Local Alignment of the BABAR Silicon Vertex Tracking Detector

D. N. Brown [a], A. V. Gritsan [b] Z. J. Guo [b] D. Roberts [c]

[a] *Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA, Dave_Brown@LBL.gov, 1-510-486-7261*

[b] *Johns Hopkins University, Baltimore, Maryland 21218, USA*

[c] *University of Maryland, College Park, Maryland 20742, USA*

Abstract

The *BaBar* Silicon Vertex Tracker (SVT) is a five-layer double-sided silicon detector designed to provide precise measurements of the position and direction of primary tracks, and to fully reconstruct low-momentum tracks produced in e^+e^- collisions at the PEP-II asymmetric collider at Stanford Linear Accelerator Center. This paper describes the design, implementation, performance, and validation of the local alignment procedure used to determine the relative positions and orientations of the 340 SVT wafers. This procedure uses a tuned mix of in-situ experimental data and complementary lab-bench measurements to control systematic distortions. Wafer positions and orientations are determined by minimizing a χ^2 computed using these data for each wafer individually, iterating to account for between-wafer correlations. A correction for aplanar distortions of the silicon wafers is measured and applied. The net effect of residual mis-alignments on relevant physical variables is evaluated in special control samples. The *BaBar* data-sample collected between November 1999 and April 2008 is used in the study of the SVT stability.

Key words:
PACS: 29.40.Gx, 06.60.Sx, 29.40.Wk, 29.85.-c

1 Introduction

Multi-wafer silicon (Si) tracking and vertex detectors have become an essential part of modern High Energy Physics experiments. Because of the short

Email addresses: `Dave_Brown@LBL.gov` (D. N. Brown), `gritsan@pha.jhu.edu` (A. V. Gritsan), `guozj@pha.jhu.edu` (Z. J. Guo), `roberts@umd.edu` (D. Roberts).

ionization drift distances and the sub-μm feature placement accuracy of silicon wafer sensors, individual silicon wafers can provide μm-scale position resolution over areas of a few tens of square centimeters. In order to extend this precision over the square meter areas covered by modern detectors, the relative positions and orientations of the constituent silicon wafers must be well known (1).

Si tracking detector construction techniques define the wafer positions and orientations only nominally. Because Si tracking detectors are typically located in extremely confined spaces near the interaction region, the Si wafer positions cannot be measured using conventional survey techniques once the detector has been installed. Lab-bench measurements during construction using mechanical or optical techniques can determine wafer positions very accurately, but because the wafers can shift due to mechanical and thermal stress during and after detector installation, and because silicon charge collection effects can distort the effective position of a wafer from its geometric value, these measurements are not sufficient. Because of these effects, the wafer positions and orientations must be determined primarily using signals readout from the silicon detectors themselves when traversed by particles in-situ.

This note describes the procedure developed and used for the *BaBar* Silicon Vertex Tracker (SVT) *local alignment*, whereby the positions and orientations of the wafers are determined. Our procedure uses track data recorded during normal *BaBar* running, filtered and prescaled to produce a fixed sample that roughly uniformly illuminates all the wafers, and constrains all the local alignment degrees of freedom in a statistically independent and systematically complete way. Tracks are fit using SVT hits and constrained using a subset of Drift Chamber (DCH) and beam energy information selected to not impose any significant systematic bias on the local alignment. To avoid statistical bias, we select an independent subset of information from each track. We combine track-based information with direct measurements of the relative positions and orientations of Si wafers made during detector construction, resulting in a statistically correct and systematically robust measure of the consistency (χ^2) of a wafer's position and orientation within the detector. We use an iterative technique to determine the relative wafer positions that minimize the $\sum \chi^2$ of all wafers. The resultant local alignment is then validated against several possible systematic effects. Each of these functions are described in detail in the following sections. The related but simpler problem of determining the rigid-body position and orientation of the SVT within the *BaBar* detector (*global alignment*) is not covered in this note.

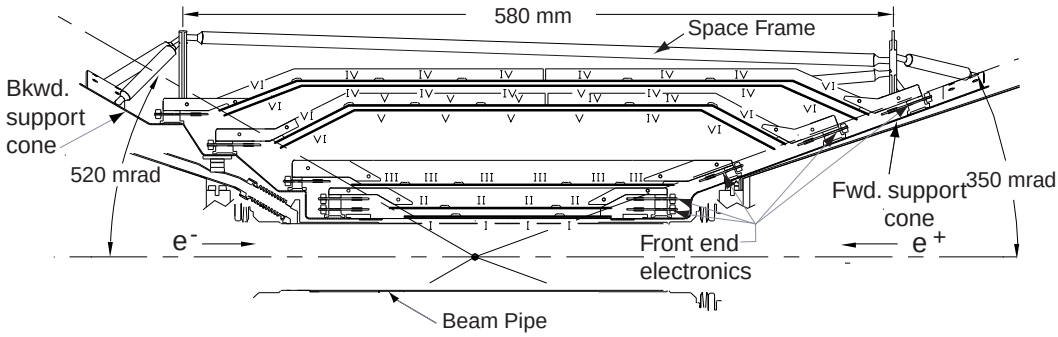

Fig. 1. Longitudinal section of the *BaBar* SVT . The roman numbers label the six different types of wafers.

2 The Silicon Vertex Tracker

The *BaBar* SVT was designed primarily to provide precise reconstruction of charged particle trajectories and decay vertices near the e^+e^- interaction point of PEP-II (2), as required by *BaBar*'s diverse physics goals. Additionally, the SVT provides the precise θ angle measurement needed to perform charged particle identification using *BaBar*'s Cherenkov detector (DIRC) and Electromagnetic Calorimeter (EMC) (3).

The SVT is composed of 340 separate Si wafers, arranged in 5 co-axial roughly cylindrical layers, see Fig. 1. Each layer is composed of between 6 and 18 modules, arranged symmetrically around the cylinder (z) axis, held in place by a rigid carbon-fiber frame. Each module is in turn composed of between 4 and 8 individual Si wafers, which are glued to supporting kevlar ribs extending in the z direction. See Table 1 for the geometrical parameters of each layer. There are six different wafer shapes, including a trapezoidal shape used to form the arch modules discussed below. The smallest wafers are 4.2×4.1 cm^2, and the largest (in layer 3) are 4.4×7.1 cm^2.

The modules of the inner three layers are planar, while the modules in layers 4 and 5 are arch-shaped. This design reduces the amount of material and improves the point resolution for particles originating from the interaction region compared to a planar module. The modules in the inner three layers are tilted by 5° in azimuth (ϕ), allowing an overlap region between adjacent modules, see Fig. 2. This arrangement is advantageous for alignment and provides full ϕ coverage. The outer layers cannot be tilted because of the arch geometry. To have an overlap and avoid coverage gaps in ϕ, the outer two layers are divided into two sub-layers at slightly different radii.

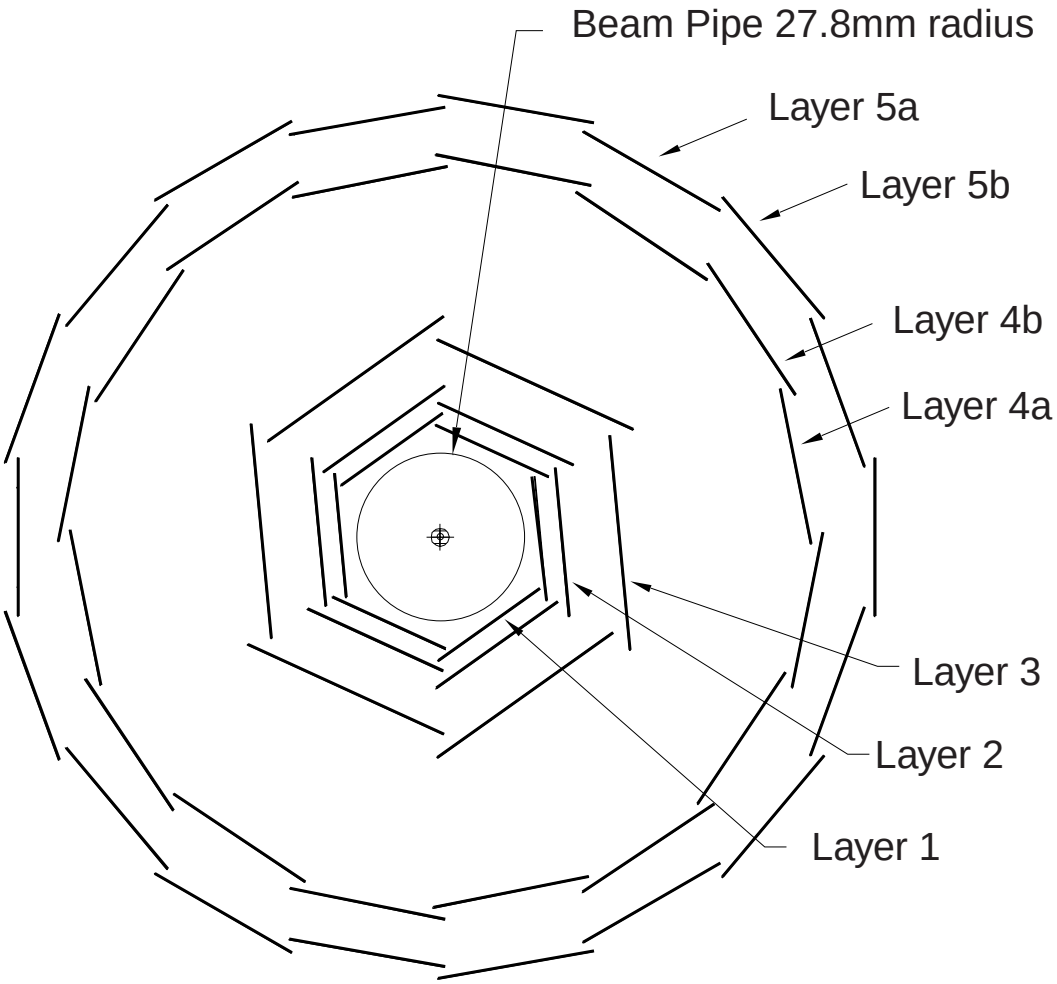

Fig. 2. Transverse section of the *BaBar* SVT .

The SVT support structure is a rigid body made from two carbon-fiber cones, connected by a space frame, also made of carbon-fiber epoxy laminate. While in operation the SVT is mounted on the innermost magnets of the PEP-II beamline, supported by an assembly fixture that allows for some relative motion with respect to PEP-II. Because the SVT is mounted independently of the rest of the *BaBar* detector, movement between the SVT and the rest of the detector can occur. During operation the SVT is cooled to remove the heat generated by the electronics and is kept in a humidity controlled environment.

3 Parameterizing the SVT Local Alignment

To describe the SVT local alignment we consider each Si wafer as an independent rigid body located and oriented near its nominal construction position and orientation. We describe the true position of a wafer by its displacement

and rotation relative to its nominal position. The local alignment of the full SVT thus requires 2040 geometric parameters, which includes 6 redundant global degrees of freedom for the position and orientation of the SVT within *BaBar*. As described in detail in Sec. 7, we also model aplanar distortion of the inner 3 layers of wafers, adding 84 more parameters.

We choose to define the local alignment parameters in the *local* wafer coordinate system, a right-handed Cartesian system with coordinates uvw, defined uniquely for every wafer. The u axis is defined to lie in the nominal plane of the Si wafer pointing in the direction of increasing ϕ. The v axis is defined to lie in the nominal plane of the Si wafer, orthogonal to $\hat{u}$. The w axis is defined as the normal direction to the nominal plane of the Si wafer, pointing roughly outwards from the IP. The local wafer coordinate system origin is defined as the geometric center of the wafer.

The displacement component of the local alignment is given by the vector (δ_u, δ_v, δ_w), which describes the true position of the Si wafer relative to its nominal position in that wafer's nominal local coordinate system. Similarly, we describe the orientation of the wafer as the vector (α_u, α_v and α_w), which defines small right-handed rotations about the $\hat{u}$, $\hat{v}$ and $\hat{w}$ axes (respectively) of the nominal local coordinate system of the given wafer, in units of radians. The *BaBar* detector reconstruction software is written so that this convention of local alignment can be easily and efficiently applied to the reconstructed position of SVT hits (5).

The readout strips on the wafers in the barrel region of the SVT are oriented parallel to the local coordinates. The strips on opposite faces of each wafer are oriented orthogonally to each other, providing 90° stereo coverage. The readout strips on the wedge wafers have a pitch which varies slightly along their length, resulting in strip directions which change slightly with position, but which are still roughly parallel. Hits reconstructed in the SVT using the strips parallel to $\hat{v}$ are referred to as u hits, as that is the dimension they constrain. Roughly speaking, these hits measure the ϕ position of the traversing particle. Similarly, hits reconstructed using strips parallel to $\hat{u}$ are referred to as v hits, and they measure the z position of the traversing particle.

The estimated Lorentz shift in the position of u hits induced by *BaBar*'s solenoid is accounted for in the SVT hit reconstruction. Any difference between this estimate and the actual Lorentz shift is absorbed into the δ_u parameter, however this can introduce systematic errors as described in Sec. 6.

4 Goals and Requirements of the SVT Local Alignment

The goal of the SVT local alignment is to determine the local alignment parameters with sufficient accuracy that the remaining misalignments contribute negligibly to the final uncertainty in the physics quantities extracted using the tracks reconstructed in the SVT. For instance, to observe *CP* violation in $\Upsilon(4S) \to B^0\overline{B}^0$, the SVT must be able to precisely measure the roughly $250\,\mu$m average separation between the B meson decay vertices. A full detector simulation study (8) showed this requires an average resolution of no worse than $10\,\mu$m for u hits and $20\,\mu$m for v hits. To insure that the local alignment does not dilute these measurements, we require the statistical precision of the local alignment contribute no more than 15% in quadrature (or 1% net) to the vertex resolution. This implies knowing δ_u to roughly $1.5\,\mu$m, δ_v to roughly $3\,\mu$m, and α_w to roughly 2 μ radians. To achieve this statistical precision requires a data sample equivalent to 400 typical tracks per Si wafer. This many tracks/wafer is recorded in less than an hour of normal *BaBar* data taking. Thus meeting the required statistical precision is not a challenge. The real challenge of the local alignment procedure is to control the systematic errors to the required level. To understand the issues involved in controlling the systematics we must first examine how the local alignment parameters are constrained by data.

A track passing through the full SVT and originating from the interaction point (IP) will generally generate 2 hits (1 u and 1 v) in each of 5 layers. As a track's trajectory is well-described as a 5-parameter helix (9), a single track will constrain 5 degrees of freedom in the local alignment. However, because tracks scatter as they pass through material, the most statistically powerful local alignment constraints will be on the relative positions of wafers in adjacent layers. Similarly, lab-bench measurements of relative wafer positions are useful only for nearby wafers, as mechanical and thermal stress uncertainties grow quickly with relative distance.

Thus, the track-based and lab-bench measurements used by the local alignment procedure effectively only constrain the relative positions of nearby wafers. Many independent local constraints may of course be added together to completely constrain the local alignment, but that procedure raises the risk of introducing a *global distortion*. An example global distortion which correctly defines the relative position of nearby wafers but distorts the SVT as a whole is shown in Fig. 3. In Table 2 we summarize the main global distortions in a system with cylindrical geometry, such as the *BaBar* SVT. In Fig. 4 we illustrate the effect of four global distortions with a natural scale of $50\,\mu$m on the position of individual wafers.

If uncorrected, global distortions would produce unacceptable systematic bi-

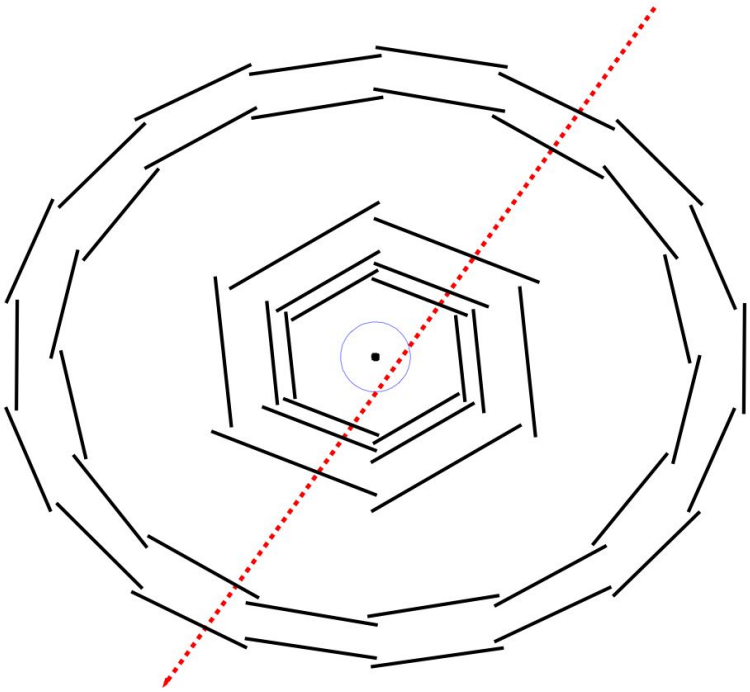

Fig. 3. Graphical illustration of one of global distortions that affect the relative position of nearby wafers only to second order. The *elliptical* effect shown has a greatly exaggerated scale compared to what is used for validation studies.

ases in physics measurements. For instance, an uncorrected *curl* distortion would introduce an artificial charge-dependent momentum asymmetry to reconstructed tracks, given the use of magnetic bending to define the charge and transverse momentum of a track. Likewise, a *radial scale* or *z-scale* distortion would systematically change the measured distance scale of the detector, dis-

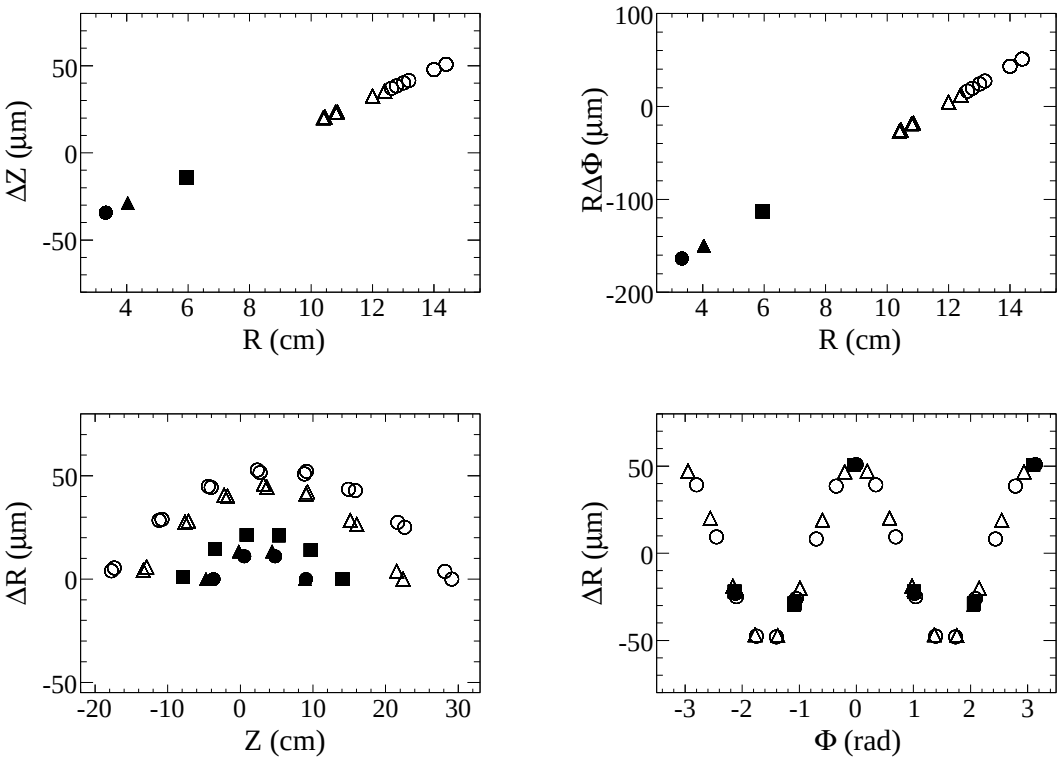

Fig. 4. Four typical global systematic distortions out of nine discussed for a cylindrical system in Table 2: telescope, curl, bowing, and elliptical effects. Each point represents displacement of an SVT wafer shown with filled circles for layer one (•), filled triangles for layer two (▲), filled squares for layer three (■), open triangles for layer four (△), and open circles for layer five (○). The typical scale of distortion is chosen to be 50 μm.

torting lifetime measurements. To estimate specific requirements on how well we must control the different global distortions, we consider the implications of a subset of relevant *BaBar* physics measurements. For instance, to make competitive measurements of the τ lepton and B meson lifetimes, the absolute distance scale must be understood to better than 1 part in 1000. This implies controlling the *radial scale* and *z-scale* distortions to less than 1/1000, or roughly 5 μm over the size of the SVT. Similar arguments can be used to motivate requiring that the local alignment constrain the scale of all of the global distortions listed in Table 2 to better than 5 μm or ten times smaller than the effects shown in Fig. 4. The verification that our local alignment procedure satisfies this requirement is given in Secs. 9 and 10.

A further requirement on the the local alignment procedure is that it be capable of following the time-dependence of actual changes in the detector. We observed some slow relative motion of the SVT wafers related to humidity changes, and due to stress changes during periods of active access to the detector, and during changes to accelerator operation. The timescale for observable changes was about a week, which implies that the local alignment

procedure should function on less than a week's accumulation of *BaBar* data.

It is also important to be able to quickly detect when the local alignment changes, to avoid extensive reprocessing after the initial *BaBar* reconstruction pass. To obtain feedback on possible alignment changes in a timely way requires that computing the alignment constants take no more than 24 hours. Because it's run frequently, the procedure must also be efficient in its use of computer resources.

5 The Local Alignment Data Sample

The data used to perform the SVT local alignment are selected from those collected during normal physics running of the *BaBar* detector. The *BaBar* physics trigger accepts a mix of events including hadronic final states of $e^+e^- \to q\bar{q}$ and $e^+e^- \to \Upsilon(4S)$, $e^+e^- \to l^+l^-$ events, and cosmic rays which pass near the nominal IP. The IP consistency requirement of a few cm was set in the trigger configuration during the first three years of *BaBar* data-taking. It was relaxed for the cosmic tracks with large impact parameters for studies. However, in our final analysis we adopt uniform approach to all data periods and apply the same IP requirements discussed below.

Events which contain tracks useful for the local alignment are identified, and the relevant tracks and hits are saved. These samples are passed to the minimization procedure described in the next section. Details of the data selection are presented in the following subsections, and shown graphically in Fig. 5.

5.1 Event Pre-selection

A pre-selection of events of eventual interest to the local alignment procedure is integrated into the *BaBar prompt reconstruction* procedure that runs shortly after the events are recorded (10). All triggered events are first passed through a minimal background rejection procedure which removes most beam-gas interactions and scales down $e^+e^- \to e^+e^-$ interactions. A pseudo-random prescaling is then applied to $e^+e^- \to e^+e^-$, $e^+e^- \to \mu^+\mu^-$ and cosmic ray triggers, which results in a roughly uniform illumination of the detector.

Pre-selected events are reconstructed using the standard *BaBar* reconstruction program. Tracks are found using standard pattern recognition algorithms in both the SVT and DCH. Tracks found in the DCH (SVT) are extrapolated into the SVT (DCH) respectively, and hits consistent with the original fit are added. Tracks are fit using a Kalman filter algorithm (9) that accounts for

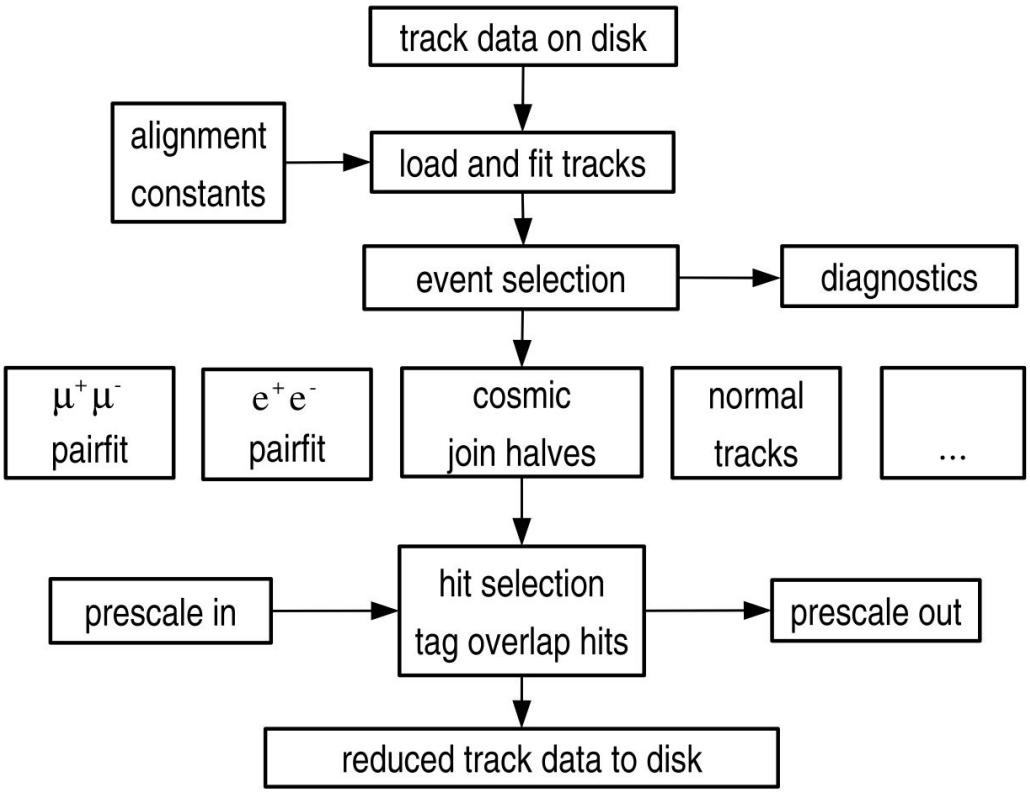

Fig. 5. Diagram of the sequence for the event, track, and hit selection, including calculation of prescale factors. See text for details.

differing hit resolutions, detector material, and magnetic field inhomogeneities.

Prompt reconstruction uses the most recent local alignment parameters available at the time of processing. If, at the end of the local alignment procedure, we observe local alignment change that might have affected the event selection, we repeat the procedure using the updated local alignment parameters. Selection iteration was necessary only when restarting the detector after a shutdown or detector opening. A single iteration was sufficient to select an unbiased sample, because the typical changes are not large. The one case when several iterations were required was the startup of *BaBar* when the starting alignment was from the optical survey.

Pre-selected events are written to a dedicated calibration *stream* (file). The calibration stream persists the events in the standard *BaBar* *mini-dst* event format (4), which records the reconstructed tracks and their associated hits. In particular, this format records the *local* centroid of all SVT hits associated with a track. This allows measuring and applying a new local alignment without first having to remove the effect of the alignment used in prompt reconstruction.

The rarest process used in the local alignment turns out to be cosmic rays,

whose rate is independent of beam luminosity. Based on the cosmic ray trigger rate, and the number of tracks required to satisfy the statistical and systematic constraints, two days of cosmic data is sufficient to meet the alignment goals. We therefore define a local alignment data sample based on a fixed calendar period of around 48 hours.

5.2 Event Categorization and Final Selection

A separate procedure makes a final selection of data useful for local alignment. This reads back the calibration stream, and reconstitutes the SVT and DCH hits using the current local alignment and calibration. From these the full Kalman filter track fit is rebuilt, using the recorded hit assignments. These tracks are used to make a final event categorization and selection.

To insure a reliable momentum measurement, only tracks with at least 10 DCH hits, and at least 2 v and 3 u SVT hits (the minimum to fully constrains all 5 track parameters) are selected for use in the local alignment. To minimize multiple scattering effects, we also require a transverse momentum of at least 1 GeV/c. To cut down on background from secondary (material) interaction products, we accept only tracks whose point of closest approach to the *BaBar* z axis is within 1.5 cm of the IP in the plane transverse to the z axis, and between -7 cm and $+9$ cm of the IP along the z axis.

Events are categorized and finally selected based on the multiplicity and properties of their selected tracks. We define three categories of events in the local alignment; $e^+e^- \to \mu^+\mu^-$ or $e^+e^- \to e^+e^-$ ($e^+e^- \to l^+l^-$) events, cosmic ray events, and 'normal' events. The definitions and selections of these categories is described below.

Events with exactly two selected tracks are tested as potential $e^+e^- \to l^+l^-$ events. Tracks in $e^+e^- \to \mu^+\mu^-$ candidates are required to have associated EMC signals consistent with a minimum-ionizing particle. Tracks in $e^+e^- \to e^+e^-$ candidates are required to have associated EMC energy deposition consistent with the reconstructed track momentum. All $e^+e^- \to l^+l^-$ events are required to have a total energy (computed from track momentum) consistent with the known combined energy of the initial e^+e^- beams, and to be back-to-back in the transverse plane.

Candidate $e^+e^- \to l^+l^-$ events that pass the above cuts are refit using a special form of the *BaBar* Kalman filter track fit which constrains the pair of tracks to have the same four-momentum as the initial e^+e^- system, within the independently-estimated errors on the beam particle momenta. If the χ^2 of the pair fit is larger than 50, it is considered a failed fit, and the individual tracks in these events are passed down to the 'normal' track selection algorithm de-

scribed below. The most common cause of failed pair fits is initial and/or final state radiation. When successful, the pair fit improves the track momentum resolution by more than a factor of 10. More importantly, the pair fit creates a correlated system in which information passes from one track through the IP to the other track. This allows the local alignment procedure to constrain the positions of wafers relative to those on opposite sides of the detector. We can also use $e^+e^- \to l^+l^-$ events to determine the beam momenta parameters simultaneously with the local alignment parameters, without using the independent beam energy estimate. This provides both a systematic check on the alignment procedure, and a precise way to measure the beam boost. This techinque is discussed in Sec. 11.

Because the *BaBar* track finding algorithm assumes all particle originate at or near the IP, a single cosmic ray passing through *BaBar* is initially reconstructed as 2 tracks, splitting the cosmic ray trajectory through the detector roughly in half. Cosmic ray event candidates are selected as having two well-measured oppositely charged tracks which match in angle and position at the their point of closest approach to the IP. These tracks are also required to have associated EMC signals consistent with a minimum-ionizing particle. If these criteria are satisfied, the hits from the upward-going track are added to those on the downward going track, and the combined track is refit.

Events which fail selection as $e^+e^- \to l^+l^-$ or cosmic ray events are taken together with multi-track events as candidates for the 'normal' category. In this category, we select all the tracks which pass the usual selection cuts, plus an isolation cut. Any event with at least one selected track is classified as 'normal'.

5.3 Track Selection and Preparation

Selected tracks from events which pass final event selection in any category are themselves labeled according to their event category. In the case of 'normal' events, tracks are further categorized according to whether or not they have hits in adjacent wafers of the same layer in the *overlap* region (see Fig. 2). As these overlap tracks have a very short extrapolation distance between the same-layer hits, they provide a powerful constraint on the relative position of adjacent modules, and so are especially valuable in the local alignment procedure. Non-overlap tracks in 'normal' events remain categorized as 'normal'.

To balance the impact of the largely-independent global distortion systematic constraints afforded by the different track categories, we perform a final track selection which roughly equalizes the statistical power of the tracks in each category for every individual wafer. Because the wedge wafers (see Fig. 1)

subtend a region of polar angle θ where the rate from $e^+e^- \to \mu^+\mu^-$ and other physics events changes rapidly with θ, we further divide these into two roughly equal parts.

To allow better control of the propagation of systematic misalignment effects from the DCH into the SVT alignment tracks, we refit all tracks using the following technique. First, we split the tracks into two, one with all the SVT hits and one with all the DCH hits. Each of these associated but separate tracks are refit using the standard *BaBar* Kalman filter fit. The parameters and covariance matrix of the DCH-only track fit are sampled at the point where that track enters the SVT detector volume, and these parameters and covariance are then used to *constrain* the SVT-only track fit. Mathematically, the parameter constraint is identical to the effect of having left the DCH hits on the track. However, by *masking* some of the parameters in the constraint, the information content of the DCH-only fit can be filtered. In particular, by masking off all but the ω parameter (inverse curvature) of the DCH-only fit in the constraint, we can greatly improve the momentum resolution of the constrained SVT-only track, without introducing any dependence on possible systematic distortions in the position or orientation of the DCH. We use the DCH-only fit ω constraint when fitting the $e^+e^- \to l^+l^-$, cosmic ray, and overlap category tracks.

Because the sum of the local alignment parameters for all wafers include the 6 *global* degrees of freedom, the local alignment procedure could introduce a global *drift*. Because the ω constraint does not depend on the relative position or orientation of the SVT and DCH it cannot constrain this global drift. To minimize the global alignment drift, we use all 5 DCH-only fit parameters to constrain the fit of the 'normal' tracks. As these tracks have the lowest statistical power, this introduces only a modest and acceptable dependence on DCH alignment distortions.

5.4 Hit Selection

From selected tracks, we select SVT hits which provide information useful for local alignment. First, hits with questionable timing or cluster shape are disabled, and the tracks which held them are refit. Then the remaining hits are filtered to make the sample uniform over the detector, over several track categories, and over the time window in which the data sample was accumulated. Once the alignment procedure is close to convergence, a final outlier removal cut is applied. Details of the hit selection are shown in Fig. 5 and described below.

For each hit on each track in each track category in every wafer we use a

pseudo-random prescaling algorithm to select roughly 100 (200) hits in the outer two (inner three) layers, respectively, for use in the χ^2 minimization. More hits are used for the inner layers to balance the larger number of wafers in the outer layers. The pseudo-random prescale is seeded on the unique event time, and so is fully repeatable but effectively random. Because the data sample has a large number of tracks of each category, only a small number of hits per track are selected, reducing the number of correlated measurements used when computing the alignment χ^2. The unselected good hits are still used for the track fit. Cosmic events have the smallest number of tracks in the sample, and thus have the largest prescale factor, corresponding to using roughly 5 hits per track.

The hit selection is done using three passes over the data, interleaved with the χ^2 minimization described in section 8. In the first pass, the hit prescale constants are determined by dividing the desired number of hits/wafer/category by the number observed. In the second pass, these prescale values are used to pseudo-randomly select hits to be used in the alignment χ^2. Selected hits are persistently tagged so that the same hits are used each iteration of the χ^2 minimization. In this pass we apply very loose requirements on the hit residuals, removing only the very worst outliers, so as to avoid biasing the alignment parameters when the initial alignment is far from optimal. After the χ^2 minimization has partially converged, we repeat the hit selection procedure using the improved alignment parameters, applying a tighter hit residual cut to suppress outlier hits which can distort the χ^2. Selected hits are again tagged and passed on to a final pass of χ^2 minimization.

The hit residuals used in the alignment procedure are computed as the distance of closest approach between the given hit, defined as a line in space, and the track trajectory, defined as a piecewise helix in space, after removing the effect of that hit from the track fit. The residual is signed by the cross-product between the track direction and a nominal hit direction. The error on this unbiased residual is computed as the square-root of the quadratic sum of the projections of the hit error and the track covariance matrix onto the residual measurement. The hit error is estimated as a function of the hit's pulse-height and width, and the direction of the track. The hit position error functional form and parameters were tuned using *BaBar* data.

The hit residuals and their estimated errors are combined to compute a *hit* χ^2 for a particular wafer:

$$\chi^2_h \equiv \sum_i^{\text{hits}} \epsilon_i^T \mathbf{V}_i^{-1} \epsilon_i, \tag{1}$$

where ϵ_i is the residual for hit i. This χ^2 is relative to the *test* set of alignment parameters used when the track was fit.

6 The Optical Survey Alignment

The components of the SVT were optically surveyed to determine their relative positions at several stages during construction. First, individual modules were surveyed on the lab bench during their construction. Then, each layer was surveyed after its modules were mounted on the support structure, starting with the innermost layer and going out. By studying their reproducibility, these surveys were estimated to have a precision of roughly 5 μm in the wafer plane and 20 μm out of the plane.

By averaging and combining the raw survey measurements, they were converted into a *survey alignment*, describing the relative positions and orientations of all the wafers in the SVT . The survey alignment was used as the initial condition when the local alignment procedure described in this paper was first performed. The survey alignment is also used as an additional constraint on relative wafer positions in the local alignment procedure itself, as described in detail below.

A comparison of the wafer positions determined using the procedure discussed in this paper with the optical survey alignment is shown in Fig. 6. The differences are shown after removing the overall global shift and rotation between the two alignment descriptions. These figures, plus tests we made in the early days of *BaBar* demonstrate that the survey alignment by itself does not meet the local alignment requirements defined in Sec. 4. This is understandable, as mechanical stresses and other operational effects will alter the relative wafer positions of the installed detector compared to lab-bench measurements.

The survey alignment does however contain useful information. This is demonstrated by comparing Fig. 6 with Fig. 7, which itself compares the measured wafer position with the SVT nominal geometry. Clearly the optical survey is a better approximation to the true local alignment than the nominal geometry. We also expect some aspects of the survey alignment to remain accurate even in the installed detector. In particular, the relative positions of adjacent wafers in a module should be well described by the survey alignment, as there is little room for stress-induced motion between them that would not destroy the module. Information on the relative position of wafers within a module is orthogonal to that provided by hit residuals, which relate the relative positions of wafers in different modules. It is therefore desirable to add the survey alignment information to the local alignment procedure. However, only the reliable parts of the survey alignment should be used, and the survey information must be appropriately combined with the hit residuals information.

We add survey information to the local alignment procedure by constructing a *survey residual* for each wafer. The survey residual compares the wafer's posi-

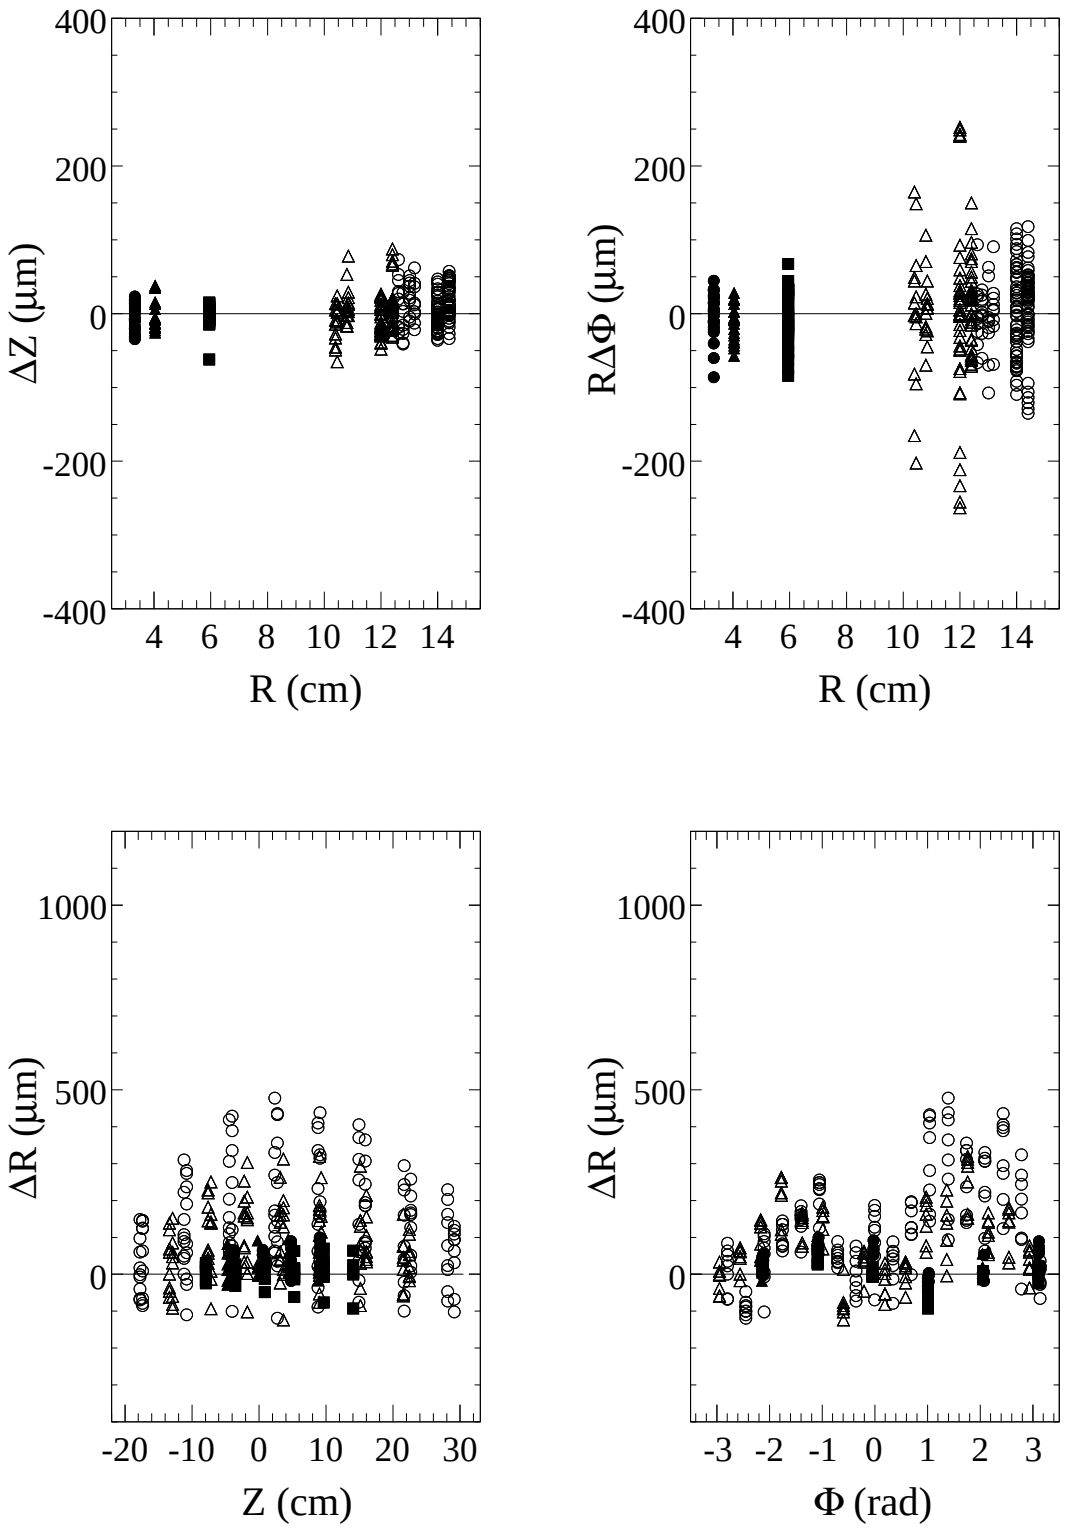

Fig. 6. Differences in the SVT wafer positions between those measured by the procedure described in this paper and the optical survey alignment, projected in the dimensions illustrated in Fig. 4. Each point represents a single wafer, with filled circles for layer one (●), filled triangles for layer two (▲), filled squares for layer three (■), open triangles for layer four (△), and open circles for layer five (○).

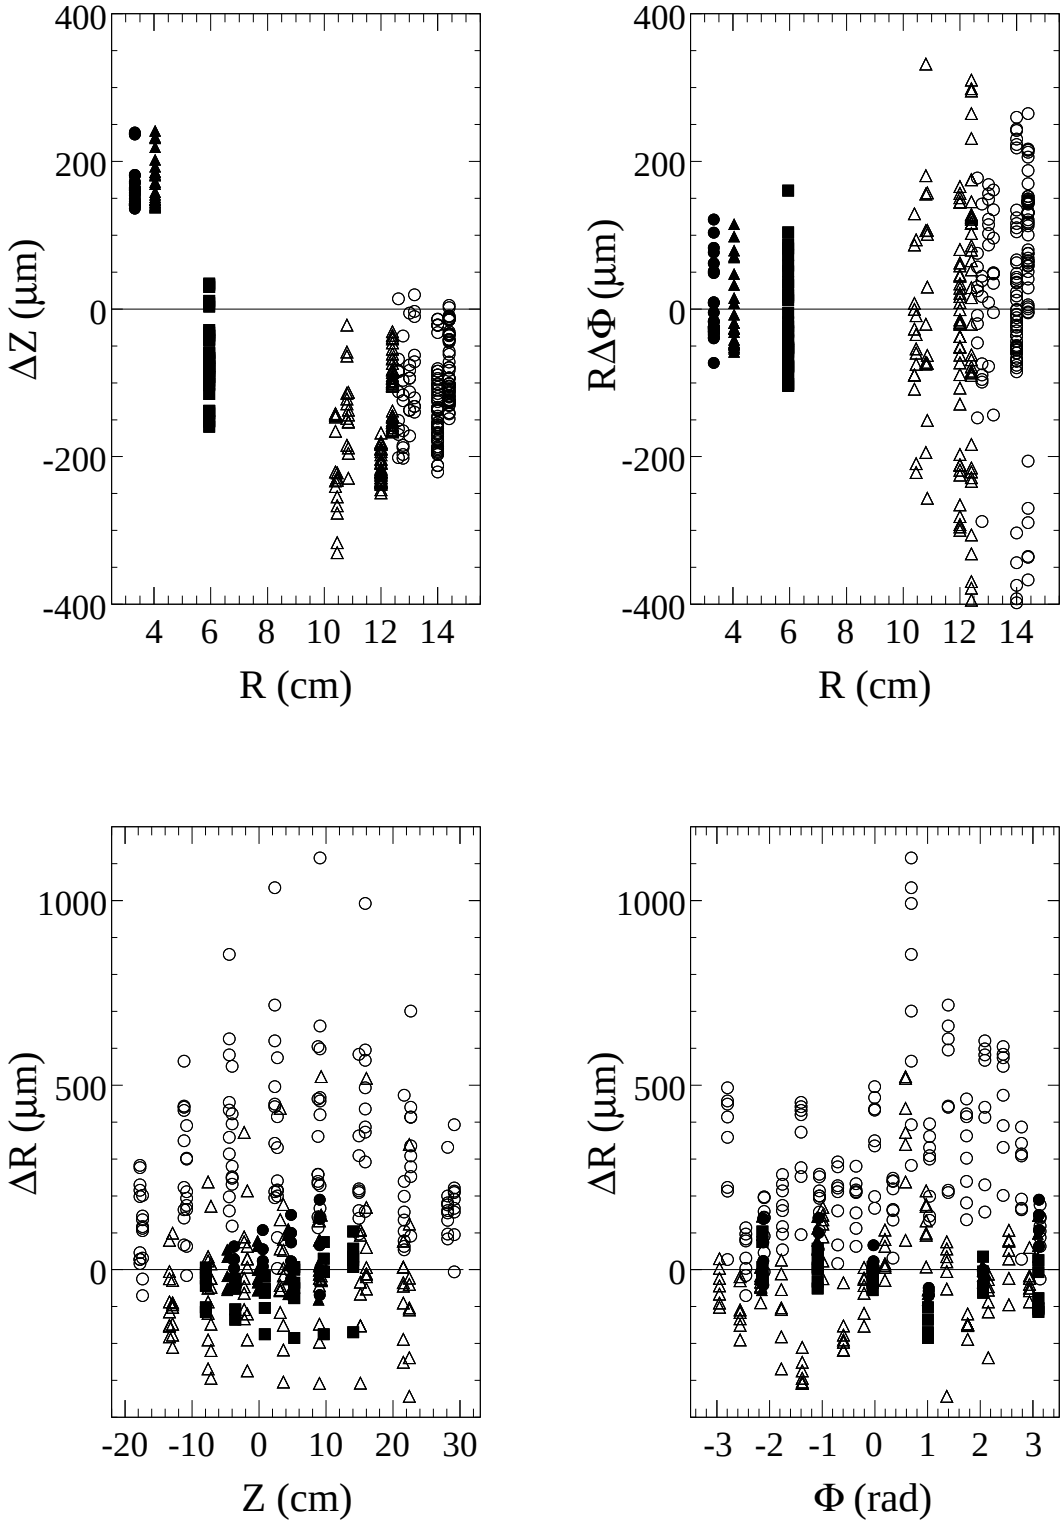

Fig. 7. Differences in the SVT wafer positions between those measured by the procedure described in this paper and the nominal geometry, projected in the dimensions illustrated in Fig. 4. The symbols used are defined in the caption of Fig. 6.

tion relative to the other wafers in its same module as predicted by the survey alignment with the relative position given by the *test* alignment, analogous to the hit residual definition given in Sec. 5.4. We assign an error to this residual based on the survey alignment precision, corrected for systematic effects which decrease the accuracy. We then construct a survey χ^2 contribution from this residual and error, which is combined with the hits χ^2 to form a total local alignment χ^2.

The optical survey wafer residual is computed using the survey measurements in the module containing that wafer. We represent the N wafers of a given module ($i = 1, .., N$, e.g. $N = 8$ in layer five) spatially by a set of 9 points ($j = 1, .., 9$) lying in the wafer plane, at fixed positions in the wafer coordinate system, located roughly at the ends and midpoints of the sides of the rectangle defined by the wafer's active area.

The wafer points in a module can be translated into the global BaBar coordinate system using either the survey alignment, or the test alignment. We use the test alignment to describe the nominal global coordinate positions of these points $\vec{r}_{ij}$. We describe the difference between the survey alignment and the test alignment transformation of these points into the global coordinate system as $d\vec{r}_{ij}$.

We use the difference vectors to solve for the translation vector $\vec{R}$ and rotation vector $\vec{\Omega}$ which minimize the total distance between the points in global coordinates as described by the two alignments (test and survey). Together these represents the optimal transformation between the two alignments. The optimal translation $\vec{R}$ is given by the average of the vector differences between the measurements

$$R_l = \left(\sum_{j,i}^{n \times N} d\vec{r}_{ij} \right)_l \Big/ \left(\sum_{j,i}^{n \times N} 1 \right) , \tag{2}$$

where $l = 1, 2, 3$, representing the spatial coordinates. The optimal rotation $\vec{\Omega}$ is defined implicitly by the equation

$$\sum_{k=1}^{3} \Omega_k \sum_{j,i}^{n \times N} \left(\delta_{kl} (\vec{r}_{ij})^2 - (\vec{r}_{ij})_k (\vec{r}_{ij})_l \right) = \sum_{j,i}^{n \times N} \left(\vec{r}_{ij} \times d\vec{r}_{ij} \right)_l , \tag{3}$$

where k, l are spatial coordinates and δ_{kl} is the Kronecker delta. Solving for $\vec{\Omega}$ requires inverting a 3×3 matrix. The components of the rotation vector $\vec{\Omega}_k$ represent right-handed rotations about the respective coordinate axis. These

rotations are calculated with respect to the average nominal global position $\sum \vec{r}_{ij}$, assuming a small angle approximation (or equivalently small $d\vec{r}_{ij}$). A small number of iterations of this procedure was found to be sufficient to solve for the optimal transformation even when the rotation angles are large. This method of describing and solving for coordinate transforms is based on the formalism of rigid body rotation. It is equivalent to minimizing a χ^2 constructed from $d\vec{r}_{ij}$, if we assign an equal error to all dimensions of all points.

To compute the survey alignment residual, we consider each wafer in the module in turn, referred to as the *wafer under consideration*, with index $i = I$. For that wafer, we first calculate the transformation that relates the overall module position predicted by the survey alignment with that predicted by the test alignment. To avoid direct bias, we exclude the wafer under consideration when computing this transformation, requiring $i \neq I$ in Eqs. (2) and (3). To reduce the impact of potential module deformation occurring after survey, we use only those wafer points j, i which are less than 15 cm from the center of the wafer under consideration when computing the sums. This cutoff was found to be sufficient to remove systematic bias due to module deformation, while still providing enough points to give a statistically meaningful constraint.

We apply the module-level transformation to all the points of the survey alignment, effectively overlaying the survey alignment on the test alignment for this module's position. We then apply Eqs. (2) and (3) to the wafer under consideration, taking the transformed survey points and the test alignment to compute the $d\vec{r}_{ij}$, and using only the points on the wafer under consideration by requiring $i = I$. The resulting $\Delta\vec{R}_I$ and $\Delta\vec{\Omega}_I$ effectively define a 6-dimensional *survey residual*, representing the difference between the position and orientation of the wafer under consideration described by the test alignment versus the survey alignment, relative to the rest of the module in the region around the wafer under consideration.

We define a 6×6 survey covariance matrix $\mathbf{V_s}$ to represent the intrinsic error in the survey residual. We use the same covariance matrix for all wafers. We approximate the survey covariance matrix to be diagonal in the local wafer coordinates, and set the elements according to the values found in the survey consistency tests discussed above. Because the module design restricts relative wafer motion in the module plane, we take these values literally for the in-plane errors (translation in u and v, and rotation about w). To account for potential bowing, twisting, or other aplanar distortions only weakly constrained by the module design, we increase by a factor of ten the estimated errors on the remaining degrees of freedom (translation in w and rotations about u and v).

We construct a survey χ^2 from the 6-dimensional survey residual $\epsilon_s \equiv (\Delta\vec{R}_I, \Delta\vec{\Omega}_I)$ and the survey covariance matrix as

$$\chi_s^2 \equiv \epsilon_s^T \mathbf{V}_s^{-1} \epsilon_s. \tag{4}$$

The use of χ_s^2 in the alignment procedure is discussed in detail in Sec. 8.

The use of the optical survey information in the local alignment procedure assumes implicitly that the positions measured optically on the surface of the wafer correspond to the hit positions reconstructed in the data. In particular, the survey constraint could introduce a bias into the alignment procedure if the Lorentz shift is different for different wafers in a module. We have not studied this effect in *BaBar*, but we estimate it to be less than the estimated survey alignment errors, given the similarity of wafers in a module. Similarly, variations in the thickness of the wafer, which can change the effective charge integration depth, are accommodated by the large out-of-plane errors we assign to V_s.

7 Wafer Curvature

Initial tests of the local alignment algorithm with *BaBar* data showed a smooth but substantial variation of residuals as a function of the local (u, v) hit position in some wafers, especially those in the inner layers. These variations were visible even after the local alignment procedure had converged. These effects were not seen in Monte Carlo simulation of the *BaBar* data or the alignment procedure.

An example of these effects is given in Fig. 8. This plots a projection of the u and v hit residuals (ϵ_u and ϵ_v) from a large sample of high momentum *BaBar* tracks, as a function of the u position of the track. The figures show the average value of the projection, which is defined so as to effectively interpret the residuals as a local deviation in the w position of the wafer:

$$\delta w_u \equiv \epsilon_u / \sin(\theta_{uw})$$
$$\delta w_v \equiv \epsilon_v / \sin(\theta_{vw}),$$

where θ_{uw} (θ_{vw}) is the angle between the track direction and the wafer normal in the uw (vw) plane, respectively. The fit to a parabola is reasonably consistent with the data, given that the errors used are statistical only. The points at the edges are excluded from the fit as they are biased by hits in the overlapping wafers in the same layer. The large uncertainty and fluctuations in the average u hit residuals near $u = 1.2$ cm occurs because tracks from the IP incident at that point are nearly normal to the wafer in the uw plane, and so have very large error in δw.

An incorrect local alignment would result in a linear dependence of δw on

u, with an offset being an incorrect translation and a non-zero slope being an incorrect rotation around the v direction. The clear non-linear dependence shown in Fig. 8 indicates instead an *aplanar* wafer distortion, not described by the standard six local alignment parameters. This geometric interpretation is supported by the fact that compatible effects are seen using either u or v residuals. We interpret Fig. 8 to say that this wafer is bowed in the uw plane, with a *sagitta* of roughly -40 μm, or 15% of the wafer thickness. Bowing in the u direction is possible, as the support ribs constrain against bowing only in the v direction. The large χ^2/NDOF (number of degrees of freedom) of the parabolic fits indicates that simple bowing may not be the only aplanar distortion present, as is discussed further in Sec. 12.

We see evidence for bowing in all SVT wafers. The observed bowing is roughly proportional to the u size of the wafers, with the largest effect in layer three. While no particular factor has been identified which causes wafer bowing, the SVT detector builders agree that bowing at the observed scale is possible (11).

The observed wafer curvature produced a systematic bias on the transverse impact parameter as a function of azimuth for high-momentum tracks. These biases were beyond the tolerance of the local alignment requirements. Furthermore, because the wafer bowing was the same direction for all of layer three, it caused an effective bias in the average w positions of roughly 30 μm, well beyond the goal of $< 5\,\mu$m radial distortion given in Sec. 4. Thus we determined that the aplanar wafer distortions must be measured and corrected for the local alignment procedure to meet its requirements.

We model the aplanar distortions as a quadratic dependence of the wafer w displacement δw on the u position of the measurement,

$$\delta w(u) = (u^2 - u_0^2)/2R, \tag{5}$$

where R is the curvature radius of the wafer, related to the sagitta $\mathcal{S}$ by $1/R = 2\mathcal{S}/L^2$, L being the u half-width of the wafer, and u_0 is a convenience parameter set to $u_0 = L/\sqrt{3}$ in order to keep the w center of gravity of the wafer independent of the curvature radius R. We do not model a first-order term as that is redundant with the α_u alignment parameter. We measure R for each inner-layer wafer by fitting the average u and v residuals dependence on u according to this model, minimizing a residual-based χ^2 to find the best R value. We do not fit for curvature in the outer layers because their smaller u size makes the effect of their curvature negligible, and other aplanar distortions are found to dominate, as discussed further in Sec. 12.

Figure 9 shows the sagitta values obtained from fits to the different wafers in layer three, as a function of the wafer positions in global z. The curvature values are generally smallest near the ends of the module, where they are fixed

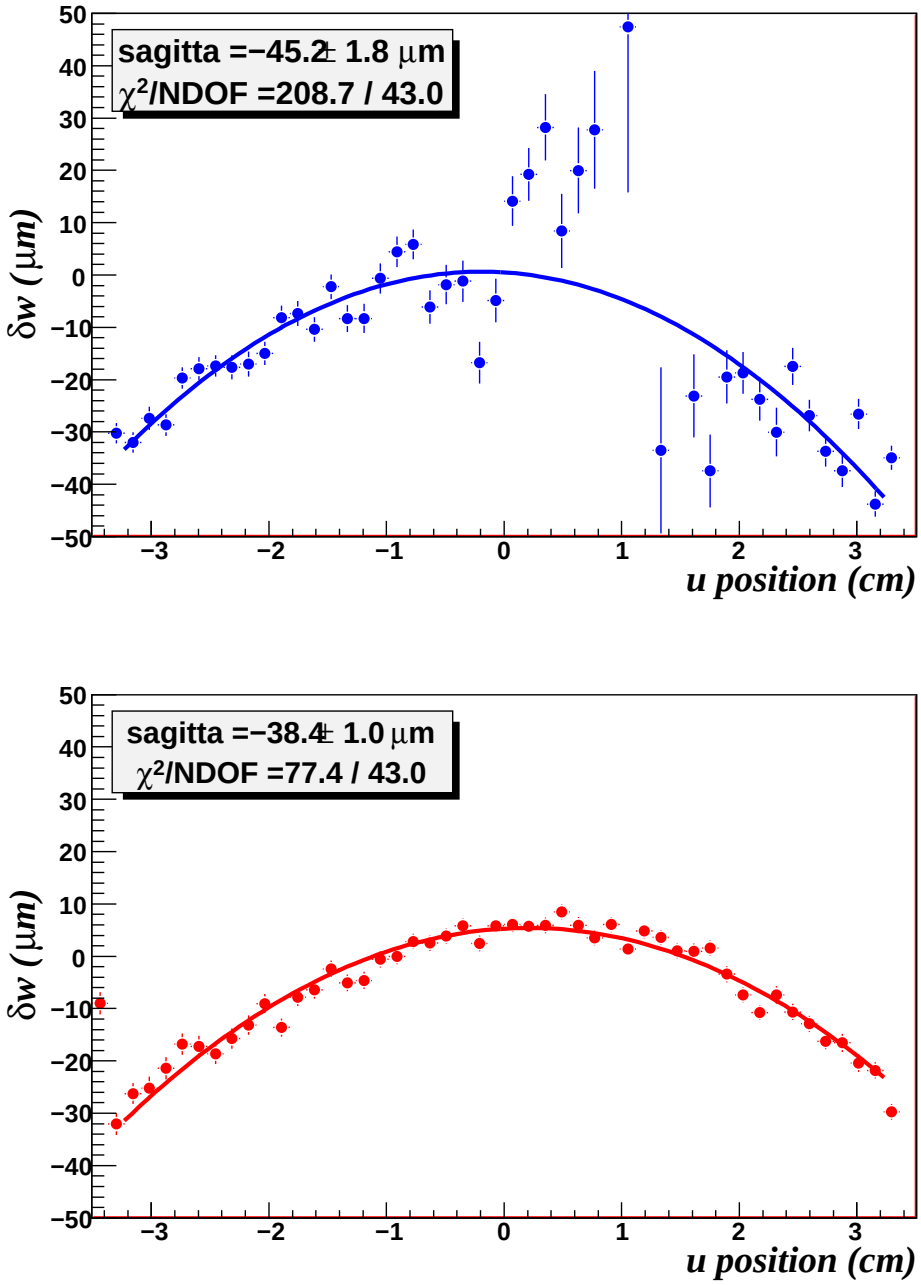

Fig. 8. Average δw projection of the u (top) and v (bottom) hit residuals in a layer 3 wafer as a function of the local u position for high-momentum tracks in the *BaBar* data.

to rigid hybrids, and are largest in the middle of the detector, where there is less mechanical constraint, consistent with expectations.

Because the curvature parameter measurement depends on residuals, it is sen-

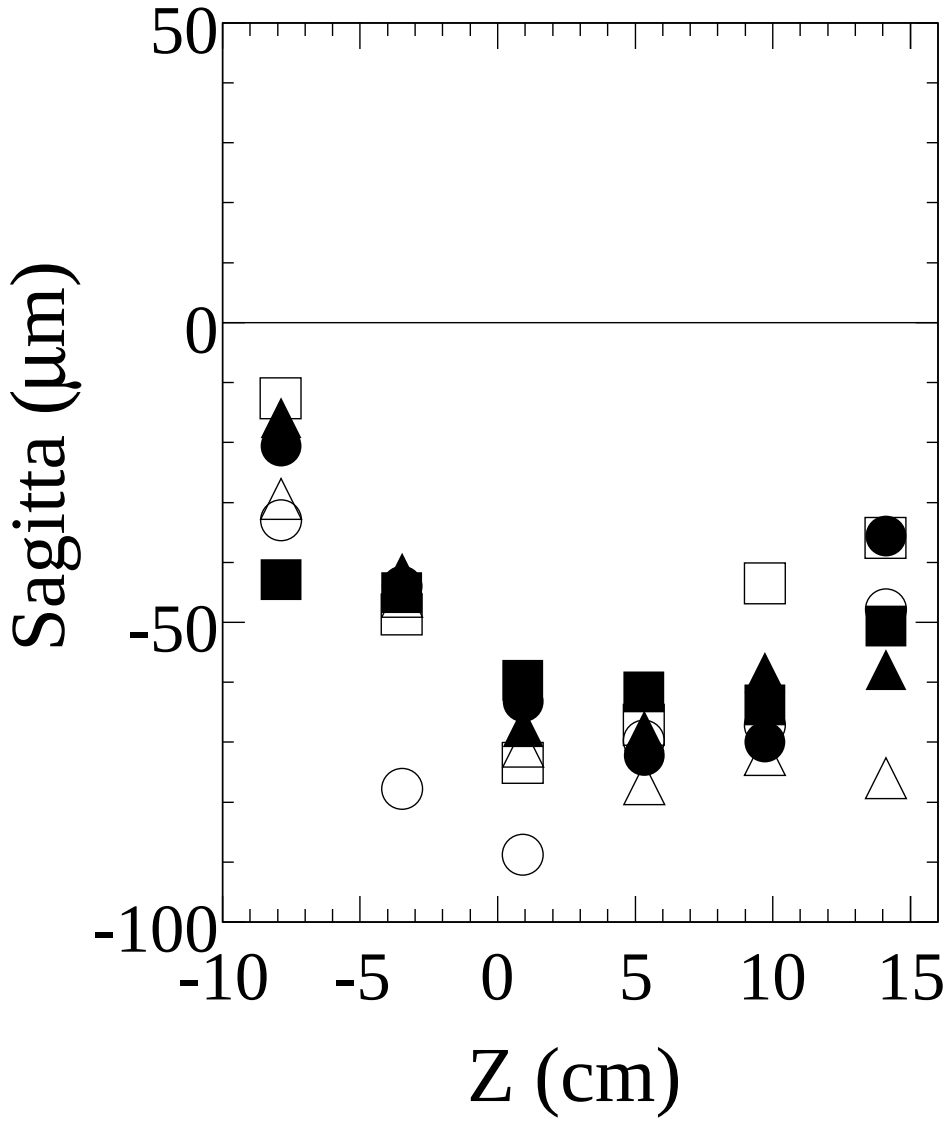

Fig. 9. Results of the fits for the wafer sagitta in layer three. Symbols represent different modules starting from $\phi = 0$ in the order of increasing azimuth: •, ■, ▲, △, □, ∘.

sitive to the local alignment. Similarly, the local alignment procedure depends on what value of curvature we assign to the wafers. This correlation forces us to fit simultaneously for both the curvature parameters and the local alignment parameters. The organization of the simultaneous fit for the local alignment and wafer curvature is discussed in Sec. 8. We do not observe any time dependence to the curvature parameters, so those parameters are normally held fixed when fitting for the local alignment.

To avoid biasing the local alignment procedure or *BaBar* physics, we must correct for wafer curvature during track reconstruction. We correct the u hits for curvature by displacing them in w by the δw amount predicted by Eq. (5), given the measured u position of the hit. We correct the v hits by modeling them with as a three-piece piecewise linear trajectory, where the endpoints of the three equal-length linear segments are chosen to lie at the δw positions described by Eq. (5), given the u coordinates of those endpoints. This trajectory is used when computing the track-hit residual, thereby naturally correcting for δw.

8 Local Alignment Minimization Procedure

To obtain the best estimate of the true alignment parameters, the local alignment procedure combines all the available information into a total χ^2. We extract the optimal local alignment parameters by minimizing this χ^2 as a function of the local alignment parameters. Conceptually, this requires computing the dependance of every residual on each of the 2040 alignment parameters (six parameters for each wafer), and then minimizing the total χ^2 in this 2040-dimensional space. However, each residual depends primarily on the alignment parameters of its hit's wafer. Additionally, minimizing such a large number of dimensions is computationally challenging, raising issues of performance and accuracy. We therefore choose to simplify the local alignment minimization procedure by dividing the total χ^2 into 340 separate wafer χ^2 functions, and minimizing each independently for that wafers alignment parameters. We then iterate to account for the secondary dependence of a residual on some other wafers alignment parameters (wafer correlation), and stop iterating when the alignment parameters for all wafers stabilize.

While our iterative procedure is less direct than a simultaneous minimization of all parameters, we feel it offers numerous advantages over that technique. For one, standard algorithms can be used to efficiently invert the 6×6 matrices involved. Likewise, computing the derivatives of residuals with respect to a single wafer's alignment parameters is straightforward and fast. Because of this, the derivatives can be recalculated between iterations, naturally accounting for second-order effects which are generally ignored in a simultaneous solution. Importantly, our iterative procedure provides access to intermediate states of the alignment, allowing us to monitor the convergence process directly. This gives us confidence in the final result, and allows us to test the sensitivity of the procedure to physical or computational effects, as is described in Sec. 9. Additionally, our iterative procedure allows us to incorporate wafer curvature (Sec. 7) and beam boost (Sec. 11) determination during the alignment minimization, thus correctly handling correlations between these parameters and the alignment parameters. Finally, organizing the minimization by wafer

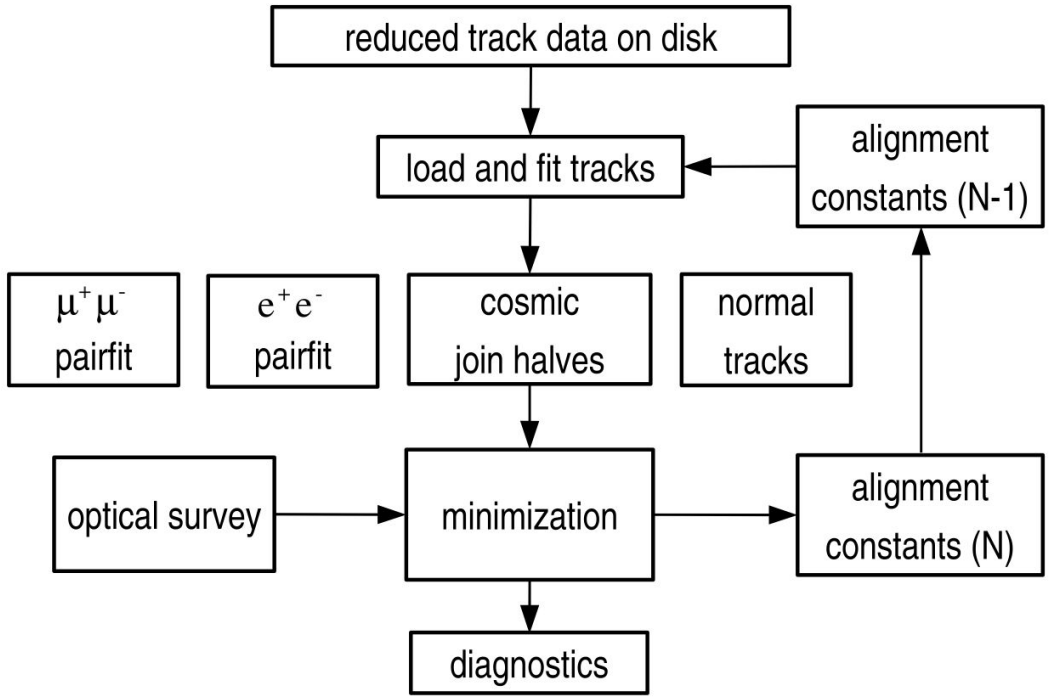

Fig. 10. Diagram of minimization sequence. Iterations over index N is performed until diagnostics shows convergence.

naturally allows for a modular software design well suited for modern Object Oriented programming languages.

The local alignment minimization sequence is shown schematically in Fig. 10. We start with the tracks and hits that were selected as described in previous sections. The tracks are fit using the current (test) estimate of the alignment parameters, from which the hit χ^2 are computed for the selected hits. To this is added the survey alignment χ^2, also relative to the test alignment parameters, to form a wafer χ^2. We then minimize each wafer's χ^2 with respect to the *change* in that wafers local alignment parameters, holding the parameters of every other wafer fixed. After minimizing every wafer's χ^2, we update the alignment parameters for all 340 wafers by adding the computed parameter change to the original alignment parameters estimate. We then use that updated local alignment to fit the tracks and evaluate the survey information in the next iteration, and repeat the process until the alignment converges. After convergence, the alignment parameters are stored in the *BaBar* conditions database (6).

The wafer χ^2 used in the alignment minimization is defined as the sum of the hit and survey χ^2 defined in Sec. 5.4 and Sec. 6 respectively. We express that explicitly in terms of a small change in the wafer's alignment parameters $\mathbf{\Delta p}$ with respect to the test alignment used in computing the hit and survey residuals:

$$\chi^2 \equiv \sum_i^{\text{hits}} \epsilon_i^T(\Delta\mathbf{p})\mathbf{V}_i^{-1}\epsilon_i(\Delta\mathbf{p}) + \epsilon_s^T(\Delta\mathbf{p})\mathbf{V}_s^{-1}\epsilon_s(\Delta\mathbf{p}). \tag{6}$$

We analytically minimize χ^2 by finding the value of $\Delta\mathbf{p}$ which gives a null first derivative with respect to all components of $\Delta\mathbf{p}$, evaluated to first order in $\Delta\mathbf{p}$. This requires computing the first derivatives of the hit and survey residuals with respect to $\Delta\mathbf{p}$. The hit residual derivatives are calculated analytically, given the hit direction and the direction of the track. The derivative formulas are presented in Appendix A expressed as the Jacobian $\mathbf{J}_k$ matrix:

$$\mathbf{J}_k = \partial\epsilon_k/\partial(\Delta\mathbf{p}). \tag{7}$$

Because the survey alignment residual is defined directly in terms of the alignment parameters themselves, its Jacobian is simply the identity 6×6 matrix. Wafers which are electronically dead or which have dead readout views naturally have their missing parameters constrained by their survey alignment.

The matrix equation which results from setting the derivative of the wafer χ^2 to zero can be inverted to solve for the change in the six alignment parameters $\Delta\mathbf{p}$:

$$\Delta\mathbf{p} = \left[\sum_j^{\text{all}} \mathbf{J}_j^T\mathbf{V}_j^{-1}\mathbf{J}_j\right]^{-1} \left[\sum_k^{\text{all}} \mathbf{J}_k^T\mathbf{V}_k^{-1}\epsilon_k\right]. \tag{8}$$

The matrix sums are computed by iterating over all the hit residuals and the survey residuals for the wafer. The new wafer alignment parameters are taken to be $\mathbf{p} = \mathbf{p}_0 + \Delta\mathbf{p}$, where $\mathbf{p}_0$ are the test alignment parameters. Upon the subsequent alignment iteration, these updated parameters become the test parameters, and the procedure is repeated. The χ^2 value used in convergence testing is computed directly from Eq. (6), and thus is one iteration behind the parameter computation.

A wafer is said to have converged when its total χ^2 changes by less than a given threshold between iterations, typically set to 0.01 absolute. While this value may seem small for a χ^2 that typically has a few hundred degrees of freedom, we found that a low threshold was necessary for the procedure to be sensitive to small global distortions, as discussed in Sec. 9. The entire local alignment procedure is said to converge when all but at most two wafers are converged. This allows for a trivial oscillation observed between dead or partially-dead wafers constrained only by survey information. The local alignment procedure typically converges after roughly 100 iterations.

As discussed in Sec. 7, the wafer curvature parameters must be fit simultaneously with the local alignment parameters. This is done in a dedicated variant of the local alignment procedure, where we introduce a fit for the curvature parameters between each normal local alignment parameter iteration, holding the alignment parameters fixed. The updated curvature parameters are used in the subsequent alignment minimization iteration. A similar procedure is applied when fitting for the beam boost, as described in Sec. 11.

Because both the derivative calculations and the χ^2 minimization are analytically computed, the minimization procedure is reasonably fast. The processing time for a single alignment iteration is limited by the time it takes to refit the tracks. This time is much reduced compared to the normal *BaBar* reconstruction, as the DCH information is applied as a single constraint instead of 40 separate hits. A single iteration of the alignment minimization on a standard alignment set, without the curvature fit, takes roughly twenty minutes on a modern multi-GHz intel-processor based Linux computer. The entire local alignment procedure typically converges in roughly twelve hours.

The local alignment minimization procedure is written within the *BaBar* software framework, using standard access to event and conditions data. Iterations are controlled using TCL/Tk (7), through either an interactive GUI or with a TCL script submitted in batch. Bookkeeping, diagnostics, input/output, and job management are also controlled through TCL/Tk.

9 Validation of the Alignment

We validate the performance of the SVT local alignment procedure through self-consistency tests using *BaBar* data, where we compare the results of the aligned detector with apriori expectations. We also compare the performance of the aligned *BaBar* detector with perfectly-aligned Monte Carlo simulation. The following tests were performed using data from a typical period of *BaBar* running, where the SVT alignment parameters were determined according to the procedure described above, applying all the calibrations and corrections as normally done when reconstructing *BaBar* data. The validation data sample has minimal overlap with the data used for alignment production, due to the pre-scaling used in the alignment procedure.

We use the track residuals from SVT hits to perform a basic test of alignment self-consistency. The residuals themselves are shown in the plots on the left of Fig. 11, with the u residuals on top and the v on the bottom. These show residuals from the inner three layers of the SVT, using tracks in $e^+e^- \rightarrow \mu^+\mu^-$events selected to be within 0.2 radians of vertical in the readout-view projection. This selection of hits give the best resolution in the SVT, and so

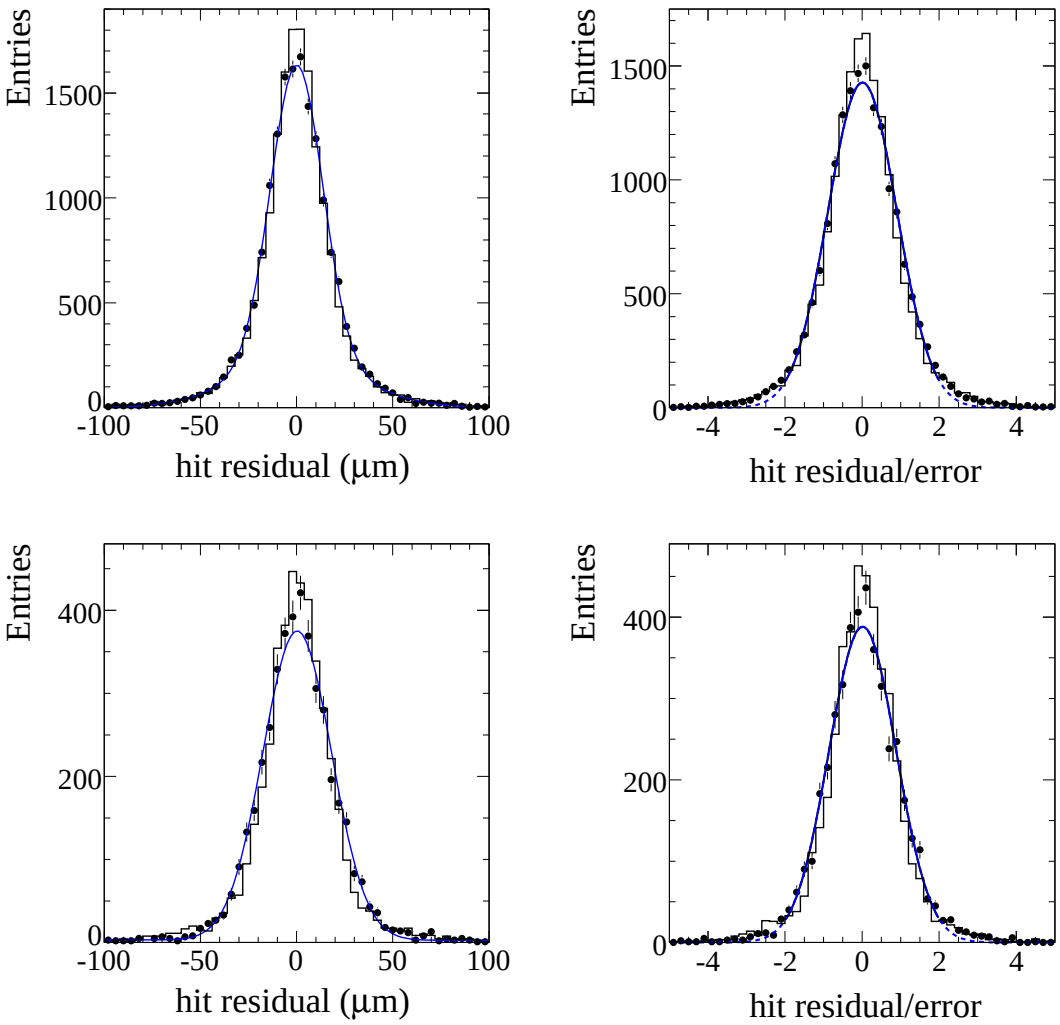

Fig. 11. Single hit residuals (left) and normalized residuals (right) in the u (top) and v (bottom) readout view. The *BaBar* data are shown as points, Monte Carlo simulation as histograms. The (blue) smooth curves are the results of a Gaussian fit to the data.

provides the most sensitivity to misalignment. The residual distributions are centered at zero, and the data and Monte Carlo distributions show very similar shapes. Both data and Monte Carlo distributions show some non-Gaussian tails, as expected from scattering tails and hit resolution variation. The RMS of the data and Monte Carlo distributions agree within a few percent. A Gaussian fit to the core of the distributions gives mean values consistent with zero, and sigma values of 14 (13) μm for data (Monte Carlo) u hits, and 18 (16) μm for data (Monte Carlo) v hits, respectively. The data-Monte Carlo width differences are consistent with the predicted accuracy of the simulation of SVT hit resolution.

The normalized residual (pull) distributions for the same tracks and hits are shown in the plots on the right of Fig. 11, with the u residuals on top and the v

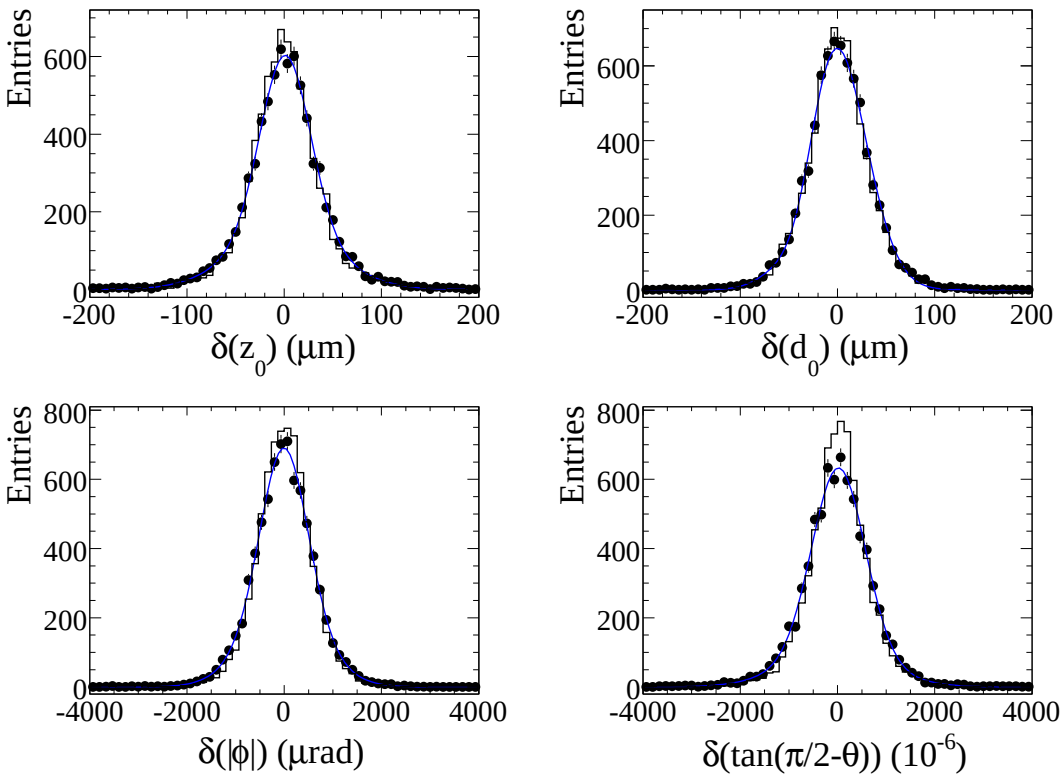

Fig. 12. Difference between the impact parameter along the beam axis (z_0), transverse to the beam axis (d_0), in azimuthal angle (ϕ), and in polar angle ($\tan(\pi/2-\theta)$) for split cosmic tracks. The *BaBar* data are shown as points, Monte Carlo simulation as histograms. The (blue) smooth curves are the results of a Gaussian fit to the data.

on the bottom. Again we see good agreement between data and Monte Carlo. The core of the data distribution is well-described as a unit-width Gaussian centered at 0. The non-Gaussian tails seen in both data and Monte Carlo are roughly consistent with known approximations in our estimate of the residual errors, which (for instance) do not take into account dead electronics channels or Moliere scattering of the tracks.

A higher-level self-consistency test comes from fitting the incoming and outgoing branches of a cosmic ray as two separate tracks. Each of these tracks has a similar number of hits as a typical *BaBar* physics track. Both tracks are fit independently, but because they represent the same particle, they should have equivalent parameters at the point where they meet if the alignment is correct. In this study we use cosmic ray events selected as described in Sec. 5. Fig. 12 compares the impact parameters and angles of the split cosmic tracks. These tracks were fit using their SVT hits plus a DCH curvature constraint, which reduces the large uncertainty in d_0 and ϕ coming from their correlation with curvature, which is poorly measured in the SVT alone due to its small lever arm. All tracks in these plot have a momentum above 2.0 GeV/c. The plots show good agreement between *BaBar* data and Monte Carlo. Fitting the

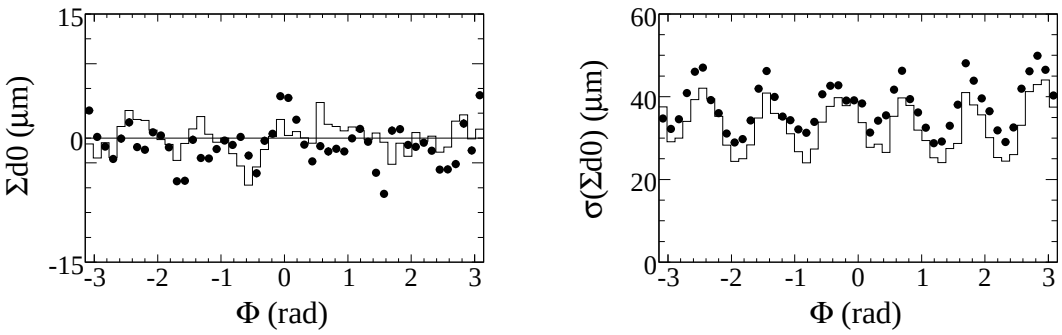

Fig. 13. ϕ-dependence of the mean (left) and σ (right) of the $e^+e^- \to \mu^+\mu^-$ track d_0 mismatch. The *BaBar* data are shown as points, Monte Carlo simulation as

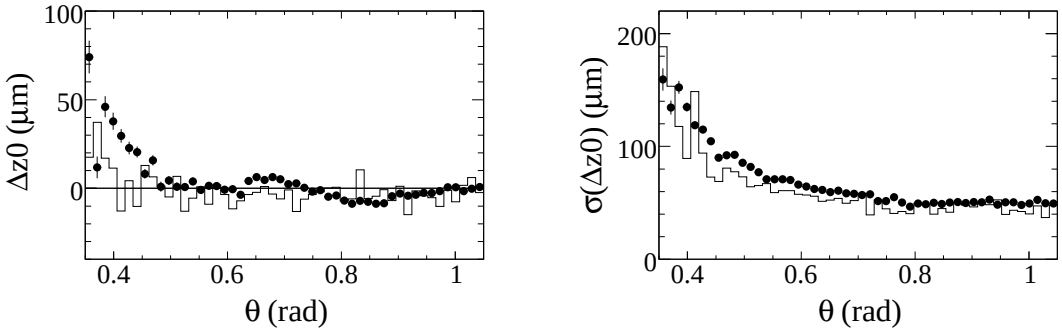

Fig. 14. θ-dependence of the mean (left) and σ of the (right) $e^+e^- \to \mu^+\mu^-$ track z_0 mismatch. The *BaBar* data are shown as points, Monte Carlo simulation as histograms.

BaBar data distributions to a Gaussian we extract single-track resolutions of 29 μm for z_0, 24 μm for d_0, 451 μrad for ϕ, and 512×10^{-6} for $\tan(\pi/2 - \theta)$, by scaling the fitted Gaussian sigma by $1/\sqrt{2}$.

Another validation test comes from comparing the reconstructed origin points of the two tracks produced in $e^+e^- \to \mu^+\mu^-$ events. Because these tracks are known to originate at the same point, the difference in their reconstructed parameters can be used to measure the impact parameter resolution, and to look for systematic biases left by the alignment procedure. In this study we use $e^+e^- \to \mu^+\mu^-$events selected approximately as described in Sec. 5. To test the performance of the entire *BaBar* tracking system, these tracks are fit with both SVT and DCH hits. This brings in the possibility that misalignments inside the DCH or between the DCH and SVT may affect our results.

Figure 13 shows the ϕ-dependence of the $e^+e^- \to \mu^+\mu^-$ tracks transverse impact parameter mismatch on the left, and its resolution on the right. The plotted points are the mean and the σ of a Gaussian fit to Σd_0 in each ϕ bin, respectively. The d_0 mean shows some structure at the level of a few microns RMS, roughly consistent for data and Monte Carlo. We believe this structure comes from track fit biases due to dead electronics in the inner layers of the

SVT , which are partially simulated in the Monte Carlo.

The d_0 resolution (σ) shows a periodic variation due to the six-fold symmetry of the inner layers of the SVT (see Fig. 2), which modulates the extrapolation of the hit error according to the distance from the innermost hit to the production point. This periodicity is well-reproduced in Monte Carlo. The Monte Carlo underestimates the d_0 resolution by roughly 10%, consistent with the underestimation of the individual u hit residual core resolution.

Because PEP-II produces a boosted final state, we cannot simply compare the $e^+e^- \to \mu^+\mu^-$ track longitudinal impact parameters as we did the transverse impact parameters, since the tracks are not back-to-back in the lab frame in the longitudinal projection. We can extract some information about the longitudinal impact parameter by constraining the production point to the event-average beamspot position, which is well measured in the transverse plane. However, this couples the statistical and systematic errors of the beamspot determination with the alignment validation. In addition, because the beamspot is large (roughly 100 μm) in the PEP-II bend (x) direction, the comparison has meaningful precision only for vertical tracks, where the beamspot constraint is limited by its measurement resolution of roughly 10 μm.

Figure 14 shows the polar-angle dependence of the $e^+e^- \to \mu^+\mu^-$ longitudinal impact parameter mismatch on the left, and its resolution on the right. The plotted points are the mean and σ of a Gaussian fit to Δz_0 in each θ bin, respectively. To select vertical tracks, we use only events with track azimuthal angles $|\phi \pm \pi/2| < 0.2$. The agreement between data and Monte Carlo is reasonable. The observed discrepancy of the z_0 mean at small θ may be due to systematic effects in the beamspot position determination, which become amplified at small angles. It may also be related to remaining aplanar distortions in the outer layer wafers at large $|z|$, as discussed in section 12. The Monte Carlo underestimates the z_0 resolution by roughly 15%. This difference is partly explained by the 10% difference in intrinsic v hit residual core width, and by the fact that the Monte Carlo does not model the beamspot position measurement resolution.

10 Validation of the Alignment Systematics

We test the ability of the local alignment procedure to remove systematic distortions by introducing a coherent misalignment of the SVT wafers, and then running the alignment procedure taking that misalignment as the initial condition. These global distortions are particularly difficult to remove as the residuals used in the alignment procedure typically depend on them only to second order. We test nine distinct distorted initial conditions, as described

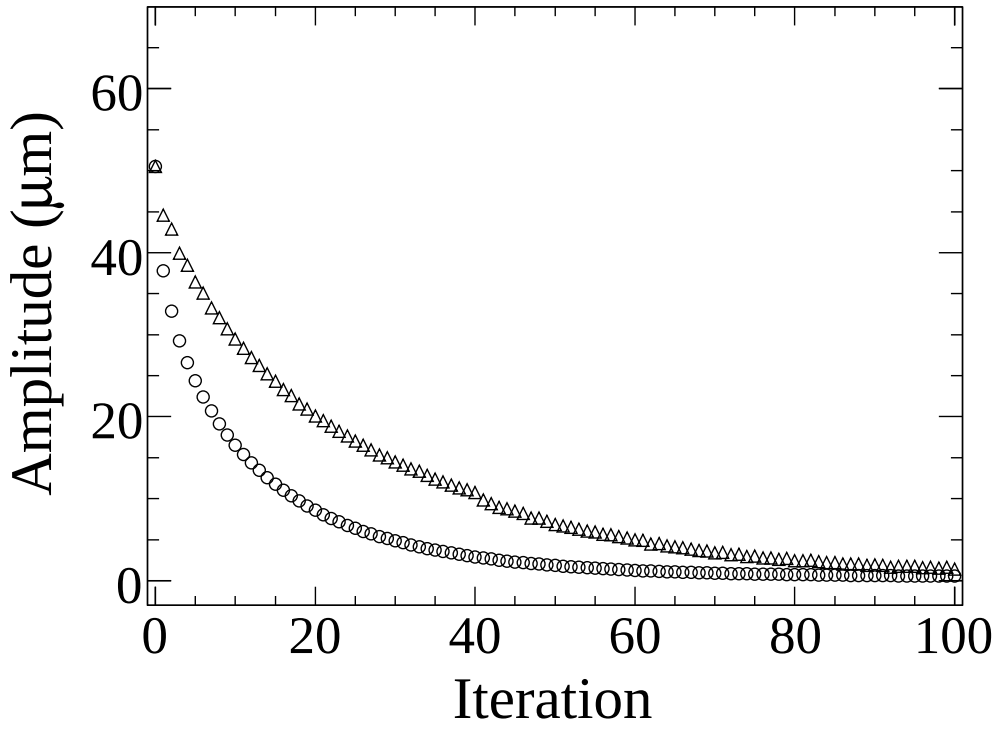

Fig. 15. Amplitude of the remaining z-expansion distortions as a function of iteration during the alignment procedure with data starting with the 50 μm amplitude z-expansion distortion initial condition for data. Circles ($\circ$) illustrate procedure with survey information, while triangles ($\triangle$) illustrate procedure with this information removed.

in Table 2 and Fig. 4. We set the initial scale of these distortions to 50 μm.

In Fig. 15 we show how amplitude of the z-scale of the SVT converges back to the initial value after a 50 μm misalignment is applied to a standard alignment set data sample. For comparison, convergence with and without optical survey measurements is shown. The survey information demonstrably provides an important constraint on systematic distortions, which improves the convergence of the alignment procedure, and reduces the systematic error of the final alignment.

We fit for the amplitude of the distortion remaining as a function of iteration for each of the nine tests. In the data we compare wafer positions to the converged set of alignment parameters prior to introducing systematic distortions. In all nine cases we find the alignment procedure is capable of reducing global distortions to a negligible level. We also perform tests with Monte Carlo where we compare wafer positions to the true positions known from MC generation. In Fig. 16 we show four representative initial misalignments in Fig. 4. The rate of convergence of the nine global distortions, defined as the decay constant of

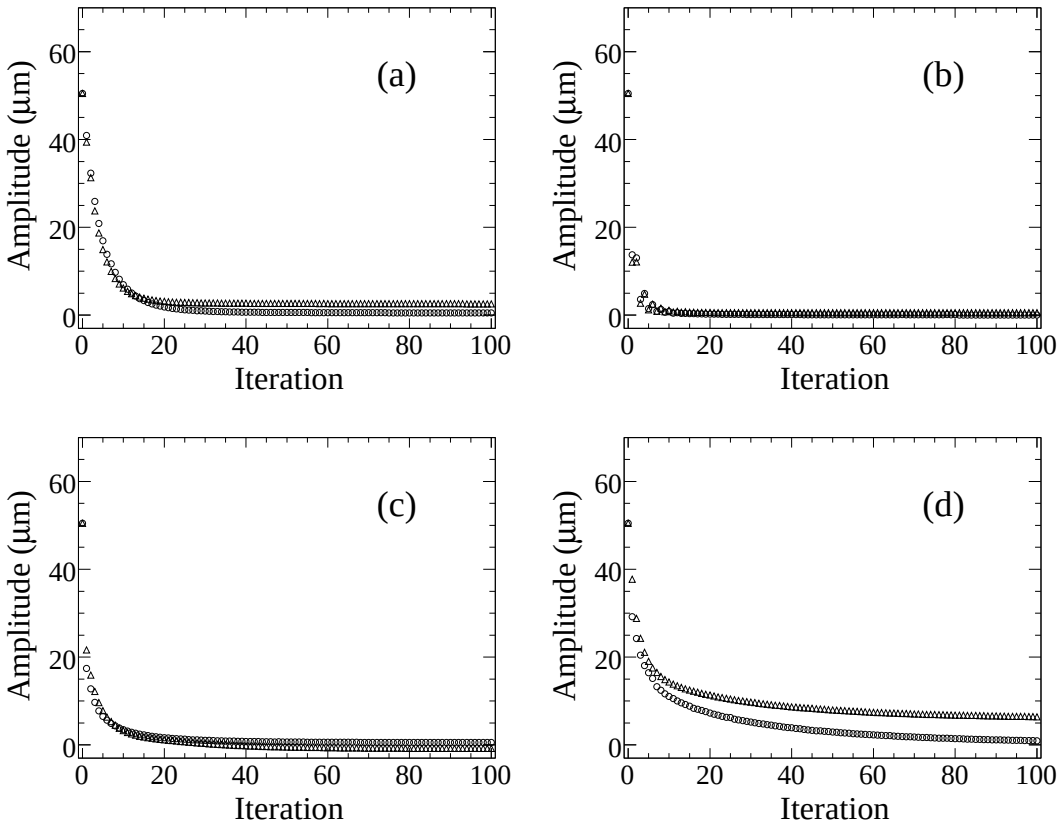

Fig. 16. Amplitude of the remaining distortions as a function of iteration during the alignment procedure starting with the 50 μm amplitude distortion initial condition : (a) telescope, (b) curl, (c) bowing, and (d) elliptical effects. Circles (○) show results with data and triangles (△) represent Monte Carlo.

an exponential fit to the scale of the misalignment per iteration, is given in Table 3.

Overall, global distortions are the most weakly constrained deformations and it was found empirically that the order of 100 iterations were necessary to solve for these deformations. Fig. 17 shows the number of wafers which are not converged as a function of iteration. The increase after iteration six is due to the residual requirement applied after partial convergence of the procedure, as discussed in Sec. 8. The convergence requirement was chosen empirically to allow convergence of the global distortions shown in Fig. 16.

Figure 18 plots the time-dependence of some global distortions of the SVT , compared to the initial day-one alignment. The plot covers roughly 50 time periods when the internal structure of the SVT was suspected to be changing due to mechanical stress during detector access or humidity changes. The y-axis plots the amplitude of the change in a particular distortion, obtained using a method similar to that shown in Fig. 4. The number of days since the initial alignment is shown on the x-axis. The first two points at negative time compare the initial alignment to the ideal geometry (day -300) and the

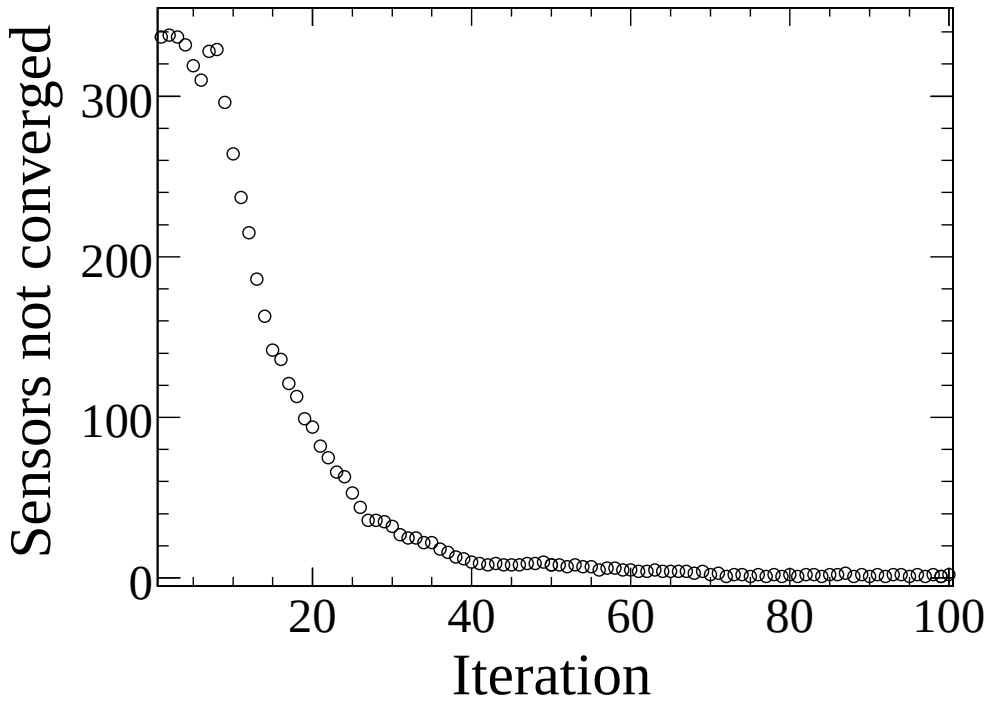

Fig. 17. Number of unconverged wafers as a function of iteration starting with a $50\,\mu$m amplitude elliptical expansion distortion initial condition relative to final alignment using *BaBar* data.

survey geometry (day -200), as shown in Figs. 7 and 6 respectively. The large bowing in the outer layers around day 700 is due to an accidental humidity increase, which caused the carbon-fiber support structure to expand.

Our study of potential distortions of the SVT wafer positions in the alignment procedure places limits on systematic uncertainties in physics measurements, such as particle lifetime or the Cabibbo-Kobayashi-Maskawa (CKM) matrix parameter measurements through time-evolution studies of the B meson decays. There are two sources of systematics: any misalignment due to time-variations, such as those shown in Fig. 18, and due to imperfections in the alignment procedure with residual misalignments remaining, both statistical and systematic. We minimize the former by having about 50 independent time periods which follow major changes in the detector. The systematic distortions are controlled in the validation plots, such as those shown in Figs. 12 and 13. We have validated that all major systematic distortions shown in Table 2 would be visible in the above validation plots.

To facilitate studying the impact of potential remaining misalignment on physics analysis, we create special alignment sets with intentional misalignments, which describe remaining misalignments possible by either of the mech-

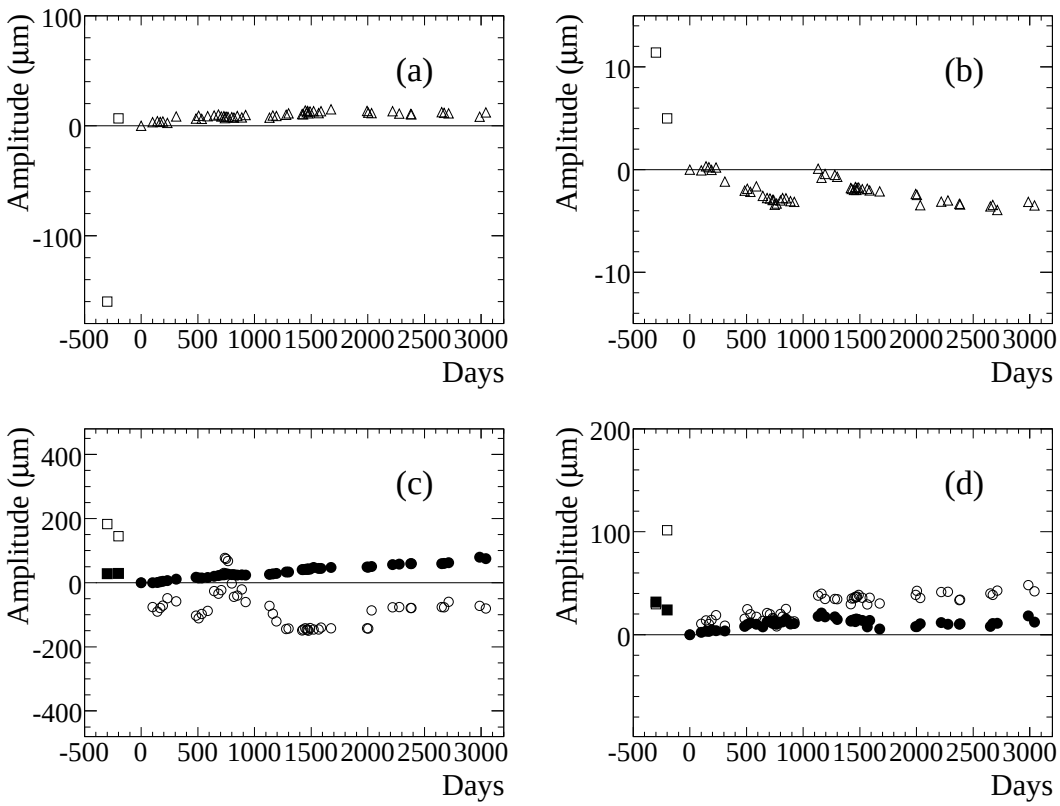

Fig. 18. Time-dependence of the global distortions of the SVT when compared to day-one. Each point represents an alignment set and the day since the first alignment set is shown on the x-axis, while days -300 and -200 correspond to ideal and survey geometries. Amplitudes of four distortions are shown: (a) telescope, (b) curl, (c) bowing, and (d) elliptical effects. In (a) and (b) triangles ($\triangle$) represent average effects for all wafers, while in (c) and (d) the inner three layers and the outer two layers are shown separately with the filled ($\bullet$) and open ($\circ$) circles, respectively.

anisms above. The impact of remaining misalignment on a physics analysis is evaluated by reconstructing the tracks in simulated events using a misaligned parameter set, followed by the normal analysis chain. The difference between that result and the one produced from analyzing the same events with perfectly aligned track reconstruction is taken as the misalignment systematic error. The misalignment error is rarely the dominant systematic error of a *BaBar* analysis.

11 Fit for the e^+e^- Beam Energies

In the joint fit of the $e^+e^- \to \mu^+\mu^-$, the same four-momentum for the initial state (from the beam energy monitoring) and the final state is assumed. However, while beam angle measurements were found to be precise in the PEP-II beam monitoring, measurements of the beam energies were not stable to better

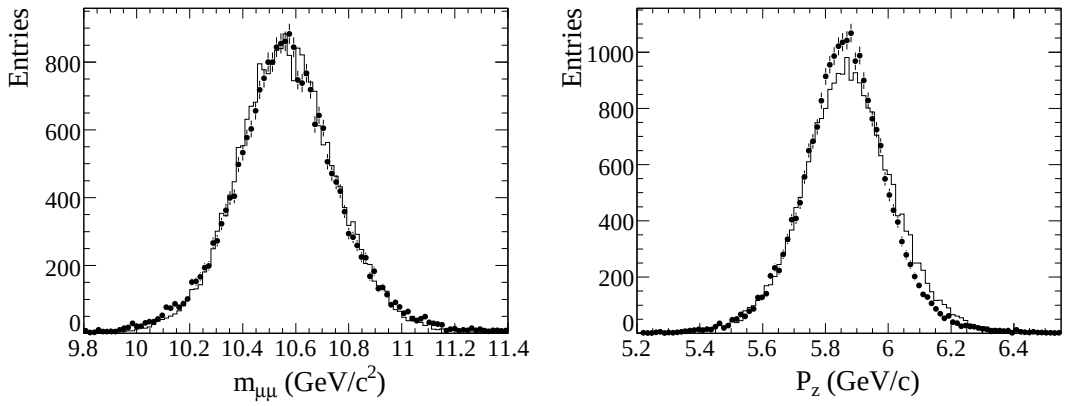

Fig. 19. Distribution of the dimuon $\mu^+\mu^-$ invariant mass (left) and total measured momentum of the two muons (right) in data (points with error bars) and MC (histogram).

than a couple of MeV. In addition, initial and final state radiation can systematically change the total energy and momentum of the $e^+e^- \to l^+l^-$pair compared to the original e^+e^-. Assuming the wrong momentum constraint in the pair fit could introduce a *telescope* distortion, as described in Table 2. This effect was indeed observed in early local alignment validation. We use the cosmic tracks to constrain this effect. Because these cross the entire detector, they would detect a telescope distortion as a kink in the track polar angle going from one side to the other. Because of the cosmic track constraint, we can use the SVT local alignment procedure to determine simultaneously the alignment and the boost of the two muons.

In Fig. 19 we show the distributions of the dimuon invariant mass and total measured momentum of the two muons (boost). In the alignment procedure, we fit the average $\mu^+\mu^-$ boost and use it in the four-momentum pair constraint instead of the initial state boost. We then iterate, allowing a simultaneous extraction of the boost value and the geometrical constraints on the telescope global distortion. This simultaneous fit makes the local alignment algorithm more stable, and allows for monitoring the e^+e^- beam boost. In principle, we can also use the $e^+e^- \to \mu^+\mu^-$ events to measure the total e^+e^- invariant mass. However, because of the high momentum, and therefore poor relative momentum resolution, of the tracks in $e^+e^- \to \mu^+\mu^-$ events, the resolution is poor. A better estimate of the e^+e^- invariant mass can be made using $\Upsilon(4S) \to B\overline{B}$ decays, when one B meson decays fully hadronically.

12 Residual Aplanar Distortions

As described in Sec. 7, we see clear evidence for substantial aplanar distortions of wafers in the inner layers of the SVT . The dominant effect can be charac-

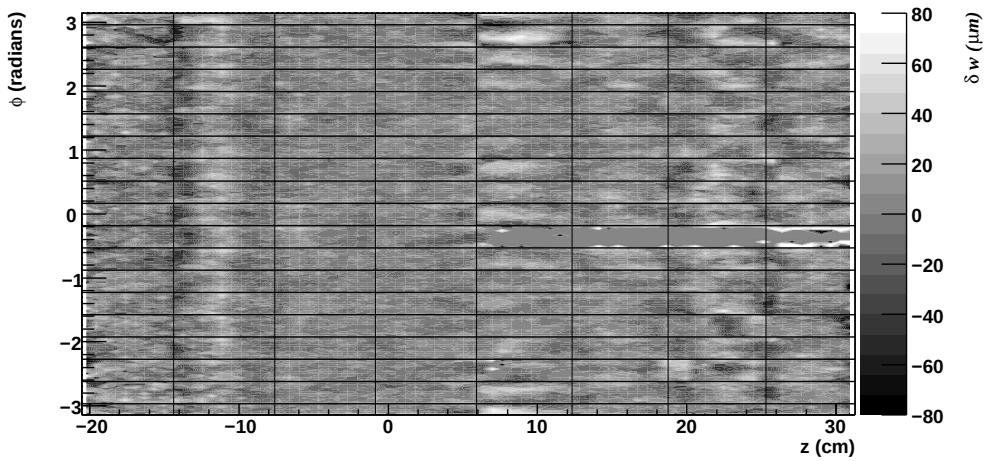

Fig. 20. Average δw projection of the v hit residuals in layer five as a function of the global ϕ and z position for high-momentum tracks in the *BaBar* data. The overlaid lines show the approximate wafer edges. The forward module near $\phi = -0.25$ radians shows no data due to its failed v view readout electronics.

terized as a bowing of the wafers along their u coordinate, which has limited external support. We also studied more general forms of aplanar distortions using the same techniques discussed in Sec. 7, where we interpret the average residuals as due to a local δw distortion. From these studies we discovered that more general forms of aplanar distortion are present in the *BaBar* SVT.

An example is shown in Fig. 20. This plots the average δw of layer five v hit residuals for high-momentum tracks in *BaBar* data as a function of their position. The distortions are substantial, but they cannot be described by simple u bowing of individual wafers. Some patterns are evident, such as the v bowing of the wafers centered around $z = -10$ cm. A similar pattern is seen at the forward end of layer five, and at both ends of layer four, but not in similar locations in layers $1 \rightarrow 3$. This suggests that these distortions may be related to the bend at the end of the arch modules (see Fig. 1). Fig. 20 also shows irregular distortions in many layer five wafers near their edge at approximately $z = 5$ cm. This is where the forward and backward module halves, which were constructed separately, join. Similar distortions occur where the layer four module halves join. No such distortions are seen in the middle of the inner layer modules, which were built in one piece.

The aplanar distortions in the outer layers are thought to be responsible for some of the remaining irregularities seen in the validation plots, for instance the variation of $e^+e^- \rightarrow \mu^+\mu^-$ miss distance resolution with ϕ and θ, as shown in Figs. 13 and 14. Unfortunately, the irregularity of these distortions makes them difficult to correct, and we do not attempt to do so in the *BaBar* local alignment procedure. A full δw map of a vertex detector would in principle

be possible given a large and diverse collection of data. In particular, tracks originating at many different positions would be necessary to avoid the lack of constraint on δw when the projected track incident angle is normal to the wafer.

13 Conclusions

We have described the procedure used to determine *BaBar* SVT local alignment. We have shown that this procedure satisfies the requirements placed on the SVT performance by the *BaBar* physics goals. We have demonstrated that this procedure is robust against global distortions that could otherwise introduce unacceptable systematic biases in *BaBar* tracking and physics data.

14 Acknowledgments

The work presented in this note could not have been accomplished without the help of many people. The first *BaBar* SVT alignment procedure had major contributions from Gerald Lynch and Jochen Schiek. We also wish to thank Stefan Kluth, Amir Farbin, Vincent Lillard, Gennadiy Kukartsev, Jurgen Krosberg, Chung Khim Lae, and Luke Winstrom for generating the alignment constants ultimately used in physics analysis. We also wish to thank Eric Charles, Fred Goozen, Natalia Kuznetsova, and Marzia Folegani for providing the optical survey measurements, Gennadiy Kukartsev for initial studies of the boost fit algorithm, Brandon Giles for his work on a-planar distortions, the SVT group for building and operating this beautiful detector, the tracking group for their support with validation studies, and *BaBar* and PEP-II for providing the data. The authors acknowledge the support from US Department of Energy, National Science Foundation, and A. P. Sloan Foundation. Part of this work was supported by the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 and by the U.S. National Science Foundation Contracts NSF-PHY0555519 and NSF-PHY0644849. A.V.G. received support from the A. P. Sloan Foundation.

A Derivative calculations for the SVT Local Alignment

To minimize the χ^2 used in the SVT local alignment procedure, we need the first derivatives of the track-hit residuals with respect to the six degrees of freedom of the wafer alignment defined in Sec. 3, being specifically the

three translations along the local wafer coordinate axes ($\hat{u}$, $\hat{v}$, and $\hat{w}$) and the three small rotations about those axes. In *BaBar*, residuals are defined as the distance in space between the track and the hit trajectory at their point of closest approach (POCA), signed by the cross-product of the track direction ($\hat{t}$) and the hit trajectory direction ($\hat{h}$). The hit trajectory is defined as a line segment in the wafer plane, with a direction given by the strips used in this hit (generally $\hat{u}$ or $\hat{v}$). Thus a barrel-module u hit (which constrains the u position of a track) has hit trajectory direction $\hat{h} = \hat{v}$, and a v hit has hit trajectory direction $\hat{h} = \hat{u}$.

The derivatives of a residual ϵ with respect to wafer translations d and rotations α can be expressed as:

$$\frac{\partial \epsilon}{\partial d_i} = \hat{D} \cdot \hat{i} \tag{A.1}$$

$$\frac{\partial \epsilon}{\partial \alpha_i} = \hat{D} \cdot (\hat{i} \times \vec{H}), \tag{A.2}$$

where $\hat{D} \equiv \hat{t} \times \hat{h}/|\hat{t} \times \hat{h}|$, $\vec{H}$ is the position of the hit at POCA, relative to the geometric center of the wafer, and $i \in \{u, v, w\}$. These derivatives are coded in the *BaBar* local alignment procedure using the CLHEP (12) class library geometry methods. The exact derivative calculations and implementation were tested by comparison with numerically-computed derivatives.

References

[1] For review of alignment approaches, see S. Blusk *et al.*, proceedings of the first LHC Detector Alignment Workshop, report CERN-2007-004.

[2] PEP-II Collaboration, PEP-II Conceptual Design Report, SLAC-R-418 (1993).

[3] The *BaBar* Collaboration, B. Aubert *et al.*, Nucl. Instrum. Methods **A479**, 1 (2002).

[4] D. N. Brown, representing the *BaBar* collaboration, the New *BaBar* Analysis Model CHEP 2004 presentation 347.

[5] D. N. Brown, representing the *BaBar* collaboration, the Detector-Model package for modeling detector geometry, material, and alignment http://www.slac.stanford.edu/BFROOT/repo/DetectorModel/

[6] I. Gaponenko, D. N. Brown, representing the *BaBar* collaboration, CDB - Distributed Conditions Database of *BaBar* Experiment CHEP 2004 presentation 316

[7] J. K. Ousterhout, Tcl and the Tk toolkit, Addison-Wesley.

[8] The *BaBar* Collaboration, *BaBar* Technical Design Report SLAC-R-95-457.

[9] D. N. Brown, representing the *BaBar* collaboration, the *BaBar* Track Fitting Algorithm, presented at CHEP 2000.
[10] The *BaBar* Prompt Reconstruction System presented at CHEP 1997.
[11] Private communications with G. Battignani and R. Kerth.
[12] L. Lnnblad *et al.*, a Class Library for High Energy Physics http://proj-clhep.web.cern.ch/proj-clhep

Table 1
Geometric parameters of five SVT layers comprised of 340 silicon wafers. The radial range for layers 4 and 5 includes the radial extent of the arched sections. The radius refers to the closest transverse radius. The length (L) and width (W) are along z and ϕ, respectively.

layer	wafers in module	modules in layer	radius (mm)	z (L) (mm)	ϕ (W) (mm)
1	4	6	32	42	41
2	4	6	40	45	49
3	6	6	54	44	71
4	7	16	91–127	54-68	43-53
5	8	18	114–144	68	43-53

Table 2
Main systematic distortions in a system with cylindrical geometry and multiple layers. Distortions in r, z, and ϕ are considered as a function these coordinates.

	Δr	Δz	$r\Delta\phi$
vs. r	radial scale	telescope	curl
vs. z	bowing	z-scale	twist
vs. ϕ	eliptical	skew	squeeze

Table 3
Decay time (in units of iterations) for the main systematic distortions in the SVT local alignment procedure. The initial distortion is 50 μm in all cases. The remaining distortion is quoted in the table. Distortions in r, z, and ϕ are considered as a function these coordinates.

	Δr	Δz	$r\Delta\phi$
vs. r	radial	telescope	curl
decay (iterations)	5.6	5.1	1.3
distortion (μm)	0.7	0.5	0.1
vs. z	bowing	z-scale	twist
decay (iterations)	2.6	11.2	12.0
distortion (μm)	0.6	0.6	0.1
vs. ϕ	eliptical	skew	squeeze
decay (iterations)	11.8	33.6	32.0
distortion (μm)	0.9	4.9	4.5

LEGAL DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.