

The Optical Alignment System of the ATLAS Muon Spectrometer Endcaps

C. Amelung^{1,2}, J. Bensinger², A. Dushkin², M. Gardner², K. Hashemi², E. Henry³,
B. Kaplan^{2,7}, P. Keselman^{2,8}, M. Ketchum², U. Landgraf⁴, A. Ostapchuk^{5,9},
J. Rothberg⁶, A. Schricker^{1,2,10}, N. Skvorodnev², H. Wellenstein²

¹*CERN, 1211 Geneva, Switzerland*

²*Brandeis University, Waltham, MA 02454, USA*

³*Harvard University, Cambridge, MA 02138, USA*

⁴*Universität Freiburg, 79104 Freiburg, Germany*

⁵*Max-Planck-Institut für Physik, 80805 München, Germany*

⁶*University of Washington, Seattle, WA 98195, USA*

⁷*now at Yale University, New Haven, CT 06520, USA*

⁸*now at University of California, San Francisco, CA 94122, USA*

⁹*now at RWTH Aachen, 52062 Aachen, Germany*

¹⁰*now at Piezocryst Advanced Sensorics GmbH, 8020 Graz, Austria*

Abstract

The muon spectrometer of the ATLAS detector at the Large Hadron Collider (LHC) at CERN consists of 1182 muon precision chambers, arranged in three concentric cylinders in the barrel region, and in four wheels in each of the two endcaps. The endcap wheels are located between 7 m and 22 m from the interaction point, and have diameters between 13 m and 24 m. Muon chambers are equipped with a complex optical alignment system to monitor their positions and deformations during ATLAS data-taking. We describe the layout of the endcap part of the alignment system and the design and calibration of the optical sensors, as well as the various software components. About 1% of the system has been subjected to performance tests in the H8 beam line at CERN, and results of these tests are discussed. The installation and commissioning of the full system in the ATLAS cavern is well underway, and results from approximately half of the system indicate that we will reach the goal of a 40 μm alignment accuracy, required for reconstructing final-state muons with the required momentum resolution at the highest expected energies.

Contents

1	Introduction	4
1.1	Motivation	4
1.2	Muon Spectrometer Layout	4
1.3	Alignment Strategy	6
2	Alignment Sensors	7
2.1	BCAM-type Sensors	10
2.1.1	BCAMs	10
2.1.2	Chamber Laser Sources	13
2.1.3	3D Sources and 3D Cameras	14
2.1.4	TGC Locator Sources	15
2.2	RASNIK-type Sensors	16
2.2.1	Proximity Cameras and Masks	16
2.2.2	Three-Element RASNIK Sensors	18
2.3	Sensor Electronics	18
2.3.1	LWDAQ System Design	19
2.3.2	Power Consumption	20
2.3.3	Resistance to Radiation	21
3	Alignment System	22
3.1	Chamber Geometry	23
3.1.1	MDT Chamber Deformations	23
3.1.2	MDT Sensor Mount Calibration	26
3.1.2.1	Coarse Calibration by Photogrammetry	26
3.1.2.2	Fine Calibration Using Calibration Tools	27
3.1.3	CSC Chamber Geometry	28
3.2	Reference Grid	29
3.2.1	Grid Layout	29
3.2.2	Alignment Bars	32
3.2.3	Bar Calibration and Shape Reconstruction	34
3.3	Proximity Measurement Systems	35
3.3.1	MDT Proximity System	35
3.3.2	CSC Proximity System	37
3.4	Atmospheric Effects	37
4	Data Handling	39
4.1	LWDAQ Data Acquisition Software	39
4.2	Data Acquisition Control and Data Flow	40
4.2.1	LTX Communications Program	40
4.2.2	PVSS Control Software and Databases	41
4.3	ARAMyS Alignment Reconstruction Software	42
4.3.1	Method of Alignment Reconstruction	43
4.3.2	Simulation of Alignment System Performance	45
4.4	Alignment Data in the Track Reconstruction	46

5	Test and Validation	46
5.1	H8 Test Beam Setup	47
5.1.1	Relative Alignment Tests in H8	47
5.1.1.1	Tests with the Muon Simulator	47
5.1.1.2	Tests with Straight Muon Tracks	48
5.1.2	Absolute Alignment Tests in H8	49
5.1.2.1	Tests Using a Survey	50
5.1.2.2	Tests Using the Muon Beam	50
5.1.3	Conclusions	50
5.2	Installation and Commissioning in ATLAS	51
5.2.1	Chamber and Bar Installation	51
5.2.2	First Commissioning Results	51
5.2.3	Outlook	53
6	Alignment Beyond the Optical System	54
7	Summary	55
	Acknowledgments	55
	References	56

1 Introduction

The ATLAS muon spectrometer [1, 2] is located outside the calorimeter and has been designed to provide a good stand-alone momentum measurement of muons up to the highest expected energies: the transverse momentum p_T should be measured with a resolution of $\Delta p_T/p_T = 10\%$ at $p_T = 1$ TeV. Four different technologies of muon chambers are used, two for precision measurements, two for triggering. The precision measurement of muon tracks is performed in most of the spectrometer by monitored drift tube (MDT) chambers, composed of 6 or 8 layers of cylindrical aluminum drift tubes glued onto a spacer support structure. In the very forward region, where particle densities and backgrounds are the highest, cathode strip chambers (CSCs) are used instead. The MDT and CSC chambers are together referred to as precision chambers. They are complemented by trigger chambers – resistive plate chambers (RPCs) in the barrel, and thin gap chambers (TGCs) in the endcaps.

1.1 Motivation

Muon chambers are placed in an air-core toroidal magnetic field, which has the advantage of causing only slight multiple scattering due to the small amount of material present between chambers. A consequence of this design is the relatively low magnetic field strength that can be reached. For example, the bending of a 1 TeV muon track in the magnetic field is such that the track sagitta varies between 0.5 mm at pseudorapidity¹ $\eta = 0$ (in the barrel) and 1 mm at $\eta = 2$ (in the endcap). Consequently, in order to measure the momentum of a 1 TeV muon to 10% at all angles, the error on the sagitta measurement must be less than 50 μm in the bending direction of the magnetic field, transverse to the MDT tubes and wires. Each track is detected in three almost equally spaced wheels of precision chambers. Thus, the design MDT single-tube resolution of 80 μm (averaged over the tube radius) results in a sagitta error of 40 μm , and the additional error from the alignment of the MDT chambers should not exceed that value. As long-term stability in ATLAS cannot be guaranteed at such small scales, a continuously running alignment monitoring system has been developed. The system is based on optical and temperature sensors, and on alignment bars, which are up to 9.6 m long instrumented aluminum tubes used as precision reference rulers. It is designed to detect slow chamber displacements, occurring at a timescale of hours or more. The information from the alignment system is used in the offline track reconstruction to correct for the misalignment of the precision chambers – no physical adjustments are made to the chamber positions after their initial positioning. This paper is a description of how this is achieved in the endcap region ($1.0 < |\eta| < 2.7$); the barrel system ($-1.0 < \eta < 1.0$) is described separately [3].

1.2 Muon Spectrometer Layout

The precision chambers in each of the two muon spectrometer endcaps are arranged in four wheels, or stations (fig. 1), abbreviated as EI (endcap-inner), EE² (extra), EM (middle), and EO (outer). Muons are detected in three stations: those in the range $1.0 < |\eta| < 1.4$ traverse the EI, EE, and EM wheels, and those in the range $1.4 < |\eta| < 2.7$ the EI, EM, and EO wheels. The CSC chambers cover the range $2.0 < |\eta| < 2.7$ in the EI wheel. Each wheel consists of

¹The pseudorapidity, η , is defined as a function of the polar angle, θ , which is measured with respect to the beam line: $\eta = -\ln(\tan(\theta/2))$.

²Due to budget constraints, the EE chambers will not be part of the initial configuration of ATLAS. Their construction is underway, and installation in ATLAS is foreseen for 2009.

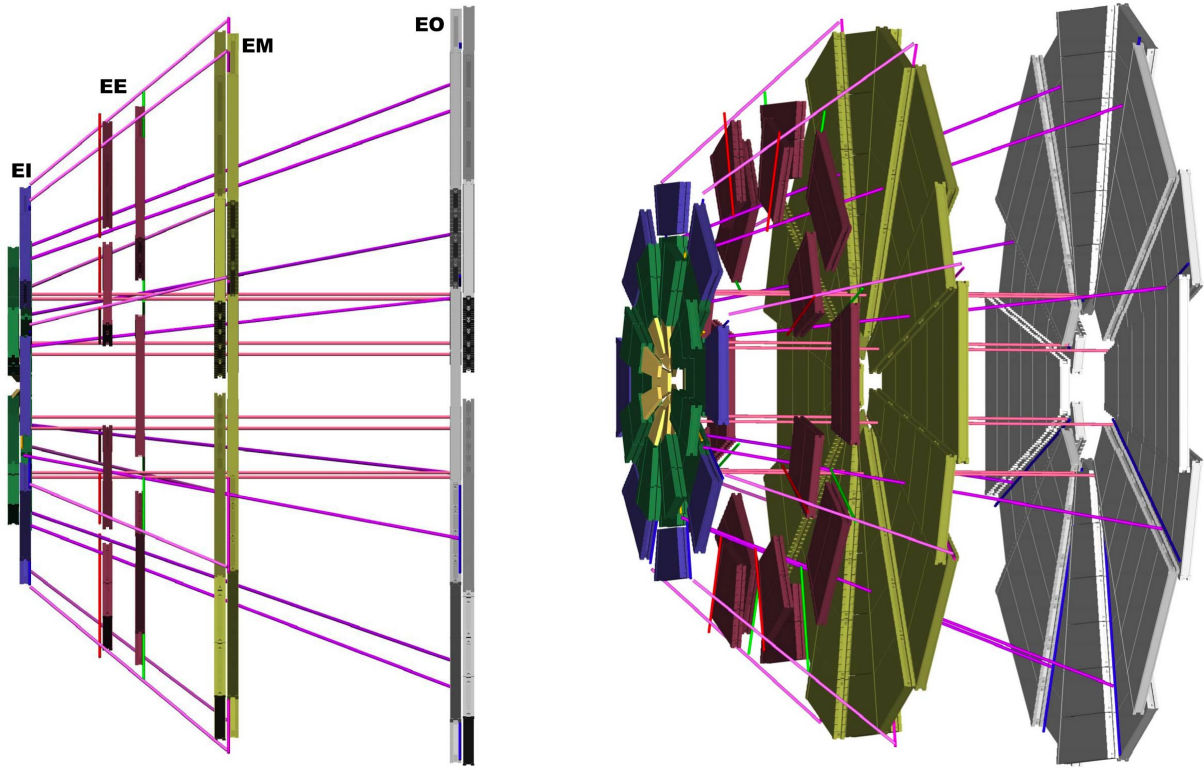


Fig. 1: Layout of the precision chambers in one muon spectrometer endcap. Left: side view; right: isometric projection. The beam line is horizontal, with the interaction point on the left, outside the picture. Shown are the four stations EI (CSCs in yellow, EIL1/2 and EIS1/2 in green, EIL4 in blue), EE (red), EM (brown), and EO (gray). Alignment bars are arranged as the spokes of the wheels (hardly visible). Those alignment lines that link bars in different wheels to each other are shown as well, while the vast majority of alignment lines, linking chambers and bars within a wheel, is not shown. The vessel of the endcap toroid magnet fills the space between the EI and EM stations, up to the inner radius of the EE chambers.

16 sectors (numbered from 1 to 16) of trapezoidal chambers. Those in odd-numbered sectors cover 28° in azimuth and are called large (L) chambers, those in even-numbered sectors cover 17° in azimuth and are called small (S) chambers. Radially, each sector consists of up to six individual MDT chambers, numbered from 1 to 6 with increasing radius, and one CSC chamber in the EI station. Some properties of the MDT chambers are listed in table 1. Mechanically, the EO stations (“end walls”), the EM stations (“Big Wheels”), as well as the inner part of the EI stations (CSC, EIL1/2 and EIS1/2 – the “Small Wheels”), have been assembled in wheel-shaped support structures. The chambers covering the outer part of the EI station (EIL4), and the EE chambers are embedded individually in the barrel toroid magnet structure. In order to minimize the systematic error due to the detector design for certain physics measurements, e.g. forward-backward asymmetries, the two endcaps of the spectrometer (referred to as A-side and C-side) are mirror images of each other.

The EM and EO wheels are approximately 22.6 m and 24.1 m in diameter, respectively. The EI station (Small Wheel and EIL4 together) has a diameter of 12.7 m, of which the Small Wheel alone contributes 9.3 m. The EI, EM, and EO stations are positioned at average distances from the interaction point along the beam line of 7.5 m, 14.1 m, and 21.6 m, respectively. For the typical length scale of the spectrometer endcaps of 16 m, an alignment accuracy of $40 \mu\text{m}$ thus corresponds to $2.5 \cdot 10^{-6}$ (2.5 ppm) of the scale.

name	cham- bers	tube layers	location in $ Z $ (mm)	radial position	tubes/ layer	chamber length in R (mm)	tube width along ϕ (mm)	spacer height (mm)
EIS	32	2×4	7261	1–2	36–42	1096–1276	898–1546	121
EIL	32	2×4	7675	1–2	36	1096	1321–2311	121
	16	2×4	7642	4	42–54	1276–1637	1281–3071	121
EES	32	2×3	10276	1–2	40–48	1216–1457	2014–2734	121
EEL	30	2×3	11322	1–2	40–48	1216–1457	2446–4441	121
EMS	80	2×3	13878	1–5	64	1937	835–3643	170
EML	80	2×3	14294	1–5	56–64	1697–1937	1186–5746	170
EOS	96	2×3	21424	1–6	48–56	1457–1697	1249–3985	170
EOL	96	2×3	21840	1–6	48–56	1457–1697	1681–6241	170

Table 1: Dimensions and positions of the 494 MDT chambers in the endcaps. The Z axis of the ATLAS coordinate system points along the beam direction, X toward the center of the LHC, and Y upward. The radial coordinate is defined as $R = \sqrt{X^2 + Y^2}$, and ϕ is the angle in azimuth measured with respect to the X axis. The location in Z refers to the center of the chamber. The number of tubes per layer may vary between the different radial positions of chambers within a wheel. Two types of small attachments to MDT chambers, referred to as EIL3 (attached to EIL2) and EIL5 (attached to EIL4), are not counted as individual chambers. Two EEL chambers cannot be installed due to space constraints, reducing their total number to 30.

1.3 Alignment Strategy

The basic concept of the muon spectrometer barrel and endcap alignment systems is:

- Chambers have internal alignment sensors that monitor their individual distortions.
- There is a global alignment system that monitors the positions of the chambers with respect to each other.

In the barrel region of the ATLAS muon spectrometer, the correction to the track sagitta that is due to chamber misalignment is, at least conceptually, measured directly. This is done with sensors on projective alignment lines, i.e. lines pointing to the ATLAS interaction point, which link MDT chambers along potential muon trajectories. In the endcap region this is not possible because the endcap toroidal magnet cryostat obstructs most of the projective alignment lines that would be required. Instead, the global alignment system in the endcap establishes a reference grid using alignment bars whose relative positions are measured by a system of bar-to-bar alignment sensors. Chamber positions are then related to the locations of the nearest bars with chamber-to-bar sensors. This information is used to provide corrections to the nominal chamber positions before calculating track sagittae. An overview of the layout of the alignment system in the endcap is shown in fig. 2.

There are two different approaches to locating muon chambers with an alignment system. One approach would be to simply turn on the sensors at some point in time, and use them to track the motion of the chambers from that point on. This provides no information on where the chambers are in space, but it does take the time dependence out of the system, so that the locations of chambers at any time could be determined once their initial positions are known. These initial positions would then have to be determined by some other method, e.g. track-based alignment algorithms. We refer to this approach as the relative alignment concept. A second and more ambitious approach is to use the alignment sensors to directly determine the

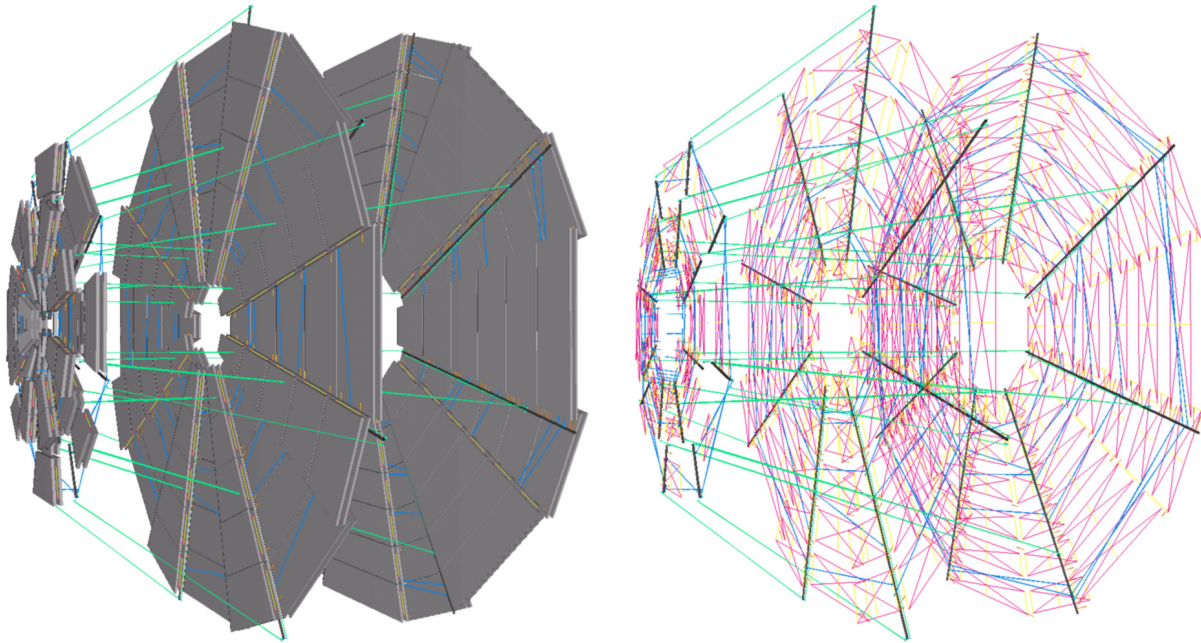


Fig. 2: Layout of the alignment system in one muon spectrometer endcap, with and without the precision chambers visible. Only the EI (left), EM (center), and EO (right) stations are shown, the EE station has been omitted. The several thousands of thin colored lines represent the alignment sensor lines: polar BCAMs (green), azimuthal BCAMs (blue), RASNIK proximity sensors (orange), in-plane RASNIKS (red), chamber temperature sensors (yellow). Shown as thick black lines are the alignment bars, inside which there are in-bar RASNIKS and temperature sensors (not visible). The different types of sensors are discussed in the following sections.

locations of the muon chambers in space with respect to each other (or to use them to measure the sagitta corrections, which is equivalent). This allows the measurement of muon momenta in the spectrometer without requiring any additional method for determining the initial positions of the muon chambers. We refer to this approach as the absolute alignment concept, and it is this approach that we have chosen as the baseline for the muon endcap alignment system. Of course, the system can be used in the relative approach as well, which is considered as a fall-back solution.

Design and implementation of the endcap alignment system and its sensor hardware are described in sections 2 and 3. Software aspects are discussed in section 4. About 1% of the system was set up in the H8 test beam at CERN and the basic operations of the system were validated, as discussed in section 5. At the time of writing, about half of the system has been installed and commissioned in the ATLAS cavern, and first results from this activity are presented in section 5 as well. In section 6 we discuss some aspects of alignment that go beyond the optical alignment system of the endcaps.

2 Alignment Sensors

All the optical sensors used in the ATLAS muon endcap are variations on the 3-point straight-line monitor [4]. In each case, a source of light is imaged through a lens onto a charge-coupled device (CCD) acting as a screen. The source of light is either a RASNIK (Red Alignment System




Fig. 3: Details of CCD images of the two sensor types used in the endcap alignment system: RASNIKS (left) and BCAMs (right). The BCAM image has been magnified with respect to the RASNIK image by a factor of four, in order to make the individual CCD pixels visible.

of NIKHEF) [5] mask, back-illuminated by an array of infrared light-emitting diodes (LEDs), or a pair of BCAM (Brandeis CCD Angular Monitor) [6] laser diodes acting as point-like sources and emitting visible light. We will refer to these two types of sensors as RASNIK-type and BCAM-type sensors, respectively. Typical CCD images of each of them are shown in fig. 3. The CCD and lens may be installed as two separate objects, or joined together in a rigid unit, in which case they form a camera. The various arrangements of CCDs, lenses, and sources of the sensors used in the alignment system are shown in fig. 4. In addition to optical measurements, it is also vital to determine the thermal expansion of the monitored objects. This is achieved by temperature measurements in both chambers and alignment bars³.

Alignment sensors may make two types of measurements, absolute and relative ones. The difference is best understood by the analogy to measuring a muon, where one can ask two questions: (1) where is the muon in space, i.e. where is the track? and (2) what is the muon momentum, i.e. what is the curvature of the track? The first question is about the absolute positions in space of the points on the track, while the second question is about the relative positions to each other of the points on the track, or the straightness of the track.

When a camera views a light source, it observes the image of that source on its own CCD. This enables us to determine the position of the source with a precision given by the uncertainty of the calibration constants of the camera. We refer to this as an absolute measurement made by the camera as it gives our best knowledge of the absolute position of the light source with respect to a reference defined by the camera (its optical axis). When a camera views two or more light sources, it makes an absolute measurement of each of them, but in addition one can ask: how well do the sources line up? The relative positions of the source images on the same CCD give information on the relative positions of the sources to each other. They are typically more accurately measured than the individual absolute positions, because many of the calibration constants of the camera cancel out. We refer to this as a relative measurement made by the camera. Another possibility of making a relative measurement is to separate CCD and lens from each other, in which case the sensor will measure how well lens and light source line up, i.e. it measures their relative positions.

³Depending on the location, different temperature sensors are used: Pt-1000 (1000 Ω Platinum) resistive sensors in all alignment bars, the Small Wheel chambers, and the EML1/EMS1 chambers; TMP27 (Analog Devices) in the EM and EIL4 chambers; NTC type DC95 (Thermometrics) in the EO and EE chambers.




Fig. 4: The various arrangements of CCDs, lenses, and sources in the alignment system sensors: (1) a RASNIK system consisting of CCD (red), lens (blue), and mask (black) as three separate elements; the mask image is focused on the CCD; (2) a RASNIK proximity camera, with CCD and lens integrated together in a solid camera body; (3) a BCAM camera viewing a pair of laser diodes (green); the field of view of the camera is shown by the yellow cone; (4) a 3D camera viewing a 3D source, two pairs of laser diodes at different distances; (5) top view of a pair of BCAMs with laser diodes integrated into the camera body, viewing each other; (6) side view and (7) top view a triplet of BCAMs, each viewing its two partners; the BCAM in the center has two cameras, facing both ways and mounted side by side in the camera body, and is shifted away from the line between the two outer BCAMs; CCDs have been omitted in the side view. Details of the sensors are described in the following sections.

There are errors associated with the absolute and relative measurements made by a sensor. Both of them have a systematic error in the sense that making many measurements instead of just one cannot reduce the error, because the systematic deviation remains constant between measurements. We refer to these as the absolute and relative accuracy of the sensor. The absolute accuracy is typically associated with the sensor, source, and mount calibrations, and thus changes when the sensor or source is replaced or its platform re-calibrated. The same is true for the relative accuracy, but for relative measurements many of the calibration constants cancel out to first order, and thus their uncertainties do not contribute. Therefore, the relative accuracy of a sensor is usually better than the absolute one. In addition, any measurement made by a sensor has an intrinsic resolution, which does not depend on the device calibrations and which causes a deviation that varies between repeated measurements made with the same device. This is typically a statistical error.

2.1 BCAM-type Sensors

The BCAM-type devices use as light source a pair of laser diodes. The image of each laser diode on the CCD is a circular spot, the centroid of which is determined as the intensity-weighted average of all pixels above a certain threshold. Imaging a single laser diode gives the bearing angle of the source, i.e. its angle relative to the optical axis defined by the lens-CCD combination. By imaging a pair of diodes with known separation, the distance to the source and the rotation around the optical axis can be determined in addition.

A BCAM-type device consists of a camera and at least two laser diodes. The latter can either be integrated in one physical unit with another BCAM device (section 2.1.1), or they can be separated, like in the chamber laser sources, 3D sources, and TGC locators (sections 2.1.2–2.1.4).

2.1.1 BCAMs

A BCAM (fig. 5) consists of one (“single-ended”) or two (“double-ended”) cameras, combined with two or four light sources into one solid camera body. The body is an anodized aluminum chassis, made out of a single piece of aluminum for stability and ease of assembly. In operation, the camera in one BCAM looks at light sources on another BCAM, and the camera in the other BCAM looks back at the light sources in the first camera. A BCAM camera consists of a plano-convex lens and a CCD image sensor⁴ at a longitudinal distance of 76 mm. The BCAM light sources are pairs of laser diodes with a transverse separation of 16 mm. The CCD has 344×244 square pixels at $10 \mu\text{m}$ pitch, of which 20×244 at one edge are inactive. The laser diodes are near-ideal point sources, in that their emitting surface is only tens of microns across. With an uncollimated output power of a few milliwatts they are not harmful to the human eye, but they are visible, which makes diagnosis of BCAM problems far easier than with infrared light. The center of the light spot on the CCD can be determined to about $0.4 \mu\text{m}$ ($1/25$ of the pixel width). To improve the resolution, the lens aperture has been chosen such that the spot is slightly blurred by diffraction, so that the image of the laser spreads over several pixels: typically around 10–15 pixels are above threshold. BCAMs are mounted on alignment bars using a 3-ball kinematic mount as shown in fig. 6. The mount is a 3-2-1 or cone-slot-flat mount, consisting of three stainless-steel balls of 6.35 mm diameter on the platform, and three depressions in the BCAM base, in the shape of a cone, a slot, and a flat surface. The mount defines a local coordinate system for the BCAM mounted on it.

⁴We use as CCD the TC255P (Texas Instruments), and as laser diodes the LDP65001E (Lumex) and the DL-3147-021 (Sanyo).




Fig. 5: A black double-ended polar BCAM with the cover removed (there are several different types of BCAMs: “polar” and “azimuthal” BCAMs differ in the height of the chassis, black and blue versions are mirror-images of each other). On the left, the lens holder and lens of one camera are visible, with two laser diodes in holes next to them. The CCD of that camera is on the right. Another camera and another pair of laser diodes are facing the opposite direction. The electronic boards for CCDs and lasers are mounted at the bottom and the front and back walls of the chassis, connected through white flat cables. Below the lens holder there is a socket to connect the CAT-5 cable for the readout system.




Fig. 6: A blue azimuthal BCAM with its sensor platform. The 3-ball kinematic mount defines the local coordinate system for the BCAM. In operation, the platform is clamped onto an alignment bar.




Fig. 7: The calibration stand for BCAMs. The blue BCAM being calibrated is mounted in the roll cage in the foreground. The black BCAM in the middle of the optical table views the laser diodes of the blue one, for their calibration. For the camera calibration, the black BCAM is replaced by the 4-source block shown at the rear.

The calibration of a BCAM is best understood if one views the camera as a point in space, called the pivot point, through which all rays of light hitting the CCD pass. This point is near to, but not exactly at, the center of the lens. Any particular ray hitting the CCD corresponds to a vector at the pivot point pointing toward the source of that ray. In the coordinate system defined by the BCAM mount, the calibration constants for a BCAM are the X , Y , and Z coordinates of the pivot point, the distance from the pivot point to the CCD, the direction cosines of the optical axis (the line connecting the center of the CCD to the pivot point), and the rotation of the CCD about that axis. From this information we can compute the location of the spot on the CCD from any source of light. For calibration, the BCAM is mounted in a roll cage which allows it to be placed in four different orientations about the optical axis of the system (0° , 90° , 180° , and 270°) as shown in fig. 7. The roll cage was measured with a coordinate-measuring machine (CMM), so the relative locations of the 3-ball base for each orientation are known. In each orientation, the BCAM views a block containing four laser diodes with known relative positions, at two different distances from the BCAM. The distance between the two source block positions is measured by an interferometer. From these measurements the calibration constants of the BCAM camera are determined. To calibrate the laser diodes in a BCAM, we mount it in the roll cage and take images of its lasers in the four orientations using another BCAM positioned in front of the roll cage, thus obtaining the transverse positions of the two laser diodes. The coordinate along the optical axis is taken as the center of the lasers by construction.

In a pair of BCAMs, one BCAM measures the absolute angle of the two laser diodes on the other BCAM with respect to its own optical axis with an absolute accuracy of $50 \mu\text{rad}$. It also measures the difference in angles, i.e. the relative angle of those two laser diodes, with a (much




Fig. 8: A chamber laser source with its 3-ball kinematic mount. The mount is made from extruded aluminum, designed to register on the 30 mm diameter aluminum tubes of the MDT chambers. The local coordinate system defining the location of the laser source is also shown. The source itself has two pairs of laser diodes pointing in opposite directions, glued onto a base plate, with the electronics in-between them.

better) relative accuracy of $\sqrt{2} \cdot 20 \mu\text{m}/d$, where d is the distance to the lasers. The intrinsic resolution of each individual spot measurement is $5 \mu\text{rad}$. In a triplet of BCAMs positioned approximately, but not exactly, on a straight line, an additional measurement can be made by each of the two outer BCAMs, by measuring the relative angle of two laser diodes, one on each of the two other BCAMs, with the same relative accuracy and intrinsic resolution as above. This makes the triplet of BCAMs sensitive to deviations from straightness. In practice, the absolute accuracy of the BCAM is limited by the calibration uncertainty of the optical axis inclination, as well as by the accuracy of the CMM measurement of its mount⁵. The relative accuracy is due to the uncertainty of the laser diode positions. The intrinsic resolution is limited by atmospheric refraction and systematic effects of the centroid calculation method. The field of view of a BCAM camera is determined by the dimensions of the active area of the CCD, the CCD-lens distance, and the lens-laser distance, d : about $\pm 20 \text{ mrad} \cdot d \times \pm 15 \text{ mrad} \cdot d$. We refer to this as the dynamic range of the sensor, i.e. the maximum displacement of the light source for which it remains visible to the sensor.

2.1.2 Chamber Laser Sources

Chamber laser sources (fig. 8) are mounted on sensor platforms glued onto the surfaces of MDT chambers. As will be explained in section 3.3.1, these sources serve to constrain a particular degree of freedom of the MDT chambers where a pair of chambers rotates, in a correlated fashion, about an axis defined by the alignment bars. In local jargon these sources are therefore called “saloon-door sources”, because they monitor a behavior where the chambers swing like saloon doors.

⁵Note that in order to measure the inclination of the local coordinate system defined by the mount to $50 \mu\text{rad}$, a measurement accuracy of $2\text{--}3 \mu\text{m}$ for each of the three spheres is required, close to the limit of typical CMMs.




Fig. 9: Left: A 3D source for the CSC chamber alignment. The two pairs of laser diodes are mounted on two different circuit boards, separated by 50 mm along the optical axis and rotated with respect to each other by 90°. They are viewed by a camera from the left. Right: A 3D camera. The light enters from the left through a 150 μm pin-hole, which is covered by a glass microscope slide cover to protect it from dust. The active area of the CCD is shown in pink. The right half of the tube contains the readout electronics.

The chamber sources contain two pairs of laser diodes like the BCAMs, but no cameras. Very few sources contain only one pair of diodes. Nearly all laser sources are mounted on 3-ball kinematic mounts. Only very few of them mount instead, for historical reasons, on V-shaped rail mounts like the ones used for proximity cameras (fig. 11). The calibration of the laser sources starts from the small plates holding the pair of laser diodes. In a method similar to that used for BCAMs, the location of the lasers is determined relative to the edges of the plate. One (“unidirectional sources”) or two (“bi-directional sources”) such plates are then mounted on a source base plate, and their locations relative to the 3-ball mount or rail mount are measured using a CMM. From these measurements the locations of the lasers in the mount coordinate system can be determined.

2.1.3 3D Sources and 3D Cameras

A limitation of the BCAM is its insensitivity to rotations of the laser diode around its center. Even for a pair of diodes, the camera remains insensitive to rotations around the line joining them. For the BCAM case, this limitation is overcome by combining the pair of laser diodes with a second camera, looking back at the first one. Another way to solve the same problem is to use four laser diodes instead of two, and to arrange them as two pairs with a 50 mm separation between them along the optical axis of the camera: the 3D source (fig. 9). The 3D sources are used in the alignment of the CSC chambers. A source is calibrated on a stand with three cameras looking at it from three different angles, with the source placed sequentially in ten different positions and rotations.

As the CSC chambers and the distances between them are small, the distance between the 3D source and its camera also had to be small. As a result, with a lensing system one would have had to deal with lens distortions. For the camera, rather than developing a lens correction program, we elected to use a pin-hole camera, which does not suffer from such distortions: the 3D camera (fig. 9). Its body is an anodized aluminum tube of 195 mm length. The 150 μm




Fig. 10: TGC locator sources are positioned just outside of the polar alignment corridors, visible as the large elliptical cut-out in the center. In this view the sources are shown from the side away from the direction of the laser. Additional small cut-outs in the circuit boards are to allow the laser diodes in other wheels to be visible as well.

pin-hole aperture is placed at a distance of 10–30 mm from the CCD. For the different camera-source distances required, the CCD-aperture distance is adjusted such that the separation of laser images on the CCD is approximately equal to half the width of the CCD. They therefore all have approximately the same dynamic range of ± 8 mm. A 3D camera is calibrated by viewing a calibrated 3D source at different distances from the camera along the optical axis.

The images of the laser diodes on the CCD of a 3D camera are significantly larger than typical BCAM images, with the spots extending over up to hundreds of pixels. Absolute accuracy and intrinsic resolution are comparable to those of the BCAM. The relative accuracy, i.e. the calibration of the relative positions of the four lasers in the 3D sources, is better by a factor of three thanks to the elaborate 3D source calibration method which involves a larger number of measurements than the BCAM and chamber laser source calibrations.

2.1.4 TGC Locator Sources

The endcap muon trigger is provided by seven layers of TGCs in each endcap, arranged in three wheels of TGC chambers surrounding the EM wheel of MDT chambers. The TGCs have their own internal alignment system to determine chamber positions within each wheel in the two coordinates transverse to the beam. The MDT alignment system is supposed to determine the positions of each of the TGC wheels with respect to the MDTs. For this purpose, use is made of the existing cut-outs in the TGCs for the polar alignment lines (section 3.2.1), and two pairs of laser diodes (fig. 10) are mounted on the TGCs on each low- η and mid- η polar corridor. Each of these sources is visible to at least one polar BCAM of the MDT alignment system, so that the location of the TGC plane can be determined. The laser circuit boards were designed to fit precisely in the holes machined into the TGC chambers. The overall resolution locates the TGCs with respect to the MDTs with a precision of about 1 mm, matching the position resolution of the trigger chambers of a few millimeters.

2.2 RASNIK-type Sensors

A RASNIK mask is a glass plate with a chessboard pattern printed on its front surface. The CCD dimensions are of the order of millimeters, while the mask has a size of the order of centimeters. Consequently, as the lens is usually placed somewhere halfway between mask and CCD, the image projected onto the CCD corresponds to only a small fraction of the mask. In order to determine from the incomplete image its location on the mask, the chessboard pattern is modified in every ninth column and row (the “code squares”) to encode coarse position information. Fine position information is obtained from the large number of black-white transitions in the image. Masks with square sizes in the range 85–340 μm are used. Due to the large number of features on the masks, the analysis of RASNIK images is highly redundant, and thus rather insensitive to image deteriorations like, e.g., pieces of dust or even hairs on the CCD or the mask. By de-coding the coarse position information in several adjacent columns or rows of code squares, their consistency can be checked. By selecting only a part of the image to be analyzed, the effect of image distortions can be minimized further if needed.

RASNIK instruments have the advantage of a large dynamic range, selectable by the size of the mask, and an intrinsic resolution of a few microns transverse to the optical axis which is independent of the dynamic range. For an equally spaced CCD-lens-mask system, a transverse motion of the lens relative to the line between mask and CCD of 1 μm results in a 2 μm motion of the image on the CCD. Longitudinal position information comes from the measured magnification of the image, and the position resolution is typically of the order of 10^{-4} times the distance between CCD and mask. If the lens is rigidly fixed to the CCD, then the device works as a proximity measurement between the two parts. A rotation of the mask relative to the CCD about the optical axis can also be measured, with a resolution of better than 1 mrad.

RASNIK-type devices used in the endcap alignment system are either RASNIK sensors with separate elements, where the mask, lens, and CCD are mounted independently and used as a 3-point straightness monitor (section 2.2.2), or RASNIK proximity sensors, with the lens and CCD fixed together (section 2.2.1). For proximity sensors, the CCD-mask distance is 0.25–0.85 m, while RASNIKS with separate elements have been successfully operated up to about 10 m CCD-mask distance (the longest actually used distance in the ATLAS endcap is 6.2 m).

2.2.1 Proximity Cameras and Masks

Proximity cameras are between 180 mm and 270 mm long anodized aluminum tubes of 30 mm diameter containing an infrared filter, the lens, the CCD, and the electronics of the camera (fig. 11). They are attached to MDT chambers on rail mounts glued to the chamber surface. Proximity masks contain, in an anodized aluminum housing, a 30 mm \times 30 mm RASNIK mask, an array of infrared LEDs, and the electronics to back-illuminate the mask⁶. They sit on 3-ball kinematic mounts which are either clamped to alignment bars or glued to MDT chambers (fig. 12). Proximity cameras are tailored individually to the required distance of camera and mask mount for a given location (between 150 mm and 600 mm), by adjusting the positions of the lens and the CCD such that the mask is in focus and the magnification is about $m = 0.3$.

The lens and CCD in the camera have a fixed relationship, and so the magnification is a direct measure of their distance from the mask. By measuring the magnification to a precision of 10^{-4} , longitudinal displacements can be measured to better than 120 μm . When mounted, the proximity cameras are oriented with their axes along the MDT wire direction, perpendicular

⁶We use as infrared LEDs the HSDL-4400 (Hewlett-Packard).




Fig. 11: Left: A proximity camera. The light passes through the 2–3 mm aperture integrated into the cylindrical lens holder. Another (larger) aperture is placed in front of the IR filter at the front of the camera tube on the left, to prevent stray light from getting in. The active area of the CCD is shown in pink. Right: The rail mount for the proximity camera is an extruded aluminum V-profile, designed to register on the MDT tubes. There are two pins defining the transverse (X and Y) coordinates of the mount, respectively, and the profile itself defines the longitudinal (Z) coordinate of the mount, as indicated by the precision cylinder above the rail.




Fig. 12: A bar-mounted proximity mask with its platform and bar clamp. The 3-ball kinematic mount defines the local coordinate system for the mask. There are also chamber-mounted proximity masks, which sit on a 3-ball mount like the one shown in fig. 8 and have a slightly different housing and base plate.

to the bending direction for muons. The precision requirement for the chamber positions is only about $250\ \mu\text{m}$ in this direction. In the coordinates transverse to the optical axis, proximity cameras and masks are designed to provide an absolute accuracy of $20\ \mu\text{m}$ and $50\ \mu\text{rad}$ for position and rotation. The calibration of proximity cameras is performed by taking images from a proximity mask at several different distances along the optical axis, over as large a range as possible (fig. 13). Proximity masks are calibrated by comparing camera measurements of the masks to those of a calibrated master-mask, and by CMM measurements of the surface of the




Fig. 13: The calibration stand for the proximity sensors. This stand is used for calibrating both cameras and masks. It consists of a rail mount for the camera, and about thirty 3-ball mounts at different distances. Masks are mounted on a heavy brass interface-plate, which can be set on any of the 3-ball mounts.

RASNIK glass plate. The relationship between mechanical (CMM) and optical (on the RASNIK mask) coordinates of the master mask is established by subsequently taking camera images of the RASNIK pattern on the mask and of the stylus of the CMM positioned in front of the mask.

An important requirement is that the proximity sensor must remain in focus over the expected displacement range of the camera with respect to the mask of 10–20 mm along the optical axis. This is achieved by placing a 2–3 mm aperture in front of the lens, yielding a depth of field of several centimeters. Transverse to the optical axis, the dynamic range is ± 15 mm, given by the size of the RASNIK mask.

2.2.2 Three-Element RASNIK Sensors

Each endcap MDT chamber is equipped internally with four RASNIK sensors consisting of single elements, embedded in the structure supporting the MDT tubes. Alignment bars are equipped with four RASNIK sensors also, mounted inside the bar tube. The CCDs, masks, and electronics are essentially the same as used in proximity cameras and masks. Details of these systems, their layout, and their calibration are discussed later, in sections 3.1.1 and 3.2.2.

2.3 Sensor Electronics

Alignment devices are controlled and read out with the LWDAQ (Long Wire Data Acquisition) system [7]. The LWDAQ can retrieve up to three images per second from a camera at the other end of a 130 m cable, and is designed to be robust, convenient, and inexpensive. It was designed for use in the ATLAS endcap alignment system, but it has since been used in several other experiments, such as ALICE and CMS, and in other parts of ATLAS.

LWDAQ device	elements	type	
in-plane CCD	1	CCD	
in-plane mask	1	LED array	
thermometer	8	Pt-1000 sensors	
in-bar sensors	}	2	CCDs
		2	LED arrays
		3 or 4	Pt-1000 sensors
proximity camera	1	CCD	
proximity mask	1	LED array	
BCAM	}	1 or 2	CCDs
		2 or 4	laser diodes
chamber source	2 or 4	laser diodes	
3D camera	1	CCD	
3D source	4	laser diodes	

Table 2: Types of LWDAQ devices used in the alignment system, with the number and type of device elements contained in each of them.

2.3.1 LWDAQ System Design

The phrase “long-wire” refers to the CAT-5 cables used to connect LWDAQ devices to their multiplexers and drivers. These cables can be up to 130 m long. A device can be connected directly to a driver with a single cable, or it can be connected through a multiplexer, along with other devices. The multiplexer is then connected to the driver with a single cable. The single cable running from the driver to the multiplexer provides both power and signals to the multiplexer and all the devices connected to it, with no power source other than the driver required. All LWDAQ cables are interchangeable, and all eight wires take part in the connection: four are used for analog and digital power, and four for two-way communication. A device may have one or several device elements (table 2), and the driver can control all of these elements either sequentially or simultaneously. Only one element at a time, however, can return information to the driver, and the returned data may be digital or analog. Controlling several device elements through one device may be seen as a second level of multiplexing in the system.

In ATLAS, multiplexers are connected to drivers by custom-designed solid-conductor cables (“root cables”). These cables are almost identical to standard CAT-5 cables, except that only four of their conductors are twisted into two twisted pairs. The remaining four conductors are not twisted, and are used for power. These cables meet the solid-conductor CAT-5 specification on the two twisted pairs, and their diameter is reduced compared to a cable with four twisted pairs. Devices are connected to multiplexers with stranded-conductor CAT-5 cables (“branch cables”). The CAT-5 specification for stranded-wire cables is ten times less strict when it comes to signal dispersion than it is for solid-wire cables. The disparity in the specification results from the greater difficulty in controlling the dielectric properties experienced by a signal traveling down a stranded wire, and from the desire to make stranded cables flexible and narrow so they can be useful for short connections. All root and branch cables are shielded and halogen-free.

The ATLAS LWDAQ drivers reside in VME crates in the service cavern of the experiment. Each driver provides eight RJ-45 sockets for root cables. In ATLAS, all devices are connected through multiplexers to the drivers, never directly (fig. 14). The ATLAS detector requires about 800 ten-slot multiplexers, with 800 CAT-5 cables running out of the detector and into the service cavern, where 100 drivers receive them. In most places, active LWDAQ repeaters are inserted into the root cables between the drivers and the multiplexers. The repeaters improve the quality of the outgoing digital control signals, thus increasing the maximum cable length to about 180 m.

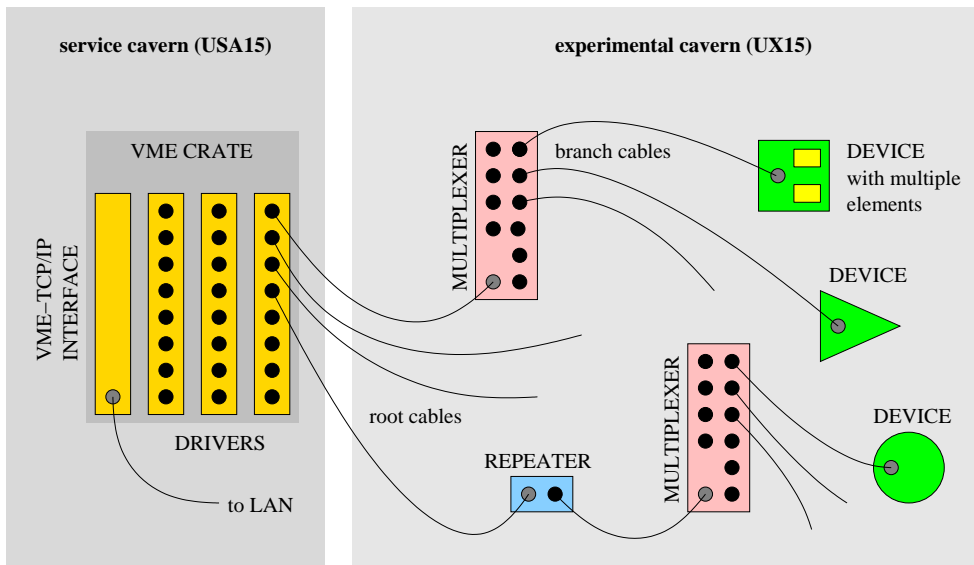


Fig. 14: Cabling scheme of the LWDAQ system. Root cables connect the LWDAQ multiplexers (installed on chambers and bars in the experimental cavern UX15) to the LWDAQ drivers (residing in VME crates in the service cavern USA15). Optionally, an LWDAQ repeater may be inserted in the root cable. The different LWDAQ devices connect through branch cables to the multiplexers, and may contain one or multiple device elements.

At the same time, the repeaters act as patch panels and make it possible to turn the power to the devices off by a software command.

The endcap alignment system acquires two main types of data: temperatures and images. To acquire temperatures, the LWDAQ driver instructs the thermometer device to return a sequence of analog voltages, which it digitizes and stores in memory. To acquire images, the driver must control both the source of light and the image sensor. It clears the CCD of residual charge, flashes the light source, reads the image pixels out of the CCD one by one, digitizes them, and stores them in memory. Table 3 lists the total number of different types of alignment devices and the images and temperatures acquired by them.

2.3.2 Power Consumption

The LWDAQ supplies power to up to ten devices through a single root cable. There are three power supplies delivered by the cable: ± 15 V and $+5$ V. All devices can operate with as little as ± 13 V in place of their ± 15 V supplies. The resistance of the solid-wire conductors in a 130 m root cable is roughly 10Ω . An active CCD sensor consumes 50 mA from ± 15 V, and a flashing RASNIK mask consumes 100 mA. If mask and CCD are connected to the same multiplexer, sharing the power delivered by the same root cable, their ± 15 V power will drop to ± 13.5 V, and both circuits will still function correctly. The LWDAQ operates sequentially, i.e. there will never be more than two devices involved in data acquisition at any one moment in time.

All LWDAQ devices are either “active” or “asleep”. When active, they can consume over 100 mA from the ± 15 V supplies, and up to 20 mA from the $+5$ V supply. When asleep, they consume no more than 5 mA from the ± 15 V supplies, and no more than 5 mA from the $+5$ V supply. At any given point in a data acquisition cycle, almost all devices will be asleep. With eight sleeping devices and two active devices on one multiplexer, the $+5$ V power consumption will be at most

device class	type	physical	logical	function
in-plane	RASNIK	1984	1984	MDT deformation
MDT temperature	TEMP	96	768	MDT expansion
in-bar	RASNIK	} 176	352	bar deformation
bar temperature	TEMP		608	bar expansion
radial	BCAM	96	256	bar deformation
polar	BCAM	208	1856	bar-bar link
azimuthal	BCAM	736	1472	bar-bar link
proximity	RASNIK	2384	1192	MDT-bar and MDT-MDT link
laser source	BCAM	584	2208	MDT-bar link
3D sensor	BCAM	192	384	CSC-bar and CSC-CSC link
CSC temperature	TEMP	16	96	CSC electronics monitoring
total		6648	11176	

Table 3: Classes of alignment devices in the endcap system (temperature sensors on MDT chambers other than the Small Wheel and EML1/EMS1 ones are not read out by the LWDAQ system and have been omitted). The “physical devices” column lists the number of LWDAQ devices; the “logical devices” column lists the number of acquired sensor images or temperatures, respectively. The device classes and their functions are described in section 3.

80 mA. At the end of a 130 m cable, the +5 V supply will drop to 3.4 V (the cable drops 0.8 V in the +5 V wire and 0.8 V in the return wire). All devices can operate with the +5 V supply as low as 3.1 V. At the time of installation in ATLAS, devices consume only between 1 mA and 2.5 mA from +5 V when asleep, but from irradiation experiments (see below) we expect an increase in their sleeping current consumption as the devices suffer damage from ionizing radiation. This increase is of the order of 0.5 mA/krad.

2.3.3 Resistance to Radiation

A selection of devices was exposed to high-rate ionizing radiation from a gamma-ray source, where they received a dose between 1 krad and 100 krad during a few hours. We performed some experiments with circuits connected to a power supply, and others with circuits simply placed in bags or aluminum boxes. Band-gap reference circuits were found to be vulnerable to doses of more than 20 krad, and CMOS analog switches, if exposed to radiation while under power, failed at around 30 krad. The remaining circuits, including CCDs, laser diodes, and LEDs, continued functioning up to and above 50 krad, although logic chips began to consume more power starting at around 10 krad. We removed band-gap references from all our devices, and allowed analog switches only when they would be under power for less than 1% of the time.

The effect of 10 krad delivered in a few hours is different from the effect of 10 krad delivered in ten years. We observed recovery from initial damage in the first few weeks after rapid irradiation, which could be accelerated by annealing the circuits at 60 °C. Circuits under power might suffer permanent damage by excessive current consumption in a rapid radiation exposure, while circuits exposed without power would be working perfectly a week later. We concluded that the best way to judge the resistance of a circuit was to expose them without electrical power, wait for a month, and test them. After that, they were annealed at 60 °C for a month and tested again. If they passed the first test, the initial damage by the dose was tolerable, and if they passed the second test, the long-term effects of the damage were tolerable also.

According to simulations [8], the maximum ionizing radiation dose the circuits will be exposed to in ATLAS is approximately 7 krad. In most places, the dose is far less. The irradiation

experiments suggest that the LWDAQ devices will consume no more than 5 mA after receiving 20 krad of ionizing radiation over the course of a year or longer. Even if the simulations underestimate the radiation dose by a factor of three, all devices will be fully operational after ten years in ATLAS. If the simulations are wrong by a factor of ten, approximately 5% of the circuits may begin to draw excessive current after the first few years of ATLAS, and may have to be replaced.

The LWDAQ devices must also endure neutron radiation. According to simulations [8], the highest neutron dose in the endcap is roughly 10^{12} neutrons per square centimeter (1 Tn/cm^2). Neutron irradiation experiments on all LWDAQ circuits and components were performed with doses up to 10 Tn/cm^2 . The only components that suffered damage were the CCDs and the infra-red LEDs. Both of them are critical parts of the system, and are often in places where they are even more difficult to replace than the device electronics that controls them. We therefore designed the LWDAQ system to endure ten times the maximum simulated neutron dose, not merely three times.

The CCDs suffer an increase in dark current from neutrons [9]. After absorbing 10 Tn/cm^2 , their pixels are completely filled up by dark current in 50 ms at 20°C (compared to 4000 ms before irradiation), and thus each image must be captured in less than 50 ms. Allowing 10 ms for image exposure leaves 40 ms for readout. The LWDAQ has a readout rate of two million pixels per second, and thus it can transfer the about 80,000 pixels of a CCD image in just 40 ms. The LEDs, on the other hand, lose roughly 90% of their optical output power after 10 Tn/cm^2 . The effect of neutron radiation upon the CCD already requires the system to be able to operate with exposure times of no more than 10 ms. If an LED is operating at 10% of its rated output power, and still is to provide a good image in 10 ms, this means that it must be able to provide a good image in 1 ms when it is operating at 100% power. Thus we have ensured that a flash time of 1 ms or less is adequate in the places we expect to receive the highest dose of neutron radiation, before radiation damage begins. Another consequence is that the LWDAQ must be able to generate such short flashes of light in a reliable way, which means it can turn on and off light sources, and switch between one device and another, in a fraction of 1 ms.

As part of the strategy to increase the radiation resistance of the LWDAQ, as much functionality as possible is located outside the high-radiation areas of the detector, in the LWDAQ drivers. For example, all timing signals required by a device are provided by its driver, with the exception of the short pulses required to decode the serial transmissions from the driver. Devices do not digitize analog signals, but transmit them as analog signals directly back to the driver. To preserve the integrity of these analog voltages, devices transmit them as low-voltage differential signals (LVDS), and all ATLAS-resident LWDAQ cables are shielded. Likewise, the driver transmits its commands as LVDS logic levels. The fast data readout and rapid response time of the LWDAQ increase its tolerance to radiation damage further, and at the same time they decrease the time taken to read out all devices in the system, thus achieving the design goal of reading out the entire endcap alignment system in 15–20 minutes.

3 Alignment System

The design of the alignment system can be seen as a hierarchy with three levels:

- At the first level, the shape, i.e. deformation and expansion, of each individual chamber is determined from the measurements of the RASNIKS and temperature sensors built into it (section 3.1).

- At the second level, the global alignment system establishes the precise grid of alignment bars by determining their individual shapes and their positions with respect to each other using the measurements of the BCAMs mounted on them (section 3.2).
- At the third level, the positions of MDT chambers, measured in pairs, with respect to the nearest alignment bars are derived from the measurements of proximity cameras on the chambers looking at masks on bars and on the other chamber of the pair, as well as of BCAMs on bars looking at laser sources on the chambers (section 3.3.1).
- Also at the third level, the positions of all CSC chambers with respect to each other and to the alignment bars are determined from measurements of 3D cameras on the chambers looking at 3D sources on other chambers and on bars (section 3.3.2).

3.1 Chamber Geometry

For a high-precision measurement of muons, there are several requirements to be fulfilled by the precision chambers themselves: the initial internal geometry of the chambers has to be accurately known; deviations from the initial geometry have to be monitored throughout the lifetime of the experiment; and the locations of alignment sensors with respect to the sensitive detector elements have to be accurately known.

3.1.1 MDT Chamber Deformations

An MDT chamber consists of two sets of 3 or 4 layers of 30 mm diameter single-wire drift tubes each, referred to as the two multi-layers. They are separated by a spacer support structure, or “spacer frame”. The lengths of the tubes vary between 0.9 m and 6.2 m, and each layer contains between 36 and 64 tubes. The spacer structure consists of two “long beams” oriented along the MDT tubes, and three “cross plates” approximately orthogonal to the tubes, following the trapezoidal shape of the chambers (fig. 15). The cross plates are referred to as central, high-voltage, and readout cross plates, matching the corresponding sides of the MDT chamber. The chamber construction method guarantees the placement of the drift tubes such that all cathode wires are positioned with a precision of better than $20\ \mu\text{m}$ r.m.s. with respect to the nominal positions. For about 15% of all MDT chambers this has been explicitly verified using an X-ray tomograph. However, after removal from the construction tooling, the expected chamber distortions, particularly for the largest chambers, are much greater than $20\ \mu\text{m}$, and cannot be neglected.

Therefore, there are in-plane alignment sensors embedded in the spacer frame to monitor these distortions from the birth of the chamber throughout its lifetime. (The phrase “monitored” in “monitored drift tube chambers” is a reference to the in-plane alignment system.) The system consists of four RASNIKS: two RASNIK lines run parallel to the MDT tubes, two go diagonally across (fig. 16). They share CCDs and masks, so that the total number of elements is two CCDs, four lenses, and two masks. Lenses are attached to the central cross plate, CCDs to the high-voltage one, and masks to the readout cross plate. The RASNIKS are read out (“zeroed”) while the chamber is still in the construction jig, and thereby calibrated at this point. Another calibration parameter, which is not determined by this procedure, is the rotation of the mask columns and rows about the optical axis of the RASNIK. For deformations below 1 mm, assuming this parameter as zero is adequate given the assembly precision of about 15 mrad r.m.s. The in-plane system monitors the mechanical deformations of the spacer frame, and thus of the tube layers, and the thermal expansion is provided by temperature sensors built into the




Fig. 15: An MDT chamber, viewed from two different angles, showing the top multi-layer of MDT tubes (gray) separated from the rest of the chamber. The spacer frame (green) contains the in-plane system. The RASNIK optical lines of the system are shown as red lines. The detail figures show: a proximity camera on a rail mount, and a survey target platform (top left); the housing of the in-plane CCD electronics, and a temperature sensor with readout electronics, on the high-voltage side of the chamber (center left); a magnetic-field sensor on its platform (bottom left); a proximity mask on a 3-ball mount (top right); a chamber laser source on a 3-ball mount (center right); the in-plane mask electronics, and another temperature sensor, on the readout side of the chamber (bottom right).

chambers. Each chamber has between 8 and 28 temperature sensors, some are glued onto the MDT tubes, others onto the cross plates and long beams.

The following set of deformation and expansion modes has been foreseen in the reconstruction programs (some parameters may have values that are sufficiently small to be ignored, and some are relevant only for the alignment system, not for track reconstruction):

- b_z – a bow of the MDT tubes in the chamber plane;
- b_p and b_n – a bow of the tubes out of the plane, varying from the short side to the long side;
- s_p and s_n – a bow of the cross plates out of the plane, varying from the readout side to the high-voltage side;
- t_w – a twist of the chamber, i.e. a rotation in opposite directions of the two outer cross plates around the tube direction;
- p_g – a parallelogram-like deformation, i.e. a rotation in the same direction of all three cross plates in the plane;




Fig. 16: Schematic representation of the in-plane RASNIK layout. Two CCDs are attached to the high-voltage cross-plate, two masks to the readout cross-plate, and four lenses to the central cross-plate. Two RASNIK lines are parallel to the long and short sides, two go diagonally across. CCDs and masks are rotated to bisect the angle made by the two RASNIK lines they share.

- t_r – a trapezoid-like deformation, i.e. a rotation in opposite directions of the two outer cross plates in the plane;
- e_g – global expansion;
- e_p and e_n – local expansions, different for the readout and high-voltage sides.

All bow-type deformations are assumed to be parabolic, which is a sufficiently accurate approximation. Typical observed magnitudes are as large as a few hundred microns for some of the parameters, particularly for b_p/b_n and t_w , the out-of-plane bow of the tubes and the twist. Others appear to be immeasurably small, like s_p/s_n , the out-of-plane bow of the cross plates.

Each chamber has four “tension rods” along the two long beams of the spacer frame. By tightening some of them after chamber installation, the value of the out-of-plane bow of the tubes can be tuned, separately for the long and short sides. By doing this, the bows can be easily reduced to within $\pm 100 \mu\text{m}$. This is desirable because a non-concentricity of the MDT wire due to a tube bow distorts the electric field, and without correction a bow of $150 \mu\text{m}$ would cause an apparent shift of the hits by up to $20 \mu\text{m}$. The second deformation mode with large magnitudes, the twist, cannot be changed by adjusting the tension rods.

The deformation modes p_g and t_r can take non-zero values only due to a particular detail of the construction of the endcap MDT spacer frame: a thin web (“flexo”) decouples a small motion of the two MDT multi-layers along the tube axis from the spacer structure. This allows for parallelogram-like or trapezoid-like deformations of the spacer frame which are not constrained by the much stiffer MDT tube layers, and thus for two deformation modes that are seen by the in-plane system but do not affect the tube layers. Consequently, these modes are ignored both for the global alignment reconstruction and for tracking. A trapezoid-like deformation could also be caused by a temperature difference between the long and short sides of the chamber.

By definition, all parameters are zero when the chamber is still in the construction jig. The only exception are most of the EO chambers, which have been constructed with a deliberately built-in bow of the tubes in the plane, i.e. a non-zero value of the parameter b_z . This bow was

chosen for each individual chamber to follow the expected gravitational sag of the MDT wire, in order to minimize the non-concentricity of the wire.

3.1.2 MDT Sensor Mount Calibration

For the sensors that link the chambers to the global alignment system, the positions with respect to the MDT wires need to be accurately known. On one of the outer tube layers, there are four alignment sensor mounts near the corners of each chamber. In addition there are four photogrammetry target stickers on plates (see the next section for an explanation of photogrammetry and the use of these targets), as well as platforms for mounting magnetic-field sensors. There are two types of alignment sensor platforms: rail mounts and 3-ball mounts.

While MDT tubes are glued together to form MDT chambers with a precision better than $20\ \mu\text{m}$, alignment sensor platforms are glued onto the tubes with an accuracy of only a few hundred microns and a few milliradians. In addition the gluing technique allows for human error, the most common one being to accidentally shift a platform by one tube ($30\ \text{mm}$) sideways. The same is true for the photogrammetry target sticker and magnetic-field sensor platforms. For the alignment system to do its job, the spatial relationship of the sensors to the tubes (and thus the wires) must be known to a much better accuracy, typically better than $20\ \mu\text{m}$ and $50\ \mu\text{rad}$ – except for the coordinate along the tubes, where the requirement is only about $250\ \mu\text{m}$. The requirement for the survey target plates is more relaxed ($200\ \mu\text{m}$), as is the one for the magnetic-field sensor platforms ($1\ \text{mm}$ and $1\ \text{mrad}$).

The MDT chambers therefore need to be calibrated. This is done in two steps:

- A coarse measurement of all three position coordinates is performed by photogrammetry for all types of platforms (section 3.1.2.1).
- A fine measurement of the two coordinates transverse to the tubes is made for 3-ball and rail platforms by using two dedicated calibration tools (section 3.1.2.2).

As the fine measurement is a re-measurement of quantities already determined to a lesser accuracy by the coarse measurement, the two can be used to cross-check each other.

3.1.2.1 Coarse Calibration by Photogrammetry Photogrammetry [10] is a surveying technique where photos of retro-reflective coded, i.e. numbered, targets are taken from many different angles and positions with a digital camera and a flash. The required number of photos ranges from about 20 to several hundreds, depending on many factors like the size of the setup and the number and size of the targets. A dedicated software⁷ analyzes the photos, and determines by a global numerical minimization algorithm the location from which each photo was taken, the internal calibration parameters of the camera, and the positions of all the targets in an (arbitrary) coordinate system. Photogrammetry is often used by the CERN survey group for optical surveys of large installations, like wheels of MDT chambers, where the positions of the survey target stickers on the chamber surfaces are measured. It can be combined with theodolite measurements, e.g. to link the photogrammetry coordinate system to survey marks in the area.

For the purpose of calibrating MDT chambers, adapters were designed that can be mounted on each type of alignment sensor and magnetic-field sensor platform, with cylindrical precision holes for one or two photogrammetry targets. Survey target platforms do not require any adapters,

⁷We use the AICON 3D studio software.




Fig. 17: Photo of an MDT chamber during the photogrammetry measurement. Five coded targets on a cross at the center define the photogrammetry coordinate system. Scale bars with two targets at their ends, positioned in various places, constrain the length scale. Adapters for two targets each are mounted on the sensor platforms near the four corners of the chamber. This type of chamber, EML2, has four rail mounts. The four single targets on small plates near the long and short sides are the survey targets. Magnetic-field sensor platforms are missing.

as the stickers on them *are* photogrammetry targets. The adapters were calibrated using a CMM, to precisely relate the positions of the photogrammetry targets to the platform below. Aluminum scale bars with lengths between 0.5 m and 1.5 m, carrying one survey target sticker at each end, are used to constrain the length scale of the photogrammetric measurement. Aluminum was deliberately chosen as the material for the scale bars, so that they would thermally expand in the same way as the MDT tubes and spacer frame. This ensured that the result of any photogrammetry measurement, performed at temperatures ranging from 15 °C to 25 °C at different times, would be independent of the room temperature, without any corrections needed. The scale bars were calibrated using an optical CMM at 20 °C, the standard temperature for chamber production. With about 25 photos per chamber (fig. 17), an accuracy of better than 50–150 μm in all three coordinates is achieved, roughly scaling with the chamber size.

3.1.2.2 Fine Calibration Using Calibration Tools As photogrammetry is insufficient to reach the required accuracy for rail and 3-ball platforms, two dedicated measurement tools based on RASNIK sensors were designed for their calibration (fig. 18). The two tools share the principle of operation, and differ only in their appearance which is adapted to the specific shape of the platform type. Each tool is calibrated (zero-ed) by reading out the RASNIKS on a dedicated calibration stand, consisting of two aluminum cylinders as a mock-up of MDT tubes, and a platform in-between. The cylinders and the platform were measured with a CMM.

The repeatability of measurements on the calibration stand and on MDT chambers is very good, about 2 μm r.m.s. in the individual RASNIK measurements. By performing CMM-controlled movements of the sensor platform with respect to the fixed cylinders, using a modified version of the calibration stand, we found that the tool can follow those movements to better than 10 $\mu\text{m}/40 \mu\text{rad}$ r.m.s. Photogrammetry and platform calibration data are analyzed together,




Fig. 18: The calibration tool for rail mounts. It consists of an inner brass part that sits, in the same way as an alignment sensor later does, on the rail mount, and an outer aluminum part that sits on the MDT tubes near the mount, and picks up their position and orientation. Four RASNIKs, with CCDs and lenses mounted on the inner part, and masks on the outer part, determine the location of the inner part (representing that of the mount) with respect to the outer part (representing the MDT tubes, i.e. the chamber coordinate system).

transforming the photogrammetry measurement from its arbitrary coordinate system to the MDT chamber local coordinate system. This also estimates the accuracy of the photogrammetry measurement by comparing it to the calibration tool measurements, and has helped in the calibration campaign to spot a small number of gross mistakes due, e.g., to improper seating of the calibration tools. In such cases, the measurements were repeated. Typically, photogrammetry measurements and platform calibrations take about 45 minutes each per chamber.

A limitation of this calibration technique is its use of the tube surfaces as a reference, rather than of the MDT wires, which was impossible. Obviously, the tubes need to be straight to a sufficient approximation. The specification on the tube straightness given to the manufacturer was rather loose, at the level of $100\ \mu\text{m}/\text{m}$, but the observed straightness in the finished chambers is better by one order of magnitude. The suitability of the tube surfaces as a high-precision reference has been investigated by performing consistency checks with data collected during ATLAS commissioning (section 5.2.2). We conclude that the target accuracy is reached within a factor of 1–1.5, and that therefore the tube surfaces are just barely accurate enough.

3.1.3 CSC Chamber Geometry

The large and small CSC chambers have four and two rail platforms, respectively, and two 3-ball platforms each (fig. 19). These are similar but not identical to those on MDT chambers, the rails being slightly shorter and the spheres of the ball platforms being embedded into a large aluminum plate covering about the width of the chamber. Unlike MDT chambers, CSC chambers do not have any high-precision surface. The only precision references are four cylindrical pins of 6.35 mm diameter penetrating the chambers. CSC chambers are calibrated by measuring the sensor platforms with respect to the precision pins directly with a CMM.




Fig. 19: A small CSC chamber mounted on the CMM for calibration. Near the two corners at the front there are two cylindrical adapters temporarily mounted on the rail mounts, taking the places of 3D cameras to be mounted there eventually; two 3-ball mounts for 3D sources are embedded in the aluminum plate on the top.

The layers composing a CSC chamber are mounted on a rigid steel structure, and are not expected to deform. They are therefore not equipped with a system to monitor deformations. They do have temperature sensors, but these do not measure the temperature of the active detector area, but rather that of the readout electronics and the cooling water circuits.

3.2 Reference Grid

The backbone of the global alignment system is a set of precision reference rulers, called alignment bars. These bars establish a precise grid in space, relative to which the positions of the precision chambers can be measured in a second step. The positions of alignment bars with respect to each other are determined by sensors mounted on the bars, looking at each other. The positions of the precision chambers with respect to the bars are determined by sensors on the chambers, looking at light sources on the bars or on neighboring chambers, and vice versa.

3.2.1 Grid Layout

The positions of alignment bars with respect to each other are determined by two classes of alignment sensors:

- pairs of azimuthal BCAMs on bars in neighboring sectors, within a wheel;
- triplets of polar BCAMs on bars in the same sector, in different wheels.

There is a third related class, triplets of radial BCAMs, all of which are mounted on the same bar. These are not sensitive to the positions of bars, but rather to bar deformations, like the in-

bar RASNIKS (section 3.2.2). Having a precision inferior to the in-bar RASNIKS, they primarily serve as a (very useful) debugging tool. In each endcap, there are six types of alignment bars:

- EI bars in the Small Wheels;
- EIL4 bars in the EIL4 region;
- EEL and EES bars in the EE region;
- EM bars in the Big Wheels;
- EO bars in the EO wheels.

Alignment bars are located approximately parallel, and next to, one of the sides of the large MDT chambers, overlapping with the small MDT chambers. There are eight bars of each type in one endcap, placed in even-numbered sectors, and thus a total of 48 bars per endcap.

Alignment bars in different wheels are linked by polar lines (fig. 20). A polar line consists of a triplet of polar BCAMs, each one on a different bar, arranged approximately on a straight line, in such a way that each of the three BCAMs can monitor the positions of the two others. The polar BCAM in the middle is therefore always double-ended, with cameras facing both ways. There are three polar corridors in each sector, referred to as high- η , mid- η , and low- η , according to their location in $|\eta|$. The latter two are at least approximately projective, at angles of about 30° and 45° from the beam line, respectively. The former one, however, is parallel to the beam line, passing through one out of eight empty tubes in the endcap toroid magnet cryostat at a radius of about 3m, and thus is far from being projective. On the high- η corridor there are three polar BCAMs, on the EI, EM, and EO bars, and thus one triplet EI-EM-EO; on the mid- η corridor there are six BCAMs in total, forming the triplets EI-EM-EO, EIL4-EEL-EM, EIL4-EES-EM; and on the low- η corridor there are four BCAMs, with triplets EIL4-EEL-EM and EIL4-EES-EM. Within those triplets, each pair of bars is linked by two pairs of polar BCAMs. These pairs determine the absolute bar positions with respect to each other, in addition to the relative position measurements made by the triplets.

In-between each pair of adjacent alignment bars within a wheel there are as many azimuthal lines (fig. 20) as there are large-small MDT chamber pairs (two per chamber in EIL4, EEL, and EES, however), and thus always two or more. At least one azimuthal line in-between each pair of bars is not parallel to the others. An azimuthal line consists of a pair of azimuthal BCAMs, one on each bar, where each one monitors the position of its partner. Two azimuthal lines are sufficient to determine the positions of the two bars with respect to each other. Where there are more than two lines, they are eventually needed for the proximity measurement (section 3.3.1), but just provide redundant information here.

By considering all eight bars in a wheel, and eventually all 48 bars in one endcap, it becomes clear that the grid system is highly redundant, and thus it deteriorates only mildly in case of single BCAMs failing. With this layout of the grid, two types of measurements can be made:

- the positions of any two (not necessarily adjacent) alignment bars with respect to each other can be determined to about $300\ \mu\text{m}$ in all three coordinates;
- along polar lines, the position of one alignment bar with respect to the straight line defined by two others can be determined to about $30\ \mu\text{m}$ in the two transverse coordinates.

If polar lines were exactly projective, as in the barrel, the second feature would be sufficient for a $40\ \mu\text{m}$ alignment system. As this was not possible in the endcap due to the geometry of the endcap toroid cryostat, the first feature of the alignment system is equally important in order

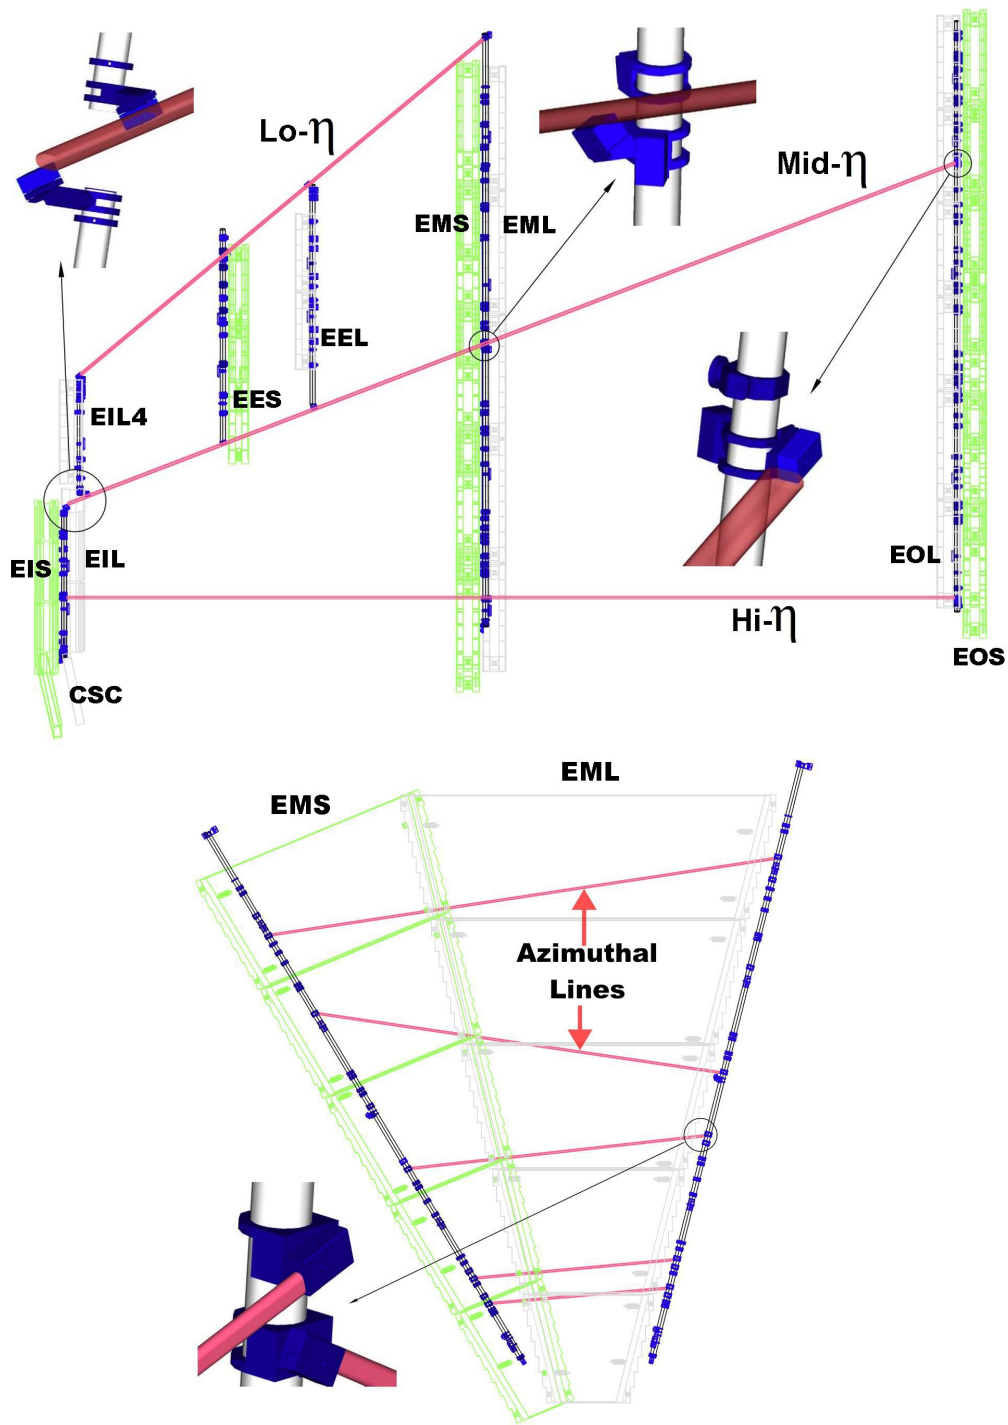


Fig. 20: Layout of the polar (top) and azimuthal (bottom) BCAM lines between alignment bars. The figure at the top shows the bars (black) of one sector (from left to right: EI, EIL4, EES, EEL, EM, EO), as well as the adjacent large (gray) and small (green) chambers. Polar BCAMs (blue) are mounted on the bars in three polar corridors (red); from top to bottom: low- η , mid- η , high- η . Details of some BCAMs on the mid- η corridor are also shown. All EMS and some TGC chambers have cut-outs for the high- η and mid- η lines. The figure at the bottom shows two neighboring bars in the EM wheel, as well as two sectors of large and small chambers. Azimuthal BCAMs on five azimuthal lines link the bars. The detail shows one of them; the second one in the detail views a partner BCAM on the other neighbor of that bar, which is not shown.

bar type	length (m)	ID (mm)	OD (mm)	chambers
EI	2.4	72	80	CSC, EIL 1-2, EIS 1-2
EIL4	1.9	72	80	EIL 4
EEL	3.7	72	80	EEL 1-2
EES	3.3	72	80	EES 1-2
EM	9.6	77	85	EML 1-5, EMS 1-5
EO	9.6	77	85	EOL 1-6, EOS 1-6

Table 4: Types of alignment bars. The table lists the length, the inner and outer diameter of the bar tube, and the types of adjacent chambers linked to each bar by the proximity system. The long bars have a larger diameter than the short ones to increase their rigidity.

to reach the $40\ \mu\text{m}$ specification. In addition to alignment purposes, the knowledge of absolute chamber positions at a level of about $500\ \mu\text{m}$ is required for efficient pattern recognition at high background rates, and at a level of 1 mm for reconstructing the magnetic field using the measurements of hall probes on the chambers.

3.2.2 Alignment Bars

Alignment bars [11, 12, 13] are large precision objects, consisting of up to 9.6 m long instrumented aluminum tubes (fig. 21). Early on, a design based on carbon-fiber bars had been considered, but was discarded because of their dimensional changes with humidity. Another option, quartz bars, was rejected because they are fragile and hard to machine.

On the outside of the alignment bar tube, platforms with 3-ball mounts are clamped onto the bar, on which sensors (BCAMs, proximity masks, and 3D sources) can be mounted. There also are rings with precision holes for up to six standard survey targets. On the inside, RASNIK elements and temperature sensors are mounted on disks and attached to the bar tube. The layout of the in-bar RASNIK lines is shown in fig. 22, and the dimensions of the different types of bars are given in table 4. The bars are supported by two kinematic mounts (a “gimbal mount” and a “slip-ring mount”) near their Bessel points to minimize shape compliance.

The purpose of an alignment bar is to establish a well-controlled spatial relationship between all the sensors mounted on it: their positions on the bar should be known to better than $20\ \mu\text{m}$, their rotations to better than $50\ \mu\text{rad}$. While the short EI and EIL4 bars can be regarded in practice as rigid bodies, the elastic deformations of the long EM and EO bars cannot be neglected. Thus, a strategy similar to the one for MDT chambers is used to obtain the bar shapes: the thermal expansion of the bars is monitored by up to eight temperature sensors, deformations are monitored by four in-bar RASNIKS.

As the in-bar RASNIKS provide information about the bar shape only at a few points along the bar, namely those points where CCDs, lenses, and masks are attached to the bar tube, an interpolation procedure is required for points in-between these points. It turns out that this is a complex problem, as, for example, interpolation by arbitrary functions like splines or polynomials produces errors far beyond the specification. For the largest expected deformation of about 4 mm, the $20\ \mu\text{m}$ requirement corresponds to an interpolation accuracy as low as 0.5%. The method that has been finally adopted makes intelligent use of the analytic solution of the differential equations describing the bar deformation as a function of discrete and continuous forces, and will be explained in the next section. (For MDT chambers, where deformations are smaller by 1–2 orders of magnitude, the quality of the interpolation is much less of an issue.)




Fig. 21: An EM alignment bar. Clamped onto the bar tube are platforms with sensors mounted on them, as well as the kinematic mounts. Detail views show, from bottom to top: a polar and a radial BCAM; the gimbal mount; an azimuthal BCAM; two proximity masks; the slip-ring mount; and a ring for inserting survey targets, together with an LWDAQ multiplexer and two proximity masks on a special platform. Attached to the inside of the tube are disks holding components of the in-bar RASNIK system. Detail views show, from bottom to top: two CCDs; a lens; a lens and two masks; and a disk without optical elements. The remaining four views are the same as the first four, in reverse order. The eight in-bar temperature sensors are attached to these disks.




Fig. 22: Schematic representation of the in-bar RASNIK layouts. Top: the layout of the bar prototypes used in the H8 test, with three RASNIKS (one long, two short ones). Bottom: the layout used in all ATLAS bars except the EIL4 ones, with four RASNIKS (two short, two medium-length ones). This layout has several advantages over the prototype: easier calibration with the CMM, larger dynamic range for bar deformations, redundancy due to the added RASNIK line. The EIL4 bars have a layout with two RASNIKS, resembling the H8 layout with one of the short RASNIKS removed.

3.2.3 Bar Calibration and Shape Reconstruction

Alignment bars are assembled with a moderate accuracy of about 1 mm and 1 mrad, and thus have to be calibrated. Due to their size, this requires an unusually large CMM: the one that was used can measure points in a volume of $6.5 \text{ m} \times 1.4 \text{ m} \times 1.2 \text{ m}$. A local (relative) precision of $3 \mu\text{m}$ and a global (absolute) precision of $10 \mu\text{m}$ have been achieved despite the large measuring volume, thanks to the temperature in the climatized room housing the CMM being stabilized at $22.2 \pm 0.1 \text{ }^\circ\text{C}$, and an elaborate procedure to re-calibrate the machine beyond what was provided by the manufacturer. This re-calibration was done by measuring corrections for imperfections of the CMM mechanics in a finer grid, and resulted in an improvement of the machine accuracy by a factor of three.

For the calibration, the positions of all the spheres of the 3-ball mounts are measured with the CMM. Together with the (sufficiently constant) in-bar RASNIK and temperature sensor readings during the measurement, they are referred to as the initial shape of the bar. The exact shape of the actual bar *tube* is neither measured nor used, as it is not needed for this procedure. The RASNIKS have to be calibrated as well, i.e. the rotations of the mask columns and rows relative to the bar coordinate system have to be determined. This is done by a number of CMM-controlled deformations of the bar, including readout of the RASNIKS in each shape. Each bar is measured with four different deformations, induced by rotating the bar around the bar axis by 180° or by attaching weights of a few kilograms. This provides sufficient redundancy to detect a faulty measurement or other problems. Since the length of the EM and EO bars is greater than that of the longest edge of the measurement volume of the CMM, they have to be measured in two portions, each covering about $2/3$ of the bar. This results in a total of eight required measurements, each with about 100 sphere points, and takes about one week of working time for each of these bars. The initial bar shape is constructed by performing a best fit

to all the measurements, joining together the measurements of single portions. As the triggering force of the CMM probe head of 0.07 N would deform the EM and EO bars significantly during the measurement process, and in addition excite vibrations, these bars are supported at several additional points during the measurement. These supports are carefully adjusted in such a way that no extra deformation is induced.

For reconstructing the bar shape in ATLAS, the difference of RASNIK and temperature values with regard to the initial shape measurement is used. All known changes of forces acting on the bar compared to the initial state are taken into account, e.g. the weight of mounted sensors or changes in the direction of gravity when the bar is inclined. These forces are used to calculate the expected deformation of the bar. Any remaining unaccounted-for forces and effects are absorbed into three effective forces, assumed to be acting at the endpoints of the bar and at the center. Examples are the weight and tension of cables connecting to sensors, and uncertainties of certain material parameters, such as the Young's modulus and the geometrical moment of inertia of the bar tube. The effective forces are determined using the RASNIK measurements, and serve as a measured correction to the expected deformation. The quality of the interpolation depends on this correction being small, because a systematic error is made in absorbing unaccounted-for effects into effective forces. This error is roughly proportional to the magnitude of the correction to the expected deformation. A lot of care is taken to avoid or reduce any unaccounted-for forces. For example, any contact is avoided between sensor cables on the bar and other objects (MDT chambers, or the large support structures of the wheels), and the cables connecting the bar multiplexers to the readout system are routed along the bar to one of the kinematic mounts, and only there they are attached to the support structures. Experimentally, the unaccounted-for forces are indeed found to be small, of the order of a few Newton for the long EM and EO bars. The measured correction to the expected deformation is about 200 μm in r.m.s. for the long bars, and closer to 20 μm r.m.s. for the shorter EI and EIL4 bars. The measured total deformation, dominated by the change in gravitational sag due to the bar inclination, is one order of magnitude greater in both cases. Dedicated deformation studies using the CMM have shown that the systematic error of the interpolation procedure is of the order of 5% of the correction, and thus about 0.5% of the total deformation, well in line with the specification.

3.3 Proximity Measurement Systems

The MDT proximity measurement system links the two MDT chambers in a large-small chamber pair to each other and to the nearest alignment bars. The CSC proximity system links all 16 CSC chambers in one Small Wheel to each other and to the eight alignment bars in that wheel.

3.3.1 MDT Proximity System

The connection between the 247 MDT chambers in one endcap and the reference grid is primarily established by proximity sensors. Typically, a large and a small MDT chamber are linked to each other and to the two adjacent alignment bars via a network of sensors as shown in fig. 23. Proximity cameras on the two outer edges of the chambers in the pair view proximity masks on the alignment bars. In addition, there is one camera-mask pair where the camera is on the large chamber and the mask on the small chamber. Finally, on the remaining corner of the large and the small chamber there is one chamber laser source each that is viewed by azimuthal BCAMs on the alignment bars. These sources are the reason why the number of azimuthal lines matches in each wheel at least the number of chamber pairs. This design locates both chambers in the pair to about 30 μm and 100 μrad with respect to the adjacent bars.




Fig. 23: Layout of the MDT proximity system in a typical pair of sectors. Shown in the center and on the right is a large-small MDT chamber pair, plus another small chamber to the left. Alignment bars overlap with the small chambers. Detail views show the three components of the proximity system: two proximity cameras on each chamber viewing masks on bars (top left); one proximity camera on each large chamber viewing a mask on the small chamber (top right); a pair of azimuthal BCAMs on the bars viewing one chamber laser source on each of the two chambers (bottom). The azimuthal BCAMs are, at the same time, part of the reference grid system, and the azimuthal line passes through the 10 mm gap in-between the two laser sources. All the sensors of the proximity system are located in the space between large and small chamber layers in a wheel.

The layout is different in the EIL4 and EE regions, where there are two proximity cameras on one side of a chamber, viewing masks on the alignment bar, and two chamber laser sources on the other side, viewed by azimuthal BCAMs. Due to the absence of a neighboring chamber (there are no EIS4 chambers, and EEL and EES chambers have a large separation along the beam line), a chamber-chamber proximity sensor could not be used here. In two exceptional cases, namely the EIL4 sectors 7 and 15, there are no proximity cameras at all on the chambers, and instead four laser sources on the four corners. For these chambers, the distance to any alignment bar in the direction along the MDT tubes is too large for proximity sensors to work.

For all layouts, chamber laser sources are bi-directional only to increase the redundancy of the system, not to improve its accuracy. It is sufficient if only one of the azimuthal BCAMs sees

the pair of laser diodes on its side of the chamber source. This is useful when chambers and bars are misaligned by relatively large amounts: while it is relatively easy to move the partner BCAM and the laser sources of a large-small chamber pair into the field of view of *one* azimuthal BCAM (usually the one at the greater distance from the chamber sources), it is often difficult to achieve the same in addition for the partner azimuthal BCAM. Unidirectional sources are used in places where it is a priori clear that this would be even impossible. The redundancy of the system in this aspect is crucial: the sagitta accuracy for a chamber whose laser sources are not seen by any of the azimuthal BCAMs deteriorates dramatically (section 4.3.2). This is called the “saloon-door effect”, because it is due to insufficient constraints on a correlated rotation of the chambers in a pair around the axes of the nearest alignment bars, similar to the way saloon doors swing [14]. Like in the polar BCAM case, each of the azimuthal BCAMs makes both an absolute measurement of the positions of its partner and of the laser source, and a relative measurement of one laser diode on the chamber source and one on the partner BCAM. The latter measurement is the only relevant one in this case.

Two types of endcap MDT chambers have small attachments, which are usually not regarded as separate chambers. These are referred to as EIL3 (attached to EIL2 at larger radii) and EIL5 (attached to EIL4 in sectors 1 and 9 at smaller radii). The attachments do not have alignment sensors of their own, and while the EIL3 chambers are glued to their EIL2 counterparts in a well-defined spatial relationship, the EIL5 chambers are only loosely attached in an ad-hoc fashion. It is thus impossible to derive their absolute position from that of the corresponding EIL4 chamber, but there is at least hope that changes in position can be tracked.

3.3.2 CSC Proximity System

The design of the proximity system for the CSC chambers differs from the one for the MDTs. CSC chambers were chosen for the area near the beam pipe in the EI wheel because of the high track densities and radiation levels in this region. As a consequence of the latter, all the sensors for the alignment system had to be placed at the outer edges of the CSC chambers, where radiation is tolerable. Proximity cameras and masks, like the ones used for MDTs, would not have provided a sufficiently accurate and redundant handle on the chamber positions in such a configuration, and thus 3D cameras and sources were developed for this region. In addition, the laser sources are more radiation hard than the LEDs used for the masks.

The alignment system for the 16 CSC chambers in one endcap is shown in fig. 24. Large and small CSC chambers are linked to each other by pairs of 3D cameras and sources. Each chamber is fitted with two 3D cameras and two 3D sources, monitoring the positions of its two neighboring chambers. In addition, there are two 3D cameras on each of the large CSCs viewing 3D sources on the EI alignment bars, thus establishing the link to the reference grid. There are no sensors linking the small chambers directly to the grid. There are also no logical sub-divisions of the system into large-small chamber pairs like in the MDT system; rather, the 16 CSC chambers of one wheel are all linked to their neighbors and to the bars in a symmetric way.

3.4 Atmospheric Effects

Turbulence and temperature gradients cause density variations in air, changing its refractive index. If this change is perpendicular to an optical line, the light is bent, causing both stochastic and systematic errors in an optical sensor measurement. When sensor measurements are separated by tens of seconds, the error due to turbulence is stochastic. Turbulence is thus not a




Fig. 24: Layout of the CSC proximity system. Shown are two small chambers (green) and one large chamber (yellow), together with the two nearest alignment bars (gray). Two 3D cameras (red) on the large chamber view 3D sources (cyan) on the two bars; in addition, there are two cameras on the large chamber viewing sources on the two small chambers, and vice versa. All these sensors are mounted at the outer radius of the CSC region, where radiation levels are tolerable.

serious problem for the alignment system, because its effects can be greatly reduced by taking the average of several measurements [15], should this turn out to be necessary. Static thermal gradients, however, can cause large, systematic errors in optical measurements. For instance, a laser beam traveling 4 m through air with a gradient of 10 K/m will be displaced by 80 μm .

Static thermal gradients are easy to eliminate in the laboratory if an optical line is not enclosed: by blowing air across the line with fans. In ATLAS, any enclosures of the optical paths are avoided unless they are absolutely necessary. In this way we rely in most places on the flow of cooled air in the cavern: 60,000 m^3/h of air at 17 $^\circ\text{C}$ are blown into the cavern through 12 inlets near the floor, and extracted through two outlets at the ceiling. Where enclosures of optical lines are unavoidable, they are made as large as possible, and the enclosing material and its thickness are chosen to conduct heat well. This greatly reduces the temperature gradient from one side of the enclosure to another. In detail, there are the following cases in the endcap:

- The RASNIK lines inside alignment bars are the obvious example of enclosed optical paths. The bar tubes have at least 72 mm inner diameter, and they are made of 4 mm thick aluminum, thus conducting heat very well.
- Polar lines pass in most places from one wheel to the next through open air, where according to simulations the temperature gradients are too small to cause problems. The only place where polar lines are enclosed is where the high- η lines pass through empty tubes in the endcap toroid cryostat. These tubes are over 200 mm in diameter for almost their entire

length, and they have thick and highly insulated aluminum walls.

- Azimuthal and chamber-source lines pass between vertical MDT chambers, whose tubes are made of aluminum. Similarly, in-plane RASNIK lines pass between vertical MDT chamber multilayers. Static thermal gradients do not form easily upon vertical surfaces.
- Proximity sensors and 3D sensors have too short distances between CCD and light source for thermal gradients to cause any noticeable effects.

With the design choices that were made, none of these cases is likely to exhibit a problem. We have looked for and studied the effect of thermal gradients in various test setups, and have observed no significant errors caused by any atmospheric effect.

4 Data Handling

Various software packages take care of the different aspects of the endcap alignment system, from the readout of the raw sensor data to providing precision chamber alignment results:

- The LWDAQ software sends commands to the alignment sensors and receives images from them, and analyzes the images (section 4.1).
- The analysis results and other parameters are sent through the LTX program to the PVSS software that controls and monitors the activities of LWDAQ. PVSS also sends the sensor measurements on to an Oracle database, where they are stored (section 4.2).
- The ARAMyS software reads the measurement data and reconstructs the alignment, i.e. the chamber and bar deformations, expansions, positions, and rotations in space (section 4.3).
- Its results are sent back to PVSS, and to another Oracle database, from where they are read into the ATLAS detector model, which provides corrections to the nominal chamber locations and shapes to the track reconstruction packages (section 4.4).

4.1 LWDAQ Data Acquisition Software

The LWDAQ software is written mostly in Tcl/Tk, with time-critical portions in Pascal. It can run on Linux, Mac, or Windows PCs, and implements the low-level routines to interact with the alignment sensor electronics. It provides a graphical user interface for each type of LWDAQ device through “instrument panels”, where the parameters of an image or temperature acquisition can be set, e.g. the driver and multiplexer slots through which the device is connected, the exposure time, etc. Some parameters are optionally adjusted dynamically, like the exposure time required to obtain an image in the desired intensity range. Acquisitions can be performed manually (e.g. for debugging), or controlled by a script running in the “acquisifier panel”. In the latter case, the acquisition script contains all the parameters needed for each acquisition, and the LWDAQ software executes the individual acquisition steps sequentially, by setting the parameters in the instrument panel, followed by the acquisition. The “system server” of LWDAQ establishes communication between LWDAQ and the outside, by receiving commands and sending data back through TCP/IP channels. The LWDAQ software can even run in an entirely remote-controlled mode, without the graphical user interface.

Each acquisition step produces an image or a temperature measurement. Images are analyzed on-line by the LWDAQ software. By default, the software does not store or transmit the image

itself (more than 80,000 bytes), but only the results of the analysis and some diagnostic and image quality information (up to a few hundred bytes). However, images may be stored locally for debugging purposes. There are two forms of image analysis, the BCAM analysis [16] which locates light spots in an image, and the RASNIK analysis [17] which identifies a chessboard pattern with its code squares. The BCAM analysis takes a few milliseconds to run, while the RASNIK analysis takes about ten times longer. The slowest step in the data acquisition process is neither the sensor exposure nor the image analysis, but rather the transfer of the image data by the TCP/IP-VME interfaces from the drivers over the local area network to the computers, which takes about 0.25 s per image and limits the acquisition speed to three images per second.

4.2 Data Acquisition Control and Data Flow

In ATLAS, the LWDAQ software runs on two dual-processor Linux PCs. In normal operation, the configuration parameters for the devices to be read out and analyzed are stored in an acquisition script which is downloaded via a software command originating in a user panel of the PVSS software. The scripts are generated from information in the alignment configuration database, which contains a complete hardware and software specification of the alignment system components: crates, driver addresses, driver slots, multiplexer sockets, device elements, as well as exposure times, image analysis parameters, etc.

4.2.1 LTX Communications Program

The acquisifier process of the LWDAQ system receives sensor images and temperature values from the driver boards through TCP/IP-VME interfaces in the VME crates. Several instances of the LWDAQ acquisifier run on each of the two Linux PCs. Each acquisifier instance, or channel, is controlled by its own script. Acquisifier channels can run in parallel provided there is no interference at the front-end hardware, i.e. they access different crates. For those cases where hardware is shared, sequential operation is required: trying to access the same crate simultaneously from different acquisifier channels is possible in principle, but inefficient due to the time lost by one channel while waiting for the competing channel to finish an acquisition. The time required to read out a full cycle of devices, expected to be about 15–20 minutes, is optimized by choosing an efficient parallel and sequential configuration of acquisifier processes.

The system server of an LWDAQ instance receives commands via a TCP/IP listening port and can be commanded to download a script, start and stop running, and to open a TCP/IP socket for results upload. Each acquisifier process performs the image analysis on its BCAM and RASNIK images. Error conditions found by the image analysis code and by additional checks within the acquisifier script are flagged, and error codes are appended to the results data stream. The image analysis results and temperature values, along with a time stamp and the error code are forwarded to the configured TCP/IP upload socket. The system server communicates via TCP/IP with a control and communications program, LTX, running on one of the two Linux PCs. The front-end command and data paths are shown in fig. 25. The LTX program initializes sockets, sends the appropriate script name and path to each instance of the system server, and sends commands for script downloading and starting. At the same time LTX receives status messages from the system server and receives the image analysis and temperature results. Each instance of the system server communicates via a different upload socket and the values are received in separate threads within LTX.

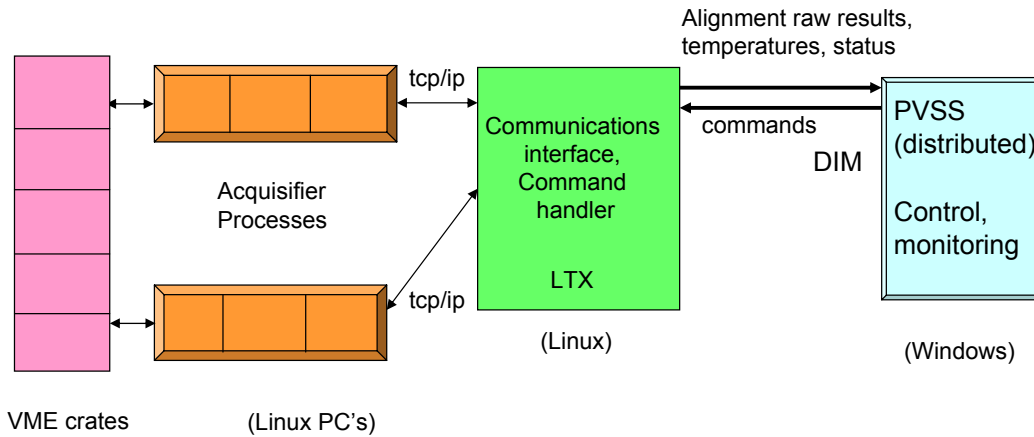


Fig. 25: LWDAQ Controls. The LTX process controls several LWDAQ instances, each of which in turn controls a group of sensors connected to drivers in a VME crate. The PVSS program performs supervisory control and monitoring.

4.2.2 PVSS Control Software and Databases

The LTX program receives its commands from a user interface which is part of a PVSS SCADA (Supervisory Control and Data Acquisition) control program running on a Windows XP machine. PVSS [18] is a commercial product which provides a framework for user interface panels, displays, scripts, and data structures. PVSS projects are constructed with user-defined panels and scripts which access “data-points”, the persistent data storage elements. PVSS is used as the control and monitoring framework for all components of ATLAS, as well as for the other LHC experiments. Individual projects run in a distributed mode, meaning that the internal data structures (data-points) are visible to other projects. A higher-level PVSS project receives state and status information from each subsystem and displays these for monitoring purposes in the ATLAS control room. LTX sends acquisifier status information as well as the analysis results up to the muon endcap alignment PVSS project. The bi-directional communication between LTX (on Linux) and the PVSS program (on Windows) uses the DIM (Distributed Information Management System) [19] software package, which itself is based on TCP/IP. The implementation of DIM within PVSS has all the required functionality and makes inter-process communication across platforms very simple. When values are received via DIM, they are automatically inserted into the appropriate PVSS data-points. PVSS functions provide queued software interrupts when a data-point element changes. Similarly, when certain data-points in PVSS are updated, command parameters are sent out to LTX for control of the acquisifier processes. Data-points and DIM service names are configured and associated with each other using a PVSS configuration script. The data flow is shown in fig. 26.

Upon receiving RASNIK, BCAM, or temperature results, the PVSS program updates one of three buffers. Each type of data goes into its own buffer, since the data structures for the three sources are different. When any of the buffers is full, the buffer contents are appended to a table in the Oracle on-line database of ATLAS. Results from each individual device are entered into a row of the corresponding database table, along with a time stamp and a numerical key to the device name. An auxiliary table associates device names with the numerical keys. Data can be extracted from the database using Java-based programs which allow data selection by time range and other parameters. For a requested time range, the data from all three data tables are made available together for use by the alignment reconstruction (section 4.3).

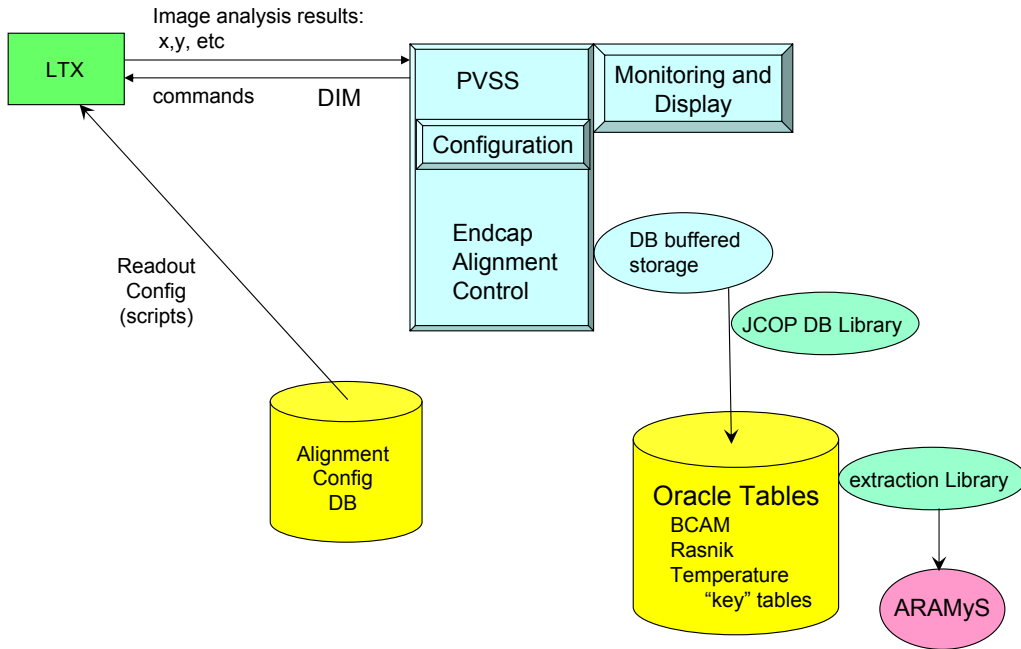


Fig. 26: Data flow in the alignment data acquisition. The LTX process sends scripts, produced from information in the alignment configuration database, to the LWDAQ instances. The readout values received from them by LTX are sent through PVSS to an Oracle database, from where they are read back for alignment reconstruction.

Several configuration panels and user interfaces were implemented in the PVSS framework:

- one user interface configures the script names and port numbers for the parallel and sequential acquisifier processes;
- a configuration panel permits changing the database buffer sizes and the database which receives the data;
- a panel allows the user to store current readout values as nominal values which are used for data checking; limits for comparison are also stored;
- a choice of continuous running or acquiring a fixed number of cycles is provided.

The PVSS control program displays the process status for each acquisifier instance and displays the device names and errors codes as results are received. It keeps statistics of number of devices read and error conditions per data acquisition cycle. Based on the number of devices read and the number of error conditions, an overall status flag is set. This status information is made available to the higher level PVSS FSM (Finite State Machine) process which monitors all subsystems for reporting to the ATLAS control room.

4.3 ARAMyS Alignment Reconstruction Software

For the reconstruction of the alignment information, i.e. the precision chamber positions and rotations, as well as their deformations and expansion, from the data measured by the alignment sensors, the alignment reconstruction software ARAMyS (Alignment Reconstruction and Simulation for the ATLAS Myon Spectrometer) [20, 21] has been developed. It has been extensively

tested in simulations of a complete ATLAS muon endcap, as well as with real data, ranging from simple calibration stands with only two sensors to complete muon wheels, comprising up to 80 chambers, 8 alignment bars, and many hundreds of sensors.

4.3.1 Method of Alignment Reconstruction

The basic elements of any alignment problem are several coordinate systems, characterized by six parameters for their positions and rotations in space. Points, e.g. the endpoints of MDT wires, characterized by three parameters for their positions, can be defined in any given coordinate system. A set of points defined in a common coordinate system represents a rigid body. Alignment becomes an issue once there are points defined in different coordinate systems, as then the relative positions and rotations of the coordinate systems affect the relative positions of the points in space. Alignment sensors link points from different local coordinate systems and provide measurements depending on their locations relative to each other. The alignment accuracy is limited by the accuracies of the individual sensors. The concept of an alignment sensor is not limited to optical devices: a temperature sensor on an object with a known thermal expansion coefficient can be regarded as a length sensor linking the two endpoints of the object, and an optical survey can be treated as the measurements of many pseudo-sensors determining the positions of survey targets in the global (or survey) coordinate system.

The alignment, i.e. the positions and rotations of all local coordinate systems, is reconstructed by comparing measured values from the sensors to expected (calculated) values for an assumed set of local coordinate system positions and rotations, numerically minimizing the quantity

$$\chi^2 = \sum_i \sum_j \frac{(X_{ij,\text{measured}} - X_{ij,\text{expected}})^2}{\sigma_{ij,\text{absolute}}^2 + \sigma_{ij,\text{relative}}^2 + \sigma_{ij,\text{intrinsic}}^2}$$

by variations of these assumed parameter values. The sums run over all sensors and all the coordinates X_{ij} measured by each sensor.

A complication to the alignment problem emerges if the objects in question are not rigid but can mechanically deform or thermally expand. In this case, an object-shape function is introduced which, depending on several shape parameters, returns for the position of any point on the object a small correction representing the deformation or expansion of the object. The values of the shape parameters are determined as part of the minimization process, by measurements of alignment sensors linking points on the same object.

The program ARAMyS implements the above concept. It is written in C, and uses for the χ^2 minimization in the alignment fit the standard package MINUIT [22]. The use of MINUIT makes useful information accessible for the user, e.g. about convergence and quality of the fit, and the errors and correlation matrices of all the fitted parameters. MINUIT has been used and tested for decades, and is a mature and reliable piece of software. Input to ARAMyS are:

- a description of the geometry, containing all the local coordinate systems and points, i.e. essentially the complete calibration data of alignment sensors and their mounts;
- the network of alignment sensors linking some of the previously defined points, as well as the assumed absolute and relative accuracies and intrinsic resolutions of the sensors;
- the measurements of the alignment sensors from the detector.

After the χ^2 minimization, ARAMyS outputs:

- The reconstructed positions and rotations of all local coordinate systems with respect to the global system, as well as the chamber and bar deformation and expansion parameters. This information is stored in the ATLAS offline Oracle database, for use by the track reconstruction routines (section 4.4).
- Additional information for data quality checking and debugging purposes, like the χ^2 /ndf of the alignment fit, and the contributions of individual sensors to χ^2 , plus all the information from MINUIT.

Alignment reconstruction by global χ^2 minimization is a process of complexity $O(N^3)$ for N fitted parameters, where each object, bar or chamber, contributes $N = 6$ parameters for position and rotation, and around $N = 9$ deformation and expansion parameters. Consequently, CPU performance is an issue for large alignment systems like the one of ATLAS. One approach to this problem is to make use of factorization, i.e. the feature of a given alignment system design that the alignment can be reconstructed by splitting the problem into subsets of objects to be aligned sequentially, without noticeable loss of accuracy or consistency. A generalized version of this approach is to use special minimization algorithms for sparse matrices: these are applicable when most of the elements of the correlation matrix of the problem are zero, which is the case for typical alignment problems. Factorization is a non-trivial property of an alignment system: the muon endcap alignment system *is* approximately factorizing, but the barrel system is not.

The alignment of the two ATLAS muon endcaps, comprising about 10,000 fitted parameters in total, can be reconstructed by performing 864 small fits of 9 or 12 parameters each (bar/chamber shapes and locations of MDT chamber pairs with respect to bars, respectively), and two large fits of 384 parameters each (locations of bars and CSC chambers with respect to each other). The fit of CSC locations has to be included in the fit of the bars because the CSC alignment system is, in contrast to the MDT system, not factorisable from the reference grid system. This process takes, on a 3 GHz dual-Pentium PC, about 2 minutes. For the initial configuration of ATLAS, without EE chambers and corresponding bars, this time is only 1 minute. Another viable option, and currently the favored one, is to perform a combined fit of bar and CSC locations and shapes, and combined fits of MDT chamber pair locations and shapes. This takes about 4 minutes for the initial configuration and reduces the expected false sagitta widths by about 10%. The improvement is due to the fact that the system is only approximately factorizing, and that e.g. azimuthal BCAM measurements do have some sensitivity to the bar shapes, so that factorizing the two is not strictly appropriate. A further reduction of the false sagitta widths by 2% could be obtained by fitting all locations and shapes at the same time. This, however, increases the CPU time dramatically to about 4 hours for the initial configuration⁸.

Another useful option in ARAMyS is to introduce additional fit parameters beyond the standard set of chamber and bar positions, rotations, deformations, and expansions. For instance, one could consider the calibration method for one type of sensor to have a systematic offset or another kind of distortion, i.e. an effect that is common to all sensors of that type. By introducing and fitting an additional parameter modeling this, the effect could be both measured and taken into account in the alignment reconstruction. In a similar fashion, quantities like e.g. a twist common to all alignment bars of a type, or a systematic deviation of the MDT chamber tube pitch from the nominal value could be studied.

⁸In this mode, MINUIT is used to fit several thousands of free parameters simultaneously. This works perfectly fine, and is possibly one of the largest applications of MINUIT ever. The authors of MINUIT actually claim that “the technological limitations of MINUIT can be seen around a maximum of 15 free parameters at a time” [22].




Fig. 27: Performance simulation of the endcap alignment system in the initial configuration, without EE chambers. Assuming design resolutions for all sensors, 600 statistically independent incarnations of the system were simulated. The small inset shows the histogrammed false sagitta distribution; the large plot shows color-coded the position-dependent variations of its r.m.s. width. Small differences between sectors are due to statistical fluctuations. The black area is a region with false sagitta widths between 55 μm and 160 μm , where, due to the missing EE chambers, the sagitta of a track has to be reconstructed using two MDT chambers only, resulting in a poorer accuracy. For a single incarnation of ATLAS the local fluctuations in the false sagitta are equal, in r.m.s., to the width of the false sagitta distribution.

4.3.2 Simulation of Alignment System Performance

ARAMyS can also be used, in the design phase of an experiment, to predict by Monte-Carlo simulation the performance of an alignment system based on a given network of sensors and their expected precision. For this application, the sensor measurements from a real setup are replaced by the expected measurements as calculated by ARAMyS, randomly smeared by an amount that reflects the intrinsic resolution and the absolute and relative accuracy of each individual sensor. The alignment is then reconstructed by using these simulated measurements, and a figure of merit is computed from the difference between true and reconstructed chamber positions.

For the ATLAS muon spectrometer, the figure of merit is the width of the false sagitta distribution. The false sagitta is the reconstructed deviation from straightness of a straight track traversing a triplet of precision chambers. The ARAMyS simulation predicts a slightly position-dependent false sagitta width in the range of 25–55 μm over the full endcap, well in line with the specification (fig. 27). To quantify the redundancy of the system, one can also simulate setups with one or several sensors removed. Typically, the effect of a single sensor failing is rather small, not exceeding a tolerable local increase in the false sagitta width by a factor of two. An

exception is the case of the failure of a chamber laser source (or any of the azimuthal BCAMs looking at this source), where the false sagitta width increases by a factor of six for the chamber (or chamber pair) in question.

4.4 Alignment Data in the Track Reconstruction

The output of the ARAMyS software are the precision chamber locations and shapes. These data are stored in the ATLAS offline Oracle database. The geometry model package in the ATLAS software framework reads the alignment data from the database, and provides the information to the offline muon track reconstruction packages. Alignment data are also used by the software that reconstructs from the measurements of magnetic-field sensors on MDT and CSC chambers the three-dimensional magnetic-field map of ATLAS, for which ARAMyS provides the locations of the sensors. The same is true for the reconstruction of the three-dimensional temperature map of the ATLAS cavern.

During ATLAS running, alignment data are required for the offline event reconstruction with a latency of less than 24 hours. While medium- and long-term variations are likely to exceed $40\ \mu\text{m}$, variations on a scale as short as a day are expected to be small: day-night temperature variations for instance are minimal in the ATLAS cavern, 100 m below ground. They do therefore not affect the trigger efficiency, so that the alignment data used by the ATLAS on-line software will not require frequent updates.

There are two options for implementing chamber misalignment and distortion in the track reconstruction (and simulation) packages of the ATLAS software:

- by shifting, rotating, and deforming the active detector elements (or entire chambers) according to the alignment data, followed by reconstructing the hits in these elements;
- by reconstructing hits using the nominal locations and shapes of the detector elements, and subsequently shifting and rotating the hits (or the track segments) according to the alignment data.

The first approach is conceptually simpler, but technically more difficult to implement. For instance, care has to be taken to avoid overlapping detector volumes (including portions of empty space typically surrounding them) in some software detector geometry representations, like the one of GEANT. ATLAS has adopted a hybrid solution, in which the misalignment is implemented by shifting and rotating the active elements, but the distortions are implemented by shifting the reconstructed hits. The distortions are small, and thus the latter step can be omitted for pattern recognition, and be only applied for the final track fit.

5 Test and Validation

Many tests of the alignment system and its components have been performed at various levels in order to validate their functionality. The most relevant milestones reached so far are reported below. The ultimate test, reconstructing muon tracks in ATLAS runs without magnetic field and verifying their straightness, will only be possible once the installation of the ATLAS muon spectrometer endcaps has been completed.

5.1 H8 Test Beam Setup

In 2002–04, the ATLAS muon collaboration installed a large-scale test setup of the muon spectrometer barrel and endcap in the H8 test beam line at CERN. The endcap part of the setup consisted of three large-small pairs of ATLAS MDT chambers (EIL1, EIS1, EML2, EMS2, EOL3, EOS3), and six alignment bar prototypes (two in EI, EM, EO each), arranged to approximately represent part of an octant of the endcap. Chambers and bars were equipped with a full set of alignment and temperature sensors. This is the smallest muon endcap system unit that can be aligned and used for track reconstruction. The chamber positions were chosen such that the muon beam from the SPS illuminated mainly the large chambers, traversing them at 15° angle. Chambers and bars were mounted on large support structures of aluminum and steel, designed to be conceptually similar to the corresponding structures in ATLAS. The setup and many test results have been described elsewhere [12, 23, 24]; we will briefly summarize a few key results related to the endcap alignment system here.

5.1.1 Relative Alignment Tests in H8

For testing the relative alignment concept with the H8 setup, two alternative approaches were taken: the use of an auxiliary hardware device called muon simulator, and the use of real muons from the test beam, in runs where no magnet was present between the endcap stations. Both approaches provided an external reference at a level well below $40\ \mu\text{m}$ to compare with the alignment system results.

5.1.1.1 Tests with the Muon Simulator The muon simulator consists of a camera with the optical axis pointing along the straight line corresponding to a hypothetical muon track, and of light sources on MDT chambers whose positions can be measured. From the observed relative movements of the light sources, the sagitta variations can be extracted directly. As the positions of the camera and of the light sources with respect to the MDT wires are known only at the level of centimeters, this device does not provide an absolute reference and can track only variations of the alignment, i.e. can only be used to test the relative alignment concept. It provides equally precise reference measurements in the bending direction and along the wires. In this respect the muon simulator is superior to using real muon tracks.

A muon simulator was present in the test setup for most of the beam time. We used BCAMs for the camera as well as for the light sources. Eight BCAMs used as the camera were placed on a tripod, and three BCAMs whose laser diodes were used as light sources were mounted on extension plates attached to the central cross plates of the large chambers. They were thus modeling a straight muon track that just missed the chambers by 100–200 mm. Of the eight cameras, five could be adjusted such that all light sources were visible to them. The intrinsic error of the sagitta measurement obtained by averaging over the results from the five cameras, estimated from the data, was about $10\ \mu\text{m}$.

The bending direction for muons, transverse to the MDT wires, was inclined by 14° with respect to the horizontal direction in the H8 setup. Static thermal gradients in the hall were vertical to a good approximation, i.e. nearly transverse to the bending direction. The influence of atmospheric effects on optical measurements of this coordinate (as made by both the alignment system and the muon simulator) was thus negligible.

The comparison of measured sagitta variations in the bending direction between alignment system and muon simulator is shown in fig. 28. They agree with an r.m.s. resolution of $14\ \mu\text{m}$ over

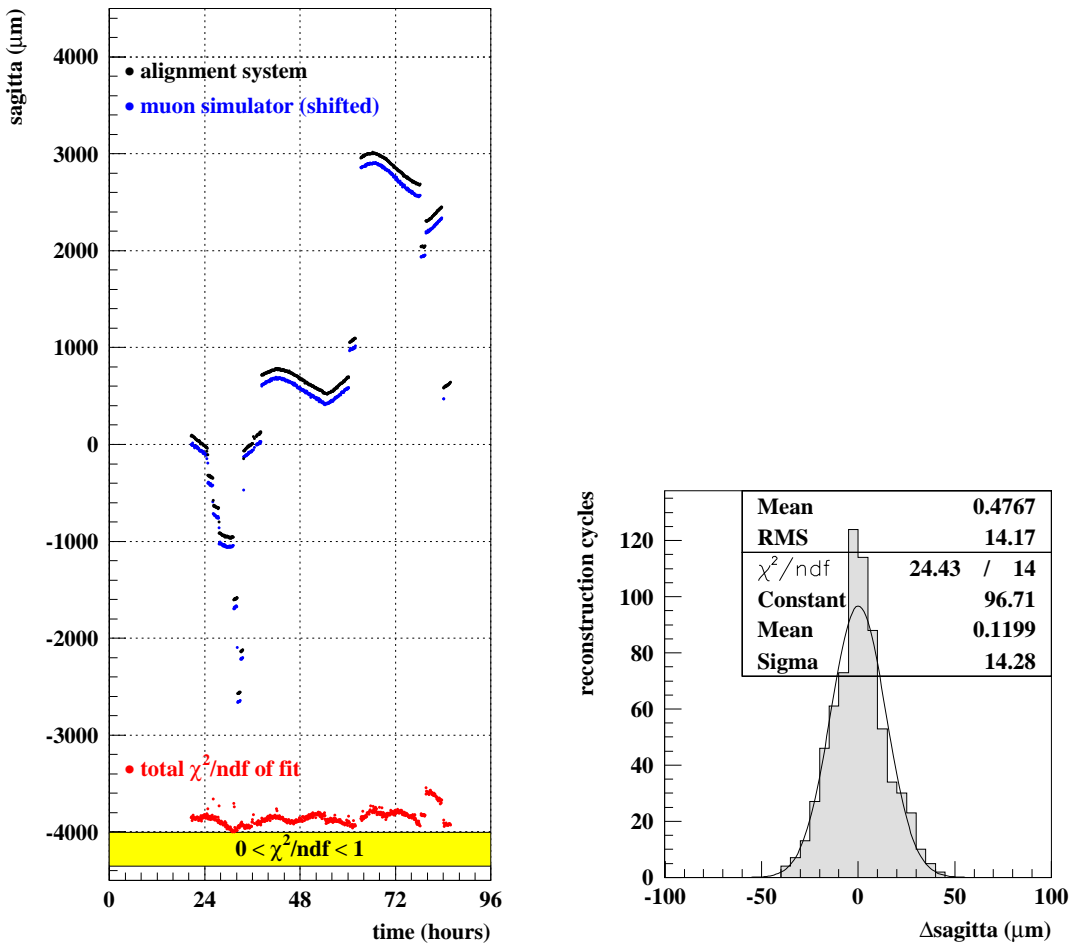


Fig. 28: Left: sagitta variations reconstructed by the alignment system (black, upper points) and measured by the muon simulator (blue, lower points, shifted by $100\ \mu\text{m}$ for clarity of presentation), over a period of 2.5 days. Smooth variations are temperature-induced, steps come from chamber shifts and rotations. The variation of the total χ^2/ndf of the alignment fit is indicated at the bottom (red); the yellow area shows the range $0 < \chi^2/\text{ndf} < 1$. Right: sagitta difference between alignment system and muon simulator. Since this is a test of the relative alignment concept, the mean value of the histogram is arbitrary. Alignment system and muon simulator sagittae agree to $14\ \mu\text{m}$ r.m.s.

a continuous period of 2.5 days. During this time, daily temperature-induced sagitta variations of up to $500\ \mu\text{m}$ were complemented by artificial variations of up to $5\ \text{mm}$ from shifting and/or rotating different chambers. The total χ^2/ndf of the alignment fit varies typically in a range of $0.9\text{--}1.5$. The observed resolution of the alignment system agrees well with the expectation from simulations for the relative alignment mode, which is around $15\text{--}20\ \mu\text{m}$ (the sensor absolute accuracies, which do not contribute here, dominate the absolute alignment accuracy of about $40\ \mu\text{m}$). The same comparison in the direction along the wires (not shown) yields an r.m.s. resolution of about $125\ \mu\text{m}$, safely below the required resolution of around $250\ \mu\text{m}$.

5.1.1.2 Tests with Straight Muon Tracks Tests using straight muon tracks were performed using the same alignment data set. Obviously, fewer data points can be obtained from tracks, as sufficient statistics has to be accumulated for each point. The details of track reconstruction are beyond the scope of this paper and are discussed elsewhere [24, 25]. A complication compared to

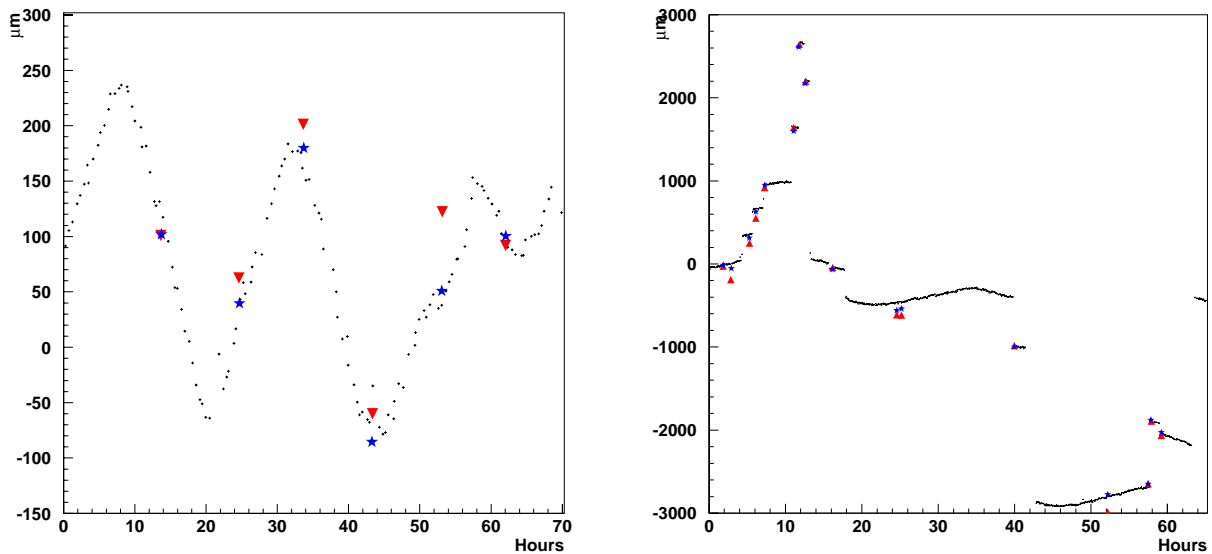


Fig. 29: Sagitta variations reconstructed by the alignment system compared with straight muon tracks from the test beam. The prediction of the alignment system is shown as small crosses (black), triangles (red) indicate the raw track measurements, and stars (blue) show the track sagitta after various corrections. Each track data point corresponds to the average value of the reconstructed sagitta for an ensemble of straight muon tracks. Left: sagitta variations due to normal day-night temperature changes. Right: sagitta variations for controlled chamber movements. Alignment system and track sagittae agree to about $20\ \mu\text{m}$ r.m.s. Note that on the right plot, some of the alignment data points are obstructed by track data points, in particular in those periods where many controlled movements were performed at short intervals.

the ATLAS case was the absence of trigger chambers to measure the coordinate of hits along the MDT wires, which is crucial for this type of analysis. The wire coordinate was obtained instead by using the MDT chambers of the neighboring barrel setup. The wires of that setup were vertical and thus rotated by 14° with respect to those of the large endcap chambers, providing a stereo-measurement of the same tracks. From this information, a track-by-track correction of the measured sagitta depending on the wire coordinate was obtained.

The resulting track sagittae before and after all corrections are compared with the alignment system data in fig. 29. These data cover the same period in time as the muon simulator data, with temperature-induced variations and controlled chamber shifts and rotations. The agreement between the tracking (after corrections) and alignment data is excellent, at the level of $20\ \mu\text{m}$.

5.1.2 Absolute Alignment Tests in H8

In principle, for tests of the absolute alignment an external reference with an accuracy better than $40\ \mu\text{m}$ is required, which can be provided only by muon tracks. Optical surveys have a typical accuracy of around $300\text{--}500\ \mu\text{m}$. However, the endcap alignment system is designed to provide most chamber and bar positions with respect to each other with an accuracy of only about $300\ \mu\text{m}$, and only few and very specific chamber positions (namely those of triplets that can be traversed by a muon track originating at the ATLAS interaction point) with respect to each other with an accuracy as high as $40\ \mu\text{m}$. It is therefore possible to perform at least a limited test by using an optical survey. Absolute alignment tests with both a survey and muon tracks as a reference were performed using the H8 data.

5.1.2.1 Tests Using a Survey The most elegant way of comparing an alignment reconstruction result to the result of an optical survey is to include the survey as a set of additional constraints to the chamber and bar positions in the χ^2 minimization, using the r.m.s. uncertainties quoted by the surveyors. This option is implemented in ARAMyS by survey “pseudo-sensors”, imagined to be measuring the positions of targets on chambers and bars in the global coordinate system. The positions of the targets in the local chamber and bar coordinate systems were determined earlier as part of their calibrations (sections 3.1.2 and 3.2.3, respectively). The χ^2 /ndf of the fit, and the individual contributions of the survey sensors to χ^2 , can then be used as a measure of the compatibility between the survey and the alignment sensor measurements.

The absolute alignment tests in H8 using an optical survey as a reference were successful [12], but have by now been superseded by a much larger amount of data collected during the installation and commissioning of the endcap precision chambers and their alignment system in ATLAS, as reported in section 5.2.2. We will therefore not discuss them here.

5.1.2.2 Tests Using the Muon Beam The absolute alignment tests using the muon beam as a reference yielded an inconclusive result. With the final set of sensor and mount calibrations, a single alignment sensor (the chamber laser source on the EML2 chamber) contributed an abnormally high amount to χ^2 in the alignment fit, thus indicating that its measurements were incompatible with those of the other sensors. As the final sensor mount calibrations could only be performed after the end of the test beam period, and after unmounting the device in question from the chamber, it was impossible to investigate the reason for the problem or to repeat the test. Excluding the measurements of the device in question from the fit yielded an average false sagitta for an ensemble of straight muon tracks of a few hundred microns: about an order of magnitude worse than the specification, but compatible with what one would expect according to simulations for a system with that device missing.

5.1.3 Conclusions

The relative alignment tests in H8 demonstrated that the alignment system works according to its specifications in this mode. The “successful failure” of the absolute alignment tests with beam indicated that there was room for improvement in the treatment of calibrations.

The main lesson learned was that at every step in the process from assembling a single sensor to commissioning the entire alignment system of an endcap, whatever could be checked and analyzed at that point should be done, and that no possible test or data analysis should be postponed until later. More specifically, it was decided that:

- For each type of alignment sensor and sensor mount there should be two independent calibration methods, shown to produce compatible results.
- Prior to its installation, each device should be put on a calibration test stand, where a quick measurement would be taken and checked for compatibility with the previously determined calibration constants.
- Any kind of calibrations of sensors and their mounts should be analyzed and implemented for alignment reconstruction shortly after the calibration, so that problems could be detected by running ARAMyS on (even partial) data from e.g. a single bar, chamber, sector, or wheel. This applies in particular to the MDT sensor mount calibrations (section 3.1.2).

5.2 Installation and Commissioning in ATLAS

The installation and commissioning of the muon spectrometer endcaps of ATLAS started in 2005 with the assembly of MDT chambers and alignment bars in 2×16 single sectors of the Big Wheels. Five chambers, and, in case of a large sector, one alignment bar, were assembled in vertically stored sector frames. Later, the sectors were assembled into the two Big Wheels in the ATLAS cavern, and the Small Wheels were assembled and surveyed in a surface building, to be lowered into the cavern as a whole. At the time of writing, the two Small Wheels and the two Big Wheels have been fully commissioned, representing about half of the entire system, and the results presented here are based on data from those four wheels.

5.2.1 Chamber and Bar Installation

Before a chamber or bar is installed on its support structure, alignment sensors are mounted on it and their serial numbers and locations are recorded in a database. Prior to this, the calibration of the sensors has been checked on a test stand. After mounting the sensors, their proper functioning is verified. The sensors inside chambers and bars, in-plane and in-bar RASNIKS, as well as temperature sensors, are also checked, and their measurements are used to reconstruct with ARAMyS the shapes of the chambers and bars, respectively. Only a crude plausibility check of the reconstructed deformation is possible at this point, with the chambers being horizontal and supported by foam, and the bars being held in cradles near the kinematic mounts on a cart. After completing the installation of bars and chambers in a sector or wheel, their multiplexers are connected to driver boards, and all alignment sensors are read out. Malfunctioning components are identified and replaced, until a stable running system is reached. At this point, many or most alignment sensors have their partner device installed, and thus their measurements can be used to obtain from ARAMyS reconstructed chamber and bar positions, which can be compared to results from surveys. While small sectors alone do not allow any meaningful tests of alignment reconstruction beyond determining chamber shapes, the alignment bars in the large sectors do permit to reconstruct the positions of the chambers with respect to the bar. In case of a fully assembled wheel, the alignment system is complete except for the polar BCAMs.

5.2.2 First Commissioning Results

A recurrent observation when starting up the alignment system of a newly assembled wheel was that many azimuthal lines were obstructed, while most or all of the other sensors worked fine. This is a feature of the layout of the MDT proximity system, with the azimuthal lines passing through the 10 mm gap between the laser sources on the large and small chambers of a pair (fig. 23). The displacement of a chamber along the beam direction, toward the azimuthal line, by an amount as small as 5 mm is sufficient to cause the small plate holding the lasers of the chamber source to overlap with the azimuthal line. The floppiness of the Big Wheel support structures in this coordinate made this a problem particularly for those wheels, where about 30% of all azimuthal lines were blocked initially. The remaining lines were, however, sufficient to reconstruct the positions of all chambers in the wheels to better than a millimeter. This was sufficient for the alignment system to provide the data for re-positioning the affected chambers.

Except for this particular displacement mode, the alignment system has proven to be, as designed, tolerant to mispositioning, with a dynamic range reaching (for proximity sensors) or even exceeding (for BCAMs) ± 15 mm in all directions. The only other parameter that needed to be adjusted frequently were the rotations of alignment bars around their axes, which are difficult

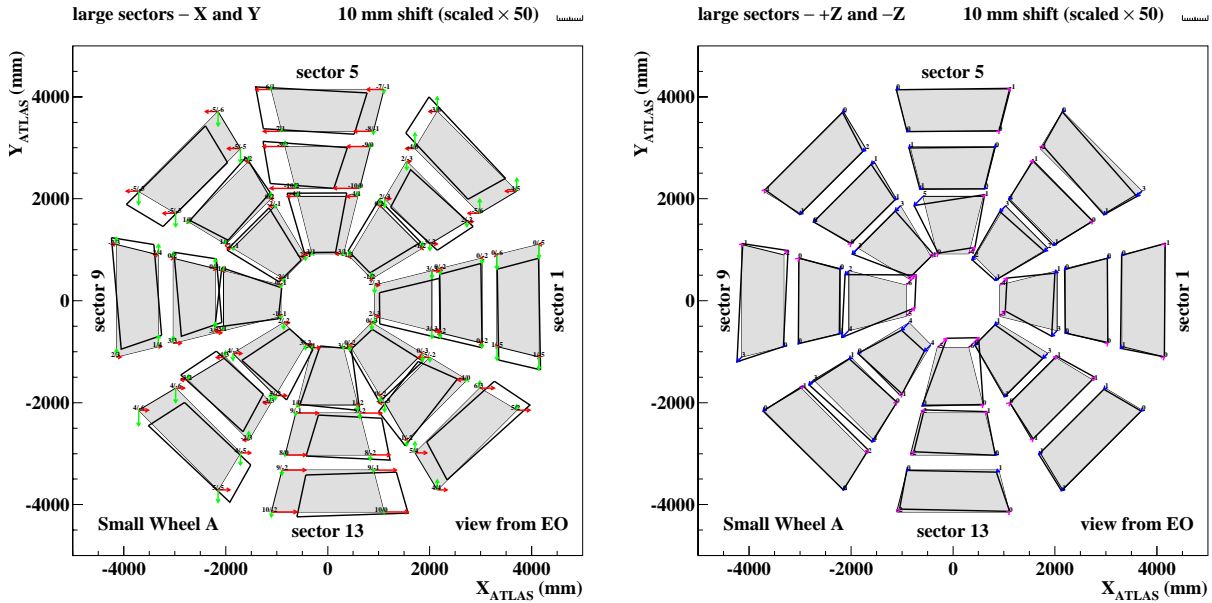


Fig. 30: Large CSC and MDT chamber positions in the A-side Small Wheel as seen from the outside, reconstructed by the alignment system. Grey areas show the nominal positions, black contours the actual positions, where the displacements of the four corner-points have been exaggerated by a factor of 50 to make them visible at this scale. Red and green arrows (left) are pointing along the ATLAS- X and ATLAS- Y axes, respectively; blue arrows are pointing out of the plane along ATLAS- Z , magenta arrows into it (right). The numbers near the arrows give the shifts in X/Y and in Z in mm, respectively. Small chambers are not shown. Their positioning quality is similar. The plots show the final chamber positions after all necessary adjustments.

to get within the required range just by accurate assembly. As an example, fig. 30 shows the large-chamber positions in the A-side Small Wheel after all adjustments.

After the required re-positioning, 98–99% of all relevant sensors were working in all wheels. (Chamber laser sources are counted as 100% working if they are seen by at least one azimuthal BCAM.) The missing ones are mostly due to mechanical damage during installation, some due to dust and dirt, and only a handful of devices are excluded because they provide apparently inconsistent measurements (see below). Running the alignment reconstruction on these data yields χ^2/ndf values for the four MDT wheels of 0.9, 1.1, 1.1, and 1.4, using design resolutions for all sensor types. As the alignment system is redundant and over-constrained, observing $\chi^2/\text{ndf} = 1$ is equivalent to the observed resolutions being in line with the design resolutions (for ndf, the number of degrees of freedom, of the order of a few thousands). For this case, one expects a false sagitta width of about $40\ \mu\text{m}$, as shown by simulations; otherwise the expected width scales like $\sqrt{\chi^2/\text{ndf}}$.

In addition, one expects the individual contributions to χ^2 , the pull distribution, to follow an approximately Gaussian distribution⁹. The observed pull distribution has the expected shape (fig. 31), and an r.m.s. width of about 0.91. A width around 0.83 is expected from simulations

⁹The expected pull distribution is, in general, not exactly Gaussian. The total pull distribution is the sum of many individual, Gaussian, pull distributions: one for each coordinate measured by each sensor type. The expected widths of the individual distributions may be less than unity, for instance due to correlations between sensor measurements which are not taken into account by the simple form of χ^2 being minimized. The sum of many Gaussian distributions with different widths is not a Gaussian distribution anymore.

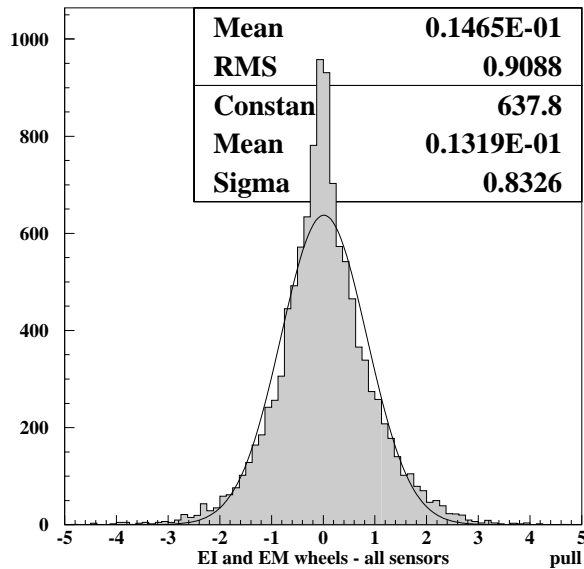


Fig. 31: Distribution of the pulls (contributions to χ^2) of all alignment sensors in the four fully assembled EI and EM wheels of ATLAS. Only a very small number, about 0.1%, of functioning sensors has been excluded from the fit or has pulls outside the plotted range. The black line shows the fit of a Gaussian – only to guide the eye, as one actually expects the distribution to be slightly non-Gaussian.

for the particular mode in which the alignment fit was performed here, i.e. all bar/chamber deformations and positions fitted at the same time. Only a handful of isolated devices, of the order of 0.1%, have huge pulls without an apparent reason, and are therefore tentatively excluded from the fit.

These satisfying results do not change significantly when an optical survey is included as additional constraint in the alignment fit, assuming the r.m.s. accuracy of the survey as provided by the surveyors. This shows that the alignment system measurements are not only self-consistent, but are also well compatible with the survey results.

From the experience in the H8 tests and the first commissioning results we conclude that, after having seen about half of the MDT endcap alignment system at work, it performs according to our expectations. The CSC alignment system was never tested in H8, and the analysis of the first commissioning data indicates that there are some open questions yet to be resolved.

5.2.3 Outlook

The last missing link are the polar BCAM lines, which could not be tested so far. There is nothing fundamentally different between an azimuthal line (with a laser source in-between the two azimuthal BCAMs) and a polar line (with a third BCAM replacing the laser source), and therefore we do not expect any fundamental problems related to BCAMs or alignment bars there. Thermal gradients in static layers of air are not likely to affect the polar lines, as discussed in section 3.4. Finally, the contingency of the middle BCAM in a polar triplet being displaced in such a way that it blocks the line of sight between the two outer BCAMs, like the laser sources on chambers blocking many of the azimuthal lines before, is a remote one, because the clearance between the middle polar BCAM and the line between the two outer BCAMs is about three times larger than in the azimuthal case.

6 Alignment Beyond the Optical System

When opting for an alignment system with a relatively slow readout speed, an implicit assumption is made: with one complete readout cycle lasting about 15–20 minutes, one expects any chamber or bar displacements during ATLAS data-taking to occur on a timescale of at least hours. Movements on scales that are short compared to the readout interval are undetectable. This is particularly true for vibrations – the resonance frequency of a long alignment bar, for instance, is of the order of a few Hertz. The presence of vibrations could be detected by continuously reading out a single alignment device at the maximum speed of about three images per second, but apart from confirming their presence, the measured data would not be of any use for dealing with the vibrations. This may be seen as an intrinsic limitation of the alignment system – on the other hand, if precision chambers would be moving at very short time scales, this would affect not only the alignment system, but the functioning of the muon spectrometer as a whole. To some limited extent the alignment system readout could be modified to accommodate shorter-timescale displacements: sensors in relatively unstable regions of the detector could be read out more frequently than those in relatively stable regions, thus making the system sensitive to movements on a scale of maybe 5 minutes in the unstable regions.

The barrel and endcap alignment systems provide the internal alignment of the three parts of the spectrometer separately. They do not produce information about the absolute positions of the barrel and the endcaps in the global ATLAS coordinate system, or with respect to each other, i.e. no external alignment. The ATLAS system is defined, in the construction phase of ATLAS, by a network of survey marks in the ATLAS cavern, and will eventually be replaced by a system defined by one of the inner tracking detectors (not the muon spectrometer). The connection to the cavern system can be obtained from optical surveys of a few MDT chambers – those in the EO region are particularly suited, as they are visible when the detector is closed for data-taking. The connection to the final ATLAS system will be obtained using tracks.

Due to budget constraints, the small chambers in the barrel are either not linked by optical sensors to other chambers at all, or not by a sufficient set of sensors to reach the required accuracy. The same is true for some chamber types in the barrel-endcap overlap region, particularly the barrel chambers covering the gaps between EIL4 chambers. All these chambers can only be aligned with respect to the rest of the spectrometer with muon tracks. Whether their positions are sufficiently stable in time for this to work remains to be seen.

Even for chambers that have alignment sensors, alignment with tracks will be performed during normal ATLAS data-taking, in order to cross-check the alignment results obtained from the optical systems. Due to the lack of resolution of MDT chambers in the coordinate along the wire, and due to the fact that muon tracks collected during ATLAS physics runs are projective and thus have a limited angular spread, tracks are not sensitive to all degrees of freedom of the chambers in the same way as the sensors are. This makes alignment using *only* tracks challenging. One possibility to overcome the angular-range limitation to some extent is to use cosmic muons. Eventually, however, the alignment of the muon spectrometer should be determined by a global optimization using sensor and track data at the same time. This could be done, e.g. by minimizing a global χ^2 with contributions from optical sensors and muon tracks, and from both barrel and endcaps. The best possible alignment of the entire spectrometer will be obtained by exploiting all available information at the same time.

Eventually, the alignment of the muon spectrometer with respect to the inner tracking detectors will have to be determined as well, mainly in order to be able to match tracks between the detector components. Owing to the already very good stand-alone momentum resolution of

the muon spectrometer, the momentum resolution can, in most regions, only very slightly be improved by reconstructing combined inner-detector-muon-spectrometer tracks. ATLAS has no optical alignment system linking the muon spectrometer to the inner detector, and thus determining the muon spectrometer alignment with respect to the inner detector will have to rely entirely on tracks.

7 Summary

The precision chambers in the ATLAS muon spectrometer endcap are equipped with a sophisticated optical alignment system. We have described the design of the sensors and of the system, as well as the software architecture. The alignment system has been subjected to numerous tests, from which we conclude that reaching the design sagitta accuracy of $40\ \mu\text{m}$ at the start-up of ATLAS is well within reach. The installation and commissioning of the system in the ATLAS cavern is underway, and has produced encouraging results.

Acknowledgments

The alignment of the barrel and endcaps of the ATLAS muon spectrometer has been a project of the Saclay, NIKHEF, MPI Munich, Freiburg University, and Brandeis University groups in the ATLAS muon collaboration, and we gratefully acknowledge the numerous contributions of our barrel colleagues to the endcap system, as well as the friendly collaboration over all the years. The results presented in this paper could not have been obtained without the work of many colleagues not listed as authors, and we would like to express our thanks to them here. W. Andreazza (CERN) provided highly competent and efficient technical support in many occasions. T. Klioutchnikova's (CERN) work in finding and fixing countless interferences between optical alignment lines and physical obstacles was invaluable. M. Bradshaw (Brandeis) helped develop and improve the sensor electronics. D. Levin (University of Michigan) performed the tracking part of the H8 data analysis leading to fig. 29. Several members of the ATLAS muon group at Freiburg University made significant contributions, most prominently J. Tobias, who programmed the institute's CMM and calibrated all the alignment bars, K. Handrich and T. Raufer who worked on alignment bar design and reconstruction aspects, M. Kollefrath who dedicated much of his time to optimizing the CMM accuracy, and finally the institute's machine shop, where all the sensor platforms for the bars were machined. H. Kroha (MPI Munich) and other members of the MPI group contributed to the design of the alignment bars. Many members of the CERN survey group, most prominently D. Mergelkuhl, A. Behrens, and A. Wiart, surveyed sectors and wheels for us many times, helped to design and set up the system for photogrammetry calibration measurements of MDT chambers, and were constantly helpful and willing to provide information. Members of the CERN ATLAS muon group, C. Fabjan, K. Mair, S. Palestini, and I. Trigger worked in the H8 test of the alignment system and contributed to its success. The Brandeis students D. Nash, D. Pomeroy, and B. Way, as well as S. Pintaric (CERN) and N. Nation (Boston University), together performed photogrammetry and calibration measurements on hundreds of MDT chambers. They and other students at Brandeis, among them A. Asen, R. Banks, S. DeRocks, I. Dybottom, and P. Graves, also assembled and calibrated the thousands of alignment devices. The Washington students C. Dowell, M. Olson, and B. Reynolds contributed to the data acquisition and control software. The CERN metrology group performed the CMM measurements of CSC chambers and made the photo in fig. 19. The workshops at Tufts and Harvard University produced parts for the alignment system. The

TGC locator sources and electronics were laid out and manufactured at the Weizmann Institute of Science. Making the project a success would have been virtually impossible without the dedicated and high-quality work of F. Mello in the Brandeis machine shop.

References

- [1] ATLAS Collaboration, The ATLAS Experiment at the CERN Large Hadron Collider, submitted to JINST.
- [2] ATLAS Muon Collaboration, ATLAS Muon Spectrometer Technical Design Report, CERN/LHCC/97-22 (1997).
- [3] C. Guyot et al., The Alignment System of the Barrel Part of the ATLAS Muon Spectrometer, ATLAS Note ATL-COM-MUON-2008-002 (2008).
- [4] G. Chiefari et al., Muon Detection in the L3 Experiment at LEP, Nucl. Instr. Methods **A 277** (1989) 187.
- [5] H. Dekker et al., The RASNIK/CCD 3-dimensional Alignment System, Proceedings of the 3rd International Workshop on Accelerator Alignment, CERN, Geneva (1993); H. L. Groenstege, The RASNIK/CCD 3D Alignment System, ATLAS Note ATL-MUON-94-063, ATL-M-PN-63 (2004); H. v. d. Graaf et al., RASNIK Technical System Description for ATLAS, NIKHEF Note ETR 2000-04 (2000).
- [6] K. Hashemi and J. Bensinger, The BCAM Camera, ATLAS Note ATL-MUON-2000-024 (2000); D. Daniels, K. Hashemi and J. Bensinger, BCAM Calibration, ATLAS Note ATL-MUON-2000-026 (2000).
- [7] K. Hashemi, Long-Wire Data Acquisition Specification (2007), available at <http://alignment.hep.brandeis.edu/Electronics/LWDAQ/LWDAQ.html>.
- [8] M. Bosman et al., Estimation of Radiation Background, Impact on Detectors, Activation and Shielding Optimization in ATLAS, ATLAS Note ATL-GEN-2005-001 (2005); M. Dentan, Overview of the ATLAS Policy on Radiation Tolerant Electronics, available at <https://edms.cern.ch/document/901767/1>; M. Dentan and P. Farthouat, ATLAS Policy on Radiation Tolerant Electronics, available at <https://edms.cern.ch/document/113816/2>.
- [9] K. Hashemi and J. Bensinger, Irradiation of the TC255 CCD by Fast Neutrons, ATLAS Note ATL-MUON-98-253 (1998); K. Hashemi and J. Bensinger, Irradiation of the TC255 CCD by Fast Neutrons Part Two, ATLAS Note ATL-MUON-2000-011 (1999).
- [10] C. Lasseur, Three Years of Digital Photogrammetry at CERN, Proceedings of the 6th International Workshop on Accelerator Alignment, Grenoble, France (1999).
- [11] A. Schricker, The Alignment System of the ATLAS Muon Endcap Spectrometer, Ph.D. thesis, Vienna University of Technology (2002), available at <http://atlas.web.cern.ch/Atlas/documentation/thesis/schricker/thesis.pdf>.
- [12] C. Amelung et al., Reference Bars for the Alignment of the ATLAS Muon Spectrometer, Nucl. Instr. Methods **A 55** (2005) 36.
- [13] C. Amelung and A. Schricker, Calibration of Alignment Bars for the ATLAS Muon Spectrometer Endcap, ATLAS Note ATL-MUON-2005-009 (2002).
- [14] C. Amelung, Muon Spectrometer Alignment in ATLAS: Reaching the Design Performance in the Endcap, ATLAS Note ATL-MUON-2001-018 (2001).
- [15] K. Hashemi, Some Unoriginal Thoughts on Thermal Gradients and Turbulence in Air, available at http://alignment.hep.brandeis.edu/ATLAS/Thermal_Gradients.pdf.

- [16] K. Hashemi, BCAM Analysis (2007), available at http://alignment.hep.brandeis.edu/Devices/BCAM/User_Manual.html#Analysis.
- [17] K. Hashemi, RASNIK Analysis (2007), available at <http://alignment.hep.brandeis.edu/Devices/RASNIK/>.
- [18] PVSS homepage at <http://www.etm.at>.
- [19] C. Gaspar and M. Donszelmann, DIM: A Distributed Information Management System for the DELPHI Experiment at CERN, Proceedings of the 8th Real-Time Computer Applications in Nuclear, Particle and Plasma Physics (1993); DIM homepage at <http://cern.ch/dim>.
- [20] C. Amelung, Alignment Reconstruction and Simulation for the ATLAS Muon Spectrometer: ARAMyS Manual, available from the ARAMyS homepage at <http://cern.ch/amelung/aramys.html>.
- [21] C. Amelung, ARAMyS – Alignment Reconstruction Software for the ATLAS Muon Spectrometer, Nucl. Phys. **B** (Proc. Suppl.) **172** (2007) 132.
- [22] F. James and M. Roos, Comput. Phys. Commun. **10** (1975) 343; F. James, MINUIT Reference Manual, CERN Program Library Long Writeup D506 (1994); F. James and M. Winkler, MINUIT User's Guide, CERN (2004).
- [23] C. Amelung et al., Results from Tests of the Muon Endcap Alignment System in the H8 Beam Line 2002, ATLAS Note ATL-MUON-2004-009 (2003).
- [24] C. Adorisio et al., System Test of the ATLAS Muon Spectrometer in the H8 Beam at the CERN SPS, ATLAS Note ATL-MUON-PUB-2007-005 (2007), accepted by NIM.
- [25] D. Levin, Impact of the ATLAS Muon Spectrometer Alignment System on Tracking at the CERN H8 Test Beam, Nuclear Science Symposium Conference Record, 2004 IEEE Vol. 1 16–22 (2004) 662, DOI 10.1109/NSSMIC.2004.1462279.