

Design and Performance of the SLD Vertex Detector, a 307 Mpixel Tracking System

K Abeⁱ, A Arodzero^f, C Baltay^l, J E Brau^f, M Breidenbach^h, P N Burrows^d, A S Chou^h,
G Crawford^h, C J S Damerell^{g*}, P J Dervan^a, D N Dong^d, W Emmet^l, R L English^g,
E Etzion^a, M Foss^h, R Frey^f, G Haller^h, K Hasukoⁱ, S S Hertzbach^c, J Hoeflich^h,
M E Huffer^h, D J Jackson^g, J A Jaros^h, J Kelsey^d, I Lee^d, V Lia^d, A L Lintern^g,
M X Liu^l, S L Manly^l, H Masuda^h, A K McKemey^a, T B Moore^l, A Nichols^g, T Nagamineⁱ,
N Oishi^e, L S Osborne^d, J J Russell^h, D Ross^d, V V Serbo^k, N B Sinev^f, J Sinnott^l,
K Skarpaas VIII^h, M B Smy^b, J A Snyder^l, M G Strauss^c, Su Dong^h, F Suekaneⁱ,
F E Taylor^d, A I Trandafir^c, T Usher^h, R Verdier^d, S J Watts^a, E R Weiss^j, J Yashimaⁱ,
H Yutaⁱ, G Zapalac^k

^a*Brunel University, Uxbridge, Middx, UB8 3PM, UK*

^b*Colorado State University, Fort Collins, CO 80523, USA*

^c*University of Massachusetts, Amherst, MA 01003-4525, USA*

^d*Massachusetts Institute of Technology, Cambridge MA 02139-4307, USA*

^e*Nagoya University, Nagoya 464, Japan*

^f*University of Oregon, Eugene, OR 97403-5203, USA*

^g*Rutherford Appleton Laboratory, Chilton, Didcot, OX11 0QX, UK*

^h*Stanford Linear Accelerator Center, Stanford, CA 94309, USA*

ⁱ*Tohoku University, Sendai 980-77, Japan*

^j*University of Washington, Seattle, WA 9819501560, USA*

^k*University of Wisconsin, Madison, WI 53706, USA*

^l*Yale University, New Haven, CT 06511, USA*

Nuclear Instruments and Methods

* Corresponding author. Tel +44 1235 446298, Fax +44 1235 446733, e-mail c.damerell@rl.ac.uk

This work was supported by US Department of Energy contracts: DE-FG03-93ER40788 (CSU), DE-FG02-92ER40715 (Massachusetts), DE-AC02-76ER03069 (MIT), DE-FG06-85ER40224 (Oregon), DE-AC03-76SF00515 (SLAC), DE-AC02-76ER00881 (Wisconsin), DE-FG02-92ER40704 (Yale); National Science Foundation grant PHY-92-03212 (Washington); the UK Particle Physics and Astronomy Research Council (Brunel and RAL); and the Japan-US Cooperative Research Project on High Energy Physics (Nagoya, Tohoku).

Abstract

This paper describes the design, construction, and initial operation of SLD's upgraded vertex detector which comprises 96 two-dimensional charge-coupled devices (CCDs) with a total of 307 Mpixels. Each pixel functions as an independent particle detecting element, providing space point measurements of charged particle tracks with a typical precision of $4\ \mu\text{m}$ in each coordinate. The CCDs are arranged in three concentric cylinders just outside the beam-pipe which surrounds the e^+e^- collision point of the SLAC Linear Collider (SLC). The detector is a powerful tool for distinguishing displaced vertex tracks, produced by decay in flight of heavy flavour hadrons or tau leptons, from tracks produced at the primary event vertex. The requirements for this detector include a very low mass structure (to minimize multiple scattering) both for mechanical support and to provide signal paths for the CCDs; operation at low temperature with a high degree of mechanical stability; and high speed CCD readout, signal processing, and data sparsification. The lessons learned in achieving these goals should be useful for the construction of large arrays of CCDs or active pixel devices in the future in a number of areas of science and technology.

PACS: 29.40.Gx; 29.40.Wk; 85.60.Gz; 42.79.Pw; 13.20.-w; 12.25.-k; 13.35.Dx

Keywords: Vertex detector; Particle tracking; Pixels; Charge-coupled device (CCD); Heavy flavour; Charm quark; Bottom quark; Tau lepton

Table of Contents

1 Introduction

- 1.1 Summary of VXD2 Design Features
- 1.2 Experience with VXD2
- 1.3 Advances in CCD Technology
- 1.4 VXD3 Design Overview
 - 1.4.1 The 3-Barrel Layout
 - 1.4.2 Impact Parameter Resolution
- 1.5 Physics Motivation for Upgrade

2 Mechanical and Thermal Design

- 2.1 Overview
- 2.2 Ladder Design
 - 2.2.1 Mechanical/Thermal Details
 - 2.2.2 Electrical Details
- 2.3 Detector Support Structure
- 2.4 Detector-related Mechanical Equipment
- 2.5 Cryostat and Cooling System

3 CCD Design and Production

- 3.1 Design Specifications
 - 3.1.1 Imaging Area
 - 3.1.2 Readout Register
 - 3.1.3 Output Amplifier
- 3.2 Prototype Evaluation
- 3.3 Results from Production Phase
- 3.4 Radiation Damage Tests

4 Optical Survey and Detector Position Monitoring

- 4.1 Ladder Survey
- 4.2 Barrel Survey
- 4.3 Survey Data Analysis
- 4.4 Summary of Detector Survey
- 4.5 Detector Position Monitoring

5 Detector Electronics and Readout

- 5.1 Charge Transfer on CCDs
- 5.2 Digitization of CCD Output Signals
- 5.3 Data Transfer to FASTBUS
- 5.4 Data Acquisition and Processing
 - 5.4.1 Data Acquisition Modules (LFM)
 - 5.4.2 Data Processing Overview
 - 5.4.3 Cluster Processing
 - 5.4.4 Performance of Data Acquisition System

6 Detector Performance in SLD

- 6.1 Operational Experience
 - 6.1.1 Hardware Performance
 - 6.1.2 Beam Background and Data Volume
 - 6.1.3 Observation of Radiation Damage Effects
- 6.2 Tracking Efficiency
- 6.3 Alignment and Resolution in Tracking and Impact Parameter

7 Lessons Learned and Future Prospects

- 7.1 Lessons Learned from VXD3
 - 7.1.1 CCD Design
 - 7.1.2 Ladder Design and Support Structure
 - 7.1.3 Flex-circuits and Front-end Electronics
- 7.2 Future Prospects

Acknowledgements

Abbreviations, Acronyms and Symbols

Tables

References

Figure Captions

1 Introduction

After the invention of the charge-coupled device or CCD in 1970 [1], the first paper to explore the possibility of using such devices as high precision detectors of minimum-ionizing particles (hereafter referred to as min-I particles) appeared in 1981 [2], followed shortly thereafter by the experimental verification of their performance in a CERN test beam [3]. A pair of these devices, covering the full spectrometer aperture with a total area of less than 1 cm², was subsequently used for the identification of charmed particles in a fixed target experiment [4, 5].

It was realized as early as 1981 [6] that these devices offered the possibility of extraordinary physics performance in the e^+e^- linear collider environment, significantly superior to what would be achievable with silicon microstrip detectors. The essential reasons for the harmonious match between the linear collider and CCD technologies were:

- (a) very small beam spots, hence a well defined primary vertex for every event.
- (b) small diameter beam-pipe, necessary for precision vertexing and compatible with the limited practicable area coverage with CCD detectors.
- (c) small-scale hence less massive detector; degradation of the track extrapolation due to multiple scattering could be greatly reduced.
- (d) long interval between bunch crossings. While this was not sufficient for complete readout, the background integrated during readout would be only ~ 10 bunch crossings as opposed to $\sim 10^5$ bunch crossings at a circular machine.
- (e) highly segmented pixel structure, which would comfortably absorb high background per bunch crossing, likely to be found in a linear collider. Using coarsely segmented devices such as silicon microstrips, the inner barrel would need to be pushed out to much larger radius, with consequent degradation in the performance.

The SLD experiment at the SLAC Linear Collider (SLC), provided the first home for these detectors in a colliding beam environment. After tests with a prototype detector VXD1, consisting of a few ladders, the 120 Mpixel detector VXD2 was installed for physics runs starting in January 1992. These early detectors were assembled from a standard commercial imaging CCD of area approximately 1 cm² and pixel size $22 \times 22 \mu\text{m}^2$. While such devices were quite appropriate to the fixed target environment, they were barely adequate for a collider detector with nearly 4π solid angle coverage.

As a result of rapid advances in CCD technology over the past 10 years, the opportunity arose for replacing VXD2 with a much more powerful vertex detector. The proposal to build this upgraded detector VXD3 was approved in March 1994, and the detector was installed in January 1996. Fig. 1 shows the contrasting layout of CCDs around the SLD beam-pipe between VXD2 and VXD3. The geometrical advantages afforded by the much larger VXD3 CCDs are immediately apparent. This paper, which provides a detailed description of this state-of-the-art CCD-based vertex detector, is the first complete technical report on any of these detectors used in high energy physics. For this reason, we include a brief introduction on the characteristics of, and operating experience with VXD2.

1.1 Summary of VXD2 Design Features

The overriding technical purpose of any vertex detector is to measure the tracks in an event with sufficient precision that, extrapolating them to the interaction region (IR), it is possible to distinguish between those from the primary vertex (PV) located at the interaction point (IP), and those from secondary or tertiary vertices (SV or TV) due to the decay of particles containing heavy quarks (bottom or charm), or tau leptons. Once this is achieved, a whole range of physics topics related to flavour tagging is opened up. In practice, the majority of tracks are generally rather low in momentum, so that the measurement precision at the IR is limited by multiple scattering in the material of the vertex detector, rather than by its intrinsic measurement precision. In these circumstances, the usefulness of the detector for physics is dependent on achieving a small inner layer radius, a spacing between layers similar to that radius (allowing a good lever-arm for extrapolation to the IR) and the smallest possible layer thickness.

In VXD2, compromises were forced on the detector design due to background conditions at SLC. Design details of this 4-barrel detector are described in [7]. The basic detector elements (8-CCD ladders) had to be drastically redistributed as the design radius of the beam-pipe was increased from 10 mm to 25 mm. The resultant problems are clearly seen in Fig. 1. Barrel 1 (13 ladders) covered only a little more than 50% of the azimuth, leaving Barrel 2 to fill the gaps. Similarly Barrel 3, despite being only 4 mm beyond Barrel 2 in radius, covered not much more than 50% of the azimuth with 17 ladders, with Barrel 4 filling the gaps. This assembly thus guaranteed only two hits on a track, with sometimes a radial separation of only 4 mm between them. As well as the problems with azimuthal coverage, and despite using eight CCDs on a ladder, the polar angle coverage was also barely adequate.

Constructing each ladder out of eight small CCDs implied a rather complex mother board design. On each side of the mother board was a two-layer trace pattern plus a ground plane. Ceramic (alumina) substrates were used, and patterned using thick film technology. CCDs

were attached using a thermoplastic adhesive; rework capability was essential given the finite probability of one of the eight devices failing the acceptance tests. The overall thickness of each ladder was 1.15% radiation length (X_0). Electrical contact to the CCDs was made by wire bonding, and to the ends of the ladders by custom designed micro-connectors. To minimise electrical interference, the drive pulses were fed to one end of each ladder, and the analogue biases and signal outputs were routed through the other end.

Front-end electronics, mounted within the inner cylinder of the Central Drift Chamber (CDC) at either end of the VXD cryostat, consisted of fast drive circuits for the linear or readout register (R drive) at the north and preamplifiers at the south end. The slow drive pulses for the parallel or imaging register (I drive), as well as timing signals and analogue biases, were fed in on fine coax cables from the electronics counting house on top of SLD. In order to minimize the effects of radiation damage, explained in Section 1.2, the detector was operated at a temperature of 190 K. It was housed in a very low mass foam cryostat, and cooled by boiloff gas from liquid nitrogen. This coolant gas was piped to and from the detector by means of vacuum jacketed pipes.

The pixel readout rate achievable with this detector was restricted to 2 MHz. The analogue shaping time of approximately 300 ns gave noise performance for each pixel of $<100 e^-$ (rms), adequate for efficient detection and precise centroid finding of the clusters from min-I particles, which typically deposit $\sim 1200 e^-$.

1.2 Experience with VXD2

During a rushed installation phase for the detector, two micro-connector fingers (out of 1400 total) failed to make good contact, and these correctable faults were only discovered after the cryostat was sealed and access to the detector was lost. Thus two of the 60 ladders were lost from the beginning of the experiment. This represented the most significant failure throughout the three years operation of the detector. A third bad contact led to the loss of a single CCD. A small number of CCDs developed intermittently high readout noise, an effect due to imperfect contact of micro-connector fingers as a result of thermal cycling. However, these tended to recover in time. No failures of CCDs or front-end electronics were encountered throughout the life of the detector. Occasional failures of accessible FASTBUS modules were of course repaired without problems.

The SLC background environment is predominantly electromagnetic radiation: low energy electrons spiralling around the beam direction, a broad X-ray spectrum resulting from multiple bounces of synchrotron radiation photons, and associated fluorescence. This can lead to two types of radiation damage; see [8] for a general review. Firstly, electromagnetic radiation can

generate electron-hole pairs in the gate oxide, leading to flat-band voltage shifts which may necessitate a change in the device operating bias conditions. Such effects were observed at a rather low level. Secondly, higher energy electromagnetic radiation can cause displacement damage within the bulk of the silicon crystal, leading to electron trapping centres. Due to the large distances over which the signal charge needs to be transported in a CCD (up to several centimetres), it is not surprising that these devices are very sensitive to bulk damage effects. The main effect is an increase in the average charge transfer inefficiency (CTI), as signal electrons are trapped during charge transport. Operating at reduced temperature is effective at increasing the emission time of bulk traps to such an extent that their effect on CTI can be greatly reduced. The intrinsic rate of generation of bulk traps is proportional to the non-ionizing energy loss (NIEL) of the radiation [9]. Electromagnetic radiation is much less damaging than hadronic radiation due to its lower NIEL factor. Despite this, and the choice of operating temperature of 190 K, close to the optimal value for CTI suppression, VXD2 did suffer a measurable amount of radiation damage by the end of its life. This was seen as a signal reduction of up to 10% on the inward facing CCDs of Barrel 1, interpretable as an effective CTI of approximately 1.7×10^{-4} , compared with the pre-irradiation value of $< 10^{-5}$. The radiation levels when running normally were negligible, but could be very much higher under beam tuning conditions, when the beam would be intercepted by scanned wires, flooding the vertex detector with scattered particles.

The closely spaced 4-barrel system achieved 2-hit coverage for all tracks (on average, a track would physically traverse 2.3 CCDs). The space-point precision on each hit was approximately $5.4 \mu\text{m}$. However, due to the poor lever-arm between measurements, the impact parameter precision at the IR for a track of momentum p and polar angle θ w.r.t. the z -axis (beam direction) was limited to

$$\sigma_{r\phi} = 11 \oplus \frac{70}{p \sin^{3/2} \theta} \mu\text{m}$$

$$\text{and } \sigma_{rz} = 38 \oplus \frac{70}{p \sin^{3/2} \theta} \mu\text{m},$$

similar to that of a state-of-the-art silicon microstrip vertex detector at LEP.

Nevertheless, the flavour tagging at SLD has benefited from two major advantages with respect to LEP. Firstly, the sub-micron beam spot size and stability in the $r\phi$ view implies that one can determine the interaction point (IP) extremely accurately at any given time by averaging the fitted IP over the last ~ 20 interactions. As a result, precise determination of a decay particle flight path direction and distance can be determined if the decay vertex is observed, even if

there are no good tracks from the primary vertex for that event. Furthermore, the very fine segmentation of the CCD detector (2500 pixels per mm²) means that, despite the high hit density on Barrel 1 due to background and particles within a jet, the level of cluster merging is truly negligible. Thus, the Monte Carlo generated distribution of impact parameters agrees quite accurately with the experimental one over four orders of magnitude, without the need for large empirical correction factors commonly required with less highly segmented detectors. This agreement forms the foundation for a more elaborate interplay between data and Monte Carlo. Simulated hadronic decays of Z^0 s can be used to develop and evaluate procedures for such requirements as flavour identification within a jet. Confirmation of these procedures is achieved by applying them to real data, and comparing appropriate kinematic and other distributions between the data and the Monte Carlo. While this approach is quite standard, the freedom from large correction factors makes the systematic precision exceptionally high, in this analysis. It has therefore become possible to develop a topological vertex finding method using VXD2 [10], a procedure that goes well beyond the impact parameter method of flavour tagging widely used in the past. The impact parameter method for tagging b -jets (for example) has limited purity due to an irreducible presence of charm jets in the event sample. The topological approach takes advantage of the precisely known primary vertex in SLD for every event. The cleanly measured space points in the vertex detector then permit a search for secondary and/or tertiary vertices. The effective mass of the SV + TV tracks is called the vertex mass, and this can be partly corrected for missing neutrals (e.g. neutrinos) by looking at the transverse momentum imbalance with respect to the vertex flight path direction (PV to SV). This kinematic information provides an effective discriminator between charm and bottom flavoured jets.

Applying this technique to the VXD2 data has led to flavour identification with very small systematic errors, e.g. in the measurement of the b fraction in hadronic Z^0 decays, $R_b = \Gamma_{b\bar{b}} / \Gamma_{had}$, [11]. The topological approach also optimises the measurement of decay length, and this has recently led to enhanced systematic precision in the determination of (for example) B lifetimes [12]. As discussed in Section 1.5, there are physics processes for which even higher precision is essential. Apart from some fine tuning (which is ongoing in the analysis of the VXD2 data) it became apparent that any major performance upgrade would be dependent on replacing this detector. This in turn would be made possible only by major advances in the CCD technology underpinning the detector design. Fortunately, by 1993 the required progress in this area had been made.

1.3 Advances in CCD Technology

A major advantage of using CCDs is that one benefits from their multi-disciplinary user base. Progress is made by manufacturers with the support of one user community, which may then

feed through to help others. High energy physics provides a rather small and intermittent market for scientific CCDs, insufficient to sustain a major R&D programme. The design of the upgrade detector VXD3 was based largely on major progress having been made over 10 years by the CCD manufacturers for other customers. Subsequently, in understanding the VXD3 CCDs it was possible to make modest contributions to advancing the frontiers of CCD development for the common good.

VXD2 was based on the same standard imaging CCDs as had been used for the original beam tests in the early 1980s. By 1993, the most important advance was the availability of affordable fully customized devices, allowing the user to tailor the design to the specific detector requirements rather than being obliged to design a detector around some off-the-shelf CCD product. The next most important advance was the availability of much larger scientific grade devices. Using 5 inch diameter wafers, the SLD requirement for an active area of $80 \times 16 \text{ mm}^2$ was easily met. This reflects an enormous advance in quality control on the part of the manufacturers: achieving the necessary device uniformity to transport signals of a few hundred electrons reliably across more than 2000 pixels for each of 1600 separate columns represents a major achievement.

Having full flexibility regarding the layout of circuitry around the imaging area, it was possible to almost eliminate the inactive regions along the long edges of the device, and locate all significant additional circuitry (readout registers, output amplifiers, bond pads for external connections, gate protection circuitry, etc) along the short edges of the device. By slightly tapering the imaging area at the ends of the device, all the subsidiary circuits could be accommodated within the width already available.

Increasing the device active area by an order of magnitude was essential in designing a detector with improved polar angle and azimuthal coverage, but a comparable increase in readout time would have been entirely unacceptable. Modern CCD architectures allow a number of output circuits to tap off the signals along a readout register, with corresponding reduction in readout time. This however increases the number of channels of external analogue and A/D circuitry, so a balance must be reached. In the case of VXD3, it was decided to design the CCD with four outputs (one at each corner of the device). Each output could be made more sensitive than that on the VXD2 devices by reducing the load capacitance and using a second stage source follower to buffer the tiny first stage FET from the external load capacitance. With a correspondingly improved noise performance (equivalent noise charge or ENC) the output circuit was designed to run with a pixel clocking rate of 10 MHz as opposed to 2 MHz for VXD2. The combination of an increased number of outputs, operated at higher speed but with no degradation in noise performance, was the key to servicing a three times larger overall active

area for the VXD3 detector, while still keeping the local electronics within the space available for it.

Finally, high performance vertex detectors with 10^8 to 10^9 pixels can only be contemplated if the power dissipation per pixel is extremely low. If one were obliged to install liquid cooling within the active volume (as would be necessary with such a density of active pixels, for example), this would seriously increase the detector thickness, with a consequent degradation in performance. In a CCD system, the power dissipation is due to the drive pulses that transport the signal charges within the devices, and the on-CCD output amplifiers. Numerous advances (low resistance bus lines and gate electrodes, reduced gate overlap, etc) have been made by the manufacturers, greatly reducing the power dissipation per unit area for given clocking conditions. For the output amplifier, the most important factor in reducing power for given performance has been the improved first stage responsivity (by a factor 3 relative to the CCDs of VXD2).

The design specification for the VXD3 CCDs was prepared by the SLD group, with the detailed CCD design being carried out by EEV [13], the manufacturer which won the production contract.

1.4 VXD3 Design Overview

These advances in CCD technology permitted an upgrade detector design (Fig. 1) with the following main advantages with respect to VXD2:

- (a) extended polar angle coverage, to benefit from the large polarized asymmetry in physics processes in the most valuable regions of high $|\cos\theta|$.
- (b) full azimuthal coverage in each of three barrels, to achieve a self-tracking capability independent of the drift chamber, and consequently improved overall tracking efficiency.
- (c) optimized geometry with stretched radial lever-arm and reduced material in each layer, for significantly improved impact parameter resolution.

Some consideration was given to achieving these aims in a mixed barrel/endcap or barrel/lampshade geometry, but this was easily shown to be non-optimal. Mechanical and electrical connections implied that the endcap detectors would suffer from additional material in the critical region in front of them. In contrast, there was only a minor penalty in precision in extending the barrel coverage, because CCDs (with an active silicon thickness of only $20\ \mu\text{m}$)

suffer very little spatial resolution degradation due to fluctuations in ionization for tracks at shallow entrance angles, by comparison with the much thicker microstrip detectors. Thus it was certainly preferable simply to extend the barrel length to match the CDC tracking coverage, without the need to consider a mechanically very complex endcap design. The ladder length was set by the wafer sizes used by manufacturers of scientific CCDs (4-5 inches diameter); the $80 \times 16 \text{ mm}^2$ CCD dimensions permitted a reasonable number of devices per wafer. One CCD is mounted on each side of a supporting substrate (see Fig. 2) forming a 2-CCD ladder as the basic detector component having an active length of 16 cm, 1.7 times longer than the VXD2 8-CCD ladder.

Many of the SLD physics analyses (particularly the b mixing and heavy quark asymmetry measurements) can benefit from an extended polar angle coverage, due to the large analysing power at high $|\cos\theta|$. The limit of the angular coverage is defined by the CDC tracking volume, which covers the region from 20 to 100 cm in radius and has an active half-length of ~ 90 cm. Efficient tracking pattern recognition and reasonable momentum measurement require a minimum of four of the ten CDC superlayers, which defines the polar angle limit of $|\cos\theta| = 0.85$. These considerations motivated the placement of the outermost layer of VXD3 at a radius of around 48 mm, with the innermost layer at a radius of 28 mm, limited by the SLC beam-pipe. This design provided a very adequate lever-arm for precise track extrapolation to the IP.

1.4.1 The 3-Barrel Layout

Many advanced applications of SLD are dependent on topological vertex reconstruction not only to tag B hadron decays, but also to determine the charge of the B hadron or resolve the $b \rightarrow c$ cascade vertex charge structure. This requires a high efficiency in correctly assigning *all* relevant tracks to their corresponding production vertices. These algorithms can only succeed when tracking quality is well under control, since a small deficiency at the single track level can quickly multiply to disable such an attempt. This demands precise impact parameter resolution for a wide momentum range and highly efficient tracking with minimal hit misassignment.

The tracking strategy adopted for VXD2 was to reconstruct the CDC track first, extrapolate it to the VXD and then search for the best VXD hit combination to form the complete track. However, Monte Carlo studies indicate that around 5% of the prompt tracks in a hadronic Z^0 decay within the CDC active volume are either not tracked by the CDC or fail linking to the VXD, mainly due to contamination of wrong CDC hits distorting the extrapolation. This is a consequence of track merging in a dense hadronic jet environment, compounded by the relatively small CDC outer radius and a rather moderate solenoid field of 0.6 Tesla. This

problem becomes progressively worse once the track $|\cos\theta|$ increases beyond 0.7, when the available tracking length is shortened. Being an effectively two-layer device, VXD2 could not offer independent assistance to alleviate these pattern recognition deficiencies. Considering prompt tracks in hadronic Z^0 decays, for those tracks which linked to VXD2 hits, the links with two VXD hits are 99.4% correct for $p_T > 0.4 \text{ GeV}/c$, but only 95.4% correct for $p_T < 0.4 \text{ GeV}/c$. However, in the case of occasional hit inefficiency or the presence of a dead ladder, some tracks only have one VXD hit. In this case, the linking purity is significantly worse, at 96% correct for $p_T > 0.4 \text{ GeV}/c$ tracks, and 91% correct for $p_T < 0.4 \text{ GeV}/c$. The low- p_T tracks suffer in addition from the drastic deterioration of the impact parameter resolution. The allowance of one-hit links also introduces another source of confusion, as strange particle decay products or tracks from particle interactions with detector material can also pick up a background hit to produce a false link. The links with ≥ 3 VXD hits on the other hand show a much cleaner linking purity of 99.9% for tracks with $p_T > 0.4 \text{ GeV}/c$, and 99.5% for $p_T < 0.4 \text{ GeV}/c$.

These observations led to the choice of a full three-barrel coverage for VXD3, as sketched in Fig. 1 and shown in more detail in Figs. 3-5. This permits linking with a minimum of two VXD hits even given the allowance for small hit inefficiencies or some dead channels, providing a strong insurance for linking purity, impact parameter resolution uniformity, and reduction of fake links. An even more attractive prospect is that for the majority of VXD3 tracks with ≥ 3 hits, a VXD hit vector is relatively easy to reconstruct stand-alone, while fake combinations only amount to less than 30% of all vectors before any matching with CDC tracks. The very fine granularity of the CCD pixels is ideal for resolving the otherwise difficult merging track cases. These high precision VXD hit vectors in 3-D are powerful additions to the global tracking pattern recognition capability. The detailed implementation of this new strategy to recover the deficiencies in pattern recognition with the CDC alone is already developed, and working well on recent SLD data.

The custom design of the VXD3 CCDs allows the active region to extend across the full ladder width of 16 mm without bond wires protruding out of the ladder surface along the long edge. This feature permits adjacent ladders in the same barrel to be placed in a 'shingled' layout with a small cant angle of $9-10^\circ$, thus providing azimuthal coverage overlap in the range of $300 \mu\text{m}$ to 1 mm, depending on layer and CCD location on the inner or outer surface of the ladder (see Fig. 3). This layout not only establishes complete azimuthal coverage for each barrel, but also provides important assistance in tracking-based alignment, by using the tracks passing through the overlap regions, establishing direct constraints between adjacent ladders.

1.4.2 Impact Parameter Resolution

One of the most unsatisfactory features in the VXD2 design was that the ladders were forced to be very closely packed around a larger beam-pipe than originally anticipated. The consequence was considerably compromised impact parameter resolution, due to the very short radial lever-arm between hits (best case 12 mm, worse case 4 mm to extrapolate over 30 mm to the IP). The much larger ladders in the VXD3 design allow significantly improved radial lever-arm while still having an increased solid angle coverage.

The combined CDC+VXD track impact parameter resolution has rather varied contributions from the CDC information, depending on the track momentum and viewing projection. The CDC provides its best help in the $r\phi$ view, high momentum regime. The CDC ϕ resolution at its inner radius is 0.3 mrad at high momentum. This, combined with the innermost VXD hit, supplies the dominant contribution to the impact parameter resolution for VXD2 and is similar to the stand-alone vertex detector contribution for VXD3. The CDC polar angle resolution of approximately 2.0 mrad, on the other hand, is only comparable to the stand-alone polar angle resolution of VXD2. The CDC therefore still helps in the rz view at high momentum in the VXD2 case, but for VXD3 the stand-alone vertex detector capability with extended lever-arm dominates the impact parameter resolution in the rz view. In the momentum regime of a few GeV/ c and below, the contribution from the CDC angular information becomes much less significant in both views for both VXD2 and VXD3, due to the multiple scattering in the material between the VXD and the CDC. In the rz view at all momenta, and the $r\phi$ view at low momenta where one relies on the VXD stand-alone capability, the improvements from the extended radial distribution of layers in VXD3 are most significant. From these basic considerations, it is also clear that one should avoid any need for fallback on links based on single VXD hits, as their impact parameter resolution at low momentum would degrade severely, beside the concern of pattern recognition impurity. This requirement is well met by the VXD3 three-barrel layout.

The mean radii of the three layers were chosen to be 28.0, 38.2 and 48.3 mm, with 12, 16 and 20 ladders respectively. The radius of Layer 1 was constrained by the beam-pipe radius at 23.2 mm which was set by the risk of increased background at a smaller radius. The Layer 3 radius is chosen based on the angular coverage consideration mentioned at the beginning of Section 1.4. There would potentially be further impact parameter resolution gain in the rz view at high momentum by increasing the Layer 3 radius. However, this would be at the cost of losing full three-barrel coverage at high $|\cos\theta|$. This would result in not only a loss of resolution when the coverage falls back to the two inner layers, but also loss of the self-tracking capability which is especially needed for global pattern recognition at high $|\cos\theta|$. A further radially expanded detector would also increase the cost due to the need for more ladders

and electronics, as well as a more expensive detector mechanical support structure. The position of Layer 2 close to midway between Layers 1 and 3 is optimal for adequate lever-arm in case of hit inefficiency in either Layer 1 or Layer 3.

Beside the geometry optimization, the impact parameter resolution with VXD3 is further improved due to the reduction of detector material. In the case of VXD2, this amounted to 1.15% X_0 of material per ladder. The advances in technology associated with the various components required for a ladder, have permitted a much improved design (see Section 2) which dramatically reduces the material to 0.40% X_0 per ladder. The VXD3 beam-pipe titanium liner and the beam-pipe itself are also thinner than for the VXD2 design. Various materials used in the cryostat surrounding the vertex detector are replaced with lighter alternatives whenever possible. The overall material thickness comparison between VXD2 and VXD3 is shown in Table 1.

The detailed geometry and material distributions of both VXD2 and VXD3 are simulated in detail with GEANT [14]. The impact parameter resolution of tracks in hadronic Z^0 decays as well as in $Z^0 \rightarrow \mu^+ \mu^-$ and $Z^0 \rightarrow e^+ e^-$ events are studied for individual momentum and $\cos\theta$ regions. The expected impact parameter resolutions based on Monte Carlo for tracks at $\cos\theta = 0$ with VXD3 are shown in Fig. 6 in comparison to VXD2 Monte Carlo and data. The impact parameter resolutions for VXD2 data are extracted from raw impact parameter distributions after correcting for heavy flavour decays and IP resolution effects. It should be noted that the VXD2 Monte Carlo simulation generally describes the data remarkably well. The same Monte Carlo framework applied to VXD3 indicates a factor of two improvement from VXD2 in the entire rz view impact parameter resolution. Improvement in the $r\phi$ view is also a factor of two at low momentum, while somewhat less at very high momentum. A rough approximation of the expected VXD3 impact parameter resolution is

$$\sigma_{r\phi} = 9 \oplus \frac{33}{p \sin^{3/2} \theta} \mu\text{m}$$

$$\text{and } \sigma_{rz} = 17 \oplus \frac{33}{p \sin^{3/2} \theta} \mu\text{m}.$$

These dramatic resolution improvements combined with a significantly increased solid angle coverage and full three-layer tracking, were expected to bring a powerful enhancement to the SLD heavy flavour physics programme.

1.5 Physics Motivation for Upgrade

The physics accessible with precision vertex detection includes some of the most topical and fundamental measurements being made at the Z^0 . The accuracy of these measurements directly reflects the impact parameter resolution of the detector, and to a lesser extent its track reconstruction efficiency. Charge tagging efficiency is also important. With its improved impact parameter resolution, extended solid angle coverage, and enhanced pattern recognition capability, VXD3 improves the accuracy of several key heavy quark measurements to the point where theory leads one to expect new effects to emerge. The enhancement in purity and efficiency of jet flavour identification, in evolving from VXD2 to VXD3, has recently been quantified [15].

The measurements of the couplings of the b -quark to the Z^0 have become an especially interesting chapter of precision electroweak testing because they are experimentally tractable (lifetime tagging efficiently isolates clean b samples) and theoretically interesting (vertex corrections induced by new physics should perturb Standard Model predictions at the 1% level). Measuring R_b with enough precision to see evidence for new physics, pushes experimental technique to its limits. The measurement requires an extremely *clean* b tag, to minimize the systematic errors arising from charm contamination, yet *efficient enough* to permit self-calibration of the tagging efficiency with a double tag, without paying too high a price in statistics. VXD3's factor of two improvement in impact parameter resolution at low momentum, translates into a significantly improved efficiency for identifying secondary vertex activity. The resulting improvements in topological vertexing allow improved discrimination of charm and bottom. VXD3 can tag a b jet with 45% efficiency and 99% purity, which doubles the statistical power of the current detector. This will lead to an overall error in R_b of less than 0.7% with a data sample of 0.5 million Z^0 s.

The measurement of the parameter A_b which describes parity violation in the Z^0 coupling to b -quarks, will also benefit from the improved vertexing capability of VXD3. A_b is measured with the polarization-enhanced forward-backward asymmetry, which requires both b -tagging and b -quark charge determination in hadronic Z^0 decays. The increased solid angle coverage provided by VXD3 permits vertex-tagging b jets in the range of $0.75 < |\cos\theta| < 0.85$, beyond the present VXD2 acceptance, where the asymmetry is nearly maximal and the cross-section is enhanced. The result is a 15% reduction in the statistical error in A_b . The improved resolution of the new detector can increase the efficiency of a high purity b tag by roughly 30%, leading to an additional 15% reduction in the statistical error. The b -quark charge can be determined topologically when the charge of the B hadron or the subsequent charm meson is non-zero. The enhanced topological vertexing possible with VXD3 will incrementally improve

the determination of the net charge of the tracks associated with the b decay and will improve the ability to distinguish which tracks are associated with the charm decay.

The most exciting physics possibility afforded by VXD3 is the direct observation of B_s mixing through the measurement of the oscillation frequency. Such a measurement would 'complete' the determination of the CKM parameters by allowing the first precise measurement of V_{td} and it could provide the first sharp test of the CKM explanation of CP violation. LEP measurements have shown that the oscillation frequency x_s is high (>13). This puts a premium on the excellent decay length resolution possible with VXD3, which is needed to observe oscillations in the range expected by the Standard Model, $10 < x_s < 25$. The improved resolution also improves event selection efficiencies, and most significantly will allow the final state quark charge to be assigned topologically. In this technique, one first selects neutral B decays then asks for a decay pattern consistent with opposite charges at the B decay and D decay vertices. This selection procedure enhances the fraction of B_s mesons in the sample, and the sign of the B vertex- D vertex charge dipole gives the charge of the decaying b -quark. In conjunction with other techniques, this method will allow the exploration of the oscillation parameter space up to $x_s \sim 15-20$ with a 0.5 million Z^0 sample. This is well beyond what has been possible at LEP, or would be possible with VXD2. With an even larger data set, VXD3 would have sufficient resolution to see mixing up to $x_s \sim 30$.

With VXD3, the art of topological vertexing is being advanced. Already with VXD2, the physics applications [11, 12] have included secondary and tertiary vertex identification, vertex charge determination to isolate charged and neutral B meson decays, and charge dipole separation to extract quark charge. With both improved resolution and enhanced pattern recognition capability, it will be possible to clarify the full pattern of vertex topologies. From this improved technical base, one can hope to identify gluon splitting, search for charmless and double-charm B decays, isolate charm from bottom decays, and search for topological anomalies. As regards efficiency and purity of flavour tagging, charm is more difficult than bottom, and consequently will show the greater improvement with VXD3. The high resolution together with the sub-micron SLC beam spot will likely enable the development of new methods and new measurements.

2 Mechanical and Thermal Design

2.1 Overview

The mechanical design phase of the VXD3 detector was intimately coupled to the plans for optical survey and position monitoring (Section 4) and to the plans for tracking-based global and internal alignment (Section 6.3). The purpose of this overview section is to explain the general principles of the mechanical design and their connection to these other aspects.

The detector operating conditions required the placement of 96 flexible thinned CCDs in a cryogenic environment (operating temperature ~ 190 K). Obtaining a high level of mechanical stability in a low mass structure was the main design goal. The general approach to satisfying these conflicting requirements was as follows.

The CCDs were attached to thin beryllium substrates, forming ladders (Fig. 2). These fragile ladders were attached to a series of three concentric beryllium annuli, which were clamped to the overall beryllium support structure, which consists of a closed cylinder external to the detector barrels (Figs. 3-5). This structure was made from an instrument grade of hot isostatically pressed beryllium (I-250 from Brush Wellman) which was found to have exceptional isotropy and a high specific modulus (elastic modulus/density). The support structure could be split to allow the two complete half-detectors to be assembled around the SLC beam-pipe. To obtain a high level of geometrical repeatability of the support structure at assembly, all mating surfaces were lapped flat and all adjacent components were match drilled, reamed, and pinned. This extremely rigid support structure transformed the 48 flexible ladders into a highly stable 3-barrel assembly, with each ladder clamped to its rigid annuli at either end. This demountable structure permitted a succession of assemblies with one barrel at a time, allowing a full optical survey of the CCD geometry to be built up layer by layer. The mesh construction of the support shell, Fig. 5, permitted survey of each installed barrel from the outside. The VXD support structure when fully assembled thus consisted of an essentially rigid cylinder which was mounted to the beam-pipe by means of a low-stress, three-point kinematic mount. This assured that the detector shape during the survey (when the detector was mounted on a short dummy beam-pipe) would be essentially identical to that in the final assembly on the real SLD beam-pipe.

In SLD, the detector is buried within a low mass foam cryostat (Figs. 4, 5). The inner section of beam-pipe on which it is mounted is mechanically isolated from the SLC final focus system by means of an atmospherically balanced bellows system. This reduces (but does not eliminate) occasional disturbances due (for example) to alignment adjustments to the final focus. Despite the fact that the vertex detector is hidden within the cryostat, any positional

variations can be accurately inferred by monitoring the inner section of beam-pipe to which it is rigidly attached; this is discussed in detail in Section 4. Within stable epochs, the position (global alignment) of VXD3 relative to the CDC is deduced by tracking. Tracking data are also used to determine residual corrections to the internal detector geometry, in order to account for effects beyond the precision of the optical survey, as discussed in Section 6.3.

In summary, the mechanical construction of VXD3, combined with a series of optical surveys, followed by tracking-based alignment after installation, formed the basis of a coherent strategy for determining the true geometry of this delicate instrument running at cryogenic temperatures, to very high accuracy. The principal mechanical parameters of the detector assembly are summarized in Table 2.

2.2 Ladder Design

The material within the active tracking volume of the vertex detector (out to a polar angle cutoff given by $|\cos\theta|=0.9$) consists of the CCDs plus the ladder motherboards on which they are mounted. Further out in radius one encounters the vertex detector support shell, the vertex detector cryostat barrel, the inner barrel of the CDC and so on. None of these structures at larger radius is critical in terms of the impact parameter precision which characterizes the capability of the vertex detector for physics. In contrast, the ladder motherboards are extremely relevant for this, and their thickness as well as that of the CCDs needs to be reduced to the lowest possible level. However, the thinned CCDs are extremely flexible, and rely on these substrates for mechanical stability. In addition, the electrical connections to each end of the CCD must be made by traces on these motherboards. Finally, the ladder temperatures even under steady operating conditions are not precisely controlled. There are thermal gradients associated with the power dissipation on the CCDs and the flow of cooling gas, and in the event of electrical failure a CCD might need to be shut down, reducing the operating temperature of its particular ladder. Thus, tolerance to temperature excursions is essential. Designing a motherboard to satisfy these conflicting requirements is a complex compromise. After trying a number of ideas (some of which are mentioned in Section 7) a solution was found which works extremely well, as well as being 35% thinner than the original design.

The basic support structure was a thin beryllium beam, $21\text{ cm} \times 1.6\text{ cm} \times 0.381\text{ mm}$ (0.015 inches thick) having an initial flatness after machining (as measured by the sagitta when held at one end on edge) of less than $100\text{ }\mu\text{m}$. Onto this was bonded, in one operation for both sides, a pair of kapton/copper flex-circuits: $13\text{ }\mu\text{m}$ kapton, carrying $\frac{1}{2}$ ounce ($17.8\text{ }\mu\text{m}$) rolled, annealed copper traces, passivated by a kapton cover layer. These circuits ran the full width and length of the ladder, with 'pigtail' extensions to carry the electrical traces to micro-connectors, approximately 5 cm beyond each end of the ladder. An acrylic thermoplastic

adhesive was used to attach the flex-circuits to the motherboards under pressure, at a temperature of approximately 180 C. The CCDs were attached to the ladders by adhesive pads, and wire bonded from each end of the CCD to gold-plated pads (electroless gold with an electroless nickel barrier) on the traces of the flex-circuits. Fig. 7 shows an exploded view of the ladder assembly and Table 3 lists the contributions to the ladder thickness of 0.40% X_0 in the active volume of the detector.

2.2.1 Mechanical/Thermal Details

The arrangement of the beryllium substrate sandwiched between two essentially identical flex-circuits provided a balanced structure with respect to thermally induced stresses. Note the presence of the dummy traces (Fig. 7) which optimally balanced the real traces carrying connections to the inner end of the CCD on the opposite face of the ladder. On cooldown, the thin kapton/copper skins were obliged to follow the contraction of the beryllium substrate. Internal forces within this assembly were considerable, but the tendency to bowing was minimized by the balanced structure.

However, this balance was in principle disturbed by the attachment of the CCDs on either side of the ladder, one at each end. Furthermore, the CCD elastic modulus was too great to allow a rigid connection to the beryllium, given the considerable difference in thermal expansion coefficient between these materials. This situation was dealt with by attaching the CCDs with eight relatively thick (200 μm) pads of NuSil, a phenyl silicone adhesive [16]. This adhesive, having a low elastic modulus and a very low glass transition temperature (well below 170 K), provided a compliant bond which ensured excellent mechanical stability with low shear stress for repeated thermal cycling of the detector. Thus on cooldown, the ladder motherboard felt very little stress arising from the CCDs attached to it.

There were two sources of potentially damaging thermal stress which needed to be considered. Firstly, the sandwich described above would never be perfectly balanced. Measurement of unsupported samples cooled to 200 K revealed shape changes which could be described by a quite uniform change of curvature (bow) along the length of the ladder of magnitude in the range 50 to 100 μm . This of course would be quite unacceptable for any assembly required to achieve tracking precision of a few microns. However, the ladders in the assembled detector were not free, but attached at each end to the essentially rigid support structure, via the ladder block/annulus block assemblies sketched in Fig. 8. These blocks were themselves attached to their substrates with thin layers of adhesive, and their rigidity in this assembly was such that a ladder which would bow by 100 μm when unsupported, would have its sagitta reduced to 0.8 μm when attached to these block assemblies. Thus the tendencies of the unsupported ladder substrates to flex on cooldown were suppressed extremely effectively in the overall support

structure. However, these attachment blocks, while solving one potential thermal problem, tended to create a different one. The support structure can be considered to be a homogeneous beryllium assembly. However, the operating temperature of the overall detector is by no means uniform. Due to the thermal dissipation on the CCDs, the ladders are slightly warmer than the parent structure, and the outermost ladders are warmer than the inner ones due to the flow direction of the coolant gas. A minute excess expansion of a rigidly clamped ladder relative to its supports would induce massive bowing. This problem was avoided by allowing one degree of freedom (longitudinal sliding) between one pair of ladder/annulus blocks (at the south end) for each ladder (see Fig. 8). Including this low-friction sliding joint in the overall assembly guaranteed extremely good shape stability of the ladder substrates under thermal cycling. Stress relief between the dissimilar materials used in this assembly (ceramic blocks on beryllium substrates) was again achieved by controlled-thickness layers of the phenyl silicone adhesive mentioned above. In this situation an adhesive layer of 100 μm thickness was sufficient to provide the necessary stress relief.

For the CCDs themselves, they could be expected to track precisely the movement of their ladder substrate, as long as their eight adhesive pads were of uniform diameter (~ 1 mm) and thickness. This was achieved by a collaboration between the SLD group and the CCD manufacturer. One of the most important decisions for the success of this project was how to handle the interface between these two groups. The original contract called for the delivery to SLAC of unpackaged devices. It was however apparent that the issues of static damage, wire bond quality control, and mechanical damage could all cause problems in this scenario. It was therefore decided to switch to an arrangement where EEV were supplied with prepared ladder motherboards, and their work continued through the die attach, wire bonding and final room temperature electrical testing of mounted CCDs, as discussed in Section 3. These are areas in which their experience was essential in achieving a reliable product.

In order to accommodate the die attachment within the standard EEV operation procedures, it was necessary to develop a very well defined and straightforward assembly procedure. This was done by the SLD group, and the outcome was a set of jigs which could be used to carry out the following operations:

- | | |
|--|--|
| precisely hold ladder substrate | (vacuum hold-down) |
| dispense eight controlled quantities
of adhesive in precise locations | (automated adhesive dispenser, plus templates) |
| set CCD1 in position | (spacers, template, weight for flat hold-down) |
| after curing, wire bond CCD1 | (standard deep access K&S bonder used for
routine CCD wire bonding) |

repeat with ladder inverted for CCD2

These operations were backed up by QC procedures to verify the adhesive quality and quantity, to tune up and check the wire bonder settings, etc.

In order to meet the schedule, as many as eight assembly jigs were in use at one time, giving a throughput capacity of eight ladders per week. The wafer fab and die preparation facilities at EEV were able to produce tested devices at a rate matched to this. In practice, due to one or two initially low yielding batches, limitations in the supply of ladder motherboards, and a low but finite level of re-work due to CCDs rejected during the SLAC cold tests, the entire production of approximately 50 ladders at EEV occupied a period of 15 weeks.

2.2.2 Electrical Details

Within the active regions of the ladders, the flex-circuits were reduced to the absolute minimal thickness, consisting of one kapton layer plus the copper traces connected at each end of the CCD, plus a kapton cover-layer. Over the outer 1 cm of ladder length, this flex-circuit was overlaid with a copper mesh ground plane, which extended (via the micro-connectors) to the outer stripline and from there to the front-end electronics board. This ground plane was essential in helping with the analogue/drive crosstalk reduction on the striplines, and it was found (not surprisingly) essential that this be continued along the length of the ladder. This was achieved by making a good connection (using conductive epoxy [17]) from the flex-circuit ground to the beryllium ladder substrate, in a region at each end of the ladder (see Fig. 7). Thus the south-end CCD on every ladder had its flex-circuit grounded to the inner surface of the motherboard, and the north-end CCD had its flex-circuit grounded to the outer surface of the motherboard. In order to make these connections, the completed ladders were removed from their transport boxes in a clean room, held carefully while the exposed beryllium surface was etched to improve electrical contact, and bonded with epoxy. After a short curing time under a heat lamp, the ladders were returned to their transport boxes and allowed to cure fully before testing. This epoxy procedure was not performed in the original motherboard fabrication because at that time it had not yet been shown to be necessary. There was some concern that having a common ground for the north and south front-end electronics could cause problems, but this has worked perfectly well.

Throughout the R&D and production phases, the crosstalk between the drive and analogue electronics caused some difficulties. The VXD3 design was in this respect more demanding than that for VXD2; the coupling between these two sets of signal paths all the way from the front-end electronics boards through the stripline/pigtails to the CCDs not surprisingly caused

problems. Achieving R drive pulses which were well balanced and which settled within 50 ns proved quite possible. The difficulty was that these 10 V pulses tended to induce long term ringing on the extremely sensitive analogue signal circuit. This is discussed further in Sections 5 and 7.

2.3 Detector Support Structure

The purpose of the support structure was primarily to hold the 48 ladders in the three-barrel assembly (Figs. 3-5). Beryllium was chosen for this structure (as also for the inner section of beam-pipe, the gas shell and the ladder substrates) because of its high elastic modulus (44.0×10^6 psi) and its long radiation length (35.3 cm). Other options considered were carbon fibre (ruled out due to the anisotropy of a fibre matrix and due to the detector operating temperature) and aluminium (ruled out due to its serious micro-creep characteristics). As already mentioned the ladders (inherently quite flexible) were mechanically stabilized by the system of spring-loaded ceramic vee blocks (Fig. 8), with one end pinned and the other end free to slide in response to slight thermal variations within the detector. When these blocks were bonded with NuSil to the ladders (the 'ladder blocks') and to the beryllium rings or annuli (the 'annulus blocks'), jigs were used to ensure a minimal thickness glue line of $100 \mu\text{m}$. For the annulus blocks, a master ladder was used to locate them within the support structure while they were being bonded to the annuli. This assembly procedure allowed the adhesive to compensate for the buildup of manufacturing tolerances. All 18 parts of the support structure were linked by dowels located in match-reamed holes, clearance $5 \mu\text{m}$, and the mating surfaces were lapped flat to $1.3 \mu\text{m}$, to ensure stress-free repeatable assembly. This procedure relied on all mating surfaces being extremely clean. Naturally, with exposed CCDs in the detector, all assembly procedures were carried out in clean room conditions. However for the assembly of the support structure, particular additional cleaning procedures were necessary. Using mirror-cleaning techniques developed by laser scientists, the mating surfaces were always wiped clean and inspected before any mechanical assembly. Microscopic contaminants missed by the inspection procedure could be relied on to flatten or bed in to the mating surfaces when the relevant screws were tightened to the specified torque levels.

2.4 Detector-related Mechanical Equipment

The equipment inside the inner bore of the CDC (of radius 20 cm, hence the 'R20 module') shown in Fig. 9, has the vertex detector at the centre of a two metre long vacuum chamber (the innermost section of SLC beam-pipe). This vacuum chamber is divided into five segments. The central, double-walled beryllium segment (beam-pipe and gas shell, see Fig. 4) of length 34 cm provides the support to which the detector is attached via a low-stress, three-point kinematic mounting system. This uses three small ceramic spheres in a cone/vee/flat

arrangement to provide a stable support while also absorbing thermal expansion differences between the beam-pipe and support structure. Engagement of the mounting system is preserved by light springs acting between these two structures. The vertex detector was divided into two halves to permit installation onto the R20 module in a clean room adjacent to the SLD detector, without disturbing the various flanges and bellows of the baked-out beam-pipe assembly in close proximity to the detector. Figure 10 shows the VXD3 upper module before mating to the lower module around the SLC beam-pipe.

The central section of beam-pipe is clamped to each end of the CDC inner barrel, as indicated in Fig. 9, by means of a pair of aluminium support cones. This system achieves excellent positional stability of the VXD relative to the CDC through long epochs of standard running conditions of SLC/D. Only between runs (when the end-doors of the detector are opened, disturbing the final focus system) or during runs when movements of the final focus quadrupoles or detector are deliberately made for beam tuning purposes, is there a significant disturbance to this alignment. As discussed in Section 4, the VXD/CDC relative alignment is continuously monitored by means of the detector position monitor system (shown also in Fig. 9).

Even door opening and final focus quadrupole motion has relatively little effect on the VXD alignment, since the sections of the SLC beam-pipe within the R20 region are coupled by compliant bellows systems. In particular the sections of vacuum chamber adjacent to the central beam-pipe are conical and float between two atmospherically balanced ten-convolution bellows. This linkage provides positioning freedom and seismic isolation. However, the heavy conical M3 masks which surround the beam-pipe (Fig. 9) are rigidly cantilevered from the final focus quadrupole cryostats north and south, and are therefore obliged to follow any movements of these systems.

The VXD/CDC link itself is not merely a passive clamp. It was necessary to allow for background tuning, including positioning of the critical M4 masks of radius 20 mm (see Fig. 4) which provide the final defence of the innermost beam-pipe from background. To allow this central section to be moved in x and y , it was linked to the CDC at each end by means of a three-point clamp, two of which were adjustable by remotely operated jacks.

2.5 Cryostat and Cooling System

The main requirements of the cooling system are to establish the desired operating temperature for the detector, and to maintain the CCDs reasonably close to this temperature when all ladders are electrically powered and being read out at full speed.

As discussed in Section 1.2, the operating temperature of a CCD detector for HEP applications is determined by the need to minimize radiation-induced degradation in charge transfer efficiency. Choosing a sufficiently low operating temperature results (under the conditions of background charge present in the detector) in effective suppression of the electron traps induced by displacement damage in the buried channel of the device.

The most critical requirement for the coolant gas is that it be of extremely low humidity. Ice formation on the detector could cause numerous problems. The same system is used as for VXD2, and uses liquid nitrogen from the SLD main supply at 40 psi. On demand, this is used to fill a small primary dewar at atmospheric pressure, in which the newly supplied liquid boils vigorously for a short time, as its temperature falls to its boiling point at atmospheric pressure. The quiescent liquid in the primary dewar is used (on demand of a level sensing system) to top up a tall (approximately 2 metre) secondary dewar mounted at a lower level, which is also open to atmospheric pressure above the liquid level. By maintaining the level in the tall dewar constant to $\pm 5\%$, it is possible to obtain a steady flow of liquid from the base of this dewar via a needle valve set to deliver the appropriate rate of approximately 1 g/s to a two-stage boiler system. This consists of two copper cans mounted in a vacuum enclosure, each filled with mesh for good thermal conduction, and wrapped externally with heating tapes. The average power dissipation of the lower can is continually adjusted so as to boil the liquid and raise the temperature of the gas to 85 K in a hardware-controlled feedback loop. The power dissipation of the second can is similarly controlled by a feedback loop to raise the temperature of the exit gas to a set value of around 135 K.

This cooling system is located outside the SLD detector, and the gas is fed in to the vertex detector using a series of linked stainless steel vacuum-jacketed cryo pipes. To achieve a sufficiently low heat leak, these pipes need to be pumped continuously. Minor vacuum leaks which would create negligible thermal leakage in pipes used for liquid transfer, could create a seriously excessive temperature rise in the slow moving coolant gas, and had to be eliminated. For VXD2, the vacuum jacketed pipes continued all the way to the cryostat. In the case of VXD3, the final section of pipe within the CDC barrel, which encroached slightly within the polar angle aperture of the endcap tracking system, was changed to a very low mass foam-jacketed aluminium pipe. The foam material used (as for the cryostat itself) is a polyisocyanurate material using HCDC 141B as filler gas [18]. This pipe had much worse thermal insulation properties than the vacuum jacketed pipes, but over this short length (approximately 1 m out of 15 m total) this resulted in an acceptable heat load on the inlet coolant gas, simply necessitating a somewhat lower temperature setting for the upper heater can of the supply system than was used for VXD2.

The one metre long inlet pipe to the cryostat (see Fig. 9) is sealed with flexible rubber at each end, thus allowing the vertex detector and cryostat (both of which are firmly attached to the inner section of beam-pipe carrying the M4 masks) to be freely translated by the R20 jack system for beam tuning. The inner (aluminium) section of the inlet pipe enters the cryostat, where it is flexibly coupled to the gas shell with a bellows linkage. The inlet gas thus pressurizes the thin cylindrical volume between the beam-pipe and the gas shell, which is sealed, apart from an array of 72 small holes through which jets of coolant gas emerge radially into the detector volume. This gas streams out over the ladders from Barrel 1 to Barrel 3, between the neighbouring ladders within each barrel, and then out through the holes in the support shell. Only a small percentage emerges through the end-plates of the support structure. Finally, after absorbing additional heat flowing in through the walls of the cryostat, the gas exits through the south cryo pipe system (similar to the inlet pipe system on the north, but with larger diameter and/or triple pipes to minimize the overpressure in the delicate cryostat).

The exit gas is led through an electrical heater on the outer periphery of the detector, returning the temperature to 20 C, after which the flow rate and humidity are monitored before the gas is released locally to the atmosphere.

The six sections of the cryostat (two half end-plates each end, and two half cylinders above and below) are sealed with NuSil adhesive. These sealed joints are penetrated by 48 striplines each end and by the beam-pipe. The joints incorporate heater wires near their outer edges, used to maintain them at a sufficient temperature to guarantee freedom from condensation on the outer surface of the cryostat. This is easily achieved, since the R20 volume is moderately well sealed at each end with plastic and adhesive tape, and is blown with dry nitrogen at room temperature. Hence the dew point in this enclosure is normally extremely low (around -20 C). The largest heat leak into the cryostat is due to the beam-pipe (which changes from beryllium to thin stainless steel at a brazed joint at $|z| \approx 15$ cm). By using extra turns of heater wire round the beam-pipe, its temperature at the outer wall of the cryostat can be maintained at approximately 15 C, entirely adequate to avoid condensation or damage to adjacent delicate equipment, a wire beam-position monitor on the south end, and a laser beam-size monitor on the north (Fig. 9).

The cryostat assembly also provides a re-entrant Faraday cage system (see Fig. 4) that screens the detector from RF associated with the beam current, any other signals coupled to the beam-pipe, or pickup from other electronics within the SLD enclosure. It should be noted that, as regards external sources of electrical interference, the heart of SLD is much quieter than a typical unshielded laboratory; the magnet and the 4000 ton iron flux return provides a highly effective electromagnetic shield. The elements on the Faraday cage begin with the beryllium gas shell which is isolated electrically from the beam-pipe. The gas shell is connected at each end by a series of aluminium fingers (complete in azimuth) to the end-plates of the cryostat,

which are faced with 0.5 mm thick aluminium. The upper and lower sections of these plates are electrically connected by thin copper jumpers. Similarly, the upper and lower cryostat barrel sections are faced with thin aluminium sheet and electrically linked together with compressible conductive tubing along the length of the barrel. Wires from these three sections of the Faraday cage penetrate the cryostat axially at each end, resulting in a ring of wires from the end-plates inside the radius of the emerging strip-lines, and another ring of wires from the barrel section outside this area. The striplines themselves are sheathed in grounded mesh screens, to which these wires are connected. In this way the Faraday cage is extended to the front-end electronics (see Fig. 9), where a more or less complete set of plates continues the shielding around the electronics. These connections continue on sheaths around the power supply cables (one per front-end board) leading to the power supplies on the periphery of SLD. Those on the south are left floating at the power supply end, whereas those on the north are grounded to the steel frame of the counting house, which is a solid external detector ground. Care was taken to ensure that the Faraday cage system is electrically floating other than for this single point ground connection.

3 CCD Design and Production

The development of the CCDs for the detector was initiated by informal discussions with a number of manufacturers to establish reasonable goals for all the key parameters, based on their experience with manufacturing scientific grade devices. This was followed by a detailed design specification, care being taken not to write in requirements that were only available from one or two manufacturers. In parallel with this, interested manufacturers were requested to supply 'engineering grade' devices for evaluation. This involved a considerable amount of testing by the SLD Collaboration, and also demanded a sufficiently flexible drive and readout system to be able to handle device types quite different from those eventually required. As it turned out, this testing period was extremely important since it revealed several unfamiliar types of defects, as a result of which the detailed specification was considerably tightened.

The Request for Production (RFP) divided the project into two phases. The first phase (prototype design and manufacture) resulted in a small number of devices for detailed testing. This was followed by the production phase of approximately 100 devices. In order to avoid the high charges made by all CCD manufacturers for performing detailed low temperature tests on their devices, the RFP specified that the devices should only have been screened for reasonable room temperature images before delivery. This streamlined test procedure minimized the cost of manufacture, but the RFP included the condition that devices found not to satisfy the stringent cold test requirements should be replaced without cost. In practice, the failure rate was low, so this procedure worked out extremely well. Of the 12 manufacturers originally approached, four produced estimates for an early Budgetary Inquiry, and three remained to bid against the RFP. The award of the contract was driven less by the cost of design and prototyping than by the quoted production costs, which varied by a factor of two among the different manufacturers.

3.1 Design Specifications

The RFP specified the design requirements for the imaging area, readout area, output amplifier, wire bonding and radiation hardness. It also requested the bidders to respond with their own estimates of the critical parameters for their particular fabrication processes. This response (in the case of the successful bidder, EEV) evolved into the detailed CCD design specification during the prototype design process. It was pleasing to discover that in almost all cases the parameters that emerged from detailed design studies were superior to those specified in the original RFP.

The buried channel CCD was manufactured on a p^+ substrate of resistivity less than 20 m Ω cm (for short carrier lifetime). This was covered by a 20 μ m thick, p -type epitaxial

layer of resistivity $20 \Omega \text{ cm}$, with charge collection into an n -channel by a combination of drift and diffusion. With $20 \times 20 \mu\text{m}^2$ pixels, the effective detector element is a cube of volume $(20 \mu\text{m})^3$, the p/p^+ interface providing a 100% reflective surface for signal electrons diffusing within the epitaxial layer [8]. The $650 \mu\text{m}$ thick inactive p^+ substrate was mechanically lapped down to $180 \pm 20 \mu\text{m}$ overall device thickness by the manufacturer after dicing, in order to achieve an adequately low degradation in tracking precision due to multiple scattering.

Other than the bond pads and fiducial marks (described later) the top CCD surface was protected in the final stage of processing by a polyimide layer of thickness approximately $2 \mu\text{m}$.

The CCD type number was CCD32, the general layout of which is shown in Fig. 11. This design is of conventional stitched format with six 'middle' sections, and the readout circuitry in two 'end' sections. The image section bus lines (see Fig. 12) are continuous along the full length of the array with the opposite directions of charge transfer being achieved by changing the polysilicon electrodes sequence. The image sections have a 3-phase clock configuration, whereas the readout registers at each end are of 2-phase configuration for high speed clocking with minimal feedthrough to the analogue outputs. Standard anti-static gate protection devices were incorporated.

In each corner of the array the pitch of the outermost 100 columns (2 mm wide) tapers down from $20 \mu\text{m}$ to $18 \mu\text{m}$ over a length of 1 mm (50 rows) to give room for the charge detection amplifiers and six blank register run-off elements, as shown schematically in Fig. 12. With a single 'dump' column provided to drain off any peripheral-generated charge, in the main part of the array the distance between the edge of the outermost pixel and the substrate ring is $215 \mu\text{m}$. By positioning the saw cut off-centre in the $200 \mu\text{m}$ wide 'scribe channel', the specified edge spacing of less than $300 \mu\text{m}$ from the edge of the sensitive area to the edge of the die was achieved in all cases.

The design of the output amplifier is shown schematically in Fig. 13. A two-stage direct-coupled circuit is employed. The node capacitance is minimized to achieve highest responsivity. The various circuit capacitances are adjusted to give both sufficient bandwidth and adequate damping with the system load capacitance C_L of 40 pF . The second-stage load resistor R3 is off-chip, on the front-end electronics board, to minimize power dissipation inside the detector cryostat. For the production CCDs, the first and second stage transistors were both surface channel devices.

Fiducial marks, in the form of a $56\ \mu\text{m}$ squares of metallization, were provided at intervals of 3 mm horizontally and 2.5 mm vertically within the sensitive area of the device. Windows were left in the polyimide surface passivation in order to allow a clear view of the metal edges, free of any distortions that could be induced by a transparent cover layer of uncertain topography.

The performance parameters estimated by the manufacturer on the basis of simulations, or measurements of similar devices, are summarized in Table 4. Some key features of the three most important sections are as follows.

3.1.1 Imaging Area

The total well capacity of the pixel is $3.5 \times 10^5 e^-$, more than adequate for min-I and all probable background hits. With drive pulse connections in all four corners of the chip, imaging area transfer rates well in excess of 200 kHz (pixel to pixel) could be achieved, more than adequate for any fast clear requirements in this detector. The associated power dissipation on chip (for 10 V clock drive pulses) is 1.3 W for continuous 200 kHz operation, approximately half of the originally specified maximum.

3.1.2 Readout Register

The charge storage capacity for the 2-phase readout register is $4 \times 10^5 e^-$. With 10 MHz drive pulses, 20 ns rise and fall times, there is negligible waveform degradation along the length of the register bus. With 10 V clock amplitude, the on-chip power dissipation is 25 mW per register (with two registers per CCD), 10 times lower than in the original specification.

3.1.3 Output Amplifier

The CCD output node capacitance is approximately 40 fF which, together with a voltage gain of 1.75, gives an overall responsivity of about $3\ \mu\text{V}/\text{electron}$. The output impedance is approximately $260\ \Omega$. Simulations indicated the output signal to be overshoot-free and with entirely adequate settling time for a 10 MHz data rate. The power dissipation on chip for each channel is about 45 mW. The output noise spectra are discussed in Section 7, in conjunction with other options. The simulated sampled noise for correlated double-sampling operation [19] at a readout frequency of 10 MHz is about $100 e^-$, entirely adequate for fully efficient min-I detection.

3.2 Prototype Evaluation

The CCDs coming out of wafer fab at EEV were subject to a number of tests. First they were DC tested at wafer level, mainly to weed out cases of inter-gate or gate-to-substrate short circuits. Then they were 'image-probed', still at wafer level. This procedure consisted of making contact to the bond pads at one end of the CCD using a probe card, projecting an image consisting of a grid of fine bright lines on a dark background onto the CCD surface, slightly skewed to the pixel array, reading out the image from the two quadrants at that end of the device, and carefully inspecting the image on a TV monitor. Subsequently the wafer would be turned through 180° and the tests repeated for the other half of each CCD. Devices were rejected if any blemishes (hot spots, dead columns, slipped columns etc) were observed in these room temperature images.

The DC tests and image probe tests were repeated after thinning and dicing (now at the die level). Accepted dice were then fitted to the ladder motherboard as described in Section 2.2 and subjected to additional tests. Each ladder-mounted CCD was again DC tested before being connected to a burn-in rig in which all the DC testable features (such as gate isolation) were subjected to a prolonged burn-in at elevated voltages, and screened for small leakage currents which would lead to rejection of the device. The accepted devices on ladders were then given a final image test (with electrical connections being made via the ladder pigtails, no longer using probe cards, of course) before being delivered to SLAC.

The ladders were sent to SLAC in 'ladder boxes' which served the purpose of mechanical support, dust protection, and electrostatic protection. These were provided with mylar windows to allow irradiation of the CCDs each side of the ladder with 5.9 keV X-rays from a ⁵⁵Fe X-ray source. The ladder pigtails penetrated the ends of the ladder box, allowing electrical connections to be made at SLAC for testing the detector, without needing to expose the ladders to the atmosphere. Care was taken in handling the electrical connections to avoid any static discharge which could damage the CCDs.

Each ladder inside its ladder box, on arrival in SLAC, was installed in a light-tight, insulated anti-static plastic and foam cryostat in which it could be cooled by a combination of liquid and gaseous nitrogen, then maintained at an operating temperature of 220 K. Control and monitoring of the cryogenic system was provided via CAMAC modules interfaced to a microcomputer. The pigtails were connected using commercial micro-connectors to striplines inside the cryostat. A movable X-ray source permitted studies on the CCDs in the presence or absence of signals. A fast local data acquisition system (similar to that used for SLD) provided very high statistics studies of each CCD. This system permitted measurements at the detector operating temperature, where the dark current was reduced to a negligible level, and where the

system was consequently sensitive to blemishes that were orders of magnitude smaller than those detectable with the room temperature studies carried out at EEV.

Measurements carried out included the readout noise of each CCD channel, with the device clocked and unclocked. If this was satisfactory, the main additional test consisted of readout of blank frames (checking for hot spots etc) and of frames populated with X-ray hits. By accumulating a large number of frames, it was possible rapidly to generate a data file with a very high hit density (approximately 10^6 hits for each of the four channels on the CCD) and hence to search for anomalies in the response within individual columns of the device. Since the charge transfer down each column is independent, this level of checking is necessary if one is to screen for local anomalies in the CCD leading (for example) to a partly or completely blocked column. These very detailed tests, developed during the prototype phase, were then followed for all the production devices.

For the most part, the detailed CCD testing confirmed what had been found in previous experiments, namely that a careful room temperature series of tests can screen out the vast majority of devices that will give noticeable performance problems with small signals at low temperature. This is a quite remarkable fact, since at room temperature one is dealing with devices close to saturation in terms of dark current, and with signal levels at least 100 times larger than are relevant for min-I detection. However, the cold tests were by no means redundant, for a number of reasons.

Firstly, they provided the first measurements of the overall system performance at the operating temperature. By measuring the noise with the devices unclocked, it was found that the noise of the CCD output circuit was generally within specifications, i.e. entirely adequate. However, the crosstalk (not on the CCD itself) between the R drive pulses and the analogue output could raise the clocked noise to dangerous levels. This is discussed further in Sections 5 and 7.

The second important outcome of the cold tests was that they revealed a low level problem in some more-or-less isolated regions of the imaging area, on a small fraction of the devices. What was found was a finite probability of a *potential trap*, usually localized to one pixel within the column, of depth sufficient to swallow part or all of the ^{55}Fe signal charge of approximately $1600 e^-$. The effect observed was that signals from 'below' the position of the trap (i.e. closer to the R register) would show a normal cluster charge distribution, while those from 'above' the trap position would be degraded or completely missing. Since such an effect would require hundreds of hits per column to be clearly seen, the overall data sample required per CCD was in excess of 10^6 ^{55}Fe hits. Fortunately, as described above, the fast data acquisition system was able to accumulate the necessary statistics in a short period of time. The diagnosis of these traps was not made during the prototyping stage. The true

understanding had to wait till late in the production stage, when one rogue batch of devices was seriously plagued with this problem, making definitive diagnosis possible, as discussed in the Section 3.3.

The fact that the full custom design of these complex devices by the manufacturer was correct at the first attempt is a very good indication that even on a tight timescale, customers needing specialized CCD designs can find expertise matched to this exacting task. The simulation tools available for such design work are extremely powerful.

3.3 Results from Production Phase

The detailed tests developed for the prototype ladders, and described above, were also used to evaluate the production phase ladders, in a setup which simulated the actual operating conditions of the final detector. The CCDs and their associated electronics were tested for gain, noise, and their response to external sources of ionizing radiation.

The complete test procedure was designed to measure the performance of all devices with respect to the specifications in the production contract. This included three distinct detector configurations: warm (room-temperature) tests of CCD response to an electronic calibration signal using charge injected onto the output node via the reset drain bias line; cold tests without a radioactive source to measure detector noise and identify local current sources on the device due to defects in the CCD; and cold tests with radioactive sources to measure the CCD response to both monoenergetic X-rays and min-I particles. All quantitative analyses of these tests were performed in the data acquisition hardware to maximize speed, minimize data transfers and simplify operation. The total time for a standard series of tests, including cooldown and warm-up of the detector, and full analysis of the test results, was about five hours per ladder.

Warm tests were performed first on all devices using a 100 mV reference signal to measure the gain of the CCD output circuit and to identify any dead channels due to broken wire bonds or ladder traces. Any device containing dead channels was rejected; only one of 52 production ladders (comprising 2 of 416 total channels tested) was rejected for this reason. Devices were then cooled to detector operating temperature and read out after an integration time of 1 s, to look for individual pixels and larger areas which had consistently high charge generation ('hot spots'). From these tests, it was found that a typical device had a few (0.3) small ($>300 e^-$) hot spots per channel. Devices with more than three pixels each generating above $1200 e^-$ of charge per readout cycle, or any single pixel or cluster generating more than $5000 e^-$ of charge, were rejected. Two channels exceeded this criterion and those two CCDs were rejected. The magnitude of the spurious charge generation generally diminished with lower

temperatures, though the two devices which were rejected still had large hot spots at 200 K. All hot spots identified in these tests were masked off for analysis of subsequent tests.

The noise of the devices was then measured both with and without R clock drive signals; any device with greater than $65 e^-$ rms unlocked noise was rejected. The mean unlocked noise signal was $(41 \pm 5) e^-$ averaged over all channels. One channel had an rms unlocked noise of $90 e^-$ and that device was rejected. It was found that the rms noise was generally insensitive to detector temperature in the operating region (190-230 K). The clocked noise as measured in the test setup was typically larger by a factor of 1.5, with most of the increase due to R clock signals feeding through to the CCD output signals on the striplines. Performance in the actual detector has been somewhat better.

Devices which passed initial warm and cold tests were then subjected to intensive testing with radioactive sources. The chief concerns here were that the CCDs transfer charge out efficiently from their entire active region and that there were no local defects which could trap charge. Several instances of the latter were discovered in prototype testing. To screen for these charge traps in production devices, a high statistics sample of source-test data (~ 1 hit/pixel) was generated for every channel. A strong ^{55}Fe source was placed close to the CCDs, and 400 full-frame exposures of one second each were taken in order to achieve the desired precision. The source geometry was chosen to ensure relatively uniform illumination of the CCDs. For fully acceptable performance, the observed number of signals from each row and column of the device needed to be constant within statistics. It was possible with this setup to measure inefficiencies at a level below 0.1% per channel. Four devices with an average inefficiency greater than 0.5% per CCD were rejected. The mean inefficiency averaged over all devices was 0.1% per channel.

A small amount of charge can be lost in each clock cycle due to inefficiencies in charge transfer. Over the course of 2000 row and 400 column advances a very small loss in individual transfers can result in a significant integrated signal loss. The efficiencies of individual I clock and R clock transfers were measured by calculating the peak position of the ^{55}Fe signal (dominated by a monoenergetic 5.9 keV photon) as a function of R and I address, and determining the slope of the peak position versus distance from the output node. For an accepted channel, it was required that the peak position should not move by more than 5% over the entire range of R addresses and 10% over the entire range of I addresses. One defective channel had a 20% loss of charge in the R register transfer. The mean charge-transfer inefficiency of all channels was about 1×10^{-5} per transfer, or 0.4% (2%) over the entire R (I) register range.

Using the measured peak position of the ^{55}Fe signal and the output circuit gain measured from the charge injection signal, it was possible to calibrate the detectors for output responsivity in μV per electron collected, and electrons collected per ADC count. The mean measured responsivity was $2.9 \pm 0.2 \mu\text{V}$ per electron, and the mean measured conversion factor was 27 ± 5 electrons per ADC count. A typical min-I particle traversing the device generates $1200 e^-$, or about 40 ADC counts on average; the long tails of the charge deposition distribution are accommodated by the full range of the 8-bit ADC. When compared to the average rms clocked noise discussed above, the average signal/noise achieved in the test lab was around 20:1.

Late on in the production phase, one batch of CCDs were encountered for which the incidence of column traps was unusually severe. In some cases, several neighbouring columns were blocked at the same address. Figure 14 shows examples of a few bad columns on one CCD. The sub-standard devices were replaced by the usual procedure (return from SLAC to EEV for re-work). A few of the worst devices were then taken to RAL for detailed microscopic examination in the bad regions, which had been pinpointed in the SLAC cold tests. This information was essential, because even at relatively high magnification the imaging area looked perfectly normal. However, by inspecting the bad regions at the highest magnification, the cause of the problem became immediately apparent (see Fig. 15). Due to localized over-etching, the nominal *overlap* between gates was transformed into a ragged *gap*. Such a feature would then expose the buried channel to fixed charges in the gate oxide, interface oxide or polyimide passivation. This charge could of course create a local disturbance to the channel potential, sufficient to block or trap small signals. Due to the effective filling in of these disturbances by background charge (from dark current) at room temperature, the image testing at EEV was inevitably unable to discover these defects; room temperature images from these devices were perfect.

The fact that CCDs normally work so well in terms of charge transfer of signals as small as a few electrons is due, not to some near-magical uniformity of the channel potential, but to the Poole-Frenkel effect [20] which effectively drags electrons out of small irregularities between the gates of either the imaging or readout register. It is in fact quite reassuring to observe the scale of device imperfections that do cause problems for small signal operation; such processing faults can certainly be addressed and large eliminated. In this particular case, EEV recognised the problem as being associated with a recent change from negative to positive resist for certain of their processing steps. There will inevitably be much smaller potential variations (e.g. due to slight fluctuations in the thickness of the gate oxide) on all devices, but these are evidently well below the threshold required to cause any CTE loss.

In summary, for the total production run of 52 ladders, nine were rejected for not meeting detector specifications. One had broken wire bonds on arrival (broken during shipping); two had significant hot spots; one had high rms noise; and five had significant inefficiencies from charge transfer imperfections or device defects. All of these devices were replaced by the manufacturer by removing the problematic CCDs and replacing them with new ones. These devices all passed re-testing and most are in use in the final detector. All of the devices rejected for efficiency problems came from a single production batch; several other devices from this batch are in use in the final detector. In addition, four devices which passed all electronics tests subsequently had wire bonds broken during the optical survey or assembly phase. These were re-bonded by the manufacturer, re-tested and assembled into the final detector without incident. Five production ladders were re-tested after optical survey to ensure that they had not been damaged in transit, and to verify the reproducibility of the test procedures. All five devices passed re-testing with results very similar to their initial tests.

Most of these tests were repeated in the clean room adjacent to SLD after detector final assembly, and again after detector installation into SLD, with very reproducible results. It was not possible to perform high-statistics tests of CCD efficiency in these locations due to the extra material in the cryostat and the much larger distance between the radioactive source and the CCDs. However, it was possible to verify that each channel was live and could transfer charge with reasonable efficiency.

3.4 Radiation Damage Tests

As well as four full-scale devices, each 5 inch wafer contained two small-scale CCDs of the same design, merely shorter in the I direction, which were available for test purposes. Such devices, from six of eight production batches, were used at SLAC to qualify those batches for radiation hardness. They were operated in the standard test setup before and after dosing with a strong ^{60}Co source, and acquired total doses of 10-15 krad per device. The CCDs were held at normal bias voltages throughout irradiation to mimic actual operating conditions.

The loss of gain of the CCDs due to radiation damage was measured, using the change in the peak position of the ^{55}Fe calibration source. A roughly linear loss of gain in the devices with accumulated dose was found, corresponding to 0.1-2.0% per krad, with a median value of about 0.6% per krad. The radiation-induced shift in effective bias voltages in the CCD was also measured, by monitoring the threshold value of the reset gate voltage ϕR for transmission of the test pulse on the VRD line. A shift of approximately 100 mV/krad was found in all test devices; these performance figures were within specification.

Finally, the CCDs were checked for the development of new hot spots due to radiation-induced defects in the silicon. One of the six devices tested developed several additional small ($300 e^-$) hot spots during dosing, but an equivalent increase in rms noise on a production CCD used in the detector would still be well within specifications.

4 Optical Survey and Detector Position Monitoring

An optical survey of VXD3 was performed in order to determine the internal geometry of the detector to an accuracy sufficient for beam-related tracks to be used for final alignment. A secondary goal of the survey was to measure aspects of the geometry which would be difficult to determine accurately from tracks, such as the complex shapes of the CCDs themselves and the gravitational sag of the ladders. These two aspects of the project were important factors in achieving the overall precision required.

All ladders (48) and all barrels (3) were surveyed. The ladder survey determined the geometry of each CCD surface and the relation of the two CCDs to each other, with a precision of generally $< 5 \mu\text{m}$, while the barrel survey fixed the position of the ladders to each other. The precision requirement of the survey was to determine the global geometry to $\sim 20 \mu\text{m}$ (1 pixel) rms. At this level of alignment the finding of tracks is straightforward. However, in order to achieve the performance required for the physics objectives, the ultimate spatial precision must be approximately $4 \mu\text{m}$. Tracking-based alignment procedures discussed in Section 6.3 are used to achieve the ultimate level of precision.

The OMIS II programmable optical coordinate measuring machine (CMM) [21] was used for the survey. The CMM had an aperture of $12 \times 5 \times 8$ inches, which easily accommodated both the ladder and barrel surveys, and had a nominal precision of a few μm in x and y (horizontal plane of the CMM) and $\sim 15 \mu\text{m}$ in z (vertical coordinate of CMM). The x - y scale calibration of the CMM was monitored throughout the survey process by means of a precision glass scale. A precision stepper gauge was used for checking the z calibration. Both standards were NIST-traceable and were cross checked with other standards.

Objects could be illuminated in the CMM by axial lighting, ring lighting or back lighting. It was determined that axial lighting yielded the best results and was therefore used for the measurement of most features. Back lighting was used for the profile measurements of the ladders. A variety of survey algorithms was available to measure different aspects of features. For example, an area tool was used to determine the z -coordinates and a line-scan tool was used to precisely determine the edges of sharply defined features. Simple geometric constructions, such as the intersection of two lines and circle fittings, were available.

Separate programs were written to measure the ladders and the barrels. The initialization of each program was performed on representative features requiring the most precision. In this way only four programs were needed for the ladder survey, one for each view defined by rotations of $n \pi / 2$ about the long axis of the ladder. Symmetry was employed in the barrel survey so that only 19 programs were needed for the 48 ladders.

4.1 Ladder Survey

Ladders were surveyed at room temperature in survey boxes which supported them from their ladder blocks, in precisely the same manner as in the VXD3 support structure. The survey jig was outfitted with 6 precision tooling balls which defined an internal coordinate system. Fig. 16 shows a ladder mounted on one of the survey jigs, seen in turn in each of the four views used for measurement. All the tooling balls are clearly visible in every view. The ladder geometry was reconstructed by relating three of the views to the 'standard view' defined by the ladder orientation where the north CCD was visible and approximately in the CMM xy plane, i.e. the view seen when looking at ladders mounted in a barrel. Roughly 6 hours were needed to measure a ladder.

Since the survey was not performed in a clean room, the CCDs had to be protected by enclosing the survey jig in a box with high quality optical glass windows, which had to be coated to reduce reflections when the axial illumination was used. The distortions of the glass windows were found to be small but were corrected by calibration and data analysis. In all, four survey boxes were fabricated although only three were used for production surveying.

A survey box was located on the CMM by means of 3 tooling balls mounted on the surface of each side of the box corresponding to each of the 4 views of the ladder survey. These tooling balls mated to grooves milled in an aluminium support plate which was rigidly attached to the xy table of the CMM.

A number of features were measured in the ladder survey in addition to the 6 tooling balls used to define the ladder local coordinate system. Each CCD had 186 fiducial marks which were very accurately positioned in the CCD structure. The fiducials were arranged in 6 columns every 3.000 mm along the short dimension of the CCD and in 31 rows every 2.500 mm along the long dimension. The fiducials were $56 \times 56 \mu\text{m}^2$, aluminium pads $1 \mu\text{m}$ thick and located $12 \mu\text{m}$ above the midplane of the epitaxial layer. However, with the characteristics of the OMIS II lighting and edge finding tool, the edges of the window in the polyimide passivation layer, surrounding each fiducial pad, were actually measured. This window had a dimension of $80 \times 80 \mu\text{m}^2$, and was quite accurately determined. All of the columns of fiducials and every other row of fiducials were measured for a total of 96 features. The flex-circuits bonded to the ladder substrate carried similar fiducial marks, providing a reference to the half of the ladder not occupied by a CCD, on each side. All these fiducials (3 columns \times 14 rows) were measured, albeit with less precision than those on the CCD surfaces. These were useful in mating the ladder survey data with the barrel survey data. In addition to the fiducials, the physical corners of the CCDs were measured. Although less precise than the fiducials, the

corners of the silicon die provided helpful navigation points. Finally, the profiles of each CCD were measured in the two edge-on views, using back lighting. These data provided useful cross checks to the face-on views and allowed the gravitational sag to be measured. Fig. 17 shows the survey data in top and side view of a typical ladder.

4.2 Barrel Survey

After the ladders of a given barrel had been measured, they were assembled into the VXD3 support structure in their final position and surveyed again to determine the barrel geometry. A set of 32 precision tooling balls mounted on the beryllium support structure (16 at each end) defined the barrel survey coordinate system (see Fig. 18). ScotchLite [22] was placed under each tooling ball to allow axial lighting to illuminate the circumference. As in the ladder survey, the barrel survey was not conducted in a clean room. Again, the detector was housed in a sealed enclosure (the barrel survey box). The support structure was mounted kinematically on a dummy beam-pipe just as in the final assembly in SLD. This was rotated inside the barrel survey box to allow all the ladders of a layer to be surveyed in turn. The window of the barrel survey box was made of optical glass, coated in order to reduce the reflections of the axial lighting needed for the survey. The barrel survey was carried out at room temperature in a dry nitrogen atmosphere.

The outside surface of each ladder of a layer was measured with each ladder approximately positioned in the CMM xy plane. There were 12, 16, and 20 ladder settings measured for layers 1, 2 and 3 respectively. Associated with each azimuthal setting was a measurement of the concomitant set of visible tooling balls. There were typically 4 to 6 measureable balls on each end for each setting with at least 3 balls on each end overlapping between adjacent settings. The barrel geometry was assembled by making the sets of overlapping tooling balls congruent.

Only about 50% of the fiducials on the outside of a ladder (north CCD face up) were visible in the barrel survey since the survey had to be conducted through the shell of the support structure (Fig. 18). Thus, all visible fiducials on the CCD and flex strips were measured as well as one of the physical corners of the CCD die which served as a navigational aid. Given that a smaller number of fiducials was measured than in the ladder survey, each z -coordinate of the visible fiducials was measured 4 times resulting in an improvement in the z resolution. The complete barrel survey of all 3 layers was executed in 10 days (not including programming time). Fig. 19 shows the end and side views of the barrel survey.

4.3 Survey Data Analysis

A number of analysis steps were conducted after the geometry assemblies of the ladder and barrel data.

The 3-dimensional distances of all combinations of the 6 tooling balls were used to determine the distortions of the ladder survey box windows in a χ^2 fitting procedure. In some cases the CMM scales had to be changed by as much as 10^{-3} mm/mm although the net effect was small owing to the small distances involved. The optical distortion of the barrel survey box was determined by measuring the glass scale standard and the stepper gauge through the glass. Distortions of order 3.5×10^{-4} mm/mm were observed. These corrections were verified in the data by using the well-defined separations of the CCD fiducials.

Since only the outer surface of the ladders was measured in the barrel survey, the position of the south CCD had to be derived from mating the ladder data with the barrel data which served as the backbone of the geometry reconstruction. This was accomplished by making the common surfaces of the ladder and barrel surveys congruent. Roughly 20 to 30 features were in common in the ladder-barrel mating. The CCD fiducials were given higher weight in the χ^2 fitting because they were much better determined than the flex-circuit fiducials. The planes of the north and south CCDs were measured face up on the CMM. Hence the gravitational sags for the north and south CCDs were of opposite sign and had to be corrected. By comparing the top and side views of the ladder survey, the gravitational sag of the ladders could be determined. It was found that the average sagitta was roughly $30 \mu\text{m}$. Hence the average correction for the south CCD was a $60 \mu\text{m}$ adjustment to be consistent with the north CCD convention. The measured gravitational sag was then projected by a $\sin\phi$ factor to compute the sag of a ladder in its final setting. Having corrected for the gravitational sag, the south CCD was constrained to the north by using the well-measured separation in the ladder profile views. Typically adjustments of 10 to $20 \mu\text{m}$ had to be made, thereby improving the precision of the z -measurements of the assembled geometry.

VXD3 was surveyed at room temperature, about 80 C warmer than the operating point. Hence the survey data had to be contracted to correspond to the lower operating temperature, by integrating the nonlinear coefficients of thermal expansion for beryllium and silicon. Given the mechanical connection between the CCDs and the substrate, it was clear that to a good approximation, on cooldown the silicon would contract uniformly about its geometric centre with the geometric centre following the movement of the substrate. Overall, the beryllium and silicon contracted by a factor 8.5×10^{-4} and 1.9×10^{-4} respectively. The differential contraction between silicon and beryllium was approximately $53 \mu\text{m}$ over the length of the CCD, inducing a shear in the adhesive pillars, zero in the middle and building up to $26 \mu\text{m}$ at

each end of the CCD. As a check of the shape of the ladder under thermal contraction, the position of the CCD with respect to the beryllium substrate was measured on a sample ladder cooled to the operating temperature. It was found that the CCD did not distort relative to the substrate under cool down; only the average distance of the CCD from the substrate contracted by approximately $5 \mu\text{m}$, due to contraction of the NuSil adhesive pillars.

For preliminary track finding, a simple model of the detector was used, with each CCD being described by the position of a corner pixel and the orientation of the average CCD plane. From Fig. 17, it is apparent that the CCD shapes are in fact complex. The thin silicon tends to curl up, and the discrete constraints (adhesive pads) prevent this only in an average sense. For the CCDs on ladders, the short dimension is concave with a shape given by a quadratic, while the long dimension has an approximately quartic shape. In order to describe these shapes precisely, a fit to the CCD surface was performed with these functional forms and all cross terms using Chebychev polynomials. The resulting fitted shapes of the CCDs of a typical ladder are shown in Fig. 20. These shapes are used to determine the perturbation of the surface from the average plane. The CCD shape correction was particularly important at small polar angles.

4.4 Summary of Detector Survey

The optical survey of VXD3 establishes the geometry to an accuracy of better than 1 pixel ($17 \mu\text{m}$ in $r\phi$ and $14 \mu\text{m}$ in z). It includes the CCD shapes (measured in the ladder survey) and corrections for the gravitational sag of individual ladders, according to their orientation in the assembled detector. This 'survey geometry' forms the basis for determining the best estimate of the true geometry, using tracking data from Z^0 events in SLD.

There are several known distortions to the survey geometry. One is the assembly of the upper and lower modules around the SLD beam-pipe. During the barrel surveys, VXD3 was repeatedly assembled around the dummy beam-pipe, and variations between assemblies of $20\text{--}30 \mu\text{m}$ were observed. Hence the true geometry can be expected to lie at some unknown point within this range. Another known distortion affects (to a small degree) all ladder shapes independently. The precision of the OMIS z -measurements was not sufficient to determine these shapes accurately in the barrel surveys, and for this reason they were assumed to be as in the ladder surveys, where the side-on measurements allowed them to be determined very precisely. However, it is known from independent trials that the transfer of a ladder from the ladder survey jig to the VXD3 support structure causes a small distortion (in the form of an imposed gentle bow or s-shape) due to imperfect matching between the annulus blocks in the two systems. Therefore, the fitting procedure used in the tracking-based alignment must permit small relative displacements of the 48 ladders, and small independent distortions of each of

these ladders, given by two 'bow' parameters, as compared with the survey geometry. Fortunately as discussed in Section 6.3, the SLD data can be used to define a suitably convergent optimization procedure for the determination of the true detector geometry.

In summary, the first essential function of the optical survey was the determination of the overall internal geometry to the precision needed for track finding in a high background environment. The second essential function was the determination of fine structures (CCD placement and shape details) which are stable and which involve too many parameters to be found from tracking. Given these foundations, the tracking-based alignment was well-placed to make the small adjustments in a limited number of parameters required to optimize the geometry to the level needed for all SLD physics.

4.5 Detector Position Monitoring

The precision survey described above resulted in knowledge of the relative positions of the internal VXD3 components at the level of better than 20 microns. The global alignment procedure (Section 6.3) is used to determine the relative alignment between VXD3 and the CDC using charged tracks measured in these detectors. For this purpose the total data sample can be divided into epochs whose boundaries either correspond to known changes in the relative alignment due to controlled movements of VXD3, or are chosen arbitrarily so as to provide epochs containing sufficient numbers of tracks to achieve alignment at the 10 μm level of precision. This procedure results *a priori* in a set of average alignment constants for each epoch and does not resolve any changes of alignment within a particular epoch; any such changes would effectively contribute to a broadening of the resolution of the average alignment.

A system has therefore been set up to allow continuous monitoring of the relative position of VXD3 with respect to the CDC. The aim is to monitor any changes in relative alignment of either a gradual or sudden nature, perhaps due to thermal or mechanical effects, and to provide the necessary information to augment the tracking-based alignment. The system is designed to measure optimally relative displacements of the order of 1–100 μm with around 1 μm resolution, and in addition to track displacements of up to 10 mm.

Two beryllium-copper wires of 250 μm diameter are strung parallel to the beam-line from supports fixed to the endplates of the CDC. The wires are located at roughly ± 60 degrees w.r.t. the vertical axis. Four probes are mounted rigidly to the beam-pipe, two at each end of the VXD3 cryostat, each close to one wire. See Fig. 9 for the general layout, which shows one wire/probe system, rotated azimuthally to the vertical position for clarity. A probe comprises an aluminium housing supporting two circular capacitive plates, of diameter 9.5 mm arranged such that the plates are orthogonal and share a common chord. The probes are

mounted so that one plate is parallel to the horizontal plane, the other is parallel to the vertical plane, and both are parallel to the beam-line. Projected into the plane perpendicular to the wire, the surfaces of the plates thus form a local x (horizontal axis) y (vertical axis) coordinate system. Each wire is strung so as to pass approximately through the geometric centre of the x - y fiducial area formed by the two plates. A schematic of the configuration is shown in Fig. 21.

Both wires are connected to a common ground at signal sources. An excitation signal is applied separately to each plate at a frequency of 15 kHz and the required voltage is monitored continually using a CAMAC-based slow readout system with data transfer to a computer. The probes and signal source units were obtained from the commercial manufacturer [23]. Since each probe and the nearby wire form a system whose capacitance depends on the distance s of the wire from the plates, this distance can be inferred from the measured signal voltages. Since the wires are mounted rigidly w.r.t. the CDC and the probes rigidly w.r.t. VXD3, changes in the monitored voltages indicate relative displacements between VXD3 and the CDC. The configuration shown in Fig. 21 yields four measurements of transverse (x - y) displacements permitting VXD3 global translations, tilts, or rotations about the beam-line, to be determined. The system is not sensitive to relative displacements along the beam-line.

The voltage-distance calibration for each probe was measured prior to installation in the detector. Initially the gain and offset of the amplifier in each signal source were adjusted to obtain a voltage response V in the range $0 < V < 10$ V for a corresponding s in the range $0 < s < 1$ cm. In a test setup a probe was mounted on a granite table and a wire was strung through the fiducial volume. After adjustment of the wire to be parallel to, and just in contact with, both plates, the probe position in the plane transverse to the wire was systematically varied, the displacement from the starting point being measured with a micrometer. At each position the voltages applied to both the horizontal plate, V_y , and the vertical plate, V_x were recorded. In this fashion the response of all probes to wire positions across the local x - y plane was measured. Because of the circular geometry of the two orthogonal plates, V_x and V_y depend on *both* the x and y coordinates of the wire. A convenient parametrisation of V_x and V_y for each of the 4 probes was obtained by fitting an *ad hoc* polynomial function of x and y to the calibration data; an example of $V_y(x,y)$ is shown in Fig. 22. The voltage can be seen to be smoothly-varying but non-linear, especially close to either plate.

Under operating conditions, at any given time i an effective position (x_i, y_i) of a wire w.r.t. a probe is determined from the measured voltages V_x^i, V_y^i by iteratively finding a solution to the parametrised responses:

$$V_x(x_i, y_i) = V_x^i$$

$$V_y(x_i, y_i) = V_y^i$$

Displacements between time i and a chosen reference time are then readily obtained by comparing the two solutions.

During the 1996 SLD run the system performed well and no unexpected displacements were observed. Four controlled movements of the R20 module w.r.t. the CDC provided an opportunity for *in situ* confirmation of the monitoring performance. As an example, the monitored voltage for one plate for the duration of one movement is shown in Fig. 23. In this case numerous systematic position changes of the R20 module were made as part of a study to minimise beam-related backgrounds in the detector; the position changes are clearly tracked by the data.

5 Detector Electronics and Readout

The VXD3 electronic system is required to efficiently identify charge depositions in the CCDs from min-I particles (a broad spectrum, with a peak at $\sim 1200 e^-$), while minimizing electronic noise, crosstalk, and power dissipation both in the CCDs and in the electronics itself. It must operate at average event readout rates up to 2 Hz with acceptable data volume (< 100 kbytes) and no deadtime. Finally, operations must not interfere with other detector subsystems (e.g. through electronic noise or heating) and must be fully integrable into the existing SLD data acquisition and control systems.

These requirements are met by a design consisting of two subsystems: front-end electronics for signal digitization, CCD clocking and bias supply, and a set of FASTBUS modules which manage data acquisition, timing and control functions. Figure 24 shows the general layout of the overall electronics system, and Fig. 25 is a schematic diagram of the full functionality for one channel. The 96 CCDs of VXD3 have 384 analogue signal outputs; these signals are processed on 16 front-end electronics (F/E) boards, mounted close to the vertex detector (see Fig. 9). Their digitized output is transmitted via optical fibres about 50 metres to the FASTBUS data acquisition modules in a counting house situated on top of the SLD detector. All SLD detector subsystems have similar FASTBUS plants. Sparsified and formatted data from the acquisition modules are then promoted to the main SLD data streams and recorded on tape. A brief overview of some details of the implementation is given in the following sections.

The CCD signals, clocks and bias voltages are brought off the ladders with flex-circuit extensions ('pigtailed', Fig. 2), then connected (using commercial micro-connectors) to 50 cm long flex-circuits in the form of microstrip lines. These striplines continue out of the cryostat and connect to small boards which provide termination and connection to the front-end electronics. There are no active components anywhere between the CCDs and the F/E boards since these are completely inaccessible after detector installation. For the same reason, the pigtail-stripline connection has to be completely reliable and was carefully tested before installation to ensure continuity.

The CCD analogue signal is the output of a two-stage source follower with an impedance of about 260Ω . (see Fig. 13). Due to limitations on power dissipation on the detector, this impedance is not matched to that of the striplines ($30 - 50 \Omega$). The input impedance of the F/E board amplifiers is designed to be $3 k\Omega$, effectively an open-end termination of the signal traces. The striplines thus represent a distributed capacitive load of around $40 pF$ whose length (propagation time $\sim 10 ns$) is comparable to the signal risetime, so ringing of the signals is not too serious. On the F/E board, the analogue signals are shaped then amplified with a

gain of about 100, digitized using 8-bit flash ADCs, multiplexed and serialised for optical transmission to FASTBUS.

Along with signal processing and transmission, the F/E boards must also provide all necessary drive pulses and biases for the 6 CCDs connected to each board. To allow sufficient flexibility to adjust for actual operating conditions and channel-to-channel variations (especially considering the inaccessibility of these boards), there are many programmable devices on every board. These provide externally controllable shapes of I clocks, controllable DC offsets for amplifiers, adjustable ADC gate delays, and generation of calibration signals. To make possible initial data loading for all these devices, there is a microcontroller on every F/E board, which communicates with the control data acquisition computer via optical links connected to its DUART port. To reduce crosstalk between different CCDs, striplines and F/E boards, all 16 boards in the system perform data readout with a high degree of synchronization (within 1 ns). To achieve such synchronization, external clocks from one source are used to control operations; these clocks are distributed to all boards via optical fibres. The control signals to start any operation (e.g. start readout of CCDs), which need to be synchronous for all boards, are transmitted by phase modulation of these clocks.

Two custom-built DC power supplies (see Fig. 24), located on the top deck of the detector just outside the counting house, provide adjustable voltages to the F/E boards, along with interlock functions and remote monitoring capabilities. The location of the F/E boards inside the CDC barrel, in a very congested area with very limited access, motivated minimizing power dissipation on the boards by using non-standard voltages. Thus the power supply voltages are all separately regulated on the F/E boards. Interlocks disable F/E board power in the event of interruption to the flow of cooling water, board overtemperature or cryostat overtemperature, to protect both the detector and the electronics from overheating.

The main task for the data acquisition modules is the reduction of the huge volume of raw data (307 Mbytes/event) to a manageable level. Compression of the data is performed in real time in a two-stage process. First, candidates for valid charge depositions ('clusters' from charged particles) are identified by an edge-finding algorithm implemented in programmable logic (the 'cluster processor'). The module's CPU then captures the data from a small region of interest around the identified cluster edge, sorts this data and formats it for the data stream. The acquisition modules also serve as mediators in communications between the central data acquisition computer and the F/E board microcontrollers.

External timing signals are generated by a standard SLD electronics FASTBUS Timing and Control Module (TCM). Such modules are used for every SLD subsystem to generate start and stop signals, derived from the accelerator beam-crossing signal. The delays between output

signals of these modules and beam-crossing time are software controllable. There is also a synchronization module, which contains a crystal oscillator for generating clocks for the F/E boards, and circuits to resynchronize the external timing signals to these clocks. The synchronization module does not have any FASTBUS accessible registers.

5.1 Charge Transfer on CCDs

To move charge to the CCD output node, a set of clocking signals is generated on every F/E board. These signals include 3-phase I clocks and 2-phase R clocks. Charge moves from row to row (through the image area) as a result of the I clock pulses. A 3-phase scheme is used here so that, depending on the sequence of these phases, a charge packet can be moved in either direction. The I clock buses on the CCD are connected in such a way that the standard clock sequence moves charge in rows 1-2000 toward the output register adjacent to row 1, and charge in rows 2001-4000 toward the output register adjacent to row 4000. Charge in the output registers is shifted by means of R clock pulses. Applying R clock pulses always shifts charge in the output registers (R registers) in the direction towards the output node as indicated in Fig. 11. The maximum frequency of R clocks in the VXD3 CCDs was designed to be 10 MHz, and this performance was verified in laboratory tests. However, due to the drive/analogue crosstalk problems already referred to, the system R clock frequency was limited to 5 MHz. Each R register contains 400 pixels linked to the image area and 12 dummy pixels, so at 5 MHz R clock frequency one row readout takes $82.4 \mu\text{s}$. A reset pulse (to reset the CCD output node to its initial level) and a triplet of I clock pulses is generated before each row readout, and this takes about $10 \mu\text{s}$. In one beam-crossing period of SLC (8.3 ms), 81 rows of data are read out, so the entire image is read out in 25 beam-crossings.

All clocks on the F/E board are derived from an external 30 MHz clock, fed onto the board with optical fibre then frequency doubled to provide a 60 MHz board synchronization clock. One of the two R clock phases has a quiescent low level (about -5.0 V), while the other has a quiescent high (+ 7.0 V). During the R transfer, these phases exchange levels. This active period has a duration of two synchronization clock periods ($2 \times 16.66 \text{ ns} = 33.3 \text{ ns}$). The capacitance of the R clock bus is about 400 pF per CCD. To charge this capacitance completely during the active period, the R clock driver impedance needs to be not more than 25Ω . With an R clock amplitude of 12 V the peak current is about 0.5 A. Since each F/E board serves 6 CCDs, and these have synchronized clocks, a pulsed current with about 3 A peak amplitude is generated on the board, on which fast, sensitive amplifiers and ADCs reside. This high current creates significant crosstalk on the board. It is helpful that the pair of R clock phases is very symmetrical, so all currents have their counterparts of the same value but opposite sign. Also, bypassing the R clock drivers with a capacitor between positive and negative power supply inputs, reduces crosstalk. Nevertheless, the R clocks induce significant

oscillation on the amplifier output, which affects the signal measurements as discussed in Section 5.2.

The I clocks are slower, of duration microseconds instead of nanoseconds. However, the capacitance of the I clock buses is much higher because, instead of one output register, the 2000 rows of pixels are driven simultaneously. The capacitance of the I clock bus is about 40 nF, about 100 times the capacitance of the R bus. Considering that about 100 times longer is available for charging this capacitance, the peak current for the I clock drivers is similar. The crosstalk from this current is not an issue because it occurs at the time interval between rows, when no measurements are being made.

To achieve good charge transfer efficiency it is essential to have control over the I clock pulse shapes; this was achieved by using RAM-DAC chips of the type ADV476 as function generators to determine the I clock waveforms. Each chip contains three sets of 256 6-bit memory cells and a 6-bit DAC to generate output signals. Desired waveforms are loaded into the memory (RAM), and then this information is applied to the DAC input.

A readout sequence is started every beam-crossing. The time needed to make a trigger decision in SLD is about 5 ms, while the beam-crossing interval is 8.33 ms. Whenever a beam-crossing is found that does not bear a trigger event, the digitized data are discarded, and a CCD fast clear sequence is activated. In this mode, only I clocks are generated, without intervals for moving charge along the R register. This allows much faster clearing of the CCD image area of accumulated charge (one row shift takes only about $10\ \mu\text{s}$ instead of $100\ \mu\text{s}$ in readout mode). At the end of a fast clear sequence (300 I transfers, stopping about $100\ \mu\text{s}$ before the next beam-crossing), one burst of 412 R clocks is generated to clear out the charge accumulated in the R register.

5.2 Digitization of CCD Output Signals

To optimize the coordinate measurements made with the CCDs, the amplitudes of the CCD output signals are used. The share of the charge received by a particular pixel is determined by the track position, so the centroid of the charge distribution gives a good approximation of the track coordinates. To measure the amount of charge in every pixel, the signals from the CCD output node are digitized. Due to the very small value of the output node capacitance, the CCD sensitivity at the output of the second stage source follower is about $3\ \mu\text{V}/e^-$. The signal generated on average by a min-I particle in the sensitive region of the CCD (thickness $\sim 20\ \mu\text{m}$) is $1200\ e^-$, so a signal of 4 mV has to be measured (and much less for full efficiency, given the fluctuations in charge deposition and the spread between pixels of the cluster). The DC level on the CCD output is of the order of 10 V. The signal is AC coupled

and fed through an amplifier of gain about 100 on the F/E board, in preparation for digitization by the ADC. Flash ADCs of the MP8775 type are used, with a maximum sampling rate of 20 MHz.

As already mentioned, there is a large signal from R clock pulses coupling to the amplifier input, much larger than the true min-I signal. However, since the feedthrough signal is almost identical for every R clock period, sampling the amplifier output signal at the same moment relative to the R clock pulse every period is very effective at suppressing the spurious signal. The difference between successive values should accurately measure the charge transferred onto the output node between measurements. This method of parasitic signal suppression requires highly stable timing between the R clock and ADC gate, and is achieved by using one device (a sequencer) to generate all clocks and gates on the board. The sequencer is based on the programmable gate array, MACH220, and uses the 60 MHz synchronization clocks to generate all logic signals. There is however some residual jitter in the interval between the R clock and ADC gate. This is caused by the fact that the sequencer contains counters required to generate a certain number of R clocks for every row readout, with the result that the currents and potentials inside the chip from one R clock period to another are not completely identical. This can lead to a regular periodic structure in the digitized data, as was observed during electronics development. Maximising the analogue signal shaping time (to ~ 30 ns), and carefully bypassing the power supplies close to the sequencer reduced these effects to tolerable levels.

In contrast to most imaging applications, the low density of signals over the CCD area in the SLD experiment permits an unusual and particularly advantageous mode of readout. It is possible to avoid the usual procedure of resetting the CCD output node to the initial level for every pixel; faster readout is achieved by resetting the node only once per row. There is a very small signal associated with thermal current generation in the CCDs, due to the low operating temperature. The charge accumulated from 400 pixels in one row due to thermal currents is much less than the signal from one min-I particle. In addition to eliminating the need for a high speed reset signal, not resetting the output node every pixel gives the advantage of simplified cluster processing. Because it is possible to sum the signals of as many pixels in the row as one chooses, by subtracting a properly delayed value of the digitized signal amplitude from the current value, the need for multiple summing can be avoided.

5.3 Data Transfer to FASTBUS

Optical links are used to transfer data from the F/E boards to FASTBUS data acquisition modules called LFMs, external to the SLD detector (see Figs. 24 and 25). Every F/E board (of which there are eight at either end of the CDC inner barrel, each serving six half-ladders) is

connected to one LFM. No data reduction is made on the F/E board, as this is a task for the LFM. Eight bits of data come from every A/D channel (with 24 channels per board), generated at a rate of 5 MHz, a total rate of $(8 \times 24) \text{ bits} \times 5 \text{ MHz} = 960 \text{ Mbits/s}$. The fast link circuits use an HP G-link HDMP-1012/1014 chip set and FINISAR optical transmitter and receiver. The G-link circuit uses a data transfer protocol which adds four more bits to every 16 bit 'frame' of the data. So the actual speed of data passing through the data links is 1200 Mbits/s per F/E board. A multiplexer, made of four chips of XILINX programmable gate arrays XC3130A-5 on every board, converts $8 \times 24 = 192$ bits of data generated simultaneously on the board into a serial bit stream. These chips are programmed to act as 16 12 bit shift registers which are loaded with 192 bits of data with 5 MHz frequency, and shift these data to the serial output with 60 MHz frequency. Sixteen such outputs are then fed into each G-link chip, which transforms these 16 input bits plus four additional bits into a single bit output stream. This chip also takes care of the data transfer protocol. Its 1200 Mbit/s output signal then modulates the FINISAR optical transmitter. On the LFM auxiliary card, the FINISAR optical receiver transforms light pulses into electrical signals, fed into the HP G-link receiver. The output of this receiver is 16 bits of data with 60 MHz frequency, which is de-multiplexed by two XILINX XC3130A-3 chips into a 48 bit data stream at 20 MHz, used by the cluster processor and row data memory.

In general, this data transfer protocol works well, though there are some specific problems of receiver-transmitter synchronization, caused by the SLD working environment. There are a few known ways to synchronize the receiver and transmitter clocks, some of which involve additional links between receiver and transmitter for feedback. These were ruled out because extra optical fibre links are expensive, and cable connections to the F/E boards must be avoided in order to maintain electrical decoupling of the detector electronics from the FASTBUS electronics. The data link is operated in the 'simplex' mode (the F/E board sends data to FASTBUS only during data readout), for which there is essentially only one clock synchronization method available. This method uses periods of time when no useful data are being transferred, for sending special frames, called fill frames, for synchronization of the receiver and transmitter oscillators. There is available approximately 3 ms every beam-crossing when the board does not transmit any data (i.e. when it is performing the fast clear of the detector). This time interval is enough to acquire synchronization even from a completely random state. When coarse 'frequency acquisition' is achieved, a fine phase adjustment of the receiver clock is made every data frame, using the additional four bits of data already mentioned. During readout of the trigger events there are no fast clear periods, therefore no fill frames during 25 beam-crossings. In general, this does not cause any problems because the receiver oscillator is already synchronized to the transmitter clocks. However, if some major disruption in the data transfer occurs during readout of trigger events, it may lead to the inability of the data link to recover. In this case a link error signal would be flagged,

preventing data from being recorded by the LFM. On the other hand, until the full amount of data expected is recorded, the LFM would not send to the F/E board the signal permitting fast clear operation. Without this, the transmitter would not be able to re-establish synchronization. In that case, intervention by the data acquisition software is required to correct the problem.

Another problem in maintaining the data link was discovered only after the boards were installed in the R20 module, and high intensity beams were sent through SLD. Such beams generate electrical signals directly on the F/E board, and these fast signals can disrupt synchronization between the transmitter oscillator and the external strobe signal. In this case, the transmitter sends a 'not ready' signal to the receiver, and they start the process of re-establishing synchronization. Fortunately, because total synchronization is not lost, recovery is rapid; usually, the error signal disappears within 3 – 5 μ s. As the CCD readout begins only about 100 μ s after beam-crossing, these short periods of link disruption do not affect data acquisition.

5.4 Data Acquisition and Processing

The VXD3 data acquisition and processing is primarily a job of data compression. An event contains only about 10-20K hit pixels of the potential 307 Mpixels. The large range spanned by these two numbers could cause problems. While the typical performance is nowhere near the worst case scenario, one must be able to handle the occasional rogue event with grace and take care that a misbehaving channel or event cannot monopolize a shared resource. Calculating the necessary resources is an adventure because errors in the estimates are likely to be in the exponent. The architecture chosen to perform this task follows the design of the previous vertex detector VXD2, with only implementation changes due to advances in electronics and experience gained with that detector. The strategy is to examine the incoming pixel stream for 2 row by 2 column areas where the total signal charge exceeds a specified threshold. Addresses of all qualifying areas are written into a common ring buffer. Each raw pixel value is simultaneously written into one of 24 separate channel buffers. The list of qualifying addresses assists the onboard processor in locating and selecting regions of interest from the raw pixel memory. Subsequently, all pixels within a region of interest which exceed a predetermined threshold are written to the output stream.

5.4.1 Data Acquisition Modules (LFM)

A single width FASTBUS slave module (the LFM, see Fig. 25) was designed to perform the data acquisition for VXD3. Each LFM controls and handles the data from one F/E board, which services 6 CCDs, or equivalently, 24 channels. As discussed in Section 5.3, a FASTBUS auxiliary card on the LFM receives the data via a fibre optics cable and performs the

necessary demultiplexing to separate the 24 channels. Major components of the LFM board are:

1. 6 Xilinx chips implementing the real time 2×2 kernel calculation
2. 8 K \times 4 bytes of accept memory
3. 24 \times 256 kbytes of pixel memory implemented in VRAM (1/channel)
4. 1 Motorola 68040LC processor running at 66 MHz
5. 8 Mbytes of general purpose DRAM

In addition to its primary acquisition duties, the LFM also acts as a master for the F/E board. In this capacity it must download the programmable devices on this board, monitor messages from the F/E board and control its data taking modes. The 68040 runs a commercial real time kernel and, when necessary, a debugger/monitor. The remaining software consists of a set of SLD standard acquisition routines and VXD3 specific processing routines.

5.4.2 Data Processing Overview

While the overall data processing scheme is straightforward, a number of factors conspire to complicate the implementation. These complications are driven primarily by two factors: the long readout time of the CCDs and the relatively austere resources available on the LFM.

Since VXD3 operates as a triggered device, the readout modes of the F/E board must be controlled. From beam-crossing time to a predetermined cutoff time all pixels are clocked into the VRAM of the LFM. If no trigger is received by the cutoff time, the F/E board begins clocking the CCDs in FAST CLEAR mode and no further data is sent to the LFM. Since this process is much faster than full readout mode, more of the CCD can be cleared thus reducing the integrated noise. The DMA pointers for the accept memory and pixel buffers are reset before the next beam-crossing, effectively discarding the data read prior to going into FAST CLEAR mode. When a trigger signal is received, the FAST CLEAR mode is disabled until a full image is acquired 26 beam-crossings later. The DMA pointers are not reset, allowing the accept memory and pixel memory to fill with data.

Each event thus consists of 26 data packets. In order to eliminate deadtime, events are allowed to overlap. Consequently, overlapping events share data packets. The data in these shared packets are identical, but the CCD row and column address ultimately assigned to a pixel will depend on its event association. For reasons based on cost and lessons learned during the operation of VXD2, neither the accept memory ring buffer nor the pixel memory buffers are large enough to handle a worst case scenario. Furthermore, the module can only detect, not prevent, data overrun. While in practice this has little or no impact on the final data, it does

mean that the processing must be carefully partitioned into multiple prioritized tasks to minimize both the risk and consequence of data overrun.

The processing occurs in three separate tasks. The highest priority task examines the accept words associated with each data packet and copies the pixels from a 6 row by 8 column region of interest centred on the accept address to less volatile memory. Typically the input data is safe for 5-6 beam-crossings or approximately 50 ms. Another lower priority task selects the over-threshold pixels from the region of interest. Both of these tasks treat the data on a packet-by-packet basis, ignoring event demarcation information. The final task assembles the 26 packets comprising a full image into an output event by assigning each pixel value an absolute CCD row and column address. The tagged data values are then sorted by channel and CCD address. The sorting procedure also eliminates any duplicate pixels which arise from overlapping regions of interest. After tacking on a statistics block and other header information the data is posted. The contributions of all 16 LFMs are gathered by an ALEPH Event Builder (AEB) and assembled into a complete VXD3 event.

The cluster processing which carries the major load of data reduction for this detector, is described in detail below.

5.4.3 Cluster Processing

There are 307 Mpixels in total in the whole VXD3 system. As 5 bytes (1 byte of amplitude information plus 4 bytes of pixel address information) are necessary for each pixel in the final data format, the total data size would become 1.5 Gbytes/event if all the pixels were recorded. However, the data size allowed for the VXD3 sub-system is at most 100 kbytes. Thus, a factor of more than 10^4 online data reduction is necessary. A cluster processor (CP3) [24] was developed to perform this role. The cluster processor selects true hit pixels with high efficiency for min-I particles, while effectively suppressing noise hits.

The cluster processor function of the LFM is shown schematically in Fig. 25. Four channels of CCD data are sent serially with 20 MHz clock speed to both the cluster processor chip and the video RAMs. The VRAMs store all the pixel amplitude data. The cluster processor selects true hits (based on a digital noise filter and an optimized cluster-finding method described below) in real time and then outputs an accept signal. When an accept signal is issued, its pixel address is stored in the Accept RAM. The CPU reads the Accept RAM and reads out only relevant pixels from the VRAMs to the main memory, from which the data are later sent to the SLD data acquisition system.

The VRAMs can store only a quarter of the whole CCD data at a time and are overwritten three times during the event readout. Therefore, the data transfer from the VRAMs to the main memory needs to be fast enough to keep ahead of the data input. The logic scheme of the CP3 was developed to use CPU power effectively in order to speed up the data transfer.

Figure 26(a) shows a schematic diagram of the CP3 and Fig. 26(b) shows its main data processing functions. The input is 8 bit wide ADC data. Since the CCD is read out with continuous charge accumulation on the output node, resetting at the end of each row, the output amplitude A_{ri} corresponds to the total charge of the pixels which have been read out since the start of the row,

$$A_{ri} = \sum_{k=0}^r a_{ki}$$

where, k and i are the pixel addresses along the R clock direction and the I clock direction, respectively. a_{ki} is the signal amplitude for the pixel (k, i) . By subtracting neighbour data, an individual pixel amplitude can be determined,

$$a_{ri} = A_{ri} - A_{(r-1)i}$$

There are a few dummy pixels in each row, before and after the image area; these are used to extract the amplitude of the pixels of the first and last columns, a_{0i} and a_{399i} .

Some of the key functions of the CP3 are described below.

Kernel Formation with Improved Extended Row Filter

A typical min-I particle deposits most of its signal charge within a 2×2 pixel cluster and the four pixel amplitudes within this area are summed to obtain the hit amplitude. This 2×2 pixel unit is called a ' 2×2 kernel'. The 2×2 kernel (r, i) is defined to be made of pixels (r, i) , $(r-1, i)$, $(r, i-1)$ and $(r-1, i-1)$. In the cluster processor, the 2×2 kernel amplitude is calculated as follows; firstly the $2(R) \times 1(I)$ kernel amplitude is calculated by $A_{r(i-1)} - A_{(r-2)(i-1)}$, then the resulting data are delayed for one row readout period using a line buffer and added to the $2(R) \times 1(I)$ kernel amplitude of the next row, $A_{ri} - A_{(r-2)i}$.

In order to reduce pick-up noise, a digital noise filter, called Extended Row Filter (ERF) [25], is applied when forming the 2×2 kernels. The basic idea of the ERF is to sample the hit amplitude twice, $(A_{ri} - A_{(r-2)i})$ and $(A_{(r+1)i} - A_{(r-3)i})$, when making the $2(R) \times 1(I)$ amplitude, and then take the lower of them. As a true hit forms a step function in the output data flow,

both calculations yield approximately the same hit amplitude for the true hit case and this filter has little effect on the final 2×2 kernel amplitude. However, for noise signals that make only one pixel high (or low), one of the two amplitudes is approximately zero and the resulting $2(\mathbf{R}) \times 1(\mathbf{I})$ amplitude is much reduced. As shown in Fig. 26(b), the noise rejection power of the original ERF is further improved in the CP3 by taking the average of two adjacent pixels when calculating $(A_{(r+1)i} - A_{(r-3)i})$. This noise filter is called Improved ERF or IERF. In summary, the 2×2 kernel amplitude k_{ri} given by the IERF procedure is calculated by

$$2 \times k_{ri} = \min \left[2(A_{ri} - A_{(r-2)i}), \left((A_{(r+2)i} + A_{(r+1)i}) - (A_{(r-3)i} + A_{(r-4)i}) \right) \right] \\ + \min \left[2(A_{r(i-1)} - A_{(r-2)(i-1)}), \left((A_{(r+2)(i-1)} + A_{(r+1)(i-1)}) - (A_{(r-3)(i-1)} + A_{(r-4)(i-1)}) \right) \right]$$

In the CP3, this calculation is done by hardware in pipeline and the resulting 2×2 kernel amplitude is sent to a comparator to be compared to a kernel threshold.

Figure 27 shows the threshold dependence on the noise hit rate. The typical cluster threshold used with VXD3 is $250 e^-$. With this setting, the ERF reduces the noise hit rate by two orders of magnitude compared with that of genuine 2×2 kernels and the IERF reduces it by another order of magnitude. By using the very powerful IERF filter, the electronic noise hit rate of the VXD3 system can be suppressed to a negligible level, in spite of the very large number of pixels.

Cluster Edge Finding

When a kernel amplitude exceeds the kernel threshold, it is considered to be a hit kernel. Multiple hit kernels can possibly be generated for only one real hit, because a pixel belongs to four 2×2 kernels. For example, a track which passes through the centre of a 2×2 kernel may produce four high amplitude pixels and each pixel may generate four 2×2 kernel hits, resulting in 16 hit kernels in total. Furthermore, when a very large cluster is produced by, for example a beam halo particle which traverses a long region of sensitive silicon (a quite common source of SLD background at small radius), hundreds of hit kernels are produced and the CPU may lose a significant amount of data not only from the hit CCD but from other CCDs on the same F/E board as well, due to lack of processing time. In order to eliminate this problem, the CP3 implements the Edge Finding (EF) method of hit cluster identification. The idea of the EF is to look for an edge of a cluster regardless of its size. This is realized in the electronics circuit as shown in Fig. 26(b). When a hit kernel is generated, its left-lower four kernels are checked (for convenience, larger R address is expressed as 'left' and larger I address is expressed as

'lower') and if none of these have hits, the hit kernel is considered to be the left-lower edge of the cluster and an accept-cluster is generated. If there is already at least one hit among the left-lower kernels, the accept-cluster signal is not generated and the hit kernel(s) are considered to be new edge candidate(s) and their left-lower kernels are further checked. The logic of the EF method is expressed as follows,

$$\text{Accept} = h_{ri} \wedge \bar{h}_{(r+1)i} \wedge \bar{h}_{(r+1)(i+1)} \wedge \bar{h}_{r(i+1)} \wedge \bar{h}_{(r-1)(i+1)}$$

where, h_{ri} is *true* if k_{ri} is equal to or larger than the threshold and, *false* if not.

On detection of a cluster-accept signal, the CPU (Fig. 25) reads out $8(R) \times 6(I)$ upper-right pixels, compares them to a pixel threshold and only the pixels whose amplitudes are larger than the pixel threshold are read out to the main memory. As the CPU has a quick block data transfer instruction, this treatment of data makes efficient use of the CPU. It has been shown that more than 99.9% of the charge desposited by true min-I hits is contained within an 8×6 pixel area [24]. The EF method has another advantage of a 'self quenching' effect for hot columns (due to high leakage current from a single pixel) if they are hot enough to form contiguous hits along the columns, because such hits are considered to be a single long cluster.

Accept Number Counter

There is a 7-bit accept number counter and an overflow latch for each channel. The counter and the latch are cleared at every beam-crossing. When the overflow latch bit is set high (≥ 128 accepts/ch/beam-crossing), further accept signals are disabled. This is to allow for the case that a sporadic high hit density region, such as hits caused by a particle shower, should not swamp other healthy data. The CPU reads the counter data and the overflow latch bit, and responds appropriately to the accept number information.

Column Suppression

It is possible to suppress a particular column entirely by setting a column suppression bit. This function could be used to suppress hot columns and is used to suppress the dummy columns which precede and follow the image area.

Internal Registers

A cluster processor has 16 internal registers each of 8 bit width. Of these, 4 registers are used as kernel threshold data, 4 registers are for the accept number counter data, and the remaining

registers are used as various control and status registers. An address is assigned to each register and the CPU can directly access the registers.

Save the Last Columns Function

The IERF would effectively kill the last two columns ($r=398,399$) because the amplitudes of the dummy pixels beyond the image area are zero. As a CCD is read out from two opposite directions, there would be four columns or $80 \mu\text{m}$ of dead region at the centre of each CCD. In order to avoid this problem, the Save the Last Columns (SLC) function is implemented. The SLC has a high speed counter and monitors the column address of the current data. When it comes to the last column, the CP3 internally sets the amplitude of the first two dummy pixels (a_{400i}, a_{401i}), to be the same as that of the last pixel in the image area (a_{399i}).

Hardware

The CP3 electrical circuit is implemented in a XILINX Field Programmable Gate Array [26]. The XC3195-5 chip is used, which has 9,000 equivalent gates. Special treatments for speed-enhanced logic and effective use of the hardware resources were necessary when designing the logic. The CP3 design uses 80% of the CLBs (Configurable Logical Blocks) of the FPGA. The configuration data are stored on the data acquisition computer and the CP3 chips are configured from the computer.

5.4.4 Performance of Data Acquisition System

The acquisition of VXD3 data typically contributes no deadtime. The data volume from each LFM varies from 2-5 kbytes, corresponding to roughly 400-1000 tagged pixels. The total event size is around 40-80 kbytes depending on beam conditions. The system is almost 100% reliable.

6. Detector Performance in SLD

6.1 Operational Experience

6.1.1 Hardware Performance

During the 15 week SLD run in 1996, the VXD3 detector system was commissioned and then operated with high (but not full) efficiency. Not surprisingly, the F/E boards, which had been completed too late to permit extensive testing before the run, were subject to a small number of problems. These were largely repaired during access opportunities through the run. Reviewing the VXD3 system elements starting from the inside, the performance through this first run can be summarized as follows.

The CCDs all performed reliably and stably, as did the in-cryostat micro-connectors. The cryostat remained free of measurable leaks, and the gas cooling performed as designed. The power dissipation on the detector was approximately 20 W, resulting in a temperature rise of only 10 C within the cryostat, with a further 10 C rise due to thermal leaks into the cold enclosure. The procedure for aligning the inner section of beam-pipe during background tuning operated smoothly; controlled movements of up to 1.0 mm were carried out and were tracked by the position monitoring system. Between these deliberate alignment changes, long epochs of excellent positional stability were observed.

The mounting of the F/E boards on the support cones was initially problematic due to very restricted space in the centre of the detector, making board removal and replacement very tedious. However with experience this process has become relatively straightforward and uneventful. Furthermore, the cooling of these boards (by means of water-cooled plates in pressure contact with the heat-sinks) was insufficient, with boards operating at temperatures of typically 55-65 C. This tended to create excessive temperature gradients in the CDC. The problem has been solved (after the 1996 run) by direct water cooling of the on-board heat-sinks. The new cooling system has reduced board operating temperatures to 35-40 C.

Due to some missing components, changed values, and component failures, some of the production F/E boards differed in performance from the prototypes, resulting in data imperfections (charge smearing, stuck low order ADC bits, imperfect optical links to the LFM, for example). Most of these effects were offline correctable or had little effect on data quality. All have been understood and corrected in the hardware during or after the run.

A serious unexpected problem encountered during commissioning was beam-related electromagnetic interference affecting the data and control links between the F/E board and the

LFM. As discussed in Section 5, the F/E boards receive their controls via modulation of the clock signals from the data acquisition system, and a phase-lock loop on the F/E board verifies these timing signals and generates a board reset signal if it cannot maintain a lock on these signals. Most boards were observed to reset themselves consistently every beam-crossing, and this was traced to beam interference in the clock signals. The self-resetting feature was disabled (though the reset signal was still monitored) but it was found that the boards would still go into undesired control modes, again synchronous with beam arrival. This was traced to beam interference affecting the decoded clock modulation and hence generating spurious control commands. A blanking signal was then implemented so that the F/E boards ignore any control commands for $10\ \mu\text{s}$ during beam-crossing, which does not affect readout operations. Later studies with small loop antennae placed around the beam-pipe confirmed that the interference was due to magnetic fields induced by the passage of the e^+e^- beams.

6.1.2 Beam Background and Data Volume

Single-pixel and cluster processor thresholds were set as low as possible, to maximize hit efficiency while maintaining acceptably low readout noise. The global single pixel readout threshold was 4 ADC counts or twice the rms noise. The cluster processor threshold can be set individually for each detector channel; about 80% of the channels had the same nominal threshold ($\sim 270\ e^-$), while 18% had a 20% higher threshold due to readout noise and 2% had a 10% lower threshold. The typical data volume read out from the detector was 30-100 kbytes under good background conditions, which corresponds to a cluster density of $2\text{-}7/\text{cm}^2$ or a pixel occupancy of a few times 10^{-5} . This hit density is dominated by background from beam-related sources. A typical accept from the hardware processor for a charged particle deposit contains 2-3 genuinely hit pixels plus, on average, one noise pixel somewhere in the 6×8 cluster search area.

Instantaneous background levels and distributions in the vertex detector are highly dependent on accelerator operating conditions and hence highly variable, though integrated background levels and distributions are fairly stable; hence only the general features of the latter are addressed here. These backgrounds are studied by examining the detector hits which are not linked to charged particle tracks found by the CDC. These tend to have somewhat larger clusters of pixels (~ 4 typical) and larger charge deposits than those left by min-I particles. Densities of these unlinked hits decrease somewhat with radius from the beam line, but are relatively flat in azimuth and distance along the beam line. Figure 28 shows the distribution of raw hits as functions of ϕ , r and z . The decreasing area under each successive radial population indicates the rapidly falling hit density through Layers 1, 2 and 3 (which contain 12, 16 and 20 ladders respectively). The ramps visible on the z distributions reflect the increase in effective integration time as the signals are read out over 26 beam crossings; data from the

central region of each CCD sees the highest accumulated background. The background is dominated by scattered and fluorescence photons generated by synchrotron radiation, plus low energy electrons spiralling around the beam direction.

Beam background distributions are very stable from one running period to another, even though the support cones in the R20 module (which support not only VXD3 but also the final M4 masking, see Fig. 9) were moved several times to centre the beam in the detector and/or the masks. Before the first adjustment there was some evidence of increased background levels in the top and bottom regions of the inner VXD layer, but this disappeared after the R20 alignment was corrected. It should be noted that radiation damage to VXD2, characterized by slightly reduced signal amplitude after three years of operation, was most noticeable in the same geometrical region.

6.1.3 Observation of Radiation Damage Effects

VXD2 provided three years experience of running in SLD, and under standard operating conditions the dose rate was certainly negligible. It was obviously much higher (but not quantified) during beam tuning. For example, the CCD signals would saturate near the beginning of every row read out during wire-scans when a carbon fibre at $|z| = 30$ cm was traversed through the reduced-intensity beam. Due to the uniquely vulnerable location of the inner layer of CCDs (closely sheathing the most transparent section of beam-pipe) no external monitor could give a useful measure of their dose. Empirically, as discussed in Section 1.2, the VXD2 detector suffered only minor radiation damage integrated through its life, confirming hopes that the general environment was quite benign.

Initial running with the VXD3 detector demonstrated that one still needs to be careful. Due to SLC problems, undamped beams (i.e. beams with much larger emittance than normal) were passed through SLD for diagnostic purposes, for extensive periods of time. Furthermore, the SLD solenoid was powered off for much of this time, allowing low energy e^\pm background to reach the vertex detector, particles which would normally be safely curled up by the solenoid field and dumped remotely. As a consequence, VXD3 suffered more radiation damage during this brief initial period than was seen throughout the operating life of VXD2. The damage was most severe on Layer 1, and negligible by Layer 3, suggesting soft charged particles rather than photons. Electrons above 200 keV can cause displacement damage, and could easily be ranged out by the first or second detector layer, while photons of the required energy would have a rather long absorption length.

The damage was perceived primarily as a reduction in min-I signal amplitude versus I address, characteristic of charge transfer degradation due to displacement damage, as discussed in

Section 1.2. The cluster charge distribution for track-associated hits is plotted in Fig. 29 for Layers 1 and 3; the increased population of small signals in Layer 1 has a significant impact on efficiency.

The situation after the non-standard running period appears to have stabilized, the absence of further degradation during the 1996 run confirming the observation with VXD2 that under normal SLC/SLD operating conditions (including standard beam tuning) the radiation environment even for Layer 1 is quite benign. The consequence of the loss of signal amplitude during the undamped beam running depends of course on the pre-irradiation situation. It happens that 10 of the Layer 1 CCDs had slightly lower gain to begin with, and these show an efficiency loss at the end of the I register (central region of the CCD, where the CTI effect is most severe) of around 15%. The detector was operated at a temperature of 220 K during the 1996 run. This was warmer than for VXD2, and was believed sufficient for at least the first year in view of the good experience with the original detector. As noted in Section 1.2, 220 K is not optimal from the viewpoint of minimizing the effect of radiation damage, and subsequent to the run it has been confirmed that the efficiency of Layer 1 can be restored to its pre-irradiation level by cooling to 185 K.

For the future, the detector will be operated at this lower temperature (still acceptable as regards thermal stress in the NuSil adhesive joints attaching the ladder blocks and annulus blocks). In addition, unnecessarily hostile beam will no longer be permitted through SLD, and the SLD solenoid will be interlocked with the presence of SLC beams. In addition, the hit rates in all 384 CCD quadrants will be monitored during SLC operation, with automatic shutoff of the accelerator if these rates become dangerous. Given these precautions, there is every reason to expect that the performance of VXD3 will show negligible further radiation damage effects through the life of the detector.

6.2 Tracking Efficiency

The tracking algorithm used for the initial VXD3 performance evaluation is the same as that used for VXD2. The tracks found and fitted in the CDC are extrapolated into the VXD region. The search for matching VXD hits starts from the outer layers, then proceeds to the inner layers where backgrounds are higher. Due to the limited precision of the CDC-measured track, combined with the effect of multiple scattering and the long extrapolation to the vertex detector, the primary search area at the VXD is rather large. For all hits within this search area (often several are found) a secondary extrapolation is made to the inner hits by fixing the track at each outer hit in turn. Due to the 3-D nature and fine segmentation of the CCD pixel device, the secondary extrapolation regions for further hit searching are sufficiently small that real track hit vectors stand out very clearly and fake combinations are easily rejected. The best set of

matching VXD hits is then combined with the CDC track to perform a joint fit using the Billoir algorithm [27], taking into account the effect of multiple scattering in detector material. In VXD3 tracking, a successful link is required to have at least two VXD hits.

To test the overall tracking efficiency, a set of good quality CDC tracks which extrapolate close to the IP and within the VXD fiducial region is selected. The rate of all VXD links and the rate of VXD links with at least three hits is examined for these CDC tracks. For the high momentum tracks in clean $Z^0 \rightarrow \mu^+ \mu^-$ events, the all-link rate is >99.8% and the rate of ≥ 3 hit links is 97%. For the tracks in Z^0 hadronic decay events, the linking rates are examined as a function of track momentum as shown in Fig. 30. Since the CDC track selection could not guarantee the origin of the track as coming from the IP, a significant fraction of the remaining apparent inefficiency (~2%) is in fact due to tracks originating beyond the VXD from interactions or strange particle decays. The additional 3% reduction in 3-hit linking rate, seen also with the mu-pair tracks, is due to genuine single hit efficiency losses. The main source of this efficiency loss is the radiation damage referred to in Section 6.1.3. Small deficiencies in the electronics during the 1996 run also contribute to a lesser extent. By running the detector at a lower temperature and correcting the electronic problems, most of these efficiency losses will be eliminated in future running.

6.3 Alignment and Resolution in Tracking and Impact Parameter

The optical survey data are used to produce the first level detector geometry. This includes the CCD shapes from the ladder survey data, together with temperature scaling and gravity sag corrections. The first level geometry is followed by a track-based full VXD internal alignment. Three types of VXD hit vectors from tracks in hadronic Z^0 decay events are used for this process. The general principle is to use the track curvature measured by the CDC, then force the VXD vector to go through two track hits precisely, and measure the residual of a third hit to this fixed vector. Apart for the minor input of the CDC curvature measurement, these vectors are pure VXD internal observables independent of the CDC. The three types of track vectors are:

- (a) doublets: tracks through the overlapping region (z) of two CCDs on the same ladder. The vector is fixed at one of the doublet hits and at the furthest-away hit on a different layer. The residual of the second doublet hit to this fixed vector is called the doublet residual.
- (b) shingles: tracks through the overlapping shingle region ($r\phi$) of two ladders in the same layer. The vector is fixed at the hit on one of the shingle ladders and at the

furthest away hit on a different layer. The residual of the hit from the second shingle ladder to the fixed vector is called the shingle residual.

(c) triplets: tracks with hits in all three layers. The vector is fixed at the Layer 1 and Layer 3 hits, and the Layer 2 hit residual is called the triplet residual.

The doublets constrain the relative locations of two CCDs on the same ladder. The shingles provide direct connections between different ladders in the same layer, and constrain the overall radius of each layer. The triplets relate ladders on different layers.

For each unique combination of the above multiplets, the $r\phi$ residuals are studied as a function of ϕ and the z residuals are studied as a function of $\tan \lambda$ for shingles and triplets, where λ is the track angle to the normal in the rz plane. Some irregularities in these residuals are observed due to the electronic problems present in the 1996 data. These local effects are corrected at this stage. Third order polynomial fits in $\tan \phi$ and $\tan \lambda$ are then performed for each unique multiplet combination in the $r\phi$ and z views respectively. Each of the 96 CCDs is then given six positional description parameters; three translations and three rotation angles. Each parameter typically only couples to one or two types of the polynomial fit coefficients. The residual values clearly depend on the positions of all three CCDs involved in any multiplet combination, but with different dependency due to the various lever-arm factors. A grand matrix equation is then built as

$$A = C \times P$$

where A contains all the fitted polynomial coefficients for the various multiplet residuals, and P contains the 576 CCD position parameters. The matrix C with dependency coefficients is then inverted using the singular value decomposition technique [28] to derive the full set of CCD alignment position parameters.

The multiplets only provide good local constraints between ladders at the same ϕ region. Some long range distortions varying slowly with ϕ can still remain without significant degradation to the total χ^2 in the multiplet. Additional constraints are obtained using $Z^0 \rightarrow \mu^+\mu^-$, e^+e^- events where the two tracks are known to be back-to-back in angle and to have originated from the same spatial point. These data have recently been included in the overall alignment program. The triplet residual distributions for all triplets are shown in Fig. 31. The single hit spatial resolution can be derived after dividing the triplet residual resolution by $\sqrt{1.5}$. The resulting values of $5.4 \mu\text{m}$ for single hits in the $r\phi$ and z views represent a global spatial resolution including alignment errors. To separate out the

contributions from intrinsic spatial resolution and alignment errors, the residuals for all doublets are plotted in Fig. 32. The derived single hit resolution of $4.6 \mu\text{m}$ for single hits in the $r\phi$ and z views is closer to (but still greater than) the intrinsic resolution.

The internally aligned VXD is then taken as a rigid body and aligned globally with respect to the CDC using tracks in hadronic Z^0 decay events [29]. The CDC-alone tracks are used to perform an event by event primary vertex finding. The CDC tracks compatible with the primary vertex are fixed to pass precisely through this point. The positions and angles of these adjusted CDC tracks at the VXD region are compared with the positions and angles of their corresponding VXD hit vectors. A total χ^2 minimization is used to determine the three translation and three rotation parameters of the VXD w.r.t. the CDC. This procedure is applied for different periods separated by known mechanical adjustments of the R20 system. Each alignment period contains 500-10,000 Z^0 events and the precision of the global alignment is $2 \mu\text{m}$ for x,y translations, $10 \mu\text{m}$ for z translation, 0.03 mrad for the roll angle and 0.10 mrad for the two tilt angles. These measurements of alignment precision are verified by subdividing long periods and checking that the results are consistent within the estimated errors.

The $Z^0 \rightarrow \mu^+ \mu^-$ events are reconstructed with the fully aligned VXD to evaluate the asymptotic track impact parameter resolution at high momentum. The miss-distance of the two back-to-back tracks near the beam-line is a clean estimator of the impact parameter resolution independent of the knowledge of the beam-spot. A single track impact parameter resolution of $14 \mu\text{m}$ in $r\phi$ and $26.5 \mu\text{m}$ in rz can be derived from the fit to the data as shown in Fig. 33. The extended lever-arm in the VXD3 geometry not only brings improvements in the impact parameter resolution, it also improves the track dip angle resolution significantly due to its superb z spatial resolution.

Due to the reduction in layer thickness compared with VXD2, the impact parameter resolution is particularly improved at low momenta. This is evaluated for Z^0 hadronic decays, and requires knowledge of the interaction point for the event. This is determined in the $r\phi$ view by averaging over approximately 20 preceding events, taking advantage of the small SLD beam-spot, and its excellent stability over time. The deduced impact parameter resolution versus momentum is shown in Fig. 34, which also indicates the comparison with the VXD3 Monte Carlo, and the improvement with respect to VXD2. The overall impact parameter resolution of VXD3 deduced from the data is given by

$$\sigma_{r\phi} = 14.0 \oplus \frac{33}{p \sin^{3/2} \theta} \mu\text{m}$$

$$\sigma_{rz} = 26.5 \oplus \frac{33}{p \sin^{3/2} \theta} \mu\text{m}$$

There remains only a small discrepancy with respect to the Monte Carlo expectation (Section 1.4) mostly in the rz view for high momentum tracks. It is known that this discrepancy is coming mainly from a problem with the CDC tracks in the combined fit. If the VXD3 vector alone is used in the analysis, the constant term is *reduced* from $26.5 \mu\text{m}$ to $15.9 \mu\text{m}$. This problem (seen also in the fitting of Monte-Carlo generated tracks in the CDC) will surely be solved in the near future. For most physics purposes, the track momenta are sufficiently low that the multiple scattering term dominates. Here, the detector is already delivering its design performance.

7 Lessons Learned and Future Prospects

CCD-based devices appear to have become established as the technology of choice in vertex detectors for current and future e^+e^- linear colliders. The combination of extremely high *segmentation*, ongoing reduction in *thickness* to levels beyond the reach of other technologies, and the intrinsically *deadtimeless* readout, combine to provide a detector system well matched to this environment, now and in the future. The long readout time associated with a CCD system is not a disadvantage in this case, since the time structure of widely separated bunches or bunch trains is well matched to the CCD readout requirements.

VXD3 represents the state-of-art vertex detector for an e^+e^- linear collider, a major advance on its predecessor VXD2. In this paper, this present phase of development has been described. It may also be useful to record a number of false trails that were pursued, and lessons learned. These can provide guidance regarding the upgrades to be developed for a detector at a future International Linear Collider. The ideal vertex detector would be a pixel-based instrument with sub-micron precision, negligible readout time and negligible layer thickness. It is believed that the VXD3 design provides a useful springboard from which it will be possible to approach closely to these ideal requirements.

7.1 Lessons Learned from VXD3

In this Section the separate components (CCDs, ladders, support structure, flex-circuits and F/E electronics) are considered in turn. Important lessons were also learned regarding the cryostat and temperature control system, the data acquisition hardware and software, but in these areas not much needs to be added to the descriptions already given.

7.1.1 CCD Design

This project demonstrated the feasibility of a rapid fully customized CCD design, with actual devices that meet or exceed all design parameters. This is most encouraging, and will lead to even more adventurous designs in the future. However, it is clear that this high level of performance is not universally available. Sample devices provided by some manufacturers were badly flawed. Therefore, as well as a tightly written specification, a most important pre-procurement step is to test devices from all interested manufacturers. Some manufacturers simply do not know if their products meet the acceptance criteria and will readily put in a bid in the hope that things may work out.

As discussed in Section 3, a problem was encountered during production that could only be found by cold testing. Fortunately, the number of faulty devices was small, but had this not

been the case, the production phase could have been severely disrupted, particularly since each test involved a trans-Atlantic trip for the ladder. In future, a safer procedure would be to plan for more complete screening of devices at the manufacturer.

The original design called for 8 μm wide supplementary channels [30] to confine the signal charge packets transversely in the imaging region, so increasing the radiation resistance (CTI degradation due to bulk damage) by a factor of about five. When the charge trapping problem was found on the prototype devices, it was feared that this could be due to irregularities in these channels, so they were omitted in the production phase. With hindsight, this design change was unnecessary; the cause of the charge trapping was faulty etching, and this supplementary channel would have improved the radiation resistance.

Two variants of the CCD output circuit were prototyped, with the first-stage transistor being either surface channel (for lowest power dissipation) or buried channel (providing lower noise but with somewhat higher power dissipation). Note the equivalent terminology; surface channel equals enhancement mode, buried channel equals depletion mode. The much larger second stage transistor was surface channel in both cases, since its contribution to the overall noise was negligible. In the case of two surface channel transistors, the output load capacitance is ~ 40 fF which, together with a voltage gain of 1.75, gives an overall responsivity of about $3 \mu\text{V}/\text{electron}$. With the alternative buried channel first-stage both the capacitance and gain are about 10% lower giving similar overall responsivity. In both cases, the output impedance is $\sim 260 \Omega$. The power dissipation on chip for each signal output is about 45 mW for the surface channel version and about 84 mW for the buried channel version. The output noise spectra and sampled noise levels (for correlated double-sampling operation [19]) are shown in Fig. 35.

The choice between the SC/SC or BC/SC output amplifiers was made during the prototype ladder testing, in favour of the former option. While the BC option did have lower noise, just as simulated, the system noise was not much improved due to the large contribution associated with the R clock feedthrough. Furthermore, the power dissipation penalty associated with the BC option (a factor of two inside the cryostat) would have necessitated the use of pulsed power (amplifiers switched on only during readout of triggered events).

7.1.2 Ladder Design and Support Structure

The VXD2 ladder motherboards were made of alumina, with two-layer interconnected gold traces plus a ground plane on each side, made with thick film technology. This complex structure was required in order to service four CCDs per side, with unfavourable bond-pad geometry (wire bonds along the long edges of the ladder). As a consequence, the ladders were inevitably rather thick ($1.15\% X_0$).

The initial plan for VXD3 was to reduce the mass by using beryllia, with single-layer aluminium traces, made possible by the careful layout of bond pads in the full-custom CCD design. This option ran into a number of difficulties, including fragility of the thin beryllia, a tendency of the substrates to crack under the high temperature metallization/passivation steps, and irregular aluminium trace widths. There were in addition concerns regarding the cryogenic use of long aluminium traces; some problems have been experienced with cracks developing in this polycrystalline material. Some or all of these problems might have resulted from the use of manufacturers without quite the right background experience. Due to these problems, the beryllium-based ladder design described in Section 2 was tried. It worked extremely well, and resulted in a ladder thickness of $0.40\% X_0$, 35% less than in the beryllia design.

There is one remaining significant weakness with this ladder design, which applied equally to the VXD2 ladders; the structure is mechanically rather flexible. As has been described, the ladder shape is stabilized by clamping to the support structure, but this stable shape will in general differ from one assembly to the next, depending on the slightest variations in the relative orientation of the mating faces of the annulus blocks. In particular, this causes the shape of any individual ladder in the ladder survey jig to change unpredictably when supported in the detector. In principle, the shape can be re-determined in the barrel survey, but making depth measurements through the support shell with the OMIS system was not extremely precise. As a result, the detailed ladder shapes have to be established from tracking data. While each shape is adequately determined by two parameters (the dip angle of each annulus block), the dependence is non-linear. These complications could be avoided if the ladder were made more rigid. In the case of VXD3, it was not possible to achieve this while preserving the low overall thickness. For the future, advances in CCD design will solve this problem, as discussed in Section 7.2.

7.1.3 Flex-circuits and Front-end Electronics

For both VXD2 and VXD3, the connections from the ladder to the F/E electronics (typically 23 traces per VXD2 ladder and 34 traces per VXD3 ladder, plus grounds) were made by copper/kapton flex-circuits in the form of microstrip lines (traces plus single-side ground). In the case of VXD2, these were attached to the ladder tongues by custom-designed micro-connectors. These delicate connections to fragile alumina presented a significant technical challenge, and (as has been discussed) were not 100% reliable. They were avoided in VXD3 by extending the copper/kapton used in the laminated ladder design by a length of ~5 cm beyond the end of the ladder, permitting the use of commercial micro-connectors. These have proved extremely reliable in the cryogenic environment, but making and verifying all these connections was a complex task in the build-up of the detector, with attendant risk factors. Consequently for the future, longer flex-circuits would be used, with the only mechanical connector being the much more robust link to the F/E board, operated at room temperature.

For VXD2, the front-end electronics differed between the two ends, drive pulses generated on the north and analogue processing on the south. Even so, the readout speed was limited to 2 MHz. For VXD3, the two-CCD layout suggested a different approach, with the north CCDs being serviced from the north and correspondingly for the other end. This resulted in intrinsically greater crosstalk between drive and analogue signals, on the F/E board, on the flex circuits and indeed on the ladders themselves. The CCDs were designed to a very high crosstalk specification, and could easily have operated cleanly in excess of 10 MHz. However, it became necessary to reduce the global readout rate from 10 MHz to 5 MHz due to the combination of these off-CCD problems. While this reduced readout rate was quite tolerable in terms of SLD backgrounds, design developments are discussed in Section 7.2 which should enable much higher clocking rates to be achieved in future, with even better noise performance, and consequently improved spatial resolution.

The use of the gas shell and inner wall of the cryostat as a Faraday cage has proved to be extremely effective at attenuating beam-related pickup. As noted in Section 5, however, direct pickup into the F/E electronics caused problems. It is probable that a more hermetic continuation of the detector Faraday cage could eliminate such problems in future.

In VXD2 the biases, I drive pulses, R drive timing pulses and analogue output signals were all connected from the FASTBUS to the F/E board on coaxial copper cables, a total of ~10,000 cables, one ton in weight. Moving all the bias and drive pulse generation, plus the A/D conversion, to the F/E board in VXD3 allowed the external connections to be made by a small number of almost massless optical fibres. This major reduction in the unwanted material carried with it the further benefit of extremely clean detector grounding. The CCDs, given their

long readout time, are particularly vulnerable to pickup. It is a reflection of the very benign conditions inside the SLD detector (the iron and coil are an extremely effective shield) that no significant pickup effects have been seen either with VXD2 or VXD3, other than those related to the e^+e^- beams.

7.2 Future Prospects

As with the evolution from VXD2 to VXD3, the next step in the development of CCD-based vertex detectors (for a future International Linear Collider) promises to be empowered by advances in CCD technology.

The first major development of this type is expected to be a further reduction in ladder thickness. Using the technology developed for back-illuminated CCDs for astronomy, the devices can be thinned to $\sim 25 \mu\text{m}$ (just short of the p/p^+ interface), by a combination of mechanical lapping and chemical etching. In order to further reduce material, all the electrical connections would be confined to one end of the CCD, at the outer edge of the ladder. Hence it is possible to avoid the need for electrical connections to be carried along the ladder motherboard within the active volume of the detector. The motherboard is in this case required merely to provide mechanical support. At the expense of a dead strip of only some tens of microns in the middle of the ladder, two-CCD ladders can be constructed with the CCDs mounted both on the same side, butted together in the middle. This single-sided ladder permits a much more rigid mechanical design (with no increase in the beryllium thickness), by using an electron beam welded omega beam structure. If one were limited to signal outputs at the ends of the R register (as for VXD3), such a design would have excessive readout time. But the already-developed procedure of subdividing a single R register with as many as 16 ports, will completely overcome this problem. A sketch of such a futuristic ladder of thickness only 0.12% X_0 is shown in Fig. 36.

As well as the multi-port R register, the readout time can be greatly reduced by a much higher R clocking rate. Already in HEP applications, the readout rate for noise performance below 100 e^- has been increased from around 10 kHz (CERN test beam in 1982), to 2 MHz (VXD2 in 1992), then to 5 MHz (VXD3 today). Future developments are likely to lead to equivalent noise performance with a readout rate of at least 50 MHz. Developments already made in output circuit noise performance are indicated in Fig. 35(c) and (f). This unprecedented noise performance results from a combination of the buried-channel first stage with further reductions in node capacitance [31]. All the advantages of this upgraded CCD performance would be thrown away if the drive/analogue crosstalk problems of VXD3 were to persist. It is proposed to reduce these to a very low level by generating the R drive pulses locally to the R register,

using special ICs bonded to the end of the ladder motherboard, just outside the active region of the detector.

As regards radiation hardness, the main problem is the sensitivity of large CCDs to displacement damage. In fixed target hadron beam experiments, a small number of devices are used, and can be replaced as required. But for the collider environment, the time and cost of building up a large detector assembly precludes this approach. Fortunately, as has been discussed, displacement damage in the SLC environment is at a relatively low level, and can be controlled adequately by the optimal operating temperature and the incidental trap-filling benefit of charge from background hits in the detector. As the beam energy and average current increase, however, the potential for significant neutron fluxes due to photodisintegration processes becomes more severe. Present estimates indicate that 'standard' CCDs should still be viable, but may be on edge. The combination of supplementary channel plus temperature optimization (170 K) can give a good factor of ten improvement. There are some exciting results regarding 'defect engineering' which suggest that material which is more oxygen-free may be a further factor five times harder [32]. An interesting alternative is to switch to p -channel devices [33]. These studies are only beginning, and there is every reason to believe that great strides will be made over the next decade. Ideally, some of the funding for these developments could be provided from HEP sources, but even if not, the widespread interest in radiation resistant CCDs will guarantee a healthy R&D effort in this direction. This is one of the benefits of working with devices having inter-disciplinary scientific application areas.

Specifically for the future Linear Collider, the problem of neutron damage needs to be carefully studied. In the case of SLC, the beam dump was the only serious potential source of neutrons. In the early days of the SLC design, this was a serious issue, but careful design of the dump effectively eliminated this potential background. For the future machine, the e^+e^- pair dumps, the beamstrahlung dumps and the main beam dumps all need careful consideration. Should any of these lead to unacceptable neutron levels, a fallback option for the vertex detector (active pixel sensors instead of CCDs) would be required. This is less favoured for physics reasons, due to the greater layer thickness ($\times 10$), poorer segmentation ($\times 50$) and poorer precision ($\times 10$) associated with this technology.

Acknowledgements

We would particularly like to acknowledge the excellent collaboration with the EEV team, including R Bell, J Bloomer, S Bowring, D Burt, R Holtom, C Hotston, P Pool, G Wood. We would also like to acknowledge the excellent engineering advice from G Tappern, and technical support at SLAC of M Freytag, P Nguyen, R Rodriguez, and L Salgado.

Abbreviations, Acronyms and Symbols

Abbreviations and Acronyms

AEB	ALEPH Event Builder
BC	buried channel (depletion mode) FET
CDC	Central Drift Chamber
CKM	Cabibbo - Kobayashi - Maskawa
CLB	Configurable logical block (of FPGA)
CMM	Coordinate measuring machine
CP	charge conjugation-parity reversal
CP3	cluster processor integrated circuit
CTE	charge transfer efficiency
CTI	charge transfer inefficiency
DUART	Dual Universal Asynchronous Receiver/Transmitter
EEV	CCD manufacturer (reference [13])
EF	edge finding feature of cluster processor
ENC	equivalent noise charge
ERF	extended row filter
F/E	Front-End (local to detector)
FPGA	Field programmable gate array
I drive, register	I = parallel or imaging register of CCD
IERF	improved extended row filter
IP	interaction point
IR	interaction region
K&S	Kulicke and Soffa (wire-bonder manufacturer)
LFM	'Last FASTBUS module' (VXD3 readout module)
M1-M4	masks for shielding SLC background
min-I	minimum-ionizing
NIEL	non-ionizing energy loss
NIST	National Institute of Science & Technology
NuSil	adhesive (reference [16])
OMIS	type of CMM (reference [21])
PV, SV, TV	primary, secondary, tertiary vertex
R20 module	Equipment inside CDC (2 m length by 20 cm radius)
R drive, register	R = linear or readout register of CCD
RFP	Request for Production
SC	surface channel (enhancement mode) FET
SLC	SLAC Linear Collider
SLD	SLC Detector
TCM	Timing and Control Module (FASTBUS)
VXD	Vertex Detector
VXD1	prototype SLD vertex detector (3 ladders)
VXD2	original SLD vertex detector
VXD3	upgraded SLD vertex detector

Symbols

A_b	Asymmetry parameter in $Z^0 \rightarrow b\bar{b}$ decay
a_{ri}	signal amplitude for pixel in column r , row i
A_{ri}	Summed signal for pixels 0- r in row i
B	Hadron containing b -quark
B_s^0	Neutral meson $b\bar{s}$
c, b, t	heavy quark flavours charm, bottom and top
D	Meson containing c -quark
e^+, e^-	positron, electron
e^-	electric charge of the electron
n, n^+	n -type or heavily doped n -type silicon
p, p^+	p -type or heavily doped p -type silicon
p, p_T	momentum, transverse momentum (relative to beam direction)
R_b	Ratio of Z^0 decay width to $b\bar{b}$ compared to all hadrons. $R_b = \Gamma_{b\bar{b}} / \Gamma_{\text{had}}$.
$r\phi$ or xy	projection onto plane normal to beam direction
rZ	projection onto plane containing track direction and beam direction
u, d, s	light quark flavours up, down and strange
V_{td}	CKM matrix element
x, y, z	Cartesian co-ordinates (z = beam axis in SLD; y axis is vertical up)
x_s	B_s^0 / \bar{B}_s^0 mixing parameter $x_s = \Delta m_s / \Gamma_s$
X_0	radiation length
Z^0	Neutral gauge boson, $M = 91 \text{ GeV} / c^2$
θ, ϕ, r	Polar co-ordinates (z = beam axis in SLD; $\phi = 0$ is in the horizontal plane)
μ, τ	muon and tau lepton
\wedge	logical and
\vee	inclusive or

	VXD2	VXD3
Titanium liner	0.28	0.14
Beryllium beam-pipe	0.28	0.22
Beryllium gas jacket	0.14	0.16
Ladder material per layer	1.15	0.40
Beryllium outer shell	0.28	0.47
Cryo nitrogen gas	0.05	0.05
Cryostat	1.90	0.70
CDC inner wall	1.80	1.80

Table 1 Average material in % X_0 seen by tracks perpendicular to the beamline. Each track traverses 2.3 ladders on average for VXD2 and 3 ladders for VXD3.

Beam-pipe liner thickness (titanium)	0.051 mm
Inner beam-pipe radius	23.2 mm
Beam-pipe wall thickness	0.76 mm
Inner gas jacket radius	24.8 mm
Gas jacket wall thickness	0.51 mm
Layer 1:	
Average CCD radius	28.0 mm
Number of ladders	12
CCD cant angle	10°
Layer 2:	
Average CCD radius	38.2 mm
Number of ladders	16
CCD cant angle	9°
Layer 3:	
Average CCD radius	48.3 mm
Number of ladders	20
CCD cant angle	9°

Table 2 Principal mechanical parameters of the detector assembly.

Component	Thickness (X_0)
Beryllium motherboard	1.08×10^{-3}
CCD + adhesive	1.60×10^{-3}
Kapton + adhesive (both sides)	0.47×10^{-3}
Metal traces (17.8 μm copper) (average, both sides)	0.90×10^{-3}
Total	4.05×10^{-3}

Table 3 Contributions to ladder thickness. The bond pads in the central region of the ladder (7.6 μm gold on a 3.8 μm nickel barrier layer) are very small (1 mm x 0.25 mm each) and contribute negligible additional material.

Basic Design Features

Substrate resistivity	< 20 m Ω cm
Epitaxial layer resistivity	20 Ω cm
Format	4 quadrant full frame
No. of pixels	800 Hor x 4000 Vert
Pixel size	20 \times 20 μm^2
Sensitive area	16 mm x 80 mm
Overall chip size	\leq 16.6 mm x 82.8 mm
Inactive edge spacing	< 300 μm
Thickness	180 \pm 20 μm
Passivation	2 μm polyimide
Image area clock type	3-phase
Readout register clock type	2-phase
No. of pre-scan elements	6
No. of amplifiers	4
Gate protection	on all gates

Performance Parameters

Clock Capacitances	
Image section to substrate	16 nF
Image section interphase	6 nF
Readout register to substrate	85 pF
Readout register interphase	30 pF
Charge Storage Capacity	
Pixel (supplementary channel)	100 $\times 10^3$ e $^-$
Pixel (total)	350 $\times 10^3$ e $^-$
Readout register	400 $\times 10^3$ e $^-$
Vertical transfer rate	> 200 kHz
Horizontal transfer rate	> 10 MHz
Output circuit responsivity	3 $\mu\text{V}/\text{e}^-$
Output impedance	260 Ω
Power Dissipation (on-chip)	
Image section (10 V clocks at 200 KHz)	1.3 W
Readout register (10 V clocks at 10 MHz)	25 mW
Each output amplifier	45 mW

Table 4 CCD design parameters.

References

- [1] W S Boyle and G E Smith, *Bell Syst Tech Journal*, **49** (1970) 587
- [2] C J S Damerell et al, *Nucl Instr and Methods*, **185** (1981) 33
- [3] R Bailey et al, *Nucl Instr and Methods*, **213** (1983) 201
- [4] C J S Damerell et al, *IEEE Trans Nucl Sci*, **33** (1986) 51
- [5] S Barlag et al, *Phys Letters B*, **184** (1987) 283
- [6] C J S Damerell et al, *Silicon Detectors for High Energy Physics*, FERMILAB publication (1981) 33
- [7] G D Agnew et al, Proceedings of the 26th International Conference on High Energy Physics, Dallas 1992 (*World Scientific, New York 1992*), Vol 2, 1862
- [8] C J S Damerell, Rutherford Appleton Laboratory preprint RAL-P-95-008 (1995)
- [9] M Huhtinen and P A Aarnio, *Nucl Inst and Methods*, **A335** (1993) 580
- [10] D J Jackson, *Nucl Instr and Methods* A388 (1997) 247
- [11] SLD Collaboration. SLAC -Pub 7482 (1997) and Physical Review Letters (to be published)
- [12] SLD Collaboration. SLAC-Pub 7266 (1997) and Physical Review Letters (to be published)
- [13] EEV Company, Chelmsford, Essex, UK
- [14] GEANT 3.21 Program, CERN Applications Software Group, CERN Program Library (1996)
- [15] C J S Damerell and D J Jackson, Proc Snowmass Workshop 'New Perspectives in Particle Physics' (1996). To be published.
- [16] CV-1142 adhesive produced by the Mc Ghan NuSil Corporation, Carpinteria, California, USA
- [17] Circuit Works 2400 Conductive Epoxy, available from Planned Products, Santa Cruz, California, USA
- [18] Nilflam insulating foam produced by the Kooltherm Company of Derbyshire, UK
- [19] M H White et al, *IEEE Journal Solid-State Circuits*, **SC-9** (1974) 1
- [20] J Frenkel, *Phys Rev*, **54** (1938) 647
- [21] RAM Optical Instrumentation, Huntington Beach, California, USA.
- [22] Product of 3M Corporation.
- [23] Capacitec Inc., Ayer, Massachusetts, USA. The 4100-S amplifier units and HPB-375 probes were used

- [24] F Suekane et al, SLD VXD3 Note-26, TOHOKU-HEP-Note-94-16; K Hasuko et al, SLD VXD3 Note-34, TOHOKU-HEP-Note-95-10; K Hasuko, master thesis, Tohoku Univ. 1996, SLD VXD3 Note-40; J Yashima et al. SLD VXD3 Note in preparation.
- [25] C J S Damerell et al, *IEEE Trans Nucl Sci*, **37** (1990) 305
- [26] XILINX, Inc., San Jose, California, USA
- [27] P Billoir, *Nucl Instr and Methods*, **225** (1984) 352
- [28] G H Golub and C F van Loan, *Matrix Computations* (Johns Hopkins University Press) (1983)
- [29] P J Dervan and E Etzion. SLD Note 253 (1996)
- [30] A D Holland et al, *IEEE Trans Nucl Sci*, **NS-38** (1991) 1663
- [31] Noise performance of CCD 42. EEV Internal Report (1996)
- [32] S Watts et al, BRUCRD-ESACCD-95-IR, ESA Contract Report (1995)
- [33] N S Saks, J M Killiany, P R Reid and W D Baker, *IEEE Trans. Nuclear Science* NS-26 (1979) 5074 and S Watts, private communication.


Figure Captions

- 1 Contrast in layout of CCDs between VXD2 and VXD3, in views transverse to and along the beam direction.
- 2 Basic 2-CCD ladder design used for VXD3. Active length = 16 cm.
- 3 Cross-section ($r\phi$ view) of VXD3.
- 4 Cross-section (rz view) of VXD3. For definiteness, the detector components are all referred to according to their orientation (north/south) in SLD, as noted on this figure.
- 5 Cut-away isometric drawing of the VXD3 detector.
- 6 Impact parameter resolution as a function of momentum for tracks at $\cos\theta = 0$ compared between VXD2 and VXD3 Monte Carlo simulations. The measured resolution from VXD2 data are also shown to agree with the VXD2 Monte Carlo.
- 7 Exploded view of ladder, showing the main components. B =beryllium substrate. F =Flex-circuit with traces t , dummy traces dt , bond pads p and grounding aperture g . C =CCD with adhesive pads a and wire bonds w .
- 8 Sprung blocks with sliding joint (vee and flat contact areas) for relief of stress due to differential contraction between ladder and support structure.
- 9 The section of SLD inside the CDC inner barrel of length 2 m and radius 20 cm, referred to as the 'R20 module'.
- 10 Photograph of the VXD3 upper module prior to mating with the lower module on the SLC beam-pipe
- 11 CCD32 basic architecture. The parallel register (I register) shifts signal charge packets to each end of the device. A serial register (R register) at each end shifts charge packets to a pair of output circuits. Chip area = 13 cm^2
- 12 The taper section in each corner of the CCD, providing room for the output circuit within the device width required for the image area bus-lines. The SS (substrate) ring provides edge passivation, and the dump channel defines the sensitive area of the detector.



- 13 Output circuit. T1 resets the node to a standard voltage VRD at the end of each row readout. T2/3 provides the 2-stage voltage sensing circuit. Load R2 is on the CCD, R3 is on the front-end board, outside the cryostat. Most of the parasitic capacitance C_L is due to the stripline connection between the ladder and front-end electronics.
- 14 Hit rate as function of I address along a few columns having significant traps. The trap location can be precisely determined by the sharp fall in hit rate.
- 15 Photomicrograph of CCD imaging area in a region of high trap density. The ragged gate edge due to over-etching is clearly visible.
- 16 Ladder mounted on survey jig, seen in each of four views used for measurement. Note the ladder/annulus block assemblies (visible in the edge-on views).
- 17 Typical ladder survey data in top view and side view. The measuring machine coordinate conventions are used.
- 18 Isometric drawing of the VXD3 detector mounted on the dummy beam-pipe. Note the precision tooling balls temporarily fitted to the support structure, during the barrel survey operations.
- 19 The end and one side view of the barrel survey. Note the 32 tooling balls which define the barrel survey coordinate system and the web pattern of the visible features through the detector support shell.
- 20 The shapes of the CCDs (active faces) of a typical ladder from the fits.
- 21 Schematic view of capacitive monitoring system.
- 22 Example of parametrised probe voltage as a function of x - y location of the wire.
- 23 Example of monitored voltage during controlled movement of the R20 module; the relative displacement between initial and final positions is 0.5 mm.
- 24 General layout of the VXD3 Electronics System
- 25 Schematic diagram of the full electronics processing chain for one channel of the VXD3 detector.

- 26 (a) Schematic diagram of the VXD3 cluster processor CP3 and (b) its main data processing functions.
- 27 Threshold dependence of the noise hit rate in the cluster processor with and without the IERF.
- 28 Background hit densities as function of azimuth ϕ , radius r and position along beam-line z
- 29 Layer 1 and Layer 3 hit cluster charge distributions.
- 30 Quality CDC track VXD linking rate, as a function of track momentum, for all links and 3-layer links.
- 31 Triplet residual distributions in $r\phi$ and z view for all triplets.
- 32 Doublet residual distributions in $r\phi$ and z view for all doublets.
- 33 Mu-pair miss distance in $r\phi$ and rz projections.
- 34 Measured impact parameter resolution as a function of track momentum for tracks at $\cos\theta = 0$ for VXD2 and VXD3, compared with the Monte Carlo simulations.
- 35 Noise spectra of (a) buried channel and (b) surface channel transistors T2 (left-hand axis). (d) and (e) show the corresponding CDS noise equivalent signals in RMS e^- (right-hand axis). Plots (c) and (f) are for currently available state-of-art output circuits.
- 36 Conceptual 2-CCD ladder design for a vertex detector at the future International Linear Collider. The beryllium omega-beam substrate is no thicker than for VXD3, but over 50 times more rigid.

VXD2



VXD3



5-97

8262A32

Fig. 1




Fig. 2




Fig. 3






Fig. 4



VXD3
(SLD)


1-97
7934A1

Fig. 5



6-97
8262A23

Fig. 6



6-97
 8262A22

Fig. 7




Fig. 8




Fig. 9





Fig. 10



7-97

8262A12

Fig. 11




Fig. 12

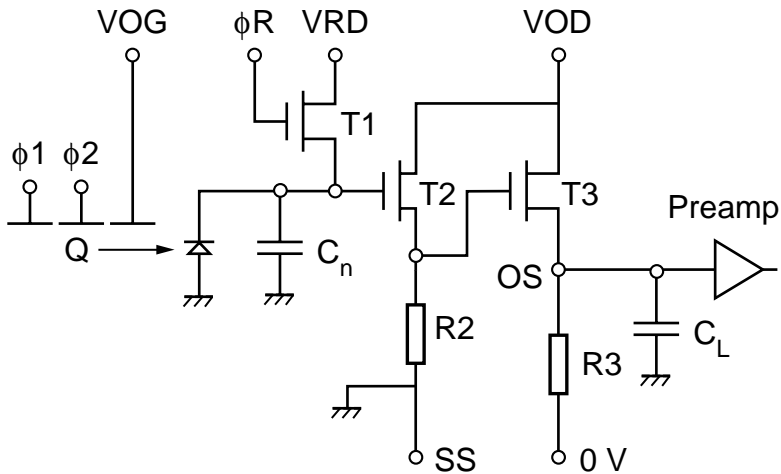


Fig. 13

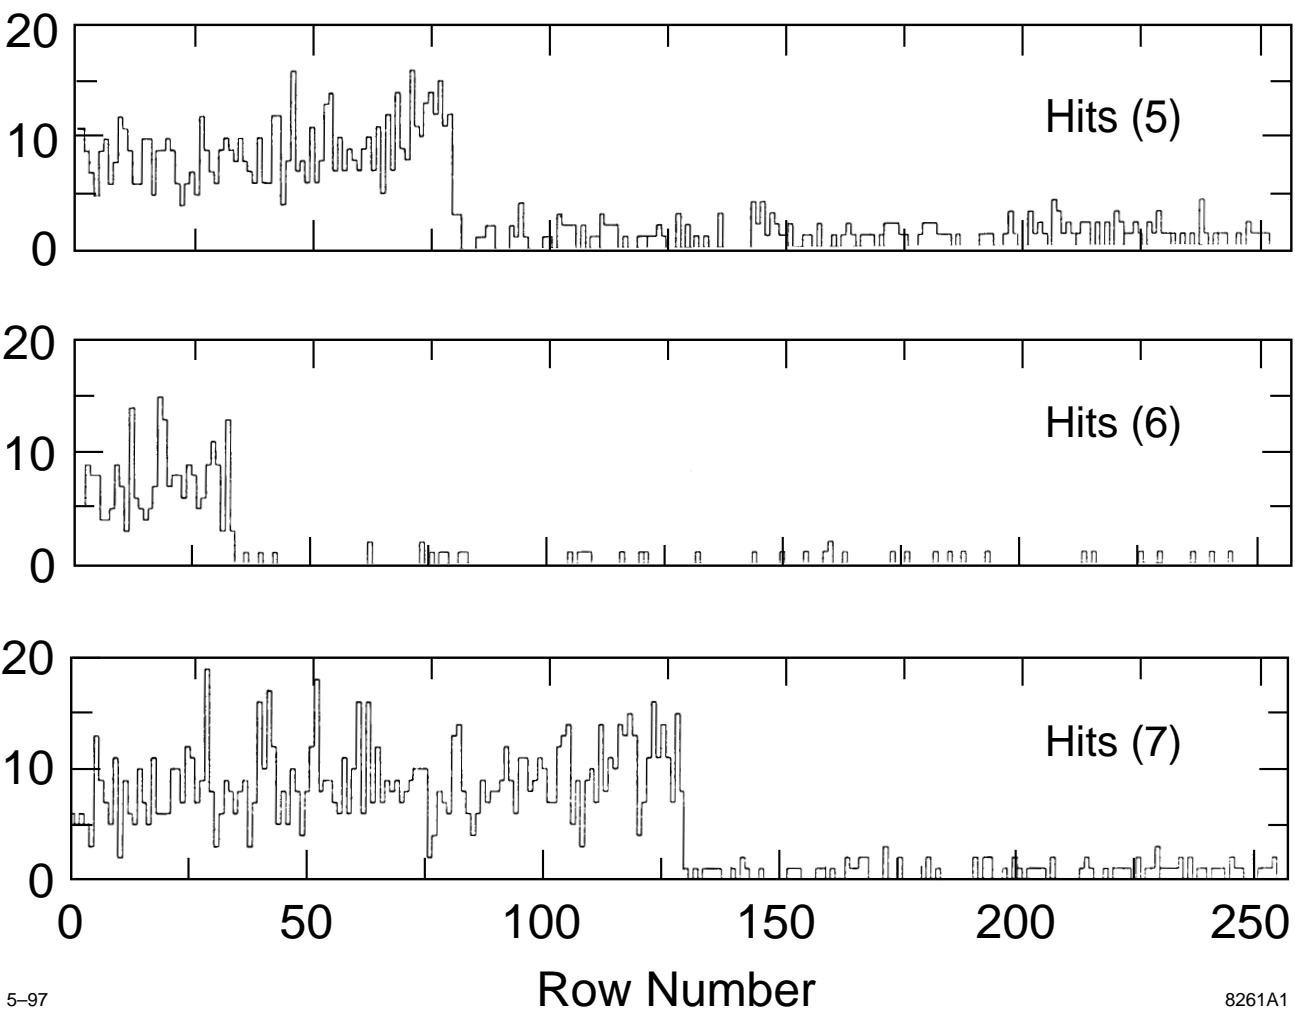
**Fig. 14**



Fig. 15




Fig. 16

Ladder Survey Data – Top View




Fig. 17a

Ladder Survey Data – Side View





Fig. 17b

VXD3 (SLD)



1-97
8262A02

Fig. 18

Barrel Survey – End View




Fig. 19a

Barrel Survey – Side View






Fig. 19b

All Dimensions in mm



North CCD



South CCD

Fig. 20

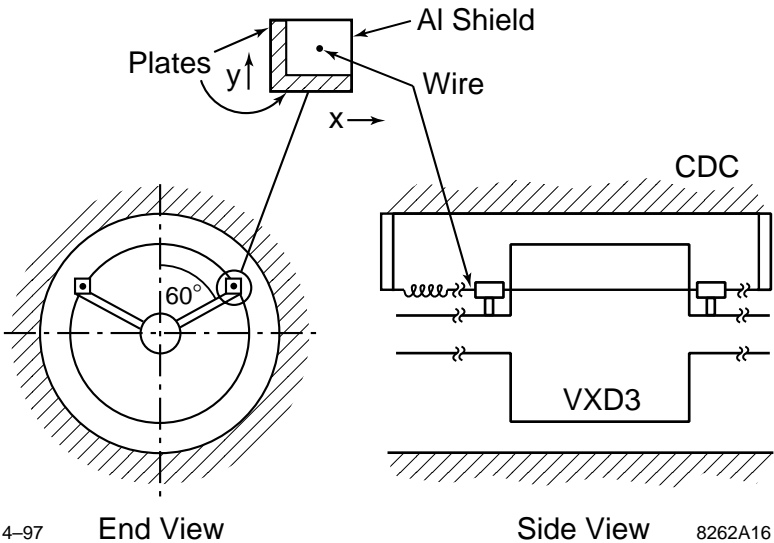


Fig. 21




Fig. 22




Fig. 23

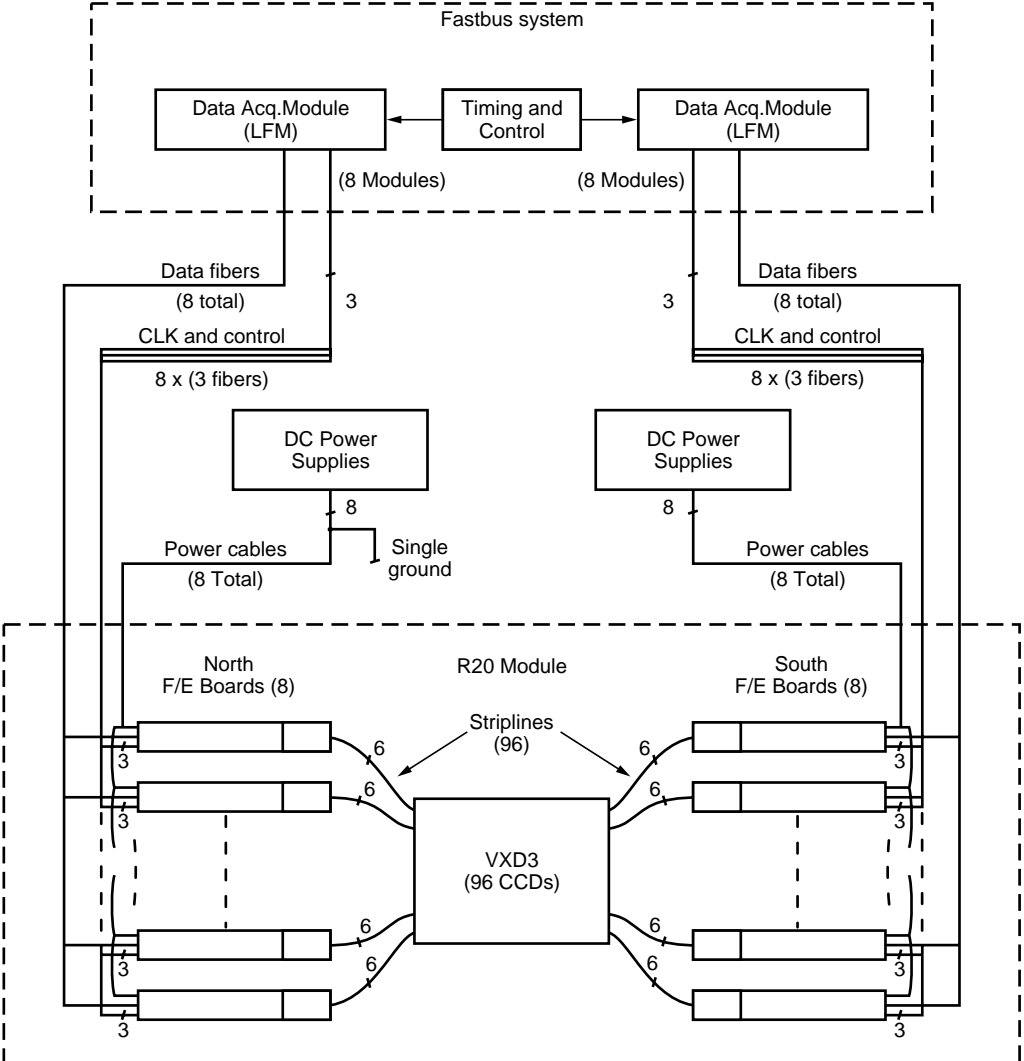


Fig. 24

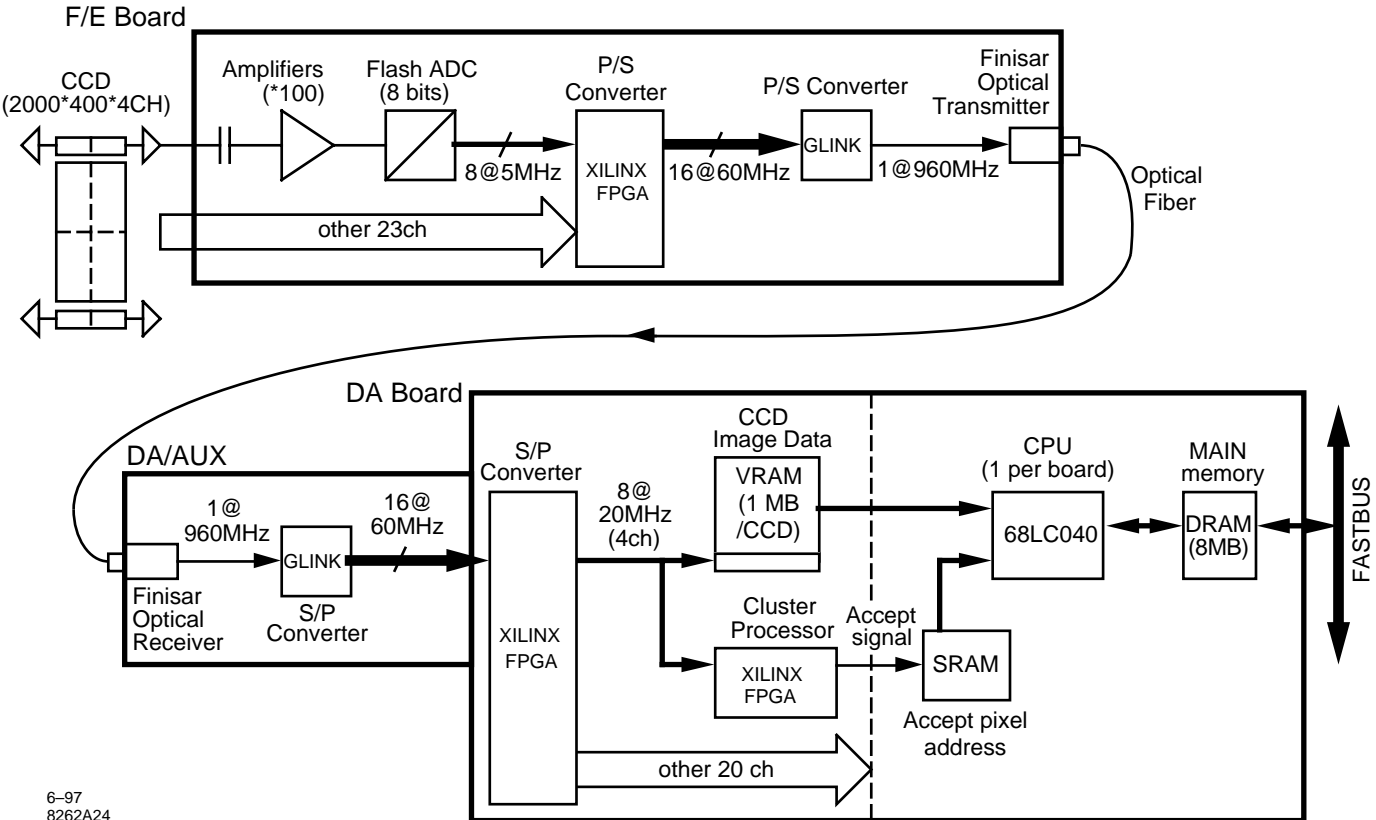
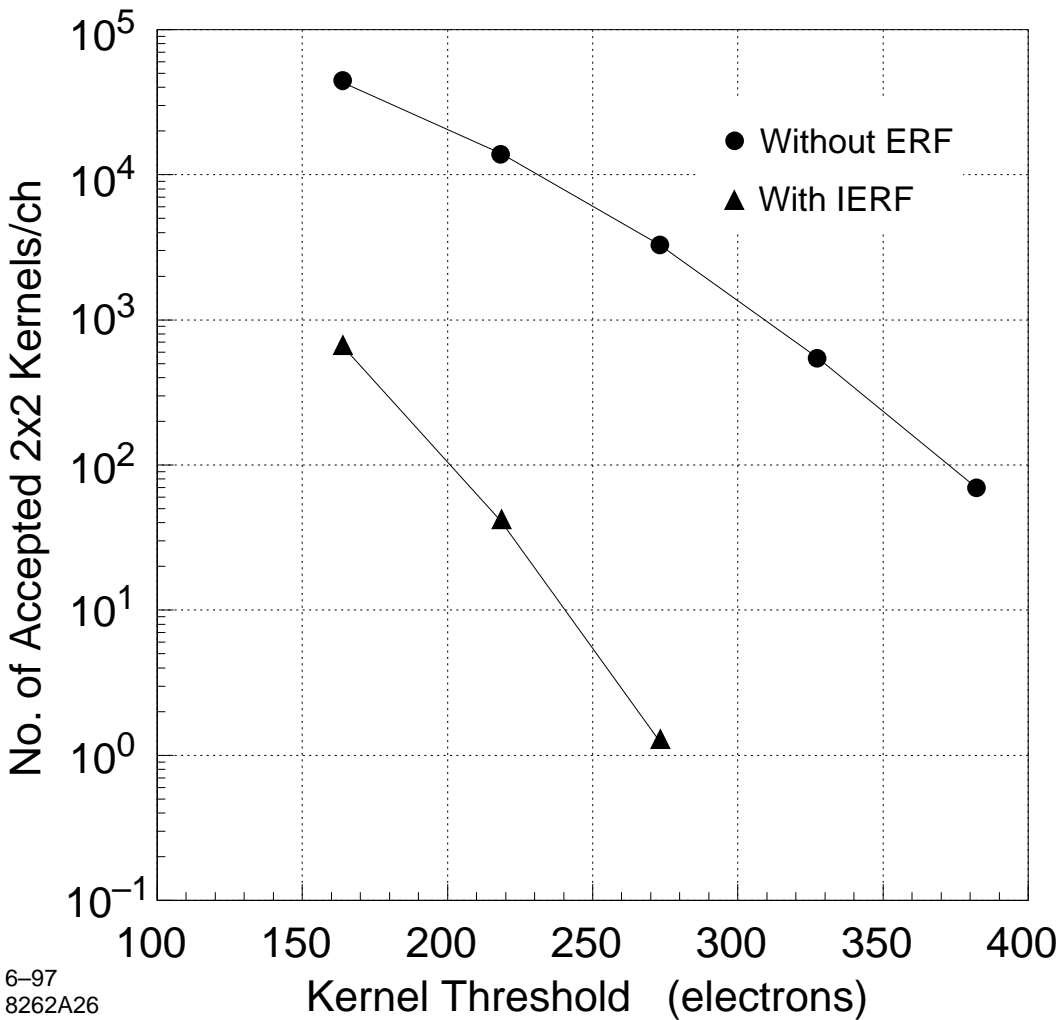


Fig. 25



6-97
8262A26

Fig. 27




Fig. 28

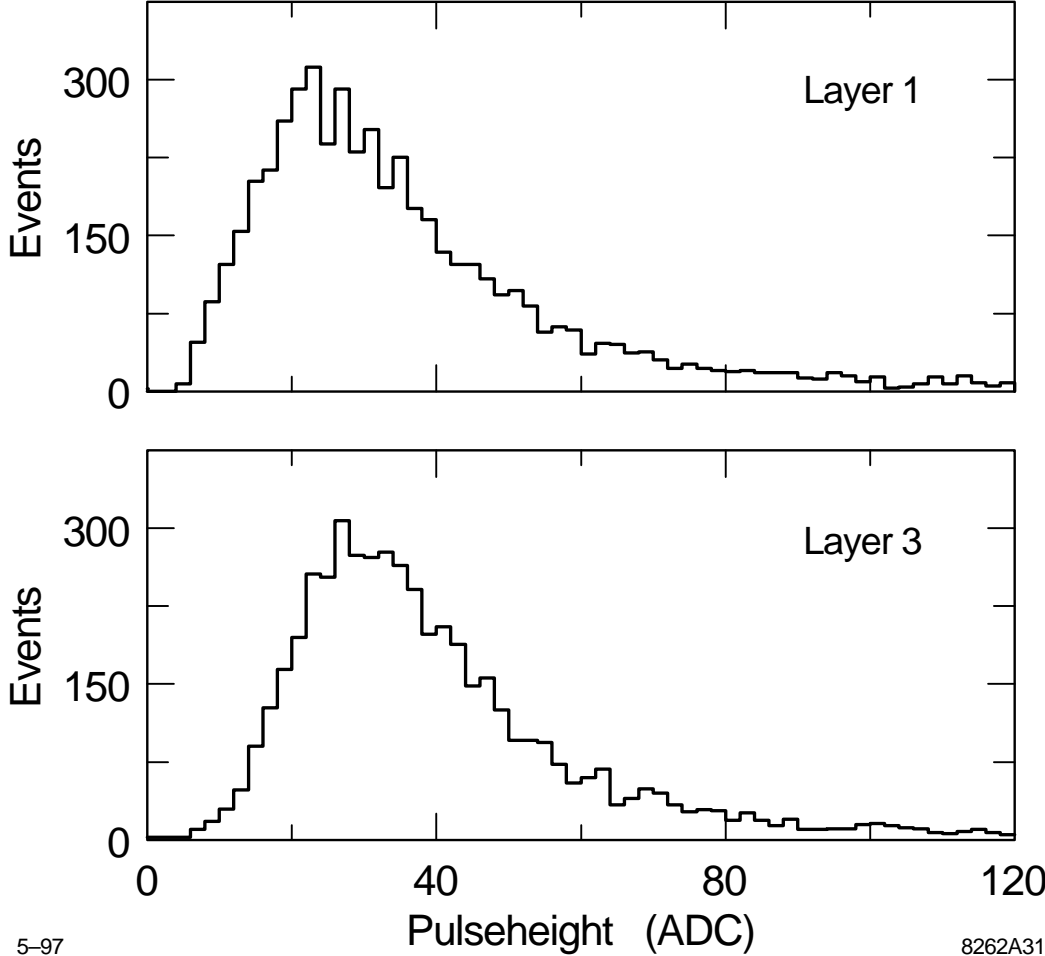


Fig. 29





Fig. 30






Fig. 31

Doublet $r\phi$ Residual



Doublet z Residual






Fig. 32

2 Prong $r\phi$ Miss Distance



2 Prong r_z Miss Distance



7-97
8262A30

Fig. 33






Fig. 34



6-97
8262A20

Fig. 35

**Fig. 36**