

Chapter 4

Inner detector

4.1 Introduction

The ATLAS Inner Detector (ID) is designed to provide hermetic and robust pattern recognition, excellent momentum resolution and both primary and secondary vertex measurements [60, 61] for charged tracks above a given p_T threshold (nominally 0.5 GeV, but as low as 0.1 GeV in some ongoing studies of initial measurements with minimum-bias events) and within the pseudorapidity range $|\eta| < 2.5$. It also provides electron identification over $|\eta| < 2.0$ and a wide range of energies (between 0.5 GeV and 150 GeV). This performance, which is required even at the highest luminosities expected from LHC collisions, is consequently at the limit of existing technology.

The ID layout, as shown in figure 4.1, reflects the performance requirements. The ID is contained within a cylindrical envelope of length ± 3512 mm and of radius 1150 mm, within a solenoidal magnetic field of 2 T (see section 2.2.4). Figures 4.2 and 4.3 show the sensors and structural elements traversed by 10 GeV tracks in respectively the barrel and end-cap regions.

The ID consists of three independent but complementary sub-detectors. The envelopes of each sub-detector are listed in table 1.2 (see section 1.2) and shown in figure 4.1. At inner radii, high-resolution pattern recognition capabilities are available using discrete space-points from silicon pixel layers and stereo pairs of silicon microstrip (SCT) layers. At larger radii, the transition radiation tracker (TRT) comprises many layers of gaseous straw tube elements interleaved with transition radiation material. With an average of 36 hits per track, it provides continuous tracking to enhance the pattern recognition and improve the momentum resolution over $|\eta| < 2.0$ and electron identification complementary to that of the calorimeter over a wide range of energies.

The high-radiation environment imposes stringent conditions on the inner-detector sensors, on-detector electronics, mechanical structure and services. Over the ten-year design lifetime of the experiment, the pixel inner vertexing layer must be replaced after approximately three years of operation at design luminosity. The other pixel layers and the pixel disks must withstand a 1 MeV neutron equivalent fluence F_{neq} , defined in section 3.3.1 (see also [52]), of up to $\sim 8 \times 10^{14} \text{cm}^{-2}$. The innermost parts of the SCT must withstand F_{neq} of up to $2 \times 10^{14} \text{cm}^{-2}$. To maintain an adequate noise performance after radiation damage, the silicon sensors must be kept at low temperature (approximately -5 to -10°C) implying coolant temperatures of $\sim -25^\circ\text{C}$. In contrast, the TRT is designed to operate at room temperature.




Figure 4.1: Plan view of a quarter-section of the ATLAS inner detector showing each of the major detector elements with its active dimensions and envelopes. The labels PP1, PPB1 and PPF1 indicate the patch-panels for the ID services.

The above operating specifications imply requirements on the alignment precision which are summarised in table 4.1 and which serve as stringent upper limits on the silicon-module build precision, the TRT straw-tube position, and the measured module placement accuracy and stability. This leads to:

- (a) a good build accuracy with radiation-tolerant materials having adequate detector stability and well understood position reproducibility following repeated cycling between temperatures of -20°C and $+20^{\circ}\text{C}$, and a temperature uniformity on the structure and module mechanics which minimises thermal distortions;
- (b) an ability to monitor the position of the detector elements using charged tracks and, for the SCT, laser interferometric monitoring [62];
- (c) a trade-off between the low material budget needed for optimal performance and the significant material budget resulting from a stable mechanical structure with the services of a highly granular detector.

The inner-detector performance requirements imply the need for a stability between alignment periods which is high compared with the alignment precision. Quantitatively, the track precision should not deteriorate by more than 20% between alignment periods.




Figure 4.2: Drawing showing the sensors and structural elements traversed by a charged track of 10 GeV p_T in the barrel inner detector ($\eta = 0.3$). The track traverses successively the beryllium beam-pipe, the three cylindrical silicon-pixel layers with individual sensor elements of $50 \times 400 \mu\text{m}^2$, the four cylindrical double layers (one axial and one with a stereo angle of 40 mrad) of barrel silicon-microstrip sensors (SCT) of pitch $80 \mu\text{m}$, and approximately 36 axial straws of 4 mm diameter contained in the barrel transition-radiation tracker modules within their support structure.

This chapter describes the construction and early performance of the as-built inner detector. In section 4.2, the basic detector sensor elements are described. Section 4.3 describes the detector modules. Section 4.4 details the readout electronics of each sub-detector, section 4.5 describes the detector power and control and section 4.6 describes the ID grounding and shielding. Section 4.7 discusses the mechanical structure for each sub-detector, as well as the integration of the detectors and their cooling and electrical services. The overall ID environmental conditions and general services are briefly summarised in section 4.8. Finally, section 4.9 indicates some initial results on the operational performance and section 4.10 catalogues the material budget of the ID, which is significantly larger than that of previous large-scale tracking detectors.




Figure 4.3: Drawing showing the sensors and structural elements traversed by two charged tracks of 10 GeV p_T in the end-cap inner detector ($\eta = 1.4$ and 2.2). The end-cap track at $\eta = 1.4$ traverses successively the beryllium beam-pipe, the three cylindrical silicon-pixel layers with individual sensor elements of $50 \times 400 \mu\text{m}^2$, four of the disks with double layers (one radial and one with a stereo angle of 40 mrad) of end-cap silicon-microstrip sensors (SCT) of pitch $\sim 80 \mu\text{m}$, and approximately 40 straws of 4 mm diameter contained in the end-cap transition radiation tracker wheels. In contrast, the end-cap track at $\eta = 2.2$ traverses successively the beryllium beam-pipe, only the first of the cylindrical silicon-pixel layers, two end-cap pixel disks and the last four disks of the end-cap SCT. The coverage of the end-cap TRT does not extend beyond $|\eta| = 2$.

4.2 Inner-detector sensors

This section describes the detector sensors of the pixel, SCT and TRT sub-systems - silicon pixel and micro-strip sensors in section 4.2.1, and straw tubes filled with a Xe/CO₂/O₂ gas mixture in section 4.2.2. As discussed in section 3.3, the detector sensors are subject to large integrated radiation doses. They have therefore been developed and controlled to withstand the expected irradiation, with a safety factor of approximately two.

4.2.1 Pixel and SCT detector sensors

The pixel and SCT sensors [63, 64] are required to maintain adequate signal performance over the detector lifetime at design luminosity (with the exception of the pixel vertexing layer, as discussed above). The integrated radiation dose has important consequences for the sensors of both detectors. In particular the required operating voltage, determined by the effective doping concentration, depends on both the irradiation and the subsequent temperature-sensitive annealing. The sensor leakage current also increases linearly with the integrated radiation dose. The n-type bulk material effectively becomes p-type after a fluence F_{neq} of $\sim 2 \times 10^{13} \text{ cm}^{-2}$. The effective doping concentration then grows with time in a temperature-dependent way. To contain this annealing and to reduce the leakage current, the sensors will, as noted above, be operated in the temperature range -5°C to -10°C . The sensors must further meet significant geometrical constraints on their thickness, granularity and charge-collection efficiency.

Table 4.1: Intrinsic measurement accuracies and mechanical alignment tolerances for the inner-detector sub-systems, as defined by the performance requirements of the ATLAS experiment. The numbers in the table correspond to the single-module accuracy for the pixels, to the effective single-module accuracy for the SCT and to the drift-time accuracy of a single straw for the TRT.

Item	Intrinsic accuracy (μm)	Alignment tolerances (μm)		
		Radial (R)	Axial (z)	Azimuth (R- ϕ)
Pixel				
Layer-0	10 (R- ϕ) 115 (z)	10	20	7
Layer-1 and -2	10 (R- ϕ) 115 (z)	20	20	7
Disks	10 (R- ϕ) 115 (R)	20	100	7
SCT				
Barrel	17 (R- ϕ) 580 (z) ¹	100	50	12
Disks	17 (R- ϕ) 580 (R) ¹	50	200	12
TRT	130			30 ²
¹ Arises from the 40 mrad stereo angle between back-to-back sensors on the SCT modules with axial (barrel) or radial (end-cap) alignment of one side of the structure. The result is pitch-dependent for end-cap SCT modules. ² The quoted alignment accuracy is related to the TRT drift-time accuracy.				

The pixel sensors required the most leading-edge and novel technology to meet the very stringent specifications on radiation hardness, resolution and occupancy in the innermost layers. The sensors are $250\,\mu\text{m}$ thick detectors, using oxygenated n-type wafers with readout pixels on the n^+ -implanted side of the detector. Despite its higher cost and complexity, this novel design involving double-sided processing was used because:

- (a) the n^+ implants allow the detector to operate with good charge-collection efficiency after type inversion, even when operated below the depletion voltage, because the depletion zone grows from the pixel side;
- (b) highly oxygenated material has been shown to give increased radiation tolerance to charged hadrons, with improved charge collection after type inversion and lower depletion voltage.

All of the 1744 pixel sensors (external dimensions $19\times 63\,\text{mm}^2$) are identical. The sensors will initially operate at $\sim 150\,\text{V}$ bias voltage, but operating voltages of up to $600\,\text{V}$ will be required for good charge collection efficiency after ten years of operation, depending on the sensor position, the integrated luminosity and the length of warm-up periods. The nominal pixel size is $50\times 400\,\mu\text{m}^2$ (about 90% of the pixels) and is dictated by the readout pitch of the front-end electronics. The size of the remaining pixels is $50\times 600\,\mu\text{m}^2$ in the regions at the front-end chips on a module. There are 47232 pixels on each sensor, but for reasons of space there are four ganged pixels in each column of the front-end chip, thus leading to a total of 46080 readout channels. A common bias

Table 4.2: External cut dimensions of the SCT barrel and end-cap (EC) sensors. The tolerance on all external dimensions is $\pm 25 \mu\text{m}$; the mask accuracy is at the level of $\pm 1 \mu\text{m}$. The inter-strip angle is that between adjacent strips of the sensor. The sensors are fabricated from 4-inch wafers.

Module type	Sensor type	Cut length (mm)	Outer width (mm)	Inner width (mm)	Strip pitch (μm)	Inter-strip angle (μrad)
Barrel	Barrel	63.960	63.560	63.560	80.0	0
End-cap inner	W12	61.060	55.488	45.735	56.9–69.2	207.0
End-cap middle	W21	65.085	66.130	55.734	69.9–83.0	207.0
	W22	54.435	74.847	66.152	83.4–94.2	207.0
End-cap outer	W31	65.540	64.635	56.475	70.9–81.1	161.5
	W32	57.515	71.814	64.653	81.5–90.4	161.5

grid ensures a DC connection to each pixel implant during wafer testing of sensor tiles, but it is only AC-coupled in production modules due to concerns about the effects of high radiation doses. To guarantee optimal post-irradiation performance and to maintain a small feature size, a p-spray isolation technology has been used [65]. Each pixel of a sensor is bump-bonded through a hole in the passivation layer to an element of the front-end readout integrated circuit as part of the module.

For reasons of cost and reliability, the 15912 sensors of the SCT use a classic single-sided p-in-n technology with AC-coupled readout strips. The sensors will initially operate at $\sim 150 \text{ V}$ bias voltage, but operating voltages of between 250 and 350 V will be required for good charge collection efficiency after ten years of operation, depending on the sensor position, the integrated luminosity and the length of warm-up periods. The sensor thickness of $285 \pm 15 \mu\text{m}$ is a compromise between the required operating voltage, the primary signal ionisation and the simplicity of fabrication. The strip pitch was determined by the required digitising precision, granularity, particle occupancy and noise performance. A strip pitch of $80 \mu\text{m}$ with two 6 cm-long sensors daisy-chained was chosen for the rectangular barrel sensors and radial strips of constant azimuth with mean pitch of $\sim 80 \mu\text{m}$ were chosen for the trapezoidal end-cap sensors. There are thus a total of 768 active strips of 12 cm length per sensor, plus two strips at bias potential to define the sensor edge. The detector dimensions are summarised in table 4.2 (see also table 4.7). Using the binary readout electronics described in section 4.4, a noise occupancy per channel of $< 5 \times 10^{-4}$ for a threshold of 1 fC is specified for fully irradiated modules. The detectors were required to operate stably at 500 V bias and to have $< 1\%$ of bad readout strips at 350 V bias before and after irradiation to $3 \times 10^{14} \text{ 24 GeV protons/cm}^2$, equivalent to the damage expected for a fluence F_{neq} of $2 \times 10^{14} \text{ cm}^{-2}$.

The SCT sensors were fabricated by two suppliers. For one supplier (7.8% of sensors), the p-spray isolation technology causes difficulties of micro-discharges at low humidity. All sensors were carefully tested, and a sub-sample was subjected to extended pixel or strip-by-strip studies [63, 64]. A sub-sample was also used for extensive post-irradiation performance studies [63, 66–68]. Apart from precautions related to the humidity sensitivity noted above, the rejected strip sensors were at the level of 1%. The pixel sensor rejection rate was somewhat higher because of a requirement to carefully control the profile and doping concentration of the p-spray isolation.

4.2.2 TRT straw tubes

Polyimide drift (straw) tubes of 4 mm diameter are the basic TRT detector elements [69]. The straw tube wall, especially developed to have good electrical and mechanical properties with minimal wall thickness, is made of two 35 μm thick multi-layer films bonded back-to-back. The bare material, a 25 μm thick polyimide film, is coated on one side with a 0.2 μm Al layer which is protected by a 5–6 μm thick graphite-polyimide layer. The other side of the film is coated by a 5 μm polyurethane layer used to heat-seal the two films back-to-back. Mechanically, the straws are stabilised using carbon fibres. After fabrication, the straws were cut to length (144 cm for the barrel and 37 cm for the end-caps) and leak-tested at 1 bar over-pressure. The straw (cathode) resistance was required to be $< 300 \Omega/\text{m}$.

For both the barrel and end-cap straws, the anodes are 31 μm diameter tungsten (99.95%) wires plated with 0.5–0.7 μm gold, supported at the straw end by an end-plug. They are directly connected to the front-end electronics (see section 4.4) and kept at ground potential. The anode resistance is approximately 60 Ω/m and the assembled straw capacitance is $< 10 \text{ pF}$. The signal attenuation length is $\sim 4 \text{ m}$ and the signal propagation time is $\sim 4 \text{ ns/m}$. The cathodes are operated typically at -1530 V to give a gain of 2.5×10^4 for the chosen gas mixture of 70% Xe, 27% CO_2 and 3% O_2 with 5 – 10 mbar over-pressure. A mixture containing 70% Ar and 30% CO_2 has been used during quality-control and cosmic-ray studies, resulting in an increased electron mobility and a degraded electron identification performance. Under normal operating conditions, the maximum electron collection time is $\sim 48 \text{ ns}$ and the operational drift-time accuracy is $\sim 130 \mu\text{m}$ [69]. Low-energy transition radiation (TR) photons are absorbed in the Xe-based gas mixture, and yield much larger signal amplitudes than minimum-ionising charged particles. The distinction between TR and tracking signals is obtained on a straw-by-straw basis using separate low and high thresholds in the front-end electronics (see section 4.4.1.3).

For the barrel straws, the anode wires (active length $\pm 71.2 \text{ cm}$) are read out from each end. Near their centre, the wires are supported mechanically by a plastic insert glued to the inner wall of the straw and split electrically by a fused glass capillary of 6 mm length and 0.254 mm diameter to reduce the occupancy. Each long barrel straw is therefore inefficient near its centre over a length of 2 cm. In the inner nine layers of type-1 barrel modules (see section 4.3.3), the wires are subdivided into three segments keeping only the 31.2 cm-long end-segments on each side active.

To guarantee stable operation, the wire offset with respect to the straw centre is required to be less than 300 μm . This is essentially a requirement on straw straightness since the wire sag is $< 15 \mu\text{m}$. To maintain straw straightness in the barrel, alignment planes made of polyimide with a matrix of 4.3 mm diameter holes are positioned each 25 cm along the module. Wires with offsets larger than 400 μm (amounting to $\sim 0.1\%$ of all wires) have been disconnected in the final barrel and end-cap acceptance tests. The nominal wire tension is 70 g and the tension of each wire has been controlled several times during the assembly and integration process. Wires with a tension in the range 47 g to 100 g were accepted in the final acceptance tests provided their tension had not decreased by more than 5 g since the time of production [70, 71].

The stable operation of TRT straws with the Xe-based gas mixture requires a re-circulating gas system with continuous monitoring of the gas quality. To avoid pollution from permeation through the straw walls or through leaks, the straws are operated in an envelope of CO_2 .

Table 4.3: Parameters of the pixel detector. The numbers of pixels, modules, staves/sectors are shown for each of the barrel layers and end-cap disks. The pixel envelope is shown in table 1.2. The active barrel length is 801 mm and the inner and outer active radii of each disk are respectively 88.8 and 149.6 mm. The quoted barrel radii are average values since the barrel staves are tilted at 20° with respect to a tangent vector at the given radius. The disk z -positions are also average values.

Barrel	Radius (mm)	Staves	Modules	Pixels
Layer-0	50.5	22	286	13.2×10^6
Layer-1	88.5	38	494	22.8×10^6
Layer-2	122.5	52	676	31.2×10^6
End-cap (one side)	z (mm)	Sectors	Modules	Pixels
Disk 1	495	8	48	2.2×10^6
Disk 2	580	8	48	2.2×10^6
Disk 3	650	8	48	2.2×10^6
Barrel and both end-caps			1744	80.4×10^6

At LHC rates, significant heat is generated in the straws by the ionisation current in the gas. The heat dissipation is proportional to the single straw counting rate and is estimated to be 10 mW to 20 mW per straw for the highest-occupancy straws (inner barrel layers) at the LHC design luminosity. To preserve gas-gain uniformity, the temperature gradient along each straw is required to be $< 10^\circ\text{C}$. The heat is evacuated differently for the barrel modules and end-cap wheels (see section 4.3.3).

At the LHC design luminosity, the straw counting rate will reach 20 MHz in the most critical detector regions, and the ionisation current density will reach $0.15 \mu\text{A}$ per cm of anode wire. The total accumulated charge after ten years of operation will reach $\sim 10 \text{ C/cm}$ in the most exposed straws. Many studies, including direct ageing tests lasting thousands of hours, have demonstrated stable straw-tube operation over their full operational lifetime [72–74]. Minute levels of pollution cannot be excluded and organo-silicone impurities, for which the relative concentrations must be kept below 10^{-11} , are particularly harmful. A gas filter has been incorporated in the TRT gas system for this reason. This filter is also effective in removing the ozone produced during gas amplification. In the case where ageing would significantly affect the detector performance, the use of a $\text{Ar}/\text{CO}_2/\text{CF}_4$ gas mixture during a few days of normal LHC operation has been shown to clean Si-based deposits from the anode wire.

4.3 Inner-detector modules

4.3.1 Pixel modules and staves

There are 1744 modules in the pixel detector [63]. As shown in figure 4.1, the pixel modules are arranged in three barrel layers (called ID layers 0-2) and two end-caps each with three disk layers. The basic detector parameters are listed in table 4.3. A total of 112 barrel staves and 48 end-cap sectors (8 sectors per disk) form the barrel and disk layers. Details of the pixel mechanical structure are described in section 4.7.




Figure 4.4: Schematic view of a barrel pixel module (top) illustrating the major pixel hybrid and sensor elements, including the MCC (module-control chip), the front-end (FE) chips, the NTC thermistors, the high-voltage (HV) elements and the Type0 signal connector. Also shown (middle) is a plan view showing the bump-bonding of the silicon pixel sensors to the polyimide electronics substrate. The photograph at the bottom shows a barrel pixel module.

A schematic view and photograph of a pixel module are shown in figure 4.4. A pixel module consists of a stack, from the bottom up, of the following components:

- (a) 16 front-end electronics chips thinned to $180\ \mu\text{m}$ thickness, each with 2880 electronics channels;
- (b) bump bonds (In or PbSn), which connect the electronics channels to pixel sensor elements;
- (c) the sensor tile of area $63.4 \times 24.4\ \text{mm}^2$ and approximately $250\ \mu\text{m}$ thick;
- (d) a flexible polyimide printed-circuit board (flex-hybrid) with a module-control chip glued to the flex-hybrid;

Table 4.4: Yield of pixel modules after bump-bonding of the electronics channels to the pixels. A sub-set of the highest quality modules were selected for the layer-0. Two materials for bump-bonding were used: indium (In) and lead-tin solder (PbSn) balls. At all stages of the fabrication, the use of PbSn gave significantly better yields.

Module yield	Number	Percentage
Module starts	2625	
Modules accepted after bump-bonding	2418	
Modules accepted after flex-hybrid glued	2312	100
Modules accepted for staves	2100	90.8
Module sub-set acceptable for layer-0	726	31.4

- (e) a polyimide pig-tail with Cu lines and a connector (barrel modules) or a wire micro-cable (end-cap modules) bonded to the flex-hybrid.

In all cases for the barrel, the top of the module faces the beam-pipe, whereas the end-cap disks have modules facing both ways. This means that the pixel sensors see in general the minimum amount of material in front of them, i.e. basically only the flex hybrid.

The sensors and electronics chips are connected by bump-bonding technology to form bare modules. Both solder (PbSn) and indium (In) bump-bonding technologies have been used to make pixel modules. Including reworked modules, the production statistics of bare modules are summarised in table 4.4. Flex hybrids with attached and tested module-control chips are glued to accepted bare modules. In total, 2312 modules were available for final electrical and mechanical characterisation.

After construction, the pixel modules were tested electrically at room temperature and at the approximate operating temperature of -10°C . Thermal cycling was performed on each module prior to completion of electrical testing. A ranking was made to separate the modules into those acceptable for the layer-0 radius (highest quality) and for the outer barrel and disk regions. Modules with the best ranking were loaded on barrel staves and disk sectors, leaving modules with the poorest ranking as spares. A brief summary of the production statistics and ranking after testing is given in table 4.4.

In the barrel region, 13 pixel modules are mounted on each stave using robotic tools and then glued (see figure 4.5). The staves are themselves mounted on carbon-fibre structures (see section 4.7.1). A bare stave consists of machined plates made of high-stiffness, thermally conducting, carbon-carbon (C-C) laminate material, an aluminium cooling tube and a carbon-fibre composite piece that captures the aluminium tube and is glued to the C-C pieces. The C-C pieces are precisely machined in a step pattern and are one-half the length of a stave. They are joined in the middle during the stave fabrication. A custom extrusion was used to make the aluminium tube with a flat surface at the interface with the C-C material. A custom fitting is welded to the end of the aluminium pipe. A thermal compound is used to conduct heat between the C-C plate and the tube. The tube is held in place by a carbon-fibre piece glued to the C-C plate. Because of corrosion, the cooling tubes associated with initial loaded staves were reworked through the insertion of a smaller



Figure 4.5: Close up of a bi-stave loaded with modules. The insert shows the U-link cooling connection between staves.




Figure 4.6: Pixel disk sector during the attachment of modules. There are also three modules on the back of the sector.

cooling tube with some deterioration of the thermal performance. The cooling tubes were replaced on bare staves. The electrical and thermal performance of each stave was measured after assembly. As illustrated in figure 4.5, the staves are joined to form bi-staves, which form the cooling unit in the barrel region. A custom-welded aluminium U-link is attached to each bi-stave. Custom low-mass cables are connected to each stave via a connector on each module and attached to the back of the stave to reach the ends of the stave.

The end-cap equivalent of the stave is a sector. The two pixel end-caps each have three identical disks. Each disk is composed of eight sectors. Six pixel modules are directly mounted on each sector, as shown in figure 4.6. The sectors are composed of thin, C-C faceplates with a rectangular aluminium cooling tube and vitreous carbon foam between the faceplates. The cooling tube is bent into a W-like shape to fit within the sector and makes contact with the faceplates with a compliant, thermally conducting adhesive. Each cooling circuit in the disk region serves two sectors.

Tested end-cap modules are positioned and glued on each sector with a precision of $1\text{--}2\text{ }\mu\text{m}$ in the plane of the module and about $10\text{ }\mu\text{m}$ perpendicular to the module plane. The module locations on each sector were optically surveyed using fiducial marks in the corners of the pixel sensors and other information. The survey precision with respect to the mounting bushings is estimated to be better than $5\text{ }\mu\text{m}$ in the plane of the module and about $15\text{ }\mu\text{m}$ in the direction perpendicular to this plane.

The spatial resolution of individual pixel modules has been measured in a test beam, for both non-irradiated and fully irradiated modules [63, 75]. At normal incidence, a spatial resolution of $12\text{ }\mu\text{m}$ is measured and approximately 80% of the tracks have a single pixel hit. The resolution is not significantly degraded after irradiation. The optimal resolution of $4.7\text{ }\mu\text{m}$ (before) and $6.0\text{ }\mu\text{m}$ (after) irradiation is obtained for incident angles of $10\text{--}15^\circ$. The pixel barrel staves overlap and are mounted at a tilt angle (the tangent to the support cylinder surface in the plane perpendicular to the cylinder axis) of -20° , defined by geometrical constraints. The Lorentz angle, however, for a magnetic field of 2T, varies between 12° (before) and 6° (after) irradiation [63, 75].

Table 4.5: SCT barrel cylinder layer parameters and number of modules per layer. There are 12 modules per row. The quoted radii and length are those of the outer surface of the support cylinder. The average active sensor radii and overall length are shown in brackets. The tilt angle is with respect to the tangent to the support cylinder surface in the plane perpendicular to the cylinder axis.

Barrel cylinder layer	Radius (mm)	Full length (mm)	Module tilt angle (degrees)	Number of modules
3	284 (299)	1530 (1498)	11.00	384
4	355 (371)		11.00	480
5	427 (443)		11.25	576
6	498 (514)		11.25	672
Total				2112

Table 4.6: The nominal z -position of the centre of each SCT end-cap disk and the number of modules on each disk (the total number of modules is 1976, summed over both end-caps). For geometrical acceptance reasons, disk 9 has only outer modules, while disks 1, 7 and 8 have no inner modules. The middle modules of disk 8 have only one sensor, again for geometrical acceptance reasons. For disks 1 to 8, the inner and outer module centres are displaced towards the interaction point by an average distance of 15.25 mm, while the middle modules are displaced away from the interaction point by 15.25 mm. For disk 9, the modules centres are displaced by 15.25 mm away from the interaction point. The modules are tiled in ϕ by ± 2.75 mm about these nominal centres.

Disk	1	2	3	4	5	6	7	8	9
$ z $ (mm)	853.8	934.0	1091.5	1299.9	1399.7	1771.4	2115.2	2505.0	2720.2
Outer	52								
Middle	40								None
Inner	None	40					None		

4.3.2 SCT modules

As indicated in figure 4.1, the SCT consists of 4088 modules [67, 68] tiling four coaxial cylindrical layers in the barrel region (called ID layers 3-6) and two end-caps each containing nine disk layers [60]. The modules cover a surface of 63 m² of silicon and provide almost hermetic coverage with at least four precision space-point measurements over the fiducial coverage of the inner detector. Tables 4.5 and 4.6 show the SCT detector parameters in detail. Details of the SCT structure are described in section 4.7.

The 2112 barrel SCT modules [67] use 80 μ m pitch micro-strip sensors [64], which are described in section 4.2. The sensors are connected to binary signal readout chips [76], which are described in section 4.4.1.2. The barrel module is shown, with its components, in figure 4.7. The module parameters are shown in table 4.7. The four sensors, two each on the top and bottom side, are rotated with their hybrids by ± 20 mrad around the geometrical centre of the sensors. They are glued on a 380 μ m-thick thermal pyrolytic graphite (TPG) base-board [77], which provides the




Figure 4.7: Photograph (left) and drawing (right) of a barrel module, showing its components. The thermal pyrolytic graphite (TPG) base-board provides a high thermal conductivity path between the coolant and the sensors.

thermal and mechanical structure. This extends sideways to include beryllia facings. A polyimide hybrid [78] with a carbon-fibre substrate bridges the sensors on each side. The two 770-strip (768 active) sensors on each side form a 128 mm long unit (126 mm active with a 2 mm dead space). High voltage is applied to the sensors via the conducting base-board.

Precision alignment criteria were applied during assembly: the assembly tolerance as well as the achieved build accuracy are shown in table 4.7. The important in-plane tolerance for positioning sensors within the back-to-back stereo pair was $< 8 \mu\text{m}$ and the achieved variance was $2 \mu\text{m}$. In the module plane, no additional distortions were measured after thermal cycling. Out-of-plane, the individual components and the assembly jiggling and gluing determine the module thickness and the intrinsic bow of the sensors determines the out-of-plane shape. A common distortion profile has been established for the sensors at the level of a few μm and a module thickness variation of $33 \mu\text{m}$ was maintained during fabrication. Following thermal cycling, the out-of plane distortions changed by a few μm (RMS). When cooled from room to operating temperature, profile deviations did not exceed $20 \mu\text{m}$, even at the sensor corners not supported by the base-board.

Figure 4.8 shows the construction of an end-cap module [68]. There are three module types, as shown in table 4.7. Each of the 1976 modules has two sets of sensors glued back-to-back around a central TPG spine with a relative rotation of $\pm 20 \text{ mrad}$ to give the required space-point resolution in $R-\phi$ and R . The module thickness is defined by the individual components and variations are compensated by the glue thickness (nominally $90 \mu\text{m}$). The TPG spine conducts heat from the sensors to cooling and mounting points at the module ends and serves as the bias contact to the sensors. Glass fan-ins attach one end of the spine to a carbon base-plate with the polyimide flex-hybrid glued to it. The modules are arranged in tiled outer, middle and inner rings.

The precision alignment criteria applied to the end-cap modules were similar to those of barrel modules. The RMS spread of the module survey measurements after construction was $1.6 \mu\text{m}$ in the back-to-back position of the stereo pair, measured transverse to the strips, and $2.8 \mu\text{m}$ in the position of the mounting hole and slot measured transverse to the strips. In the module plane, no additional distortions were measured after thermal cycling. Out of the plane, the end-

Table 4.7: SCT barrel and end-cap module specifications and the RMS build accuracy for accepted modules. The barrel out-of-plane bowing specifications and the measured results are with respect to an average module shape. The quoted hybrid and sensor power consumption is more extensively described in section 4.5.

Parameter	Description
Strips	2×768 active strips, ± 20 mrad stereo rotation
Nominal resolution	17 μm in-plane lateral (R - ϕ) 580 μm in-plane longitudinal (z or R)
Module dimensions: -barrel -outer end-cap -middle end-cap -short-middle end-cap -inner end-cap	Active length 126.09 mm + 2.09 mm dead space between sensors Active length 119.14 mm + 2.09 mm dead space, radius 438.77 – 560.00 mm Active length 115.61 mm + 2.09 mm dead space, radius 337.60 – 455.30 mm Active length 52.48 mm, radius 402.82 – 455.30 mm Active length 59.1 mm, radius 275.00 – 334.10 mm
Specified build tolerance	Barrel back-to-back in plane: $< 8 \mu\text{m}$ (lateral) $< 20 \mu\text{m}$ (longitudinal) End-cap back-to-back in plane: $< 5 \mu\text{m}$ (lateral) $< 10 \mu\text{m}$ (longitudinal) Barrel out-of-plane (module thickness and sensor bowing): $< 70 \mu\text{m}$ Barrel envelope: $< 200 \mu\text{m}$ End-cap envelope: $< 115 \mu\text{m}$ Barrel module fixation points with respect to module centre: $< 40 \mu\text{m}$ End-cap module fixation points with respect to module centre: $< 20 \mu\text{m}$
Build accuracy of accepted modules (RMS)	Barrel back-to-back in plane: $\pm 2.1 \mu\text{m}$ (lateral) $\pm 2.7 \mu\text{m}$ (longitudinal) End-cap back-to-back in plane: $\pm 1.6 \mu\text{m}$ (lateral) $\pm 1.3 \mu\text{m}$ (longitudinal) Barrel module thickness: $\pm 33 \mu\text{m}$ Barrel out-of-plane (sensor bowing): $\pm 9 \mu\text{m}$ End-cap module thickness: $\pm 15 \mu\text{m}$ End-cap out-of-plane (sensor bowing): $\pm 20 \mu\text{m}$ Barrel module fixation points with respect to module centre $\pm 10 \mu\text{m}$ End-cap module fixation points with respect to module centre $\pm 6 \mu\text{m}$
Hybrid power consumption	5.5 – 7.5 W
Sensor power consumption	Up to 460 V bias, < 1 W at -7°C

cap modules are less rigid, and are affected by variations of the spine thickness and bowing of the sensors. A common distortion profile has been established for the sensors at the level of a few μm and a module thickness variation of 15 μm was maintained during fabrication. Following thermal cycling, the out-of plane distortions changed by only a few μm (RMS). As for the pixels, the variations within the module are small compared to the subsequent module-placement accuracies.

The spatial resolution of individual SCT modules has been measured in a test beam, for both non-irradiated and fully irradiated modules [66]. At normal incidence, a combined spatial resolution of $\sim 16 \mu\text{m}$ is measured in R - ϕ ; the resolution is consistent with the binary readout of the two sensors with 80 μm strips, including a small fraction of multiple hits. The resolution is not significantly degraded after irradiation. The barrel staves are mounted at a tilt angle of $\sim 11^\circ$,




Figure 4.8: The upper photograph shows the three SCT end-cap module types (outer, middle and inner from left to right). The lower schematic shows an exploded view of the different components for a middle module, including the high thermal conductivity spine, the polyimide hybrid and the ABCD readout ASIC's.

defined by geometrical constraints and opposite in sign to that of the pixel barrel staves because of the different nature of the sensors used in each detector. The measured Lorentz angle, however, for a magnetic field of 2T, varies between 4.2° (before) and 2.7° (after) irradiation.

The barrel and end-cap sensors are specified to operate at -7°C , with a maximum variation within and between modules of 5°C , to reduce the bulk leakage current after radiation damage. The hybrid power will be $5.5 - 7.5\text{ W}$ per module, and the sensor load will reach $\sim 1\text{ W}$ per module after ten years of operation. In addition, convective loads of $\sim 0.8\text{ W}$ per module plus $\sim 0.8\text{ W}$ per module at the top of the barrel cylinders and outer disks are expected. The heat is extracted by evaporating C_3F_8 at $\sim -25^\circ\text{C}$, circulating in cooling pipes attached to each module.

For the barrel, the sensor and hybrid heat leaves via the base-board and the hybrid substrate to the large beryllia facing on the base-board, which is interfaced to an aluminium block with a $\sim 100\mu\text{m}$ layer of thermal grease and a copper-polyimide capacitive shunt shield. At full load for irradiated modules, the hybrid and sensor temperatures are expected to be approximately 14°C and 12°C above the cooling-pipe temperature, respectively. The block is itself soldered to a 3.6 mm diameter Cu/Ni cooling pipe. Each cooling loop serves 48 barrel modules.

For the end-cap, the sensor heat leaves via the spine, while the hybrid heat is transferred via the carbon-fibre hybrid substrate to a carbon-carbon cooling block, which is split to minimise heat

Table 4.8: Parameters of the TRT barrel modules and end-cap wheels (for one TRT end-cap). The quantities shown in bold are global parameters, including services and electronics. All other quantities are for individual modules and the active region. Type-1 barrel modules include two straw types, as described in the text.

	$ z _{\min}$ (mm)	$ z _{\max}$ (mm)	R_{\min} (mm)	R_{\max} (mm)	Number of modules	Number of layers	Straws per module
Barrel (both sides)	0	780	554	1082	96	73	52544
Type-1 module (inner)	400	712.1	563	624	32	9	329
Type-1 module (outer)	7.5	712.1	625	694		10	
Type-2 module	7.5	712.1	697	860	32	24	520
Type-3 module	7.5	712.1	863	1066	32	30	793
End-cap (one side)	827	2744	615	1106	20	160	122880
Type-A wheels	848	1705	644	1004	12	8	6144
Type-B wheels	1740	2710	644	1004	8	8	6144

transfer between the sensor and the hybrid. At full load, the ASIC and sensor temperatures are expected to be respectively $\sim 30^\circ\text{C}$ and $\sim 10\text{--}15^\circ\text{C}$ above the coolant temperature. A layer of thermal grease is applied between the modules and the cooling block. The blocks are soldered to a Cu/Ni cooling pipe that serves up to 33 modules.

All modules were tested electrically at room temperature and at the expected operating temperature of -7°C . Thermal cycling was performed on each module prior to completion of electrical testing. After module assembly and testing, 99.8% of the strips operate satisfactorily.

4.3.3 TRT modules

The TRT contains up to 73 layers of straws interleaved with fibres (barrel) and 160 straw planes interleaved with foils (end-cap), which provide transition radiation for electron identification. All charged tracks with $p_T > 0.5\text{ GeV}$ and $|\eta| < 2.0$ will traverse at least 36 straws, except in the barrel-end-cap transition region ($0.8 < |\eta| < 1.0$), where this number decreases to a minimum of 22 crossed straws. Typically, seven to ten high-threshold hits from transition radiation are expected for electrons with energies above 2 GeV.

The barrel TRT is divided into three rings of 32 modules each, supported at each end by a space frame, which is the main component of the barrel support structure (see section 4.7). Each module consists of a carbon-fibre laminate shell and an internal array of straws embedded in a matrix of $19\text{ }\mu\text{m}$ -diameter polypropylene fibres serving as the transition radiation material. The straws, described in section 4.2, form a uniform axial array with a mean spacing of $\sim 7\text{ mm}$. The module shells are non-projective to reduce the dead region for high p_T tracks. The main barrel parameters are shown in table 4.8. Figure 4.9 shows a quadrant of the TRT barrel during the integration of the modules at CERN. A ϕ -slice showing one outer, one middle and one inner module is highlighted.

The dimensional specifications are set by the intrinsic straw R - ϕ resolution of $130\text{ }\mu\text{m}$, implying that each wire position is constrained to within $\pm 50\text{ }\mu\text{m}$. The module shell, made of $400\text{ }\mu\text{m}$ thick carbon fibre with high thermal conductivity and flat to within $250\text{ }\mu\text{m}$, is measured to satisfy maximum distortions of $< 40\text{ }\mu\text{m}$ under full load.

The module shells also serve as a gas manifold for CO_2 which circulates outside the straws to prevent high-voltage discharges and the accumulation of xenon due to possible gas leaks which would absorb the transition radiation photons. The heat dissipated by the barrel straws is transferred to the module shell by conduction through the CO_2 gas envelope. Each module shell is cooled by two cooling tubes located in the acute corners. These tubes also serve as return pipes for the C_6F_{14} cooling circuits of the front-end electronics.

The module end with its components is shown in figure 4.11. The central element is the HV plate which has stringent requirements on flatness and cleanliness to prevent discharges and on the straw feed-through accuracy to ensure mechanical precision of the straw location. The HV plates were individually surveyed after machining. The tension plate of figure 4.11 mounted on the HV plate is a printed-circuit board holding the wire ends (and ensuring the wire tension) and providing electrical connections. It also closes the active gas volume and serves as a Faraday cage for the active module elements.

Mechanical and electrical tests on the modules and individual straws were made at successive assembly stages and following delivery of the modules to CERN. The module lengths were required to be in the range $1461.5\text{--}1462.9\text{ mm}$ and the twisting of individual modules was required to be $< 1\text{ mm}$. Stringent criteria were applied to the gas tightness, wire tension, straw straightness, and high voltage stability. The modules installed in the barrel had 98.5% operational channels.

The TRT end-caps each consist of two sets of independent wheels (see figure 4.10 and table 4.8). The set closer to the interaction point contains 12 wheels, each with eight successive layers spaced 8 mm apart. The outer set of wheels contains eight wheels, also with eight straw layers but spaced 15 mm apart. Each layer contains 768 radially oriented straws of 37 cm length with uniform azimuthal spacing. The space between successive straw layers is filled with layers of $15\text{ }\mu\text{m}$ thick polypropylene radiator foils separated by a polypropylene net.

Each eight-plane wheel consists of two basic four-plane assembly units. To assemble a four-plane wheel, straws were inserted and glued into precisely drilled holes in grounded inner and outer carbon fibre rings. The rings and the straws constitute the main mechanical structure of the wheels. The successive straw layers, interleaved by the radiators, are rotated from one layer to the next



Figure 4.9: Photograph of one quarter of the barrel TRT during integration. The shapes of one outer, one middle and one inner TRT module are highlighted. The barrel support structure space-frame can be seen with its triangular sub-structure.

by $3/8$ of the azimuthal straw spacing in a given layer. For high- p_T prompt tracks, this ensures optimal uniformity in the number of crossed straws, which varies radially from \sim six to \sim four straws across an eight-plane wheel.

Flex-rigid printed-circuit boards provide high-voltage and signal connections to the end-cap wheels through two separate layers, as shown in figure 4.12. Each of the flexible layers has conducting paths on one side which connect to the rigid part of the board. To provide a reliable electrical connection, flexible "petals" in the high-voltage layer are forced into contact with the inner straw wall through the insertion of a plastic plug [71]. A press-fit between similar but smaller petals in the signal layer and a metallic crimping pin positions and fixes the anode wires. There are 32 such boards per four-plane wheel, each serving a ϕ -sector of 96 straws.

Each sector is further segmented into three groups of 32 readout channels and 12 high-voltage groups of eight straws sharing a common fuse and blocking capacitor. The carbon-fibre ring holding the straws and flex-rigid boards, together with a third carbon fibre ring and a simpler glass fibre board provide a rigid structure around the outer wheel perimeter that also serves as a gas manifold (see figure 4.12). The inner gas manifold is made from reinforced polyimide material.

The heat dissipated by the end-cap straws is evacuated through the CO_2 gas envelope which is forced to flow along the straws from the inner to the outer radius. Each group of wheels has its own CO_2 cooling circuit, passing the gas sequentially through all the wheels of the group. Heat exchangers cooled with C_6F_{14} extract heat from the gas between adjacent wheels. The high flow rates required, 50 m^3 per hour for type-A and 25 m^3 per hour for type-B wheels, necessitates a closed-loop system capable of maintaining a small gas pressure between 0 and 5 mbar with a stability of ± 0.5 mbar inside the detector.

As for the barrel modules, all end-cap wheels passed quality control procedures during construction and after delivery to CERN, resulting in more than 99% of fully operational channels.




Figure 4.10: Photograph of a four-plane TRT end-cap wheel during assembly. The inner and outer C-fibre rings can be seen, as well as the first layer of straws and the first stack of polypropylene radiator foils beneath it. Also visible are the plastic end-plugs which are used to position and fix the straws to the C-fibre outer ring. The high-voltage petals used to connect the straws to the high-voltage lines (see text) are laid back at this stage of the assembly and will only be folded vertically to push the petals into the straws at the next stage.




Figure 4.11: Detailed view of the end of the TRT barrel modules, showing the connection of the straw ends to high voltage through the high-voltage (HV) plate, of the wires to the front-end boards through the tension-plate and of the gas inlet to the individual straws through the active-gas manifold.




Figure 4.12: Schematic view of the inner and outer ends of the TRT end-cap wheels, showing the plastic end-plugs used to position and fix the straws in the inner and outer C-fibre rings, the crimping pins holding and positioning the wires, the inner and outer active-gas manifolds, and the flex-rigid printed-circuit board used to connect the straws to high voltage and the wires to the front-end electronics.

4.4 Readout of the inner detector

The readout architecture of the ID is optimised separately for each of the three sub-detectors, but is characterised by the following common elements:

- (a) the reception of a 40.08 MHz clock signal synchronous with the LHC bunch-crossings used to time-stamp the signal generated in the low noise front-end electronics;
- (b) signal generation in the front-end electronics and storage in binary or digital buffers for approximately $\sim 3.2 \mu\text{sec}$, compatible with the L1 trigger latency of $2.5 \mu\text{sec}$;
- (c) following a L1 trigger, the subsequent transfer of the buffer content associated with the bunch-crossing or possibly several bunch-crossings to a readout driver (ROD) off the detector.

The readout of the pipelines is subject to the reception of a L1 trigger signal from the calorimeters or the muon detectors via the central trigger processor, as discussed in section 8.2.3. The ID is not part of the L1 trigger. External power supplies provide regulated voltages to the front-end electronics and to the sensors (see section 4.5).

4.4.1 Front-end electronics

4.4.1.1 Pixel front-end electronics

Each front-end readout ASIC [63, 79, 80] of the pixel detector contains 2880 readout cells of $50 \times 400 \mu\text{m}^2$ size arranged in a 18×160 matrix. Sixteen ASIC's are bump-bonded to each sensor (module). The ASIC's are fabricated using commercial $0.25 \mu\text{m}$ CMOS technology. A high level

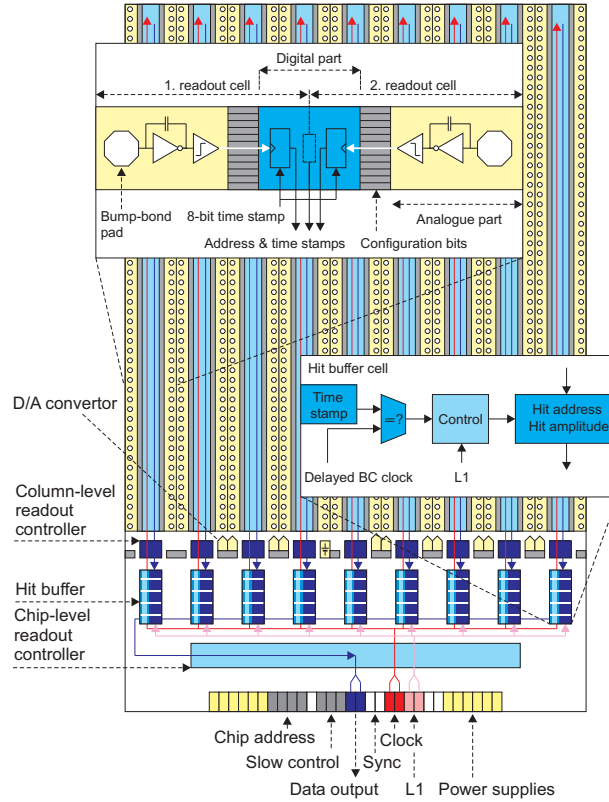


Figure 4.13: Layout and schematic description of the front-end readout ASIC for the pixel detector (see text).

of radiation tolerance is achieved by a combination of the process technology (thin gate oxide) and the use of special layout techniques (annular layout and guard rings for all NMOS transistors). The schematic and basic functionality of the pixel circuit is shown in figure 4.13.

Each readout cell contains an analogue block where the sensor charge signal is amplified and compared to a programmable discriminator threshold. The digital readout then transfers the hit pixel address, a hit time stamp and a digitised amplitude (the time over threshold (ToT)) to buffers at the chip periphery. These hit buffers monitor each stored hit by inspecting the associated time stamp.

The charge sensitive amplifier uses a single-ended folded-cascode topology optimised for a nominal capacitive load of 400 fF and designed for the negative signal expected from DC-coupled n^+ -on- n sensors. Attention has been paid to the preamplifier design because, following the irradiation expected at the LHC, the sensor leakage current (50 nA) is two orders of magnitude larger than the signal, that is itself reduced due to carrier trapping inside the silicon. The pre-amplifier has an approximate 5 fF DC feedback capacitance with a 15 ns risetime. The total analogue front-end (pre-amplifier, second stage amplifier and discriminator) has a bias current of only 24 μ A per pixel for the default DAC settings. To ensure the separation of contiguous bunch crossings, a front-end time walk of <25 ns is required. To fulfill the requirements of sensor leakage current, a compensation circuit is implemented that drains the leakage current and prevents any influence on the bias current of the fast feedback circuit used to discharge the feedback capacitor. Each pixel has several

parameters that are tuned and stored in a 14-bit control register, for example the feedback trimming and threshold levels. The digital circuitry in the readout cells generates the required hit information to measure the charge and associate the hit to the bunch crossing.

The readout is made using the column-based readout controller. The first task of the controller is the generation of the readout sequence to transfer the hit information from the pixel to the hit buffer at the end of the column. The second task is the digital processing of hit data in the front-end chip buffers. Hits are continuously compared to a delayed beam-crossing counter (see figure 4.13). When a match is found with the correct latency, the hit is flagged for readout if a L1 trigger signal is present, or deleted if there is no matching trigger. Flagged hits are transmitted to a serialiser and sent out of the chip. Hits older than the trigger latency are cleared from the front-end chip buffers.

The module-control chip [63, 81] is a digital chip running with the same 40 MHz clock. It has three main system tasks: the loading of parameter and configuration data in the front-end chips and in the module-control chip itself, the distribution of timing signals such as bunch-crossing, L1 trigger and resets (TTC functions), and the front-end chip readout and event building. The design of the module-control chip reflects the required pixel performance during LHC operation: the association of signals to a bunch-crossing, the expected bandwidths at the highest luminosity, the maximum L1 trigger rate of 100 kHz (initially 75 kHz) and the number of front-end chips, which are controlled in a module. Because of the high radiation environment, particularly in layer-0 modules, special attention has been given to ensure a single-event upset (SEU) tolerant design.

Extensive electrical tests have been made on ASIC chips, single front-end chip assemblies and full modules, before and after irradiation. Some production modules were irradiated to the end-of-life dose expected at LHC. The noise and hit efficiency are shown in figure 4.14 for a pixel module before irradiation. Similar results are also shown for a module irradiated to a fluence F_{neq} of 10^{15}cm^{-2} (and to 500 kGy ionising dose) in a 24 GeV proton beam, at an operating temperature of -4°C and after threshold retuning. Both quantities are only slightly deteriorated by the irradiation and remain within the operating specifications. Furthermore, the difference in threshold dispersion after retuning is negligible.

4.4.1.2 SCT front-end electronics

The readout hybrid of each SCT module (see section 4.3) houses 12 identical 128-channel ASIC's [76] to read a total of 1536 sensor strips per module. The ASIC is fabricated in radiation tolerant bi-CMOS DMILL technology. The successive blocks of the ASIC are shown in the circuit schematic of figure 4.15. A pre-amplifier, shaper and tunable discriminator exists for each channel. A 132-length binary pipeline stores the hit information for each channel associated to the beam crossing for a period of $\sim 3.2 \mu\text{sec}$. Following a L1 trigger, the chip compresses the data pertinent to that beam crossing and serialises it for output. An 8-deep de-randomising buffer after the pipeline ensures that the dead-time is negligible for the expected data rates.

Two critical module performance specifications are the detection efficiency ($> 99\%$) and noise occupancy ($< 5 \times 10^{-4}$), for signals from the 12 cm long silicon strips with a capacitive load of $\sim 20 \text{pF}$. These have led to the choice of a front-end discriminator threshold of 1 fC. Extensive studies have been made using ASIC chips, single front-end chip assemblies and full modules, before and after irradiation. A sample of production modules was irradiated in a 24 GeV proton




Figure 4.14: Noise distribution for normal pixels of a non-irradiated module (upper left) and of a module irradiated with 24 GeV protons to a fluence F_{neq} of 10^{15} cm^{-2} (lower left), measured after retuning at an operating temperature of -4°C . The measured efficiency as a function of the incident beam particle arrival time is also shown for the non-irradiated (upper right) and irradiated modules (lower right). The arrows indicate the efficiency at the timing plateau.

beam to a dose of $\sim 3 \times 10^{14}$ protons per cm^2 , equivalent to the damage expected for a fluence F_{neq} of $2 \times 10^{14} \text{ cm}^{-2}$. The efficiency and noise occupancy measured in a test-beam are shown as a function of the discriminator threshold in figure 4.16. At the nominal operating threshold of 1 fC, the efficiency and noise-occupancy specifications are easily met before irradiation and almost met after irradiation.

The chips are daisy-chained so that all the data of one module are read out over two serial links. Several design features provide fault tolerance. For example, any faulty chip (except for the master on barrel modules) can be bypassed in the serial data path and if one link should fail, it is possible to send the data using the remaining active link. Likewise, if the primary clock and command lines to the ASIC's on a module fail, it is possible to instead use the clock and command signals from an adjacent module.

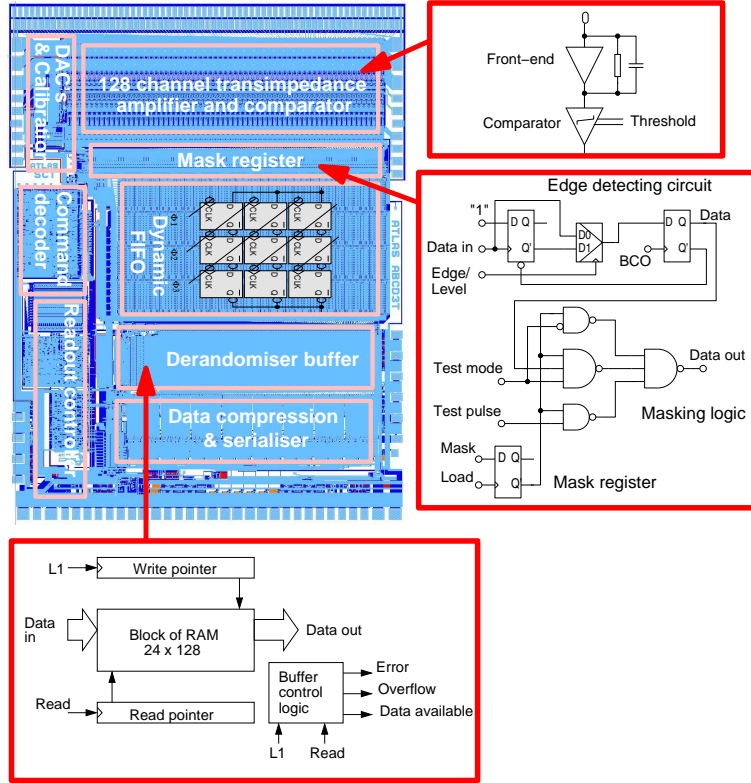


Figure 4.15: Schematic of the readout ASIC for the SCT detector, showing the successive signal processing steps.




Figure 4.16: The efficiency (circles) and noise occupancy (triangles) for SCT barrel modules measured in a test-beam before irradiation (left) and after exposure to a dose of $\sim 3 \times 10^{14}$ p/cm² in a 24 GeV proton test-beam (right). The nominal operating threshold is 1 fC. The dashed horizontal lines represent the nominal module performance specifications in terms of efficiency and noise. The vertical lines represent the range of thresholds over which these specifications are met after irradiation.

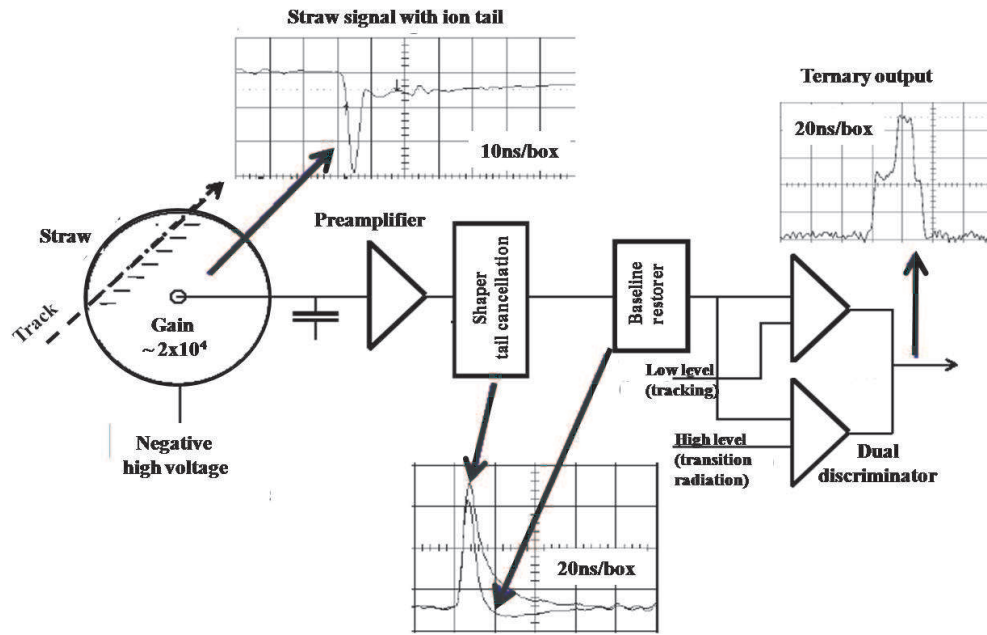


Figure 4.17: Schematic of the front-end readout of the TRT detector, showing the input signal shape and the signal shape after the amplification and shaping, the baseline restoration (BLR) and the dual-threshold discrimination which provides the ternary output corresponding to the low threshold set or both low and high thresholds set.

4.4.1.3 TRT front-end electronics

The analogue signal processing and threshold discrimination to detect signals from both minimum-ionising particles and transition radiation, as well as the subsequent time digitisation and data-pipelining are implemented in two on-detector ASIC's. The signal is shown in figure 4.17 at each stage of the TRT signal readout chain which comprises:

- an eight-channel analogue ASIC [82], called the ASDBLR, fabricated in bi-CMOS radiation tolerant DMILL technology. It performs the amplification, shaping and baseline restoration. It includes two discriminators, one operating at low threshold (typically 250 eV) for minimum-ionising signal detection and one operating at high threshold (typically 6 keV) for transition radiation detection;
- a subsequent 16-channel ASIC fabricated in commercial radiation-tolerant $0.25\mu\text{m}$ CMOS technology [83]. This ASIC performs the drift-time measurement (~ 3 ns binning). It includes a digital pipeline for holding the data during the L1 trigger latency, a derandomising buffer and a 40 Mbits/s serial interface. It also includes the necessary interface to the timing, trigger and control as well as DAC's to set the discriminator thresholds of the analogue ASIC and test-pulse circuitry for mimicking analogue inputs to the analogue ASIC.

These ASIC's are housed on front-end boards attached to the detector. There are 12 different boards for the barrel and three different boards for the end-cap. The electronics are cooled by a liquid mono-phase fluorinert (C_6F_{14}) cooling system.

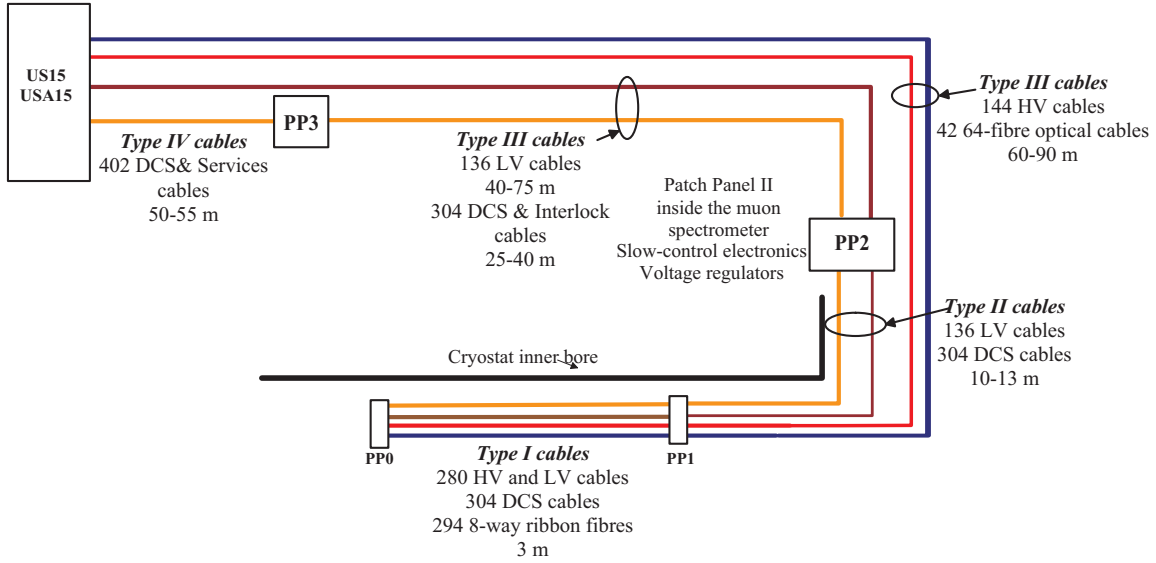


Figure 4.18: The routing of data links and power supply cables from each side of the pixel detector to respectively the off-detector electronics and power supplies in the service caverns, together with the number, type and utilisation of the cables and optical links.

At the TRT operating low threshold used for tracking (equivalent to $\sim 15\%$ of the average signal expected from minimum-ionising particles), the mean straw noise occupancy is $\sim 2\%$, but a small fraction of 1% of channels have a noise occupancy exceeding 10%, which however remains small compared with the expected maximum straw occupancy of 40%. The full front-end electronics chain was exposed to a neutron dose of $\sim 4 \times 10^{14}/\text{cm}^2$ and to a γ -ray dose of 80 kGy. Changes of up to 25% were observed in the ASDBLR gain, but with no change in the effective thresholds and noise performance after a standard voltage compensation procedure.

4.4.2 Data transmission and power-supply services and routing

The transmission of data from the ID modules to the off-detector electronics in the service cavern, as well as the digital transmission of the clock and control commands to the modules differs for the 3 sub-detectors. Figures 4.18, 4.19 and 4.20 summarise the layout and technology of the readout and control services for each of the pixel, SCT and TRT sub-detectors. The locations of key patch-panel connection boards are shown: PP0 close to the ends of the pixel detector, PP1 at the edges of the ID volume, PP2 in specifically designed parts of the muon spectrometer system and PP3 outside the ATLAS active detector volume. The numbers and lengths of lines for each module (pixel, SCT) or front-end board (TRT) are tabulated for both the barrel and end-caps. Similarly, the cables used for the readout electronic bias as well as the silicon sensor and TRT straw high-voltage lines are also shown.




Figure 4.19: The routing of data links and power supply cables from each side of the SCT to respectively the off-detector electronics and power supplies in the service caverns, together with the number, type and utilisation of the cables and optical links.




Figure 4.20: The routing of data links and power supply cables from each side of the TRT to respectively the off-detector electronics and power supplies in the service caverns, together with the number, type and utilisation of the cables and optical links.

Table 4.9: Summary of the main characteristics of the optical links used for the pixel and SCT data readout and trigger, timing and control (TTC) signals (NRZ stands for Non-Return to Zero and BPM for Bi-Phase Mark).

Sub-detector	Link	Speed (Mbits/s)	Format	Links per module	Number of links
Pixels	Data	40/80	NRZ	1 (layer-1, 2 and disks)	1458
		40/80	NRZ	2 (layer-0)	572
	TTC	40	BPM	1	1744
SCT	Data	40	NRZ	2	8176
	TTC	40	BPM	1	4088

On the other side of the link, back-of-crate (BOC) cards interface the opto-signals with the electrical signals in the readout driver (ROD). These cards are the same for the pixel and SCT detectors and they contain two kinds of electro/optical converter plug-ins: a RX plug-in with an array of PIN diodes and a data receiver, and a TX plug-in with an array of VCSELs driven by a bi-phase mark encoding ASIC. The SCT uses a radiation-hard step index multi-mode fibre for the full length from the back-of-crate cards to the detector. The pixel detector uses a long length of radiation-tolerant graded-index fibre, spliced to 7 m lengths of the radiation-hard fibre inside the detector.

4.4.2.2 TRT readout

As described in section 4.3.3, the TRT readout is segmented in 32 ϕ -sectors in order to simplify the data transmission for L2 triggers. The initial 40 Mbits/s LVDS data readout uses small custom designed twisted pair lines from the ASIC to the boards at patch panel PP2, located after the first muon chambers. The data are then serialised in a Gbit serialiser [87] and an electrical-to-optical conversion is made. The complete readout of the TRT requires 768 1.6 Gbits/s optical links. Timing and control signals are electrically transmitted to and from the TRT-TTC module.

4.4.2.3 Pixel, SCT and TRT readout drivers

As described in section 8.3.5, all readout driver modules (ROD's) receive serialised data from the detector upon reception of a L1 trigger, perform de-serialisation, error checking, data compression (in the case of the TRT), and local event-building and data-monitoring tasks. The ROD's combine the data received into a single event packet associated to the L1 trigger. The event packets are transferred in a standard data format over a high-speed optical link, the ATLAS S-link, to the readout system, which is common to all ATLAS sub-detectors.

Each pixel and SCT ROD crate is a 9U VME crate with up to 16 ROD's per crate. Each ROD is paired with a back-of-crate card, which is plugged into the crate back-plane and provides the optical to electrical interfaces [85, 88, 89]. The back-of-crate card can accept both 40 Mbits/s optical

links from pixel and SCT modules and 80 Mbits/s from pixel modules. Each ROD card services up to 48 SCT modules (96 data links) or between 6 and 26 pixel modules, depending on their geometrical location. The lower number of pixel modules serviced per ROD is due to the larger quantity of data received from pixel modules at the smallest radius, due to their higher occupancy. The number of modules which can be processed by one ROD is limited by the output bandwidth of the S-link. The pixel and SCT ROD's are identical units except for the data-treatment algorithms. The ROD crates receive clock, trigger and fast commands from the ATLAS TTC system through a TTC interface module (TIM), located in each crate which transmits these signals across the crate, using a custom back-plane, to the ROD's and back-of-crate cards. Clock and command signals are then transmitted to the detector modules using a simpler protocol over the optical fibres mentioned above.

The TRT off-detector electronics also uses two custom-designed TRT-ROD and TRT-TTC modules [90]. Each ROD receives serialised data from 1/16 of one barrel side or 1/32 of one end-cap side, using eight optical links operating at 1.6 Gbits/s. After the optical-to-electrical conversion, the ROD's perform the standard tasks described above, including a data-compression scheme which does not introduce any data losses. The TRT-TTC module provides an interface between the standard ATLAS TTC system and the TRT front-end electronics. It also feeds the ROD with all necessary L1 trigger information (event ID, bunch crossing ID, trigger type). Each TRT-TTC module interfaces to the front-end using 40 links and to two ROD's via a dedicated back-plane using the VME 9U P3 connector.

4.5 Electronics and detector power supplies and services

Both the sensors and the front-end electronics of the inner detector require substantial power for initial operation and the power load will increase during the high-luminosity operation. Table 4.10, together with figures 4.18 to 4.20, summarise the requirements for each of the pixel, SCT and TRT sub-detectors, in terms of bias voltage for the silicon sensors or high voltage for the straws (HV) and of low voltage (LV) for the front-end electronics, as well as the segmentation and routing used for the power distribution. The power distribution lines are connected at PP0 and PP1 on the cryostat wall, at PP2, which is external to the inner-detector volume and located after the first muon chambers, and at PP3 on the outside of the whole detector.

The pixel power supply system has four main components: the LV and HV power supplies, the regulator stations, and the supply and control for the opto-links. Two commercial LV supplies provide the analogue and digital parts of the front-end readout electronics. To protect the front-end electronics against transients, remotely-programmable and radiation-tolerant regulator stations are installed at PP2. Separate HV supplies are able to power the sensors up to 700 V with a maximum module current of 4 mA. The worst-case expected module operating condition is a 600 V depletion voltage with a 2 mA leakage current. The LV and HV lines are connected to respectively the low-voltage and high-voltage patch-panels which distribute the power and monitor the currents of individual lines. The supply and control of the optical link is a complex link in itself, consisting of three voltage sources and a control signal. It delivers the adequate levels for the operation of the on-detector part of the optical link.

Table 4.10: Summary of the silicon sensor bias voltage, straw-tube high-voltage and front-end electronics requirements and granularity. The different front-end electronic supply voltages are described in the text. The power consumption expected during initial operation is also shown. Following irradiation, the power consumption increases significantly in the sensors, front-end electronics and cables (see text). Additional power requirements to ensure thermal isolation in the ID volume are described in the text.

		Pixel	SCT	TRT
Bias or high-voltage supplies	Voltage (maximum)	700 V	500 V	2000 V
	Voltage (nominal)	150–600 V	150–350 V	1600 V
	Current (maximum)	4 mA	5 mA	3 mA
	Segmentation	One per module 1744 channels	One per module 4088 channels	One per ~ 200 straws 1984 channels
Front-end electronics low-voltage supplies	Voltages	1.7–2.1 V analogue 2.1–2.5 V digital 0.8, 2.5, 10 V opto-device	3.5 V analogue 4 V digital 5, 10 V opto-device	± 3 V analogue 2.5 V digital
	Current	3.7 kA	6 kA	6.5 kA
	Segmentation of bulk supply	Configurable 6–7 modules	One per module 4088 channels	1/32 of end-cap side 1/32 of barrel side
	Segmentation of regulated supplies	One regulator per module	See above	One regulator per 1/32 of wheel or per barrel module
Power	Front-end electronics power	6 kW	22 kW	22 kW
	Cables plus regulators:			
	- in ID volume	~ 4 kW	~ 5 kW	~ 3 kW
	- outside ID volume	~ 14 kW	~ 9 kW	~ 20 kW
	- total cable loss	~ 18 kW (~ 6 V drop)	~ 14 kW (1.5–4.6 V drop)	~ 23 kW (3.8 V drop)
	Total power (initial operation)	~ 24 kW	~ 36 kW	~ 44 kW

The SCT [91] maintains independent electrical services to each of the 4088 modules. Each module receives multi-voltage LV channels providing power and control signals to the readout chips, the optical electronics, timing and control electronics, as well as HV to the silicon sensors. The SCT LV power module provides the analogue and digital voltages for the SCT front-end ASIC's and optical-link components. Each LV power module also includes the hybrid temperature readout as well as the digital control lines. The HV bias provides a stable and controlled voltage of 0–500 V with a current limit of 5 mA. The LV and HV modules are based on DC-DC converters with a single power line serving each module; twelve identical 4-channel LV and six identical 8-channel HV boards are housed in each crate. Each board in the crate is connected through a system-interlock card to the detector control system (see section 4.8.4). The four crates of each rack are powered by four commercial power units (including one spare) connected in parallel to provide a 48V DC supply. As an example for the barrel SCT, the power distribution lines are in four parts: low-mass polyimide tapes for the innermost region of the detector, which are connected to PPB1 on the cryostat wall, thin cables along the cryostat wall, which are spliced to medium-size cables outside the cryostat, and, after PP3, thick cables going to the power-supply crates. Because of the large voltage drop in the cables (up to 4.6 V), the supplied voltage exceeds the maximum

allowed voltage of the ASIC's. This could be problematic, for example in case of a drop in the power consumption of the front-end system. To avoid such incidents, a voltage-limiter circuit has been introduced at PP3. To dampen common-mode noise, inductors on all power and low-current control and monitoring lines are also housed at PP3.

The TRT front-end electronics require three low-voltage power supplies (+2.5 V digital and ± 3 V analogue). Commercial supplies deliver power to boards located at PP2. The PP2 boards house radiation-tolerant voltage regulators delivering power to each front-end board. The size of the cables feeding the PP2 power boards and the front-end boards is a compromise between the available space and the power dissipation in the cable trays allowed by the cooling system. The TRT straw high voltage is nominally 1530 V. Sets of about 200 straws are powered by a single commercial HV source able to deliver up to 3 mA at 2 kV. A total of 1984 HV channels are needed for the whole detector. Standard multi-wire HV cables are used up to the level of PP2, then custom miniature HV cables are used to reach the detector.

The power dissipation during initial detector operation is also indicated in table 4.10. The initial power dissipation from the front-end electronics and cables is ~ 62 kW within the ID volume, and ~ 104 kW in total. However, during high-luminosity operation of the TRT and after irradiation of the pixel and SCT sensors, several kilowatts of power are dissipated by the sensors (this is sensitive to temperature for the silicon sensors). Together with additional losses in the cables and increased front-end power consumption, the total power loss within the ID volume is expected to exceed 85 kW after several years of operation.

Additional power supplies are required for the pad heaters located on the thermal enclosures surrounding the pixel and SCT detectors (see section 4.7.2) and for the heaters which are an integral component of the evaporative cooling system for the pixel and SCT detectors (see section 4.8.3). The thermal enclosures require a total power of almost 30 kW, including the heaters surrounding the beam-pipe (see section 4.8.1). The heaters for the cooling system require a total power of 126 kW, which is provided by four racks situated in the service caverns. These power supplies are controlled by programmable logic controllers, which regulate the entire system and manage, in particular, transients caused by changes in the power load as the detector electronics are switched on or off.

4.6 Grounding and shielding of the inner detector

The three components of the inner detector, the pixels, the SCT and the TRT are designed to be electrically isolated both from earth and from each other, as is generally required for all ATLAS components (see section 9.4.8). The detectors are therefore floating by design. This provides the possibility of a controlled ground connection at a location identified as IDGND between the tile calorimeter and the first layer of barrel muon chambers. A star connection of all the safety grounding cables for ID components, as well as common items (support structure, pipes for cooling and gas, heaters, and DCS sensors), ends in IDGND. The IDGND is connected to the ATLAS equipotential network by means of a dedicated 95 mm² conductor. To provide reasonable isolation of the sub-detectors and at the same time to complete the requirements of the safety grounding, the connection to IDGND from any sub-detector is made from its electromagnetic shield using a single-point connection. The power supplies are floating and




Figure 4.22: Schematic top view of the inner-detector sub-systems with their vertical support points. All supports are on the plane of the detector axis and symmetric with respect to this axis.

are referenced in most cases on the detector modules (analogue and digital grounds) and at the PP0 and PP1 patch-panels (connection of the detector-module grounds to the electromagnetic shields).

Each pixel module is referenced at PP0, and all the PP0's are grounded with both the electromagnetic shield and the beam-pipe shield at PP1. The pixel detector itself is surrounded by a Faraday cage defined by an aluminium layer surrounding the beam-pipe, and an aluminium layer surrounding the outside of the pixel support tube. The PP1 end-plates close the Faraday cage at the two ends and serve as the common grounding point for the pixel detector. There is a single grounding connection to IDGND on side-A of the detector. The SCT barrel and end-cap module assemblies are each enclosed by outer electromagnetic shield skins. The individual module supply tapes, outer supply cable shields and detector array shield skins are electrically bonded at PPB1 (for the barrel) and PPF1 (for the end-caps). The common shield nodes of the barrel and of each end-cap are single-point connected to IDGND through three dedicated grounding cables. The TRT barrel modules and wheels are fully enclosed in electromagnetic shields which are connected to IDGND using individual grounding cables.

4.7 Structure and mechanical integration of the inner detector

The complexity of the individual ID sub-detectors, together with the integration and installation requirements, the different sub-detector operating temperatures and the tight stability requirements, have resulted in a complex support scheme illustrated in figure 4.22. Most of the 66 support points are in the axial horizontal plane [92, 93].

The ID system is supported by rails fastened to the barrel cryostat inner wall. The ID sub-detectors rest on these cryostat rails and are basically supported at four points, implementing a quasi-kinematic support scheme, with simple support conditions in all directions except the vertical.




Figure 4.23: The barrel support structure of the TRT, which serves as the support for the full barrel ID, shown during initial assembly and measurements. The SCT detector is supported on carbon-fibre rails inside the carbon-fibre inner cylinder. The two-piece outer carbon-fibre cylinder is not assembled yet at this stage of the integration. The space-frame geometry at each end was designed to support individual TRT barrel modules. A number of mirrors, visible as bright spots, are photo-grammetry targets used for survey measurements.

The mechanical support for the barrel SCT system and for the barrel TRT modules is the barrel support structure, shown in figure 4.23. It is designed for high stiffness and stability, with $<10\ \mu\text{m}$ displacements under the expected temperature and humidity variations. It consists of two 21 mm thick carbon-fibre space frames, joined by inner and outer carbon-fibre cylinders. The SCT and TRT end-caps are each supported from a pair of girders, sliding on the cryostat rails. The separate 6.6 m long pixel and beam-pipe package includes the pixel support tube, which slides inside the SCT and is itself supported by the SCT barrel. Two end-plates provide the external supports for the beam-pipe and the pixel support tube.

Prior to integration as part of the full ATLAS detector, the barrel and end-caps for each of the pixel, SCT and TRT sub-detectors were separately assembled and fully tested on the surface. This section describes the overall mechanical structure of each of the sub-detectors. The subsequent




Figure 4.24: A perspective cut-away view of the pixel detector. The view shows individual barrel and end-cap modules, supported with their associated services on staves and disks within an octagonal support frame.

integration of the ID in ATLAS comprises four steps which are also described in this section: the integration of the barrel SCT and barrel TRT, the integration of the end-cap SCT and end-cap TRT (two end-caps), the integration of the barrel and end-cap pixel detectors with the beam-pipe, and finally the insertion of the pixel package.

4.7.1 Pixel structure and integration

The pixel detector and the pixel support tube (PST) within the ID are shown schematically in figure 4.1. The detector with its associated services and the vacuum inner detector beryllium beam-pipe (see section 4.8.1) are precisely located inside the PST. The pixel services (cooling, power and monitoring) are routed to the ends of the PST [92].

The active region of the pixel detector is shown in figure 4.24. The parameters of the pixel detector, with its three barrel layers and two end-caps, are listed in table 4.3 of section 4.3. The total active area of silicon is approximately 1.7 m^2 , with 112 barrel staves and 48 end-cap sectors (eight sectors per disk).

In the barrel region, the bi-staves are mounted in half-shells as illustrated in figure 4.25. Each half-shell is a thin carbon-fibre shell formed with facets to match the number of staves, with cut-outs to reduce the mass and with mounting rings that position the staves at five locations [92]. The disks of the pixel end-cap detector are bolted with precision bushings to a carbon-composite support ring. The disks are then held with four mounts within a section of the octagonal support frame to form an end-cap. An end-cap during the final stages of assembly is shown in figure 4.26, after connection of the cooling circuits. Two fully-loaded half-shells form a barrel layer. The largest layer in size, layer 2, is shown after this step in figure 4.27. Each barrel layer is inserted into the supporting octagonal frame and connected to end-cone structures with fingers to mate precisely with mounting brackets on the barrel. Capillaries and outlet cooling-tube extensions are then added.




Figure 4.25: A pixel barrel half-shell, with its cutouts, being loaded with barrel bi-staves and services.




Figure 4.26: A pixel end-cap at the last stage of assembly, after connection of its cooling circuits.

Following the integration of the beam-pipe, the barrel octagon and the two end-cap octagons, the barrel cooling pipes and cables are passed over the outside of the end-cap frame. All end-cap services are on the inside of the frame.

The PST itself consists of three sections. Each section is a cylinder with external stiffening rings and precision rails. The barrel section is made from carbon-fibre composite. Each end-section is composed of carbon fibre and fibreglass composites. The rails are carbon-fibre composites and are accurately positioned within each cylinder. The cylinders are joined at each end of the barrel by bolted carbon-fibre flanges. Heater panels (copper-on-polyimide printed-circuit boards) are glued to the cylinders. These are activated if there is a failure in the dry environmental gas around the pixel detector.




Figure 4.27: Barrel pixel layer-2, loaded with bi-staves, viewed along the axis after the joining of the half-shells.

The barrel PST section is precisely located with respect to the barrel SCT structure. Mounting points on the barrel PST receive mounts on the pixel detector support frame described below and locate the pixel detector to approximately $100\,\mu\text{m}$ with respect to the barrel SCT. An octagonal carbon-composite frame supports and positions the barrel and the two end-caps. The two end-cap sections are joined with composite bolts to the barrel frame. Mounts which position the pixel detector within the PST are located on the end-cap end-plates.

Each pixel module receives LV power directly from the LV regulators located at PP2 (see section 4.5). There are in total 88 cooling circuits for the pixel detector. Certain monitoring functions require connections within the pixel detector volume itself. The electrical and cooling connections pass through the end regions of the PST. Since the pixel detector is inserted from one end, all services, including connectors, must fit within the 230 mm radius PST inner envelope. The electrical services and cooling pipes are contained within service quarter-panels, which deliver one-quarter of the required services at each end of the pixel detector. Power and monitoring wiring is routed by individual twisted pairs and soldered to printed circuit boards (PP0) near the active detector, which contain miniature connectors. The low-mass cables from each module are plugged into these connectors. At the other end (PP1) the wires are soldered to printed circuit boards to penetrate the sealing plate. On the outside of the PP1 region, other twisted pairs are soldered to the printed circuit boards and terminated in commercial connectors. Optical transceivers are also located in the PP0 region and convert electrical signals from the detector to light transmitted by fibres to connectors at PP1. Control signals from outside the detector are also transmitted by fibres to the transceivers for conversion to electrical signals (see section 4.4).

Each cooling circuit includes a custom heat exchanger that consists of an inlet and an outlet tube glued together along the length of the service quarter-panel. These tubes penetrate the plate at PP1 and dry-gas integrity is maintained by a bellows seal that also allows for the 2–3 mm contraction of these aluminium pipes when the detector is operating. The capillaries and outlet extensions are attached at the PP0 end.

The service quarter-panels and the beam-pipe are supported by a composite beam-pipe support structure. The beam-pipe supports are adjustable from the ends of the PST to position the beam-pipe. The overall detector integration is illustrated in figure 4.28.

The pixel detector is sensitive to the high instantaneous rates that might occur during accidental beam losses [94]. For this reason, a set of small, fast and radiation-hard diamond detectors, called the ATLAS beam-conditions monitor (BCM) has been built and integrated into the pixel package to monitor the beam conditions and to distinguish lost beam particles from proton-proton interactions. The BCM, described in more detail in section 3.4.1, is shown as installed near the beam-pipe in figure 3.6.

4.7.2 SCT structure and integration

As shown in figure 4.1, the SCT consists of 4 coaxial cylindrical layers in the barrel region and two end-caps each having 9 disk layers [60]. Tables 4.5 and 4.6 of section 4.3 show the SCT detector parameters in detail.

The low-mass barrel cylinders are designed to be extremely stable to both temperature and humidity variations, and to long-term creep [93]. They are made from three-layer (0° , $+60^\circ$, -60°) carbon-fibre skins of $\sim 200\ \mu\text{m}$ total thickness over a carbon fibre/cyanate ester honeycomb core to form a 6 mm sandwich. The cylinder ends are closed with flanges, incorporating holes that are machined to high precision. Pads for the precision mounting of module brackets are attached to the surface and both the surface and a precise mounting hole are machined to within $\pm 20\ \mu\text{m}$ accuracy. A similar precision is specified for the inserts of an alignment system mounted on each barrel, and for the machined holes on the end flanges. Because of poor adhesion for a few pads, all




Figure 4.28: The pixel detector during integration of the barrel, end-caps and their services: (a) the end-cap region; (b) the barrel detector region; (c) Patch Panel 1 (PP1) region; (d) Patch Panel 0 (PP0) region and (e) region of the optical transceivers on the service quarter-panels. See text for details.

the pads were subsequently attached using small stainless steel screws and plastic bushes, with a slight precision loss. The external radius of all cylinders has been maintained to within 1 mm.

The barrel modules [67] are mounted in rows of 12, on individual carbon-fibre brackets, as depicted in figure 4.29. The module is rotated by ± 20 mrad, alternating from barrel to barrel, to align the strips of one side along the cylinder axis. The modules are attached to the support structure at three points, two on the beryllia facing (cooling side) and one on the far side [95]. When mounted on the barrel, the variance of the mounting precision in z as measured by the mounting robot is $60 \mu\text{m}$, but there has been no systematic survey of the variance in R - ϕ . To avoid HV breakdown, a distance of more than 1 mm is maintained between the sensor edges and any ground potential. The centres of adjacent modules in each row are radially separated by 2.8 mm using the tiling arrangement.

The end-cap disks support end-cap modules, as illustrated in figure 4.30, with tight stability and accuracy requirements, together with their electrical, mechanical and alignment services [96]. The 8.7 mm thick disks consist of carbon-fibre face skins ($200 \mu\text{m}$ thick) with an aramid/phenolic honeycomb core. The orientation of the carbon-fibre skins and the choice of materials minimise the effects of thermal and humidity changes. Individual modules are attached to cooling blocks held by inserts glued to the disk. A large cooling block at the hybrid end (230 mm^2) defines the position of the module, while a slot at the end of the module defines the ± 20 mrad rotation of the module when attached to the smaller block (78 mm^2), to within ± 1 mrad. The RMS spread of the surveyed




Figure 4.29: The upper photograph shows a complete SCT barrel with all its modules mounted. A blown-up detail of some of these barrel SCT modules mounted on the support cylinder, together with the module services including the polyimide signal and power cables, and the cooling tubes is also shown (bottom left). A drawing of the mounting brackets, which are attached to the barrel SCT cylinders (in this case, the innermost barrel cylinder), and of the attachment of the module and cooling pipes to the bracket, is also shown (bottom right).

module in-plane placement positions on each disk was $10\text{ }\mu\text{m}$ and the placement uncertainty of the disk within the support cylinder is estimated to be $50\text{--}100\text{ }\mu\text{m}$ ($< 1\text{ mm}$ in the z -direction).

The disks of an end-cap are supported by springs at 12 points around the support cylinder circumference. The springs are soft in the radial direction but otherwise stiff, allowing for radial expansion of the disks and cylinder. The cylinder composition is similar to that of the disks. Services leave the disks through apertures in the cylinders and run along the cylinders before exiting at the far end of the end-cap thermal enclosures. Each support cylinder is in turn supported by two flat composite panels (of similar construction to the cylinder). These panels rest on the same rails as those which support the TRT.

Prior to the mounting of SCT modules on the support structures, each barrel or disk was equipped with electrical services, optical services and cooling loops. In the barrel, each cooling loop, with two inlets and one common exhaust, services four rows of 12 modules. The loops are mounted on the module mounting fixtures and connected to the modules using thermal grease. For




Figure 4.30: End-cap SCT modules mounted on end-cap SCT disk with outer and inner modules (left) and middle modules (right).

the end-caps, each cooling loop is attached to the disk and traverses the module mounting blocks for all the modules within a disk quadrant.

Each service assembly and each cooling loop were tested and all modules of a given barrel or disk were powered and read out to verify the full module functionality [97]. Less than 0.5% of modules needed intervention. For a final test of completed barrels and disks, test systems were constructed to operate and read up to one million channels simultaneously [89, 98]. After assembly, individual barrels were transported to CERN for final integration.

The integration of individual barrels and disks differ due to their different support structures. Both the barrel and end-cap are surrounded by low-mass outer and inner thermal enclosures (see section 4.8.2). Their role is to prevent condensation during operation by maintaining a low-temperature and low-humidity N_2 environment, to prevent the out-flow of N_2 gas surrounding the SCT modules, which would affect the TRT gas-gain and performance, and to prevent the in-flow of the ID environmental gas (CO_2). They also provide a Faraday shield to protect the SCT from external electrical noise. The outer thermal enclosures and the end-surfaces are covered with resistive pad heaters to ensure thermal neutrality.

To complete the SCT barrel from four individual barrels, the outer thermal enclosure was mounted into a support cradle and the four SCT barrels sequentially inserted and fastened together using eight radial interlinks at each end, having dowel pins matching the precisely machined holes in the barrel flanges. The deformation of the cylinders was measured to be below $90\text{ }\mu\text{m}$ (RMS). Within the measurement accuracy of $20\text{ }\mu\text{m}$ (lateral) and $40\text{ }\mu\text{m}$ (along the axis), the four barrel SCT cylinders are concentric with their axes aligned. All services were sealed into slots of the




Figure 4.31: End-view of the TRT barrel structure, during the final attachment of cooling and electrical services.




Figure 4.32: A completed TRT end-cap during the final service integration, showing (from the left) twelve type-A wheels and eight type-B wheels surrounded by their services and supporting rings.

outer thermal enclosure feed-through on the outer circumference of the SCT barrel. Inlet cooling capillaries, outlet exhaust cooling pipes and ground-reference connections were added between individual barrels and thermal-enclosure bulkheads. Finally the outer and inner thermal enclosures were sealed and the air-tightness of the barrel enclosure was checked.

For the end-cap, each assembled disk was inserted into the overall carbon-fibre support cylinder [96]. Polyimide power tapes, optical fibres and cooling pipes were connected at the outer disk circumference, and run along the outside of the cylinder to a patch-panel at the end. Each completed end-cap assembly was fully tested and characterised. After assembly, each of the end-cap cylinders was transported to CERN for final integration and testing.

Finally, the SCT is equipped with a geodetic grid of 842 interferometers to monitor real-time deformations at the scale of a few μm . Distances between nodes attached to the structure are measured simultaneously using frequency-scanning interferometry [62] to a precision of $< 1 \mu\text{m}$. The three-dimensional grid node positions can be reconstructed to better than $5 \mu\text{m}$ in the critical direction [99], thus offering access to short time-scale, low spatial frequency detector deformations, which may otherwise only be weakly constrained with data.

4.7.3 TRT structure and integration

As indicated in section 4.3.3 and figure 4.1, the TRT occupies the outer radial regions of the inner detector. There are three module layers with axially oriented straws in the barrel region [70] and 20 wheels with radial straws in each of the end-cap regions [71]. The active regions of each detector are shown in table 4.8 (see section 4.3.3).

Each of the 96 barrel modules is supported at each end by the barrel support structure, which also provides the required overall module stability. It is a 21 mm thick carbon-fibre disk, machined to a triangular strut array and attached to two thin inner and outer carbon-fibre cylinders (see figure 4.23 and figure 4.31).




Figure 4.33: Insertion of SCT barrel into the TRT barrel. The three module types of the TRT barrel are clearly identified. The SCT outer thermal enclosure is visible, together with the barrel services extending on support frames from each end.

As described in section 4.3.3, each of the TRT end-caps consists of twelve type-A wheels with a straw-layer spacing of 8 mm and of eight type-B wheels with a spacing of 15 mm. The geometry of the lower-density type-B wheels maintains the required number of straws crossed by a particle from the interaction point as well as keeping the material of the active detector approximately constant as a function of η .

Each eight-plane end-cap wheel consists of two basic blocks of four-plane wheels mounted in an inner and outer electrically grounded carbon-fibre ring. The rings and the straws serve as the mechanical support structure of the wheels. Each eight-plane wheel is covered with a thin metal-clad polyimide membrane on each front side connected at the inner radius and providing a signal-return path with defined electronics ground. Openings at the outer radius of the electronic shield allow a path for the CO_2 cooling gas.

As shown in figure 4.32, the type-A and type-B wheels are assembled in two independent groups and the individual wheels are held together using tie-rods between two solid carbon-fibre membranes of 10 mm thickness. The end-membranes of a stack are supported on rails fixed to the ATLAS barrel cryostat. Each group is sealed at the inner radius by a glass-fibre cylinder covered with a thin copper-clad polyimide foil. This cylinder serves as both an electrical Faraday shield and a mechanical link keeping the end-membranes parallel.

The end-cap CO_2 gas circulating between and within sectors is cooled by heat exchangers positioned at the outer ends of the straws between each second eight-plane wheel of type-A and each eight-plane wheel of type-B before flowing to the next wheel of the end-cap. The leak-tightness of both end-caps for the CO_2 flow has been verified. To prevent pressure changes of the cooling gas which might cause mechanical stress and deformations (leading to discharges) of the straws, a set of passive safety valves has been installed on each group of wheels.

Table 4.11: Mechanical position of the axes for the installed barrel and end-cap assemblies with respect to their nominal positions relative to the inner warm vessel. The nominal positions are given in figure 4.1. In addition, for reasons of space, each end-cap has been placed 5 mm further away from the interaction point than its nominal position. All dimensions are in mm.

Barrel	x (side A)	y (side A)	x (side C)	y (side C)
Pixel	−0.3	−1.4	−0.1	−1.0
SCT	+0.50	−1.70	+0.4	−1.2
TRT	+0.2	−1.9	−0.2	−1.7
End-cap (side A)	x at z_{\max}	y at z_{\max}	x at z_{\min}	y at z_{\min}
SCT	+0.4	−1.1	+0.1	−2.2
TRT	+1.0	−2.1	+0.3	−2.0
End-cap (side C)	x at z_{\max}	y at z_{\max}	x at z_{\min}	y at z_{\min}
SCT	−0.3	−2.1	+0.4	−1.3
TRT	−0.2	−1.3	−1.0	−1.0

Supplies for high-voltage, active gas, electronics cooling and the routing of the detector signals to the back-end system are organised in sectors covering 1/32 of all the wheels of one end-cap. Along a given sector, all pipes and cables are brought in cable trays from outside the ATLAS cryostat, to the PPF1 patch-panel and distributed from there to the individual wheels.

4.7.4 Integration and installation of the inner-detector components

The completed SCT barrel and end-caps were finally inserted into the corresponding TRT sub-detectors. The barrel procedure is described here: the end-cap integration followed a similar procedure. The SCT was supported on a cantilever frame (foreground of figure 4.33). The completed TRT was transferred into the final support and lifting frame (background of figure 4.33). The TRT was guided on rails over the SCT. During the movement the mechanical alignment and electrical isolation of the sub-detectors were verified. After insertion, the SCT was positioned on rails inside the TRT with a precision of $\sim 250 \mu\text{m}$. The final survey of the SCT barrel with respect to the TRT barrel shows a displacement of their axes by -0.29 mm (horizontal) and 0.2 mm (vertical) on side A, and -0.55 mm (horizontal) and 0.45 mm (vertical) on side C. The precision of this survey was $\pm 0.15 \text{ mm}$. The barrel positions are in good agreement with the global alignment found in a later cosmic run.

The pixel barrel and end-cap integration into the pixel support tube has been described above. The following elements were finally installed on horizontal rails inside the inner warm vessel of the barrel cryostat: the SCT/TRT barrel, the two SCT/TRT end-caps and finally the pixel package (see figure 4.22). The result of the mechanical survey of the components is shown in table 4.11 with respect to the inner warm vessel on which the rails are fixed. The z -positions of the pixel and barrel assemblies are well-centred. However, the z -positions of the end-caps are displaced away from the interaction point by 4.88 mm (side A) and 5.35 mm (side C) from their nominal positions. A more complete description of the survey strategy and measurements and of the positioning accuracies achieved for the major components of the ATLAS detector is given in section 9.3.2.3.




Figure 4.34: Cross-section of the beryllium vacuum pipe with its various layers for bake-out in situ. Dimensions are in mm.

4.8 Inner-detector environment and cooling services

4.8.1 Beam-pipe interface and operational aspects

The vacuum inner detector (VI) beam-pipe has a total length of 7100 mm, an inner radius of 29 mm and a nominal outer radius of 34.3 mm. The interfaces between it and the other components of the beam-pipe in the ATLAS experiment are described in section 9.8. To reduce the amount of material to an absolute minimum, the VI beam-pipe has been manufactured from beryllium with a thickness of 0.8 mm. Figure 4.34 shows the structure of the completed VI beam-pipe as it was inserted into the pixel detector. Several additional layers of material have been added around the beryllium wall to provide the possibility for bake-out in situ (see section 9.8).

The VI beam-pipe is supported by the pixel detector (see section 4.7.1) and adjusted after insertion to be straight and on the nominal beam axis to within 3 mm. Each end was surveyed, then positioned to this target, and with an optical device developed to straighten the pipe, the mid-supports, internal to the pixel package, were externally adjusted to straighten the beam-pipe between its two end-flanges. It was then surveyed again and the final values recorded show that the beam-pipe is within 1 mm of the centre of the pixel package.

Other than the envelope and mechanical interfaces, the beam-pipe has environmental interfaces to the pixel package. The beam-pipe has an electromagnetic interference shield, namely a 50 μm thick aluminium foil wrapped around its outer radius, which is also shown in figure 4.34.




Figure 4.35: Schematic breakdown of the environmental conditions inside the inner-detector volume: in yellow are shown the TRT volumes, in light blue the SCT volumes, and in green the pixel volume inside the pixel support tube. Also shown is the routing of the cold tubes bringing C_3F_8 coolant to the pixel and SCT volumes, as well as the points where CO_2 gas is flushed into the ID volume (see text).

This foil is electrically terminated to the pixel-system ground, where the pixel gas envelope is sealed to the beam-pipe, at the flange of the PP1 bellows seal at each end of the pixel package. These bellows close both the pixel gas volume and the Faraday cage to the beam-pipe.

As shown in figure 4.34, the beam-pipe has integral bake-out heaters laminated to the outer skin of the beryllium tube inside ~ 4 mm of aerogel insulation. With the beam-pipe at bake-out temperatures of up to 230°C , a heat load of approximately 100 W/m is dissipated into the pixel volume from a surface with a temperature of up to 100°C . The beam-pipe will be baked out only when the pixel detector modules are not operating. During bake-out, it is expected that the complete pixel cooling system would be operational, in order to ensure that no portion of the detector becomes too hot. In addition, there is a dedicated interlock system which will monitor all of the layer-0 module temperatures, as well as cooling-system faults, and will switch off the beam-pipe bake-out heaters in case of any danger to the detector. The cooling of the vertexing layer remains essential to ensure that its modules remain cold, which will prevent damage, particularly after significant exposure to irradiation.

4.8.2 Inner-detector environmental gas

The environmental conditions under which the ID sub-detectors operate are rather different both in terms of temperature and of environmental-gas requirements, as shown in figure 4.35.

The pixel and SCT silicon detectors operate at a low temperature of approximately -7°C to minimise the impact of radiation damage, while the TRT detector operates at room temperature. The cooling systems required to operate these detectors in a safe and stable way at LHC design luminosity are described in section 4.8.3. This section describes the environmental gases required and the technical solution implemented to maintain a stable environment in the overall inner-detector volume.

The environmental gas for the silicon detectors is dry nitrogen at temperatures of between -7°C and 0°C to avoid condensation on the cold surfaces. The nitrogen, chosen because it is inert, has a small flow and does not provide any cooling power. In contrast, the TRT detector requires an overall CO_2 envelope gas acting as a barrier between the active part of the detector and the environment, to prevent contamination from nitrogen and possibly other sources of pollution. As described in sections 4.3.3 and 4.7.3, the TRT barrel modules are ventilated with a flow of $\sim 3 \text{ m}^3/\text{h}$ of CO_2 . The cooling of the end-cap straws requires on the other hand the operation of a dedicated cooling gas system providing CO_2 flows up to $50 \text{ m}^3/\text{h}$.

The different operating temperatures and gases, as well as the need for thermal and gas neutrality between the different sub-detectors, has resulted in the definition of the independent units shown in figure 4.35, with specific thermal enclosures and environmental gas flows wherever required. Details are given for each sub-system in sections 4.7.1, 4.7.2 and 4.7.3. Because of space constraints, the thermal enclosures must achieve thermal neutrality with minimum material. They are covered by pad heaters to maintain the outer temperature of the cold volumes above the dew point, and, over the area facing the TRT, to minimise heat losses and avoid cooling of the warm TRT structures. The humidity and temperature of the gases are constantly monitored as described in section 4.8.4.

To avoid condensation, primarily around the cold pipes of the silicon-detector cooling system, and to minimise the amount of water vapour present in the overall ID volume, the latter is sealed at each end by two end-plates and the gaps between the various ID sub-detectors are flushed with CO_2 .

As shown in figure 4.36, the end-plates consist of a large segmented double-skin aluminium plate, extending up to a radius of 2.5 m, which is sealed around the beam-pipe and at the 32 service feed-throughs located at the outer radius of the cryostat flange. The ID services, exiting from the PPF1 region and from the pixels and the beam-pipe at smaller radius, are routed between the two skins and form the core of this sandwich-type structure.

The inner portion of these end-plates seals the ID up to the LAr cryostat bore diameter of 2.3 m; it consists of a structural cross which is embedded in the core and holds the support nose and the four beam-pipe supports. The nose in itself provides the fixation points for the pixel support tube as well as an independent sealed environment for the pixel patch-panels. Additional local supports have been added between the nose and the cryostat rails to provide the required stiffness of the assembly during pixel insertion. The plates on the back of this portion of the end-plates are equipped with cooling circuits to extract the heat load coming from the pixel services and to preserve the thermal neutrality of the ID environment. The whole surface of the front skin, which seals the cold pixel environment from the ambient air of the cavern, is covered by electrical heater pads to avoid possible condensation.

The front skin of the outer portion of the end-plates, with a diameter between 2.3 m and 5 m, consists of the outer wall of the cryostat flange itself. In contrast, the outer skin is made by several




Figure 4.36: Details of the inner detector end-plate (see text).

panels sealed together for easier assembly. These panels are equipped with dedicated openings for local maintenance access to the evaporative heater and splice-box areas during short access periods to the detector.

The distribution of the CO₂ flushing points inside the inner-detector volume has been optimised to achieve a nominal dew point of -30°C in about three hours with a total CO₂ flushing rate of 5 m³/h. Safety valves are installed at the outer edges of the end-plate gas seal to prevent the differential pressure between the inside of the ID volume and the outside from exceeding ± 0.5 mbar.

4.8.3 Inner-detector cooling

As described in table 4.10 and section 4.5, the cooling systems of the ID [100] must remove up to ~ 85 kW of heat at the LHC design luminosity from the ID volume. The pixel and SCT sensors operate at $\sim -7^{\circ}\text{C}$ while the TRT operates at room temperature. The temperature stability within the ID volume must be maintained within $\pm 2^{\circ}\text{C}$. In addition, the ID power cables must be cooled and maintained at approximately room temperature as they exit the active detector volume. The cooling systems must be robust and have a high reliability over the lifetime of the experiment.

As explained in section 4.3.3, the heat from the TRT straws is removed in the barrel by the same room-temperature C₆F₁₄ coolant as that used for the front-end electronics. For the end-cap wheels, where the geometry is quite different, the heat from the straws is removed by a forced flow of CO₂ gas. The heat exchangers between the TRT end-cap wheels and the front-end electronics




Figure 4.37: Schematic of the evaporative cooling for the barrel SCT. The heaters are located just outside the ID volume and the pressure regulators are located in the cooling racks at the periphery of the ATLAS detector.

boards are cooled by a standard mono-phase cooling system using a room-temperature C_6F_{14} fluorinert coolant. An identical cooling system is also used to cool the pixel, SCT and TRT cables outside the ID volume (and also in part those inside the ID volume), as well as the electronic boards housed in the PP2 crates (see section 4.4.2).

The evaporative cooling system adopted for the pixel and SCT detectors operates in a way similar to those used in the refrigeration industry. Nevertheless, the safe operation of a large number of sensitive and expensive semiconductor detectors and front-end electronics in the difficult and high-radiation environment of the LHC imposes new and non-standard constraints on the cooling system [101]. The particular choice of coolant (C_3F_8 or octafluoropropane) [102] is motivated by the requirement that it be non-flammable, electrically non-conductive, chemically inert in the event of a leak, and very stable against the high radiation doses expected at the LHC. In addition, the absolute pressure of the two-phase fluid as it flows in the cooling structures on the detector is between 1.7 and 6 bar, which is reasonably well matched to the requirement of minimal material for the cooling structures.

While different in detail, the general features of the individual cooling circuits of the pixel, barrel and end-cap SCT are conceptually the same: as an example, figure 4.37 shows the layout of one barrel SCT circuit. Each of the 204 circuits serving the detector has a fixed flow. The fluid is delivered to the detector in liquid form at room temperature, at an absolute pressure between 11 and 14 bar, from four distribution racks situated on the detector platforms in the underground cavern. Each rack serves one quadrant of the inner detector. The delivery pressure is regulated at the distribution racks by pneumatically-controlled pressure regulators. The fluid travels approximately 30 m to the detector via small (4 mm inner diameter) copper tubes.

Inside the detector volume, the liquid enters a small-diameter capillary and the pressure starts to fall. The fluid emerging from the capillary is in two-phase form and starts to boil, lowering the temperature of the thin-wall cooling structures connected to detector modules. The pressure of the

coolant, which in turn determines the evaporation temperature, is set by a pneumatically-controlled back-pressure regulator in the exhaust line. The back-pressure regulator is also mounted on the distribution racks remote from the detector. The operating pressure on the cooling pipework is typically varied between 6.4 bar (absolute pressure for an evaporation temperature of 14°C) for warm commissioning operation to 1.7 bar (absolute pressure for an evaporation temperature of -25°C) for normal operation. The capillary size has been chosen to deliver a given mass flow which depends on the particular detector requirements. Typical mass flows for the different detectors vary between 2.9 g/s for pixel disks to 8.9 g/s for the barrel SCT.

Under normal operation, the fluid in the pipes which emerge from the detectors into the environmentally-controlled inner-detector volume is cold, at a temperature of -25°C. This cold exhaust is used to pre-cool the incoming liquid to maximise the available cooling capacity on the detector structures. This is done in a counter-flow heat exchanger, which is mounted outside the detector structure but inside the cryostat, where the inlet and exhaust pipes are allowed to come into direct contact. Before the residual two-phase fluid is allowed to exit the environmentally-controlled area via the large diameter (12–14 mm inner diameter) pipes, it must be heated above the dew-point of the cavern to avoid condensation on the exhaust pipes as they return to the rack. This is achieved by an in-line heater, where the fluid is boiled and super-heated. The heater element is a hot wire which is in the flow of the C₃F₈ cooling fluid. The power-rating for the heaters varies between 250 W for the pixel disks to 1 kW for the barrel SCT. The 110 V power to the heaters is pulsed (on–off) with a regulated feedback mechanism to maintain a constant 20°C temperature of the exhaust gas. The temperature of the liquid is measured by one of a pair of redundant thermocouples mounted either in the liquid flow or directly on the exhaust pipe. A hardware protection system (see section 4.8.4) protects the heaters against the possibility of over-heating by switching off the supplied power when the temperature measured on the body of the heater exceeds a pre-set value (typically 55°C).

4.8.4 Inner-detector controls, safety and interlocks

The ID-specific detector control system (DCS) operates within the general context of the overall ATLAS DCS, as described in section 8.5. It controls the detector powering and monitors and controls the environmental parameters, in particular the pixel and SCT cooling, and the TRT gas gain. The sub-system-specific environmental monitoring is similar for all three sub-systems and is described briefly below.

The evaporative and mono-phase cooling systems are controlled and monitored by the ID DCS. Critical temperature, pressure and flow parameters are monitored and controlled, including the pressure regulators, the heater temperatures and the current. Additional sensors monitor and control the thermal enclosure heaters, the electronics rack temperatures and voltages, and the cable and patch-panel temperatures.

Within the pixel detector volume, nearly 400 thermistors are installed. They monitor the temperatures of the cooling loops, the service panels and the gas volume. In addition, to protect against moisture, several radiation-tolerant humidity sensors are mounted inside the pixel detector and read out via the pixel DCS. Finally, each detector module and other sensitive units such as the regulator stations and opto-boards are equipped with thermistors, which are connected to the pixel

hard-wired interlock system. Thanks to the high granularity of this system, even small groups of channels can be switched off in case of risky situations. The pixel-detector interlock includes a signal from the beam-conditions monitor system (see section 3.4.1).

For the SCT [103], the environmental monitoring measures 735 temperature and humidity sensors across the detector. Three types of temperatures are measured: sensors are located at the exhaust of each of the 116 cooling loops; sensors are attached to the mechanical structure of the detector to monitor possible deformations due to temperature changes; and sensors are located next to the laser interferometric survey monitoring system to measure the gas temperature inside the detector volume. Radiation-hard humidity probes are installed in various locations and are used to calculate the dew point and hence to avoid condensation on the modules. The monitored cooling-pipe temperatures also trigger the interlock, protecting the SCT modules if the cooling stops. It is implemented in hardware without the use of DCS micro-processors. The custom-built hardware compares the temperature sensor values to a preset threshold and signals the appropriate power supply cards to turn off in ~ 1 second if the threshold is reached.

More than 3000 temperature sensors are distributed in the TRT to monitor the temperature of the active detector, as well as that of the front-end electronics and of the cooling circuits. A hardware interlock system cuts the power delivered to the front-end electronics if NTC thermistors located on the front-end boards detect a temperature above a critical hard-wired threshold. In addition, approximately 200 parameters of the closed-loop active gas system are monitored. The gas gain is continuously measured by a set of reference straws located outside the ID volume, and an automated algorithm adjusts the high voltage on the detector to preserve the stability of the gas gain as the environmental parameters (temperature and/or pressure) vary.

4.9 Performance status of the integrated inner detector

4.9.1 Electrical performance of the integrated detector

The electrical performance of individual pixel, SCT and TRT modules was monitored throughout assembly and some test modules were irradiated to the total dose expected after ten years of LHC operation (see section 4.4). Following delivery to CERN, the SCT barrels and end-caps were fully connected at the surface to the SCT DAQ. The equivalent noise charge, noise occupancy at the nominal 1 fC operating threshold and the number of defective channels were measured and compared with previous data. After the SCT barrels and end-caps were fully integrated into the TRT, connectivity checks were repeated.

Before installation in the ATLAS pit, two opposite sectors of the barrel SCT and barrel TRT were tested [98]. The connected sectors comprised 1/8 of the TRT and 468 modules of the SCT barrels as shown in figure 4.38. A partial test was also made of one end-cap (side C), in which one quadrant (247 modules) of the SCT and 1/16 of the TRT wheels were connected. The tests emulated the final pit configuration, in particular the service routing and detector grounding. One pixel end-cap was also operated under realistic conditions and cosmic-ray data were studied.

The goals of the combined tests for SCT and TRT included the commissioning of the DCS system and the operation of the cooling system, as well as the measurement of the noise performance of the combined detectors under a wide range of operating conditions, to ensure the absence

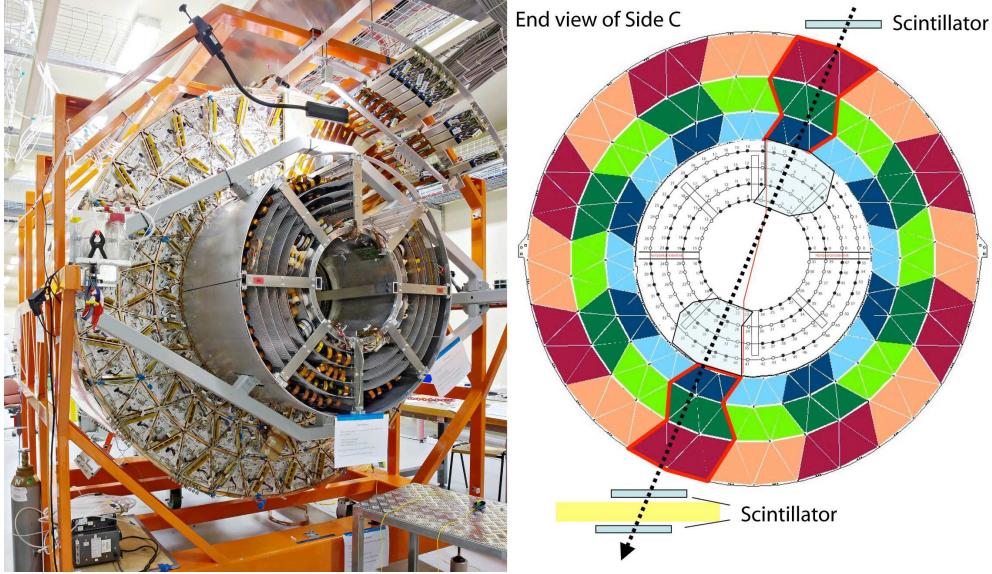


Figure 4.38: Photograph of the ID barrel set-up for cosmic-ray studies (left) and the highlighted configuration of module groups chosen for this test (right).

of cross-talk, to study timing and synchronisation issues and to obtain experience in the combined operation as preparation for commissioning. Cosmic-ray tracks passing through the SCT/TRT barrels and end-caps were also used to study their performance. Table 4.12 shows the rate of defective SCT channels, after full integration with the TRT at the surface. Table 4.13 shows the rate of defective TRT channels after installation in the experiment.

The fraction of defective channels in the barrel and end-cap pixel components after the final integration of the detector in the surface assembly building is 0.33%, as summarised in table 4.14.

The noise levels of the SCT, TRT and pixel sub-detectors were measured in many different configurations from data collected using random triggers.

For the SCT detector, the noise performance and occupancy were compiled for the individual barrels and end-caps. The distribution of measured noise occupancies of all active module sides of the barrel and end-cap cosmic ray runs are shown in figure 4.39 and, at a threshold of 1 fC, has a mean of $< 5 \times 10^{-5}$. The timing of the L1 trigger was optimised with respect to the scintillator trigger by histogramming the hit coincidences between the two sides of each module as a function of the timing offset. The corresponding equivalent noise-charge values extracted from threshold scans are shown in figure 4.40 for all active barrel and end-cap (side C) modules. All numbers are within specifications.

The SCT noise was measured in extreme conditions: varying the trigger rate from 5 Hz to 50 kHz, with and without TRT operation, with the thermal enclosure heater pads on, off and switching between the two states, with several grounding schemes, and also while the TRT was being read out. No increase of noise was observed in any of the tested configurations. Similar tests were also performed for the TRT barrel and the TRT end-caps: the noise was verified before and after the insertion of the SCT, in a configuration in which the TRT analogue ground was connected

Table 4.12: Numbers, types and percentages of defective channels in the SCT after integration with the TRT in the surface assembly building.

SCT	Total channels	Not bonded	Dead	Not reached	Partially bonded	Noisy	Other defects	Total defects
Barrel	3244032	803	3595	1664	400	70	152	0.206%
End-caps	3035136	811	6464	-	352	230	5	0.259%

Table 4.13: Numbers and percentages of defective channels in the TRT after installation in ATLAS.

TRT		Total channels	Dead channels	Total defects
Barrel	Side A	52544	926	1.8%
	Side C	52544	1050	2.0%
End-cap	Side A	122880	2115	1.7%
	Side C	122880	1993	1.6%

Table 4.14: Percentages of defective pixel channels after the final integration of the detector in the surface assembly building.

Pixels	Defective channels (%)			
Barrel	Layer-0	Layer-1	Layer-2	Average
	0.07	0.40	0.29	0.28
End-cap side A	Disk 1	Disk 2	Disk 3	
	0.14	0.23	0.52	0.30
End-cap side C	Disk 1	Disk 2	Disk 3	
	0.12	2.19	0.31	0.87
Pixel average				0.33

to the SCT power return, for different SCT thresholds, and also while the SCT was being readout. The straw noise occupancy was in all cases close to 2% at a threshold of 250 eV.

The pixel detector was not included as part of the combined SCT-TRT barrel cosmic test since it is installed as a complete unit into the ID. However, one pixel end-cap was oriented with the disks in the horizontal plan, such that a reasonable cosmic-ray rate of ~ 1 Hz could be obtained, using a simple scintillator trigger. A prototype of a fraction of the complete internal pixel services was used to bring power and cooling to the end-cap, and to provide optical readout. The external cables, power supplies and readout were close to those used in the final detector. The complete end-cap with its 144 modules was also operated and tested.

The modules in the pixel end-cap were tuned for a threshold of ~ 4000 e . The average noise for each active module was ~ 170 e and the distribution is shown in figure 4.41. The average noise seen in the cosmic ray test was about 10% lower than that measured during individual module testing under similar, but not identical, circumstances. The observed pixel occupancy per BCID (beam crossing ID) included fixed pattern noise from a small fraction of the channels. Most of these hot (noisy) pixels were previously identified during individual module tests using an ^{241}Am source.




Figure 4.39: Distribution of average noise occupancy for all active module sides of the barrel and end-cap SCT (outer or middle end-cap modules on side C), as obtained at 1 fC threshold.




Figure 4.40: Distribution of the noise in electrons, normalised to a temperature of 0°C, for all active modules (or front-end chips) in the barrel and end-cap SCT (side C).




Figure 4.41: Average noise in electrons for each active module in the pixel end-cap cosmic ray test.




Figure 4.42: Pixel occupancy for active modules in one end-cap disk during the pixel end-cap cosmic ray test as described in the text. The value BCID = 5 corresponds to the peak of the cosmic-ray timing distribution and the value BCID = 4,6 corresponds to adjacent time bins. The occupancy for other BCID values represents a measurement of the random pixel noise occupancy.

The pixel occupancy for the active modules in one of the end-cap disks is shown in figure 4.42 after masking of the hot pixels. About 5×10^{-5} of the active pixels were masked for the modules shown. BCID 5 corresponds to the peak of the cosmic ray timing distribution with small tails before (BCID 4) and after (BCID 6). The pixel occupancy per BCID for other BCIDs was found to be $10^{-9} - 10^{-10}$ and is indicative of the random noise occupancy for these operating conditions.

4.9.2 SCT and pixel cooling performance after integration

A satisfactory cooling performance of the pixel and SCT detectors is crucial to their operation. The heat dissipated inside the ID detector volume is at least 40 kW and eventually 50 kW, with a temperature stability requirement within the SCT and pixels of $\sim 2^\circ\text{C}$ (see table 4.10). There is real risk of thermal runaway once the sensor leakage currents become sizable. Data following the detector installation are not yet available. However, data collected following integration on the surface have been analysed.

The performance of the pixel end-cap sectors was measured during the cosmic ray tests. The evaporative cooling was operated with a base temperature (no power) of approximately -25°C . Temperatures on each module were monitored. The average increase in temperature on a module for nominal power (23 W for the six modules on a sector) was $9\text{--}10^\circ\text{C}$ (operation at -16 to -15°C) for the eight sectors on a disk and projected to be $13\text{--}14^\circ\text{C}$ for end-of-lifetime power (about 38W). The barrel pixel system has not been measured in such a realistic way. However, extensive measurements have been made on single and bi-staves. The difference in temperature between the coolant (nominally -25°C) and a module on a bi-stave is expected to be $\sim 15^\circ\text{C}$. Due to some corrosion problems with the aluminium pipes, some staves in layer 2 had to be reworked and the temperature gradient may be as high as $\sim 23^\circ\text{C}$ in these cases.

The various phases of integration have allowed a preliminary study of the SCT cooling performance. For the combined barrel test, a cooling pipe exhaust temperature of approximately 10°C was used. The average temperature on the barrel module hybrids (read by a thermistor on each module side) was $\sim 10^\circ\text{C}$ above the coolant temperature (unpowered) and $\sim 17\text{--}19^\circ\text{C}$ (powered). When averaged over the operating cooling loops, the mean and spread of the hybrid temperature for each module along a barrel cooling loop is shown in figure 4.43. For the end-cap, this temperature difference was slightly higher, $\sim 15^\circ\text{C}$, with respect to the selected exhaust temperature of 1°C .




Figure 4.43: The mean and RMS (indicated by error bars) of all hybrid temperatures for a given module position (from the cooling inlet to the cooling outlet), averaged over all cooling loops in the combined SCT-TRT run. The exhaust cooling pipe temperature was $\sim 10^\circ\text{C}$.

4.10 Material distribution of the inner detector

The performance requirements of the ATLAS inner detector are more stringent than any tracking detector built so far for operation at a hadron collider. The harsh environment and the pile-up from multiple interactions per bunch crossing make a high detector granularity mandatory, with electronics, readout services and cooling within a detector volume that must have good mechanical stability. The overall weight (~ 4.5 tonnes) and material budget of the ID (in terms of radiation




Figure 4.44: Mapping of photon conversions as a function of z and radius, integrated over ϕ , for the ID. The mapping has been made from 500,000 minimum bias events (~ 40 minutes of data-taking at 200 Hz), using $\sim 90,000$ conversion electrons of $p_T > 0.5$ GeV originating from photons from π^0/η decays.

length X_0 and interaction length λ) are therefore much larger than those of previous tracking detectors. The consequences of this are quite serious and are currently the focus of many studies (see section 10.2):

- (a) many electrons lose most of their energy through bremsstrahlung before reaching the electromagnetic calorimeter;
- (b) approximately 40% of photons convert into an electron-positron pair before reaching the LAr cryostat and the electromagnetic calorimeter;
- (c) even in the case of low-energy charged pions, a significant fraction will undergo an inelastic hadronic interaction inside the inner detector volume.

A detailed modelling of the ID material has been implemented in the simulation. Figure 4.44 shows a map of generated photon-conversion vertices in the ID volume, integrated over azimuth, for electrons with $p_T > 0.5$ GeV from photons originating at the primary vertex in minimum-bias events. The sample shown consists of 500,000 events; at a collection rate of 200 Hz, such a data sample can be accumulated in approximately one hour. It should be noted that many structural elements, for example the end-plates of both the barrel and end-cap, are azimuthally discrete and have been modelled faithfully in the simulation. Figures 4.45 and 4.46 show the integrated radiation length, X_0 , and interaction length, λ , traversed by a straight track as a function of $|\eta|$ at the exit of the ID envelope. The most striking feature is the onset of non-active service and structural material




Figure 4.45: Material distribution (X_0 , λ) at the exit of the ID envelope, including the services and thermal enclosures. The distribution is shown as a function of $|\eta|$ and averaged over ϕ . The breakdown indicates the contributions of external services and of individual sub-detectors, including services in their active volume.




Figure 4.46: Material distribution (X_0 , λ) at the exit of the ID envelope, including the services and thermal enclosures. The distribution is shown as a function of $|\eta|$ and averaged over ϕ . The breakdown shows the contributions of different ID components, independent of the sub-detector.

at the interface of the barrel and end-cap regions. This includes cooling connections at the end of the SCT and TRT barrels, TRT electrical connections, and SCT and TRT barrel services extending radially to the cryostat, to the PPB1 patch-panel, and then along the cryostat wall. Another service contribution is from the pixel services at $|\eta| > 2.7$, which leave the detector along the beam-pipe; their extended range in $|\eta|$ can clearly be seen. A large fraction of the service and structural material is external to the active ID envelope, therefore deteriorating the calorimeter resolution but not the tracking performance. Table 4.15 lists the contribution to X_0 as a function of radius for different elements of the ID and for straight tracks at $|\eta| = 0$ and $|\eta| = 1.8$.

The material breakdown is particularly important at small radius. The pixel barrel radiation length for perpendicular incidence is approximately 10.7% for the three pixel layers. This can be broken down into: electronics+bump-bonds (1.4%), sensors (1.1%), hybrids (1%), local support structures with cooling (5.4%), cables (0.3%) and global supports (1.5%). The corresponding number for the SCT barrel layers is 11.8% when averaged over the active area. This amounts to 2.96%

Table 4.15: Integrated radiation length (X_0) from interaction point, estimated as a function of radius R for $|\eta| = 0$ and for $|\eta| = 1.8$, using the detailed description of the inner-detector material implemented in the simulation. The quoted radii are the maximum radii over which the integration is made. The data are averaged over ϕ .

	$ \eta = 0$		$ \eta = 1.8$	
	Radius (mm)	X_0	Radius (mm)	X_0
Exit beam-pipe	36	0.0045	36	0.014
Exit pixel layer-0	57	0.037	57	0.105
Exit pixel layer-2	172	0.108	172	0.442
Entry SCT	253	0.119	253	0.561
Entry TRT	552	0.205	621	0.907
Exit TRT	1081	0.469	907	1.126

per layer, including 1.33% for modules, 1.15% for services and 0.48% for the support structure. The module budget includes the sensors (0.61%), the base-boards and BeO facings (0.20%), and the hybrids with their components and the carbon-carbon bridge (0.32%). The equivalent material per disk is 3.75%: modules (1.81%), services (1.47%) and disk support structure (0.47%). Services include the electrical services, cooling blocks and pipes. Off-barrel or off-disk services, especially at the interface between the barrels and end-caps and in the forward direction for pixel services, and items such as the thermal enclosures and outer support cylinder of the end-cap SCT, are not included.

An attempt has been made to check the assumed material distribution in the ID by making a comparison with the measured weight of different ID sub-detector units, and with a bottom-up inventory of the ID components. Table 4.16 shows the status of that comparison. There is good agreement between the measured and estimated weights of the component sub-detectors and their services. There remains a discrepancy between the measured weights of the two SCT end-caps.

Table 4.16: The measured weight of individual ID sub-detectors (where available) together with bottom-up estimates of the ID weight from component measurements, before and after detector integration/installation. Those items which were not measured are labelled as not available (n/a). The post-integration weight includes all services up to, but not including, the PP1 patch-panels. The services of the installed pixel layers have not been estimated independently for the barrel and end-caps. The measured weights should be compared to the estimated weights before integration except in the case of the TRT end-caps for which the weights were measured only after integration of all services. The quoted estimates from simulations should be compared to the post-integration weights.

Sub-detector	Measured weight (kg)	Estimated weight (kg)		Weight in simulation (kg)
		Pre-integration	Post-integration	
SCT barrel	201 ± 20	187 ± 5	222 ± 6	222
TRT barrel	707 ± 20	677 ± 3	703 ± 3	700
SCT+TRT barrel	883 ± 20	864 ± 6	925 ± 7	922
SCT end-cap A	207 ± 10	174 ± 5	225 ± 10	225
SCT end-cap C	172 ± 10	174 ± 5	225 ± 10	225
TRT end-cap A	1118 ± 12	n/a	1129 ± 10	1131
TRT end-cap C	1120 ± 12	n/a	1129 ± 10	1131
Pixel barrel	n/a	20.2	n/a	18.3
Pixel end-cap A	n/a	7	n/a	6.1
Pixel end-cap C	n/a	7	n/a	6.1
Pixel total	193.5 ± 5	33	201	197