

# The ATLAS inner detector and flavour tagging performance

R. Hawkings, on behalf of the ATLAS collaboration

CERN, EP division, Meyrin, CH 1211 Genève 23, Switzerland, e-mail: richard.hawkings@cern.ch

Received: 1 August 2003 / Accepted: 9 February 2004 /

Published Online: 13 July 2004 – © Springer-Verlag / Società Italiana di Fisica 2004

**Abstract.** The design of the ATLAS inner detector is described, paying particular attention to recent design changes and their impact on performance. The status of construction is briefly discussed, followed by calibration and alignment issues and a summary of current flavour tagging studies.

## 1 Introduction

Physics at the LHC imposes unprecedented requirements on tracking detectors, due to the high interaction rate, particle density and received radiation dose. At the LHC design luminosity, around 700 tracks are expected per bunch crossing within the detector acceptance, leading to very high detector occupancies. The time between bunch crossings is only 25 ns, requiring detectors and electronics with fast response and on-detector pipelined memories to store the data from each collision for several  $\mu\text{s}$  until a level-one trigger decision is made and the data read out or discarded. Inner detector elements will receive fluences of between  $10^{13}$  and  $10^{14}$  equivalent 1 MeV neutrons/cm<sup>2</sup>/year and are required to continue working for at least ten years in this environment. Finally, material in the tracking detectors must be minimised to avoid compromising the calorimeter performance for challenging electromagnetic channels such as  $H \rightarrow \gamma\gamma$ .

The ATLAS solution to this challenge is based on three different technologies—pixel detectors, silicon strips and a straw-tube tracker with transition radiation detection capacity. The layout of the ATLAS inner detector is shown in Fig. 1. The pixel detector is closest to the beamline (radius 5–25 cm), and consists of three barrel layers, and three endcap disks on each side of the interaction point. The pixel modules have a spatial resolution of 12  $\mu\text{m}$  in the  $r$ - $\phi$  plane, and 60  $\mu\text{m}$  in the  $r$ - $z$  plane. The semiconductor tracker (SCT) occupies the radial space between 25 and 50 cm, and consists of four barrel layers and nine disks in each endcap. Each layer is equipped with two sets of single-sided silicon strip detectors, one measuring in the  $r$ - $\phi$  plane and one with a 40 mrad stereo angle, leading to resolutions of 16  $\mu\text{m}$  in  $r$ - $\phi$  and 580  $\mu\text{m}$  in  $r$ - $z$ . Finally, the transition radiation tracker (TRT) forms the outermost tracking detector, and consists of many layers of 4 mm diameter straws. In the barrel, the straws run axially and are arranged in three layers of modules, whilst in the endcaps the straws run radially and are arranged in wheels of 8 or 16 layers of straws each. This detector gives up

to 36  $\phi$  coordinate measurements per track, each with a spatial resolution of around 170  $\mu\text{m}$ . The straws are interleaved with radiator material, leading to an electron identification capability through the detection of transition radiation photons.

Due to resource and schedule constraints, not all the detector elements will be installed initially. The middle pixel barrel layer and endcap disks, and the last four TRT wheels in each endcap (the ‘C’ wheels giving coverage beyond  $|\eta| = 1.7$ ), will be staged and installed at a later date ( $\eta$  is the pseudo-rapidity of the particle, defined as  $\eta = -\ln \tan(\theta/2)$  where  $\theta$  is the angle with respect to the beamline). The effect on the performance of the inner detector is discussed below.

## 2 Detector performance and recent design changes

The baseline reference for the ATLAS inner detector design and performance is the Physics TDR [1], with more detail being given in the inner detector and pixel TDRs [2, 3]. The detector will be installed inside the inner bore of the liquid argon calorimeter cryostat, forming a cylinder approximately 7 m long and 1.1 m in radius. The detector sits in a nominal 2 T magnetic field provided by a superconducting solenoid integrated into the calorimeter cryostat. However, the field is rather non-uniform due to the restricted length of the solenoid coil, dropping off to 1 T at the end of the tracker volume. This complicates the track finding and fitting at high  $|\eta|$ .

All three subdetectors provide coverage up to  $|\eta| = 2.5$  to maximise the physics acceptance. A major challenge is that of pattern recognition in the very dense LHC tracking environment, especially in the cores of jets and at high luminosity. The inner detector provides seven precision space points per track (three from the pixel detector and four from the SCT combined  $r$ - $\phi$  and stereo measurements), combined with up to 36 TRT straw hits. The TRT is particularly important for pattern recognition and

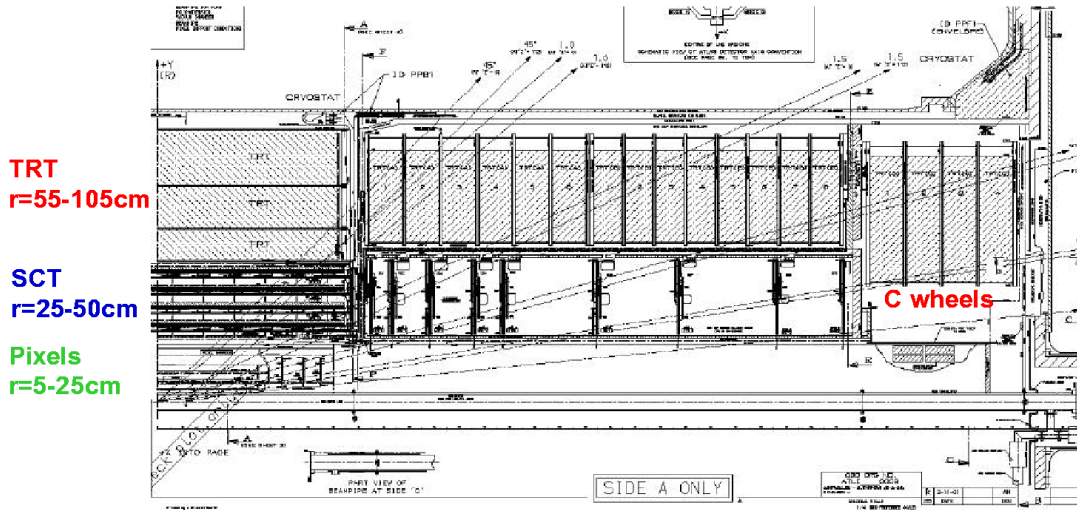


Fig. 1. Layout of the ATLAS inner detector, showing the pixel, SCT and TRT subdetectors

reconstructing long-lived neutral decays (*e.g.*  $K_s^0$  and  $\Lambda$ ) that occur far from the primary vertex. The combination of precision space points from the silicon detectors and ‘continuous’ tracking from the TRT provides very robust pattern recognition performance even inside jets—here the reconstruction efficiency for hadrons with transverse momentum  $p_T > 1$  GeV is 80–90 %, and the fraction of reconstructed tracks which are fake is only 0.2–0.3 %.

The TRT has been optimised for tracking performance rather than transition radiation detection, but still provides a very useful electron identification capability. A rejection factor of order 100 can be achieved against charged pions for 80 % electron efficiency, varying as a function of track momentum and  $|\eta|$ .


Since the time of the Physics TDR, the main changes have been to the pixel detector. ATLAS has adopted a ‘fully insertable’ layout that allows the complete pixel detector and beampipe package to be installed and removed (*e.g.* for upgrades) from the end without disturbing the rest of the inner detector. In particular, this means that all pixel services now run along the inner bore of the SCT at low radius. Material at low radius has the most significant effect on tracking performance, but these services are partly outside of the tracking  $|\eta|$  acceptance. The radius of the innermost pixel detector has also been increased from 4.3 to 5.0 cm to accommodate a larger diameter beampipe, and the material in the pixel modules has increased by a factor 1.5 after the final engineering design. Numerous other minor changes have been made as the designs of the subdetectors and their associated services have been completed. The material distribution in the inner detector is shown in Fig. 2, where it can be seen that the combined radiation length of all detectors is around  $0.3 X_0$  at  $\eta = 0$ , rising to around  $1 X_0$  at  $|\eta| = 2$ . It should be noted that the contribution of the inner detector services (shown separately in Fig. 2) occurs mainly at high radius and/or high  $z$ , where its effect on both the tracking and calorimetry performance is minimized.

The design changes result in some loss of detector performance with respect to the design presented in the Physics TDR. The most significant changes are to the single track transverse momentum and impact parameter resolution. Figure 3 shows these resolutions for the Physics TDR layout and the current layout implemented in the latest round of Monte Carlo simulations (known as data challenge DC1). The resolutions are parameterised as a function of  $p_T$  as the sum in quadrature of terms  $a$  and  $b/p_T$ . The parameter  $a$  gives the asymptotic resolution at infinite momentum, whilst  $b$  represents the dependence on multiple scattering at low momentum. The multiple scattering terms have increased by factors of 1.2 for the momentum and 1.6 for the transverse impact parameter, mainly as a result of the increased material and the change in radius of the first pixel layer. The effect of this degradation on the b-tagging performance is discussed in Sect. 5.2.

The effect of staging the middle pixel layers and TRT endcap ‘C’ wheels has also been studied. Removal of the middle pixel layer results in a significant loss in b-tagging performance—about 25 % more light quark background for the same b-tagging efficiency. Removal of the TRT ‘C’ wheels, which provide coverage in the region ( $1.7 < |\eta| < 2.5$ ) results in around a 50 % degradation in momentum resolution in this region, which is then only covered by the pixel and SCT detectors at lower radii. Some loss in pattern recognition capability and robustness will also occur, but it should be noted that this will happen during the initial low-luminosity running phase of LHC, where the track occupancy will be lower. This loss in performance typically results in about 15–20 % more integrated luminosity being needed for physics scenarios that lead to discoveries in the first phase of LHC running.

### 3 Construction status

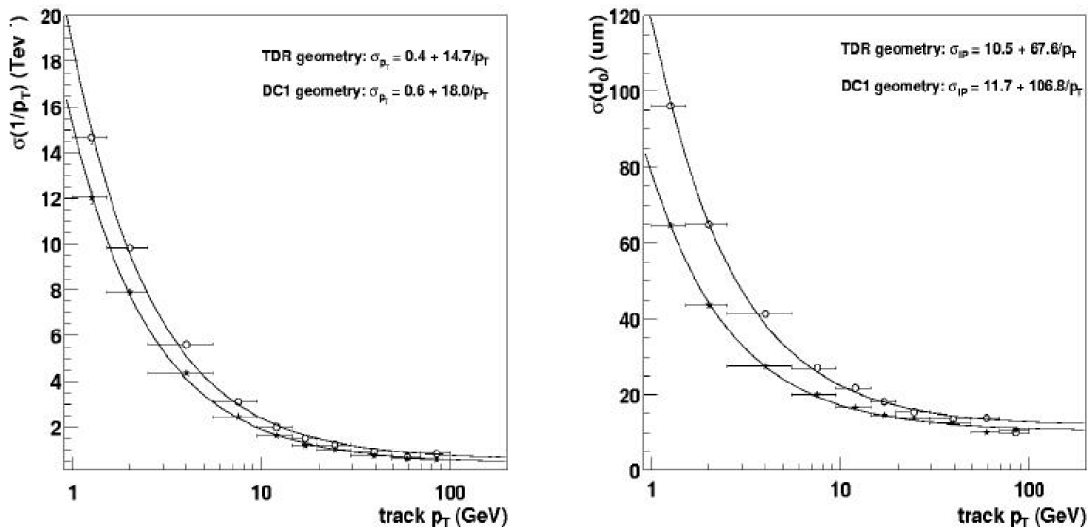
All of the inner detector subdetectors are now entering the major production stage, with prototyping of the main



**Fig. 2.** Estimated thickness of the inner detector material in radiation lengths  $X_0$ . The separate contributions from the beampipe, pixels, SCT, TRT and services external to the detector volumes are indicated

detector components largely completed. For the pixel detector, both the sensors themselves and the detector support structures are in production. The electronics, based on deep sub-micron (DSM) technology, is undergoing the last design iterations, and chips from ‘pilot’ production runs are being used to produce the first ‘final’ modules to gain time and experience. These modules are performing well, and tests with modules irradiated to ten years of LHC operation show good results, with for example very low noise probabilities of  $10^{-7}$  per pixel—corresponding to ten noise hits per event in the complete ATLAS pixel system. An intensive program of irradiation, beam and system tests (operating several modules and associated services together) is now being undertaken, together with the production of 10–15% of the final modules, which will be used for production debugging, development of quality assurance procedures and tests, and for equipping the first disks.

Most of the SCT sensors and front-end electronics have now been produced, and the focus is switching to module assembly—around 4000 are needed for the entire detector. Barrel modules are now being produced in several sites around the world. Early endcap modules had problems with electronics noise when several were operated together, but this has now been solved with a redesign of the module circuit board. The modified modules have been shown to work after irradiation, in test beams and system test setups, and preparations for full module production and assembly are now underway. Meanwhile, the assembly of the first barrel modules onto their carbon fibre support cylinder is beginning—this also requires the finalisation of structural and cooling system components. The first barrel has around 400 modules, and a major focus is now on preparations for large scale tests of the completed barrel towards the end of this year.



**Fig. 3.** Transverse momentum ( $1/p_T$ ) and impact parameter ( $d_0$ ) resolution integrated over  $\eta$  as a function of track  $p_T$  for the Physics TDR and DC1 (*latest simulation*) geometries, together with fitted parameterisations (*see text*)

All the TRT straws have now been produced, and the electronics is in production, albeit with some delays to the analogue part. The barrel straws have an electrically-insulating glass wire joint in the centre, allowing the two ends to be read out independently. Problems emerged with ageing of this joint after irradiation, which lead to a pause in module wire stringing whilst studies of alternative wire joint techniques were undertaken, together with a reconsideration of the TRT active gas mixture. After extensive tests, the baseline gas mixture has been changed to 70 % Xe, 27 %  $\text{CO}_2$  and 3 %  $\text{O}_2$ , and barrel module production has been restarted with the original wire joint technology. The new gas mixture does not significantly affect the performance of the TRT. A new concept of ‘cleaning runs’ has been introduced, to remove any silicon deposits which may build up on the wires over time. One day of running with around 10 %  $\text{CF}_4$  added to the active gas mixture is enough to remove these deposits, during which time the TRT would operate as a tracker but without transition radiation detection. This operation could be performed typically once per year. Several endcap wheel modules have already been produced and delivered to CERN, and the focus is now on continuing the series production of all wheels in Russia. In the next year, the barrel and endcap modules will start to be assembled to form the complete TRT sections at CERN.


A major surface integration facility is now being prepared at CERN to receive the various ID parts. This facility (the ‘SR building’) is located in the ATLAS surface building area above the cavern and will be used over the next few years to assemble and test the complete barrel and endcap inner detector parts (SCT/TRT) and the combined pixel/beampipe package before they are lowered into the ATLAS cavern for installation in the experiment. Significant infrastructure, including a class 100 000 clean room, cooling systems, data acquisition and control systems, is being prepared to enable the first inner detector

parts to be tested and assembled together once they start arriving in the second half of 2003.

## 4 Alignment and calibration

The ATLAS physics goals impose very stringent requirements on the alignment and calibration of the inner detector. A requirement that alignment uncertainties degrade the tracking resolution by no more than 20 % of its nominal value leads to target  $r$ - $\phi$  alignment precisions of  $7 \mu\text{m}$  for the pixels,  $12 \mu\text{m}$  for the SCT and  $30 \mu\text{m}$  for the TRT, which is adequate for most physics studies. But a precision measurement of the W boson mass to 15 MeV requires that the  $r$ - $\phi$  alignment of the silicon detectors be understood even more precisely, to the level of  $1 \mu\text{m}$ . Achieving these goals over the scale of several metres, with several tens of kilowatts of power being dissipated from the detectors and electronics, will be a major challenge. The physics goals also require that the magnetic field be known to a precision of 0.02 %, which will be achieved with a combination of initial field mapping, fitting the field shape and online monitoring with NMR probes installed inside the inner detector volume.

The alignment precision goals will be fulfilled with a multi-step strategy beginning with precision mechanics and optical survey at many stages of the construction and assembly process. As part of the tests in the SR building, the complete barrel and endcap SCT will undergo a precision X-ray scan which will locate all modules with a precision of  $10 \mu\text{m}$  in  $r$ - $\phi$  and  $50 \mu\text{m}$  in  $r$ - $z$ . The TRT may also participate in this survey procedure. The SCT also has an in-situ laser interferometer system with several hundred grid lines, that monitors the positions of various points on the SCT support structures and some modules to a three-dimensional precision of  $10 \mu\text{m}$  in real time. This



**Fig. 4.** Typical distributions of the track impact parameter significance (*left plot*) and jet weight (*right plot*), for jets containing u and b quarks

system will be operated during the X-ray survey on the surface, then again after installation and during running, allowing any distortions in the SCT support structures to be tracked and the X-ray survey to be corrected. Finally, alignment will be performed with tracks, exploiting both hadrons passing through the overlaps between adjacent detectors and high momentum muons from gauge boson and heavy flavour decays. The data taking rate is such that enough statistics are gathered to align the detector to a few  $\mu\text{m}$  precision in less than one day—the challenge will clearly be to understand the detector stability and associated systematics.

## 5 Flavour tagging

Flavour tagging is one of the most important tasks of the inner detector, and is an essential tool in many parts of the ATLAS physics program. Studies so far have concentrated on tagging b-jets, which are important *e.g.* for tagging top quark decays via the decay  $t \rightarrow bW$ , and for searches for any new particles coupling preferentially to heavy quarks, for example Higgs bosons. Therefore b-tagging is critical for many discovery physics channels, both at low and high luminosity.

Jets from b quarks can be tagged by exploiting the long lifetime and high mass of b hadrons, which lead to decay tracks with large impact parameters with respect to the primary vertex, and reconstructible secondary vertices. Semileptonic b decays can also be tagged by searching for soft electrons or muons in the hadronic calorimeter. Recent work has concentrated on evaluating the lifetime-based b-tagging performance with the latest inner detector layouts, and on improvements to the b-tagging algorithms and techniques. This is an ongoing process, and the results presented below show the state of the studies in May 2003.

### 5.1 Basic b-tagging algorithm

The b-tagging algorithm employed by ATLAS is based on track impact parameter significances  $S$ , defined as the distance of closest approach of a track to the primary vertex (or beamspot) position, divided by its error. The impact parameter significance is signed using the direction of the nearest jet—positive if the track crosses the jet axis in front of the primary vertex, and negative otherwise. As shown in Fig. 4, tracks from light quark jets tend to have small values of  $|S|$ , originating mainly from resolution effects and forming a symmetric distribution centred around zero. By contrast, tracks from b decays show a large tail towards positive values, corresponding to tracks with significant impact parameters.

These distributions are used to transform  $S$  into a ratio of probabilities that the track originated from a light quark jet or a b jet. The probabilities for all tracks in the jet (typically 5–10 with  $p_T > 1\text{ GeV}$  in b jets, with around half arising from decay products of the b hadron) are combined into a jet weight, also shown in Fig. 4, which forms the main discriminating variable for distinguishing between b and light quark jets.

This procedure can be performed both in the transverse and longitudinal planes. The information from the transverse plane is more powerful, as the track impact parameter resolution is superior, and the primary vertex position is constrained to the beamspot (formed by the overlap region of the colliding beams), which has a transverse width of only  $15\ \mu\text{m}$ . This is referred to as two-dimensional b-tagging. In the longitudinal ( $z$ ) direction, the beamspot has an RMS length of around 6 cm, so the  $z$  coordinate of the event primary vertex must first be determined (typically with a resolution of around  $30\ \mu\text{m}$ ) to make use of the longitudinal impact parameter information. At high luminosity, many separate interactions will

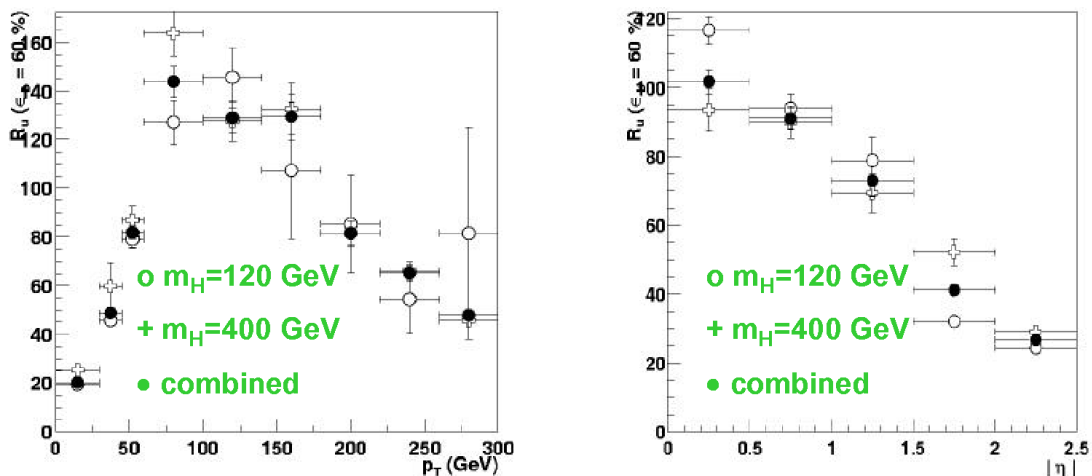


Fig. 5. Rejection against u-jets for 60% b-tagging efficiency in WH events, as a function of jet  $p_T$  (left) and  $|\eta|$  (right)

happen within the same bunch crossing, so the primary vertex associated to the high- $p_T$  interaction must also be selected. The jet weights from both planes can be combined to give a three-dimensional (or ‘2D+z’) b-tagging algorithm.

## 5.2 b-tagging performance

The performance of two- and three-dimensional b-tagging algorithms on different inner detector layouts has been studied using associated production  $pp \rightarrow WH + X$  events with Higgs masses of 120 and 400 GeV. Event samples have been simulated forcing the Higgs to decay to  $b\bar{b}$ ,  $c\bar{c}$  and light ( $u\bar{u}$ ) quark pairs. Although the latter decay is not expected, it provides a convenient source of light quark jets with the same kinematic properties as the corresponding b-jets, allowing the light quark rejection to be easily studied.

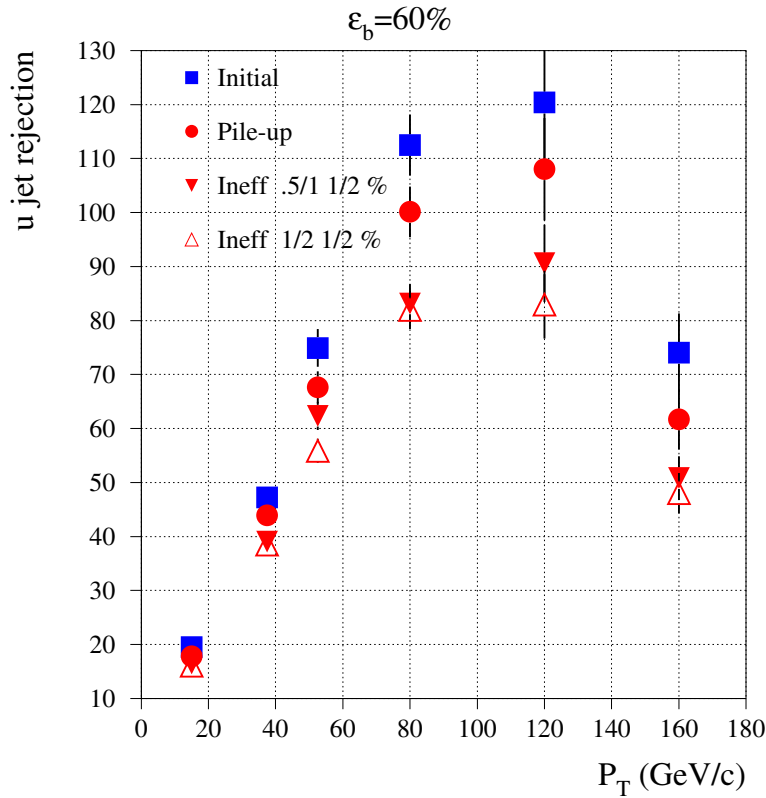
Comparisons of the performance of the latest (DC1) detector layout using two- and three-dimensional b-tagging algorithms with the ‘reference’ performance figures from the Physics TDR [1] with two-dimensional b-tagging are shown in Table 1. The performance is studied in terms of the rejection factors (*i.e.* the reciprocal of the efficiency) achieved against light quark and charm jets, for a constant 50% b-jet tagging efficiency, integrated over  $|\eta| < 2.5$ . Compared to the Physics TDR, a significant deterioration is seen in the rejection of light quark jets at low Higgs masses. These low  $p_T$  jets have many low momentum tracks, and are strongly affected by the increase in multiple scattering in the DC1 detector layout. At high Higgs masses, and for rejection against charm jets, the deterioration is less significant. The performance here is limited more by pattern recognition errors leading to misreconstructed tracks, and by the finite lifetime of charm hadrons. In all cases, moving to the three-dimensional algorithm results in a significant improvement, and some increased performance with respect to the Physics TDR.

The u-jet rejection factors of the three-dimensional algorithm are shown as a function of  $p_T$  and  $|\eta|$  in Fig. 5.

There is a significant dependence on  $p_T$ , with the best performance being achieved for jets of around 100 GeV. Below this, the performance falls off due to the increase in multiple scattering at low  $p_T$  and the fraction of tracks that fall below the 1 GeV reconstruction cutoff. Above 100 GeV, the jets become narrower and more collimated, and pattern recognition errors start to dominate. There is little intrinsic dependence on the boson mass, with both 120 GeV and 400 GeV Higgs decays leading to similar jet rejections at the same  $p_T$ . The dependence on  $|\eta|$  is also significant, with the rejection falling off in the forward detector region as the tracking performance degrades and the jet energy rises. For comparison, tagging algorithms using soft electrons or muons give typical rejection factors of around 300 against light quark jets and 40 against charm jets, for a b-tagging efficiencies of around 7% each. The efficiency of these tags is limited by the b semileptonic branching ratio of around 10% to each lepton type.

As well as the WH events discussed above,  $t\bar{t}H$  events have also been studied. These events are much more complicated, with additional jets from the top decays leading to a typical 20% reduction in light jet rejection due to merging of and overlaps between the heavy and light quark jets. At fixed  $p_T$ ,  $|\eta|$  and jet isolation, similar b-tagging efficiencies are observed for the b-jets from the Higgs and top decays.

These  $t\bar{t}H$  events have also been used to study the b-tagging performance with the initial detector layout with only two pixel layers. The light quark rejection is typically 25% worse than that for the three layer layout at the same b-tagging efficiency. The effect of various additional degradations is shown as a function of jet  $p_T$  in Fig. 6. The effect of ‘pile-up’—the additional 4.6 minimum bias events expected in every bunch crossing at an initial ‘low’ luminosity of  $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ —is rather small, since the performance is dominated by the track density inside the jet cores. Some additional loss in efficiency can be expected if the pixel modules are not fully efficient, but have a few % of dead channels—the effect of this is largest in the innermost pixel layer, and can be partially compen-



**Fig. 6.** Rejection against light quark jets for 60% b-tagging efficiency in  $t\bar{t}H$  events, for the initial detector layout, with low-luminosity pile-up events added, and with additional detector inefficiencies. The inefficiencies simulated are 0.5% per detector module and 1% per readout chip in the innermost pixel layer, and 1% and 2% in the other layers (denoted ‘0.5/1 1/2%’); and 1% per module and 2% per readout chip in all pixel layers (denoted ‘1/2 1/2%’)

**Table 1.** Rejection power against light  $u$  and charm  $c$  jets for a constant 50% b-tagging efficiency and Higgs masses of 120 and 400, GeV. Results are given for the Physics TDR and DC1 detector layouts, using two- and three-dimensional b-tagging

u/c jet rejection	PhysTDR (2D)	DC1 (2D)	DC1 (3D)
$R_u(120 \text{ GeV})$	$359 \pm 23$	$189 \pm 8$	$321 \pm 14$
$R_u(400 \text{ GeV})$	$141 \pm 9$	$144 \pm 8$	$195 \pm 9$
$R_c(120 \text{ GeV})$	$11.5 \pm 0.4$	$9.8 \pm 0.1$	$12.3 \pm 0.2$
$R_c(400 \text{ GeV})$	$13.6 \pm 0.5$	$10.6 \pm 0.2$	$12.6 \pm 0.3$

sated for by choosing the ‘best’ detector modules for this layer.

Another area of current study concerns the calibration of the b-tagging performance. Here,  $t\bar{t} \rightarrow WbW\bar{b} \rightarrow q\bar{q}b\ell\nu\bar{b}$  events can be used to calibrate the b-tagging efficiency and light quark rejection from the data. The decay of one top to  $\ell\nu\bar{b}$  can be used to tag the decay (using the high  $p_T$  lepton and b-tag on one jet), leaving an unbiased sample of  $q\bar{q}b$  three-jet systems, where the b-jet can be identified via a kinematic fit to determine the top mass. This allows large samples of both b and light quark jets to be isolated, providing an invaluable environment to study the b-tagging algorithm performance.

### 5.3 Algorithm enhancements

The impact parameter based method does not use all available information, as it simply looks for a number of

tracks inconsistent with coming from the primary vertex. However, in a genuine b-jet, the b hadron decay tracks should all come from a common secondary vertex (neglecting the flight distance between the b and subsequent charm hadron decays), and this decaying b hadron should have high mass. Exploiting these properties was crucial in achieving the maximum b- (and charm-) tagging performance at LEP and SLC, and first studies have started in ATLAS along similar lines.

Explicit secondary vertex reconstruction has been studied using a Kalman filter algorithm, identifying the primary and all two-track secondary vertices and then merging the latter to make a single b hadron vertex. Using this technique, a secondary vertex can be reconstructed in 70% of b and only 8% of u jets. Two discriminating variables are formed—the mass of the secondary vertex, and the fraction of the jet energy asso-

ciated to the secondary vertex. These are then combined in a likelihood-based method with the jet weights from the impact parameter-based tag. This already results in a factor 1.5–2 improvement in light quark jet rejection over the three-dimensional method based on impact parameters alone, and further improvements are expected as the studies continue.

#### 5.4 $\tau$ -tagging

The tagging of  $\tau$ -jets is useful in many physics processes, for example the decay of supersymmetric Higgs particles  $A/H \rightarrow \tau\tau$ , usually studied in the ‘mixed’ decay mode where one  $\tau$  decays hadronically and the other leptonically. The main selection tools are based on lepton identification and large missing transverse energy, but secondary vertex based selections can contribute some additional rejection power.

Selecting  $\tau$  decays based on vertex properties is very challenging, since a highly boosted  $\tau$  typically decays in a very narrow cone, with only one or three charged particles. This leads to small impact parameters and low multiplicity secondary vertices with a very small opening angle. The resolution attainable on such vertices is typically  $50\ \mu\text{m}$  orthogonal to the  $\tau$  flight direction and 2 mm along the flight direction. Requiring a three-track  $\tau$  decay vertex separated from the beamspot by at least 4 mm in the transverse plane results in an efficiency of 44% for three-prong  $\tau$ s and a rejection against light quark jets in excess of  $10^3$ . This can be applied in the  $A/H \rightarrow \tau\tau$  analysis to give a 10% enrichment of the signal sample with the same level of background.

## 6 Conclusions

The design of the ATLAS inner detector is now mature, with only the last phases of detailed engineering design,

mainly for services, off-detector components and integration, still ongoing. The full detector is expected to be installed ready for the start of LHC operation in 2007, except for the middle pixel layers and last four TRT wheels in each endcap, which will be staged for resource and schedule reasons. This initial detector is capable of meeting the physics needs of the initial low luminosity running with only a modest loss in performance, but the full detector is still needed for high luminosity LHC operation.

Most detector components are now in production, and assembly is beginning. As is inevitable in such a large and complex system, some problems have been encountered, but solutions are being found and production and assembly is proceeding largely on schedule. The focus will shift in the next years from production to assembly, integration and commissioning, and preparations for these phases are already beginning.

In parallel, work continues on software, alignment, calibration and performance studies, to follow changes in the engineering design and evolution of the ATLAS software. Despite the recent design changes, the performance of the inner detector is being maintained to meet the ATLAS physics goals, showing that the design is robust. New ideas in *e.g.* b-tagging algorithms continue to be explored, and the ATLAS inner detector community looks forward to beginning to exploit the detector for physics in 2007.

## References

1. ATLAS collaboration: detector and physics performance technical design report, CERN/LHCC/99-14, May 1999
2. ATLAS collaboration: inner detector technical design report, CERN/LHCC/97-16 and 17, April 1997
3. ATLAS collaboration: pixel detector technical design report, CERN/LHCC/98-13, May 1998