Alignment of LHCb tracking stations with tracks fitted with a Kalman filter

Louis Nicolas, Adlène Hicheur, Matthew Needham, Jan Amoraal, Wouter Hulsbergen and Gerhard Raven


Fig. 1.   The LHCb detector.

Abstract—The LHCb detector, operating at the Large Hadron Collider at CERN, is a single arm spectrometer optimised for the detection of the forward b anti-b production for b physics studies. The reconstruction of vertices and tracks is done by silicon micro-strips and gaseous straw-tube based detectors. In order to achieve good mass resolution for resonances the tracking detectors should be aligned to a precision of the order of ten microns. A software framework has been developed to achieve these goals and has been tested in various configurations. After a description of the software, we present alignment results and show in particular for the first time that a global χ^2 solving for alignment using a locally parametrised track trajectory can be achieved.

I. Introduction

THE LHCb tracking system (see Fig. 1) consists of a Vertex Locator (VELO) and a large area silicon detector (Trigger Tracker) located upstream of the dipole magnet of the experiment and three stations (the T-stations) located downstream of the magnet. The latter are divided into the Inner and Outer Tracker (IT and OT), The three IT+OT stations are composed of 3 stations × 4 boxes × 4 layers × 7 = 336 silicon micro-strips IT ladders and 3 stations × 4 layers × 22 modules = 264 straw-tube OT modules, respectively.

Performance studies are concentrated on so called 'Long tracks' that traverse the entire spectrometer [1]. For these tracks a momentum resolution of between $3 - 4‰$ is achieved with simulated data. To achieve this precision with real data

the position of the sensors has to be determined with an accuracy well below the hit resolution, 60 μm and 200 μm, for IT and OT, respectively.

II. Alignment method

The method used to fit for the alignment parameters is the minimisation of a chi-square, based on residuals **r** of reconstructed tracks, with respect to both alignment and track parameters (global chi-square or closed form, [2]–[4]). The minimisation leads to an expression of the form for the alignment parameters **a**:

$$\mathbf{a} = -\left(\sum_t \frac{\partial \mathbf{r}_t}{\partial a}^T V^{-1}(V - HCH^T)V^{-1}\frac{\partial \mathbf{r}_t}{\partial a} \right)^{-1} \times$$
$$\left(\sum_t \frac{\partial \mathbf{r}_t}{\partial a}^T V^{-1}(V - HCH^T)V^{-1}\mathbf{r}_t \right)$$

The sum is made over the tracks in the sample. $\frac{\partial \mathbf{r}_t}{\partial a}$ represents the derivatives of residuals with respect to the alignment parameters. V is the measurements (diagonal) error matrix. $V - HCH^T$ is the covariance matrix of the residuals, H is the derivative of the residuals with respect to the track parameters and C is the track parameter covariance matrix. The novelty in this study is that we use a Kalman filter track fit in the alignment procedure, following the recipe in [5] to compute the global covariance matrix C. This allows us to use the standard LHCb track fitting tools and track model that are based on a Kalman filter in the alignment procedure, including all the complexity of such a track model (eg. energy loss correction, multiple scattering, etc.).

III. Software and procedure

The software is developed in the LHCb Gaudi framework [6]. The core alignment algorithm, which processes the track sample and fills the algebra objects, is used in conjunction with a set of tools performing tasks such as track selection in input, spectral analysis, regularisation of the problem and solving, in output.

The general sequence starts with pattern recognition, followed by track fit, track quality selection, computation of the alignment derivatives, solving and update of the alignment constants. To obtain an analytic solution, the method assumes a linear dependence of the residuals with respect to alignment parameters. To resolve non-linearities, we iterate over the sequence until convergence. Convergence is obtained when the χ^2 shows no significant change in subsequent iterations.

Fig. 2. Distribution of the track fit χ^2/dof for selected tracks using the ideal geometry.


Fig. 3. Distribution of the track fit χ^2/dof for selected tracks using the misaligned geometry.

The framework is flexible enough to allow several options for alignment: one can align for groups of sensors or decide to turn off degrees of freedom which we are not sensitive to.

IV. EVENT AND TRACK SELECTION

As described in section II, the alignment procedure minimises the total χ^2. Hence, it is sensitive to the effect of ghost tracks and particles with kinks in their trajectory due to hadronic interactions in the detector material. If these tracks are used in the alignment procedure they can lead to the algorithm converging to a false minimum. Therefore, it is important to obtain a sample of tracks with a minimum contamination of ghost and other bad tracks.

A. Cuts on track quality

One way of dealing with bad tracks is to cut on the track fit χ^2. For the long tracks it is also possible to cut on the track fit match χ^2 (χ_m^2) defined as

$$\chi_m^2 = \chi_{tot}^2 - \chi_T^2 - \chi_{Velo}^2 \qquad (1)$$

where χ_T^2 and χ_{Velo}^2 are the χ^2 of the tracking station segment, respectively the Velo segment of the track. These variables need to be used with care as selection criteria since their quality is degraded in a misaligned detector. Therefore, an evolving strategy has been developed. In the first iteration of the procedure a weak cut on the χ^2 is made which is then tightened in subsequent iterations. To develop this strategy the distributions of selected, not selected (ie tracks which are known to have interactions) and ghosts has been studied with the ideal geometry and the misalignment strategy described in section V-A. Figs. 2 and 3 show the distribution of the χ^2/dof in both the misaligned and the ideal case respectively. It can be seen in fig. 3 that cutting on the track fit χ^2/dof at 100 will not bias the sample of selected tracks. Therefore, this is chosen as the starting value in the iterative procedure. From the plot in Fig. 2 it seems a reasonable cut to apply at the end of the scheme is χ^2/dof < 10.


Fig. 4. Distribution of the track fit match χ^2 for selected tracks using the ideal geometry.


Fig. 5. Distribution of the track fit match χ^2 for selected tracks using the misaligned geometry.

Figs. 5 and 4 show the same plots for the case of χ_m^2. In this case an initial cut value of 100 is reasonable. This is reduced to 30 during the iteration procedure.

Geometry	Cut	Efficiency	Ghost rate
	No cut	100 %	18.01 %
Misaligned	χ^2/dof < 100	98.65 %	17.09 %
	fit match χ^2 < 100	93.25 %	13.89 %
	No cut	100 %	17.92 %
Aligned	χ^2/dof < 10	93.01 %	13.56 %
	fit match χ^2 < 30	84.60 %	9.26 %

Cut	Efficiency	Ghost rate
No cut	100%	17.92 %
# IT Clusters < 400	46.63 %	12.57 %
# Velo Clusters < 900	35.03 %	12.07 %
P > 10 GeV	36.26 %	18.08 %
η < 5.2	94.87 %	21.34 %

Table I summarises the performance of the cuts described above in removing ghosts for long tracks.

B. Other cuts

Several other cuts were used to reduce the ghost rate in order to obtain better convergence of the algorithm. It has been shown for the Long Tracking [7] that most ghosts are due to incorrect matches between the track segments reconstructed upstream and downstream of the spectrometer magnet. These incorrect assignments (ghosts) occur most often when there are many reconstructed segments in either of the sub-detectors. In order to reduce this source of ghosts, events with many Velo or Inner Tracker clusters are rejected. High Outer Tracker occupancy is not as big a problem as the IT case because of the lower track density. The cuts are chosen to reduce the ghost rate to an acceptable level. For the time being, 400 IT clusters and 900 Velo clusters are chosen. These cuts are summarised in Table II.

Poorly reconstructed tracks also have a negative impact on the alignment results, as discussed in section IV-A. Tracks with large multiple scattering in the detector will tend to be badly fitted. As this process most strongly effects low energy tracks, it can be reduced by cutting on the track momentum at 10 GeV (see Table II).

Finally, the ghost rate can be further reduced by cutting on the track pseudo-rapidity. The LHCb acceptance in this variable is between 1.9 and 4.9. Ghosts tend to be at large pseudo-rapidity in the Velo [8] where the occupancy is highest. A reasonable cut is at η < 5.2.

V. MONTE CARLO STUDIES

During the LHC start-up phase first alignment studies will be made with cosmic and beam-gas events with zero field. As soon as the machine commissioning is finished, proton–proton collisions at $\sqrt{s} = 14$ TeV will be the main focus.

Detector		DoF	Amplitude
	boxes	TX [mm]	1.0
IT	layers	TX [mm]	0.1
	ladders	TX [mm]	0.05 (*)
OT layers		TX [mm]	1.0
		RZ [mrad]	0.15

These collisions will allow the detector to be used to its full extent. The magnet will be switched on at this point allowing a momentum measurement. In addition, the Velo will be usable in its closed position, which may not be the case in the first phase. This in turn will give the possibility to use long tracks (ie tracks going through the VELO, the trigger tracker and the tracking stations) for the alignment. Since for the VELO a precise standalone alignment is possible (partly due to a high number of measurement planes), it is used as a reference point for the alignment of the inner and outer tracker.

The sections below describe a scenario that has been studied using long tracks of minimum bias events produced and reconstructed with the magnet on.

A. Scenario

A sample of 20 thousand minimum bias events reconstructed with the default geometry has been used for this study. The data was simulated with a centre-of-mass energy of 14 TeV and the magnetic field set to its nominal value. During the procedure, the tracks are refitted using a misaligned database. In the scenario studied here, only translations along the local x axis (measurement direction) have been applied at the level of the IT boxes, IT and OT layers and IT ladders individually. With all these elements being aligned only along the measurement direction, this gives 408 degrees of freedom to align for. The amplitude of these misalignments follows a flat distribution of 1 mm in width for the IT boxes, 100 μm and 1 mm for the IT and OT layers respectively and 50 μm for the IT ladders except for the innermost ones (the most illuminated during a run) which are misaligned with a value of up to 100 μm. In addition, the OT layers are rotated around the beam-axis, with amplitudes in a flat distribution up to 150 μrad. The summary of these values can be found in table III.

This scenario with total misalignments at the ladder level of 1 mm or more complicates the direct alignment of the ladders. Therefore for IT a multistep approach was used, going deeper in the geometry tree at each step. At first, only the IT boxes are aligned, removing the constraint to the outer tracker by requiring tracks with no OT hits. At the end of this first job, the alignment constants are written and used in the second job where the IT layers are aligned, requiring again no OT hits on the tracks. In parallel to that this, the OT layers are also aligned, with no constraint to the IT (requiring no IT hits on the tracks). Once the IT and OT layers have been aligned

Fig. 6. Convergence of the normalised total track χ^2. The number of tracks is reduced during the iterative process (due to the cut on the χ^2) from 293'000 to 284'000.

Fig. 7. (output - input) Tx misalignment for each of the IT ladders and OT layers.

separately, they are again aligned together. This step is used to align the two detectors with respect to each other after their separate inner-alignment. Finally, the IT ladders and OT layers are again aligned together.

During the whole chain, the detector elements are only aligned along the measurement direction (local Tx), while fixing Ty, Tz and Rx (the movements to which the track χ^2 is the least sensible), except during the first step where the boxes are still aligned along Ty. As described in section V, the VELO is taken as a reference point during the whole procedure. This means that no other external constraints are needed.

B. Results

In this section, we present only the results of the last step of the whole alignment procedure, namely the simultaneous alignment of the IT ladders and OT layers. There are several ways of looking at the results of the alignment job. One can look at the flow of all the alignment parameters to see whether they have converged. However, since 348 parameters are being aligned for here such a plot is not so clear.

The alignment procedure minimises the total sum of the track χ^2. Another way of looking at the results of the alignment job is to look at this χ^2 as a function of the iteration. Fig. 6 shows the convergence of the normalised total sum of track χ^2 (ie. the normalised track χ^2 averaged over the track sample). During the iterative process, the number of tracks used for the alignment reduces, due to the cut on the track fit match χ^2), from 293'000 down to 284'000 tracks. It can be seen that in this step of the alignment procedure, 3 iterations are needed to converge.

Finally, we can also look at the output of the alignment job compared to the misalignment set in the input geometry. Fig. 7 shows the difference between the input (as set in the conditions) and the output (the result of the alignment job) misalignments for each of the IT ladders and OT layers. Each point on this plot corresponds to one detector element, starting from the first OT layer (station 1, layer X1), then moving on to the IT ladders (station 1, top box, layer X1, ladder 1) until the last ladder (station 3, C-side box, layer X2, ladder 7). We


Fig. 8. Alignment resolution in the Tx direction for the IT ladders and OT layers.

see on this plot that all the detector elements have converged within 100 microns of the input misalignment. Moreover, we see in this distribution that only a few outliers are found outside ± 40 microns. The structure of the distribution hints that there are still some weak modes left that could be removed by applying additional constraints.

In order to have a better feeling of the alignment precision, we can project all of these results on the y axis of fig. 7, leading to the distribution in fig 8. The FWHM of the distribution is 30 microns, which corresponds to a Gaussian σ of 13μm, only 20 % of the IT resolution. The large number of tracks (about 290'000) used for this analysis induces that the statistical error on this result cannot be reduced more.

It can also be seen that the mean of the distribution is shifted by about 10 microns. This shift has been found to be due to tracks coming from interactions (which are poorly fitted) and from ghost tracks, which are known to worsen the alignment results and cause biases. As will be seen in the next section this bias has no impact on the performance of the detector.

Fig. 9. Dimuon mass resolution using the unaligned, aligned and default geometry.


Fig. 10. Bias on the J/ψ mass versus p/GeV.

Fig. 11. J/ψ mass resolution versus p/GeV.

Fig. 12. Track χ^2/dof for tracks in true J/ψ candidates.

C. Validation with J/$\psi \to \mu^+\mu^-$ studies.

In section V-B three methods of checking the alignment procedure convergence were discussed. However, this is not enough to ensure the results of the job are good enough for physics studies. Indeed, if the procedure converges but the detector elements are still displaced by a large amount, the physics performance of the detector will be degraded. Therefore, the quality of the alignment has been validated using J/ψ decays. The results of these studies are described in the next section.

A sample of 65'500 simulated inclusive J/ψ events has been refitted using a standard loose J/ψ selection on three different geometry databases:

1) Unaligned (before alignment job).
2) Aligned (after alignment job).
3) Default (ideal geometry).

The aligned database is the output of the alignment job. The shape of the distribution of the dimuon mass resolution is very well recovered after alignment (see Fig. 9). We see that the mass resolution and bias as a function of the J/ψ momentum (see Figs. 10 and 11), with the aligned database match very well the default results. The difference in resolution is less than 1%, even at high momentum. On the other hand, the results with the misaligned database are clearly considerably worse.

The same conclusion can be drawn for the distribution of the track χ^2/dof (Fig. 12) which is fully recovered after the alignment, validating to a high degree of confidence the procedure.

VI. A FIRST GLIMPSE AT FIRST DATA AT LHCb

During the Summer 2008, two data samples have been acquired at LHCb. The first one consists of cosmic rays triggered either by the muon stations or by the calorimeters. Several runs have been performed between August and September 2008. These cosmic rays are being used to align various parts of the detectors, including the Outer Tracker. A sample of around 50,000 cosmic tracks passing through the Outer Tracker has been collected and are being used for alignment studies. Unfortunately, the Inner Tracker being a rather small subsystem with vertical measurement planes, relatively few cosmics cross more than one of its stations. Out of 2.2 million triggered events, only 80 tracks crossing two or more Inner Tracker boxes are found.

The second type of data acquired at LHCb is beam-dump data. During injection tests of the beam from the pre-accelerator (the SPS) into the LHC, particles were dumped in large blocks of materials. These collisions produced a large number of particles crossing the LHCb detector. Unfortunately, the events are very busy and it is hard to find good quality tracks in these events.

In this section, we present the first results of the alignment of the OT half-layers using real cosmic rays. In order to constrain the global translations and shearings, the two first and two last layers are fixed. The subsystem being aligned only in the measurement direction, this leaves 16 degrees of freedom to align for. Requiring that the track doesn't share any of its hits with a neighbouring track (in order to select only isolated and hence cleaner tracks) ands that no hits are found in the IT leads to a sample of 8,000 tracks for the alignment. Fig. 13 shows the χ^2 per degree of freedom of the change

Fig. 13. Convergence of the alignment of the OT half-layers with cosmic rays: χ^2/dof of the change in the alignment parameters.


Fig. 14. Result of the alignment of the OT half-layers with cosmic rays: corrections to apply to the survey measurements.

in the alignment parameters:

$$\Delta\chi^2 = \frac{1}{2}\frac{d\chi^{2T}}{d\alpha}\Delta\alpha = -\Delta\alpha^T \mathrm{Cov}(\alpha)^{-1}\Delta\alpha \qquad (2)$$

This corresponds to the significance of the alignment correction, as explained in [5]. Anything below 1 is only statistical noise. The alignment procedure has hence converged in 2 iterations. It is important to notice that not all hits were found by the pattern recognition in the first iteration due to large initial misalignments.

Fig. 14 shows the correction that must be applied to the survey measurements, according the alignment procedure. The results here shows that the survey measurements were only good to 2 mm. However, this result is a proof that the alignment procedure, using tracks from the standard Kalman-filter track-fit, converges and is usable to align the LHCb detector

VII. CONCLUSION

In this paper, we have presented a new alignment method that is being used to align the LHCb detector using Kalman-fitted tracks coming from the standard track fit. The procedure makes use of the formulae to calculate the global track covariance matrix after the Kalman filter track fit given in [5].

We then presented a realistic scenario on which this alignment procedure as been tested. In this scenario we aligned the OT layers as well as the IT boxes, layers and ladders starting from a realistic day-1 misalignment. The tracking stations were here aligned using simulated minimum bias collisions at 7 TeV protons at the nominal magnetic field. The procedure has been adapted to the high complexity and number of degrees of freedom by aligning step by step, starting from a coarse granularity (IT boxes and OT layers separately) and moving step by step to the finest granularity (IT ladders and OT layers together). We showed, in this scenario, that we were able to align the detector to a precision good enough to not affect the J/ψ mass resolution by more than 1 %.

These results are very promising a few months prior to the first beam-gas or collisions data will be acquired at LHCb. Work is ongoing to analyse the first beam-dump and cosmics data and new ways for identifying and rejecting weak modes are under study.

REFERENCES

[1] M. Needham, *Status and expected performance of the LHCb tracking system*, Nuclear Physics B156 (2006), 217-220.
[2] V. Blobel, *Software alignment for tracking detectors*, NIM A 566 (2006) 5.
[3] P. Brückman, A. Hicheur and S. J. Haywood, *Global χ^2 approach to the alignment of the ATLAS silicon tracking detectors*, ATL-INDET-PUB-2005-002 (2005).
[4] A. Bocci and W. Hulsbergen, *TRT alignment for SR1 cosmics and beyond*, ATL-COM-INDET-2007-011 (2007)
[5] W. Hulsbergen, *The global covariance matrix of tracks fitted with a Kalman filter and an application in detector alignment*, Nucl. Instr. and Meth. A (2008), doi:10.1016/jnima.2008.11.094.
[6] G. Barrand et al., *GAUDI - A software architecture and framework for building LHCb data processing applications*, Proceedings of CHEP 2000.
[7] M. Needham, *Classification of Ghost Tracks*, CERN-LHCb-2007-128 (2007).
[8] M. Needham, *Performance of the LHCb Track Reconstruction*, CERN-LHCb-2007-144 (2007).