Internal Alignment of the SLD Vertex Detector using a Matrix Singular Value Decomposition Technique ¹

David J. Jackson^(a), Dong Su^(b) and Fred J. Wickens^(a)

^(a)Rutherford Appleton Laboratory, Chilton, Didcot, Oxfordshire OX11 0QX, UK

^(b)Stanford Linear Accelerator Center, Stanford, CA 94309, USA

Abstract

The tracking resolution and vertex finding capabilities of the SLD experiment depend upon a precise knowledge of the location and orientation of the 96 elements of the SLD pixel vertex detector (VXD3) in 3D space. At the heart of the deterministic procedure described here to align the 96 CCDs is the matrix inversion technique of singular value decomposition (SVD). This tool is employed to unfold the detector geometry corrections from the track hit residual data in the VXD3. The algorithm is adapted to perform an optimal χ^2 minimization by careful treatment of the errors and correlations in the residual measurements. The general form of the problems that might be solved with this technique is discussed. The tracking resolution obtained with the aligned geometry is compared with the starting point, based on an optical survey of the CCDs, and is shown to achieve the design performance.

submitted to Nucl. Inst. and Meth. A

---

¹This work was supported by the U.S. Department of Energy contract DE-AC03-76SF00515, and the U.K. Science and Engineering Research Council

1 Introduction

During the 1996-1998 run period around 400,000 hadronic Z^0 decays were recorded by the SLD Large Detector (SLD) at the e^+e^- SLAC Linear Collider (SLC). Due to the lifetimes of the b and c hadrons, tracks from $Z^0 \rightarrow b\bar{b}$ and $Z^0 \rightarrow c\bar{c}$ events have more than one origin: fragmentation tracks form the primary vertex at the e^+e^- interaction point (IP) while tracks from the decays of the heavy hadrons form displaced secondary vertices. The reconstruction of vertices formed by the tracks from the weakly decaying hadrons allow the properties and interactions of the heavy quarks to be studied [1]. The ability to reconstruct the secondary vertices depends on the tracking resolution. Misalignments of the vertex detector elements not only degrade the resolution but also introduce highly unpredictable systematic effects that are difficult to reproduce in the Monte Carlo simulation. This paper describes the technique developed for the internal alignment of the components of the SLD vertex detector which allowed the design performance to be achieved.

Figure 1: Layout of the CCDs in VXD3: (a) cross-section through the rz plane, (b) view from $+z$ end with each line representing one ladder supporting two CCDs.

The SLD Vertex Detector (VXD3) [2] consisted of 96 charge coupled devices (CCDs) [3] arranged on 3 cylindrical layers of supporting ladders, with radii of 2.7cm, 3.8cm and 4.8cm, around the IP as shown in figure 1. The z axis was parallel to the beam-line, with the positive direction towards the North face of the SLD detector. There were 48 CCDs facing the beam on the South half of each ladder together with 48 outward facing CCDs on the North half, with an overlap of around 1 mm, called a ‘Doublet’, in the region near $z = 0$. There was an overlap of similar size between neighbouring CCDs in the same layer, called a ‘Shingle’, in the $r\phi$ plane, figure 1(b). The design of the VXD3, installed in 1996, was motivated by the experience with the previous SLD vertex detector VXD2. In particular, the Doublet and Shingle overlaps are a crucial feature that not only secure the coverage and avoid gaps, but were also introduced with tracking based alignment in mind. Each of the 96

CCDs measured 8cm in length by 1.6cm in width and contained 3.2 million $20 \times 20 \mu\text{m}$ pixels. A charged track traversing a CCD would leave a hit consisting of a cluster of several pixels, within each of which the ionization charge was recorded. The centroid of the cluster pixel charge determined a point on the track trajectory with an intrinsic resolution estimated to be around $3.5 \mu\text{m}$. Sets of such hits on different CCDs within small solid angular regions about the IP were called VXD3 vectors, each assumed to have been produced by the passing of a single charged track. A vector was required to consist of at least three hits, with more than three being possible due to the overlaps of the CCDs. Charged tracks from the interaction region were identified and reconstructed initially in the Central Drift Chamber (CDC) [4]. The helix parameters of each track, in the 0.6 Tesla solenoid field, were used to determine the region of the VXD3 traversed by the corresponding charged particle and to associate the track with a VXD3 vector. Since the VXD3 hit efficiency was very high almost all tracks that passed through all three layers of the detector were found to associate with a VXD3 vector. (If such a vector was not found an attempt was made to associate the track with individual CCD hits close to the track trajectory).

In a preamble to the track based internal alignment the CCD geometry was corrected for deviations from the ideal geometry location and a flat surface with measurements from a room temperature optical survey. The calculated effects of gravitational sag were also taken into account. Further, relying both on predicted effects and observation of detector performance the geometry was corrected for mechanical contraction inside the VXD3 cryostat. These initial corrections determined the ‘nominal’ geometry of the detector, with typically a few 10s of microns precision, with which the first data was reconstructed. Subsequently a χ^2 was calculated to measure the deviation of the associated VXD3 hits from the CDC track trajectory taking into account the tracking resolution. Initially the VXD3 was treated as a rigid body and globally aligned to the CDC by minimizing the $\sum \chi^2$ over tracks from a large number of events with respect to the three translational and three rotational degrees of freedom of the whole VXD3. The impact parameter resolution at the IP of tracks fitted to the combined CDC and VXD3 hit data then depended primarily upon the CCD single hit resolution. This in turn was the product of the intrinsic resolution and the systematic uncertainties in the relative internal locations and orientations of the 96 CCDs with respect to the same set of six degrees of freedom. The aim of the internal alignment is to remove the latter contribution with the determination of these $96 \times 6 = 576$ geometric corrections. Before the physics analysis could fully benefit from the vertex detector the tracking data had to be used to improve the single hit resolution to $5 \mu\text{m}$ or less, close to the true intrinsic resolution.

A major goal of the project was to find a prescription for combining the alignment data from particle tracks in a deterministic way, so that a robust optimum geometry could be obtained, i.e. one which does not depend on the vagaries of the order in which parameters were optimised, nor require many subjective judgements to be made. We believe that this goal has been achieved and will describe the mathematical underpinnings of the result. A key tool which has enabled us to achieve the result was the use of singular value decomposition

(SVD). In appendix A we give a brief overview of SVD and its use as a tool in finding a least squares solution. We explain how the SVD technique was integrated into our alignment procedure to find an optimum χ^2 minimisation solution for the CCD geometry using the track data. In the following section we describe how the displacements of a CCD from its nominal location perturbs the position at which a track hit is observed. Section 3 discusses how hits from various types of track data are included and the constraints made on the detector geometry. A description is then given in section 4 of how these constraints are combined into a global set of equations which have been solved using SVD. In section 5 we describe our results with the VXD3 data and examples are given of the difference in the distributions of track residuals before and after the alignment. Finally we make some concluding remarks and discuss the general form of the problems to which this technique could be applied in section 6. The algorithm was first tested using the track data from a Monte Carlo simulation of a misaligned VXD3. This conveniently allowed us to identify problems due to the algorithm as distinct from those due to features in the real data. A better understanding of the procedure was achieved in successfully aligning the simulated detector, before applying the analysis to the real VXD3.

2 CCD Misalignments and Hit Residuals

Figure 2 shows the degrees of freedom allowed in this analysis for a pair of CCDs on one of the 48 ladders of VXD3.

Figure 2: Definition of the CCD translation δz , δr , $\delta \eta$ and rotation $\delta \alpha$, $\delta \beta$, $\delta \gamma$ corrections to be determined, indicated for two CCDs on the same ladder.

The alignment correction parameters consisted of three translations, parallel to the edges of the rectangular CCD and normal to the CCD plane, and three rotations about the normal point in the CCD plane (i.e. the point in the plane closest to the mean location of the IP):

- δz : Translation in CCD plane in z direction
- $\delta \eta$: Translation in CCD plane in direction of increasing ϕ
- δr : Translation normal to CCD plane away from the IP
- $\delta \alpha$: Pitch – Rotation axis along width of CCD
- $\delta \beta$: Yaw – Rotation axis normal to CCD
- $\delta \gamma$: Roll – Rotation axis along length of CCD

Figure 2 also shows how a hit $\circ$ in a CCD was specified by the angle λ between the line connecting the event IP to the hit and the $r\phi$ plane, and the distance L_ϕ across the CCD in the CCD's own reference system. The sign of these co-ordinates was chosen such that $\tan \lambda$ and L_ϕ increase with increasing z and ϕ respectively. Although the location of the IP varies slowly with time in the $r\phi$ plane and from event to event by $O(700\mu\text{m})$ in the z direction, deviations of the event IP from the mean IP position were small relative to the scale of the detector and assumed to have negligible effect on the alignment. The fact that the mean location of the IP was offset from $z = 0$ by several millimetres was of no significance; however it was important that the mean IP position in the $r\phi$ plane was close to the detector centre at $x = y = 0$. The alignment procedure described here also assumed that each CCD was approximately flat, with small shape corrections as measured in the optical survey having been applied (see further discussion on CCD shapes in section 5).

Figure 3 illustrates the effect on the apparent hit position within a CCD for a track of fixed trajectory when the CCD position is perturbed by movements in the six degrees of freedom. Misalignments of the CCDs cause the measured hit on the CCD to be displaced from the true track trajectory by a residual amount δ_z along the CCD length and $\delta_{L\phi}$ across its width. The sign of δ_z was defined such that it measured the z location of the hit minus the z location of the track in the plane of the CCD. The sign of δz and the other alignment corrections was chosen such that these parameters indicate the required reversal of the misalignments. For example, if the only degree of freedom of the CCD was translation along the z axis, then since $\delta_z + \delta z = 0$ (misalignment + correction = no displacement) it would be trivial to ‘unfold’ the required CCD alignment correction δz from the measured residual δ_z and apply the δz shift to the CCD to align the hit to the track. In general each of the degrees of freedom can contribute to the residual. Straight forward geometric arguments show that the δ_z residual may also be compensated by δr and $\delta \beta$ according to:

Figure 3: The apparent hit position in a CCD due to misalignment.

$$\delta_z = -\delta z + \delta r \cos \epsilon \tan \lambda + \delta \beta L_\phi. \quad (1)$$

Since $\delta \eta$ is directed out of the page in figure 3(a) it does not affect δ_z . The other two CCD corrections $\delta \alpha$ and $\delta \gamma$ are most easily approximated, retaining only first order correction terms, as a radial effect:

$$\delta r \rightarrow \delta r + \delta \alpha \frac{r}{\cos \epsilon} \tan \lambda + \delta \gamma L_\phi. \quad (2)$$

Substituting equation 2 into equation 1 and taking the simplifying approximation $\cos \epsilon = 1$ (the minimum value of $\cos \epsilon$ is actually around 0.88 for the lower ϕ edges of the layer 1 CCDs) yields:

$$\delta_z = -\delta z + \delta r \tan \lambda + \delta \alpha r \tan^2 \lambda + \delta \gamma L_\phi \tan \lambda + \delta \beta L_\phi. \quad (3)$$

As more alignment degrees of freedom are taken into account it clearly becomes less trivial to unfold the true geometry, and a larger number of residuals need to be measured in order to constrain the geometry corrections. Similar arguments based on the geometry shown in figure 3(b) leads to the single hit residual for the $r\phi$ plane:

$$\delta_{L\phi} = -\delta\eta + \frac{\delta r}{r}L_\phi + \frac{\delta\gamma}{r}L_\phi^2 + \delta\alpha L_\phi \tan\lambda - \delta\beta r \tan\lambda \quad (4)$$

where again the residual is measured in the plane of a CCD.

Although the CDC was used to identify charged tracks most internal alignment observables were calculated using the CCD hits only, independent of the VXD3 to CDC alignment. In general since the true track trajectory was unknown it was necessary to identify hits on at least three CCDs associated with a reconstructed track in order to determine a relative residual internally in the VXD3. For the initial residual types studied good quality tracks were constrained to pass through two of the CCD hits with the corresponding residual measured to the third, reference, CCD. The small curvature effect of the charged track of known momentum, measured by the CDC in the SLD magnetic field, was taken into account in the $r\phi$ plane. The relative lever-arm weights with which each of the three CCDs contributed to the observed residual were determined from the ideal geometry to within a very good approximation (since the alignment corrections are very small compared with the dimensions of the detector), in all cases the reference CCD was given a weight = +1.0. The CCD contributions to the $\tan\lambda$ terms in equations 3 and 4 were readily combined by using the $\tan\lambda$ parameter of the track itself which, again to a very good approximation, is equivalent to the $\tan\lambda$ of each of the associated CCD hits. Terms involving L_ϕ required more care since in general this value differs for each CCD even for hits from the same track. However we determined L_ϕ on the reference CCD and to a close approximation the value of this local co-ordinate on another CCD i , L_{ϕ_i} , was linearly related as $L_{\phi_i} = a_i L_\phi + b_i$ where a_i and b_i were constant factors determined solely by the ideal geometry of the CCDs concerned. Taking all of these terms into account lead to the following general functional forms for the residual measured on the reference CCD for a given track:

$$\begin{aligned} \delta_z = \sum_i w_i [& (-\delta z_i + b_i \delta\beta_i) & \mathbf{1} \\ & + (\delta r_i + b_i \delta\gamma_i) & \mathbf{\tan\lambda} \\ & + (r_i \delta\alpha_i) & \mathbf{\tan^2\lambda} \\ & + (a_i \delta\gamma_i) & \mathbf{L_\phi \tan\lambda} \\ & + (a_i \delta\beta_i) & \mathbf{L_\phi}] \end{aligned} \quad (5)$$

$$\begin{aligned} \delta_{L\phi} = \sum_i w_i f_i [& (-\delta\eta_i + \frac{b_i}{r_i} \delta r_i + \frac{b_i^2}{r_i} \delta\gamma_i) & \mathbf{1} \\ & + (\frac{a_i}{r_i} \delta r_i + \frac{2a_i b_i}{r_i} \delta\gamma_i) & \mathbf{L_\phi} \\ & + (\frac{a_i^2}{r_i} \delta\gamma_i) & \mathbf{L_\phi^2} \\ & + (a_i \delta\alpha_i) & \mathbf{L_\phi \tan\lambda} \end{aligned} \quad (6)$$

$$+ (b_i \delta \alpha_i - r_i \delta \beta_i) \tan \lambda]$$

where the sum is over the contributing CCDs and their signed weights w_i were determined relative to the reference CCD. For example, the layer 2 reference CCD of a Triplet (described in section 3) was given a weight of +1.0 while the layer 1 and layer 2 CCDs each contributed with a weight $w_i = -0.5$ according to their geometric lever arm. Corrections to the non-reference CCD contributions due to the fact that the CCDs are not parallel in the $r\phi$ plane were determined from the ideal geometry and yielded the factors $f_i (\sim 1.0)$. For most of the residual types a reduced functional form could be used and these are shown in table 1 and described in the following section.

3 Residual Types, Distributions and Fits

The internal alignment began by classifying the types of tracking constraints as shown in figure 4 with the corresponding residuals δ indicated in the $r\phi$ plane. In each case there are analogous residuals in the rz plane. There are 48 pairs of CCDs shown in figure 4 with the thick and thin lines representing the North and South CCDs respectively. The six residual types used are described in the following paragraphs. The first three, using tracks with a momentum of at least 1 GeV/ c from hadronic Z^0 decays, provided local constraints on the detector geometry.

- Shingles - require a pair of hits in adjacent CCDs that overlap within the same layer, plus a third, anchor, hit in another layer. These constrain the relative positions of CCDs within a layer, and in particular the layer radius, with the anchor weight being relatively small. Since each Shingle covers only a small ϕ range a constant average L_ϕ was determined for each CCD of each Shingle and the residual fits were made to a function of $\tan \lambda$ only, see table 1.
- Doublets - require a hit in each of the two CCDs on a single ladder, in the overlap region halfway along the ladder length, plus a third hit in another layer. These constrain the relative positions of the North and South halves of the VXD3. In this case only a small range of λ is covered, a constant $\tan \lambda$ was determined for each layer, and the fits were made to a function of L_ϕ only.
- Triplets - require three hits in different layers, either all on the North or all on the South half of the detector. These constrain the relative positions of the three layers. Each Triplet covers a range of λ and ϕ and full five parameter fits of the residual distributions as a function of $\tan \lambda$ and L_ϕ were performed.

Figure 4: Construction of the VXD3 residual types.

- Pairs - provide long range constraints on the geometry using pairs of back-to-back tracks from selected $Z^0 \rightarrow \mu^+\mu^-$ or $Z^0 \rightarrow e^+e^-$ events. Layer 2 was not used for this and the following two residual types. One track was constrained by hits in layers 1 and 3 of the North CCDs while the other track was correspondingly constrained to the South CCD hits. Due to the symmetry of the VXD3, figure 4, a reasonably low number of unique four CCD combinations needed to be considered. These events relate opposite regions of the detector with the residuals being measured as the miss-distance between the two tracks at the IP in the z direction and in the $r\phi$ plane. The ϕ angle difference provided an additional constraint while the λ angle difference was not used due to the small uncertainty in the symmetry of the SLC beam energies.
- VXD3 vs. CDC Track Angle Matching - using the individual high momentum tracks of the Pair events the track angle measured by the CDC was compared with the angle measured from the hits in layers 1 and 3 of VXD3. By multiplying the angular discrepancy in radians by the distance between the two CCDs, and correcting for the angle of incidence of the track to the reference CCD, the residual angle was converted into a residual distance on the reference CCD and treated in the same way as the other residual types. The residual uncertainty was determined from the CDC track as well as the CCD hit resolution.
- IP constraint - on average tracks measured in layers 1 and 3 should project back to the SLC interaction point, which was small and reasonably stable in the xy plane. Tracks with momentum greater than 7 GeV/ c from light flavour selected hadronic Z^0 decays were used for this constraint. In addition to the 6×96 CCD corrections a further two, δx and δy , were introduced for the offset of the IP, common for all data periods, from the nominal location of the event IP. Hence with this constraint added there were a total of 578 alignment corrections to be determined. The residual at the IP, effectively an impact parameter $\delta_{r\phi}$, listed at the bottom of table 1, had two terms. The constant term $i_1^\perp$ included a contribution from the IP offset corrections δx and δy and is given by:

$$i_1^\perp = f(\delta_{\text{CCD}}) - \delta x \sin \phi_i + \delta y \cos \phi_i \quad (7)$$

where $f(\delta_{\text{CCD}})$ was determined from equation 6 in a manner similar to that for the CCD contributions to the $p_1^\perp$ term for the Pairs and the constants ϕ_i were the average value of ϕ for each residual region.

The dashed and dash-dotted lines in the upper right quadrant of figure 4 indicate the division of the data into ten Triplet regions, while the dash-dotted lines only indicate the division into seven regions used to define the Pair, CDC angle matching and IP constraint residual regions. The boundaries were determined by the requirement that each residual region involved a unique set of CCDs. These divisions repeat in the other three quadrants by the symmetry of the detector. The 40 Triplet regions correspond to the $N_I = 80$ independent

fits listed in table 1 since North and South CCDs were considered separately for the Triplets. For the last three sets of residual types the value of L_ϕ in equations 5 and 6 was taken to be a constant determined for each CCD in each residual region, similarly as described above for the Shingles. (Dropping this assumption and performing the CDC angle matching fits to the full set of five parameters to each of δ_λ and δ_ϕ did not improve the ultimate result). The full set of functional forms is listed in table 1. The philosophy of the alignment procedure was to find the set of small corrections, one for each degree of freedom for each CCD, which gave a global minimization of the residuals over all of the residual types. With each observable being sensitive to several alignment corrections a simultaneous determination of all these parameters, being the only way to achieve a controlled solution, was essential.

Type	Functional Form	N_I	N_C
Shingles	$\delta_z = s_1^\parallel + s_2^\parallel \tan \lambda + s_3^\parallel \tan^2 \lambda$	96	288
	$\delta_{L\phi} = s_1^\perp + s_2^\perp \tan \lambda$	96	192
Doublets	$\delta_z = d_1^\parallel + d_2^\parallel L_\phi$	48	96
	$\delta_{L\phi} = d_1^\perp + d_2^\perp L_\phi + d_3^\perp L_\phi^2$	48	144
Triplets	$\delta_z = t_1^\parallel + t_2^\parallel \tan \lambda + t_3^\parallel \tan^2 \lambda + t_4^\parallel L_\phi \tan \lambda + t_5^\parallel L_\phi$	80	400
	$\delta_{L\phi} = t_1^\perp + t_2^\perp L_\phi + t_3^\perp L_\phi^2 + t_4^\perp L_\phi \tan \lambda + t_5^\perp \tan \lambda$	80	400
Pairs	$\delta_{rz} = p_1^\parallel + p_2^\parallel \tan \lambda + p_3^\parallel \tan^2 \lambda$	28	84
	$\delta_{r\phi} = p_1^\perp + p_2^\perp \tan \lambda$	28	56
	$\delta_\phi = p_1^\phi + p_2^\phi \tan \lambda$	28	56
CDC Angle	$\delta_\lambda = c_1^\lambda + c_2^\lambda \tan \lambda + c_3^\lambda \tan^2 \lambda$	56	168
	$\delta_\phi = c_1^\phi + c_2^\phi \tan \lambda$	56	112
IP Constraint	$\delta_{r\phi} = i_1^\perp + i_2^\perp \tan \lambda$	56	112
Total		700	2108

Table 1: Functional forms for fitting the various residual types. The superscripts $\parallel$ and $\perp$ indicate coefficients of fits to residuals measured parallel and perpendicular to the z -axis respectively. N_I lists the total number of independent residual fits involved while N_C is the number of coefficients determined.

For each type of residual n-tuples were accumulated containing the deviations, such as δ_z and $\delta_{L\phi}$, and co-ordinates $\tan \lambda$ and L_ϕ with each residual region and it's unique combination of CCDs identified by an integer index label. Data in the n-tuples were fitted to the functional forms given in table 1 to determine the coefficients and the covariance error matrix of each fit. The residual fits were made using MINUIT [5] with an automated

procedure to loop over the large number of residual types and CCD combinations involved. Examples of the fits for a Shingle are shown in figure 5.

Figure 5: An example of (a) δ_z and (b) $\delta_{L\phi}$ residual distributions as a function of $\tan \lambda$ for a Shingle in layer 2 before alignment. The solid lines show the fitted curves.

Each of the points in figure 5 represents the residual obtained from one track. All residual fits were made to data points obtained by binning the distributions in $\tan \lambda$ (and/or L_ϕ) and averaging over the track data in each bin to determine the data point with an error bar, taking the intrinsic detector resolution into account. After each fit the automated procedure included the search for and removal of outliers (assumed to be caused by CCD hits incorrectly associated with tracks) and a refit to the remaining data if necessary. This was particularly important for fits with relatively little track data, and in all cases a visual cross-check was made using a graphical fit display. The scatter of the points about the fitted curves in figure 5 is due to the intrinsic resolution. The effects of the misalignments to be corrected are seen as a parabola and constant gradient slope in figure 5(a) and (b) respectively, as predicted by the corresponding functional forms for Shingles listed at the top of table 1.

We wished to determine the 578 alignment corrections from the 2108 coefficients obtained from the residual fits as listed in table 1. The preceding geometrical arguments however give the reverse information, indicating how the measured coefficients depend upon the alignment corrections. The repeated application of equations 5 and 6 yields a large set of simultaneous equations which can be expressed in terms of a single matrix equation $\mathbf{Ax} = \mathbf{c}$, where $\mathbf{x}$ is a column matrix of the 578 CCD and IP degrees of freedom, $\mathbf{c}$ is a column matrix of the 2108 coefficients measured from the residual fits listed in table 1 and the 2108×578 elements of $\mathbf{A}$, called the ‘design matrix’, were determined by the ideal geometry:

where for example $\delta\eta_j$ denotes the 96 CCD $\delta\eta$ corrections to be determined. With $\mathbf{A}$ and $\mathbf{c}$ known the question is how to find the optimal solution to this equation for the alignment corrections $\mathbf{x}$. Several features of matrix singular value decomposition are described in appendix A. The important property of interest here is that inverting the possibly singular $m \times n$ matrix $\mathbf{A}$ using SVD yields the $n \times m$ matrix $\mathbf{A}^+$ such that the solution $\mathbf{x} = \mathbf{A}^+ \mathbf{c}$ minimizes the magnitude $|\mathbf{Ax} - \mathbf{c}|$. This ‘least squares fit’ minimization of $\sum_i (A_{ij}x_j - c_i)^2$ could directly be used to determine an alignment geometry. However, this solution does not take into account the error and correlation information contained in the covariance matrices of the residual fits. It is necessary to redefine the basis of parameters for each residual fit such that each covariance matrix then takes a unit diagonal form for the least squares fit

using SVD to effectively become a more optimal χ^2 fit.

Quite generally the $n \times n$ covariance error matrix $\mathbf{V}$ for a fit to the n Gaussian parameters of vector $\mathbf{p}$ can be uniquely decomposed as $\mathbf{V} = \mathbf{H}\mathbf{H}^T$ where $\mathbf{H}$ is a non-singular lower triangular matrix. The new vector $\mathbf{p}' = \mathbf{H}^{-1}\mathbf{p}$ then contains a set of n independent Gaussian variables with a unit covariance error matrix [6]. We used the CERNLIB [7] routines TRCHLU to decompose the error matrix of each residual fit, $\mathbf{V} = \mathbf{H}\mathbf{H}^T$, and TRINV to invert the triangular matrix $\mathbf{H} \rightarrow \mathbf{H}^{-1}$. Multiplying each $\mathbf{H}^{-1}$ matrix into the relevant elements of $\mathbf{c}$ transforms away the residual fit correlations with each of the new elements having an error of unity. The number of lower triangle elements in the 700 triangular $\mathbf{H}^{-1}$ matrices from the 700 residual fits, given by $\frac{1}{2}n(n+1)$, varied from three for the two parameter fits to fifteen for the five parameter fits of table 1. Symbolically we collect all these elements into a single 2108×2108 matrix denoted $\mathbf{T}$ and in order to maintain the equality in equation 8 multiply into both sides:

$$\mathbf{T}\mathbf{A}\mathbf{x} = \mathbf{T}\mathbf{c}, \quad (9)$$

yielding:

$$\mathbf{A}'\mathbf{x} = \mathbf{c}', \quad (10)$$

where the rows of $\mathbf{A}'$ and elements of $\mathbf{c}'$ have been obtained by the same linear transformation, under the operation of $\mathbf{T}$, from the rows of $\mathbf{A}$ and elements of $\mathbf{c}$ respectively. The new $\mathbf{A}'$ matrix is then decomposed according to the SVD prescription and the inverse $\mathbf{A}'^+$ constructed. Finally, the solution $\mathbf{x} = \mathbf{A}'^+\mathbf{c}'$ determines a χ^2 minimization over all the track residual data and yields the 578 alignment corrections.

5 Results of Vertex Detector Alignment

5.1 Including CCD Shape Corrections

After following the prescription discussed in the previous sections and correcting the geometry according to the SVD solution optimal resolution in the $r\phi$ plane was effectively achieved. However the rz performance was less satisfactory, particularly for high $|\tan \lambda|$ where tracks make a large angle to the CCD normal and the resolution becomes more susceptible to the radial uncertainty in the CCD hit location. From the optical survey data the shapes of the CCD surfaces had been fitted to 14 parameter Chebychev polynomials [2]. In particular, the

CCDs were observed to take a 4th order polynomial ‘W’ shape, seated on their glue attachments to the ladders, along the z -direction, with deviations from a flat surface typically a few 10s of microns in size. This systematic distortion of mechanical origin was expected to be modified under the detector cooldown from room to operating temperature. In fitting for the six alignment corrections for each of the 96 CCDs it had been assumed that the shapes were perfectly known from the optical survey. Examination of the structure of individual residual region distributions showed deviations from the assumed functional forms listed in table 1, consistent with the ‘W’ shapes not having been fully corrected – due to both the limited survey data resolution and the thermal difference between the room-temperature survey and the VXD3 cryostat. The effects seen in the residual distributions, indicated that the resolution might be further improved by applying shape corrections based on the track data.

With the δz and $\delta\alpha$ rigid body CCD corrections already accounting for two of the five variables required to parameterize a 4th order polynomial, three further shape correcting terms δq , δh and δt (representing the displacement from a flat surface at $\frac{1}{4}$, $\frac{1}{2}$ and $\frac{3}{4}$ of the distance along a CCD with both ends fixed) were introduced. By analogy with equation 2 the shape was treated as a $\tan\lambda$ dependent radial correction: $\delta r \rightarrow \delta r + f(\delta q, \delta h, \delta t, \tan\lambda)$, where the function f was determined by a few lines of straight-forward algebra. Substituting this redefined expression for δr into equations 3 and 4 and following the complete procedure described in the previous three sections with the addition of the shape correcting terms led to a new version of the matrix equation $\mathbf{Ax} = \mathbf{c}$ of equation 8. This consisted of a 4,160 $\times$ 866 design matrix $\mathbf{A}$ determined from the ideal geometry, an array of 866 unknown alignment corrections $\mathbf{x}$ (now including the δq , δh and δt displacements for each CCD) and an array of 4,160 coefficients $\mathbf{c}$ determined from the original 700 residual fits with new higher order terms accounting for the CCD shapes. Since there was no track data to constrain the shape in the high $|\tan\lambda|$ region of layer 2 and particularly layer 1 CCDs, (as can be seen from figure 1(a) tracks from the IP that traverse any part of a layer 3 CCD can only intersect about half of a layer 1 CCD) a further 866 rows, each with only one non-zero element, were appended to the matrix $\mathbf{A}$ and the array $\mathbf{c}$ was extended downwards with 866 zero elements. These extra lines were used to apply conservative restraints on the geometry by effectively combining the dummy measurements $\delta z, \delta r = 0 \pm 10\mu\text{m}$, $\delta\eta = 0 \pm 2\mu\text{m}$, $\delta\alpha, \delta\beta, \delta\gamma = 0 \pm 0.05\text{mrad}$ and $\delta q, \delta h, \delta t = 0 \pm 5\mu\text{m}$ (δx and δy for the IP were not ultimately constrained in this way) into the SVD χ^2 minimization fit. The complete 5026 $\times$ 866 element matrix $\mathbf{A}'$ and the 5026 element vector $\mathbf{c}'$ were obtained from $\mathbf{A}$ and $\mathbf{c}$ respectively by equation 9 using the 16,332 non-zero elements of the effective matrix $\mathbf{T}$ obtained in turn from the residual fit covariance error matrices as described in the previous section. The matrix $\mathbf{A}'$ was inverted using SVD and the VXD3 geometry was corrected with the 866 elements of the solution vector $\mathbf{x} = \mathbf{A}'^+ \mathbf{c}'$.

5.2 Achievement of Design Performance

The alignment procedure described in the previous sections requires only a single iteration to a given data set to determine the corrected geometry. In practice, due to a difference in run conditions and as the algorithm developed, several aligned geometries were determined. In particular, during the 1996 and 1997-98 run periods the VXD3 was operated at 220 Kelvin and 185 Kelvin respectively, with the difference in thermal contraction requiring the two periods to be considered as being essentially independent. In this section we show our results with the 350,000 hadronic Z^0 events collected during the latter period, which comprises the bulk of the VXD3 data. Figure 6 shows the Triplet δ_z and $\delta_{L\phi}$ residuals obtained with the aligned detector together with the pre-alignment residuals derived from the optical survey geometry. Since the Triplets cover the full volume of the detector, that is over the full region of track acceptance for all CCDs in all three layers, they provide a representative indication of the local alignment. The data used for these Triplet plots correspond to charged tracks selected with momentum greater than 5 GeV/c to suppress the multiple scattering contribution.

Figure 6: Triplet residuals for 1997-98 data obtained with the original survey geometry (clear histograms with thick outline) and after the alignment (shaded histograms).

The post-alignment RMS of the residual distributions was found to be around a factor of four improved over the pre-alignment RMS values. Similar plots for all residual

types, along with the mean and RMS measurements, were a major guide in debugging and refining the alignment algorithm. Indeed, only after taking into account the full residual fit error matrices with the operation of $\mathbf{T}$ as expressed in equation 9 were the post-alignment residual plots observed to stand out in a manner approaching the ideal performance with an RMS dominated by the intrinsic CCD hit resolution about an essentially zero mean as seen in figure 6. Fitting a single Gaussian curve to each of the post-alignment histograms in figures 6(a) and (b) yielded a width of $4.45\mu\text{m}$ in each case. This number, depending on the three Triplet CCDs, is divided by the geometric weight factor $\sqrt{1.0^2 + 0.5^2 + 0.5^2}$ to yield $3.63\mu\text{m}$ as the single CCD hit resolution for both δ_z and $\delta_{L\phi}$ after the alignment, well within the initial target goal of $5\mu\text{m}$ and close to the intrinsic hit resolution. Erratic variations in the residual distribution mean as a function of the Triplet index that had been observed in the initial data were reduced to the $1\mu\text{m}$ level, as can be seen in figure 7.

Figure 7: Triplet δ_z and $\delta_{L\phi}$ residual mean as a function of the ϕ -dependent Triplet index (see figure 4), for the North and South halves of VXD3, obtained before $\square$ and after $\bullet$ the alignment corrections (plot (c) made with different vertical scale).

The magnitudes of the measured geometry corrections were typically $O(10\mu\text{m})$ with the largest effects being perpendicular to the CCD plane. A second iteration of the global alignment, described in section 1, was performed following the internal alignment and yielded only minor corrections for the six rigid body degrees of freedom of the detector as a whole. The problems of internal and global alignment were essentially decoupled. It should be noted that the SVD procedure does not need to be iterated on the same data set. This was expected given the relatively minor nature of the approximations made in the analysis and was confirmed by the observed negligible effect of performing a second iteration of the internal alignment.

As a measure of the long range VXD3 alignment the Pair miss-distance at the IP in the 1997-98 data with the post-alignment geometry is shown in figure 8 for comparison with the same data using the initial optical survey geometry. The Pair residuals in this plot have been divided by $\sqrt{2}$ to indicate the impact parameter resolution of high momentum tracks at the IP. (For figures 6, 7 and 8 δ_z and δ_{rz} were measured in the direction perpendicular to the track trajectory, consistent with the definition of the impact parameter resolution, whereas for the residual fit data used for the alignment δ_z and δ_{rz} were measured along the z -axis direction, parallel to the CCD plane).

Figure 8: Single track contributions to the Pair residuals measured at the IP in the (a) rz and (b) $r\phi$ planes as a function of the track angle before and after the alignment corrections.

Given that one track in a Pair event was constrained to pass through CCD hits in layers 1 and 3 (figure 4) only and taking into account the geometric level arm, a single hit resolution of $4\mu\text{m}$ corresponds to an impact parameter resolution at the IP of

$4 \times \sqrt{(4.7/2.0)^2 + (2.7/2.0)^2} = 10.8\mu\text{m}$, neglecting the effects of multiple scattering for the high momentum track. In general the one hit resolution for the Shingle, Doublet, Triplet and Pair residuals over the whole detector was found to be consistently better than $4.0\mu\text{m}$ in both the rz and $r\phi$ plane, very close to the true intrinsic CCD resolution, and design performance was achieved.

6 Conclusions

With the internal alignment of the vertex detector complete the charged track parameters were determined using a combined fit to the VXD3 hits and CDC (Central Drift Chamber) track data. Compared with the VXD3 only Pair data discussed above the impact parameter resolution improves due to the inclusion of the layer 2 VXD3 hits and further in the $r\phi$ plane due to the CDC information. The constant and momentum/angle dependent multiple scattering terms for the final rz and $r\phi$ impact parameter resolution at the IP in the 1997-98 SLD data were determined to be:

$$\sigma_{rz} = 9.7 \oplus \frac{33}{p \sin^{3/2} \theta} \mu\text{m} \qquad \sigma_{r\phi} = 7.8 \oplus \frac{33}{p \sin^{3/2} \theta} \mu\text{m} \qquad (11)$$

where the constant terms depend primarily upon the VXD3 CCD intrinsic resolution and geometry. In order to achieve this design performance a great deal of effort outside of the SVD analysis presented here was required. This work included the initial optical survey and corrections for electronics and mechanical effects (for more details see reference [2]).

Ultimately, however, only the SVD analysis had the necessary versatility to simultaneously correct the geometry of each of the 96 individual CCDs for three translational, three rotational and three shape defining degrees of freedom. The variety of residual types, including the Pair, CDC matching angle and IP constraint data, successfully incorporated into the SVD fit further demonstrate the remarkable robustness of the technique. The quality of the solution depended in turn on the careful inclusion of the residual fit error information into the analysis as described in section 4 and equation 9 to adapt the least squares solution property of the SVD method into an optimal χ^2 minimization fit. This technique could, of course, be used for quite general applications beyond that of aligning vertex detectors. Any system in which the required solution takes the form of a perturbation described by up to $O(1000)$ parameter corrections which are small in relation to the dimensions of the system as a whole and for which constraining data exists that can be expressed in terms of a set of simultaneous equations for the parameters could in principle be treated as described in this paper. Little need for subjective judgements or ‘tuning’ is required except for treating effects where the real data is observed to significantly deviate from the ideal model and assumptions required to define the problem.

On a practical note we have demonstrated that it is possible using SVD to handle simply and reliably the matrices required for the VXD3 alignment (inversion of sparse matrices of order of 5000×1000 elements) using double precision arithmetic in modest times on a standard workstation. Although the size of the matrices involved appears daunting only $\sim 1\%$ or $\sim 35,000$ elements of the final 5026×866 design matrix $\mathbf{A}$ were given non-zero values and these in turn were generally determined in a straight-forward way from only ~ 10 parameters describing the VXD3 geometry in figure 1. Further, the large number of non-zero elements involved in the $\mathbf{T}$ matrix and $\mathbf{c}$ array were obtained from the residual data skims by automated looping over large numbers of similar fits each identified by an index label. Much of the analysis was an exercise in book-keeping.

While only one iteration of the whole analysis was required to obtain an optimal solution, we found some benefit from tuning the geometry restraints discussed at the end of the last section. Applying a trivial dummy measurement ($\delta z = 0 \pm 10\mu\text{m}$ etc.) to each alignment correction gives a fully constrained set of alignment equations. This was of value early in the analysis, for example using the Shingle residual fit data only together with the dummy constraints allowed a clean test of the code relevant for the Shingle data only without interference from the other data types. We also found some benefit in artificially reducing the residual fit errors by tuning an extra weight applied to the elements of the coefficient vector $\mathbf{c}$ and the corresponding rows of the design matrix $\mathbf{A}$ for residual types that suffered from relatively low statistics, for example the Pair data in the early run period. Such fine-tuning, applied toward the end of the analysis, did not play a major role.

A further iteration of the whole procedure was only necessary after a significant change in the data sample or an improvement in the algorithm such as the inclusion of the CCD shape correcting terms. The VXD3 alignment project began with the Monte Carlo studies as the first data was being collected in 1996, was refined as the data accumulated, and concluded in May 1999 within a year of the end of the final SLD run. Although some visible effects in the residual data remained, for example due to a slight bow shape across the CCD width observed in a few Doublets, these were all too minor to consider further extensions of the algorithm. The achievement of the design resolution indicated in figure 8 and equation 11 has greatly benefitted the rich program of heavy flavour physics analysis at SLD during the subsequent years.

A Singular Value Decomposition

SVD techniques are based on relatively recent mathematical developments and have mainly been restricted to research papers or advanced books on matrix algebra or applied mathematics (indeed the earliest text-book we have found referring to the technique was published in 1983). However in spite of this because of their power and robustness SVD techniques are becoming better known and have been increasingly used over recent years. The FORTRAN

code we used as well as much of the discussion in this appendix was derived from reference [8], a fuller discussion of the mathematical details can be found in [9] and we have also found [10] useful to obtain insight into the properties of SVD. This technique exposes a rich variety of the mathematical properties of any $m \times n$ matrix $\mathbf{A}$ through the factorization into $\mathbf{USV}^T$:

$$\begin{pmatrix} & \\ & \\ & \\ \mathbf{A} & \\ & \\ & \end{pmatrix} = \begin{pmatrix} & \\ & \\ & \\ \mathbf{U} & \\ & \\ & \end{pmatrix} \begin{pmatrix} s_1 & & \\ & s_2 & \\ & & \ddots \\ & & & s_r \end{pmatrix} \begin{pmatrix} & \\ & \\ & \\ \mathbf{V}^T & \\ & \\ & \end{pmatrix} \quad (12)$$

where $\mathbf{U}$ is an $m \times m$ orthogonal matrix, $\mathbf{S}$ is an $m \times n$ matrix with non-zero entries among the diagonal elements s_i only and $\mathbf{V}$ is an $n \times n$ orthogonal matrix. The elements $s_1, s_2, \dots, s_r$ with $r = \min(m, n)$ are called the ‘singular values’ of matrix $\mathbf{A}$. This decomposition is unique, except for the possible re-ordering of the s_i elements and the corresponding columns of $\mathbf{U}$ and $\mathbf{V}$. (In the convention adopted for our analysis code taken from [8], that is more economical in terms of computing resources, $\mathbf{U}$ is an $m \times n$ column orthonormal matrix and $\mathbf{S}$ is an $n \times n$ diagonal matrix containing the same elements s_i as above. This does not effect the relevant SVD properties.) The singularities of the matrix $\mathbf{A}$ can be diagnosed by inspecting the set of elements s_i , each of which takes either a zero or positive value. An element $s_i = 0$ corresponds to a singularity while a near-singularity is identified as an element s_i that is very small relative to the largest element of the set. By scanning the diagonal elements of $\mathbf{S}$ some degree of judgement is required to identify the ‘small’ elements and choose a cut, the use of which is explained below.

The subset of all m -dimensional vectors $\mathbf{b}$ for which an n -dimensional vector $\mathbf{x}$ exists such that $\mathbf{Ax} = \mathbf{b}$ is called the ‘range’ of $\mathbf{A}$. The columns of $\mathbf{U}$ with a column index i corresponding to the non-zero s_i elements form an orthonormal set of basis vectors that span the range of $\mathbf{A}$. This property is used at the end of this section.

If $\mathbf{A}$ is a non-singular square matrix then all of the diagonal elements of $\mathbf{S}$ are non-zero and it follows that the inverse $\mathbf{A}^{-1}$ is given by $\mathbf{VS}^{-1}\mathbf{U}^T$ using well known properties of matrices (e.g. $s_i^{-1} = 1/s_i$ for the elements of a diagonal matrix and $\mathbf{U}^{-1} = \mathbf{U}^T$ for an orthogonal matrix). In this case the inverse can be used to solve for the unknown values of the vector $\mathbf{x}$ for the set of simultaneous equations $\mathbf{Ax} = \mathbf{b}$. If there are more equations than unknowns then a ‘least length’ solution that minimizes $|\mathbf{Ax} - \mathbf{b}|$ can be obtained, even if there are singularities, by defining a pseudo-inverse of $\mathbf{A}$ (denoted $\mathbf{A}^+$) as $\mathbf{VS}^+\mathbf{U}^T$ with $s_i^+ = 1/s_i$ but setting $s_i^+ = 0$ if $s_i = 0$ or very small, below the chosen cut value (these small s_i elements are also themselves zeroed). This construction of the pseudo-inverse matrix by

setting infinities caused by singularities equal to zero has surprisingly useful consequences. If $\mathbf{A}^+ \mathbf{b}$ is used as the solution of $\mathbf{Ax} = \mathbf{b}$ then singularities (and near singularities) are identified and avoided in this way and if the set of simultaneous equations:

- under-constrain the parameters then a particular solution $\mathbf{x}$ will be obtained with the full solution space being spanned by linear combinations of the basis vectors constructed from the columns of $\mathbf{V}$ with indices i corresponding to the $s_i = 0$ elements (including the small s_i elements that were zeroed).
- just constrain the parameters then the exact and unique solution for $\mathbf{x}$ is determined.
- over-constrain the parameters then a ‘least squares’ solution that minimizes $|\mathbf{Ax} - \mathbf{b}|$ is determined.

This latter case can be demonstrated by assuming that a vector $\mathbf{x}_0$ exists and can be added to the solution $\mathbf{x}$ such that $|\mathbf{A}(\mathbf{x} + \mathbf{x}_0) - \mathbf{b}|$ has the minimum length. Expanding $\mathbf{A}$ and the solution $\mathbf{x} = \mathbf{A}^+ \mathbf{b}$ in terms of their SVD components yields:

$$\begin{aligned}
& |\mathbf{A}(\mathbf{x} + \mathbf{x}_0) - \mathbf{b}| \\
&= |(\mathbf{USV}^T)(\mathbf{VS}^+ \mathbf{U}^T)\mathbf{b} - \mathbf{b} + \mathbf{Ax}_0| \\
&= |(\mathbf{USS}^+ \mathbf{U}^T - \mathbf{I})\mathbf{b} + \mathbf{Ax}_0| \\
&= |\mathbf{U}((\mathbf{SS}^+ - \mathbf{I})\mathbf{U}^T \mathbf{b} + \mathbf{U}^T \mathbf{Ax}_0)| \\
&= |(\mathbf{SS}^+ - \mathbf{I})\mathbf{U}^T \mathbf{b} + \mathbf{U}^T \mathbf{Ax}_0|
\end{aligned}$$

where the last step follows since $\mathbf{U}$ is an orthogonal matrix that simply ‘rotates’ the vector constructed between the large brackets. The diagonal matrix $(\mathbf{SS}^+ - \mathbf{I})$ has non-zero elements only for $s_i = 0$, while $\mathbf{U}^T \mathbf{Ax}_0$ has non-zero components i only for $s_i \neq 0$ since $\mathbf{Ax}_0$ is clearly in the range of $\mathbf{A}$. It follows that the minimum is obtained when all elements of the vector $\mathbf{Ax}_0$ are zero and, since the system is over-constrained, this requires that $\mathbf{x}_0 = \mathbf{0}$. This proves the assertion that the solution $\mathbf{x} = \mathbf{A}^+ \mathbf{b}$ minimizes the length $|\mathbf{Ax} - \mathbf{b}|$. Our adaption of this least squares fit property of the solution with SVD into a full optimal χ^2 minimization is described in section 4.

Acknowledgments

While this paper has described our use of the singular value decomposition matrix technique, the successful VXD3 alignment and performance has been the responsibility of a large num-

ber of SLD Collaborators [2]. In particular, we would like to acknowledge the MIT group for performing the initial optical survey, John Jaros for co-ordinating the VXD3 tracking program, Danning Dong for work on the Monte Carlo study and global alignment, Aaron Chou for offline reconstruction work and for performing the IP position finding, Nik Sinev for correcting electronics hardware problems before the 1997 run and Chris Damerell for discussions on the mechanical and thermal properties of the VXD3 components.

References

- [1] D. J. Jackson, *A Topological Vertex Reconstruction Algorithm for Hadronic Jets*, Nucl. Inst. and Meth. **A388**, 247 (1997).
- [2] K. Abe *et al.*, *Design and Performance of the SLD Vertex Detector, a 307 Mpixel Tracking System*, Nucl. Inst. and Meth. **A400**, 287 (1997).
- [3] The CCDs were manufactured by the EEV Company, Chelmsford, Essex, United Kingdom.
- [4] M.J.Fero *et al.*, *Performance of the SLD Central Drift Chamber*, Nucl. Inst. and Meth. **A367**, 111 (1995).
- [5] F. James and M. Roos, ‘*MINUIT*’ a System for Function Minimization and Analysis of the Parameter Errors and Correlations, Comp. Phys. Comm. **10**, 343 (1975).
- [6] Particle Data Group, *Review of Particle Properties*, Phys. Rev. D**50**, 1273 (1994).
- [7] CERNLIB - The CERN Computer Centre Program Library, IT Division, CERN, Geneva, Switzerland.
- [8] W.H. Press (Editor) *et al.*, *Numerical Recipes in C(FORTRAN)*, Cambridge University Press, ISBN 0-521-43108-5 (1992).
- [9] G.H. Golub and C.F. van Loan, *Matrix Computations*, John Hopkins University Press, ISBN 0-801-85414-8 (1996).
- [10] P.E. Gill, W. Murray, M.H. Wright, *Numerical Linear Algebra and Optimization* Vol. 1, Addison Wesley, ISBN 0-201-12649-4 (1991).