A first look at jet pattern recognition using reduced tracking layouts

Alan Poppleton and *Giorgos Stavropoulos*

Abstract

The reconstruction performance of H $\rightarrow$ b$\overline{\text{b}}$ and H $\rightarrow$ u$\overline{\text{u}}$ with m_H = 400 GeV using reduced layouts (with removal of a pixel or/and SCT layer, with/without the TRT) is compared to the performance of the full layout presented in the Inner Detector Technical Design Report. The data samples and analysis procedure are the same as used for the TDR. Reconstruction is performed with the TDR version of the iPatRec program apart from a modification to flag the removed layers as inert material. The relative contributions of the pixel and SCT (discrete) and TRT (continuous) tracking to pattern recognition in dense jets can be clearly seen. Before taking any decisions the study should be extended to explore the robustness of these layouts to high luminosity pile-up and detector degradation effects. The physics benefits of the removed material should also be studied.

1 Introduction

The pixel B-layer is now accepted as a permanent addition to the Inner Detector. Since this represents an increase in the number (and cost) of tracking layers, Atlas management has requested that a corresponding saving should be made elsewhere in the ID. There is also concern expressed by the LHCC that the ID material is at the upper end of what can be considered as acceptable in front of state-of-the-art calorimetry. Removal of tracking layer(s) will help to alleviate this problem.

The most challenging pattern recognition study addressed in the TDR concerns the reconstruction and subsequent tagging of $H \rightarrow b\bar{b}$ and $H \rightarrow u\bar{u}$ with $m_H = 400$ GeV [1][2]. The purpose of this note is to report on a repeat of that study with various reduced layouts. To facilitate a comparison of the pattern recognition capabilities of these layouts, only the minimum necessary set of changes have been made from the TDR procedure. Thus the same generated datasets (DICE digitisation banks) have been used. The only change to the reconstruction code (using iPatRec[3][4]) is to flag the 'removed' layers as inert material (i.e. to take Coulomb scattering into account during track-fitting but with no hit decoding). More specifically, some known bug-fixes and improvements (in particular to the pixel clustering algorithm) have not been implemented. During analysis, any reconstructed secondary tracks originating from 'removed' layers have been ignored and the quality cuts on the number of hits on a track adapted as appropriate for each layout, otherwise the analysis procedure remains unchanged. The definitions of primary, secondary and fake tracks, spoilt and wrong hits, are taken from the TDR.

The reduced layouts are described in Section 2, with the basic pattern recognition performance presented in Section 3. Removal of discrete layers is discussed in Section 4, the TRT contribution to pattern recognition in Section 5. Robustness issues are considered in Section 6.
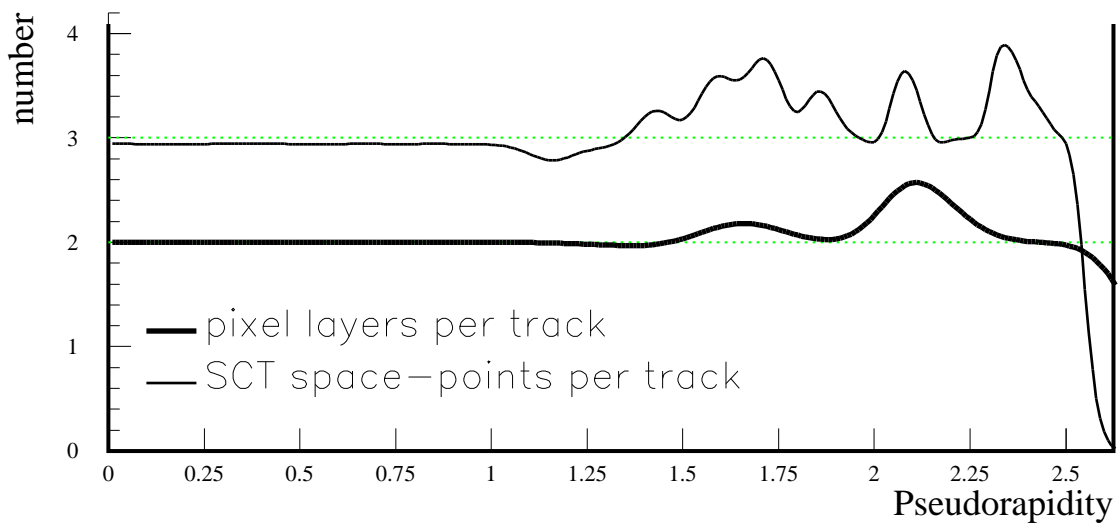
2 Reduced layouts
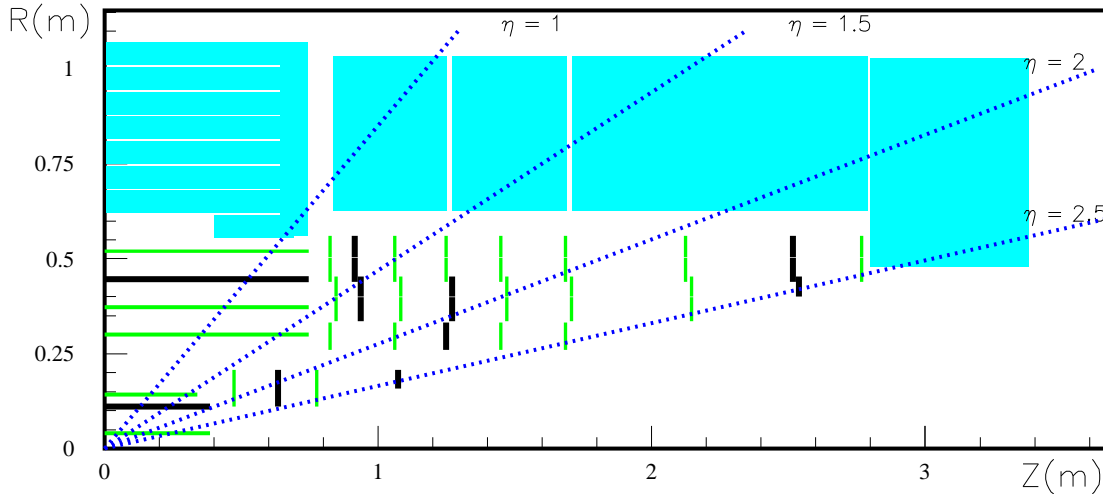
Figure 2-1 Schematic diagram of the Inner Detector layout. The 'removed' pixel and SCT layers are shown in bold.

The reduced layouts should preserve the inner and outermost layers of pixels and SCT as the removal of any of these layers would imply a descoping of the sub-detector performance. Thus if one layer is to be removed, it should be either the pixel barrel layer at radius 11.2 cm or the merging of SCT barrel layers 2 and 3 into one intermediate layer at radius 41 cm. The corresponding changes to the end-cap would retain 2 instead of 4 pixel disks or 7 instead of 9 SCT wheels with a re-optimization of the disk/wheel positions. However, to be able to re-use the existing data-sets, the reconstruction program has been given the possibility to ignore digitisations from the layers drawn in bold on Figure 2-1. For the pixels, these are the barrel layer at radius 11.2 cm and disks 2 and 4 (at Z = 63 and 107 cm); for the SCT, the barrel at radius 44.7 cm, wheels 2 and 7 (at Z = 93 and 253 cm) plus the inner and middle module rings from wheel 4 (Z = 126 cm). This arrangement satisfies the requirement of 2 pixel or 3 SCT layers hit as can be seen from Figures 2-2 and 2-3.

The TDR dataset consists of 2.4K H → b$\overline{\text{b}}$ and 20K H → u$\overline{\text{u}}$ events, giving ~4500 b-jets and ~37000 u-jets inside the ID acceptance. The full statistics have been processed, using iPatRec to reconstruct tracks with p_T > 1 GeV in a 'cone' of ΔR = 0.4 around each jet, with the following layout scenarios:

1. the TDR layout – also referred to as *3+4* (and depicted by a shaded circle symbol in the figures)

2. removal of one pixel layer – *2+4* (apex-upwards triangle)

3. removal of one SCT layer – *3+3* (apex-downwards triangle)

4. removal of both the preceding layers – *2+3* (shaded square)

5. removal of the TRT in the TDR case – *3+4$\overline{T}$* (open circle)

6. removal of the TRT in the *2+3* case – *2+3$\overline{T}$* (open square)

Figure 2-2 The mean number of pixel layers and SCT space-points versus |η| in the 'reduced' layouts
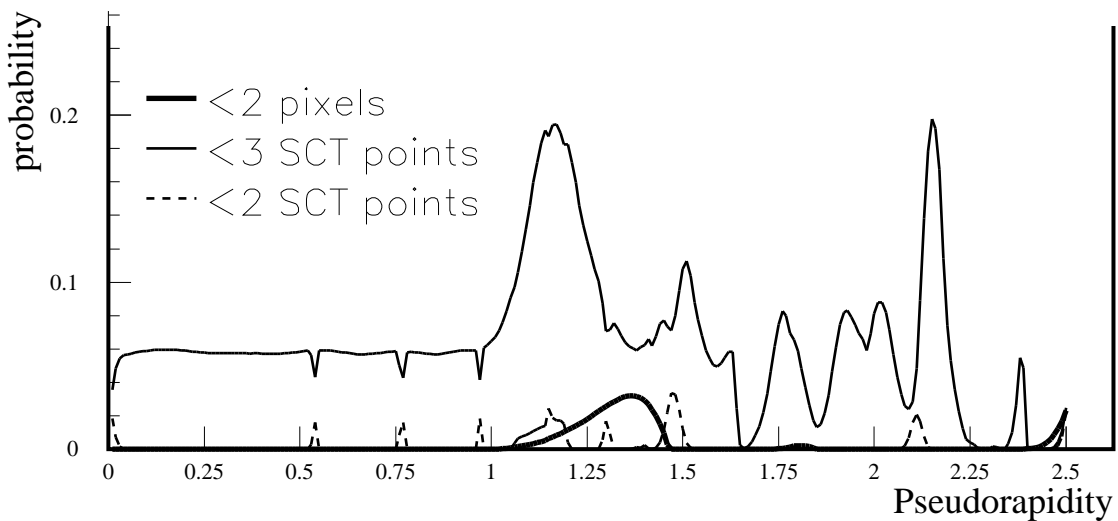
Figure 2-3 Probability of failing the 2 pixel or 3 SCT space-point specification for the reduced layouts. The most significant loss in each case is at the barrel/end-cap transition (pixels: |η| ~ 1.3, SCT: |η| ~ 1.1).

3 Basic pattern recognition performance

In the TDR, selection cuts are applied to the tracks found by iPatRec in order to give an optimal b-tagging performance. By design these cuts select tracks with the best possible impact parameter resolution and reduce the proportion of tracks from conversions, V^0 decays, secondary interactions and fakes (all of these are background contributors with significant impact parameters). The cuts are:

1. Number of precision hits ≥ 9

2. Number of pixel hits ≥ 2

3. At least one associated hit in the B-layer

4. $|d_0| \leq 1$ mm

Broadly similar cuts have been applied to enhance signal during the other TDR physics analyses. However, cuts on the number of hits need to be adapted as appropriate for each layout. For the TDR layout the average number of precision hits is 11 (3 from the 3 pixel layers and 8 from the 4 SCT layers). So for the layouts with an SCT layer removed the precision hits requirement is reduced by 2 (e.g. to ≥ 7 for *3+3*). Similarly for the layouts with a pixel layer removed, *2+4* or *2+3*, the minimum number of precision hits is respectively 8 or 6, and the minimum number of pixel hits falls to 1. For each layout there are about 320000 tracks which pass these standard selection cuts.

Secondary tracks reconstructed from interactions and conversions on the 'removed' layers have been excluded from the following analysis, but no attempt has been made to compensate for the consequent reduction in service/support material.

Table 3-1   Principle performance characteristics of the different layouts with standard selection cuts

|	TDR Layout	Layout *2+4*	Layout *3+3*	Layout *2+3*	Layout *3+4T̄*	Layout *2+3T̄*
Efficiency	88.2 %	86.4 %	89.9 %	86.6 %	88.6 %	84.0 %
Fake Rate	0.1 %	0.3 %	0.3 %	1.6 %	0.3 %	3.2%
Secondaries after B-layer	0.6 %	1.0 %	0.6 %	1.5 %	0.5 %	1.6 %
Wrong Pixel	1.4 %	4.0 %	1.7 %	5.6 %	1.2 %	7.3 %
Wrong SCT	1.2 %	1.6 %	1.4 %	3.0 %	1.5 %	4.7 %
Spoilt Pixel	6.0 %	5.7 %	6.2 %	5.7 %	6.2 %	5.6 %
Spoilt SCT	17.0 %	16.3 %	13.9 %	12.8 %	17.5 %	12.6 %
Tail $1/p_T$	3.1 %	3.7 %	3.5 %	5.0 %	1.9 %	5.7 %
Tail d_0	1.9 %	3.6 %	2.2 %	5.1 %	1.6 %	6.7 %
u-rejection	66	36	62	25	72	17

Following the selection cuts, some significant pattern recognition performance characteristics of the different layouts are summarized in Table 3-1. Efficiency refers to the track-finding efficiency for primary tracks; the fake rate is given with respect to the total number of reconstructed tracks passing the cuts; the rate of secondaries is given for those produced beyond the B-layer

(thus by definition with a wrong pixel association); the wrong and spoilt rates refer to the fraction of tracks containing any wrong/spoilt pixel/SCT hit association; the tail rates are for primary tracks reconstructed at more than 3σ from the true value; the u-jet rejection is quoted for 50% b-tagging efficiency.

Adapting the cut on the number of hits according to layout keeps the track-finding efficiency roughly constant (at between 84% and 90%), but the rate of fakes increases in a non-linear manner (0.1/0.3/1.6%) as layers are removed. The $|\eta|$-dependence of the track-finding efficiency and fake-track rate is shown in Figures 3-1 and 3-2. In the TDR, it was stressed that the domi-
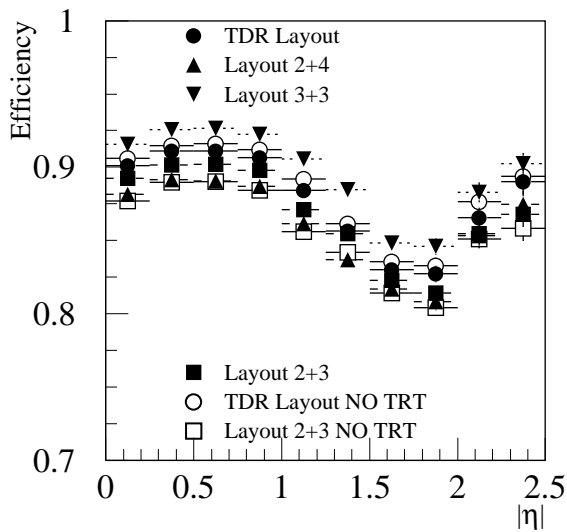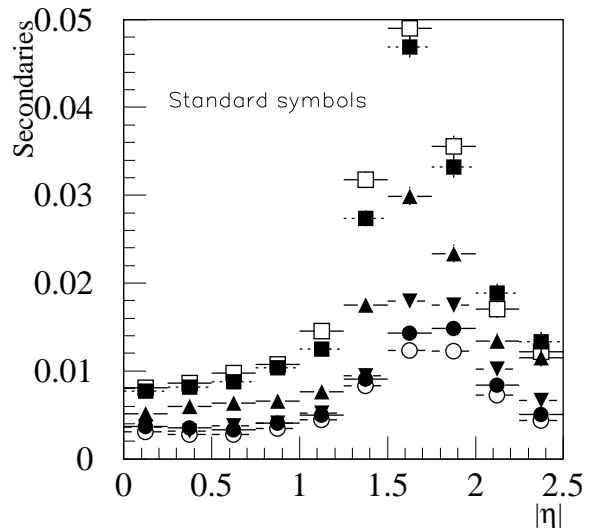

Figure 3-1 The variation of track reconstruction efficiency with $|\eta|$, after standard selection cuts, for the 6 layout scenarios.
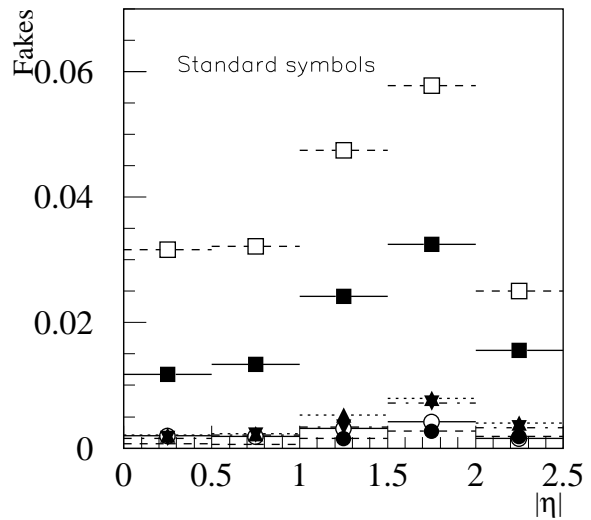
Figure 3-2 The corresponding rate of fake tracks after standard selection cuts.
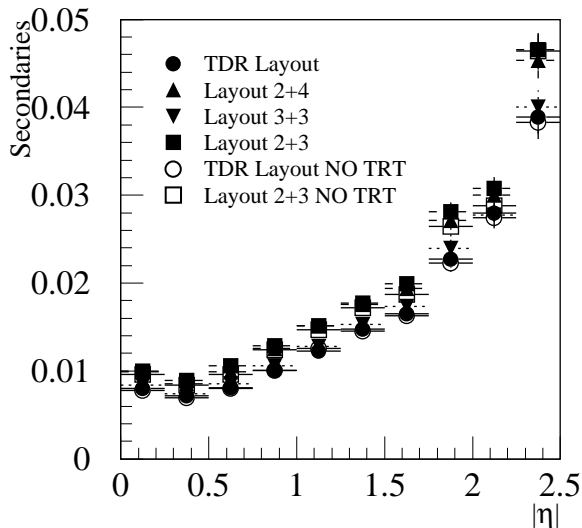
Figure 3-3 The rate of secondary tracks produced on or before the B-layer satisfying the standard selection cuts.


Figure 3-4 Corresponding rate of secondaries produced after the B-layer.

nant limitation to the b-tagging performance came from secondaries, in particular those origi-
nating after the B-layer. The rate of these, passing the cuts, increases significantly when a pixel
layer is removed. The rate of secondaries is shown as a function of $|\eta|$ in Figures 3-3 and 3-4,
with the contributions from those produced before and after the B-layer plotted separately. The
$|\eta|$-dependence of the efficiency, fake and post-B-layer secondary curves all reflect the amount
of discrete tracker material, with the effect of the notorious pixel 'dog-leg' services being partic-
ularly apparent $(1.4 < |\eta| < 1.9)$. Note that the performance of layouts with a pixel layer re-
moved are pessimistic in this region as some of this 'dog-leg' material would also be removed.
The early secondaries 'see' only the barrel geometry of the beam-pipe and B-layer material,
hence exhibit a $1/\sin(\theta)$ shape.

Wrong pixel associations (which are mainly observed in the B-layer) increase sharply when a
pixel layer is removed. This implies a need for 2 layers between the B-layer and the SCT layer
starting at R = 30 cm. Reduced layouts lead to an increase in wrong SCT associations which
seems to depend more on the total number of layers rather than whether they are pixel or SCT
(or the TRT). Spoilt hits are mainly caused by a lack of detector resolution, and occur with a rate
of ~0.5% hit clusters in the outer pixel layers, ~2% hits in the SCT and ~5% in the B-layer. Thus
for tracks, spoilt pixel rates are almost layout independent as they are predominantly due to the
B-layer, while the SCT rate depends on the number of SCT layers present. Despite the high rate
of tracks with spoilt SCT hits, there is no serious effect as the hits are de-weighted according to
cluster size during track-fitting. However the pixel layers produce broad clusters by design[5],
to improve their precision on isolated tracks, thus spoilt pixel clusters are over-weighted and
can cause a serious limitation to the overall performance.

Table 3-1 also shows that the effect of the TRT on all of the above quantities is approximately the
equivalent of 1 SCT layer.

The fraction of primary tracks in the tail of the momentum resolution increases as layers are re-
moved, apart from layout *3+4T* where the TRT appears to contribute to the tail. The impact pa-
rameter tail is particularly sensitive to the number of pixel layers because of the wrong B-layer
associations previously mentioned.
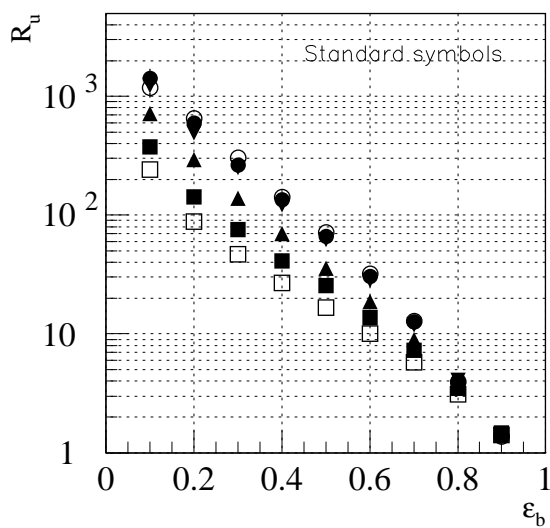
Figure 3-5 u-jet rejection versus b-tagging efficiency
obtained with the standard selection cuts.
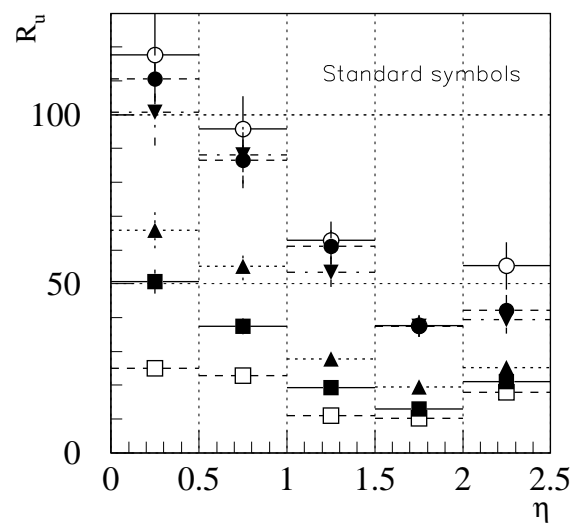
Figure 3-6 u-jet rejection at 50% b-tagging efficiency
as a function of $|\eta|$.

For a given detector resolution, u-jet rejection is again limited by the wrong B-layer associations (which include fakes and secondaries produced after the B-layer), with additional contributions from secondaries produced on or before the B-layer and from tails on the primary track impact parameter resolution. Figure 3-5 shows the u-jet rejection versus b-tagging efficiency obtained for the different layouts using the standard selection cuts.

Figure 3-6 shows the variation of the rejection with $|\eta|$ at a b-tagging efficiency of 50%. This reflects the material distribution as discussed above, with a fall in rejection of almost a factor 3 in going from the central barrel to the 'dog-leg' region. There is a clear degradation if a pixel layer or more than one layer or detector is removed. Each layout exhibits a broadly similar relative performance as a function of $|\eta|$ despite the variation in material thickness.

4 Number of discrete layers

In the previous section, it was shown that applying the standard selection procedure gave a marked degradation in performance when a pixel layer was removed. Removal of an SCT layer gave a less pronounced effect – even a small increase in track-finding efficiency. Additional understanding can be gleaned by repeating the analysis with identical cuts applied to each layout. Here, the first cut is relaxed to a 7 hit requirement to avoid undue loss of efficiency for the reduced layouts, while a 2 pixel minimum is invoked to avoid the high rate of wrong associations observed on tracks formed from the B-layer and SCT without intervening hits. Note that this 2 pixel minimum is somewhat optimistic for the reduced pixel layouts as demanding hits in both layers leaves no contingency for detector degradation effects (inefficiences). The $|\eta|$-dependence of the track-finding efficiency and fake-track rate is shown in Figures 4-1 and 4-2 (only the 4 layouts with TRT to illustrate the main points without unnecessary detail). With these cuts, the fake rate becomes almost layout independent whereas the efficiency falls smoothly as layers are removed. Removal of a pixel layer reduces the efficiency more than removal of an SCT layer.
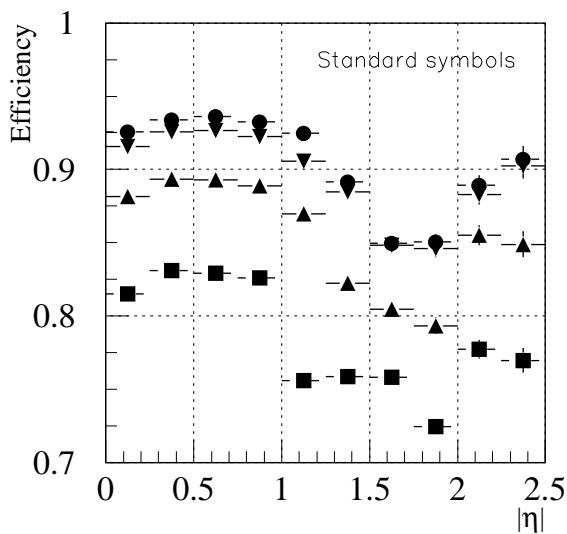
Figure 4-1 The variation of track reconstruction efficiency with |η|, using identical selection cuts.
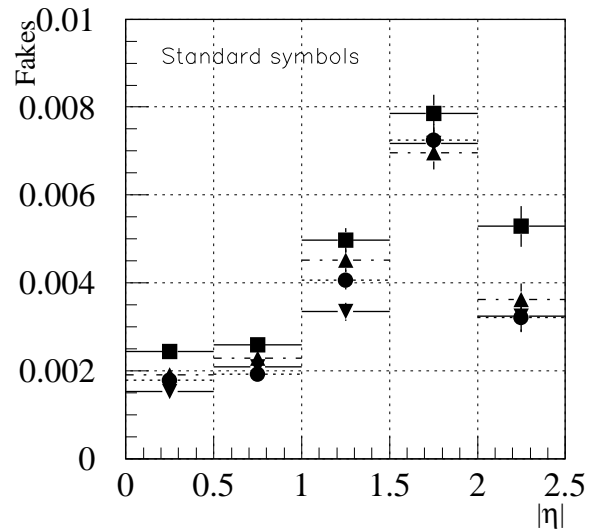
**Figure 4-2** Corresponding rate of fake tracks.

Note the low efficiency in the SCT barrel/forward transition region when both layers are removed.

The principle performance characteristics are summarized in Table 4-1. The '7 hits + 2 pixel' cuts applied to the TDR layout select more secondaries with a wrong B-layer association than the standard cuts. These secondaries have insufficient hits when an SCT layer is removed, hence the small increase in rejection power in this case. There is a significant loss of rejection when a pixel layer is removed due to the fall in track-finding efficiency and increased tail in d_0.

Table 4-1 Performance characteristics with identical cuts (7 hits + 2 pixels)

	TDR Layout	Layout 2+4	Layout 3+3	Layout 2+3
Efficiency	90.7 %	86.2 %	89.9 %	79.4 %
Fake Rate	0.3 %	0.3 %	0.3 %	0.4 %
Secondaries after B-layer	0.8 %	0.5 %	0.6 %	0.4 %
Wrong Pixel	1.7 %	2.3 %	1.7 %	2.1 %
Wrong SCT	1.5 %	1.4 %	1.4 %	1.2 %
Tails $1/p_T$	3.5 %	3.4 %	3.5 %	3.2 %
Tails d_0	2.2 %	2.7 %	2.2 %	2.5 %
u-rejection	55	48	62	44
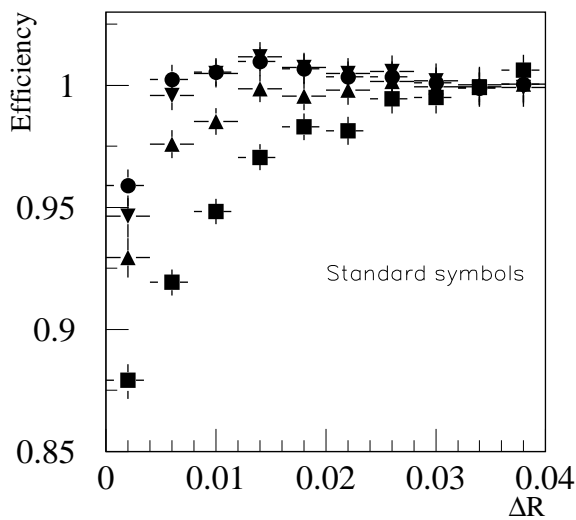
Figure 4-3 Relative track-finding efficiency against angular separation from nearest generated neighbour (ΔR).
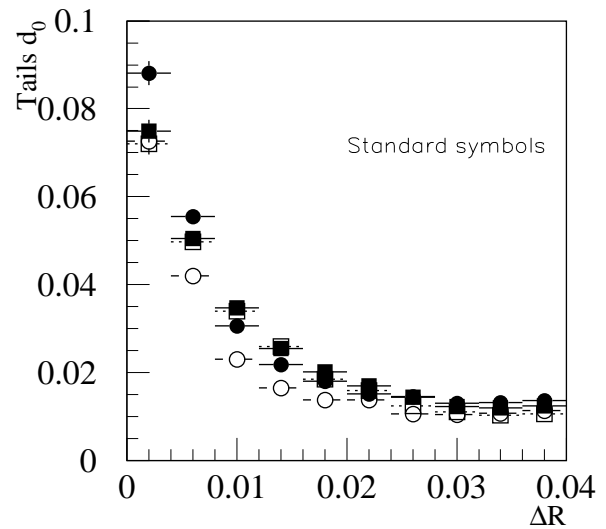
Figure 4-4 Tail on d_0 (fraction of tracks > 3σ) against ΔR.

Figure 4-3 shows the relative track-finding efficiency as a function of the angular separation of a track from its nearest generated neighbour (ΔR). The efficiency has been normalised to asymptote to unity at large ΔR, but as iPatRec is less efficient for the lower p_T tracks observed at large

ΔR, the relative efficiency tends to exceed one before dipping as the jet core is approached. More layers provide a better performance at small ΔR and, yet again, the plot is most sensitive to the removal of a pixel layer.

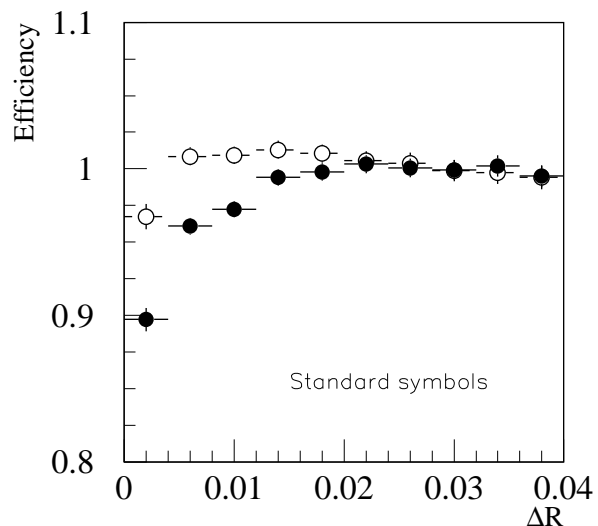Figure 4-4 shows the tail to the impact parameter measurement as a function of ΔR. A serious, but essentially layout independent (as it is associated with the B-layer), degradation is apparent at small ΔR. This should be improved, but will not be eliminated, by the modifications currently under consideration by the pixel group[6]. The reconstruction should be modified to correctly weight pixel clusters associated to more than one track (this will only be possible in the cases where the additional tracks have been successfully reconstructed – thus will be sensitive to the layout). The b-tagging analysis procedure could also be updated to compensate for the tails in this region.

From the results presented so far, it appears that the b-tagging performance improves as the number of discrete layers increases, despite the clear dependence of many of the relevant parameters on the material thickness. However, the benefit of the 4[th] SCT layer is the least pronounced in the context of an impact parameter sensitive analysis. Before reaching any conclusions on the merits of 3 or 4 SCT layers, it is necessary to examine the combined performance of the complete outer tracking system (SCT + TRT).

5 Contribution of the TRT

To satisfy the ID specifications on track parameter resolutions, it is sufficient to match the TRT to a precise point in the vertex region and to another precise point in the sagitta region. Thus, for isolated tracks, 4 discrete layers would be sufficient for reconstruction with a conservative level of robustness against detector inefficiency. The justification for the additional layers of discrete tracking is to improve the association of hits in the more demanding pattern recognition environment of a high energy jet, since the discrete tracking layers provide a much superior 2-track resolution.

In iPatRec, pattern recognition is performed starting from space-points in the most favourable region of the discrete layers (i.e. from beyond the B-layer). The TRT is associated to track candidates to improve the parameter resolutions, in particular the additional measurements at large radius improve the momentum

Figure 5-1   Relative track-finding efficiency at small angular separation (ΔR) with/without the TRT.

precision by almost a factor 3. However, when ambiguous discrete hit associations do occur, the track candidates with TRT confirmation are preferred. The confirmation criterion is that ≥ 60% of the straws crossed should have compatible drift-time hits. Around 10% of primary tracks fail to associate to the TRT, with the major conributions coming from secondary interactions in the SCT/TRT material and from tracks with insufficient drift-time hits resolved in the TRT.

Figure 5-1 compares, as a function of track separation from its closest neighbour, the relative re-construction efficiency for all tracks (i.e. with no analysis cuts) in the layout *3+4T*, to those with TRT association in the TDR layout. It is apparent that explicitly demanding the TRT reduces the efficiency for finding close tracks. The missing tracks are reconstructed however, but do not have TRT association.

In Section 3 it was shown that, for the standard selection cuts, the TRT made a positive contribution to pattern recognition performance with the 5 layer scenario (where it limited the number of wrong associations between the SCT and B-layer). However providing (and demanding) intermediate pixel hit(s) fulfilled the same purpose, and in this case the addition of the TRT proved slightly detrimental. The reason for this becomes apparent when the tails of the reconstructed parameters are examined. Figures 5-2, 5-3 and 5-4 show the fraction of tracks in the tail
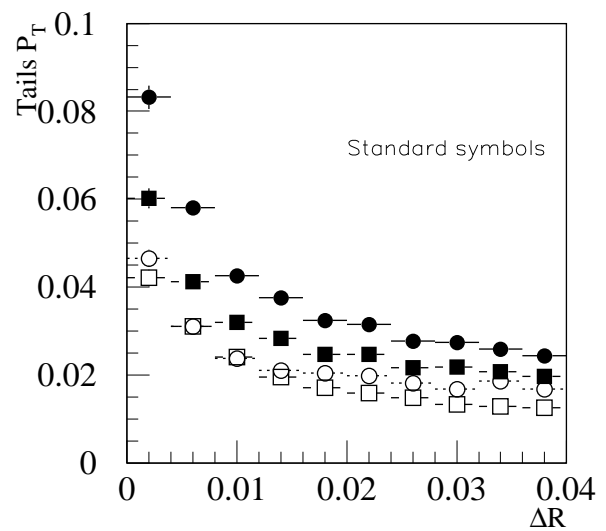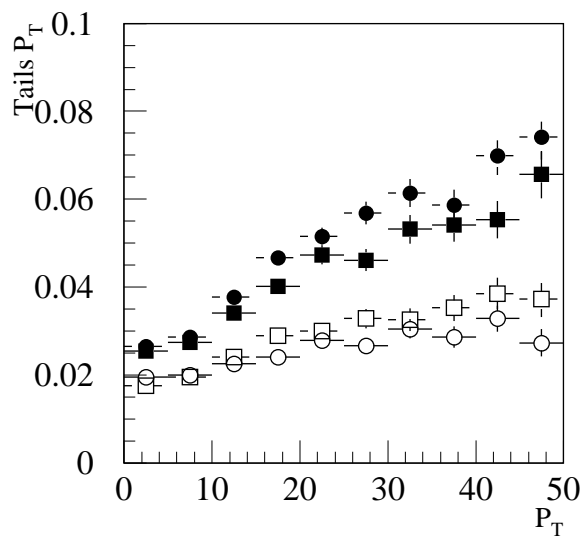
Figure 5-2 The variation of the tail on p_T with p_T shown for layouts with and without TRT.

Figure 5-3 Variation of the tail on p_T with ΔR.

of the p_T distribution as a function of p_T, ΔR and $|\eta|$, respectively. All indicate a substantial increase in the tail population when the TRT is included. This is not a material effect as the increase does not vary significantly with $|\eta|$. The rapid increase at small track separation (ΔR) and relative increase with p_T indicate that it is a 2-track resolution effect. The consequences are seen mainly on momentum resolution, Figure 5-5 shows the TRT influence on d_0 tails (hence b-tagging performance) to be less pronounced.

Additional software development is needed to identify and either re-assign or de-weight straw hits shared by more than one track. This is only possible if all the concerned tracks have been reconstructed. As the highest occupancy is in the innermost layers of the TRT, more use could be made of the SCT to unravel the confusion here. A momentum resolution sensitive analysis would be a useful test of whether 3 or 4 SCT layers are really needed.
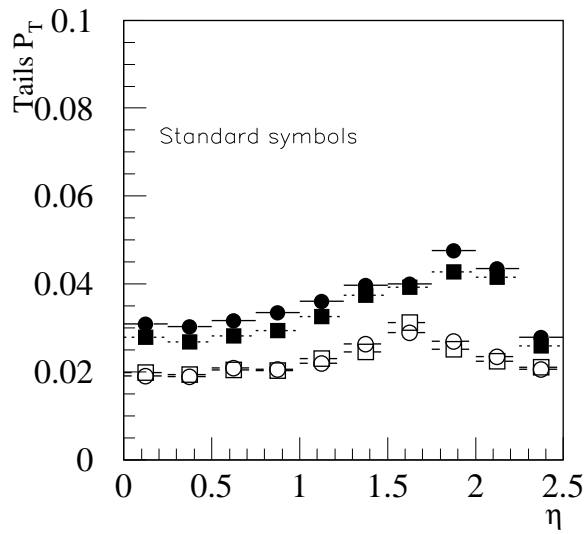
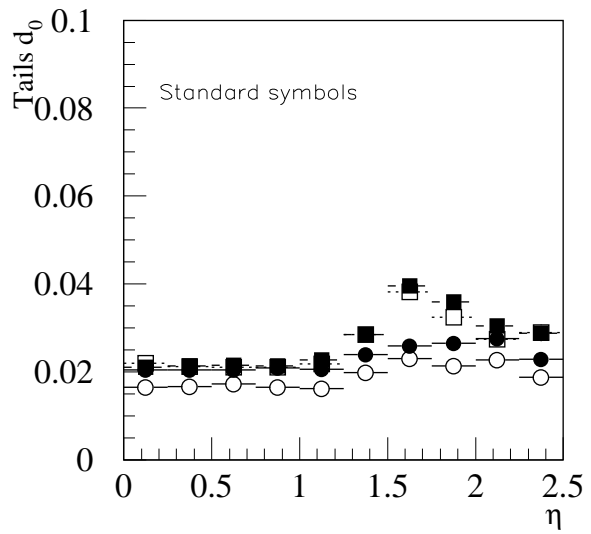**Figure 5-4** Variation of the tail on p_T with $|\eta|$.

Figure 5-5 Variation of the tail on $|d_0|$ with $|\eta|$.

6 Robustness

Another consideration is the realistic robustness of the Inner Detector. For the TDR simulations, a random efficiency of 97% was introduced at the digitisation step. Clustering effects in the pixel/strip detectors translate into a higher effective efficiency, especially for the pixel detectors which have broad clusters by design. Given the very restricted access scenario for the Inner De-
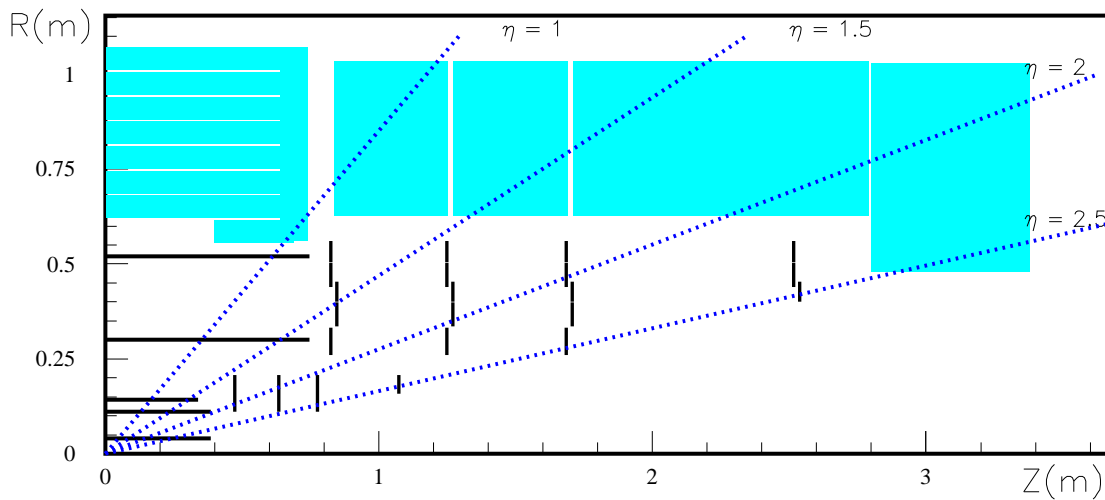
Figure 6-1 The *3+2* layout has 2 SCT barrel layers and 4 end-cap wheels.

tector, some contingency for dead detectors, read-out chips, modules, etc, should be provided. There is a need for degradation studies where inefficiency is also introduced for groups of pixels and strips. Meanwhile a more realistic estimate of the 7-layer performance is that it would be a combination of the TDR, *2+4* and *3+3* layout performances. Following this logic, a *3+2* layout has also been simulated in order to give a feeling of realism for the *3+3* performance. A schematic diagram of this layout is shown in Figure 6-1.

The standard cuts for layout *3+2* become a minimum of 2 pixels and 5 discrete hits. Again there is a correction to remove secondaries produced on the missing layers. The results are summarised in Table 6-1, with the first 4 columns repeated from Table 3-1 to facilitate comparison. The TDR layout is in the first column, with the next 2 columns indicating the effect of 1 inefficient layer. Similarly the rightmost 2 columns indicate the degradation expected for a 1 layer inefficiency on the performance of the *3+3* layout.

Table 6-1   Performance characteristics with standard selection cuts illustrating robustness

	TDR Layout	Layout *2+4*	Layout *3+3*	Layout *2+3*	Layout *3+2*
Efficiency	88.2 %	86.4 %	89.9 %	86.6 %	89.6 %
Fake Rate	0.1 %	0.3 %	0.3 %	1.6 %	1.1 %
Secondaries after B-layer	0.6 %	1.0 %	0.6 %	1.5 %	0.7 %
Wrong Pixel	1.4 %	4.0 %	1.7 %	5.6 %	2.3 %
Wrong SCT	1.2 %	1.6 %	1.4 %	3.0 %	2.2 %
Tail $1/p_T$	3.1 %	3.7 %	3.5 %	5.0 %	4.6 %
Tail d_0	1.9 %	3.6 %	2.2 %	5.1 %	2.9 %
u-rejection	66	36	62	25	45

While the *3+3* layout superficially offers a comparable performance to the TDR setup, it is considerably worse when there is a module/layer inefficiency in either the pixel or SCT detectors. Thus a 6 layer scenario including the B-layer does not appear to offer a sufficiently robust set-up for pattern recognition in dense jets.

7 Conclusions

For this event sample, the spoilt hit rate (2-track resolution) observed in the outer pixel layers is better than that in the SCT, whereas it is by far the worst in the B-layer. This reflects in the pattern recognition performance such that the removal of an outer pixel layer has a worse effect than the removal of an SCT layer. It follows that if one layer is to be removed purely on the basis of pattern recognition performance it would be the B-layer.

When the analysis cuts are adjusted to give a roughly constant track-finding efficiency for the different layouts, as layers are removed there is a faster than linear rise in the number of fake and spoilt tracks. Conversely, if the cuts are chosen to keep the fake-rate approximately constant

then there is a faster than linear reduction in efficiency with removed layers. Both cases demonstrate a greater sensitivity to the removal of a pixel layer.

Even with 3 pixel layers, the b-tagging performance is limited by wrong associations and badly resolved hits in the vertex region. Thus it is essential that 2 intermediate points are provided between the B-layer and the SCT layer at R = 30 cm.

The presence/absence of the TRT makes little difference to the b-tagging performance.

The overall performance would benefit from hardware and software developments. The pixel group are already aware that, although deliberately broadened clusters provide increased precision on isolated tracks, there is an unacceptable loss of 2-track resolution in jets. Software improvements in the reconstruction should include de-weighting pixel clusters assigned to more than one track and a more careful assignment of TRT drift-time hits for straws associated to more than one track. Of course such improvements can only be made if all the close tracks are actually found – the removal of layers causes a fall in track-finding efficiency for tracks produced with an angular separation below ~20 mrad.

The above study suggests that, in order to maintain robustness, given the harsh LHC environment and difficult access, and keep the physics contribution of the B-layer, it is advisable to have a total of at least 7 discrete layers. Further studies including degradation and pile-up are of course necessary, but can be expected to confirm this conclusion.

Acknowledgement

References

1   Atlas Collaboration CERN/LHCC/97-16 (Atlas TDR 4) Section 5.2

2   Atlas TDR 4, Section 6.7

3   R. Clifft and A. Poppleton, Atlas Internal Note, SOFT-NO-009

4   Atlas TDR 4, Section 2.5.2.1

5   Atlas TDR 4, Section 3.5.3

6   Atlas TDR 4, Section 3.1.5.1