ATLAS Muon Note
February 7, 2008
DRAFT Version 2.0.6

The Optical Alignment System of the ATLAS Muon Spectrometer Endcaps

C. Amelung [1], J. Bensinger [1], A. Dushkin [1], M. Gardner [1], K. Hashemi [1], E. Henry [2], B. Kaplan [1,5], P. Keselman [1,6], U. Landgraf [3], J. Rothberg [4], A. Schricker [1,7], N. Skvorodnev [1], H. Wellenstein [1]

[1] *Brandeis University, Waltham, MA 02454, USA*
[2] *Harvard University, LPPC, Cambridge, MA 02138, USA*
[3] *Universität Freiburg, Fakultät für Physik, 79104 Freiburg, Germany*
[4] *University of Washington, Seattle, WA 98195, USA*
[5] *now at Yale University, New Haven, CT 06520, USA*
[6] *now at University of California, San Francisco, CA 94122, USA*
[7] *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 over a thousand muon precision chambers, arranged in three concentrical 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, with an outer diameter from 13 m to 24 m. Muon chambers are equipped with a complex on-line 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 ambitious goal of a 40 μm alignment accuracy, required for reconstructing final-state muons at the highest expected energies.

Contents

1 Introduction **4**

1.1 Motivation . . . 4
1.2 Muon Spectrometer Layout . . . 4
1.3 Alignment Strategy . . . 6
1.4 Terms and Definitions . . . 6

2 Alignment Sensors **7**

2.1 BCAM-type Sensors . . . 8
2.1.1 BCAMs . . . 8
2.1.2 Chamber Laser Sources . . . 10
2.1.3 3D Sources and 3D Cameras . . . 11
2.1.4 TGC Locator Sources . . . 12
2.2 RASNIK-type Sensors . . . 12
2.2.1 Proximity Cameras and Masks . . . 14
2.2.2 Three-Element RASNIK Sensors . . . 15
2.3 Sensor Electronics . . . 15

3 Alignment System **17**

3.1 Chamber Geometry . . . 18
3.1.1 MDT Chamber Deformations . . . 18
3.1.2 MDT Sensor Mount Calibration . . . 20
3.1.2.1 Coarse Calibration by Photogrammetry . . . 20
3.1.2.2 Fine Calibration using Calibration Tools . . . 21
3.1.3 CSC Chamber Geometry . . . 22
3.2 The Reference Grid . . . 23
3.2.1 Grid Layout . . . 23
3.2.2 Alignment Bars . . . 25
3.2.3 Bar Calibration and Shape Reconstruction . . . 27
3.3 The Proximity Measurement Systems . . . 28
3.3.1 MDT Proximity System . . . 28
3.3.2 CSC Proximity System . . . 30

4 Data Handling **31**

4.1 LWDAQ Data Acquisition Software . . . 31
4.2 Data Acquisition Control and Data Flow . . . 31
4.2.1 LTX Control and Communications Program . . . 32
4.2.2 PVSS Software and Oracle Databases . . . 33
4.3 ARAMyS Alignment Reconstruction Software . . . 34
4.3.1 Method of Alignment Reconstruction . . . 34
4.3.2 Simulation of Alignment System Performance . . . 36
4.4 Alignment Data in the Track Reconstruction . . . 37

5 Test and Validation **38**
5.1 The H8 Test Beam Setup . . . 38
5.1.1 Relative Alignment Tests in H8 . . . 38
5.1.1.1 Tests with the Muon Simulator . . . 39
5.1.1.2 Tests with Straight Muon Tracks . . . 39
5.1.2 Absolute Alignment Tests in H8 . . . 40
5.1.2.1 Tests Using a Survey . . . 40
5.1.2.2 Tests Using the Muon Beam . . . 41
5.2 Installation and Commissioning in ATLAS . . . 42
5.2.1 Chamber and Bar Installation . . . 42
5.2.2 First Commissioning Results . . . 42

6 Alignment Beyond the Optical System **45**

7 Summary **46**

Acknowledgments **46**

References **46**

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 $p_T = 1\,\text{TeV}$. 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 drift tubes glued onto a 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. The drawback 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 pseudorapidity[1] $\eta = 0$ (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%, 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 which is based on optical and temperature sensors, and on alignment bars, which are up to 9.6 m long instrumented aluminum tubes used as large precision reference rulers. 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 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 (inner), EE[2] (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 16 sectors (numbered from 1 to 16) of trapezoidal chambers; those in odd-numbered sectors

[1] The pseudorapidity, η, is defined as a function of the polar angle, θ, which is measured with respect to the beam line: $\eta = -\ln(\tan(\theta/2))$. Consequently, $\eta = \pm\infty$ points along the beam line, and $\eta = 0$ is transverse to it.

[2] Due 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.

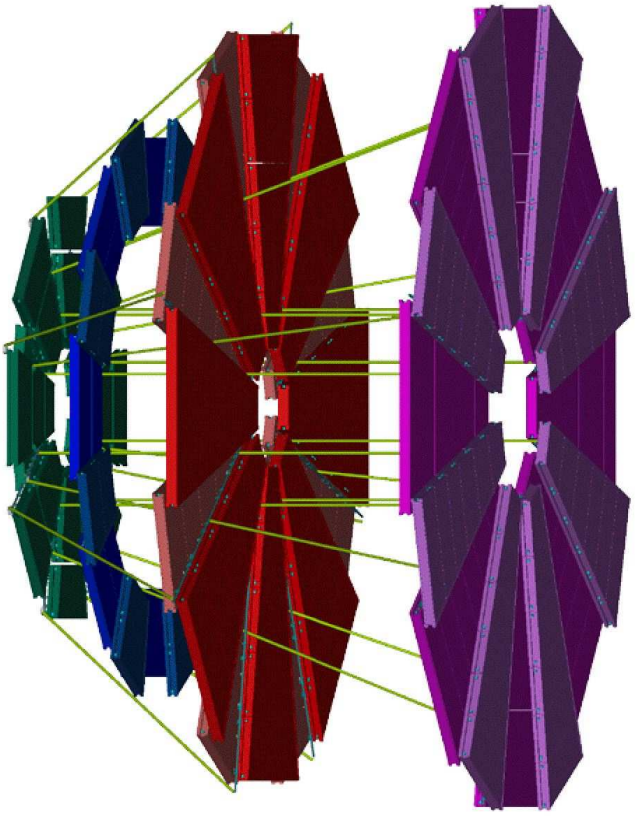

Fig. 1: Layout of the MDT chambers in one muon spectrometer endcap. The ATLAS interaction point is to the left, the beam line horizontal. Shown are the four stations EI (green), EE (blue), EM (red), EO (magenta), as well as the alignment bars (cyan) and the small number of alignment lines (yellow) linking the stations to each other. Not shown are the CSC chambers, and the vast majority of alignment lines linking chambers and bars within a station. The vessel of the endcap toroid magnet fills the space between the EI and EM stations, up to the inner radius of the EE station. There is no magnet between EM and EO. *(To be replaced by a slightly modified figure for the final version.)*

cover 28° in azimuth and are called large (L) chambers, those in in even-numbered sectors cover 17° in azimuth and are called small (S) chambers[3]. 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. Mechanically, the EO stations ("end walls"), the EM stations ("Big Wheels"), as well as the inner part of the EI stations (CSC, EIL1/2/3 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

[3]This geometry has been chosen in order to match the layout of the barrel MDT system, where small chambers overlap in azimuth with the barrel toroid magnet coils, and large chambers are mounted in-between them. This makes it possible to track muons in the barrel-endcap overlap region using a combination of barrel and endcap chambers; other than that, there is no reason intrinsic to the endcap requiring such a layout.

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 µm thus corresponds to only $2.5 \cdot 10^{-6}$ (2.5 ppm) of the scale – scaled down proportionally to the situation of a microvertex detector with a typical length scale of 1 m or less, this would correspond to alignment at the level of a few microns.

1.3 Alignment Strategy

The basic concept of the alignment system is: (1) chambers have internal alignment sensors that monitor their distortions; (2) 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 system, the correction to the track sagitta that is due to chamber misalignment is measured directly with projective (i.e. pointing to the ATLAS interaction point) alignment lines linking 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, a global alignment system establishes in the endcap a reference grid using alignment bars, with alignment sensors measuring the relationship between the bars. Chamber positions are then related to the bar locations with proximity sensors. This information is used to provide corrections to the nominal chamber positions before calculating track sagittae (rather than directly providing sagitta corrections, like in the barrel).

Design and implementation of the 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 or on surface, and first results are also presented there. In section 6 we discuss some aspects of alignment that go beyond the optical alignment system.

1.4 Terms and Definitions

There are two different ways to make use of the data provided by an alignment system. The absolute concept is the straightforward one: to ask that the alignment system provides chamber positions at any time and without using any external references. The other concept is relative: assuming that at one moment the chamber positions are known (e.g. from the use of straight muon tracks in special runs with the toroidal magnetic field switched off [but the central solenoid magnet turned on, to select high-momentum muons], or using muons from Z^0 decays recorded in normal running), the alignment system is used only to follow variations of the positions from this point on. The advantage of the latter strategy is that all effects due to the limited absolute and relative accuracies of the sensors (see below) cancel out to first order; as these come typically from the calibration of the sensors and of their mounts this means that the importance of calibrations is greatly reduced. The disadvantage of the latter strategy is that knowing the initial chamber positions does not come for free, and the alternative methods to determine them have their own problems and uncertainties.

Not to be confused with the absolute and relative concept of alignment are absolute and relative measurements made by alignment sensors, and absolute and relative positions of objects determined by them. For an absolute measurement, a sensor uses no reference but itself, and by the absolute position of one object with respect to another we mean its position in a coordinate

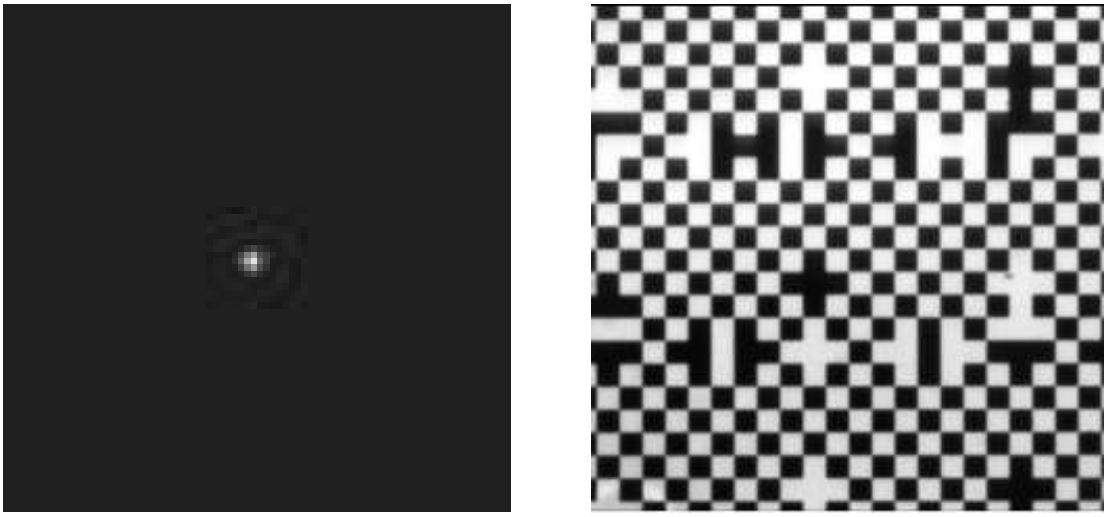

Fig. 2: Details of CCD images of the two sensor types used in the endcap alignment system: BCAMs (left) and RASNIKs (right). The BCAM image has been magnified with respect to the RASNIK image, in order to make the individual pixels visible. We flash the BCAM laser diodes of a pair sequentially, so that there is only one light spot on each CCD image.

system that is completely defined by the first object. A relative measurement, on the other hand, is a measurement in which a sensor uses at least one other sensor to define its reference, before making a measurement of a third one. Likewise, the relative position of one object with respect to two others is its position in a coordinate system defined by those two.

There are an absolute and a relative accuracy associated with the two types of measurements. Both of them are systematic errors in the sense that the error cannot be reduced by making many measurements instead of one, and that the deviation remains constant between measurements. The deviation is typically associated with the sensor or its mount, and thus changes when e.g. the sensor is replaced or its mount re-calibrated. In addition, sensors have an intrinsic resolution, which is typically a statistical error, but may have a systematic component in addition. Usually, the absolute accuracy of a sensor is greater than the relative one, and both are greater than the intrinsic resolution.

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 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. 2. 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 bars[4].

[4]Depending on 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.

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 threshold. Imaging a laser diode singly gives the bearing angle of the source relative to the optical axis defined by the lens-CCD combination. By imaging both diodes in a calibrated pair, the distance to the source and the rotation around the optical axis can also be determined.

A BCAM-type device consists of a camera (lens and CCD fixed together) 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. 3) 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, while the camera in the other BCAM in turn looks back at the light sources in the first camera. The camera consists of a plano-convex lens and a CCD image sensor[5] at a distance of 76 mm; the light sources are pairs of laser diodes with a separation of 16 mm. The CCD has 344×244 square pixels at 10 μm 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. They are bright (but not harmful to the human eye) and they are visible, with an output power of a few milliwatts, which makes diagnosis of BCAM problems far easier than with infrared light. The camera and the laser diodes, as well as the necessary electronics, mount directly into the BCAM body, and by construction are within 200 μm of their nominal positions. The center of the light spot on the CCD can be determined to about 0.4 μm (1/25 of the pixel width). To improve the resolution, BCAMs are assembled such that the spot is slightly de-focused, 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 (a 3-2-1 mount generally called a cone-slot-flat mount) as shown in fig. 3. The mount defines a local coordinate system for the BCAM mounted on it. The spheres are stainless-steel balls of 6.35 mm diameter, glued onto conical holes in an aluminum platform.

The calibration of a BCAM is best understood if one views the camera as a point in space (called the pivot point), near to, but not exactly at, the center of the lens, which all rays of light hitting the CCD pass through. Then 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 a spot on the CCD from any source of light.

[5] We use as CCD the TC255P (Texas Instruments), and as laser diodes the LDP65001E (Luminix) and the DL-3147-021 (Sanyo).

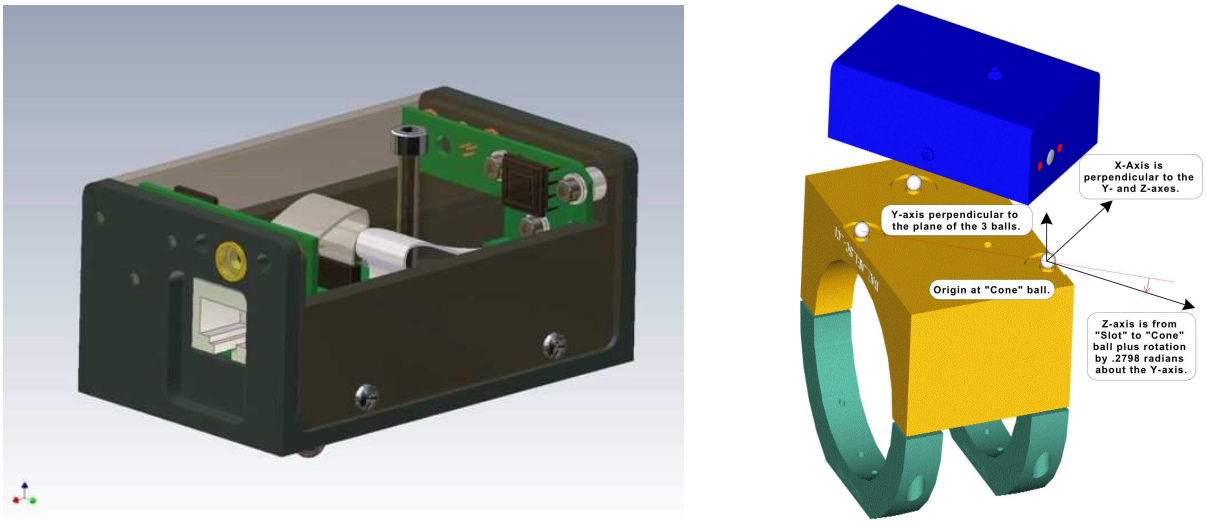

Fig. 3: Left: A black double-ended polar BCAM with the cover removed (there are several different types of BCAMs: "polar" and "azimuthal" BCAMs differ in the height of the chassis, black and blue versions are mirror-images of each other). On the left, the brass lens holder and the lens of one camera are visible, with two laser diodes in holes next to them; the CCD of the same 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. The screw at the center fixes the BCAM on its mount. Below the lens holder there is a socket to connect the CAT-5 cable for the readout system. Right: An 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.

The calibration system for BCAMs on an optical table is shown in fig. 4. The BCAM at the front of the table is mounted in a roll cage which allows it to be placed in four different orientations about the optic axis of the system (0°, 90°, 180°, and 270°). 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 know relative positions, at two different distances from the BCAM (in roughly the center of the table). The distance between the two source block positions is measured by an interferometer at the rear of the table. 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 bearing angle of the two laser diodes on the other BCAM with respect to its own optical axis, with an absolute accuracy of 50 µrad. It also measures the difference in bearing angles, i.e. the relative angle, of those two laser diodes, with an intrinsic resolution of $\sqrt{2}\cdot 5$ µrad. In a triplet of BCAMs on a straight line, an additional measurement can be made by each of the two outer BCAMs, measuring the relative bearing angle of two laser diodes on the two other BCAMs, again with a resolution of $\sqrt{2}\cdot 5$ µrad; this makes the triplet of BCAMs sensitive to deviations from straightness. An additional contribution to both the absolute and relative accuracy of 20 µm is due to the uncertainty of the laser diode positions and the pivot point position. 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

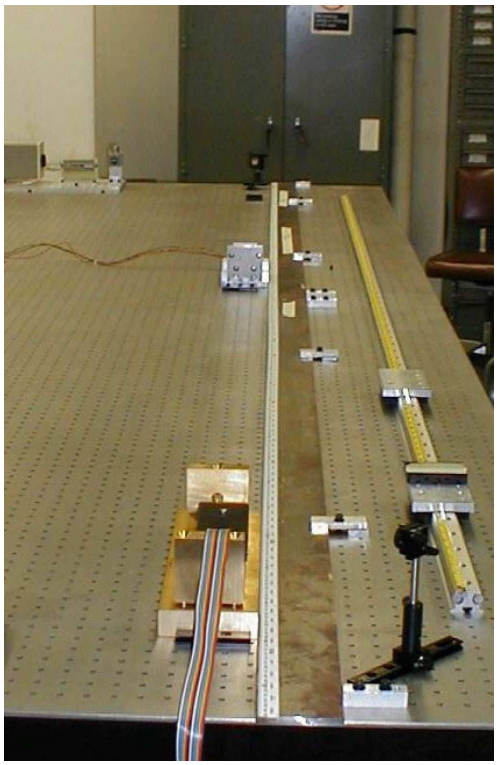

Fig. 4: The calibration stand for BCAMs. The BCAM in a roll cage is in the foreground, together with a 4-source block in the middle of the table and an interferometer at the rear.

measurement of its mount[6]; the intrinsic resolution is limited by atmospheric refraction and systematic effects of the centroid calculation method. The dynamic range of a BCAM camera is determined by the dimensions of the active area of the CCD and the CCD-lens distance: about $\pm 20\,\mathrm{mrad} \times \pm 15\,\mathrm{mrad}$.

2.1.2 Chamber Laser Sources

Chamber laser sources are mounted on sensor platforms glued onto the surfaces of MDT chambers[7]. They contain typically two pairs of laser diodes like the BCAMs, but no cameras; very few sources contain only one pair of diodes. Most laser sources are mounted on 3-ball kinematic mounts; only a few of them mount instead on V-shaped rail mounts (fig. 5).

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 can be determined.

[6]Note that in order to measure the inclination of the local coordinate system defined by the mount to 50 µrad, a measurement accuracy of 2–3 µm for each of the three spheres is required, which comes close to the limit of the CMMs used in this project.

[7]As will be explained in section 3.3.1, these sources serve to constrain a particular degree of freedom of the chambers where a pair of chambers rotates, in a correlated fashion, about an axis defined by the alignment bars. All other sensors are insensitive to this motion; in local jargon these sources are therefore called "saloon-door sources", because they monitor a behavior where the chambers swing like saloon doors.

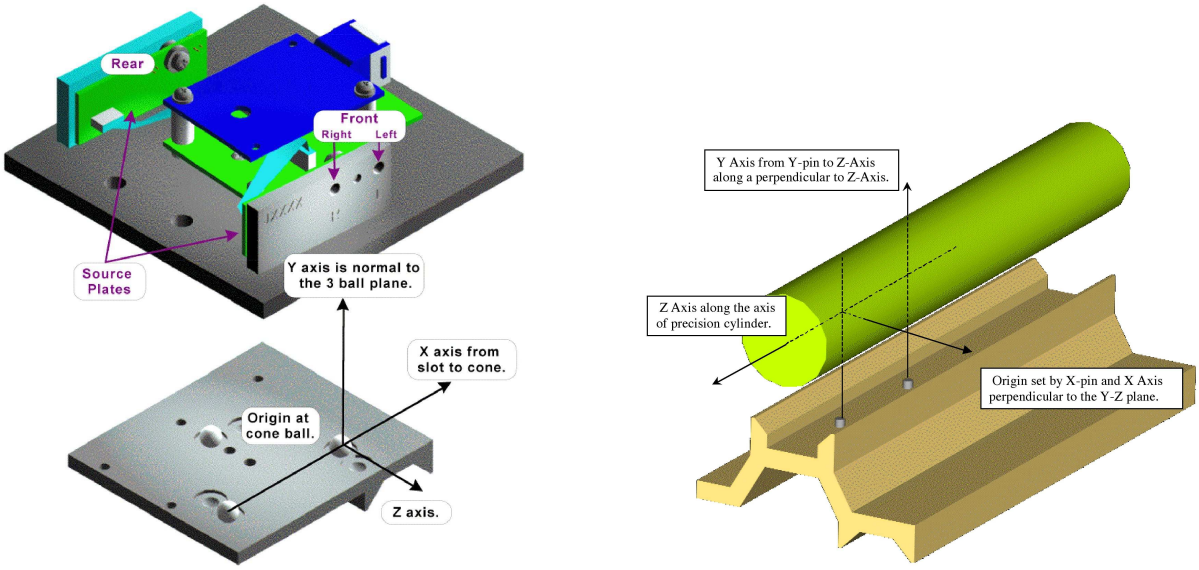

Fig. 5: Left: Chamber laser source with its 3-ball kinematic mount. The mount is made from extruded aluminum, designed to register on the 30 mm aluminum tubes of the MDT chambers. The local coordinate system defining the location of the laser source is also shown. The source itself has two pairs of laser diodes pointing in opposite directions. Right: The rail mount 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 shown above the rail. Very few chamber laser sources are mounted, for historical reasons, on this type of mount (those on EIL1, EML2, and EOS3 chambers); the mount has originally been designed for the proximity cameras (section 2.2.1).

2.1.3 3D Sources and 3D Cameras

One of the limitations 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, 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. 6). 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 develop a lens correction program, we elected to use a pin-hole camera so the position of the images was entirely determined by the geometry of the system: the 3D camera (fig. 6). Its body is an anodized aluminum tube of 195 mm length. The 150 µm pin-hole aperture is placed at a distance of 10–30 mm in front of the CCD, and has a glass cover to prevent dust from getting into it. For the different camera-source distances required, the CCD-aperture distance is adjusted so that the distance between laser images on the CCD is approximately equal to half the width of the CCD, and thus they all have approximately the same dynamic range of ±8 mm. The camera is calibrated by viewing a

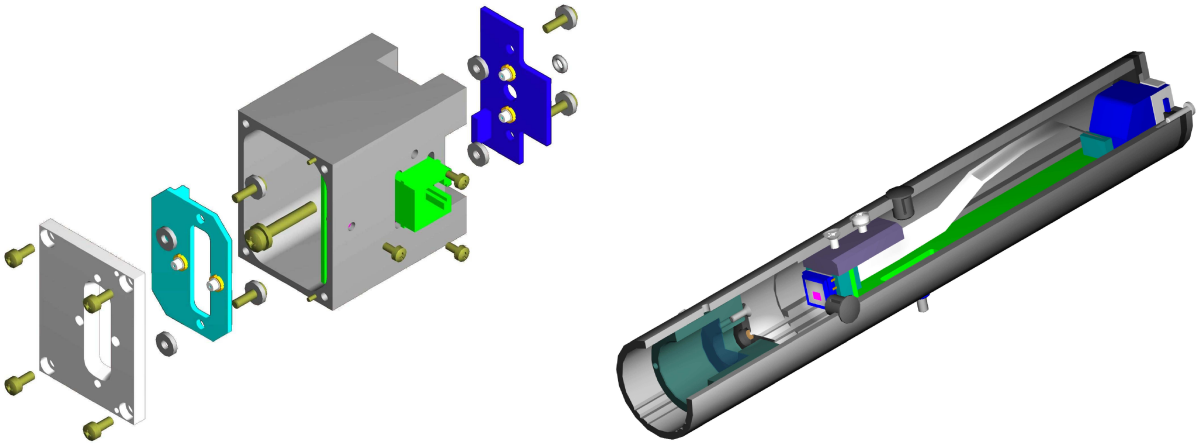

Fig. 6: Left: Exploded view of the 3D source for the CSC chambers. The two pairs of laser diodes are mounted on two different circuit boards. When assembled, the lasers are separated by 50 mm, and are viewed by a 3D camera from the left. Right: A CSC 3D camera. The light enters from the left, and the 150 µm aperture is covered by a glass microscope slide cover to prevent dust from entering. On the right are the CCD and the readout electronics.

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.

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 TGC chamber positions within each wheel in the two coordinates transverse to the beam. The task of the MDT alignment system is 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 are installed on the TGCs on each low-η and mid-η polar corridor (fig. 7). 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 it. The CCD dimensions are of the order of millimeters, while the mask has a size of the order of centimeters. 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

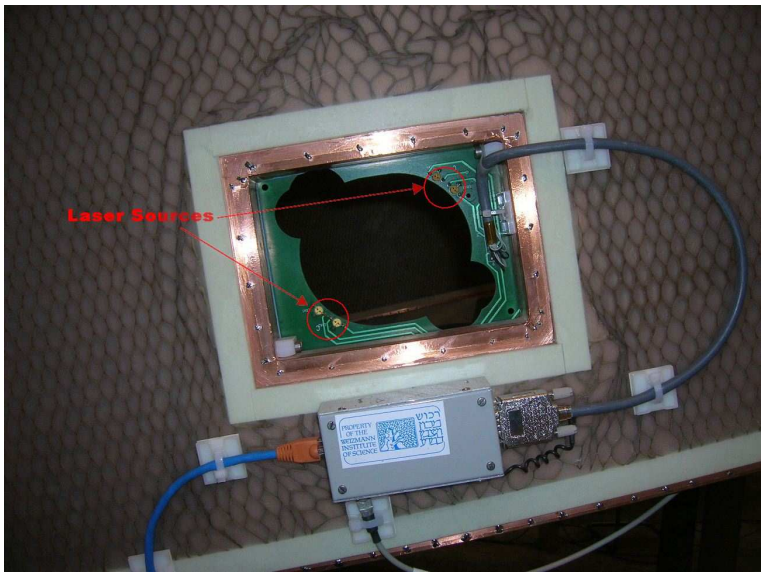

Fig. 7: TGC locator sources are positioned just outside of the polar alignment corridors (visible as the large elliptical cut-out in the center). In this view the sources are shown from the side away from the direction of the laser. Additional small cut-outs in the circuit boards are to allow the laser diodes in other wheels to be visible as well.

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.

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 (resolutions in the sub-micron range have been achieved under ideal conditions). For an equally spaced CCD-lens-mask system, a transverse motion of the lens relative to the line between the centers of the mask and CCD of 1 µm results in a 2 µm motion of the image on the CCD. If the lens is rigidly fixed to the CCD, then the device works as a proximity measurement between the two parts. Longitudinal position information comes from the measured magnification of the image, and the resolution is, under ideal conditions, about 10^{-5} times the distance between CCD and mask. A rotation of the mask relative to the CCD about the line between them can also be measured, with a resolution of better than 1 mrad.

Due to the large number of features in the images, the analysis of RASNIK images is highly redundant, and thus rather insensitive to image distortions like e.g. dust 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; for inconsistent cases, it is usually still possible to choose a best guess. By selecting only a part of the image to be analyzed, the effect of image distortions can be minimized further if needed.

RASNIK-type devices used in the endcap alignment system are either RASNIK sensors with single elements, where the mask, lens, and CCD are mounted separately 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).

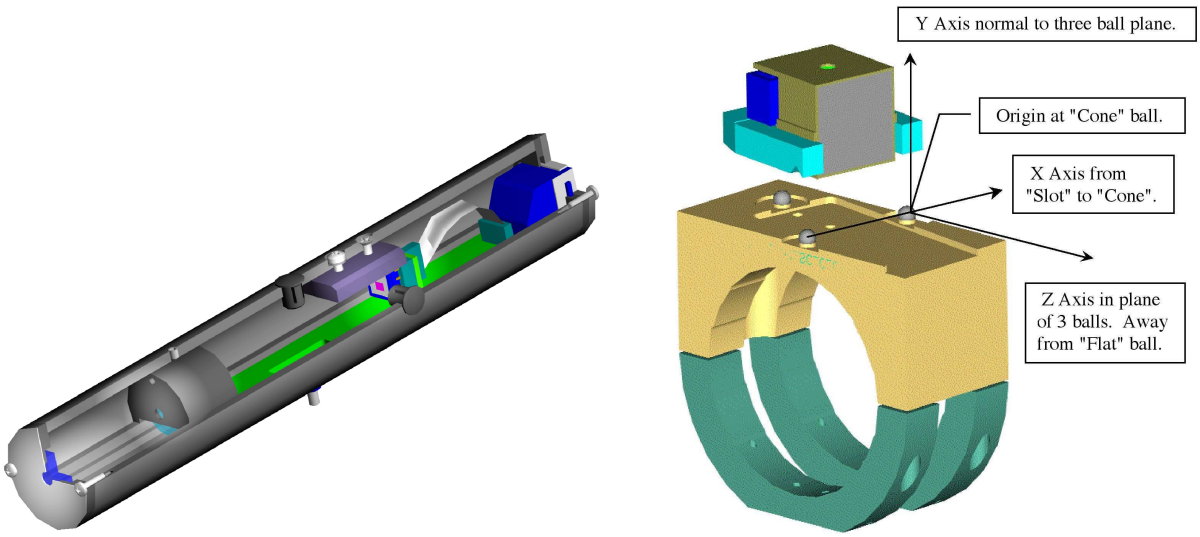

Fig. 8: Left: A proximity camera. The light enters from the left, and the lens is held in a cylindrical lens holder. On the right are the CCD and the readout electronics. Right: A bar-mounted proximity mask with its platform and bar clamp. The 3-ball kinematic mount defines the local coordinate system for the mask. There are also chamber-mounted proximity masks, which sit on a 3-ball mount like the one shown in fig. 5 and have a slightly modified housing and base plate.

2.2.1 Proximity Cameras and Masks

Proximity cameras are between 200 mm and 300 mm long anodized aluminum tubes containing the lens, the CCD, and the necessary electronics of the camera (fig. 8). 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 × 30 mm RASNIK mask, an array of infrared LEDs, and the electronics to back-illuminate the mask[8]. They sit on 3-ball kinematic mounts which are either clamped to alignment bars or glued to MDT chambers (fig. 8). Proximity cameras are tailored individually to the required camera-mask distance (typically between 100 mm and 400 mm) for a given location, by adjusting the positions of the lens and the CCD such that the mask is in focus and the image magnification is about 0.3.

Proximity cameras and masks are designed to provide, like BCAMs, an absolute accuracy of 20 µm and 50 µrad for position and rotation. We have developed a calibration stand for the proximity cameras and masks (fig. 9) that achieves this in only a few minutes per item, important since there are approximately 1200 camera-mask pairs in the system. The calibration of proximity cameras is achieved by taking images from a proximity mask at several different distances along the optical axis, over as large a range as possible. Proximity masks are calibrated by comparing measurements between them and a calibrated master-mask, and by CMM measurements of the surface of the RASNIK glass plate.

Early on, a perceived limitation of the proximity sensor was that the mask would not remain in focus over its likely range of motion along the optical axis of the camera of 10–20 mm. This is prevented by placing a 2–3 mm aperture in front of the lens, thus increasing the dynamic range to several centimeters, sufficient for our purposes. Transverse to the optical axis, the dynamic range is ±15 mm, given by the size of the RASNIK mask.

[8] We use as LEDs the HSDL-4400 (Hewlett-Packard).

Fig. 9: The calibration stand for the proximity sensors while being measured with a CMM. This stand is used for calibrating both cameras and masks. It consists of a rail mount for the camera, and about 30 3-ball mounts at different distances from it for a brass interface plate on which the mask is mounted.

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 and their calibration are discussed later, in sections 3.1.1 and 3.2.2.

2.3 Sensor Electronics

The system for controlling and reading out the alignment devices is called the LWDAQ (Long Wire Data AcQuisition) system [7]. It is a general-purpose data acquisition system designed for use with the endcap alignment system. It can retrieve three images per second from a camera at the other end of a 100 m cable, and is designed to be robust, convenient, and inexpensive.

The phrase "long-wire" refers to the CAT-5 cables used to connect devices to their LWDAQ multiplexers and LWDAQ drivers. These cables can be up to 100 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, with the multiplexer then connected to the driver with a single cable. The single cable running from the driver to the multiplexer provides the power and signals required by the multiplexer and all the devices connected to it. The LWDAQ requires no power source other than the driver. 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 many elements, 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. The returned data may be digital or analog.

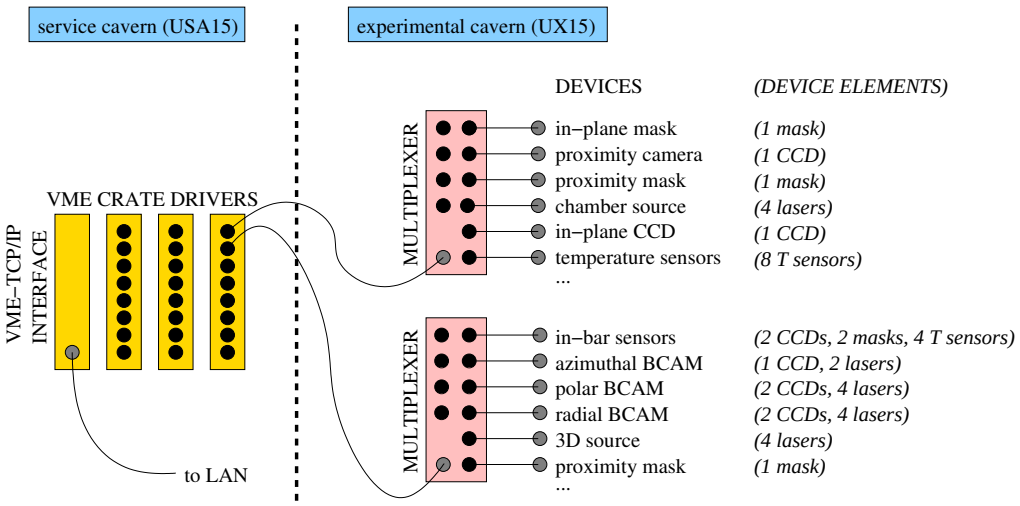

Fig. 10: LWDAQ architectural schematic. Root cables connect the LWDAQ multiplexers (installed on chambers and bars in the experimental cavern UX15) to the LWDAQ drivers (residing in VME crates in the service cavern USA15). Alignment devices in turn connect through branch cables to the multiplexers. Inside those devices that have multiple elements, short cables connect the elements to the main board.

The ATLAS LWDAQ driver resides in a VME crate, and provides eight CAT-5 sockets. In ATLAS, all devices will be connected through multiplexers to the drivers, never directly. The ATLAS detector requires nearly 800 ten-slot multiplexers, with 800 CAT-5 cables running out of the detector and into the service hall, where 100 VME-resident LWDAQ drivers receive them. In most places, active LWDAQ repeaters (acting also as patch panels) are inserted into the root cables, at a distance of about 10 m before the multiplexers. An architectural schematic of the LWDAQ system is show in fig. 10.

The endcap alignment system acquires two main types of data: temperatures and images. To acquire temperatures, the 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 starts by clearing the images sensor of residual charge. It flashes the light source. It reads the image pixels out of the sensor one by one, digitizes them, and stores them in memory. Table 1 lists the total number of different types of alignment devices and the images acquired by them.

The LWDAQ system was designed to solve several problems. One problem was how to provide electrical power to the sensors. We decided to deliver power to and signals to and from each alignment device through a single cable. One complication is that this has to be done through a multiplexer. Since at most two devices connected to the same multiplexer need to operate at one time, the power consumption can be reduced by putting to sleep any devices that are not active. With the sleeping power consumption reduced to tens of milliwatts, we can supply sufficient power to two active devices through a single 100 m solid-wire CAT-5 network cable, called "root cable". Devices are connected to their multiplexers with CAT-5 cables called "branch cables"; these can be stranded-core since their length is typically much shorter (not exceeding 10 m), and each branch cable needs to deliver power to only one device.

Another problem is the pervasive ionizing and neutron radiation to which the circuits will be subjected for the operating lifetime of the detector. The highest ionizing dose is approximately 10 krad, and the highest neutron dose is roughly 10^{12} neutrons (1 Tn). As almost all of the 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 1: 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.

severely irradiated devices will be inaccessible, they need to be resistant to radiation, rugged, and long-lived. Neutron radiation damages both the CCDs and the infrared LEDs we use in our alignment devices; the laser diodes do not exhibit any noticeable damage from the amount of radiation expected in ATLAS. The Texas Instruments TC255P CCD suffers an increase in dark current due to neutron radiation [8]. After absorbing 10 Tn, its pixels fill up with dark current in 50 ms. To capture images with these sensors, they have to be captured and read out in less than 50 ms. There are 83,936 pixels per image, and allowing at least 10 ms for image capture (exposure to light), the pixels have to be retrieved at a rate no slower than 2 million pixels per second. The Hewlett-Packard HSDL-4400 infrared LED is more resistant to neutron radiation than any other we know of. Nevertheless, it can lose up to 90% of its optical output power after 10 Tn. Using these diodes up to a dose of 10 Tn is possible only if the time for which an undamaged diode must be flashed to obtain an adequate image is less than 1 ms. In order to provide 1 ms flashes, the data acquisition system must be able to turn on and off light sources, and switch between one device and another, in a fraction of 1 ms.

To meet these requirements we kept the alignment devices simple, but at the same time versatile and fast. 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.

3 Alignment System

The design of the alignment system can be seen as a hierarchy with three levels: at level (1), 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 level (2), 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); and at level (3), 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, and of BCAMs on bars looking at laser sources on the chambers; similarly, 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).

3.1 Chamber Geometry

For a high-precision measurement of muons, there are several obvious requirements to be fulfilled by the precision chambers themselves: (1) the internal geometry of the chambers has to be accurately known, at least initially; (2) deviations from the initial geometry have to be monitored throughout the lifetime of the experiment; (3) the locations of alignment sensors with respect to the sensitive detector elements have to be accurately known.

3.1.1 MDT Chamber Deformations

The phrase "monitored" in "monitored drift tube chambers" refers to the in-plane alignment system, the RASNIKs which monitor the shape of the chamber from its birth throughout its lifetime. An MDT chamber consists of two sets of 3 (in EE, EM, and EO) or 4 (in EI) layers of 30 mm diameter single-wire drift tubes each, referred to as the two multi-layers, separated by a spacer support structure which also contains the in-plane alignment devices. The lengths of the tubes vary between around 0.9 m and 5.6 m. 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. 11). 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 μm r.m.s. with respect to the nominal positions.

After removal from the construction tooling, the expected chamber distortions, particularly for the largest chambers, are way beyond the required precision. The in-plane system monitors these distortions with four RASNIKs. Two RASNIK lines run parallel to the MDT tubes, two go diagonally across. 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[9]. They are read out ("zero-ed") 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.) While the in-plane system monitors the mechanical deformations of the spacer frame (and thus of the tube layers), determining the motion of the chamber wires due to thermal expansion requires temperature information in addition, which is provided by temperature sensors built into the 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

[9]Except for four EO chambers, where CCD and mask side have been accidentally swapped.

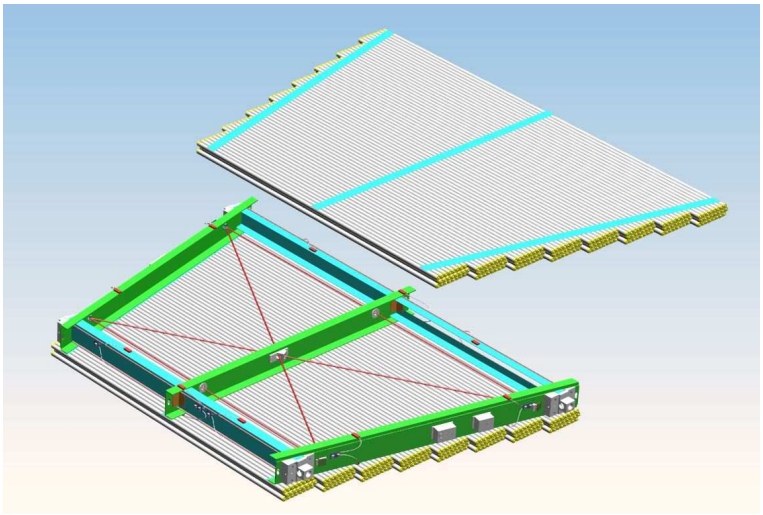

Fig. 11: Exploded view of an MDT chamber showing the top multi-layer separated from the rest of the chamber. The spacer frame contains the in-plane system; the RASNIK optical lines of the system are shown as red lines. The small red rectangles on the spacer frame show the locations of the temperature sensors on this particular chamber type (EML2). *(To be replaced by a better figure for the final version.)*

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): (1) a bow of the MDT tubes in the chamber plane; (2) a bow of the tubes out of the plane, varying from the short side to the long side; (3) a bow of the cross plates out of the plane, varying from the readout side to the high-voltage side; (4) a twist of the chamber, i.e. a rotation in opposite directions of the two outer cross plates around the tube direction; (5) a parallelogram-like deformation, i.e. a rotation in the same direction of all three cross plates in the plane; (6) a trapezoid-like deformation, i.e. a rotation in opposite directions of the two outer cross plates in the plane; (7) global expansion; (8) local expansions, different for the readout and high-voltage sides.

Typical observed magnitudes for these parameters are as large as a few hundred microns for some of them, particularly for (2) and (4), the out-of-plane bow of the tubes[10] and the twist; others appear to be immeasurably small, like (3), the out-of-plane bow of the cross plates. All the bow-type deformations are assumed to be parabolic, which is a sufficiently accurate approximation. The deformation modes (5) and (6) 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 tube multi-layers along the tube axes from the spacer structure. This allows for a parallelogram-like deformation of the spacer structure which is not constrained by the (much stiffer) MDT tube layers, and thus for a deformation mode that is seen by the in-plane system but does not affect the tube layers. 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

[10]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, bows within $\pm 100\,\mu m$ can be easily obtained. The non-concentricity of the MDT wire due to a tube bow distorts the electric field, and without correction a bow of $150\,\mu m$ would cause an apparent shift of the hits by up to $20\,\mu m$.

built-in bow of the tubes in the plane, i.e. a non-zero value of parameter (1); 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

Accurate knowledge is also required of the positions with respect to the MDT wires of the alignment sensors mounted on the chambers, which link the chambers to the global alignment system. 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), and 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 μm, 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 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 μm and 50 μrad, except for the coordinate along the tubes, where the requirement is only about 250 μm. The requirement for the survey target plates is more relaxed, at the level of 200 μm, as is the one for the magnetic-field sensor platforms, at the level of 1 mm and 1 mrad. The MDT chambers therefore need to be calibrated. The adopted procedure follows a two-step approach: (1) a coarse measurement of all three position coordinates, and, for rail platforms, also of the orientation, is performed by photogrammetry (section 3.1.2.1); (2) a fine measurement of the two coordinates transverse to the tubes, and of the two angles with respect to the tube axis, is made for 3-ball and rail platforms, by using two dedicated high-precision 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 [9] is a surveying technique where under-exposed photos of retro-reflective coded targets are taken from many different angles and positions with a commercial digital camera and a flash. A dedicated software[11] analyzes the photos and determines the positions of all the targets in an (arbitrary) coordinate system. This technique is widely 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), typically combined with theodolite measurements to improve the accuracy and to link the photogrammetry coordinate system to established 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 targets platforms do not require any adapters, 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

[11] We use the AICON 3D studio software.

Fig. 12: Photo of an MDT chamber during the photogrammetry measurement. Five coded targets on a cross at the center define the (arbitrary) 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), and the four single targets on small plates near the long and short sides are the survey targets. The magnetic-field sensor platforms are missing.

below. Aluminum scale bars[12] with lengths between 0.5 m and 1.5 m, carrying one survey target sticker at each end, were used to constrain the length scale of the photogrammetric measurement. With about 25 photos per chamber (fig. 12), an accuracy of better than 50–150 µm in all three coordinates was 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 calibration tools based on RASNIK sensors were designed for these measurements (fig. 13). They share the principle of operation, and differ only in their appearance which is adapted to the specific shape of the platform type. Each device 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 controlled movements (each of them re-measured with the CMM) of the sensor platform with respect to the fixed cylinders using a special version of the calibration stand, we found that the tool can follow those movements to better than 10 µm/40 µrad r.m.s. Photogrammetry and platform calibration data are analyzed together, 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

[12] 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 (not even measured) room temperature, without any corrections needed. The scale bars were calibrated using an optical CMM at 20 °C, the standard temperature for chamber production.

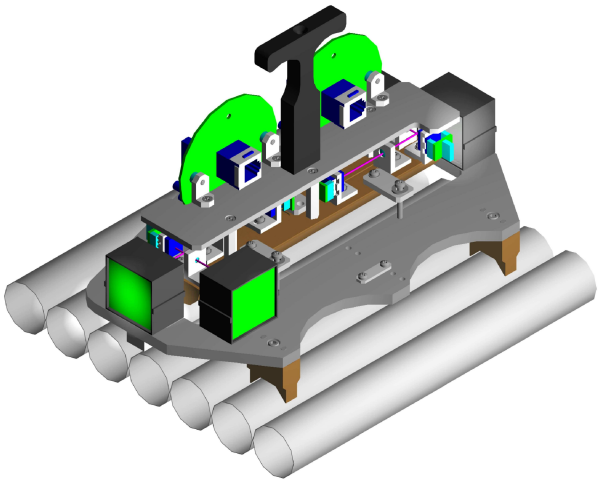

Fig. 13: The calibration tool for rail mounts. Like its 3-ball mount counterpart, it consist of two parts: (1) an inner part, made of brass, that sits (in precisely the same way as an alignment sensor later does) on the rail mount; (2) an outer part, made of aluminum, that sits kinematically on the MDT tubes near the mount, and picks up their position and orientation. Four RASNIKs, with the CCDs and lenses mounted on the inner part, and the masks mounted 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, and thus the MDT chamber local coordinate system).

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 (for an experienced person) about 45 minutes each per chamber.

A fundamental limitation of this calibration technique is its use of the tube surfaces as a reference, rather than of the MDT wires (which was impossible). The suitability of the tube surfaces as a high-precision reference has been questioned. From first consistency checks performed during ATLAS commissioning (section 5.2.2), we conclude that the target accuracy of 20 μm and 50 μrad is reached within a factor of 1–1.5, and that therefore the tube surfaces are just barely accurate enough for the purpose.

3.1.3 CSC Chamber Geometry

The large/small CSC chambers have four/two rail platforms and two 3-ball platforms each (fig. 14). 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 reference are four cylindrical pins of 6.35 mm diameter penetrating the chambers), and thus calibration tools like the ones designed for MDTs are not practical. Instead, CSC chambers are calibrated by measuring the sensor platforms with respect to the precision pins directly with a CMM. This is affordable given their small number, and technically possible thanks to their small size of less than 1.5 m × 1.2 m.

The layers composing a CSC chamber are mounted on a rigid steel structure, and are not expected to deform by any significant amount. They are therefore not equipped with a system

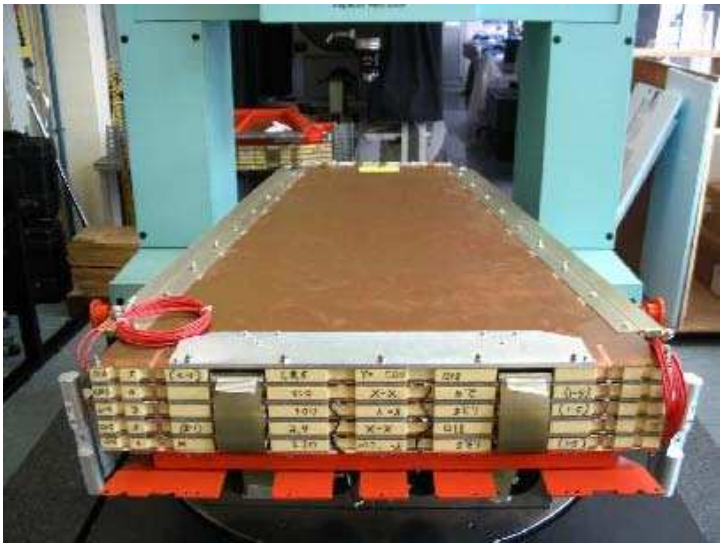

Fig. 14: A small CSC chamber mounted on the CMM for calibration. Visible at the sides are two cylindrical adapters temporarily mounted on the rail mounts for the CMM measurement; the two 3-ball mounts are embedded in the aluminum plate on the top.

equivalent to the MDT in-plane system to monitor deformations. They do have temperature sensors, however they are attached in places where they do not measure the temperature of the CSC active detector area, but rather the temperature of the readout electronics and the cooling water circuits.

3.2 The 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: (1) pairs of azimuthal BCAMs on bars in neighboring sectors within a wheel; (2) triplets of polar BCAMs, arranged approximately on a straight line, on bars in the same sector, in different wheels. There is a third related class, (3) triplets of radial BCAMs, all of which are mounted on the same bar; these do not determine the positions of bars, but are instead sensitive 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 useful debugging tool.

In each endcap, there are six types of alignment bars: (1) EI bars for the CSC, EIL1/2/3 and EIS1/2 chambers; (2) EIL4 bars for the EIL4 chambers; (3) EEL bars for the EEL1/2 chambers; EES bars for the EES1/2 chambers; EM bars for the EML1/2/3/4/5 and EMS1/2/3/4/5 chambers; EO bars for the EOL1/2/3/4/5/6 and EOS1/2/3/4/5/6 chambers. Except for the EE

Fig. 15: *(In the final version, this figure (or figures) will show the layout of polar and azimuthal BCAM lines, together with an appropriate figure caption.)*

region where the layout is more complicated, alignment bars are located approximately parallel, and next to, one of the sides[13] of the large MDT chambers, overlapping with the small MDT chambers. There are eight bars of each type, placed in even-numbered sectors (fig. 15).

Adjacent alignment bars in different wheels are linked by polar lines. A polar line consists of a triplet of polar BCAMs, each one on a different bar, arranged approximately (but never exactly) 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 $|\eta|$. While the latter two are at least approximately projective, at angles of about 30° and 45° from the beam line, respectively, the former one is exactly parallel to the beam line, passing through one out of eight empty tubes in the endcap toroid magnet cryostat at a radius of about 3 m, and thus is highly non-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. Each pair of adjacent bars is thus linked by two pairs of polar BCAMs (being part of those triplets), determining their absolute positions with respect to each other in addition to the relative measurement made by the triplets.

[13] Usually the high-voltage side, except for EIL4, where it is the readout side.

In-between each pair of adjacent alignment bars within a wheel there are as many azimuthal lines 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, and each one monitoring the position of its partner. Two azimuthal lines are sufficient to determine the positions of the two bars with respect to each other with good accuracy; 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 (i.e. by "closing the 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, the absolute positions of any two (not necessarily adjacent) alignment bars in the grid with respect to each other can be determined to about 300 µm. In addition, along polar lines, the relative position of one alignment bar transverse to the straight line defined by two others can be determined to about 30 µm. If polar lines were exactly projective (as in the barrel), the latter would be sufficient for a 40 µm alignment system; as this was not possible in the endcap due to environmental constraints (most prominently the muon endcap toroidal magnet cryostat), the former feature of the alignment system is equally important in order to reach the specification. In addition to alignment purposes, absolute chamber positions at a level of about 500 µm are also 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 [10, 11, 12] are large precision objects, consisting of up to 9.6 m long instrumented aluminum[14] tubes of 77/85 mm inner/outer diameter for the EEL, EES, EM, and EO bars, and an up to 2.4 m long tube of 72/80 mm inner/outer diameter for the EI and EIL4 bars (fig. 16). On the outside of the tube, platforms with 3-ball mounts are clamped to the bar, on which sensors (BCAMs, proximity masks, and 3D sources) can be mounted. On the inside, RASNIK elements and temperature sensors are mounted on disks and attached to the bar tube. 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 µm, their rotations (for BCAMs in particular) to better than 50 µrad. While the short EI and EIL4 bars can be regarded in practice as rigid bodies, the long EM and EO bars are rather floppy. 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[15].

However, the in-bar RASNIKs can provide information about the shape of the bar only at a few points along the bar (namely those points where CCDs, lenses, and masks are attached to the bar tube), and thus an interpolation procedure is required for points in-between these few measured points. It turns out that this is a highly non-trivial problem; for example, interpolation by arbitrary functions like splines or polynomials produces errors far beyond the specification. For a typical deformation of 4 mm, the 20 µm requirement corresponds to an interpolation accuracy as low as 0.5%. (For MDT chambers, where typical deformations are smaller by 1–2 orders

[14] Early on, a design based on carbon-fiber bars was considered, but discarded because of their uncontrollable dimensional changes with humidity.

[15] Prototypes of alignment bars as used in the H8 test (section 5.1) had three in-bar RASNIKs and 27 temperature sensors; EIL4 bars in ATLAS have only two in-bar RASNIKs and four temperature sensors.

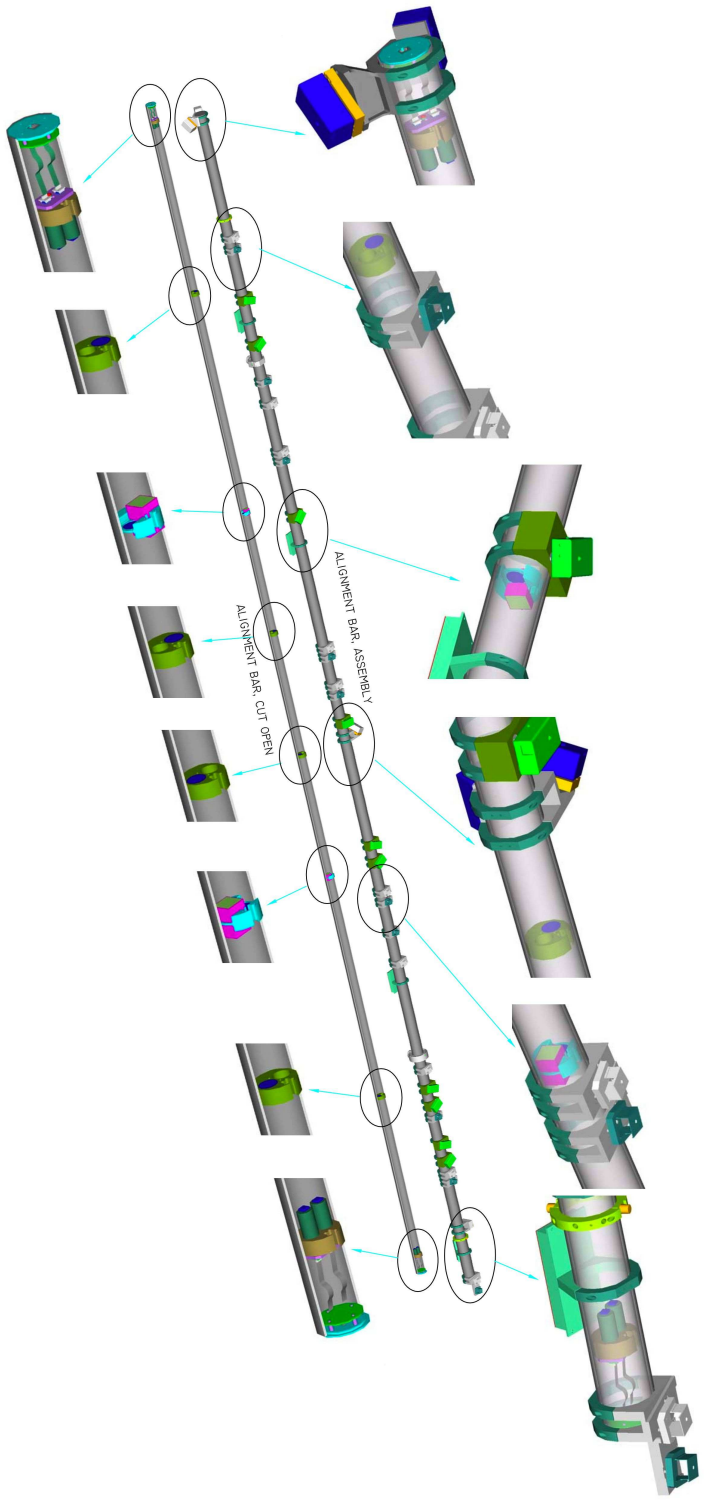

Fig. 16: Several views of an EM alignment bar. Clamped onto the bar tube are platforms with sensors mounted on them (magnified, from left to right, for the first three views): a polar and a radial BCAM, a proximity mask, and an azimuthal BCAM. Attached to the inside of the tube are disks holding components of the in-bar RASNIK system (again magnified, from left to right, for the first three views): two CCDs, one lens, one lens and two masks. In-bar temperature sensors are attached to these disks, too.

of magnitude, the quality of the interpolation is much less of an issue.) The solution 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/or continuous forces, and will be explained in the next section.

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 $6.5\,\mathrm{m} \times 1.4\,\mathrm{m} \times 1.2\,\mathrm{m}$. A local (relative) precision of $3\,\mu\mathrm{m}$ and a global (absolute) precision of $10\,\mu\mathrm{m}$ have been achieved despite the large measuring volume, thanks to the temperature in the climatized room housing the CMM being stabilized at $22.2 \pm 0.1\,^{\circ}\mathrm{C}$, and an elaborate and complex 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 (some of which had not been determined at all beforehand), 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 (to a good approximation 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.) In addition, the RASNIKs have to be calibrated as well, i.e. the rotations of the mask columns and rows relative to be bar coordinate system have to be determined, which 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), in order to have enough redundancy to detect a faulty measurement or other problems (we found in this way e.g. an in-bar RASNIK element that was not firmly attached to the bar tube). 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 measurements, each with about 100 sphere points, for each of these bars, taking about one week of working time. 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, bars are supported at several points during the measurement, 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 sensors mounted on it, or changes in the direction of gravity when the bar is inclined[16]). These forces are used to calculate the expected deformation of the bar. Any remaining unaccounted-for forces and effects (e.g. the weight of cables connecting to sensors, but also uncertainties of certain material parameters, such as the Young's modulus and the geometrical moment of inertia of the bar tube) are absorbed into three effective forces, assumed to be acting at the endpoints of the bar and at the center. These forces are determined using the RASNIK measurements, and serve as a measured correction to the expected deformation. A lot of care is taken to avoid or reduce any unaccounted-for forces: the quality of the interpolation depends on them being small, because a

[16] In ATLAS all bars are inclined, by angles between 14° and 76° with respect to the horizontal direction. In the H8 test (section 5.1), three bars were horizontal, the others inclined by 45°.

systematic error is made in absorbing unaccounted-for effects into effective forces. For example, the bars are kinematically supported near their Bessel points to minimize shape compliance, and all cables are routed along the bar to one of the support points, and only there they are attached to the large support structures of the wheels. Experimentally, the unaccounted-for forces are indeed small (e.g. of the order of a few Newton for the long EM and EO bars). The correction to the 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 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 typically 5% of the correction, and thus about 0.5% of the total deformation, well in line with the specification.

3.3 The 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 endcap to each other and to the eight alignment bars in the Small Wheel.

3.3.1 MDT Proximity System

Depending on the counting method, there are between 496 and 516 endcap MDT chambers: 64(+16) in the Small Wheels, 160 in the Big Wheels, 192 in the EO station, 16(+4) EIL4 chambers, and 64 EE chambers[17]. Numbers in parentheses refer to small attachments to "normal" MDT chambers, referred to as EIL3 (attached to EIL2 at larger radii) and EIL5 (attached to EIL4 in sectors 1 and 9 at smaller radii, and sometimes dubbed BIM chambers instead). These attachments do not have alignment sensors of their own; and while the EIL3 chambers are rigidly connected to their EIL2 counterparts in a well-defined way, the EIL5 chambers are only loosely attached in an ad-hoc fashion, and it is impossible to derive their absolute position from that of the corresponding EIL4 chamber.

The connection between normal MDT chambers 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. 17. Proximity cameras on one of the outer edges of the chambers view proximity masks on 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 small chambers there are two chamber laser sources that are 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.

The layout is slightly 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. In exceptional cases (EIL4 sectors 7 and 15), there are no proximity cameras at all on the chambers, and instead four laser sources on the four corners. The performance of both these layouts is comparable to the one of the standard layout.

[17] Owing to mechanical constraints, it is highly unlikely that the EEL2 chambers in sector 5 can ever be installed, reducing the total number effectively to 62.

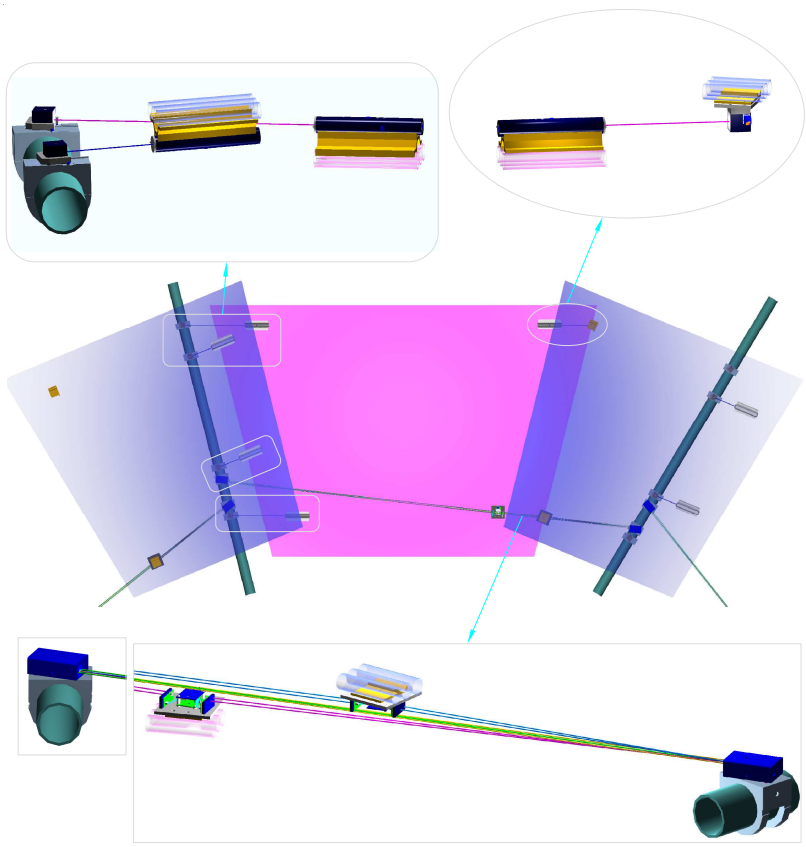

Fig. 17: Layout of the MDT proximity system in a typical pair of sectors. Shown in the center and on the right is a large-small MDT chamber pair, plus another small chamber to the left. Alignment bars overlap with the small chambers. Highlighted are the three components of the proximity system: two proximity 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 gap between large and small chamber layers in a wheel.

For all layouts, chamber laser sources are bi-directional only to increase the redundancy of the system, not to improve its accuracy: it is entirely sufficient if only one of the BCAMs on the azimuthal line sees the pair of laser diodes on its side of the chamber source. This feature is useful when chambers and bars are misaligned by relatively large amounts: while moving the partner BCAM and the laser diodes 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) is relatively easy, achieving the same in addition for the other azimuthal BCAM of the pair is often difficult. 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 any 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 axis of the nearest alignment bar, similar to the way saloon doors swing [13]. Like in the polar BCAM

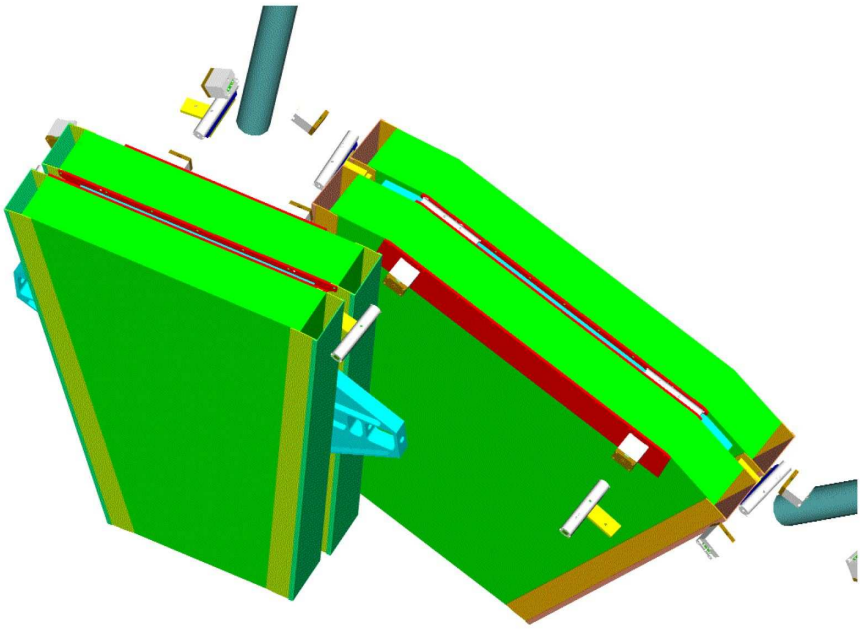

Fig. 18: Layout of the CSC proximity system. Shown is a large-small chamber pair with the nearest alignment bars. Two 3D cameras on the large chamber view 3D sources on the two bars; in addition, there is a camera on the large chamber viewing a source on the small chamber, and vice versa. All these sensors are mounted at the outer radius of the CSC region, where radiation levels are tolerable. *(To be replaced by a slightly modified figure for the final version.)*

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, though.

3.3.2 CSC Proximity System

The proximity system for the CSC chambers has a different design than the one for the MDT chambers. CSC chambers were chosen for the area near the beam pipe in the EI wheel because of the high radiation levels in this region. As a consequence, the sensors for the alignment system can only be placed at the outer edges of the CSC chambers, where the radiation is tolerable. Proximity cameras and masks would not have provided a sufficiently accurate and redundant handle on the chamber positions, and thus 3D cameras and sources were developed for this region.

The alignment system for the CSC chambers is shown in fig. 18. Large and small CSC chambers are linked by pairs of 3D cameras and sources. Each chamber is fitted with two cameras and two 3D sources, monitoring the positions of its two neighbors. In addition, there are two cameras mounted on each of the large CSCs viewing 3D sources attached at the high-η end of the EI alignment bars, thus establishing the link to the reference grid. There are no sensors linking the small chambers directly to the grid.

4 Data Handling

The data flow in the endcap alignment system is as follows: (1) the LWDAQ software takes care of sending commands to the alignment sensors and receiving images from them, and analyzes the images (section 4.1). The analysis results and other parameters are sent through (2) 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 (3) 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 (4) 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.

Each acquisition step produces an image or a temperature measurement. When the result of a step is an image, the LWDAQ software analyzes the image and produces an analysis result. By default, the software does not store or transmit the image (82 KB), 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 main forms of image analysis performed by the software. The simple form is the BCAM analysis, which finds and locates spots of light in an image. The complex form is the RASNIK analysis, which identifies a chessboard pattern and analyzes code squares. The BCAM analysis takes a few milliseconds to run, while the RASNIK analysis takes a few tens of milliseconds. The slowest step in the data acquisition process is not the analysis of the images or the exposure of the sensors, but the time it takes for the TCP/IP-VME interfaces to transfer the image data from the drivers over the local area network to the computers where they are analyzed. The transfer speed of the fastest interface cards used is 340 KB/s and each image is 82 KB, limiting the acquisition speed to three images per second. Without this limitation, up to 20 images per second could be acquired.

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

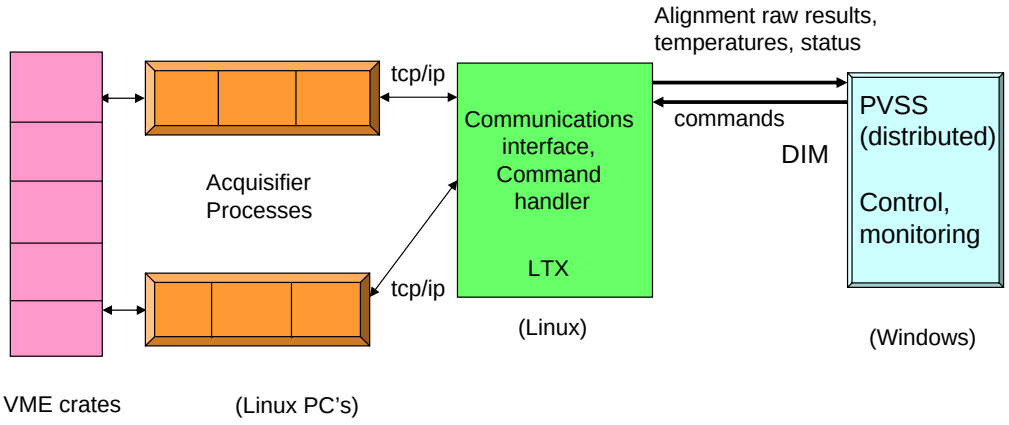

Fig. 19: LWDAQ Controls. The LTX process controls several LWDAQ instances, each of which in turn controls a group of sensors connected to drivers in a VME crate. The PVSS program performs supervisory control and monitoring.

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 Control and 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 software, i.e. of the system server and 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. 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. 19. 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.

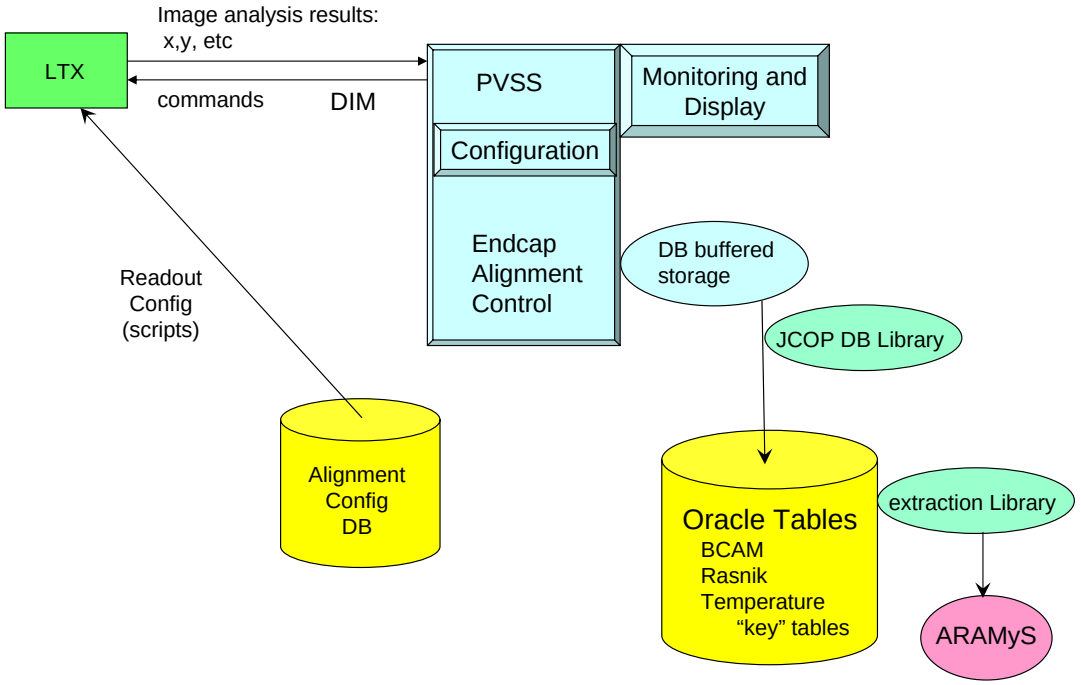

Fig. 20: Data flow in the alignment data acquisition. The LTX process sends scripts, produced from information in the alignment configuration database, to the LWDAQ instances. The readout values received from them by LTX are sent through PVSS to an Oracle database, from where they are read back for alignment reconstruction.

4.2.2 PVSS Software and Oracle Databases

The LTX control 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 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) [14] 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. 20.

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).

Several configuration panels and user interfaces were implemented in the PVSS framework: (1) one user interface configures the script names and port numbers for the parallel and sequential acquisifier processes; (2) a configuration panel permits changing the database buffer sizes and the database which receives the data; (3) 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; (4) a choice of continuous running or acquiring a fixed number of cycles is provided.

The PVSS control program displays the process status (Idle or Running) 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) [15, 16] has been developed. It has been extensively tested in simulations of the entire ATLAS muon endcap, and with real data, ranging from simple calibration stands with only two sensors to complete muon wheels of ATLAS, comprising up to 80 chambers, 8 alignment bars, and 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 relative locations. 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 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, and by minimizing numerically the quantity χ^2 defined as

$$\chi^2 = \sum_{i=1}^{n} \frac{(X_{i,\text{measured}} - X_{i,\text{expected}})^2}{\sigma^2_{i,\text{intrinsic}} + \sigma^2_{i,\text{mounting}}},$$

by variations of these assumed parameter values (the sum runs over all sensors and all the coordinates measured by each sensor).

An additional complication 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 determination of the values of the shape parameters is part of the minimization process, and is made possible by alignment sensors linking points on the same object.

The program ARAMyS implements the above concept. It consists of about 3,000 lines of code written in C. For the χ^2 minimization in the alignment fit, the standard package MINUIT [17] is used. The use of MINUIT, rather than of a private minimization routine, which might be slightly faster, 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. Also, MINUIT has been used and tested for decades, and is a mature and reliable piece of software.

ARAMyS is used in ATLAS to reconstruct the alignment of the endcap MDT and CSC chambers. A bar shape function and a chamber shape function were implemented, in order to take into account deformations and expansion of alignment bars and MDT chambers, respectively. Apart from these components, the program is not very endcap- or even ATLAS-specific, and is in principle suitable for any alignment problem.

Input to ARAMyS are: a description of the geometry, containing all the local coordinate systems, points, and sensors (including calibration data of both sensors and sensor mounts), as well as the alignment sensor measurements 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, and the chamber and bar deformation and expansion parameters. This information is stored in the ATLAS offline Oracle database, for use by the detector geometry model and the muon track reconstruction routines (section 4.4). For debugging purposes, ARAMyS provides additional information, 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 about $N = 9$ deformation and expansion parameters. Consequently, CPU performance becomes an issue for large alignment systems like the one of ATLAS. One approach to this problem is to make use of factorization: the property of an alignment system design that the alignment can be reconstructed by splitting the problem into subsets of objects to be aligned, without (noticeable) loss of accuracy or consistency. (A complementary 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 usually is the case for alignment problems.)

The endcap alignment system *is* approximately factorizing; the barrel system for instance 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)[18]. This process takes, on a 3 GHz dual-Pentium PC, only about 2 minutes. For the initial configuration of ATLAS, without EE chambers and bars, this time is further reduced to 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, which takes about 4 minutes for the initial configuration and reduces the expected false sagitta widths by about 10%. This 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 configuration.

Besides speeding up computations, the property of factorization also eases debugging, by allowing the user to start with analyzing small units, like bar or chamber shapes, and to put them together in larger entities only after fixing all the problems in the previous stage. On the other hand, fitting many parameters at the same time can help identifying single suspicious sensors which provide measurements that are incompatible with the rest: χ^2 minimization turns out to be surprisingly powerful in spotting a (small) number of inconsistent measurements, by converging to a minimum where these have significantly larger pulls than the others.

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 other 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.

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 design based on the foreseen network of sensors and their expected precision. For this application, the sensor readout values from an actual setup are replaced by the expected measurements as calculated by ARAMyS, which are randomly smeared by an amount that reflects the intrinsic resolution and the absolute and relative accuracy of the 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 mildly position-dependent false sagitta width in the range of 25–55 μm over the full endcap, well in line with the specification (fig. 21). To quantify the redundancy of the system, one can also simulate setups

[18]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.

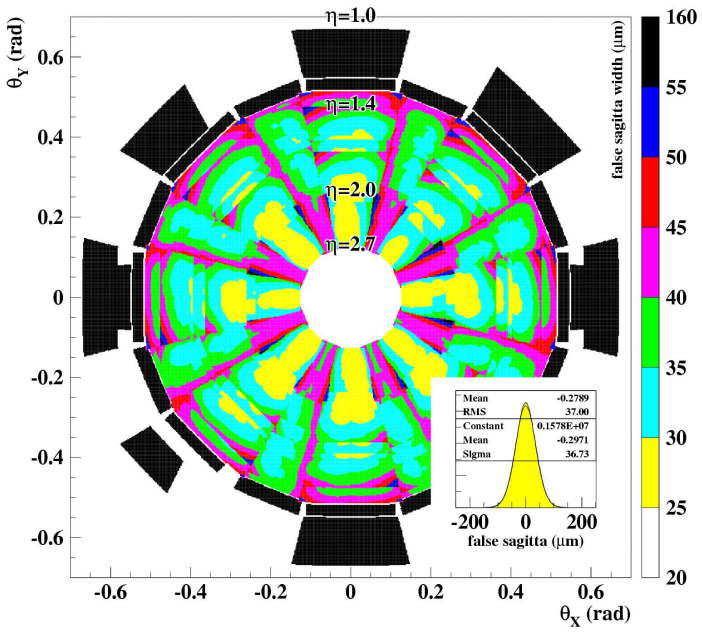

Fig. 21: Performance simulation of the endcap alignment system in the initial configuration, without EE chambers. Assuming design resolutions for all sensors, 600 statistically independent incarnations of the system were simulated. The small inset shows the histogrammed false sagitta distribution; the 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.

with one or several sensors switched off. Typically, the effect of a single sensor failing is rather small, not exceeding a (quite tolerable) local increase in the false sagitta width by a factor of two in the worst case; a striking 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 ATHENA reads the alignment data from the database, and provides the information to the offline muon track reconstruction packages MUONBOY and MOORE. 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 μm, 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): (1) 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; (2) 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 cleaner, but technically more difficult to implement, for instance because 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 (below 1 mm), 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 its and their functionality, respectively. The most relevant milestones reached so far are reported below. The ultimate test, reconstructing muon tracks in ATLAS runs without magnetic field and verifying them to be straight, will only be possible once the installation of the ATLAS muon spectrometer endcaps has been completed.

5.1 The H8 Test Beam Setup

In 2002–04, the ATLAS muon collaboration had 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 pairs of large and small MDT chambers, arranged such that they approximately represented a part of an octant of the EI, EM, and EO stations of ATLAS. In addition, there were six alignment bar prototypes (two 2.6 m short ones in the EI station, and four 9.6 m long ones in the EM and EO stations). 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 kinematically on large support structures of aluminum and steel, designed to be conceptually similar to the corresponding structures in ATLAS. The setup and the alignment and other test results have been described elsewhere [11, 18, 19, 20]; we will summarize here briefly only a few key results related to the endcap alignment system. Specifically, we will not discuss various little problems related to the sensor electronics that were identified and solved at this time, and the lessons that were learned about installing, commissioning, and running an alignment system with a large number of devices.

5.1.1 Relative Alignment Tests in H8

For testing the relative alignment concept in the H8 setup, two alternative approaches were taken: the use of a 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 of them provided an external reference at a level well below 40 µm to compare with the alignment system results.

5.1.1.1 Tests with the Muon Simulator

Conceptually, 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 directly extracted. As the positions of the camera and of the light sources with respect to the MDT wires are not known with high accuracy, this device does not provide an absolute reference and can track only variations of the alignment, i.e. 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. For simplicity we used BCAMs for the camera as well as for light sources, while a technology "more different" (ideally even a non-optical one) from that used in the alignment system itself might have been more desirable. Eight BCAMs, used as cameras, 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, 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 µm.

The comparison of sagitta variations in the bending direction (transverse to the wires, and thus inclined by 14° with respect to the horizontal direction[19] in the test setup) as reconstructed by the alignment system to those measured by the muon simulator yields an r.m.s. resolution of 14 µm over a continuous period of 2.5 days, during which daily temperature-induced sagitta variations of up to 500 µm were complemented by artificial variations of up to 5 mm from shifting and/or rotating different chambers (fig. 22). The total χ^2/ndf of the alignment fit varies typically in an acceptable range of 0.9–1.5. The observed resolution of the alignment system agrees well with the expectation from simulations for the relative alignment mode (the sensor absolute accuracies, which do not contribute here, dominate the absolute alignment accuracy of about 40 µm). The same comparison in the direction along the wires (not shown) yields an r.m.s. resolution of about 125 µm, safely below the required resolution of around 250 µm.

5.1.1.2 Tests with Straight Muon Tracks

Tests using straight muon tracks were performed using the same data set. Obviously, fewer data points can be obtained from tracks, as sufficient statistics has to be accumulated for each point. How tracks were reconstructed is beyond the scope of this paper. It was complicated compared to the ATLAS case by 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 which were vertical and thus rotated by 14° with respect to those of the large endcap chambers), as a stereo-measurement of the same tracks. This and other corrections to the track data are discussed elsewhere [19, 20]. The resulting agreement of the track sagittae after all corrections with the alignment system data is excellent (fig. 23).

[19] As the bending direction was approximately horizontal, the effect of refractive bending of light due to thermal gradients in the hall (expected to be vertical to a good approximation) on the optical measurements of this coordinate made by both the alignment system and the muon simulator was negligible.

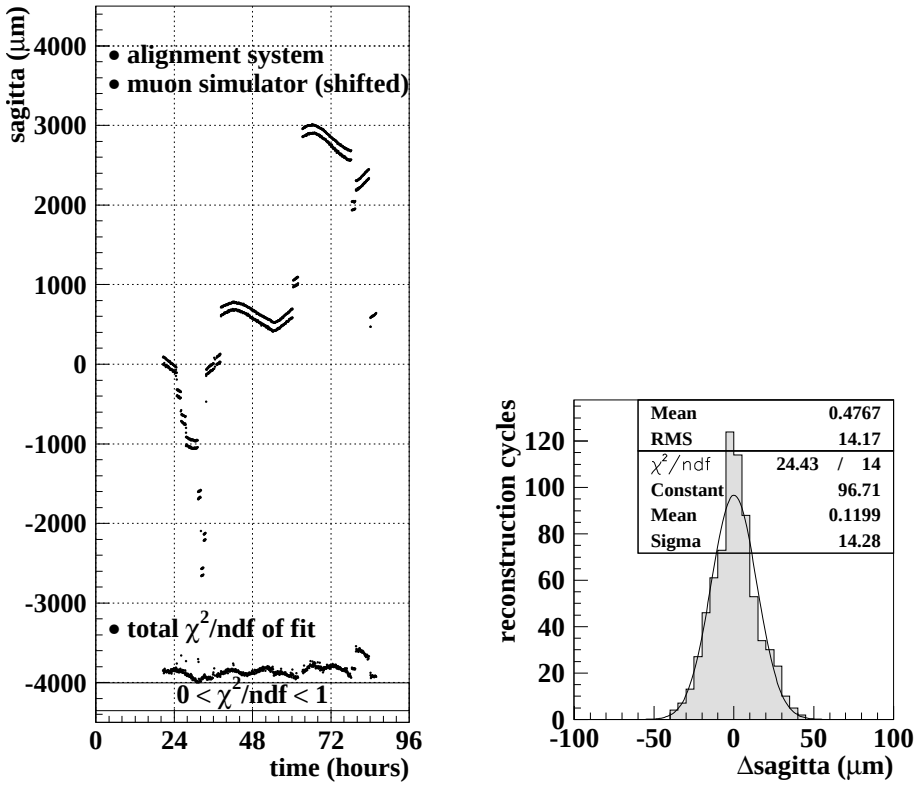

Fig. 22: Left: sagitta variations reconstructed by the alignment system (upper points) and measured by the muon simulator (lower points, shifted by 100 μm for clarity of presentation), over a period of 2.5 days. Smooth variations are temperature-induced, steps come from chamber shifts and rotations. The variation of the total χ^2/ndf of the alignment fit is indicated at the bottom; the horizontal lines show 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.

5.1.2 Absolute Alignment Tests in H8

For tests of the absolute alignment, an external reference with an accuracy below 40 μm is required in principle, which can be provided only by muon tracks. However, as the endcap alignment system is designed to provide absolute positions with an accuracy of typically about 300 μm, and only certain MDT chamber positions (namely those triplets that can be traversed by a muon track originating at the ATLAS interaction point) with respect to each other with an accuracy of 40 μm, it is possible to perform a limited, still reasonably meaningful, test by using an optical survey, with a typical accuracy around 300–500 μm. Absolute alignment tests using 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

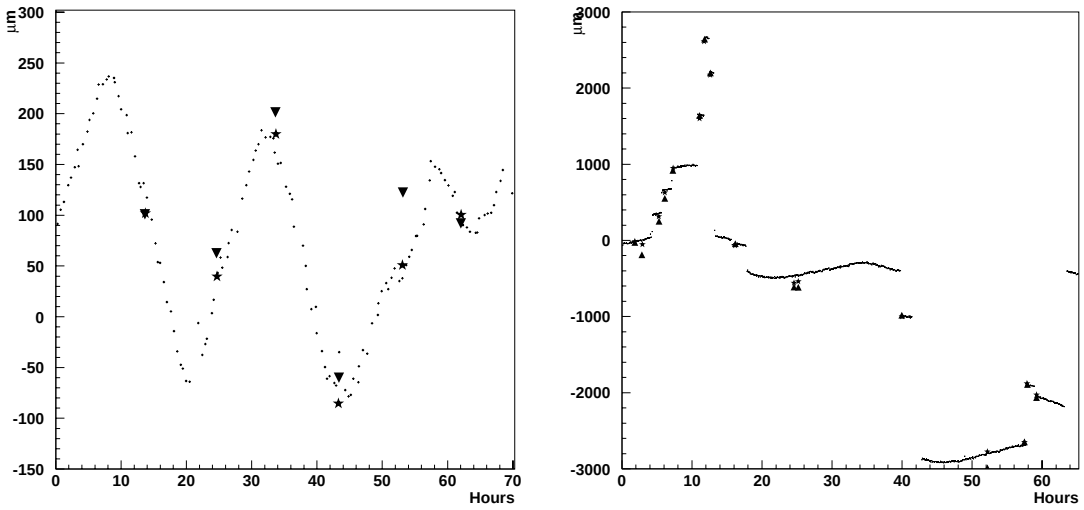

Fig. 23: Sagitta variations reconstructed by the alignment system compared with straight muon tracks from the SPS test beam. The prediction of the endcap alignment system is shown as small crosses, large triangles indicate the raw track measurements, and large stars show the track sagitta after various corrections. Left: sagitta variation due to normal day-night temperature variation. Right: sagitta variations for controlled chamber movements.

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 alignment tests using an optical survey as a reference were successful [11], 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, one alignment sensor (unfortunately precisely the most relevant one: the chamber laser source on the EML 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 had only been performed in early 2005, after the end of the test beam period, and after unmounting the device in question from the chamber, it was impossible to determine whether e.g. the device had been improperly seated on the mount, or whether the device calibration was wrong, and whether replacing it with a different device of the same type would have improved the result. 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.

The main lesson learned from this "successful failure" was that at every single step in the process from assembling a single sensor to commissioning the entire alignment system of an endcap, whatever could be checked at that point should be checked, and that no possible check should be postponed until later. More specifically, it was decided that: (1) for each type of alignment sensor and sensor mount there should be two independent calibration methods, shown to produce compatible results; (2) 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; (3) 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 a e.g. a single bar, chamber, sector, or wheel. With very few exceptions, all theses established principles were followed throughout the installation and commissioning process of the ATLAS alignment system.

5.2 Installation and Commissioning in ATLAS

After the end of the H8 test beam activities, the installation and commissioning of the ATLAS muon spectrometer endcaps started with the assembly of MDT chambers and alignment bars in sectors of what would eventually become the Big Wheels. Five chambers, and, in case of a large sector, one alignment bar, were assembled in a vertically stored sector frame; while small sectors did not allow any meaningful tests of alignment reconstruction beyond determining chamber shapes, the alignment bars in large sectors did permit to reconstruct the positions of the chambers with respect to the bar, and were compared to results from a survey. Later, the sectors were assembled into the 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.

5.2.1 Chamber and Bar Installation

Before a precision chamber or an alignment bar is installed on its support structure (sector or wheel), alignment sensors are mounted on it and their serial numbers and locations are recorded in the alignment configuration database. Prior to this, the calibration of the sensors has been checked on a calibration test stand. After mounting the sensors, their proper functioning (and that of the cables connecting them to the nearest multiplexer) 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 horizontal and supported by foam, and the bars held in cradles (near the kinematic mounts) on a cart. After completing the installations of bars and chambers in a sector or wheel, their multiplexers are connected through patch panels to driver boards, and all alignment sensors are read out. Malfunctioning devices (of which, in theory, there should be none at this point) and bad cables are identified and replaced, until a stably running system is reached.

At this point, many or most alignment sensors have their partner device (e.g. for a BCAM) or their mask (e.g. for a proximity camera) installed, and thus their measurements can be used to run ARAMyS in order to obtain reconstructed chamber and bar positions. In case of a fully assembled wheel, the alignment system is even complete, except for the polar BCAMs. At the time of writing, two Small Wheels and 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.

5.2.2 First Commissioning Results

A recurrent observation when starting up the alignment system of a newly assembled wheel for the first time was that many of the azimuthal lines were blocked, while most or all of the other

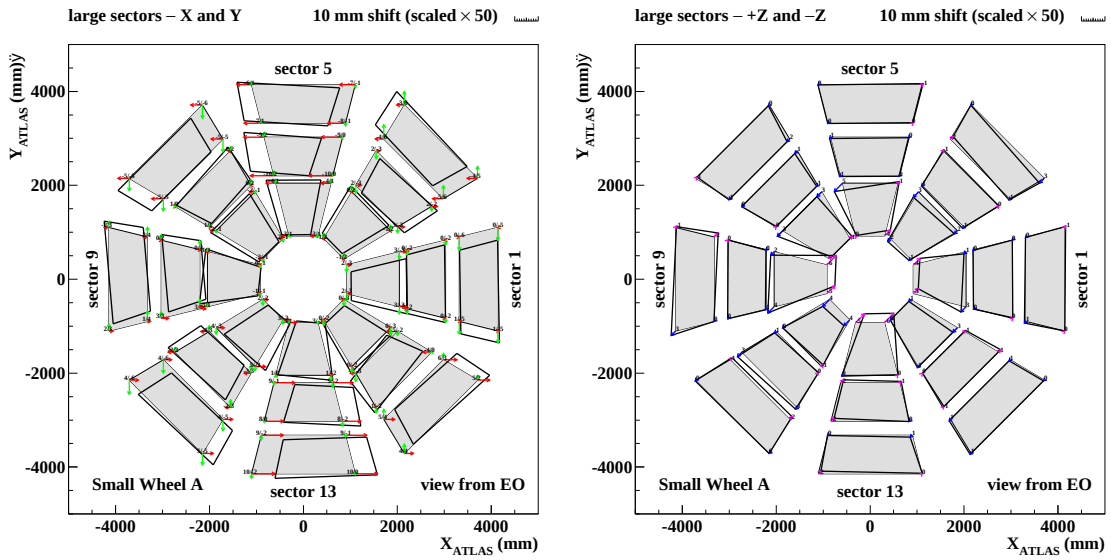

Fig. 24: Large precision chamber positions in the A-side Small Wheel as seen from the outside, reconstructed by the alignment system. Grey areas show the nominal positions, black contours the actual positions, where the displacements of the four corner-points have been exaggerated by a factor of 50 to 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, magenta arrows into it (right). Small chambers are not shown for clarity of presentation; their positioning quality is similar. The plots show the chamber positions after all necessary adjustments were made.

sensor types worked fine. This turns out to be 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. 17). The displacement of a chamber toward the azimuthal line, along the beam direction, 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 EM 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 functioning lines were, however, sufficient to reconstruct the positions of all chambers in the wheels to a level of a few hundreds of microns, perfectly sufficient to provide accurately the required adjustment data for re-positioning the affected chambers. Except for this particular displacement mode, the alignment system has proven to be, as it was designed, rather tolerant to mispositioning, with a dynamic range reaching (for proximity cameras) or even exceeding (for BCAMs) ± 15 mm in all directions. Essentially the only other parameter that needed to be adjusted frequently were the rotations of alignment bars around their axes, which are quite difficult to get within the required range just by accurate assembly. As an example, fig. 24 shows the large-chamber positions in the A-side EI wheel after all adjustments.

After all the required re-positioning, a constant fraction of 98–99% of all sensors working has been reached in all wheels. (Chamber laser sources are counted as 100% working if they are seen by at least one azimuthal BCAM.) The few 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 $\chi^2/\mathrm{ndf} = 1$ (for ndf, the number of degrees of freedom, of the order of a

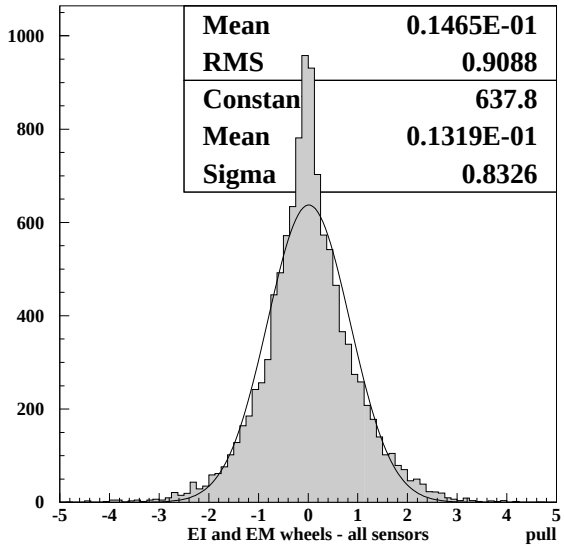

Fig. 25: Distribution of the pulls (contributions to χ^2) of all alignment sensors in the four fully assembled EI and EM wheels of ATLAS. Only a very small number (about 0.1%) of functioning sensors has been excluded from the fit or has pulls outside the plotted range. The black line shows the fit of a Gaussian (only to guide the eye, as one actually expects the distribution to be slightly non-Gaussian).

few thousands) is equivalent to the observed resolutions being in line with the design resolutions, for which case, as shown by the ARAMyS simulations, one expects a false sagitta width of about 40 µm (otherwise the expected width scales like $\sqrt{\chi^2/\mathrm{ndf}}$). In addition, one expects the individual contributions to χ^2 (the pull distribution) to follow an approximately (not exactly) Gaussian distribution. The observed pull distribution has quite the expected shape (fig. 25), and an r.m.s. width of about 0.9, where a value around 0.83 is expected from simulations for the particular mode in which the alignment fit was performed here (all bar/chamber deformations and positions fitted at the same time). The small difference is mostly due to sensor measurements with pulls being slightly on the high side, causing the region between 2–5σ to be a little bit more populated than would be appropriate; the width of the core distribution agrees rather well between data and simulation. Only a handful of isolated devices (of the order of 0.1%) have huge pulls without an apparent reason, and are tentatively excluded from the fit. These 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), showing that the alignment system measurements are 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 endcap alignment system at work, it performs fully according to our expectations so far. The last big missing piece are the polar BCAM lines, none of which can be tested at this point. However, technically, an azimuthal line (with a laser source in-between the two azimuthal BCAMs) is very similar to a polar line (with a third BCAM replacing the laser source), and it would be rather difficult to devise a mechanism that could cause a problem affecting only polar lines, but not the azimuthal lines that have already been demonstrated to work fine. 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 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 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. To some extent, the presence of vibrations could be detected by continuously reading out a single alignment device at the maximum speed of about 3 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 short time scales, this would affect not only the alignment system, but the functioning of the muon spectrometer as a whole.

The barrel and endcap alignment systems provide the individual internal alignment of the three parts of the spectrometer; they do not produce any 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 do not become obstructed when the detector is closed); 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. For all these chambers, the only method to align them with respect to the rest of the spectrometer is with muon tracks. Whether their positions are sufficiently stable in time for this to work remains to be seen.

Alignment with tracks is of course also foreseen for ATLAS running in general, in order to cross-check the alignment results obtained from the optical systems. Due to the poor resolution of MDT chambers in the coordinate along the wire, and due to the fact that muon tracks collected during ATLAS physics runs are typically projective (and thus have a limited angular spread), tracks are not sensitive to all the degrees of freedom of the chambers in the same way as the sensors are, which makes alignment using *only* tracks challenging. One possibility to overcome this limitation to some extent is to use cosmic muons; eventually, however, the alignment of the muon spectrometer should be done by performing a global fit using sensor and track data at the same time, minimizing a global χ^2 with contributions from optical sensors and from muon tracks, and from both the barrel and the endcaps, thus obtaining the best possible alignment of the entire spectrometer by exploiting all the available information at the same time. Of the many challenges that this approach poses, CPU time is likely to be a non-negligible one.

Eventually, the alignment of the muon spectrometer with respect to the inner tracking detectors will also have to be determined, albeit being less of an issue in ATLAS than it is, for instance, in the CMS experiment at the LHC, because of the superior stand-alone momentum reconstruction capability of the ATLAS muon spectrometer. This is one of the reasons why, unlike CMS, 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 μm at the start-up of ATLAS is well within reach. The installation and commissioning of the system in the ATLAS cavern is well underway, and has produced encouraging results so far.

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, A. Ostaptchouk (RWTH Aachen, formerly MPI Munich) did much of the early conceptual design and simulation work, T. Klioutchnikova's (CERN) work in finding and fixing countless interferences between optical alignment lines and physical obstacles was invaluable, and M. Bradshaw (Brandeis) helped develop and improve the sensor electronics. 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. D. Levin (University of Michigan) performed the tracking part of the H8 data analysis leading to fig. 23. The CERN metrology group performed the CMM measurements of CSC chambers and made the photo in fig. 14. 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 set up the system for photogrammetry measurements of MDT chambers, and were helpful and willing to provide additional information whenever needed. 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 M. Ketchum, D. Nash, D. Pomeroy, B. Way, as well as S. Pintaric (CERN) and N. Nation (Boston University), together performed photogrammetry and calibration measurements on hundreds of MDT chambers. The Washington students C. Dowell, M. Olson, and B. Reynolds contributed to the data acquisition and control software. Many students at Brandeis *(names missing here)* assembled and calibrated the thousands of alignment devices. The workshops at Tufts and Harvard University produced parts for the alignment system. 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] 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).

[9] C. Lasseur, Three Years of Digital Photogrammetry at CERN, Proceedings of the 6th International Workshop on Accelerator Alignment, Grenoble, France (1999).

[10] 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`.

[11] C. Amelung et al., Reference Bars for the Alignment of the ATLAS Muon Spectrometer, Nucl. Instr. Methods **A 55** (2005) 36.

[12] C. Amelung and A. Schricker, Calibration of Alignment Bars for the ATLAS Muon Spectrometer Endcap, ATLAS Note ATL-MUON-2005-009 (2002).

[13] C. Amelung, Muon Spectrometer Alignment in ATLAS: Reaching the Design Performance in the Endcap, ATLAS Note ATL-MUON-2001-018 (2001).

[14] 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`.

[15] 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`.

[16] C. Amelung, ARAMyS – Alignment Reconstruction Software for the ATLAS Muon Spectrometer, Nucl. Phys. **B** (Proc. Suppl.) **172** (2007) 132.

[17] F. James and M. Roos, Comp. Phys. Comm. **10** (1975) 343; MINUIT Reference Manual.

[18] 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).

[19] 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).

[20] C. Adorisio et al., Study of the ATLAS MDT Spectrometer using High Energy CERN Combined Test Beam Data, ATLAS Note ATL-COM-MUON-2007-013 (2007), to be submitted to NIM.