical uncertainties on the pedestals and noise would be smaller by a factor of ≈ 30 than the noise itself, and hence small enough not to degrade cluster finding performance.

It is proposed to read out the events via the VME back-plane of each FED and process them in the FED crate controller. An alternative solution would be to read out the events through the main data channel, and process them on the DAQ computer farm. Which of these two solutions is best may depend on the data rates. For a 96 channel FED, which digitises the pulse height information from each strip with a 10 bit ADC, the data volume

per event is $96 \times 2 \times 128 \times 10$ bits = 31 Kbytes. As there are likely to be 20 FEDs per crate, this implies a data size of 610 Kbytes/event/crate.

A typical VME crate controller, such as the CES RIO3, has a 1 Gbyte DRAM memory [1], which would allow it to buffer 1600 events in memory. This is more than adequate, since it is likely that each event could be analysed as it arrives, and that only a small buffer is needed.

The time required to transfer the data over the VME link to the crate controller is also an issue. Assuming a data transfer rate of 80 Mbytes/s [1], the time required to transfer 1000 events from the crate would be $1000 \times 0.61/80 = 8$ seconds. One should also consider the time required for the processor to determine the calibration constants for each of the $20 \times 96 \times 2 \times 128 = 500$k strips in the crate. Sophisticated, offline analyses running on present day computers can do this in less than one minute. This could easily be done before the start of a fill.

The calibration constants must be available off-line, so need to be merged with the data stream as start-of-run/slow-control records or stored in a calibration database.

2.2 Monitoring

The calibration constants could be occasionally (e.g., once per hour) checked during a fill, to see if they have altered or if any additional bad wires have been identified. This would help ensure data quality. One might consider updating the constants if they do change. However, this is probably unnecessary and furthermore, in the current design of the final FED, could only be done if the trigger is first stopped.

The FED crate controller could also monitor several other quantities to check data quality and APV/FED behaviour. These tasks will be addressed in the 'FED User Requirements' document (in preparation).

3 Simulation Tools

Most of the studies presented in this note were performed with the CMS Monte Carlo. This has two parts: CMSIM v120 which simulates the energy deposits of particles in the silicon wafers, and ORCA v4.40 (with some updates) which converts these deposits to a pulse height on each strip, after addition of noise and common-mode noise. ORCA also digitises the signals and simulates the various FED zero suppression algorithms. The tracker geometry used corresponds to the latest all-silicon tracker proposals [2].

Estimates of cluster finding efficiency are extremely sensitive to the assumed signal to noise (S/N) ratio. It is expected that, after irradiation, the final tracker will have S/N ratios of 13 and 15 in the inner and outer tracker respectively. However, to be prudent, it is desirable to use smaller S/N values of 10 and 12 respectively for studies of FED performance [3]. Unfortunately, in ORCA, whilst the S/N ratio usually takes a fairly reasonable value of 10.4, it drops to 6 in the inner three end-cap rings and rises to 15 in much of the inner four barrel layers ! The cluster finding efficiency is also dependent on Landau fluctuations being correctly simulated. This is achieved thanks to an update to ORCA v4.40.

In view of these defects, most of the ORCA studies presented in this note do not use the end-cap. Even so, one should treat the results with caution. In particular, those on cluster finding efficiency in the inner barrel are doubtless optimistic.

Motivated by the problems with ORCA, a second simulation tool was developed for studying cluster finding efficiencies. This is the Toy Monte Carlo described in Sect. 3.1, which provides an S/N ratio of 10.4.

Studies of occupancy in the tracker are sensitive to the assumed capacitive couplings between strips, since the larger these are, the wider the reconstructed clusters become. In the ORCA simulations used here, a charge deposited on one strip induces 6% of its value on each of the neighbouring strips. (N.B. This differs from the default capacitive couplings in ORCA, which unrealistically also assume large capacitive couplings to the next-to-neighbouring strips.)

Tracker occupancy is also sensitive to the APV time resolution, since a good resolution suppresses clusters from out-of-time bunch crossings. In ORCA the time resolution is assumed to be 12.1 ns, when the APV is running in deconvolution mode. This corresponds to recent measurements of APV25 performance.

3.1 The Toy Monte Carlo

The Toy Monte Carlo program [4] simulates a detector with 128 strips as follows:

- Noise is simulated on each strip according to a Gaussian distribution of width 3.5 ADC counts.
- Common-mode noise is neglected.
- Signal can optionally be added, either for tracks passing perpendicularly through the detector, or for inclined tracks.
  - For perpendicular tracks, which are assumed to deposit their energy on only one strip, the energy deposited is given by $\Delta_E = C[\lambda + K]$ [5], where λ is a random number from a Landau distribution. The parameters C and K are constants, whose values are chosen so as to obtain a signal to noise ratio of 10.4 and a signal distribution of width similar to that seen in test-beams [6].
  - For inclined tracks, making an angle θ to the normal to the detector, the energy deposited on each strip across which the track passes is given by $\Delta_E = C(d/w)[\lambda + K + \log_e(d/w)]$ [5], where C and K are the same constants used for perpendicular tracks, d is the length of the track's path through the silicon above a given strip and w is the thickness of the silicon. The strip pitch p was set to $0.244w$, corresponding to the outer-most barrel layer. This layer has the widest (and hence most difficult to reconstruct) clusters, for a track of given P_t. Fig. 1 illustrates the dependence of the cluster pulse height distribution on track crossing-angle.
- A small fraction (6% in accordance with test-beam data) of the charge on each strip is moved to each of its immediate neighbours, to take into account capacitive coupling.

- Finally, a clustering algorithm is applied to the simulated strips.

4 Typical Cluster Characteristics

This section presents some basic characteristics of reconstructed clusters in minimum bias events, as obtained from the ORCA simulation. The aim of this is just to give a 'feeling' for what the FED is dealing with. High luminosity conditions are assumed, which implies a mean of 24 superimposed events per bunch crossing. The clusters were reconstructed using the FED clustering algorithm recommended in Section 8.

Figure 2 shows the transverse momentum P_t with respect to the beam-axis of the track which produced each genuine, reconstructed cluster. Most clusters are produced by tracks with $0.05 < P_t < 1$ GeV/c. Tracks with $P_t < 0.65$ GeV/c can spiral inside the tracker and so produce many entries in this figure. A small fraction of clusters are produced by very low momentum ($P_t \approx 0.5$ MeV/c) tracks, which are presumably delta-ray electrons produced in the silicon.

Figure 3 shows the time at which genuine, reconstructed clusters are produced relative to the nominal readout time. Many clusters are produced near the nominal readout time. However, a significant fraction are produced several nanoseconds late, by low momentum, spiralling tracks. When the APV is operating in deconvolution mode, its good time resolution suppresses clusters from the wrong bunch-crossing: only 23% / 2.9% of genuine, reconstructed clusters are produced in the previous / following bunch-crossings respectively. (The asymmetry in these two numbers arises from clusters produced by spiralling tracks.)

Figure 4 shows the width in strips of reconstructed clusters. It shows separately *prompt* clusters (i.e., those produced within 5 ns of the nominal beam-crossing time) in two separate transverse momentum ranges: $P_t > 1$ GeV/c (useful for many physics analyses) and $1 > P_t > 0.5$ GeV/c (useful for a few physics analysis [7]). It also shows the other clusters which are produced in the correct beam-crossing but are of no physics interest, since they either have $P_t < 0.5$ GeV/c or are not prompt (i.e., they have spiralled around the tracker). Finally, it shows the widths of clusters reconstructed in the wrong bunch-crossing and of fake clusters.

Prompt clusters produced by the tracks which are most interesting for physics analysis ($P_t > 1$ GeV/c) usually have a width of 1–3 strips. However, it is desirable for some physics analysis to reconstruct tracks with momenta as low as 0.5 GeV/c, which implies that the FED should be able to reconstruct clusters with a width of 1–5 strips. This statement can be checked by a simple calculation. – The angle θ to the normal of tracks crossing silicon wafers in the barrel is given by $\theta = \arcsin(3RB/20P_t)$, where $B = 4$ T is the magnetic field, R is the distance in meters of the wafer from the beam-axis and P_t is measured in GeV/c. Fig. 6a shows a plot of θ versus P_t for tracks crossing the outer layer of the silicon barrel at $R = 1.08$ m. Knowing that detectors in this layer have a thickness of 500 μm and a pitch of 122 μm, one can calculate the number of strips across which each track passes as a function of its crossing-angle. This is shown in Fig. 6b. From these two figures, it is apparent that a track with $P_t = 1$ GeV/c crosses this layer at an angle of $\theta = 40^\circ$ and can induce signals on up to 4 strips, (even neglecting capacitive coupling to additional strips).

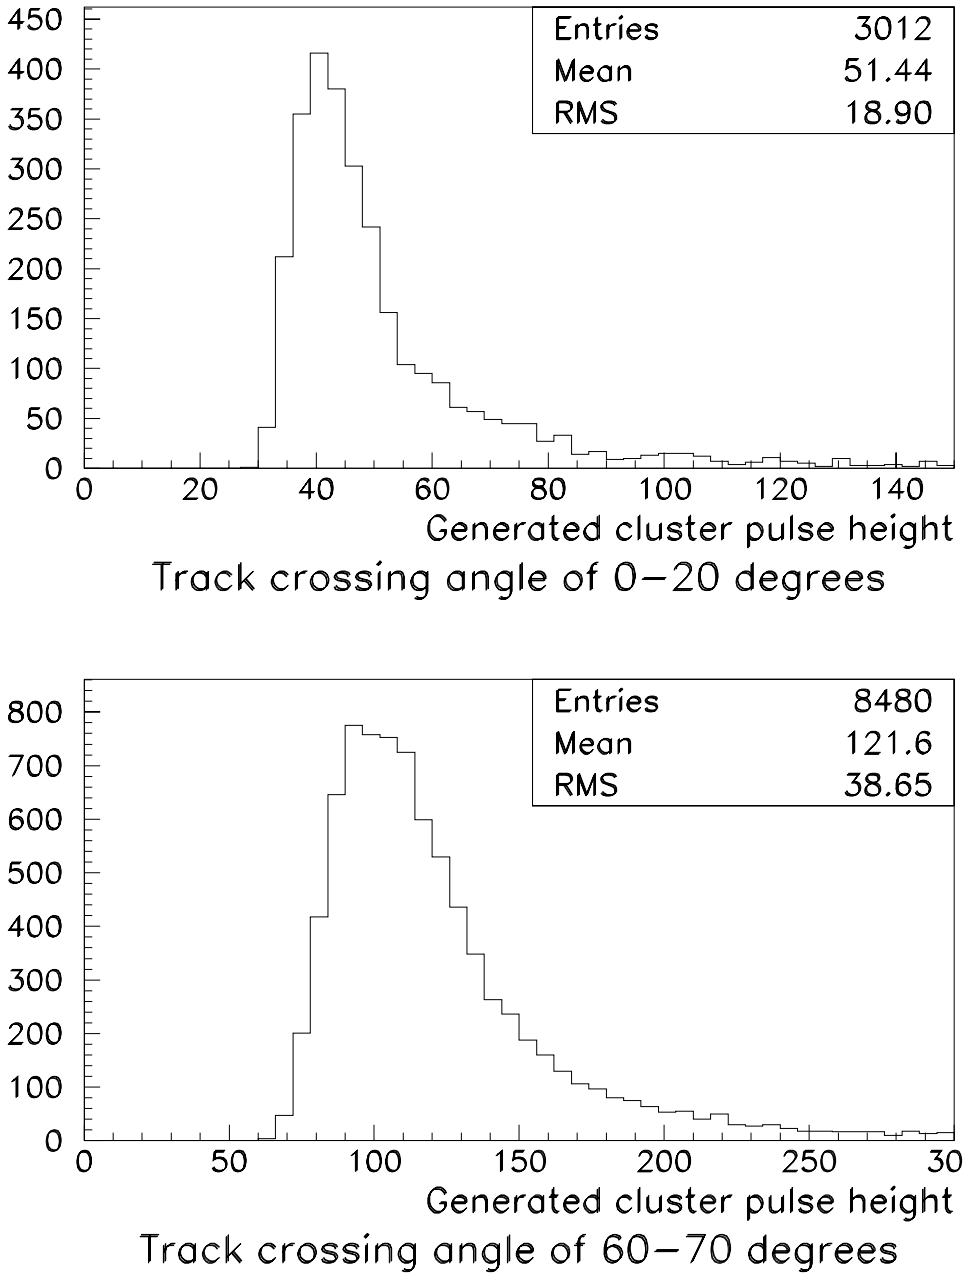

Figure 1: Comparison of simulated total cluster pulse height for two different track crossing-angles.

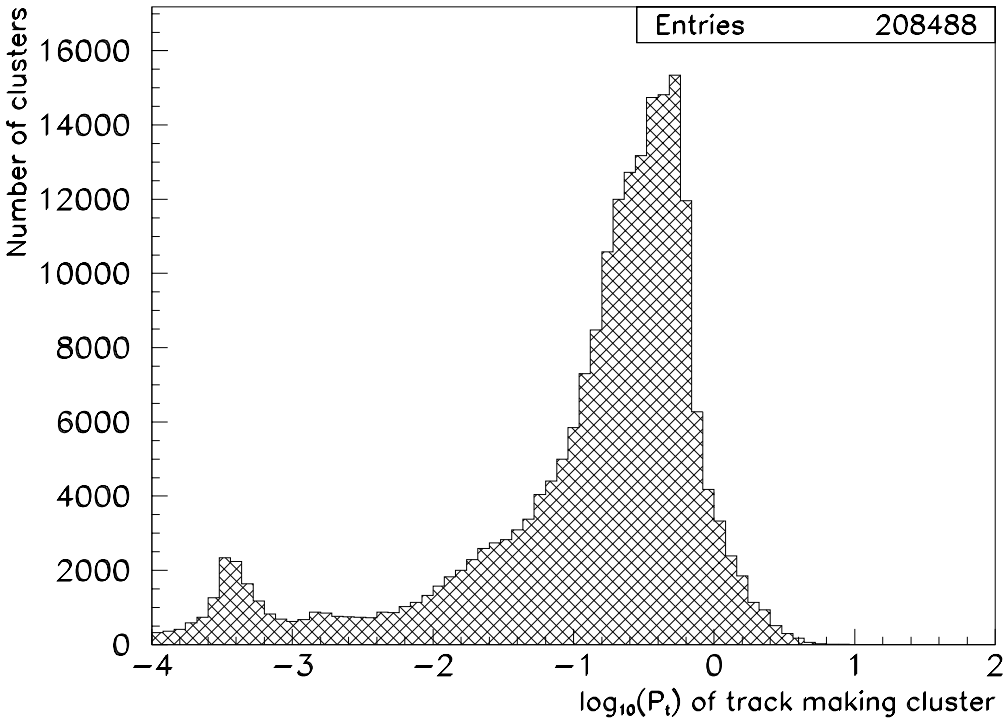

Figure 2: For every genuine, reconstructed cluster in the tracker, this shows the transverse momentum with respect to the beam-axis of the track which produced it.

5 Cluster Finding

The aim of this section is to propose a cluster finding algorithm which could be used by the FED. It will be assumed that the noise and pedestal of each strip have previously been determined in a calibration run (discussed in Sect. 2) and that the FED has subtracted the pedestals. It will also be assumed that for each event, the FED has perfectly determined the common-mode offset of each APV (discussed in Sect. 6) and subtracted that too.

Section 5.1 examines cluster finding algorithms used by other experiments and goes on to propose algorithms which might be used in the FED. In Sect. 5.2.1, the efficiency of suggested algorithms is studied using both ORCA and the Toy Monte Carlo. Section 5.2.2 examines the ghost rates.

5.1 Consideration of Possible Clustering Algorithms

5.1.1 Algorithms used Online by other Experiments

- **CDF** and **D0** use a simple approach: each strip over threshold is output, together with its immediate neighbour on either side, which improves positional resolution through charge sharing [8, 9].

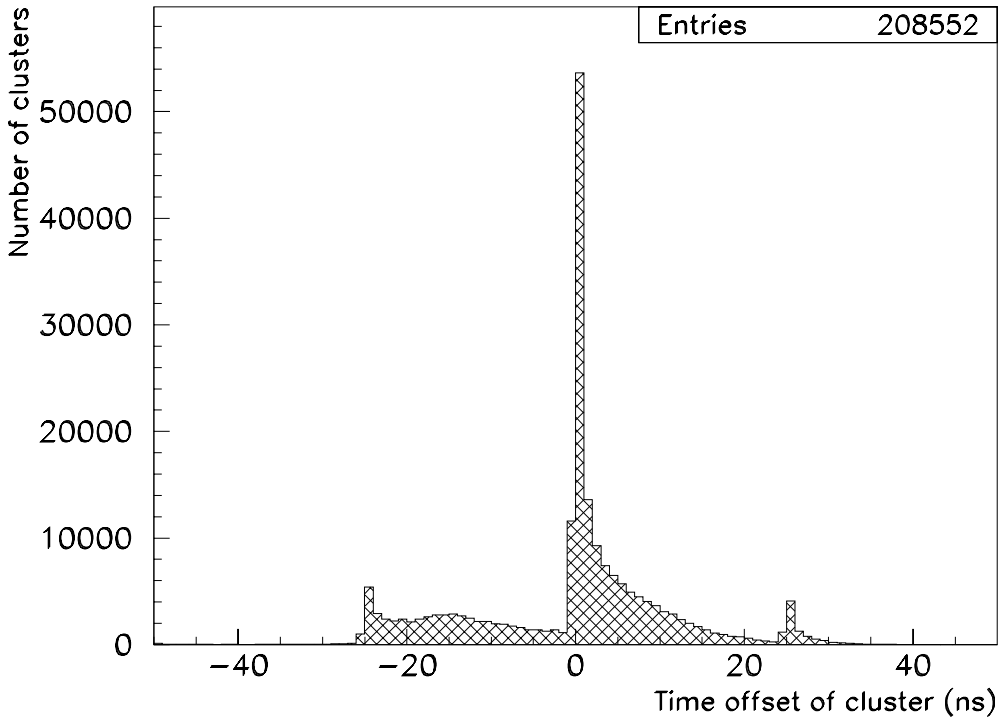

Figure 3: For every genuine, reconstructed cluster this shows the time it was produced relative to the nominal readout time.

- **ALEPH** applies the same algorithm as D0 (but the low data rates allow it to output seven nearest neighbours on either side of hit strips) [10].
- **DELPHI** accepts clusters if they contain n strips and have a total signal to noise ratio (summing over all the strips in the cluster, i.e. $\sum_i S_i/N_i$) exceeding $4\sqrt{n}$. The number of strips n is required to be in the range 1–3. Like D0 and ALEPH, they also output some neighbours of the cluster [11].
- **OPAL** add the signals on two neighbouring strips, and accepts both strips if the sum exceeds three times the noise [12].

The CDF/D0 algorithm seems feasible for the CMS tracker, but outputting many neighbouring strips as ALEPH does would probably lead to an unacceptably high data volume. The DELPHI algorithm involves floating point division, which renders it less attractive, considering that it must be performed quickly in an FPGA. The OPAL algorithm is possible, but as it only considers pairs of strips, it may produce wider clusters than necessary.

5.1.2 Proposed Algorithms for use in the FED

Three clustering algorithms are proposed here and studied below, together with the off-line TDR algorithm [13] which acts as a reference:

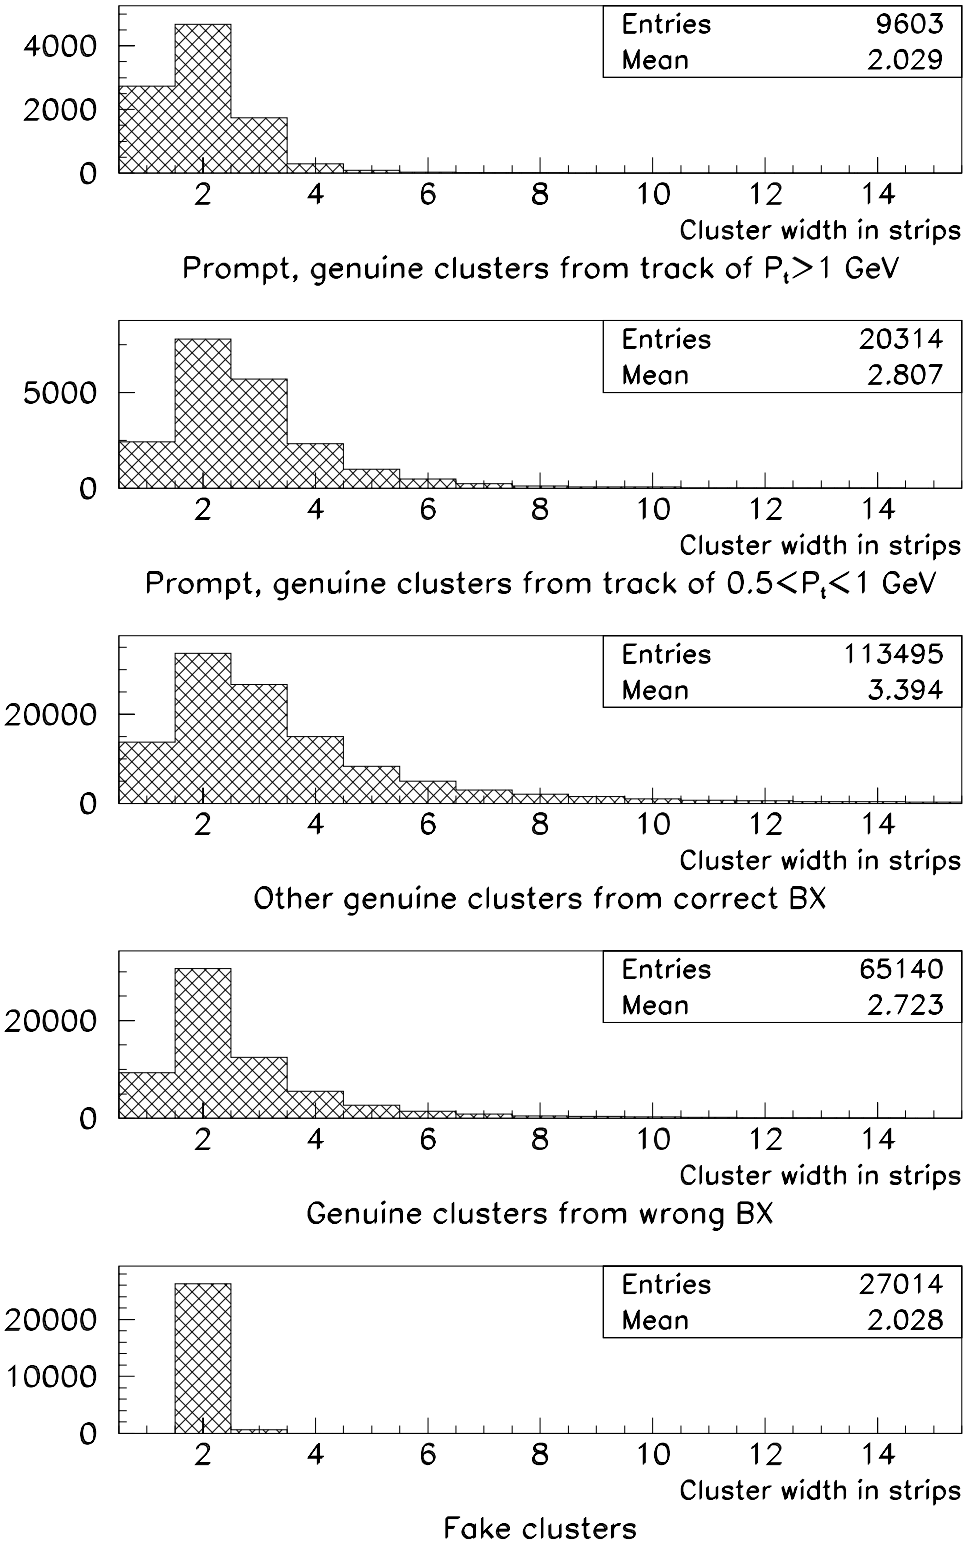

Figure 5: Width in strips of reconstructed clusters, divided up into various categories.

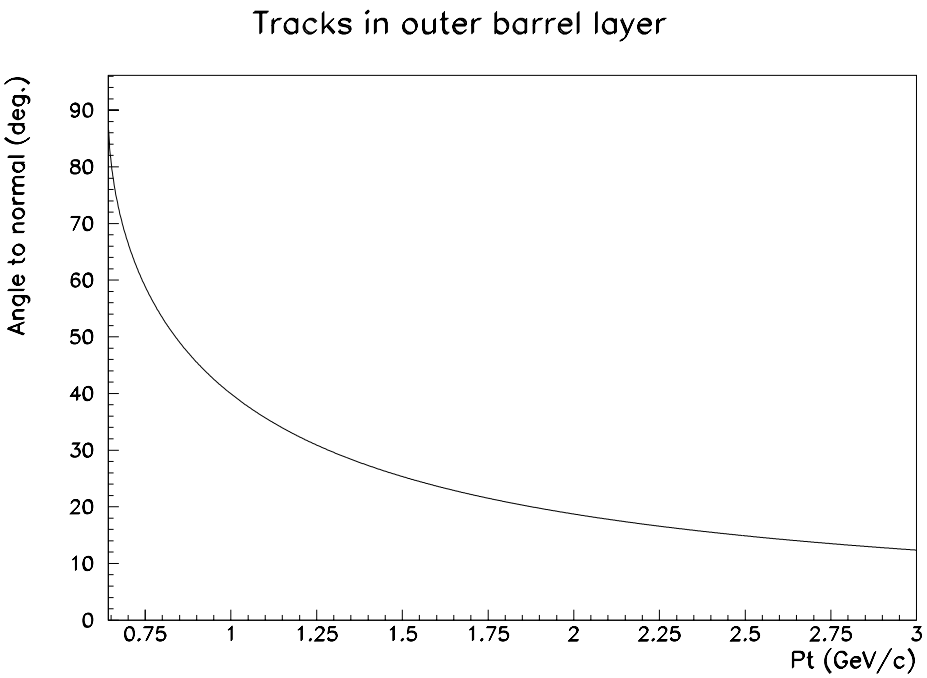

Figure 6a: Angle to normal of tracks crossing outer barrel layer as a function of their transverse momentum to the beam-axis.

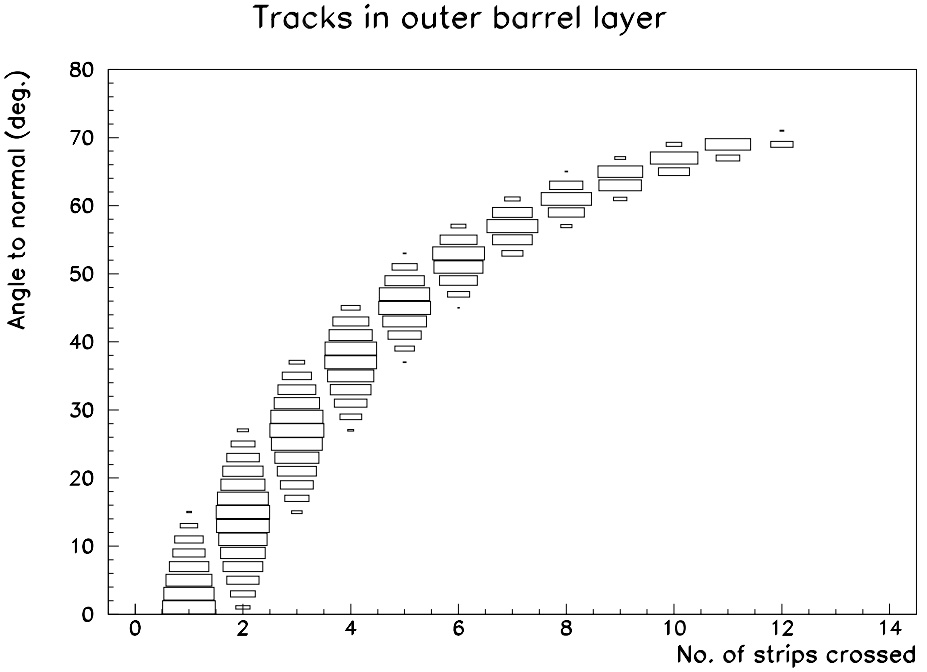

Figure 6b: Angle to normal of tracks crossing outer barrel layer as a function of the number of strips they passed over in the detector.

1. **Algorithm FED 1**: A very simple one, which simply accepts all strips with $S/N > T$, where the threshold T was here set equal to 2.7.

2. **Algorithm FED 2**: Accepts all strips with $S/N > T_l$, but rejects single-strip clusters unless they pass a more severe requirement $S/N > T_h$. Here, $T_l = 2$ and $T_h = 5$. This is motivated by the fact fake clusters are usually only one strip wide.

3. **Algorithm FED 3**: Similar to that used by CDF/D0: accepts all strips with $S/N > T_h$ and their immediate neighbour on either side if that neighbour has $S/N > T_l$. Here, $T_l = 0$ and $T_h = 3$.

4. **Algorithm TDR**: The Tracker TDR offline algorithm: requires clusters to contain a seed strip with $S/N > 3$, all other strips in the cluster to have $S/N > 2$ and the cluster as a whole to satisfy $S_{tot}/N_{rms} > 5$, where S_{tot} is the total charge and N_{rms} is the rms noise of the strips.

5.2 Cluster Finding Performance

As cluster finding performance is particularly sensitive to inaccuracies in the simulation, it is studied using both ORCA and the Toy Monte Carlo.

5.2.1 Cluster Finding Efficiency

When studying the efficiency, two alternative definitions have been used for what constitutes a reconstructed hit:

- A loose definition, stating that a hit is reconstructed, if a cluster is found containing at least one of the strips which was crossed by the particle track.

- A strict definition, requiring that all of the strips fully crossed by the track be reconstructed in a single cluster.

According to both ORCA and the Toy Monte Carlo, all three FED algorithms yield a 'loose' cluster finding efficiency in excess of 99.5% for track crossing angles of up to 60°. Hence hits are virtually always at least partially reconstructed. However, unless all strips fully crossed by a track are reconstructed in a single cluster, the reconstructed cluster position will be highly inaccurate. The 'strict' definition of efficiency is therefore most relevant.

Figures. 7a and b show the estimated strict efficiency as a function of crossing angle using ORCA and the Toy Monte Carlo, respectively. Both plots show results from all four algorithms. The ORCA results were obtained for the outer barrel layer, where the ratio of strip pitch to wafer thickness is the same as that assumed in the Toy Monte Carlo. Within rather large statistical errors, the results from the two simulations seem comparable, which is reassuring. From the Toy Monte Carlo results, it is apparent that the FED 1 algorithm looses efficiency for fairly modest angles, whereas the other algorithms do much better. The near equality of the TDR and FED 2 algorithms' performance is no coincidence: for clusters consisting either of a single strip or of three or more strips, the two algorithms are

mathematically identical. The FED 3 algorithm performs best. At large crossing angles, many of the missed clusters are actually split in two by the clustering algorithms, because a strip with low pulse height in the middle of the cluster produces a 'hole'. Algorithm FED 3 is less sensitive to this effect, as its approach of outputting the neighbours of hit strips means that it can fill in most holes.

Figure. 8 shows three plots of the strict efficiency as a function of barrel layer obtained using ORCA. The plots show the efficiency to reconstruct clusters produced by tracks in three different transverse momentum (P_t) ranges. For $P_t \approx 1.2$ GeV/c, all algorithms are statistically compatible with almost perfect performance. For $P_t \approx 0.8$ GeV/c (and to a lesser extent for $P_t \approx 1.2$ GeV/c), a sdrop in efficiency is apparent in the outer two barrel layers. This is partially due to the fact that the crossing angle is very large in these layers for such low P_t tracks and partially due to the small ratio of strip pitch to wafer thickness, which results in clusters being spread over several strips. Algorithm FED 1 performs significantly worse than the others in these layers. However, none of the four algorithms performs well there, implying that it may be necessary to use looser cuts in these two layers. For the lowest momentum tracks $P_t \approx 0.5$ GeV/c, the efficiency is virtually zero in the outer two layers, simply because the tracks circle before reaching it. The efficiency is well below unity in several other layers of the outer barrel, particularly for algorithm FED 1. The layers with the lowest efficiency are those with the smallest ratio of strip pitch to wafer thickness. Note that the high efficiency in the innermost four layers should not be taken too seriously: it is strongly biased by the high S/N ratio of 15, assumed by ORCA in much of these layers.

5.2.2 Ghost Rates

Table 1 shows the mean ghost rate per strip obtained with each of the four algorithms. This is defines as the total number of strips associated to fake, reconstructed clusters divided by the total number of strips in the detector. Predicted ghost rates from ORCA and the Toy Monte Carlo are in reasonable agreement.

Table 1: The mean ghost rate per strip for each cluster finding algorithm. Results obtained from both ORCA and the Toy Monte Carlo are compared.

Algorithm	Ghost rate (%)	
	ORCA	Toy Monte Carlo
FED 1	0.29	0.36
FED 2	0.09	0.11
FED 3	0.23	0.28
TDR	0.014	0.013

Although the FED algorithms all yield far higher ghost rates than the TDR algorithm, this is of little consequence for the DAQ. All one cares about is that the ghost rate should be small compared with the rate from genuine hits, and all three FED algorithms achieve this. Their higher ghost rates are a direct consequence of the loose cuts, which they use to obtain comparable efficiency to the TDR algorithm with simpler algorithms.

There is some variation in the ghost rate between the three FED algorithms. However,

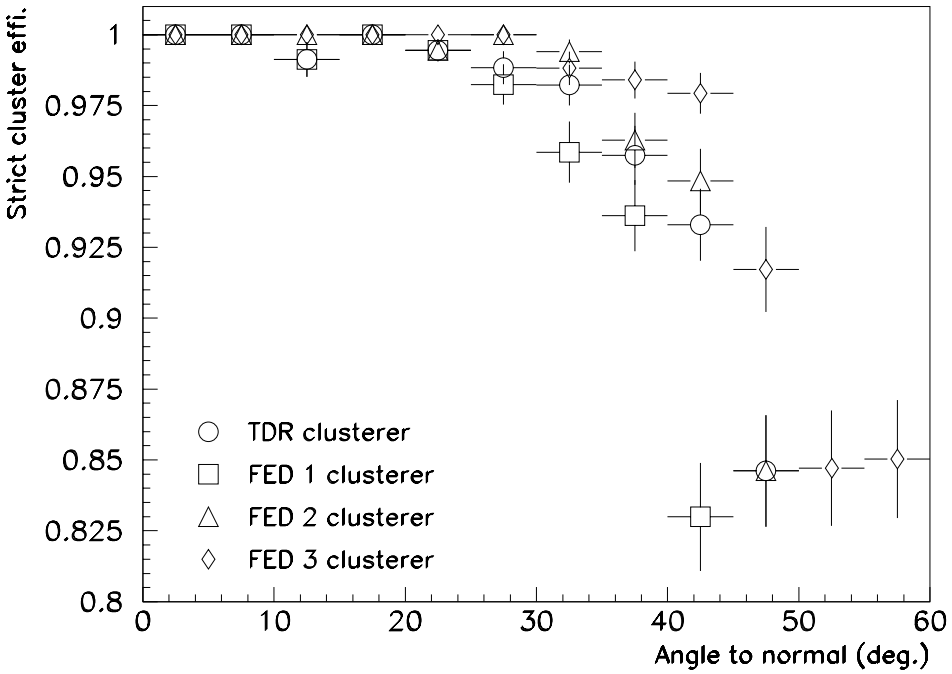

Figure 7a: ORCA prediction of strict cluster finding efficiency vs. crossing angle for all four algorithms in the outer barrel layer.

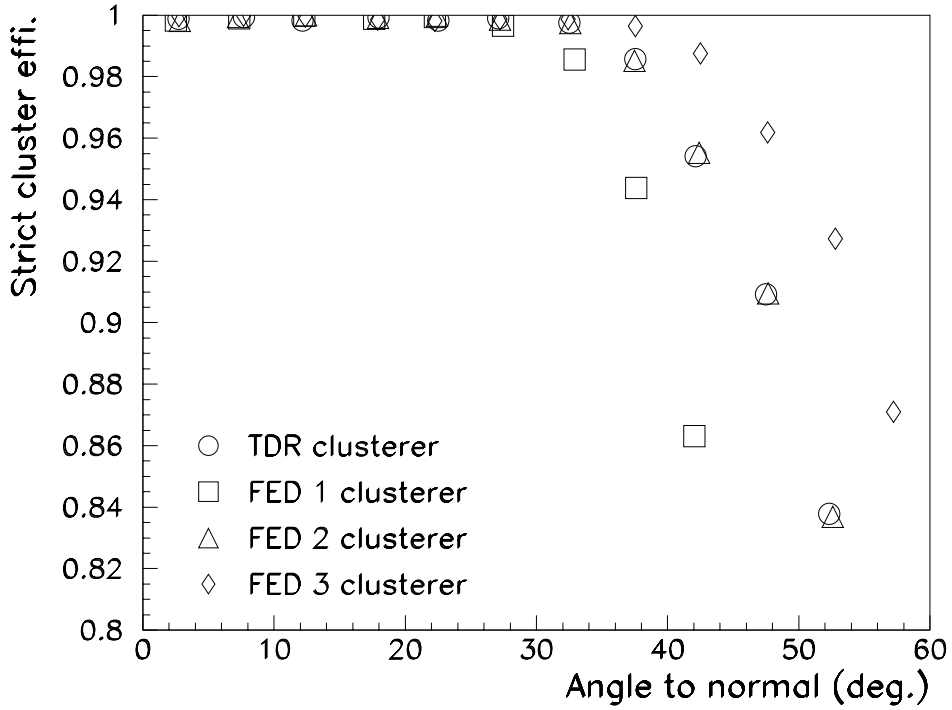

Figure 7b: Toy Monte Carlo prediction of strict cluster finding efficiency vs. crossing angle for all four algorithms in the outer barrel layer.

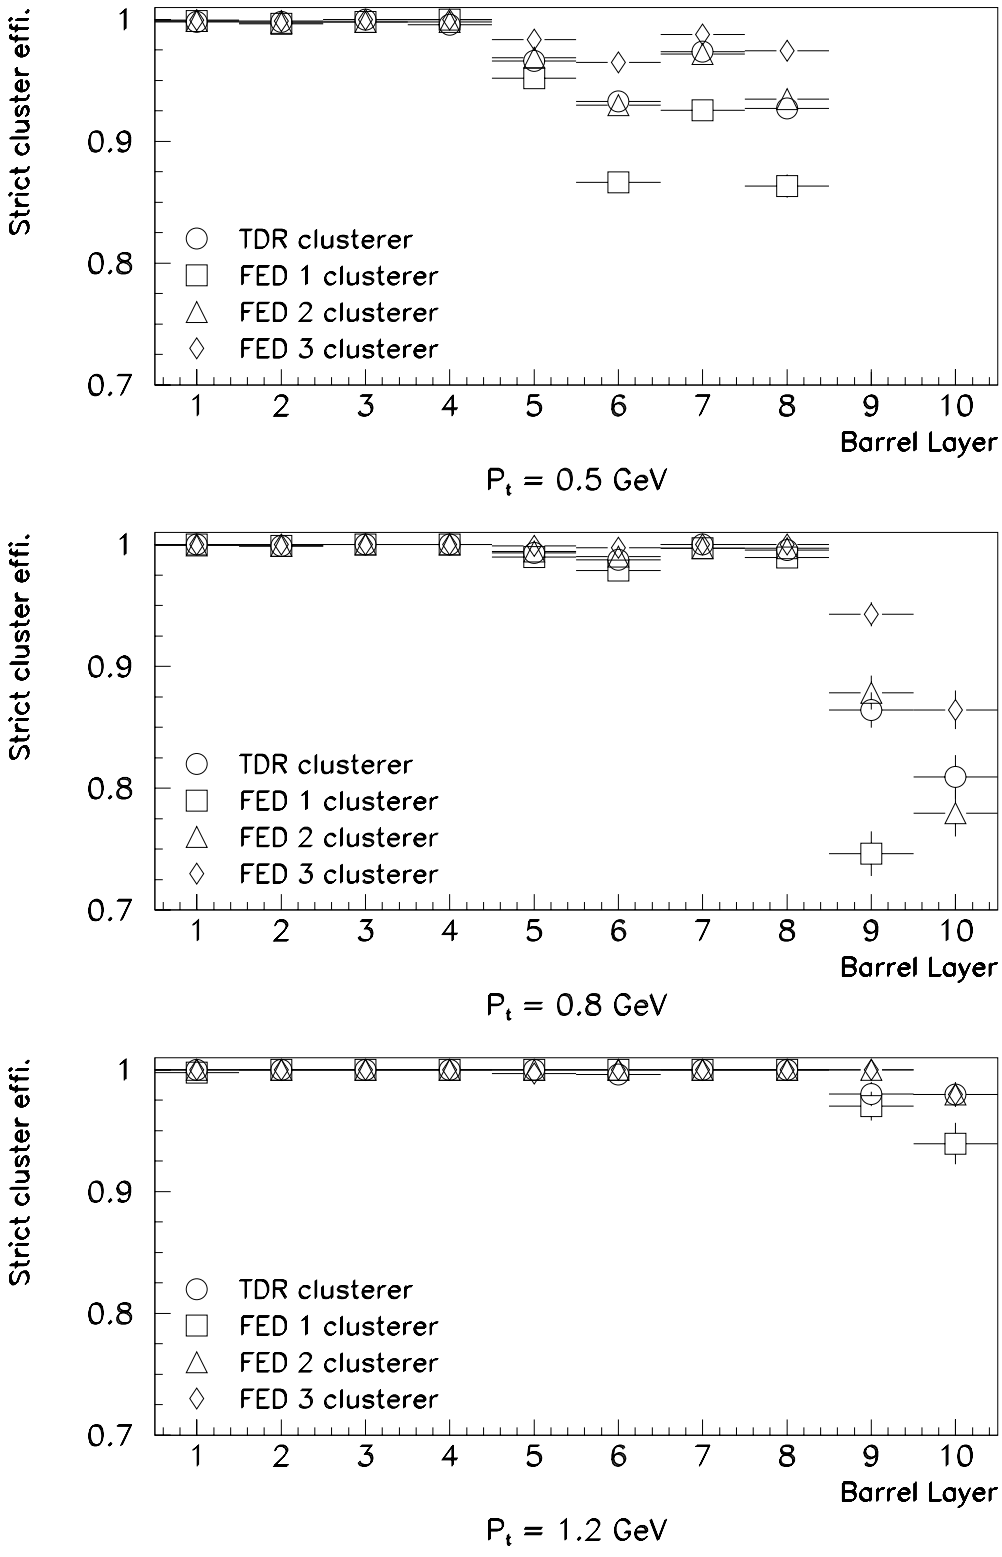

Figure 8: ORCA predictions for strict cluster finding efficiency vs. barrel layer in three P_t regions, for all four algorithms.

comparison of these values is not an useful way of judging their relative merits. A far more relevant quantity is the total strip occupancy they produce. This will be studied in Sect. 7.

Figure 9 shows the width in strips of fake clusters obtained from the four algorithms, estimated with the Toy Monte Carlo. The simple FED 1 algorithm is dominated by clusters only one strip wide, which are neatly suppressed by the extra cut in the FED 2 algorithm. This allows the latter to use looser cuts for wide clusters. The FED 3 algorithm can produce rather wide fake clusters, as a result of its tendency to output neighbouring strips.

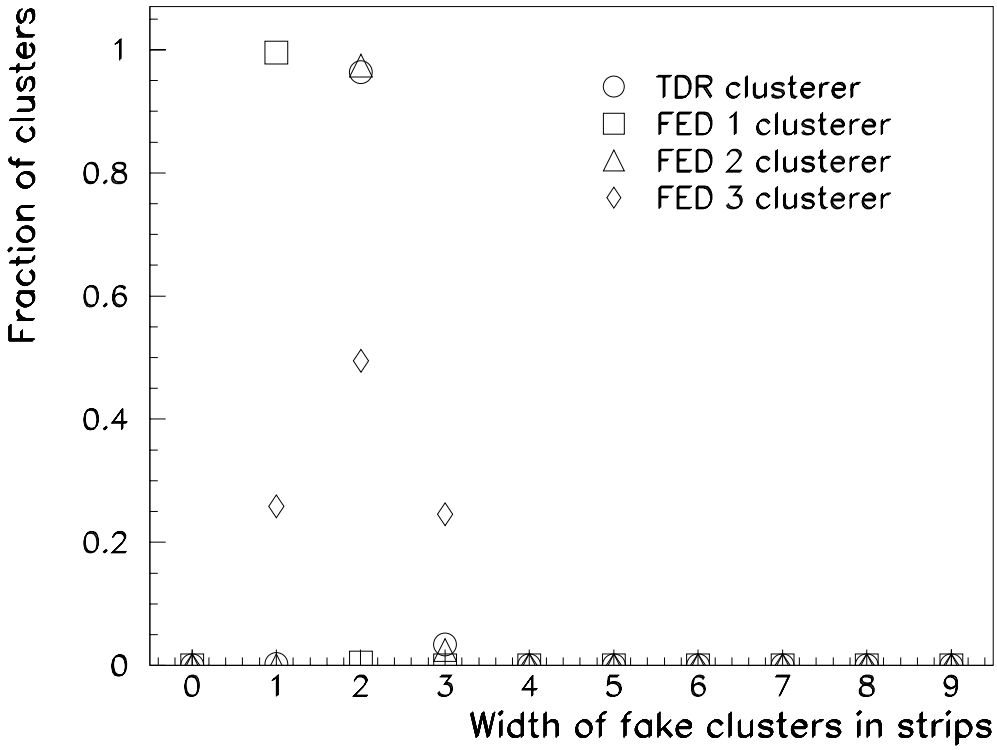

Figure 9: Width of fake clusters in strips for all four algorithms.

6 Estimating the Common-Mode Offset

The results on cluster finding performance presented in Sect. 5.2 assumed that the common-mode offset can be perfectly estimated and subtracted, It reality, it is not possible to achieve this using a simple algorithm in the FED. This section considers what might practically be achieved.

6.1 Algorithms used Online by other Experiments

D0 neglects common-mode noise [9], which would represent a gamble. ALEPH uses a rather sophisticated histogramming algorithm to evaluate it [10], which can't easily be implemented in the FED. CDF and OPAL assume that the common-mode offset is equal to the median

pulse height, which is less biased upwards by strips containing signals than the mean pulse height would be [8, 12]. This algorithm may work well in the FED.

6.2 Proposed Algorithms for use in the FED

Two basic algorithms have been considered:

- **The 'mean' algorithm**

  1. Estimate common-mode offset for each APV in the event, by setting it equal to the mean pulse height on its 128 strips.
  2. Flag strips which have $S/N > T_c$, since these are likely have a signal and this will bias upwards the common-mode estimate. For the study presented below, the threshold $T_c = 2$.
  3. Repeat step (1), ignoring flagged strips.

  Performing several iterations of this algorithm will yield more accurate results, but will be prohibitive in terms of the time used. In the study presented here, results are presented for only 1 iteration (i.e., step (1) performed once) and 2 iterations (i.e., step (1) performed twice).

- **The 'median' algorithm** This simply estimates the common-mode offset for each APV in the event, by setting it equal to the median pulse height on its 128 strips. This is less strongly biased by signals or noisy strips than the mean, so no attempt is made to reject these. An FPGA algorithm for calculating the median has been proposed [14], which is similar in complexity to the mean algorithm with two iterations.

Since the FED clustering algorithms presented in Sect. 5.1 use thresholds of a few times the rms noise, it seems likely that if the bias in the common-mode offset (defined as the estimated minus the true common-mode offset) exceeds the rms noise, the cluster finding efficiency will suffer. Figure 10a shows the probability of this happening as a function of barrel layer, when using the different algorithms for estimating the common-mode offset. In the innermost layers, where the occupancy is highest, several percent of the APVs have a bias in their common-mode offset exceeding the rms noise, when only one iteration of the mean algorithm is performed. It seems likely that this would result in a loss of cluster finding efficiency of at least a few percent. Performing two iterations solves this problem. The median algorithm performs just as well as the mean algorithm with two iterations. In the few cases where the median algorithm goes badly wrong, this is usually a result of delta ray electrons travelling in the plane of the silicon and depositing their energy on many tens of strips.

Figure 10b shows the same thing again but for simulated Pb-Pb collisions. As no heavy ion Monte Carlo was available, this plot was obtained by superimposing 250 minimum bias events. This yields roughly the same occupancies as expected in Pb-Pb collisions (with 4000 tracks per unit rapidity interval) [15]. This is a relatively optimistic scenario as some studies predict Pb-Pb occupancies which are twice this high. Figure10b shows that provided that

two iterations of the mean algorithm are performed or the median algorithm is used, it may be possible to estimate the common-mode offset in Pb-Pb events, especially in the outermost layers. However, based on the expected trigger rate during Pb-Pb runs [15], this does not appear to be essential: it should be possible to switch off the FED zero suppression and output pulse height information for all strips. Nonethless, zero suppression may be used in lighter ion collisions, as these yield occupancies at least a factor three lower than those in Pb-Pb collisions, but demand higher trigger rates.

Although the statistical precision of the ORCA studies is limited, one can obtain some confirmation of these ideas in Fig. 11. This shows, for minimum bias events, the dependence of the strict cluster finding efficiency on barrel layer, in three transverse momentum ranges. The efficiencies are displayed for the FED 2 algorithm using the different common-mode offset algorithms, and also when the common-mode is taken from its true value. The two more complex algorithms significantly improve the efficiency, especially for $P_t < 1$ GeV/c.

7 Occupancy and Data Compression

7.1 Occupancy

Figure 12 shows the dependence of the strip occupancy (defined as the fraction of strips associated with clusters) on barrel layer and on end-cap disk. The clusters were reconstructed using the FED 2 algorithm, with the common-mode offset being estimated with the median algorithm. The occupancy reaches 3.0% in the innermost barrel layer.

The strip occupancy depends strongly on the clustering algorithm. Table 2 shows its mean value averaged over the barrel for the algorithms considered this note. This table also reveals the relative weight of different contributions to the occupancy. The FED 2 algorithm achieves a slightly lower occupancy than the FED 1 algorithm. The FED 3 algorithm performs much worse. The reason for this is that as it outputs the neighbours of hit strips, the FED 3 algorithm produces much wider clusters. This statement is confirmed by Table 3, which shows the mean number of clusters per detector module. The FED 3 algorithm is not guilty of finding more clusters than any other. It is interesting to note that the occupancy is dominated by non-prompt clusters: either from spiralling low momentum tracks or from the wrong bunch-crossing. Such clusters look identical to prompt clusters, except that they have a poor signal to noise ratio due to the time-response of the APV. Unfortunately, this makes it virtually impossible to suppress them, whilst at the same time preserving a good efficiency for prompt clusters in a tracker with a low S/N ratio.

7.2 Data Compression

In the Tracker TDR [13] it is proposed that for each strip selected, in a given APV, by the FED, both the strip address and the ADC pulse height should be output to the DAQ, representing a total of 2 bytes of information. An alternative would be to output the address of the first strip in a cluster, the cluster size and the ADC pulse height of each strip in the cluster. This requires $2 + N_{strip}$ bytes, where N_{strip} is the number of strips in the cluster. This

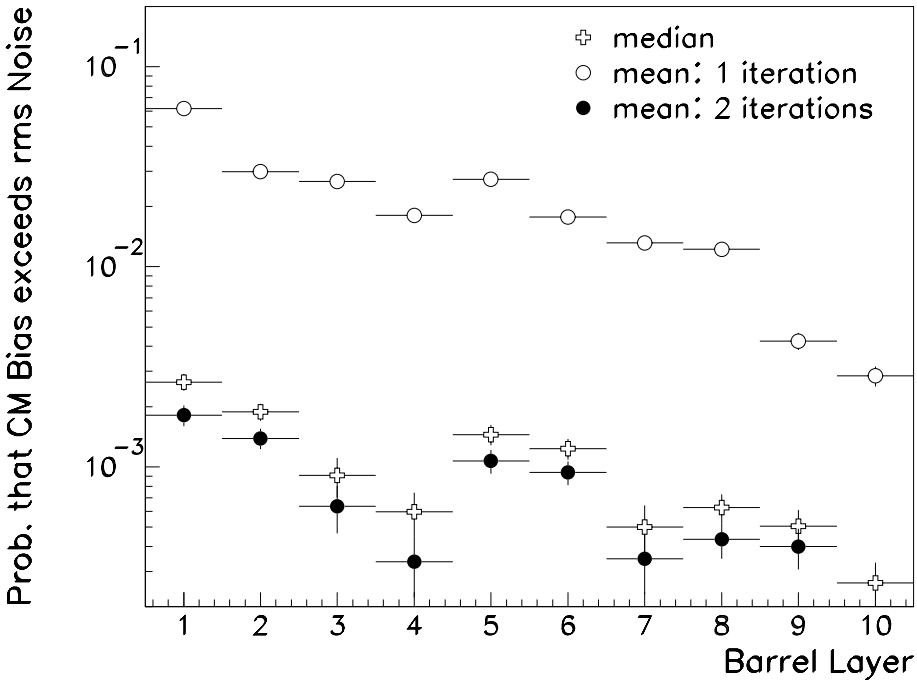

Figure 10a: Probability that the estimated common-mode offset exceeds its true value by more than the rms noise, as a function of barrel layer. The dependence on the common-mode offset estimation algorithm is shown.

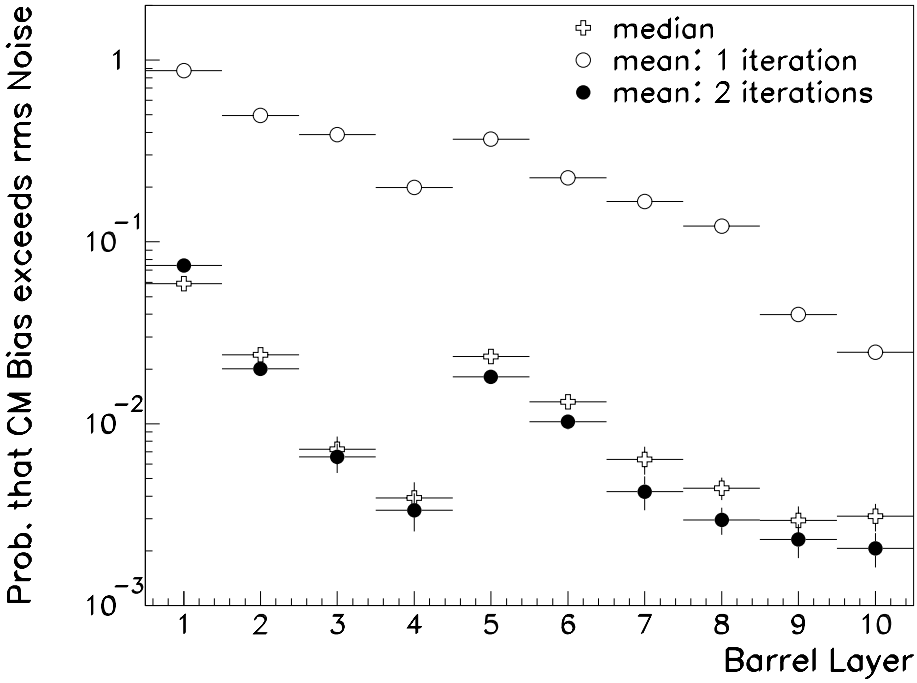

Figure 10b: Same again, but for an fairly optimistic simulation of Pb-Pb events.

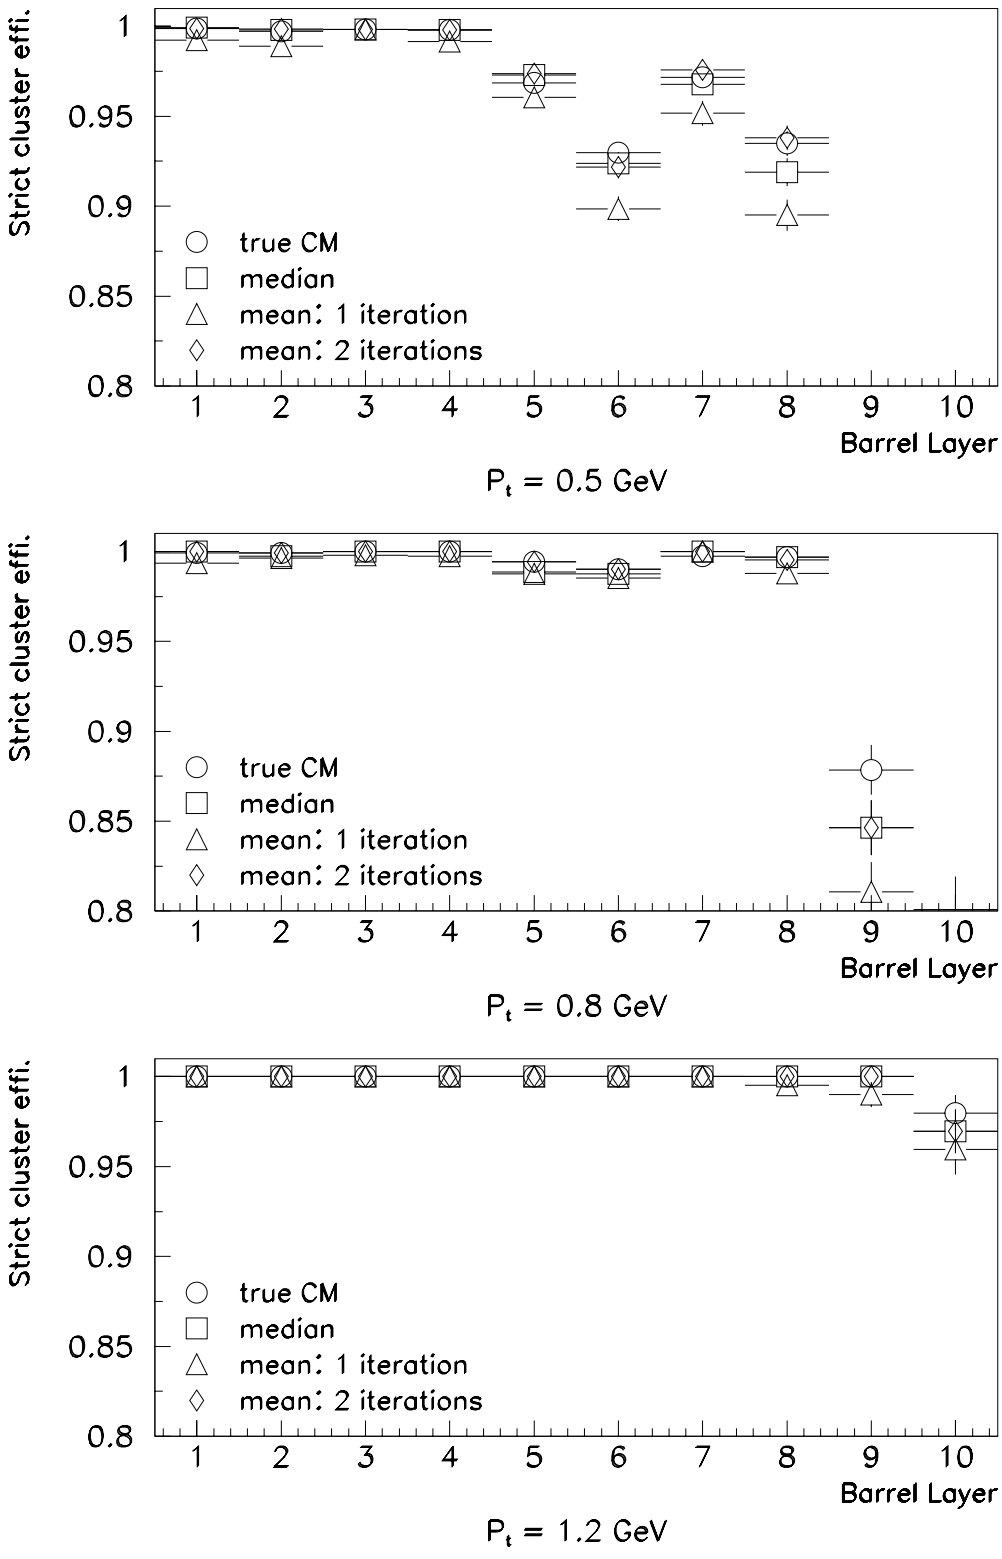

Figure 11: ORCA predictions for strict cluster finding efficiency vs. barrel layer in three P_t regions, for the FED 2 algorithm. The dependence on the algorithm used to estimate the common-mode offset is shown.

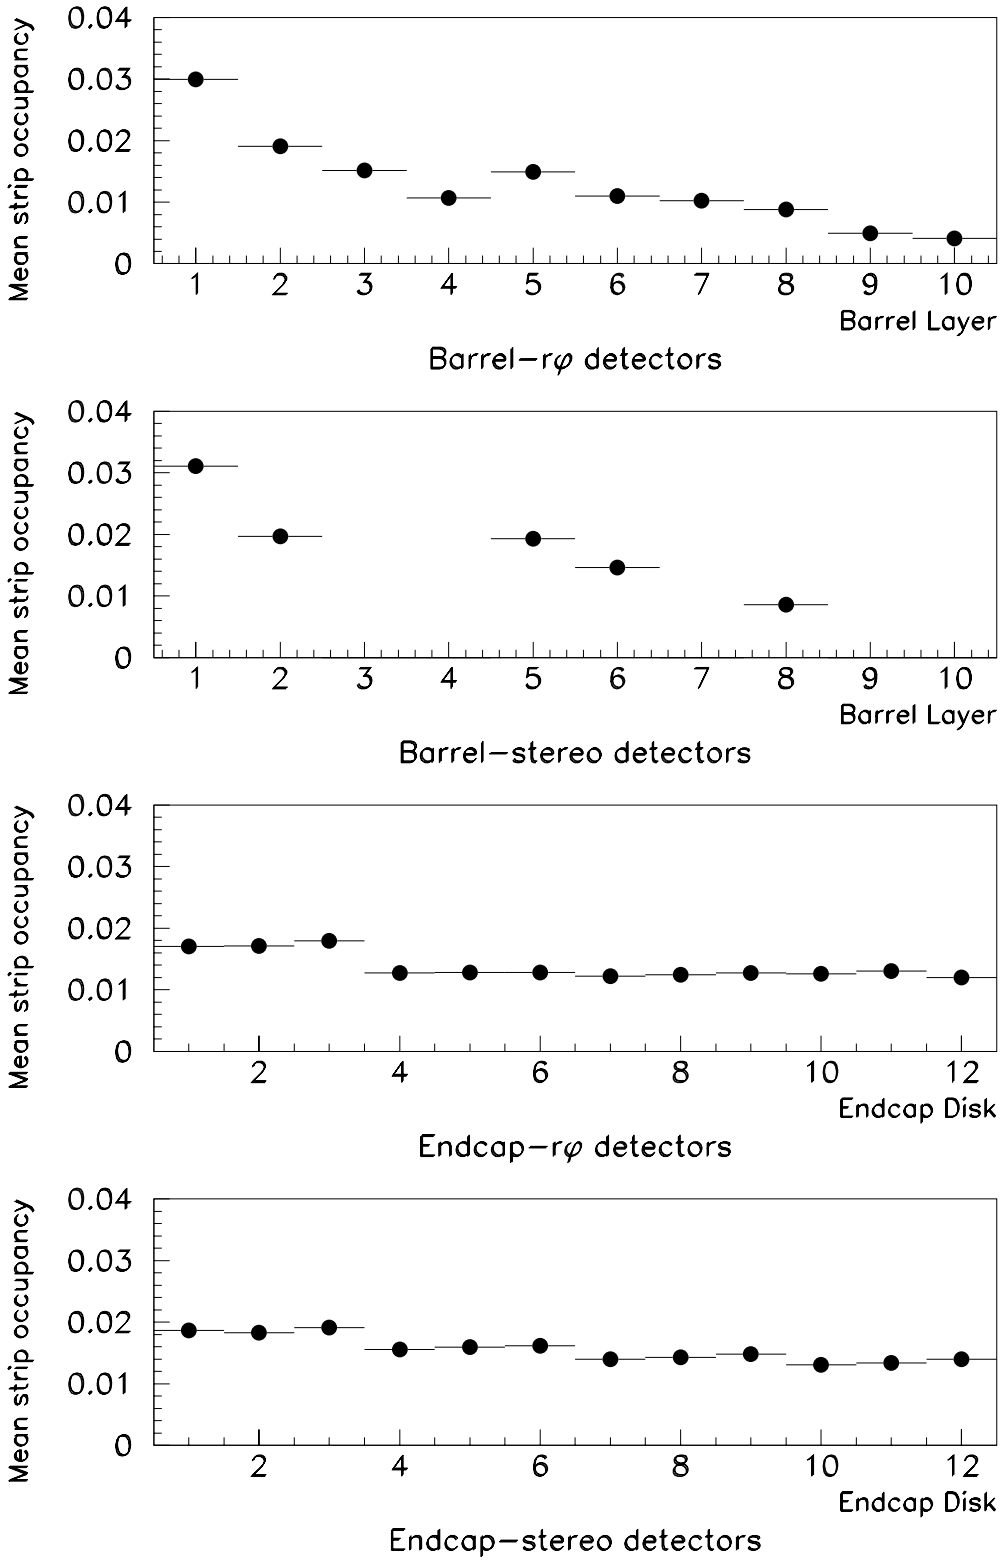

Figure 12: Expected strip occupancy in FED data, from minimum bias events during high luminosity running.

Table 2: The mean strip occupancy for various clustering algorithms. The occupancy is subdivided into four separate contributions: (i) prompt clusters (i.e., those produced within 5 ns of the nominal beam-crossing time) from tracks with $P_t > 0.5$ GeV/c, (ii) other (uninteresting) clusters produced in the correct beam-crossing, (iii) clusters from the wrong beam-crossing, (iv) fake clusters.

Clustering Algorithm	Common-Mode Algorithm	Strip Occupancy (%)				
		Prompt $P_t > 0.5$ GeV/c	Other in correct BX	Wrong BX	Fake	Total
FED 1	true CM	0.13	0.70	0.34	0.29	1.46
FED 2	true CM	0.14	0.78	0.35	0.09	1.36
FED 3	true CM	0.21	1.00	0.52	0.23	1.96
TDR	true CM	0.14	0.78	0.32	0.01	1.25
FED 2	mean: 1 iter.	0.13	0.68	0.29	0.07	1.17
FED 2	mean: 2 iter.	0.14	0.75	0.34	0.10	1.33
FED 2	median	0.14	0.74	0.33	0.10	1.31

Table 3: The mean number of clusters per detector module found by various clustering algorithms. The number is subdivided into four separate contributions: (i) prompt clusters (i.e., those produced within 5 ns of the nominal beam-crossing time) from tracks with $P_t > 0.5$ GeV/c, (ii) other (uninteresting) clusters produced in the correct beam-crossing, (iii) clusters from the wrong beam-crossing, (iv) fake clusters.

Clustering Algorithm	Common-Mode Algorithm	Number of Clusters / Detector				
		Prompt $P_t > 0.5$ GeV/c	Other in correct BX	Wrong BX	Fake	Total
FED 1	true CM	0.35	1.45	1.21	1.83	4.84
FED 2	true CM	0.34	1.30	0.78	0.29	2.71
FED 3	true CM	0.34	1.28	0.93	0.75	3.30
TDR	true CM	0.34	1.27	0.67	0.04	2.32

leads to a larger data volume if most clusters are one strip wide, but a smaller data volume if many clusters are three or more strips wide.

Figure 13 shows the mean data size per APV obtained using these two algorithms, as a function of barrel layer or end-cap disk. The clusters were found use the FED 2 algorithm, with the common-mode offset being estimated with the median algorithm. The results suggest that using the second data formatting scheme leads to 15% reduction in data volume. This confirms the results of a preliminary study by G. Pasztor [17]. One should note however, that the veracity of this statement is dependent on the clustering algorithm and thresholds used.

A data compression algorithm such as “Huffman” might be used after cluster finding to reduce data volume by an additional 25%, at the cost of extra complexity [17].

Further information on expected occupancies and data rates, together with an estimate of the total number of FEDs required in the tracker, can be found in Ref. [16]. This includes

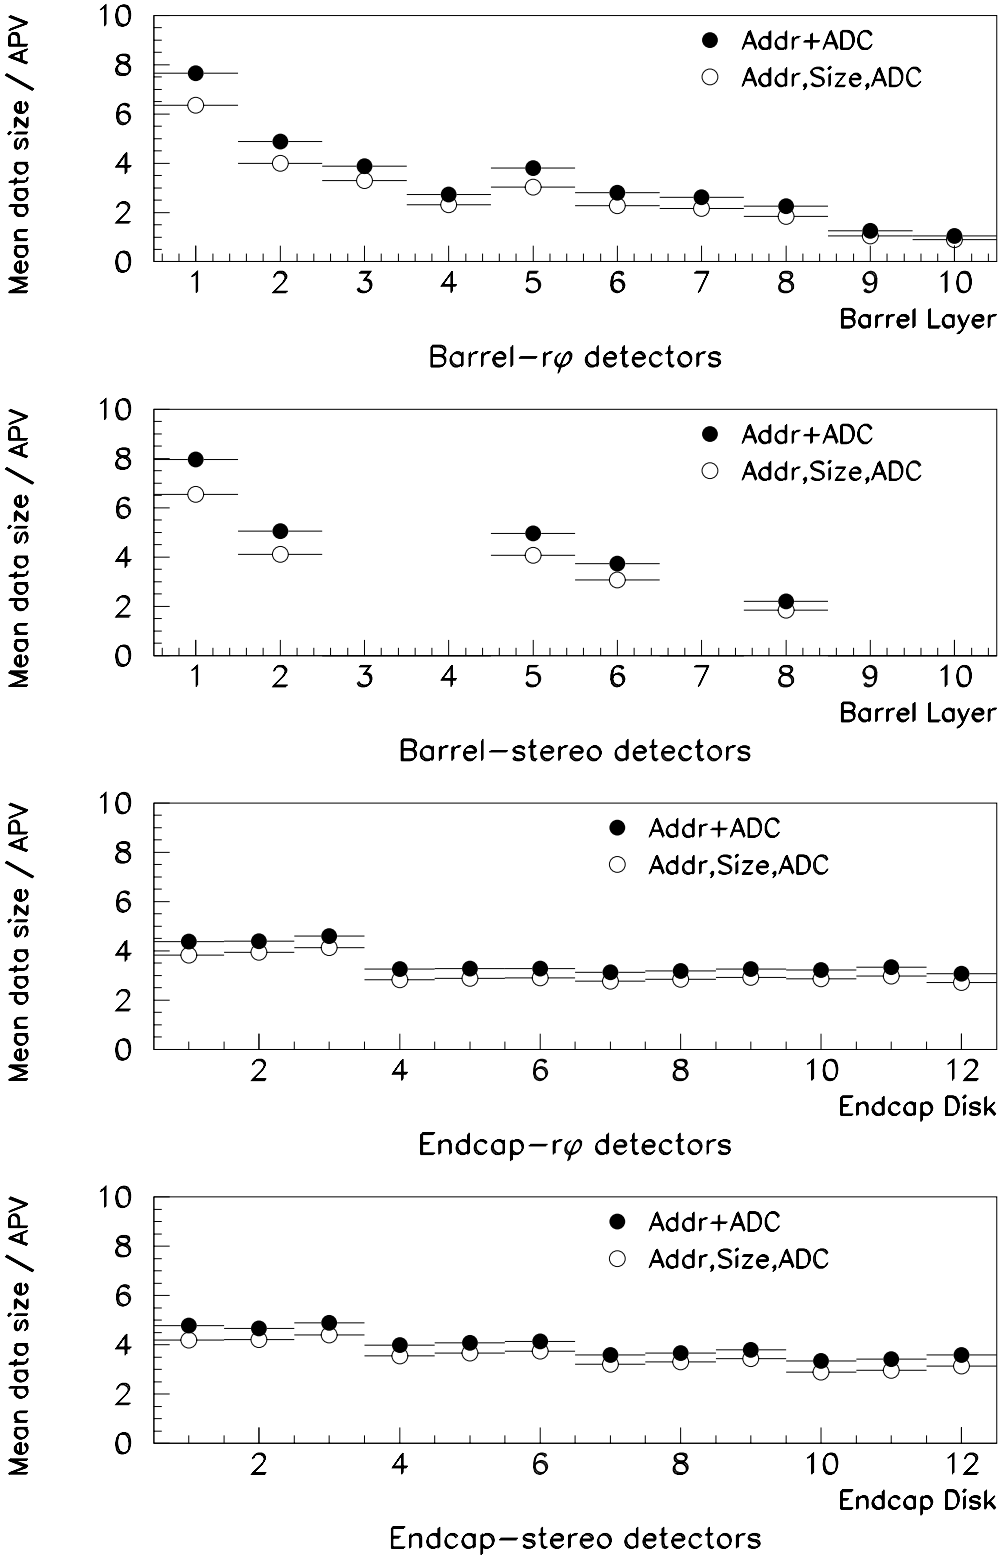

Figure 13: Mean data size per APV using two alternative data formats. Clusters were found using the FED 2 algorithm. High luminosity conditions were assumed.

information on occupancies in heavy ion collisions.

8 Conclusions

The FED 1 algorithm gives excellent performance for tracks with $P_t > 1.2$ GeV/c. At lower momenta, its cluster finding efficiency is inadequate however, especially in the outer detector layers. Nonetheless, it is possible that some of this efficiency might be recovered by using looser cuts in the outer layers, and it should be borne in mind that most physics analyses do not need such low momentum tracks. One should also remember that a pessimistic value was assumed for the detector S/N ratio. Doubtless, the FED 1 algorithm would give much better performance if the S/N ratio in the final tracker is really 13–15 as seems likely.

Nonetheless, the FED 2 algorithm gives significantly higher efficiency and somewhat lower occupancies. It would therefore be preferable to implement this in the FED if FPGA resources permit this. The clustering algorithm could be implemented in the FED by looping through the strips sequentially, and deciding if a strip should be kept, by checking if it and its immediate neighbours exceed the requested thresholds.

The FED 3 algorithm gives excellent cluster finding efficiencies, but is disfavoured on the basis of the high occupancies it yields.

When deciding if the pulse height on a given strip exceeds a threshold T, rather than checking if $S/N > T$, which would demand a floating point division, it would be preferable to check if $S > T \times N$, where the product $T \times N$ is passed to the FED as a calibration constant. Since the noise can vary from strip to strip, this would imply that $2 + 1 = 3$ registers would be needed in the FED for each strip it processes, to contain both the two clustering and the one common-mode thresholds. This would also have two additional advantages:

- Different thresholds could be used on the various strips processed by a FED, which may be useful if it reads detectors from different detector layers.
- Noisy strips could be suppressed by setting their thresholds to infinity.

Regarding the algorithm for estimating the common-mode offset, it is apparent that the median algorithm performs well. Alternatively, two iterations of the mean algorithm may be used, but in this case, it may be necessary to suppress even moderately noisy strips during the common-mode determination, as these add a large statistical uncertainty to the mean. (Offline analyses are insensitive to this effect, as they weight the pulse heights according to the inverse noise squared when estimating the common-mode offset. Unfortunately, this would not be practicable in the FED.) As the common-mode offset is a floating point number, some thought must be given to whether it is truncated.

Regarding the format of the FED data: it is probably not optimal to output the address and ADC count of each selected strip. Instead, it is better to output the address of the first strip in a cluster, the cluster size and the ADC pulse height of each strip in the cluster. This reduces the data volume by roughly 15%, with the exact reduction being dependent on the cluster algorithm and thresholds used.

Acknowledgements

I'm particularly grateful to Duccio Abbaneo for providing much of the code for the toy Monte Carlo and to Vitaliano Ciulli for providing much needed example software for ORCA tracker analysis.

References

[1] R. Halsall, private communication.

[2] *"The Tracker Layout Working Page"*, http://cern.ch/duccio/layout/lay.html .

[3] M. Mannelli, private communication.

[4] The toy Monte Carlo was derived from an earlier version developed for MSGC studies by D. Abbaneo.

[5] G. Cowan, *"Statistical Data Analysis"*, Oxford Science Publications (1998), Chapter 2.9.

[6] S. Braibant *et al.*, CMS note 2000/050, *"Test-beam results on* $\langle 100 \rangle$ *silicon prototype detectors with APV6 front-end chip readout"*.

[7] Private communication from CMS b–τ group.

[8] M. Rehn *et al.*, *"Statistical Study of SVX3D Chip Dynamic Pedestal Subtraction Threshold Level"*, CDF/DOC/SEC_VTX/PUBLIC/4852 (1999).

[9] R. Yarema *et al.*, *"A beginners guide to the SVXIIE"*, FERMILAB-TM-1892, page 4 (1996).

[10] H. Seywerd, *"Sirocco programs for the VDET II"*, ALEPH note 97-109, page 3 (1997).

[11] V. Chabaud *et al.*, *"The DELPHI silicon strip microvertex detector with double sided readout"*, CERN-PPE/95-86, pages 9-10 (1995).

[12] P. Allport *et al.*, *"The OPAL silicon microvertex detector"*, Nucl. Instr. Meth. **A324** (1993) 34.

[13] CMS Tracker TDR, CERN/LHCC 98-6 (1998)

[14] M. French, W. Gannon, private communication.

[15] A. Racz *et al.*, *"Data Acquisition for Heavy Ion Physics"*, CMS IN 2000/027.

[16] A. Caner *et al.* *"On balancing data flow from the silicon tracker"*, http://cern.ch/tomalini/readout.pdf . http://pcvlsi5.cern.ch:80/CMSTControl/documents/IanTomalin/readout.pdf .

[17] G. Pasztor, *"Data Compression in FED ?"*, http://trackercontrol.cern.ch/CMSTControl/documents/GabriellaPasztor/Gabriella.pdf .