CDF/DOC/TRACKING/GROUP/5700
8$^{\text{th}}$ August 2001

A Procedure for the Software Alignment of the CDF Silicon System

R. McNulty, T. Shears (University of Liverpool),
A. Skiba (University of Karlsruhe).

Abstract

A method to align the silicon system of the CDF Run2 detector, using the Millepede fitting programme, is described. A number of studies using simulated data have been performed. The results are presented and the critical issues are discussed.

1 Introduction

This note describes a method to align the silicon system of the CDF Run2 detector. An overview of the method and its advantages over standard techniques are given in section 2. The technical details are presented in section 3. Results obtained from simulation studies are given in sections 4 and 5 where the critical issues are discussed. Section 6 summarises our results and conclusions.

2 Overview

The alignment procedure we follow is based upon the Millepede package written by Volker Blobel [1]. Most alignment procedures [2] attempt to find the alignment parameters by performing a series of track fits which initially ignore the misalignments. From the distribution of residuals (the difference between the fitted and measured data) the alignment parameters are then deduced. In contrast, Millepede makes one fit to all the data, fitting the track parameters and misalignments simultaneously.

The standard approach would be perfectly correct if the alignment parameters were deduced from residuals which were defined as the difference between measured and truth data instead of between measured and fitted data. The fitted data is a function of the very misalignments we are trying to calculate, but it is rare that any attempt is made to parameterise it thus. The usual procedure is to hope that the misalignments cancel in the average, and the fitted track then gives a better approximation to the truth; this results in a biased fit because the fitted track has been calculated with the wrong model (viz. it assumes there are no misalignments). Sometimes this procedure is iterated in the belief that the fitted data comes closer to the truth; this may in fact compound matters making the aligned parameters even more biased.

It is formally much better if all free parameters are fitted simultaneously with one gigantic error matrix which precisely describes the relationship between the track and alignment parameters. Furthermore, provided the problem is linear, no iteration is required.

The principle difficulty with performing one all-embracing fit occurs in the number of free parameters which exist and the necessity of inverting an unfeasibly large matrix. Blobel's method gets around this by splitting the large inversion up into a series of smaller inversions.

Different classes of tracks are sometimes identified in order to constrain different parts of the detector. For example, stiff tracks might be used to align silicon wedges internally. Then tracks in the overlap regions might join the wedges together, and finally cosmics might be used to remove overall deformations. In the standard approach, these samples might be considered sequentially by the alignment procedure. With Blobel's approach, all these tracks are simultaneously presented to the alignment programme and his method naturally combines them and uses them to constrain the detector in the optimal way.

Another advantage of Blobel's method is that, for the CDF geometry, *no* external information is required. There are sufficient internal constraints with overlapping wedges and a large number of layers in order for the fit to converge holding only one ladder of silicon fixed. Thus we end up with one internally consistent aligned detector. Systematics in the COT or the knowledge of the beam line will not be mapped onto the silicon. On the contrary, the independently aligned silicon can be used to extrapolate tracks outwards to aid in the determination of pedestals, t-zero offsets, alignment, and drift velocity in the COT.

3 Formalism

We present a general description of the minimisation procedure employed, before describing the specific implementation of this scheme for the CDF silicon system. For a complete description of the formalism, see the Millepede documentation [1].

3.1 General Description

Given a set of N measurements (x_i, y_i) with x the independent and y the dependent variable, a straight line fit for slope m and intercept c can be found in the usual way by minimising the quantity

$$\chi^2 = \sum_{i=1}^{N} \left(\frac{y_i - mx_i - c}{\sigma_i} \right)^2$$

where σ_i is the measurement uncertainty. Letting the weight $w_i = 1/\sigma_i^2$, and defining the vector of parameters $\vec{a} = (c, m)$ and N vectors $\vec{x}_i = (1, x_i)$,

$$\chi^2 = \sum_{i=1}^{N} w_i(y_i - \vec{a}^T \cdot \vec{x}_i)$$

and the solution is given by $\vec{a} = C^{-1}\vec{b}$ with C a matrix and $\vec{b}$ a vector defined by

$$C = \sum_{i=1}^{N} w_i \vec{x}_i \vec{x}_i^T, \vec{b} = \sum_{i=1}^{N} w_i y_i \vec{x}_i.$$

More generally, the linear fit can be extended from two parameters to n when $\vec{a} = (a_0, a_1, a_2, ...a_n)$ and $\vec{x}_i = (1, x_i, x_i^2, x_i^3, ...x_i^{n-1})$

Millepede distinguishes between local parameters which pertain to a subset of measurements (eg. the track parameters) and global parameters which contribute to all measurements (eg. alignment parameters). Provided that the measurements can be described as *linear* functions of the local and global parameters then

$$\chi^2 = \sum_{i=1}^{N} w_i(y_i - \vec{a}^T \cdot \vec{x}_i - \vec{\alpha}^T \cdot \vec{\xi}_i)$$

where $\vec{a}$ and $\vec{\alpha}$ are local and global parameters respectively. This can be solved as above; however with (say) 5 free local parameters per track, 1000 global parameters and 10,000 tracks, this gives 51,000 free parameters. Inverting the 51,000x51,000 matrix appears impossible. Fortunately, Blobel shows that this matrix inversion can be reduced to the inversion of a matrix whose side equals the number of global parameters; in the above example this is a 1000x1000 matrix. This inversion is further simplified by the fact that this matrix is almost diagonal.

Blobel's formalism also allows for constraints on the global parameters to be included. They appear as additional terms in the Chisquared.

3.2 Particular Description

Because there is more $R\phi$ than z information and it has a better precision, better understood resolution and less tracking ambiguity, our intention is first to perform a precise alignment in the transverse plane, and then to add in the z-alignment. This note thus concerns itself with the former. The latter can be accomplished in a similar fashion.

In the transverse plane, each track appears as a circle; thus there are three free local parameters. In addition, in order to correctly account for multiple scattering, we add free parameters for the total amount of scattering occurring before each measurement.

The silicon detectors have 1752 wafers of silicon arranged on 352 separate structural ladders. The individual wafers within a ladder have been assembled, aligned, and their placement relative to fiducials are known through microscope measurements with an accuracy of a few microns (typically 1 or 2 μm) [3]. The placement of the wafers onto the support structures is less precisely known (typically 20μm). Thus, initially we will only concern ourselves with aligning ladders of silicon. (This can be refined at a later stage with the same formalism if necessary.)

The deviations of the wafers from the nominal positions in the transverse plane can be described in terms of two translations and three rotations. The translations considered are those parallel to the silicon wafer (approximating to a global $R\phi$ motion) and those perpendicular to it (approximating to a global radial motion). The rotations are firstly, in the plane of the silicon (a $R\phi$ v z twist or 'yaw' in airline terms), secondly out of the plane of the silicon about its long axis (a R v $R\phi$ twist or 'roll') and thirdly out of the plane of the silicon about its short axis (a R v z twist or 'pitch'). There is much less sensitivity to these last two rotations, so initially we only concern ourselves with the first rotation. There are thus 352×3 global parameters which we must fit for.

The next step, and the most important in evaluating the validity of this method, is to linearise the equations. Blobel's method depends on our equations having a linear dependence on both the local and global parameters. Since our tracks describe circles in the transverse plane, and since we must account for multiple scattering (which is an implicit function of the curvature) this is not absolutely true. We attempt to circumvent this problem by instead considering deviations from a reference trajectory. The deviations as a function of arc length are much more linear, provided we are 'close' to the correct solution. This in turn requires that the reference trajectory be close to the truth and the misalignments are small enough that their deviations from this reference can be described in a linear fashion. To this end, the reference trajectory is taken as the best circle fit through the misaligned silicon hits.

A transformation of co-ordinates is made from the Cartesian (x, y) frame to the (s, t) frame where s is the distance along the reference trajectory and t is the perpendicular distance of a point from the trajectory. The point (0,0) in (s, t) space occurs at that point on the reference trajectory closest to the origin in the (x, y) frame.

The quantity to be minimised as a function of the global parameters (Δx_L parallel to the silicon plane, ΔR_L perpendicular to the plane, and the 'yaw' $\Delta\psi_L$ for all ladders L) and the many local parameters, ($a_0^{(k)}, a_1^{(k)}, a_2^{(k)}$ which describe the circle fit for track k, and $\theta_j^{(k)}$ which account for the multiple scattering between each silicon layer for track k) is

$$\chi^2 = \sum_k^{\#tracks} \chi_k^2 + \sum_L (\frac{\Delta x_L}{\sigma x_L})^2 + \sum_L (\frac{\Delta R_L}{\sigma R_L})^2 + \sum_L (\frac{\Delta\psi_L}{\sigma\psi_L})^2$$

where the first term on the right-hand-side is a sum over the individual chisquared contributions for each track, and the other three terms are constraints on the size of the alignment parameters, with $\sigma x_L, \sigma R_L, \sigma\psi_L$ the estimation of the uncertainty on the shifts and twists coming from the mechanical survey.

The individual track chisquareds are constructed in a standard form which includes the effects of multiple scattering [4]

$$\chi_k^2 = \sum_{i=1}^{\#Si\ hits} 1/\sigma_i^{(k)2} \left(t_i - (a_0^{(k)} + a_1^{(k)} s_i + a_2^{(k)} s_i^2) - \mathrm{ladder}_L - \sum_j^{\#Scat\ centres} \theta_j^{(k)} (s_i - Rscat_j) \right)^2$$

$$+ \sum_j^{\#Scat\ centres} (\theta_j^{(k)} / \sigma_{\theta j}^{(k)})^2$$

where $\mathrm{ladder}_L = \Delta x_L \{(\frac{\partial t}{\partial x}, \frac{\partial t}{\partial y}) \cdot \vec{u_x}\} + \Delta R_L \{(\frac{\partial t}{\partial x}, \frac{\partial t}{\partial y}) \cdot \vec{u_R}\} + z\psi_L \{(\frac{\partial t}{\partial x}, \frac{\partial t}{\partial y}) \cdot \vec{u_x}\}$.

$\vec{u_x}$ is the unit vector parallel to the wafer and $\vec{u_R}$ is the unit vector perpendicular to it. The intrinsic error on silicon hit i on track k is $\sigma_i^{(k)}$. Each scattering centre is found at a distance $Rscat_j$; and the average scattering angle from this material as seen by track k is

$\sigma_{\theta j}^{(k)}$. Note that $\sigma_{\theta j}^{(k)}$ scales with $1/p_T$ which is an implicit function of the curvature parameter $a_2^{(k)}$. This creates a non-linearity in the fit. As will be seen later, provided the misalignments are small and the material budget well known, the p_T can be taken from the reference trajectory. However, it is safer to use the COT momentum determination to estimate this quantity.

Minimising the chisquared as given above would lead to an attempt to invert a singular matrix, since an infinity of solutions are possible where the whole detector is displaced or rotated. However, if just one silicon ladder is held fixed, this ambiguity is removed. In what follows, we hold the ladder at "12 o' clock" on the outermost layer of the SVX fixed (ie. Wedge 3, Layer 5).

4 Results

We now present our results. Due to CPU limitations and in order to have sufficient statistics in each of the silicon wafers, we limited our tests to the combined L00 and SVX detectors for which there are 204 ladders of silicon.

4.1 Proof of Principle

In the first test, 100,000 tracks were generated with a flat momentum spectrum between 2 and 10GeV, and a flat polar angle distribution. They were extrapolated to the silicon, taking account of multiple scattering in each preceding layer, to obtain the silicon hit position which was smeared with a resolution of 6μm. The detector was then misaligned with random shifts sampled from Gaussian distributions of width 20μm parallel to the wafers, 60μm perpendicular to the wafers, and 0.2 mrad twist in the plane. These values are roughly of a size that might be expected from the mechanical survey [3]. Finally, the alignment procedure is applied, using the RMS of the shifts as a constraint. Note, that in the calculation for multiple scattering, the true momentum was taken. The COT value will be sufficient when using real data. Figure 1 plots the generated misalignments against the calculated corrections. A clear correlation is observed showing that the procedure improves the detector alignment.

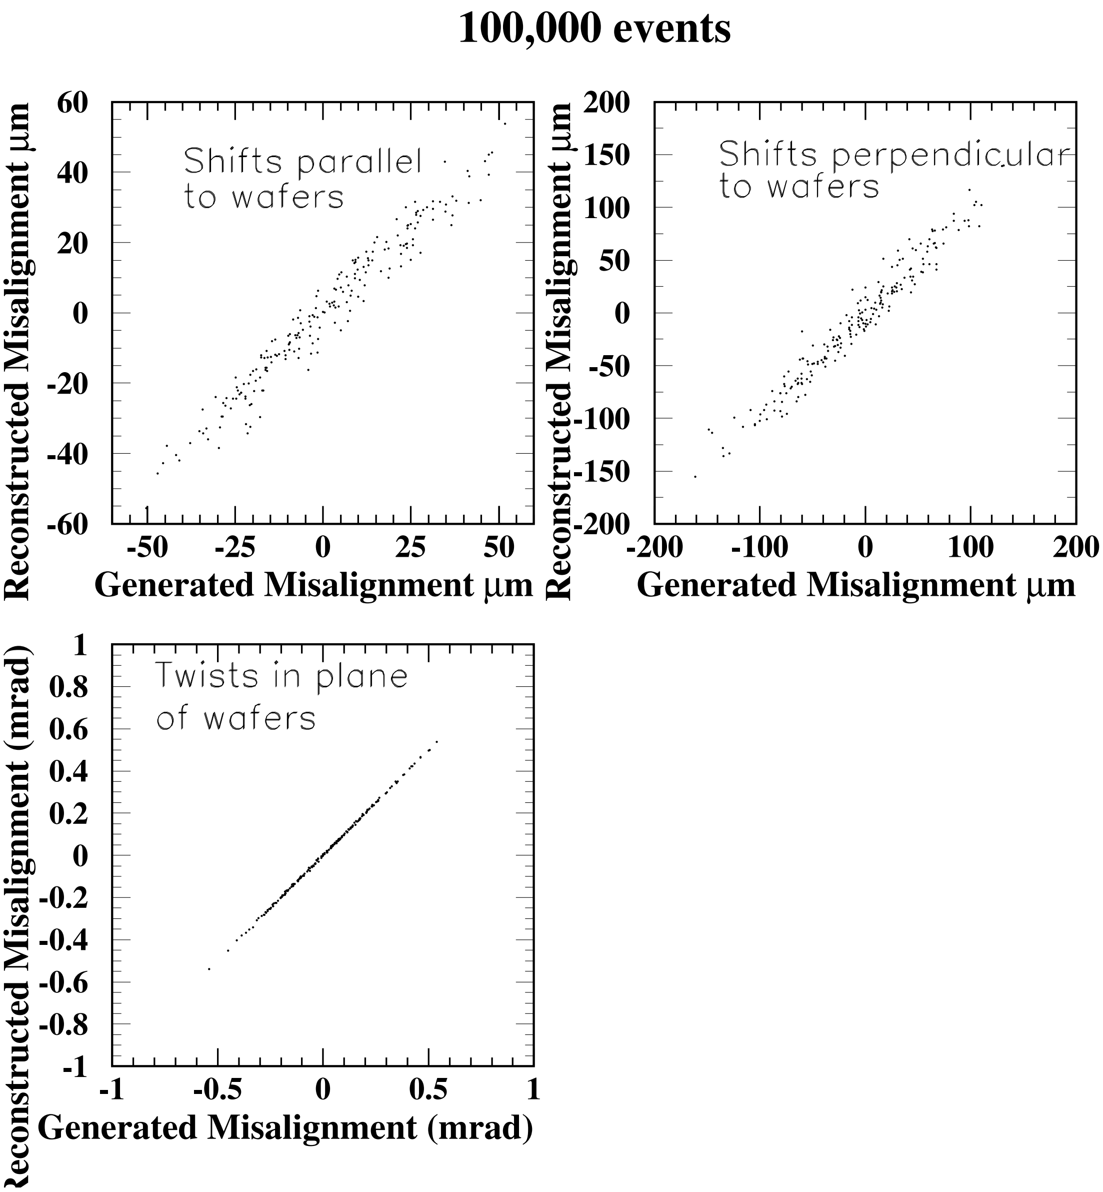

Figure 1: Generated versus Reconstructed Misalignments for the two shifts and twist applied. The RMS. of the misalignments are 20μm in the $R\phi$ direction, 60μm in the R direction, and a 0.2 mrad twist in the plane of the silicon.

4.2 The Error Estimate

The error on each alignment parameter varies depending on the geometry, the number of tracks incident on the ladder, and the uncertainty in the mechanical alignment.

First we considered the case with no constraints on the alignment parameters. The precision in calculating the alignment parameters increases roughly (though not exactly due to the one fixed ladder,) as $1/\sqrt{N}$, as can be seen in Figure 2. With 10,000 tracks, a typical precision of 100μm on the estimation of the misalignments parallel to the wafers is obtained.

Adding as constraints our knowledge of the uncertainty in the mechanical survey (which was used to simulate the misaligned tracks), the precision improves by an order of magnitude, though the dependence on the number of tracks used is much slower, as shown in Figure 3. Thus we can perform a relatively precise alignment rapidly with a rather small number of tracks (of order 1000). Using 10,000 tracks, the absolute precision in calculating the misalignments from the nominal position for ladders in the central barrel ranges from 4μm and 8μm parallel and perpendicular to the wafers for ladders in layer 1, to 7μm and 19μm in layer 5.

4.3 The Pull Distribution

The Pull Distribution, defined as (Reconstructed Alignment - Generated Misalignment)/Error, is plotted in Figure 4. We might initially expect this to be a Gaussian of unit width, but deviations from such a hypothesis are evident. This is due to the fact that the pulls from ladder to ladder are highly correlated. An examination of the full covariance matrix from the fit shows that adjacent layers in the same wedge have a correlation coefficient of 0.97 while adjacent wedges in the same layer have a correlation coefficient of 0.5. Even ladders with the same wedge and layer but in adjacent barrels have correlation coefficients as high as 0.9. This can be seen graphically in Figure 5 where the pull is plotted as a function of wedge number for layers 3,4 and 5 (all of which have the same 12 fold symmetry). Note that a high pull value for wedge 0 is visible in all layers, and that the shape plotted is broadly similar for each of the layers. This indicates a large correlation between the layers. Note also that a high value in wedge

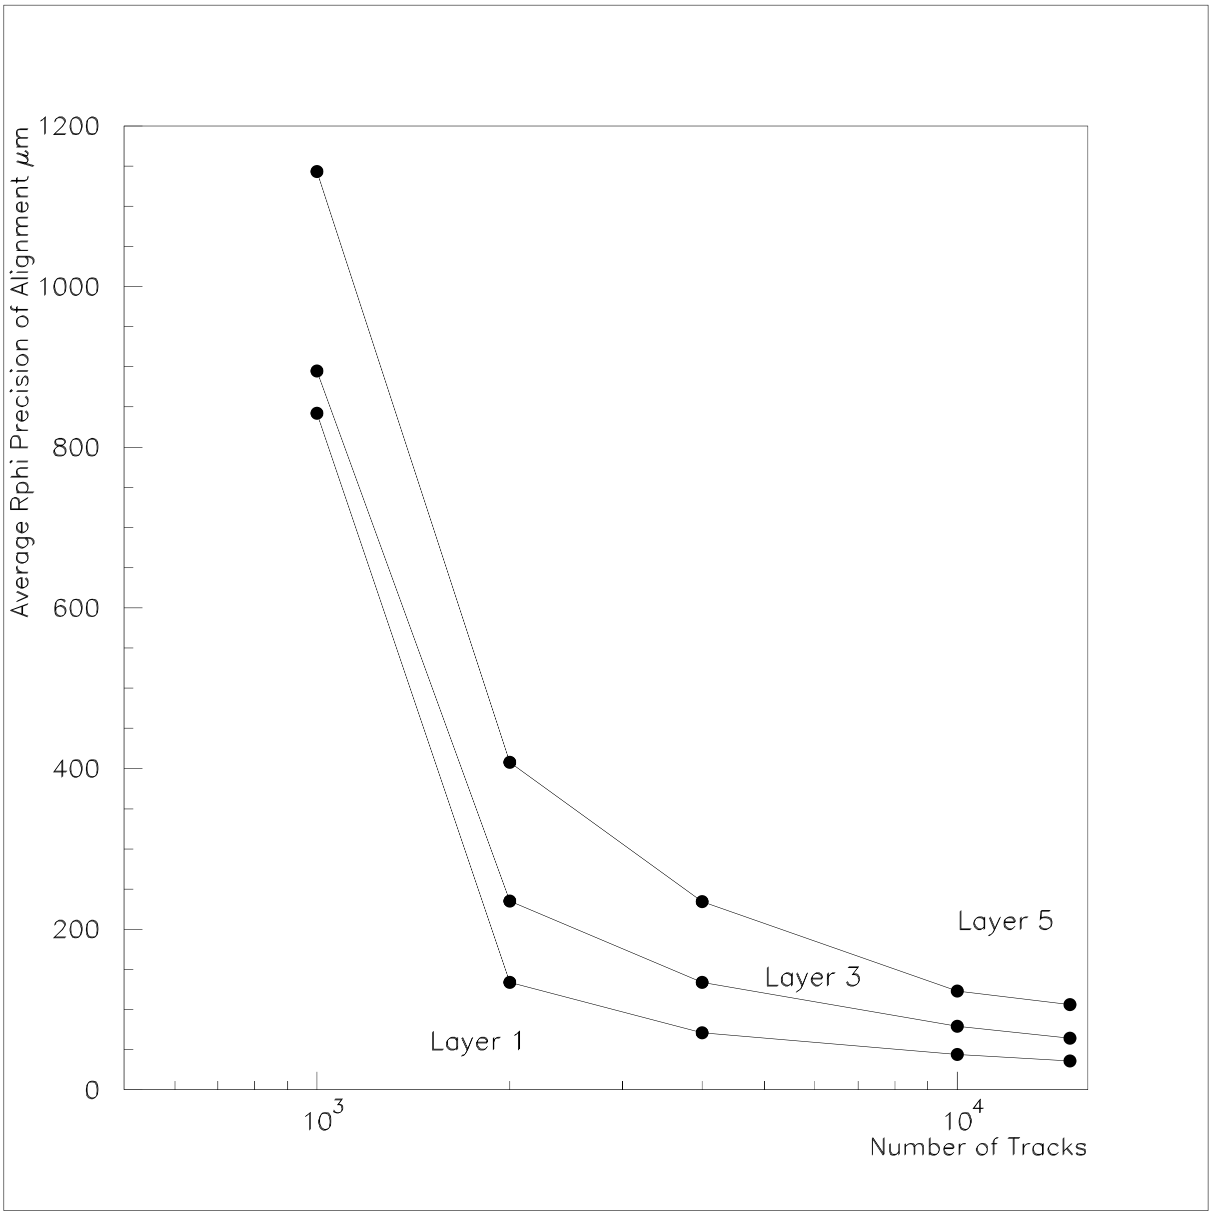

Figure 2: Uncertainty in the estimation of the misalignment shifts parallel to the wafers as a function of the number of tracks used, for layers 1,3 and 5, where no constraint is applied in the determination of the alignment parameters.

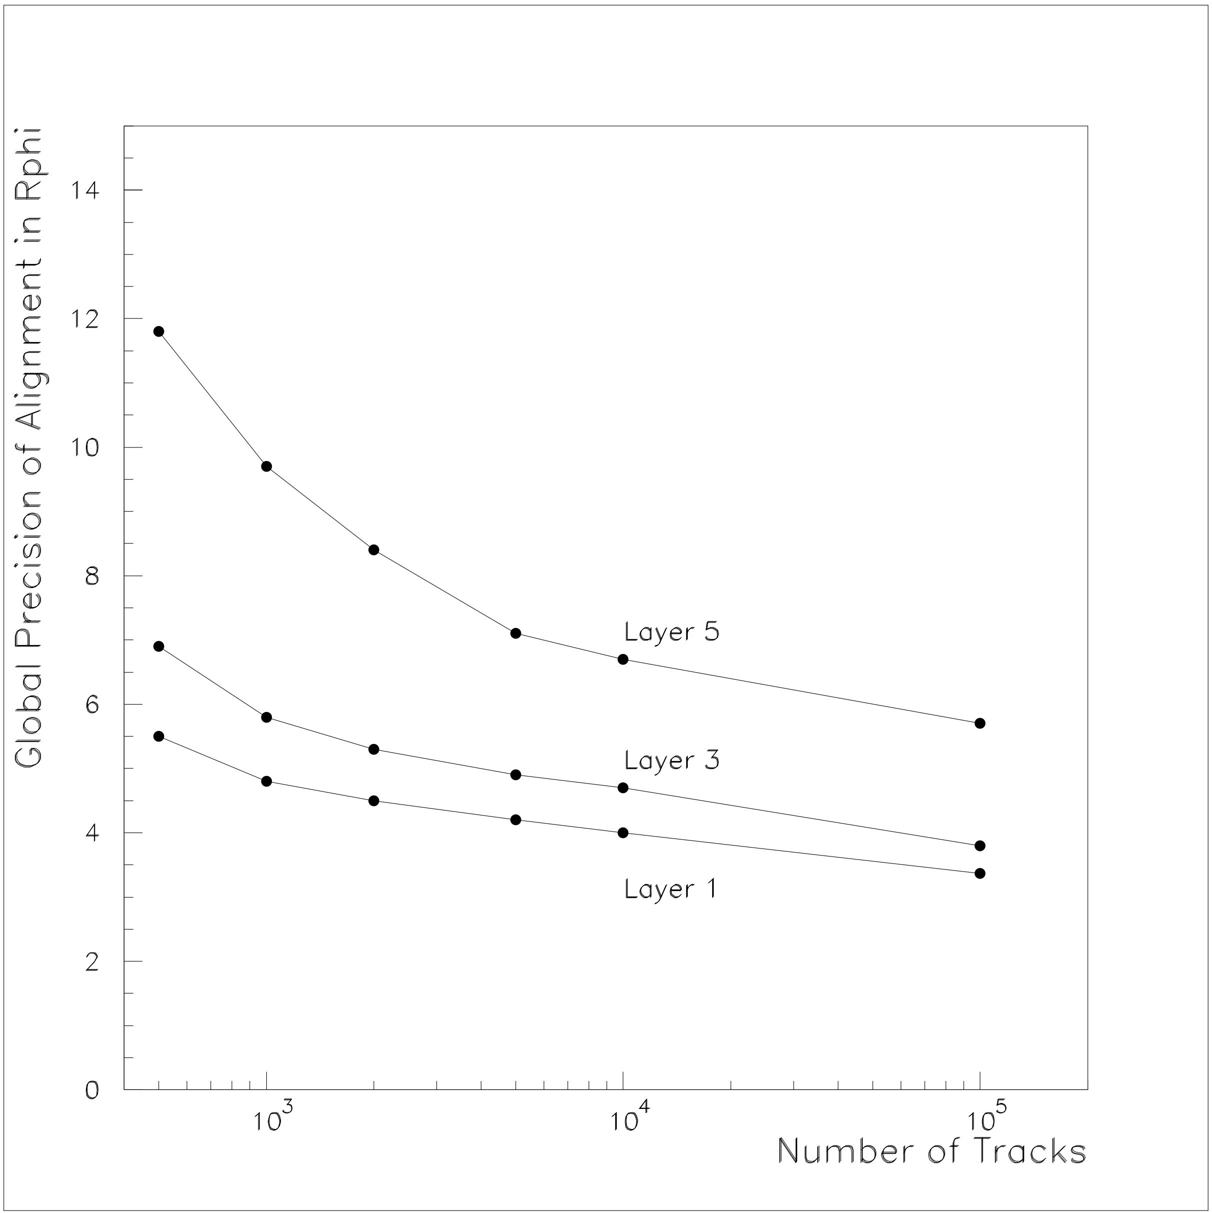

Figure 3: Uncertainty in the estimation of the misalignment shifts parallel to the wafers as a function of the number of tracks used, for layers 1,3 and 5, when constraints somewhat similar to what might be expected from the mechanical survey, are used. Units on the y-axis are in microns.

1 for any layer, causes pulls in the same direction in wedges 0 and 2, which is an indication of the correlation between wedges.

4.4 The Chisquared Probability Distribution

The best way to tell if this procedure produced a reasonable result is to examine the value of the chisquared probability. The chisquared is defined as the product of three matrices $AC^{-1}A^T$ where A is a 1x612 matrix having as elements the generated minus reconstructed alignment parameters, and C is the 612x612 covariance matrix. There are 612-3 degrees of freedom. For the above fit, the chisquared is 643.2 giving a chisquared probability of 0.18 which tells us that a reasonable fit has been performed with a trustworthy error estimate.

To really prove this point however, we must repeat the above procedure using many independent data sample and examine the chisquared probability distribution to ensure it is flat. We have repeated this for 100 sample each of 1000 tracks. The superimposed pull distributions are shown in Fig 6. Fitting each with a Gaussian gives values close to one. The chisquared probability distribution is plotted in the upper plot of Fig 7. It has a very slight bias towards 1 indicating that our calculated errors are slightly underestimated. Considering the approximations as outlined above this is an excellent result. The errors must be decreased by just 2% before a flat distribution is obtained as shown in the lower plot in Fig 7.

4.5 Impact Parameter

The alignment procedure has returned a global alignment which, with 10,000 tracks is accurate to about 6μm in $R\phi$ and 12μm in R. This is a global precision and applies through the volume of the silicon system. Locally of course, the *relative* precision is much better. The relative position of adjacent wedges will be known to sub-micron accuracy.

Of greatest interest for physics though, is the precision with which secondary vertices can be reconstructed. This can most easily be evaluated by considering the resolution of the impact parameter. The top left plot in Figure 8 shows the impact parameter distribution for tracks coming from B decays with $p_T > 2$GeV. (The lifetime information has been removed; we plot the distance between the extrapolated track and its true production point.) It has a fitted width of 37μm. The

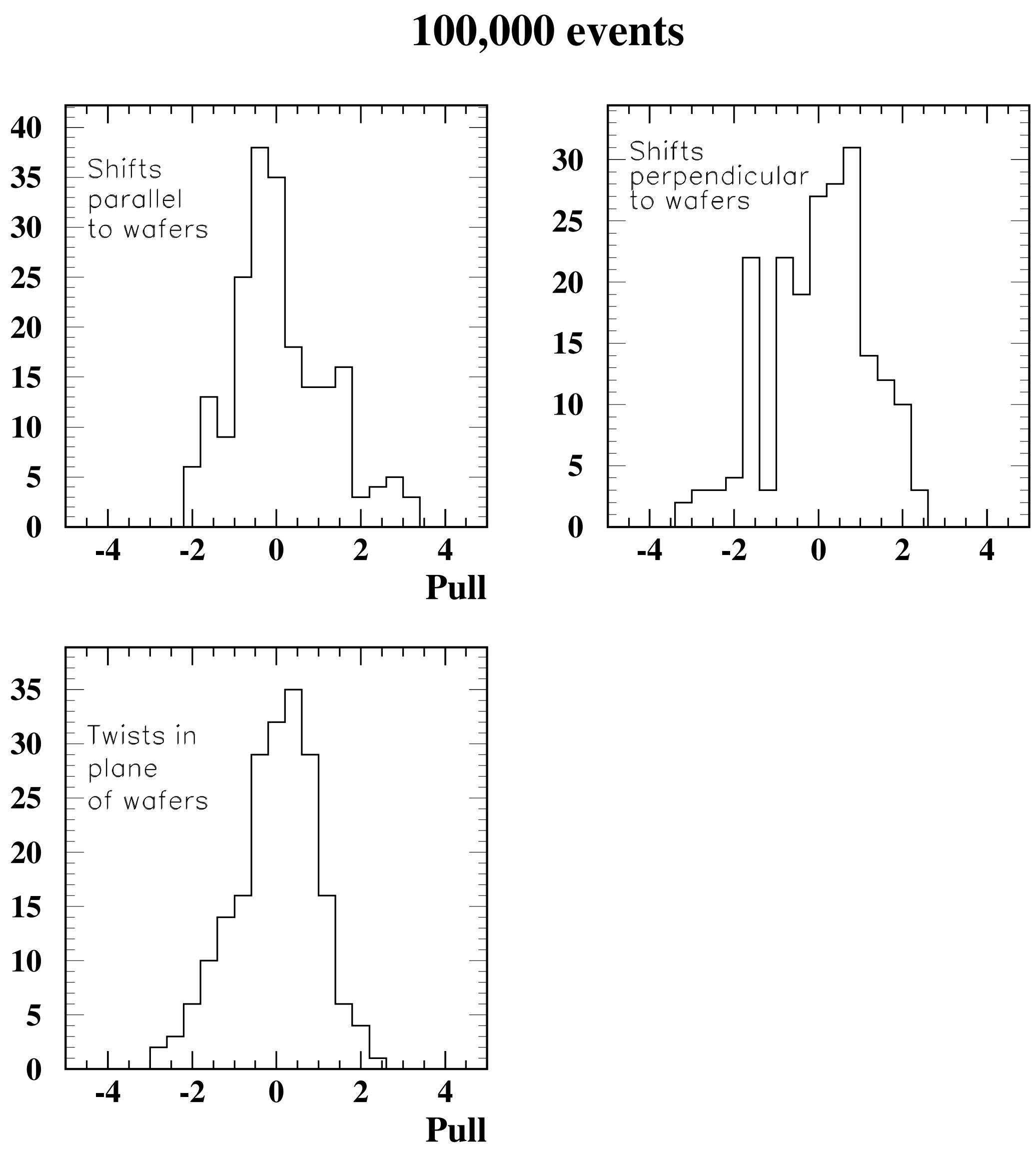

Figure 4: Pull Distributions, (Generated-Reconstructed Misalignment)/Error, for each of the misalignments introduced, where the alignment has been recovered using a sample of 100,000 tracks.

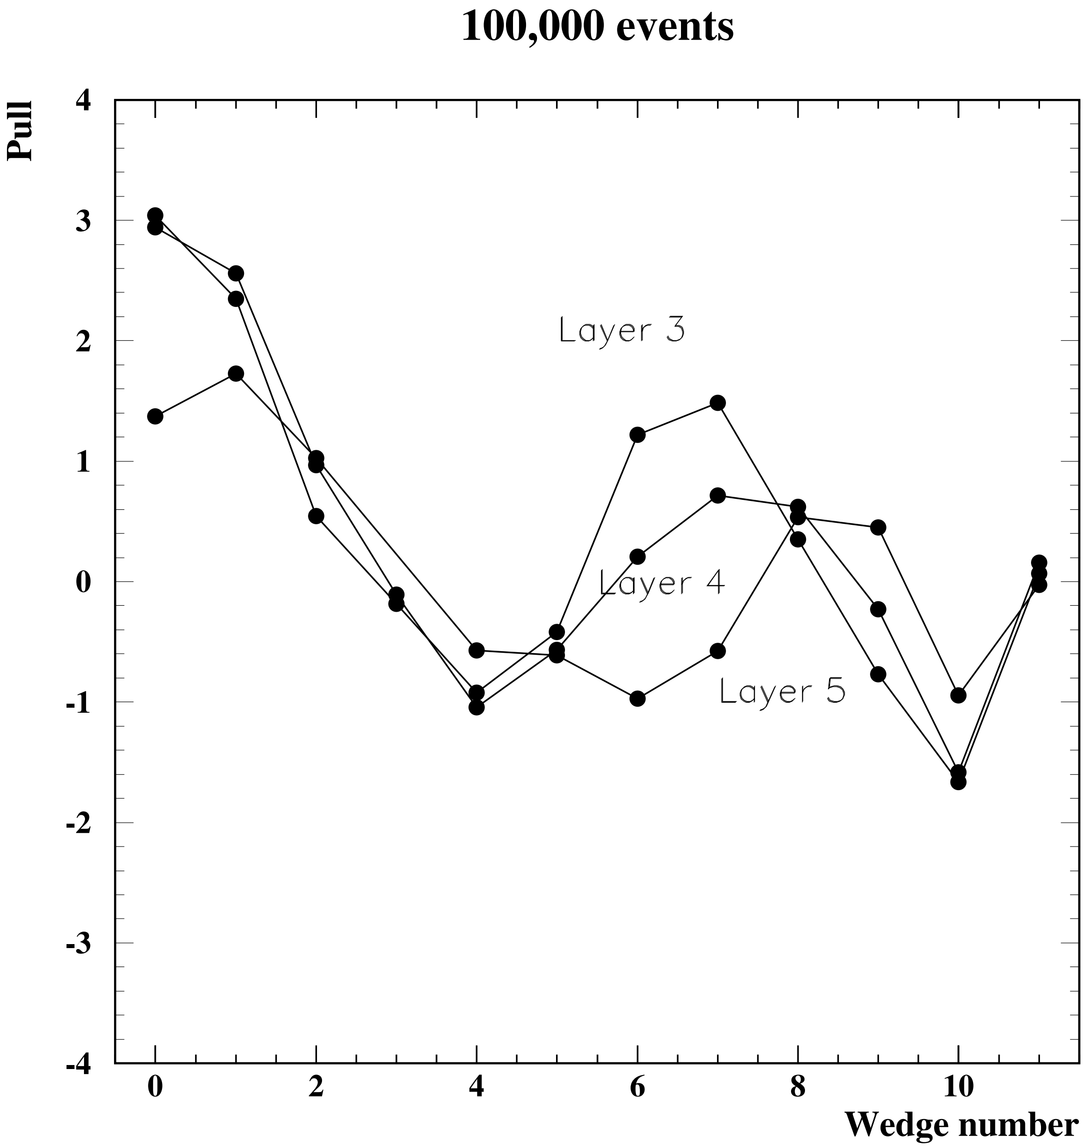

Figure 5: Pull versus Wedge Number plotted separately for three different layers. Wedge 3 in Layer 5 was fixed to its correct position.

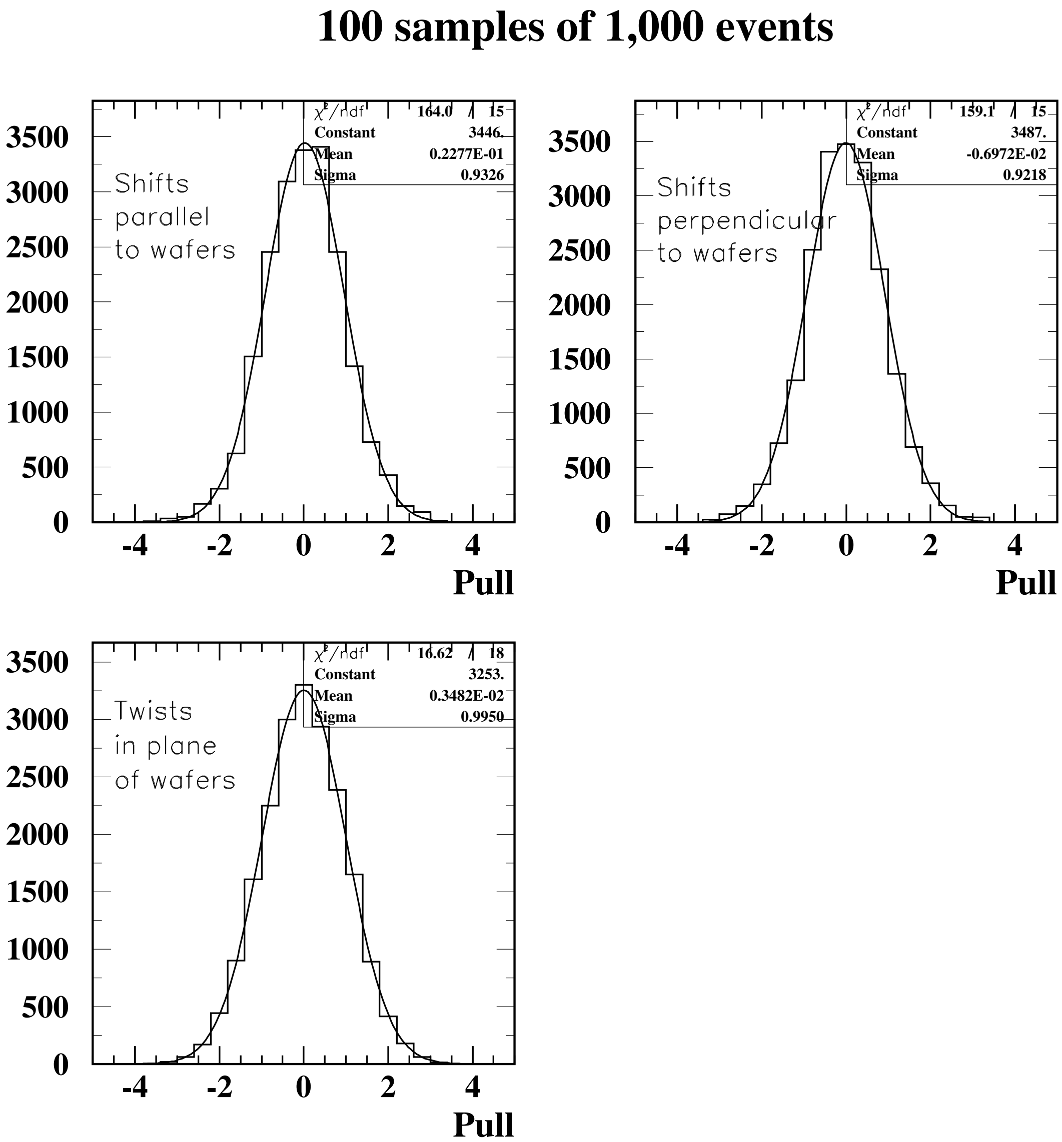

Figure 6: Pull Distributions for each of the misalignments introduced. The results from 100 different alignments, each produced using 1000 tracks, are superimposed.

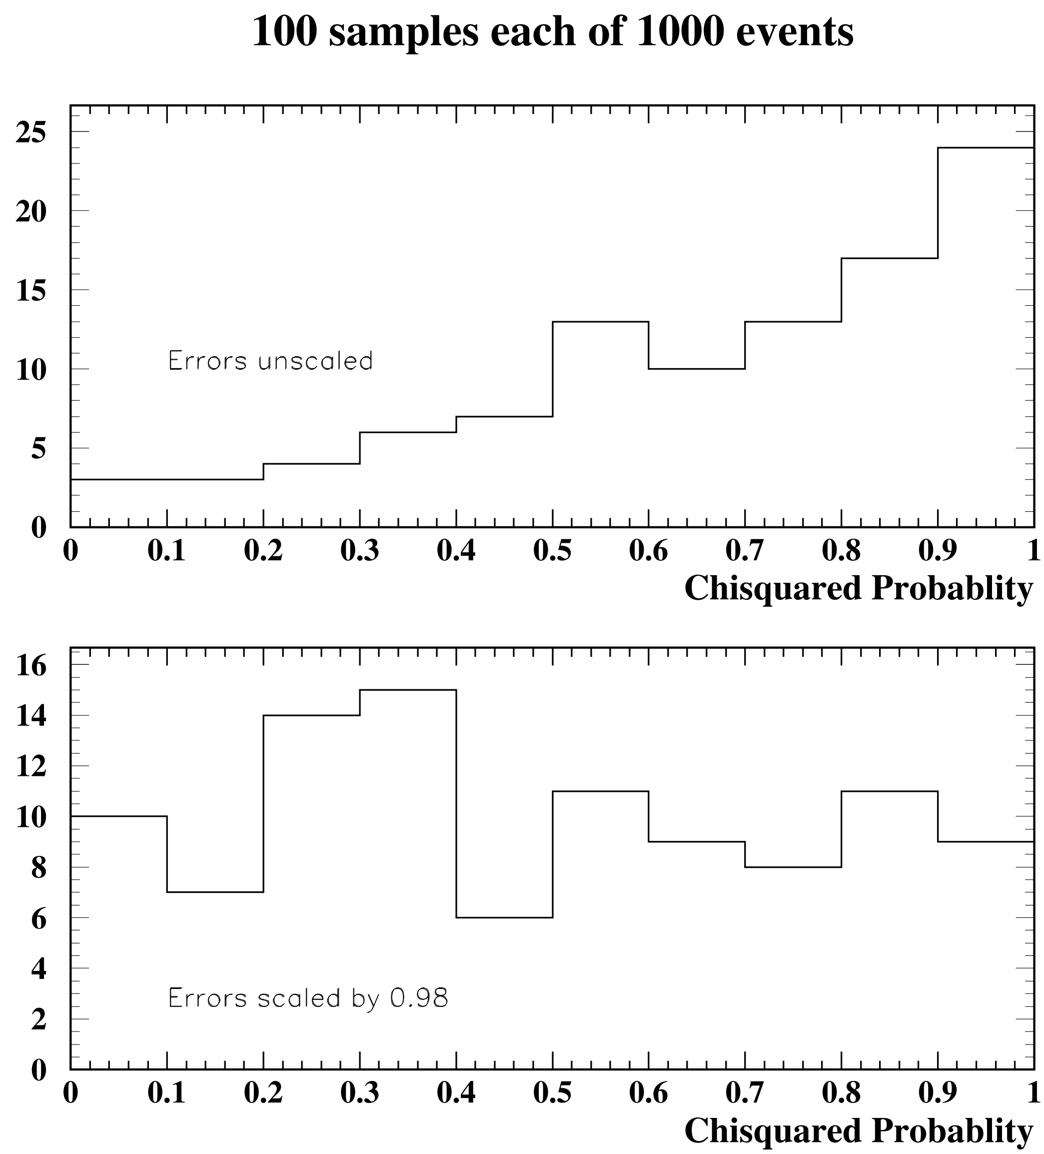

Figure 7: Chisquared probability distributions obtained by comparing the reconstructed detector to its true position. The upper plot is with the errors unscaled; the lower plot is after the errors have been decreased by 2%.

top right plot shows the same distribution after the alignment procedure, using 10,000 tracks. It is centred on zero and has a fitted width of 14.5μm. Finally the bottom left plot shows this distribution for a perfectly aligned detector. It is centred on zero and has a fitted width of 13μm. We conclude that with 10,000 tracks the alignment procedure gives an unbiased impact parameter distribution with a resolution which is just 1.5μm greater than the theoretical minimum.

One way to obtain an even better result would be to use more tracks. Another technique would be to use additional constraints. Including a sample of tracks which are known to come from a common vertex (Z boson decays, J/ψ etc.) would help constrain the focusing precision of the silicon system.

5 Stability of the Alignment Procedure

We now examine the dependence of our alignment procedure on our estimates of the precision of the mechanical survey, our knowledge of the material budget, the size of the misalignments and the misassociation of silicon hits to tracks by the pattern recognition procedure. We will see that so long as we remain in the 'linear' region of the fit, everything works fine.

5.1 Dependence of the Alignment on the Multiple Scattering Estimate

It is important to examine the effect of the multiple scattering term for two reasons. Firstly, the current amount of material in the simulation is incorrect and thus taking these values while fitting the data leads to a bias. Secondly, even if the detector was very well modelled, one can never be sure whether the particle hit the edge of a hybrid or the wall of a cooling pipe. Thus, we know there will definitely be an incorrect estimate of multiple scattering in our fit, and we wish to know how this propagates through to the alignment.

As our first check, the multiple scattering in the generated tracks was increased by a factor of two. The alignment procedure was fol-

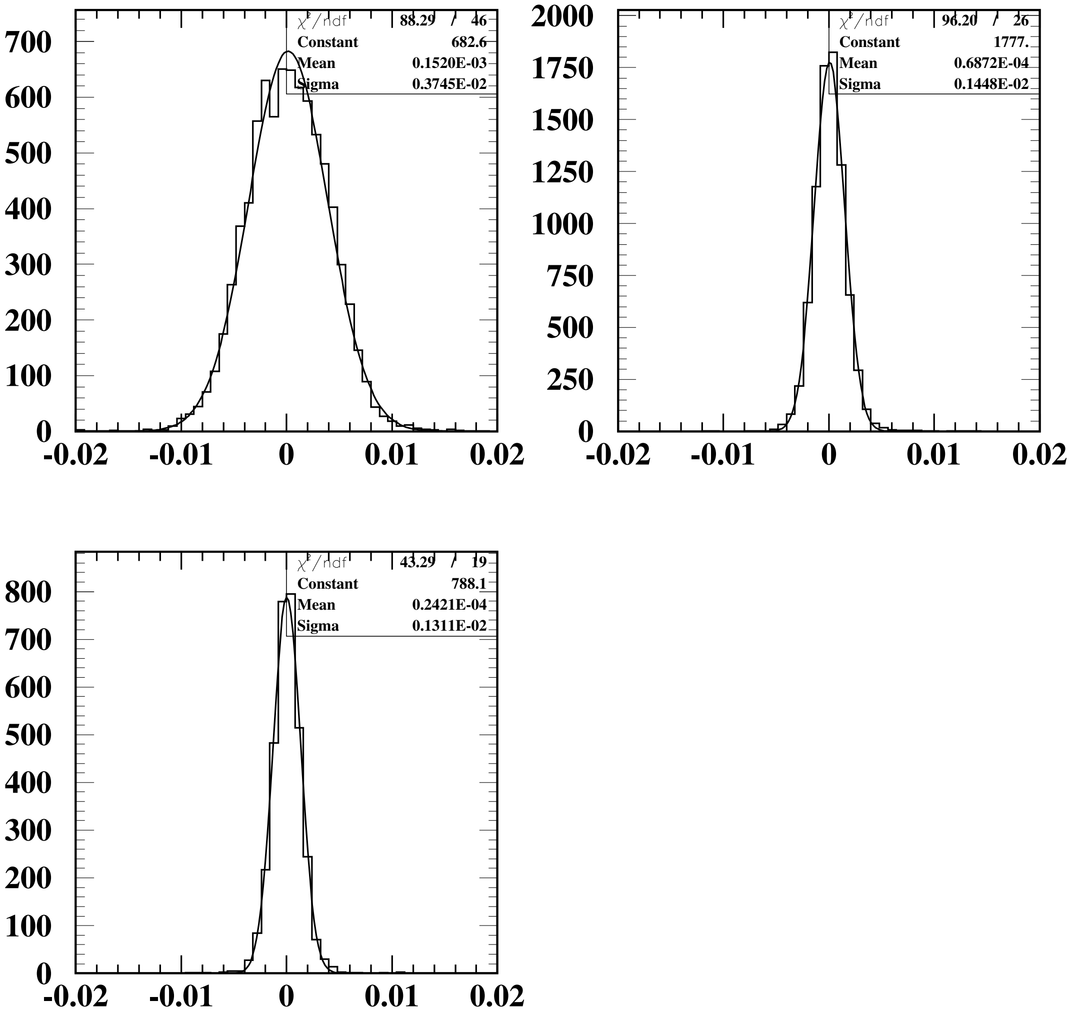

Figure 8: The Impact Parameter Distribution for the misaligned detector (top left), the misaligned detector after our alignment procedure (top right), and a perfectly aligned detector (bottom left). Units on the x-axis are in cm.

lowed but was ignorant of the increase in the material budget. The pull distributions for 20 samples of 1000 tracks is shown in Figure 9. The chisquared probability distribution requires that the errors be increased on average by 32% in order to be flat. Thus we conclude that our alignment gives reasonable central values, but the uncertainties on the alignment are underestimated. This is to be expected since our point resolution has been underestimated by underestimating the amount of scattering.

An interesting result is obtained if, instead of using the COT momentum information to determine the amount of multiple scattering, we use the reference track, where the momentum was found from a fit to the misaligned data points. Now we find the errors to be underestimated by 40%. We conclude that the precision of the result is improved if we use some external information in order to get the scattering terms correct.

5.2 Dependence of the Alignment on the Mechanical Survey Precision.

It is possible that we could have under- or over-estimated the precision of the mechanical survey. To test the effect of this, we proceed as before to generate data with random misalignments of 20μm,60μm,0.2mrad in $R\phi, R$ and Rvz. However in the alignment procedure, we inform the programme that the misalignments have an RMS twice as small as they actually do. The pull distributions which result from 20 samples of 1000 events are plotted in Figure 10. The chisquared probability distribution is only flat if the errors are increased by 30%. More worryingly, the pull distributions are skewed.

We repeat the alignment procedure, this time overestimating the size of the misalignments by a factor of two. The resulting pull distributions are plotted in Figure 11. In this case, the pull distributions are symmetric, better centred, and the chisquared probability distribution is flat when the errors are decreased by 9%.

Thus we conclude that it is safer to err on the side of underconstraining the problem.

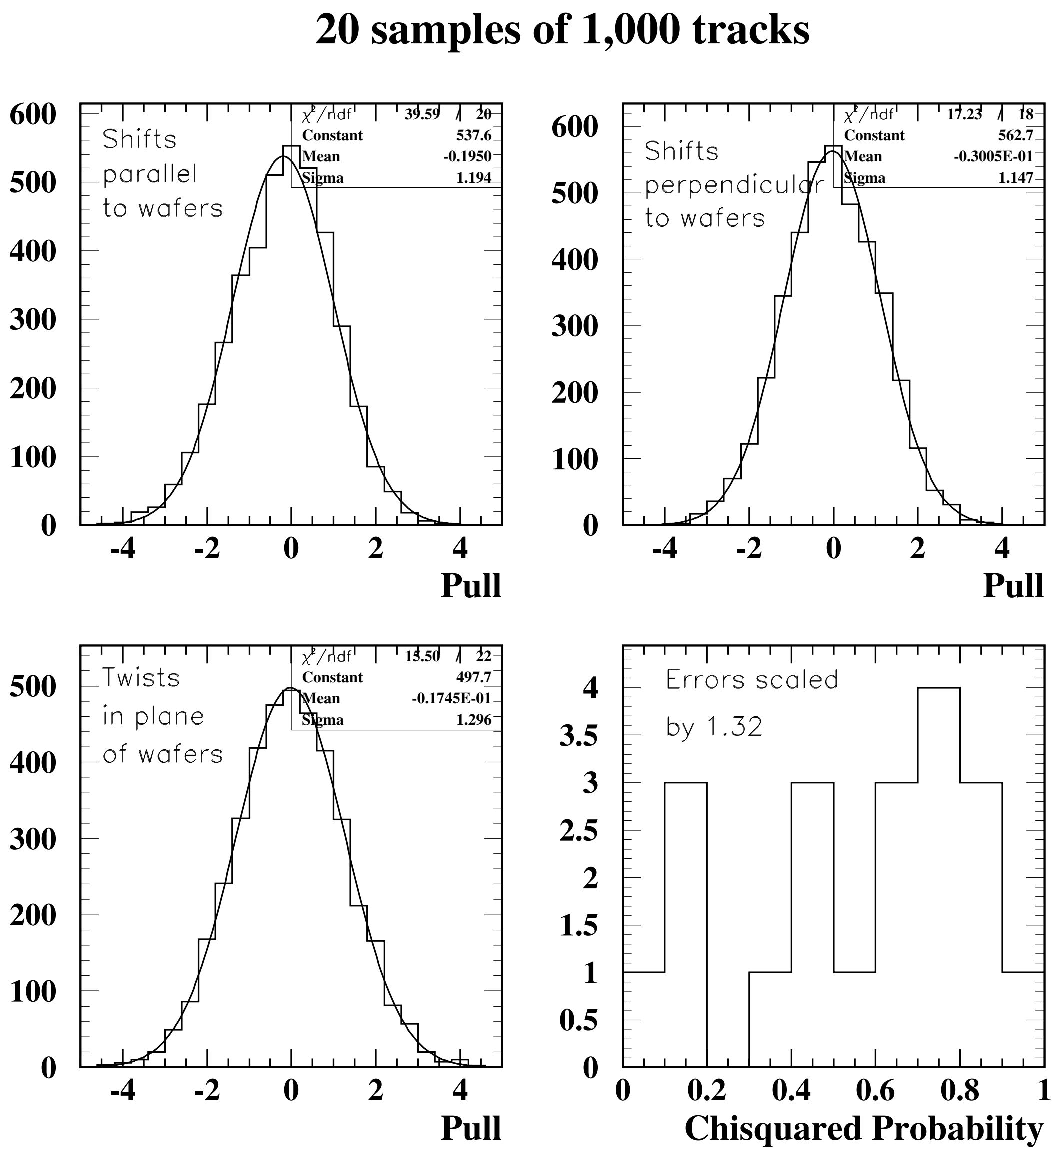

Figure 9: Pull Distributions and the chisquared probability distribution, when the multiple scattering has been underestimated by a factor of two.

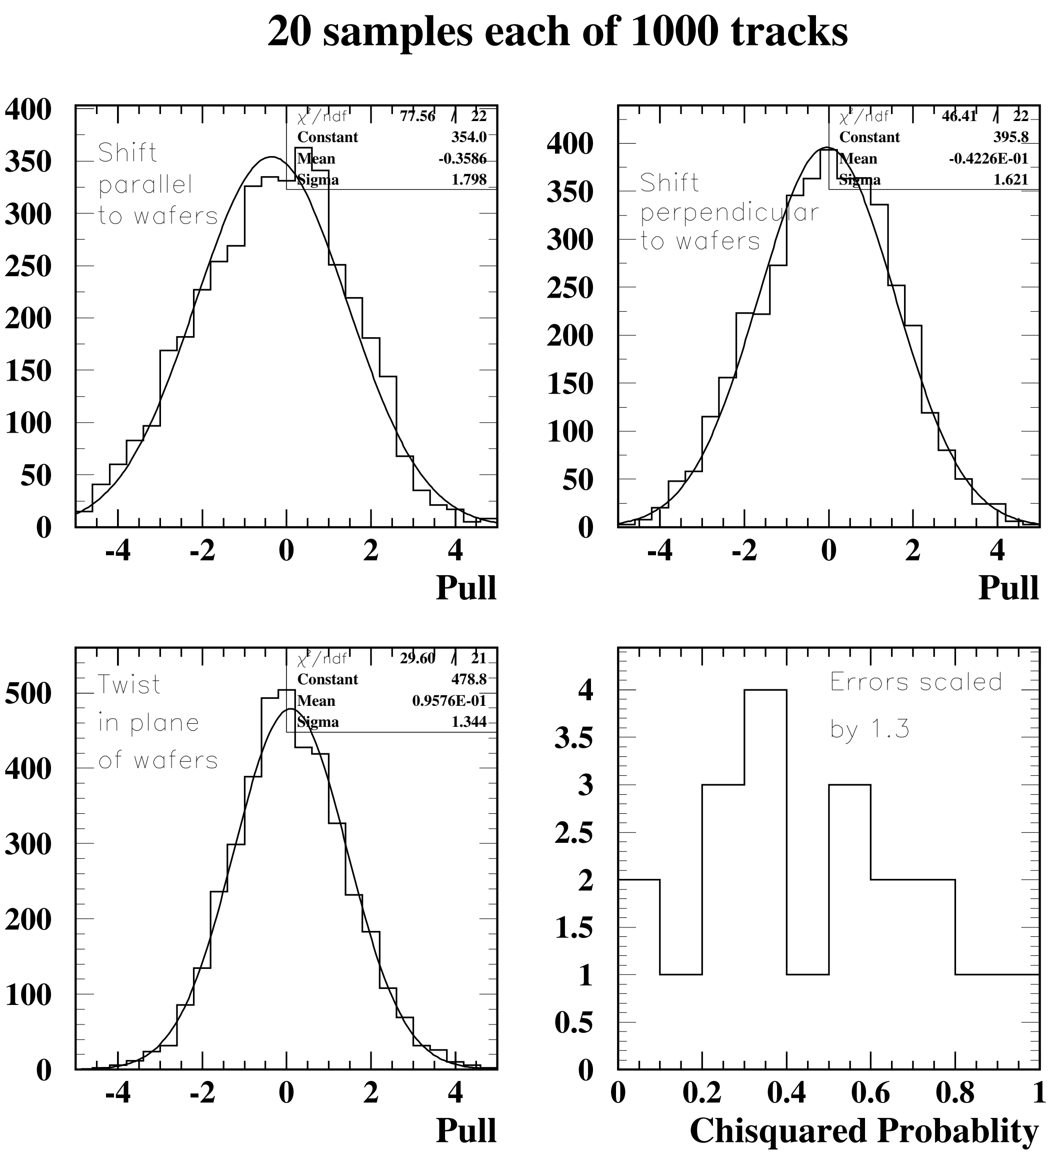

Figure 10: Pull Distributions and the chisquared probability distribution, where the estimate of the RMS of the misalignments has been underestimated by a factor of two.

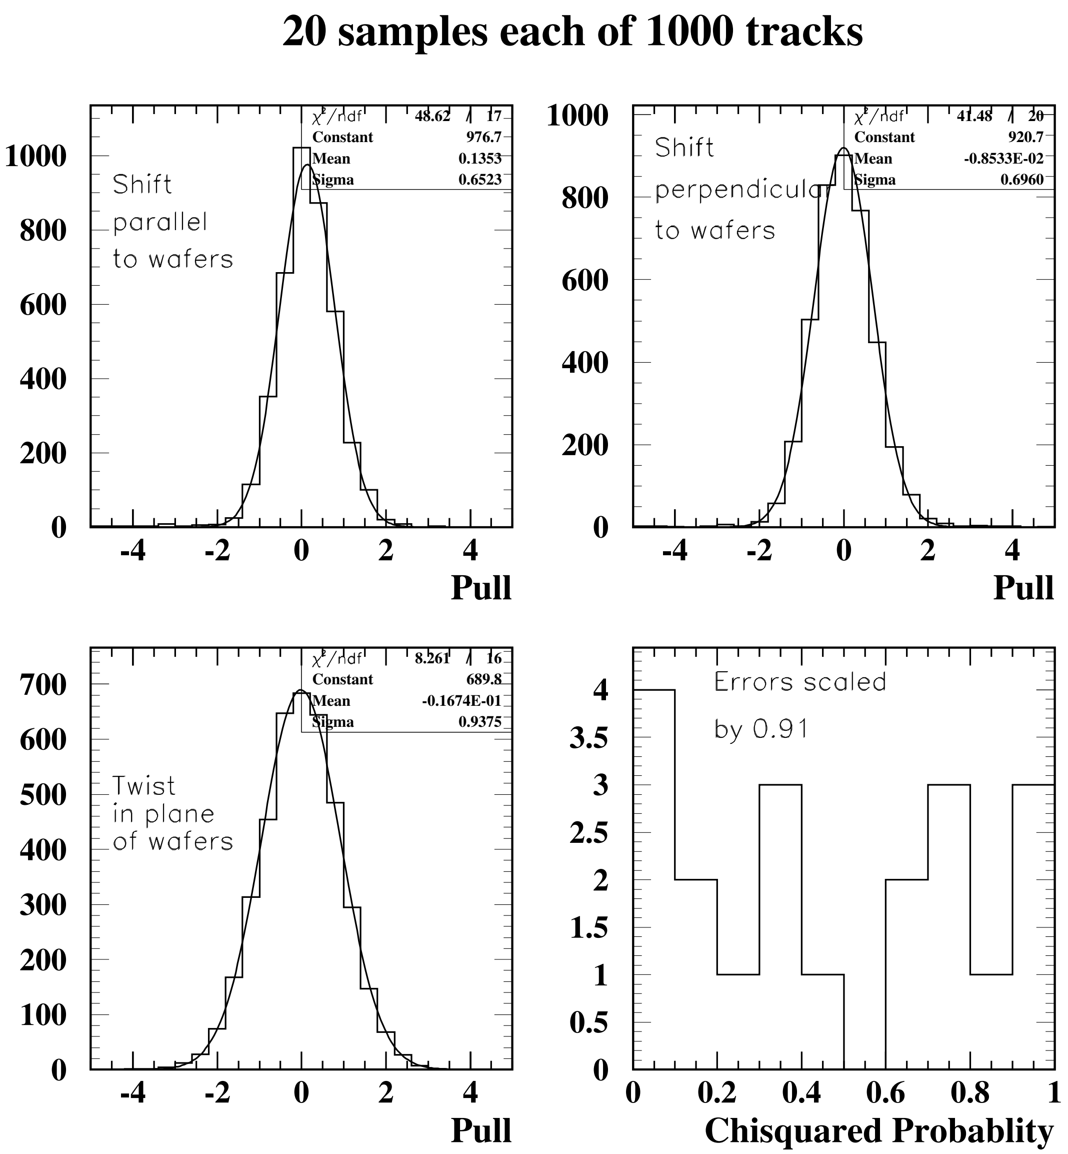

Figure 11: Pull Distributions and the chisquared probability distribution, where the estimate of the RMS of the misalignment has been overestimated by a factor of two.

5.3 Stability to Larger Misalignments.

We tested the fit's ability to reconstruct much larger misalignments (far bigger than we expect from the mechanical survey). We generated data with random misalignments of 100μm,300μm and 1mrad in $R\phi$, R and Rvz, and applied the alignment procedure with these RMS values as constraints. The pull distributions are shown in Figure 12. The errors must be scaled up by 8% in order to get a flat chisquared probability distribution. The larger scaling on the errors compared to when the misalignments were five times smaller (Figure 7) are due to the bigger deviations from the reference which start to invalidate our assumption of linearity.

If we estimate the multiple scattering from the reference trajectory instead of using the COT information, we find the errors to be underestimated by 25%. In the calculation of the reference trajectory, the large misalignments give rise to incorrect momentum determinations, which lead to an inaccurate estimate of the multiple scattering. We conclude, once again, that the precision of the result is improved if we use some external information in order to get the scattering terms correct

5.4 Dependence of the Alignment on the Misassociation of Silicon Hits to Tracks.

Finally, we examine the effect of the tracking and search algorithms, and the inevitable fact that some silicon detector hits will be misassociated and attached to the wrong tracks. This is particularly easy to do when the detector is misaligned and the normal road widths for the search do not apply. We use a special version of the Karlsruhe tracking [5] to accomplish this goal. Normally, the tracking algorithm proceeds by extrapolating tracks or track-stubs to a silicon layer and considering all silicon hits inside a certain window size. Further extrapolations to other layers allow the best track to be selected based on the total number of silicon hits on the track and the chisquared of the fit.

To deal with a misaligned detector, it is necessary but not sufficient to simply increase the window sizes, for although many candidate tracks will be collected (the correct solution among them) the best track can not be selected using the chisquared criterion. Thus an

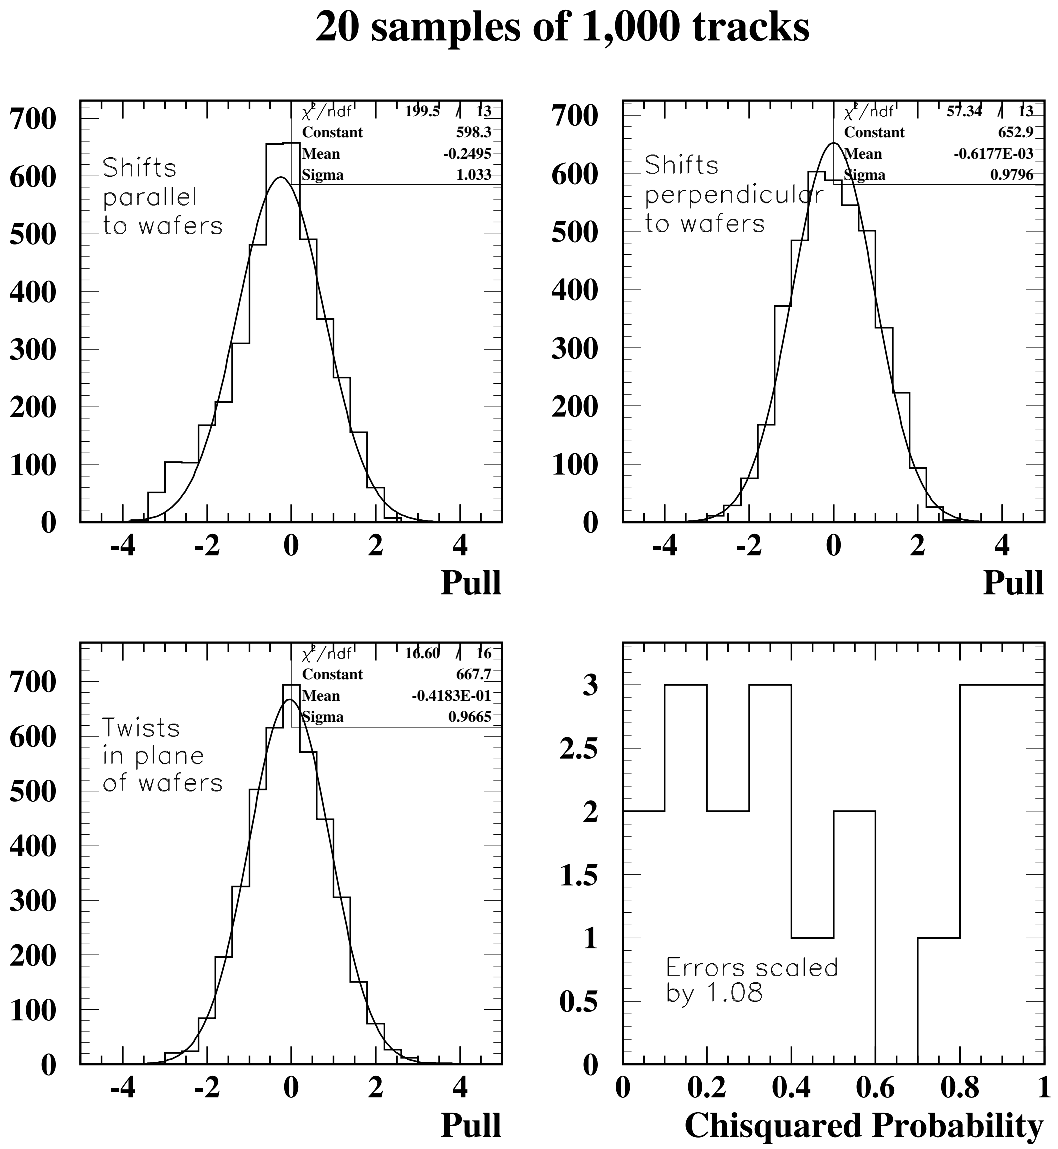

Figure 12: Pull Distributions and the chisquared probability distribution, where the RMS of the misalignments are 100μm in the $R\phi$ direction, 300μm in the R direction, and a 1 mrad twist in the plane of the silicon.

additional requirement is made that within a window of 600μm around the candidate track, there is only one single hit in a given silicon ladder. Tests using simulations of b-quark events show that even with misalignments having a RMS of 100μm in the $R\phi$ direction, this special algorithm produces tracks with the correct hits attached 95% of the time. The efficiency for tracking is of course much lower than normal. On average only about ten tracks per event are sufficiently isolated to have silicon hits attached.

The tracks found were used to align the silicon. The resultant pull distribution is shown in Figure 13. The central Gaussian-like distributions have widths which are much larger than 1, and sizeable tails (up to ten sigma) are observed. This is directly attributable to the 5% of tracks with incorrectly associated hits. That such a small contamination can seriously skew the alignment might, at first sight, be surprising. However, Blobel's alignment procedure depends on all errors being Gaussianly distributed, thus in order to minimise the global chisquared, the procedure may weigh the 5% of incorrect hits with large residuals more heavily than the 95% with small residuals.

Our initial reaction was to iterate the solution. Perhaps after a second iteration the tails in the pull distribution will be reigned in. This did not work. Not alone did the tails remain the same, but the width of the pull distribution broadened.

The procedure which gave the best results was to identify a subset of tracks for which the hit associations were correct. The usual way of doing this, would be to examine the track chisquared and remove tracks for which this was bad. This will not work straight off, since the misalignments create bad chisquareds. However, after the first iteration, it appears that we can use this criterion in order to produce a list of good tracks (not to perform an alignment). Then, with this list of good tracks, we rerun Blobel's code from scratch. In this way the 5% of incorrect hits contaminating the sample is almost eliminated. The proof that this works comes from examining the pull distributions (Figure 14) which have a unit width and no tails.

6 Conclusions

We have devised a procedure with which to align the CDF silicon

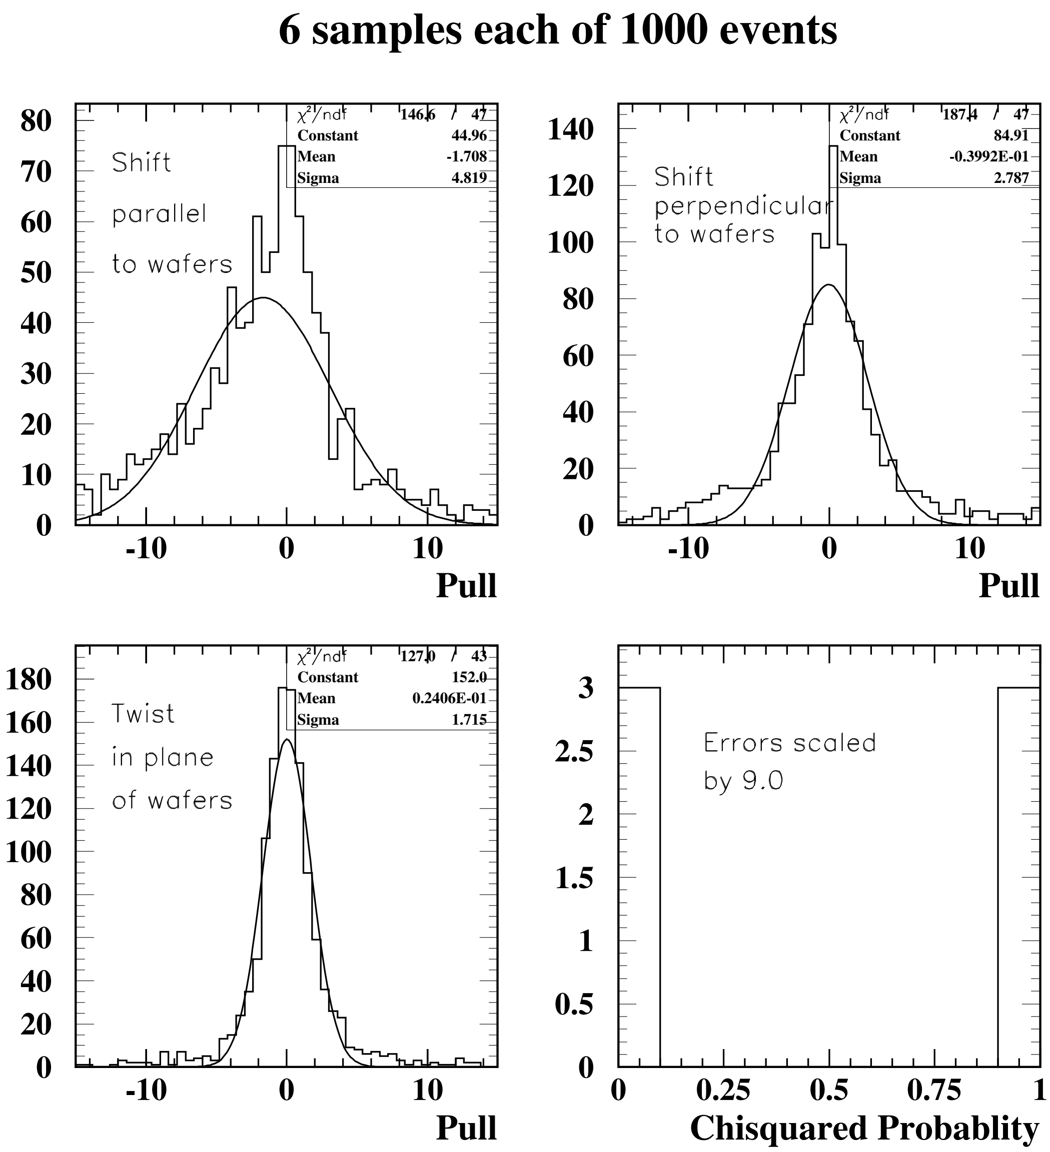

Figure 13: Pull Distributions and Chisquared Probability Distribution for six alignments performed using 1000 tracks. 5% of the tracks have an incorrect silicon hit attached.

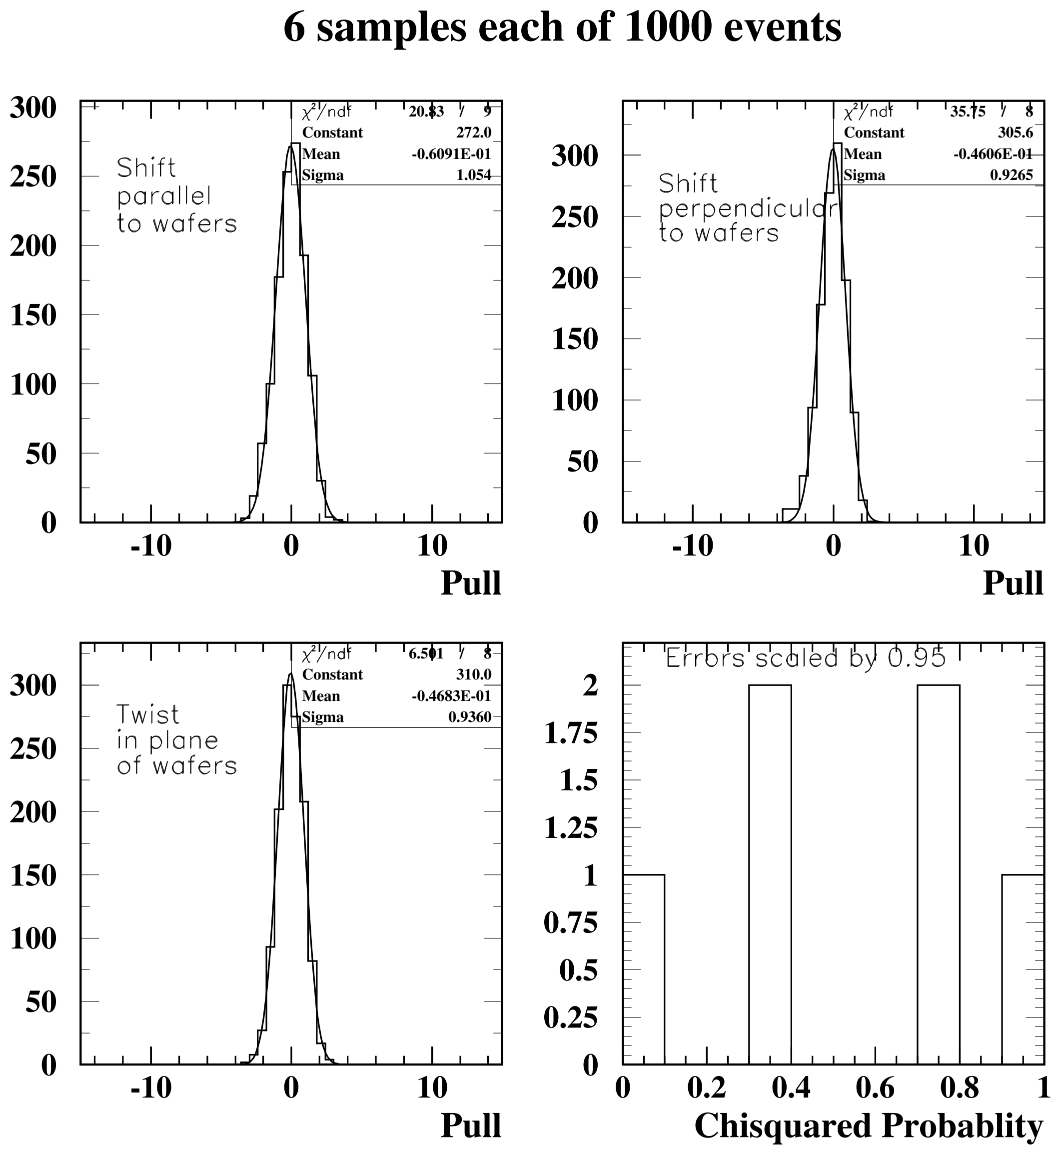

Figure 14: Pull Distributions and Chisquared Probability Distribution for six alignments performed using 1000 tracks. 5% of the tracks had incorrectly assigned silicon hits before the procedure described in the text was applied.

detectors. Starting from an estimate of the precision of the mechanical survey we have shown that with 10,000 tracks the correct positions are recovered with a global precision which is typically 6μm in $R\phi$ and 12μm in R. The error estimate is accurate to a few percent. The local alignment is known much more precisely. The impact parameter distribution is unbiased and has a resolution which is 1.5μm broader than its theoretically minimum value.

We have tested the procedure with changes in our estimates of the mechanical survey, the size of the misalignments, the material budget, and the hit misassociations, and find it to be stable provided an external reference is used to determine the size of the multiple scatters.

Acknowledgements

We wish to acknowledge the considerable help afforded us by Stephanie Menzemer in tuning the KAL tracking algorithms to produce a version which is robust against hit misassociations.

References

[1] V. Blobel, Millepede program write-up, available as a link from `http://www.desy.de/~ blobel/`

[2] V. Chabaud, A. Andreazza, P. Collins, H. Dijkstra, *Alignment of the DELPHI vertex detector*, DELPHI 95-177 MVX10.
Auke-Pieter Colijn, *Measurement of the Tau Lepton Lifetime*, Thesis, University of Amsterdam (1999).
C. H. Wang et al., *Alignment of SVX' using Run 1B data,* CDF Note 3002, (1995).

[3] A.Foland, F.Palmonari, Private Communications.

[4] G. Lutz, *Optimum Track Fitting in the Presence of Multiple Scattering* Nucl. Instr. Meth. **A273**, (1988) 349.

[5] M. Feindt, S. Menzemer, K. Rinnert, P. Schemitz, A. Skiba, *TrackingKal - a Tracking and Alignment Software Package for CDF II Using a Fast Kalman Fitter*, CDF note 5387.
M. Feindt, S. Menzemer, K. Rinnert, P. Schemitz, A. Skiba, *KalKalmanFitter - A Fast and Versatile Kalman Track-Fitter for CDF II*, CDF note 5388.