ILC Vertex Detector R&D
Report of Review Committee

1. COMMITTEE MEMBERS

Members of ILC Detector R&D Panel

Chris Damerell   RAL, UK
Hwanbae Park   Kyungpook U, Korea (chair)

External Consultants

Dave Christian   Fermilab, USA
Gerhard Lutz   MPI Munich, Germany
Masashi Hazumi   KEK, Japan
Pavel Rehak   BNL, USA
Petra Riedler   CERN, Switzerland
Steve Watts   Manchester U, UK
Yasuo Arai   KEK, Japan

Chair of ILC Research and Development Board

Bill Willis   Columbia U, USA

Regional Representatives

Europe: Chris Damerell   RAL, UK
Asia: Not available for this review
America: Tim Bolton   Kansas State U, USA

Local Vertexing Experts

Simon Kwan   Fermilab, USA
Lenny Spiegel   Fermilab, USA

Admin Support

Naomi Nagahashi   SLAC, USA
Maxine Hronek   Fermilab, USA

2. EXECUTIVE SUMMARY

The Vertex Detector Review Committee studied the programmes of 10 groups or collaborations, encompassing 9 distinct technology options now under development, as well as important general issues such as geometry options (long barrels vs short barrels plus endcaps), electromagnetic interference, and issues of machine background. Brief conclusions and recommendations for all these activities can be found in Section 6. Despite 10 years of R&D, we are still some way from the goal of knowing how to build a vertex detector that will satisfy the physics challenges. In fact, the world-wide R&D is still expanding, in response to these challenges which have attracted some of the finest detector physicists in the HEP community. Some of the groups are confident of delivering full-scale detector components (ladders) in test beams by 2012. Others will take longer, and see the opportunity to build upgrade detectors later. Given the relatively modest cost of a vertex detector system (measured in millions of dollars, not tens of millions), and the physics benefit likely to accrue from some of the suggested upgrades, this approach is certainly justified. So the most important recommendation of the committee is to urge ongoing support, at the current levels or higher if possible, for the R&D groups. Furthermore, the particle physics vertexing community world-wide is coherent and self-organising, so that everyone learns from the others in workshops and conferences.

Having said this, the ILC situation really calls for a level of coordination beyond what is currently in place. The case for this is argued in Section 6.9, and reflects conclusions which have emerged from each of the reviews held during 2007. The one-sentence summary of the recommendations of our first review, held in February 2007, read as follows: *'Form a Tracking Coordination Group to coordinate the completion of the R&D programme, so that the choice of tracking technologies for ILC detectors can be made on the basis of these results'.* While this hasn't happened, the case seems to the Detector R&D Panel to be as strong as ever. The emergence from that review of a new idea (for a silicon pixel tracker - SPT) is a good example. This technology could open the door to a tracking system with greatly reduced material budget that would be ideal for ILC physics. Without a coordination group, new ideas such as this may founder due to the political strength of the currently embedded technologies, making it difficult to win funding to explore new ideas. In the spirit that 'unity is strength' and the fact that again we are far from establishing the preferred technology (or technologies) for ILC detectors, we repeated this recommendation for calorimetry. In our opinion, the same applies to vertexing. Gathering hard information about actual performance for these many sensor options is complicated, and we suggest that the best way to achieve this will be to form a Vertexing Coordination Group (VCG). Together with the TCG, CCG and TBCG (for test beams), this structure could optimise the evolution of the R&D to the point at which experiment collaborations, when the time comes for their formation, will be able to choose the best technologies for their subdetector systems. The Vertexing Coordination Group may be something for the ILC Detector Directorate to consider implementing. One of its most important functions would be, working with the ILC Directorate, to negotiate common test facilities somewhere in the world, including a test beam with appropriate time structure, and other facilities that will allow quantitative comparisons between ladder prototypes, rather than relying on uncoordinated studies made with different test setups at different labs. Finding one laboratory (it more or less has to be CERN or Fermilab) prepared to offer support for this facility over a number of years, say 2010-2015, could be one of their most important contributions to the realisation of our common goal - vertex detectors that can do the job demanded by the physics.

3. INTRODUCTION

The WWS-OC asked the Detector R&D Panel to organise reviews of Detector R&D, to be held in conjunction with the regional workshops (3 per year). This report documents the recommendations from the third of these, on ILC vertex detectors. The aim of these reviews is to understand the activities of the numerous groups undertaking ILC detector R&D round the world, and to assist them to follow a timely path to completion of their work. In the case of vertex detectors, for each R&D collaboration, completion of work is defined as:

evaluating the performance of an assembly of full-scale ladder prototypes with respect to all relevant criteria, in sufficient detail to determine the performance of a full-scale detector constructed with this technology.

- 'an assembly of ladders' means of a size sufficient to permit reliable extrapolation to the full-scale detector. Details to be worked out with the collaborations.
- 'all relevant criteria' includes tracking precision, mechanical stability, effects of operation in a 3-5 T field, effects of coarse beam time structure (1 ms on, 200 ms off), effects of pulsed power including mechanical vibrations, cooling system, material budget, sensitivity to electromagnetic interference, etc. Details to be worked out with the collaborations.

There was agreement between the committee and collaborations that completion of the work as defined above, should be achievable by about 2012, at current funding levels, for some of the technologies. More ambitious technologies could reach maturity on a longer timescale, and be considered for future upgrades.

4. REVIEW PROCESS

Each review has followed a similar pattern. Firstly, an expert committee was formed from members of the Detector R&D Panel and expert consultants from outside the ILC community. The GDE has generously provided financial support to cover the expenses of these consultants. Secondly, guidelines for each review were published on the ILC Detector R&D Panel website at http://www.linearcollider.org/wiki/doku.php?id=drdp:drdp_home and these provided the basis for preparations by the R&D groups. Reviews have been held in conjunction with the ILC regional workshops (one per region per year), during the parallel sessions of those workshops. This has several advantages, the most important being to permit everyone including the consultants to attend the workshop plenary sessions, thereby getting an up-to-date picture of overall ILC developments. It also guarantees a good attendance from the community at the open session presentations which, along with the written reports provided beforehand by the R&D groups, form the backbone of the review. However, there is the disadvantage of further overloading the parallel sessions. While there was not unanimous support for holding these reviews during the regional workshops, there was a majority in favour, particularly from our esteemed consultants.

In the guidelines to the collaborations, we asked them to provide an overview of their goals, starting from where they are now, up to the completion of their R&D programme, ready to start construction with a frozen design and proven manufacturers. By common consent, the completion of the R&D to the point where an experiment collaboration could responsibly propose to use their vertex detector technology for the startup of ILC, should occur approximately in phase with the completion of the Technical Design for the machine, currently planned for 2012. Therefore, we were very interested, for each collaboration, to learn how they plan to reach this stage of development on this timescale.

Different groups use somewhat different terminologies. In this report, we aim to follow a consistent terminology, which in areas of possible ambiguity is as follows:

VXD: vertex detector (referred to by some groups as vertex tracker)

Barrel layers: layer 1 to layer N (innermost layer referred to by some groups as layer 0)

Endcap disks: disk 1 to disk N (disks with few micron precision, not to be confused with forward tracking disks with typically 10 micron precision)

Ladders: sensor-wide full length barrel staves or disk components. Ladders may contain single sensors, pairs of sensors, or multiple sensors (mosaics) along their length. Conventionally, for a pixel system the 'columns' run along the length of the ladder, with the 'rows' being normal to the 'beam direction'. Orientation of readout is here described as column-parallel if the signal processing edge logic is at the ends of the columns, and row-parallel if it is at the ends of the rows.

3-D detectors: A development from SOI devices; detectors in which the sensor consists of a depleted region of a high resistivity 'handle wafer', bonded to two or more tiers of CMOS electronics, all interconnected by few micron diameter metal-filled vias (not to be confused with the 'Sherwood Parker' 3-D sensors in which charge is collected by transverse drift to column electrodes that extend through the full thickness of depleted silicon, though these are made by similar technologies). The technology is also becoming known in commercial electronics as stacked devices, as in the term 'stacked memories' from Intel.

DSM: 'deep submicron', CMOS technology in which the minimum feature size (gate width) is 0.25 μm or below. The most advanced discussed for ILC VXDs is 45 nm, with industry seeing a possible end of the road at 11 nm.

EMI: electromagnetic interference, from different sources, typically beamline and detector systems at any time, and beam-related RF during the bunch trains

SW: sensitive window; the effective resolving time during the bunch train for a particular VXD technology, varying from single-bunch timing for the most precise, to integration through the entire 1 ms bunch train for the most inclusive

CDS: correlated double sampling, a form of time-variant filtering comprising baseline restoration and noise suppression, widely used in imaging systems, developed originally to eliminate ktC noise arising from the node reset in CCDs. Can be implemented in different forms with very different noise suppression capability: see Section 6.4 for clarification.

Stitching: production of monolithic devices larger than the reticle size by precise alignment of fields on the wafer (eg 6x1 stitching can produce single devices of size ~12x2 cm^2). 2-D stitching is also possible, but it has not been contemplated for vertex detectors.

5. OBSERVATIONS AND RECOMMENDATIONS (SPECIFIC PROJECTS)

5.1. MIMOSA Collaboration (Marc Winter et al)

This group has an excellent track record for development of CMOS sensors for several HEP applications, leading steadily towards the most challenging goal, ladders appropriate for ILC conditions. Their written report is welcomed for its clarity. There are possible showstoppers, and these are clearly stated.

While they refer always to CMOS sensors, they are in fact restricted to NMOS transistors in their pixels since their n-type implants are reserved for sense diodes. However, where they include integrated signal processing edge logic, already implemented in MIMOSA-8 and -16, they have the full CMOS capability at their disposal.

A recent change in their planning has been to avoid the requirement for large area stitched devices (of ladder-length or half-ladder length) due to their experience with yield of smaller devices (they quote typically 30-40% for 3 cm^2 devices). Ladders made with a mosaic of many reticle-sized sensors may be only marginally thicker than 1- or 2-sensor ladders made with stitched devices, but it is not obvious that servicing ~10 devices along the ladder length can be done without a significant increase in material budget. There is furthermore the problem of determining the positions and orientations of ten times more 'postage stamp' sensors. They are confident that edgeless dicing and advanced chip placement will permit negligible (few micron) gaps between the active pixels of neighbouring sensors.

With the available foundries, they have difficulty defining or sometimes even discovering after the event, process details such as epi layer thickness, dopant concentration and profile etc. They have, after trying 7 different commercial processes, established that the AMS 0.35 OPTO process best satisfies their needs at present. One could consider the alternative (more expensive) approach of working with a foundry which would allow them to define their process requirements, hence obtain more predictable results. A problem recently encountered (low c.c.e) may be explained by the hypothesis of excessive diffusion of the p+ implant close to the sense diode, but this may not be the correct explanation. If they were able to simulate the manufacturer's process details, the room for speculation would be reduced. However, their choice of inexpensive multi-project wafer runs in which they may be unable to establish the detailed processing parameters, is certainly defensible in terms of throughput and gaining valuable experience from many device variants. Such decisions are always balancing acts, and their approach may in the end prove more cost effective than working with a foundry that supports processing fully specified by the customer.

They set themselves the target for tolerance of ionising radiation of 150 krad/year, allowing a safety factor 3 wrt calculated backgrounds, and have produced devices that perform satisfactorily to 1 Mrad. With modest cooling, the sensors perform well after 10^{13} 10 MeV electrons, demonstrating their robustness wrt the expected pair background. This excellent performance will need to be sustained as they move to fully integrated electronics, so it remains a challenge. They have generally presented results at room temperature, where their sensors perform well. However, the radiation tolerance, for example regarding dark current, would be considerably eased by operating at reduced temperature, say -20 C. It is possible that the tracker (if based on silicon strip or pixel sensors) will also prefer cold operation, so the entire vertexing/tracking system may benefit from operating in the same gas envelope.

In order to achieve their goal of a 25 μs sensitive window (which in their case is the same as the readout time) they adopt a 'row parallel' as opposed to 'column parallel' architecture, meaning that they read transversely to the beam

direction. This has the disadvantage that their readout electronics must lie inside the active volume of the barrels, instead of being at the ends. The column parallel option is generally preferred since it puts the material associated with the readout electronics outside the volume in which ultra-precise tracking is required. However, they expect that the dead material along one edge of each ladder can be restricted to a width of about 2 mm (1 mm for 5-bit ADC, 0.5 mm for data sparsification and 0.5 mm for on-detector memory). They are pursuing a number of ADC design options, and the final outcome in performance and mechanical dimensions, within the available 20 μm pitch, is yet to be established. The larger these dead regions become, the greater the impact on performance for low momentum tracks. They are hoping to have sensors with integrated electronics operating this year (2008), incorporating the nearly-full chain of in-pixel amplification and CDS, with at the edge discriminators (not yet ADCs), zero-suppression architecture and output memories, which will be a major milestone for this approach. With active ladder widths of 5 mm for layer 1, they will read the 250 pixels across the width at 10 MHz in order to achieve their sensitive window. This readout will include frame-rate CDS, with 25 μs cycle time between successive samples of the same pixel. The committee was somewhat confused by plots of noise vs frequency, which we now understand means clock frequency, which is 16 times the pixel sampling frequency. Given that the maximum pixel sampling frequency is 10 MHz, the absence of significant noise degradation is as expected. They estimate that their 25 μs sensitive window will allow the layer 1 radius to be reduced by 15-20% compared with designs that integrate for 50 μs, which would certainly be a useful gain as regards impact parameter precision. However, work is only now beginning to evaluate the physics impact of a given hit density in layer 1, so such estimates are at present only provisional. See Sections 5.6 and 5.8 for alternative approaches.

For the outer layers, they propose larger pixels (33 μm square for layers 3-5). While it is true that the average hit density falls with radius, long-lived Bs and Ds frequently produce high hit densities on all VXD layers. Detailed physics studies will be needed, particularly of such cases as measurement of vertex charge with high energy jets, to quantify the penalty of such a relaxation in pixel size. While they have demonstrated that there is very little loss in precision with the larger pixels, the diminishing 2-hit resolution could be a significant disadvantage for the reconstruction of long-lived decays. An advantage of the graded-size approach is reduced overall pixel count (0.3 Gpixel in contrast to ~1 Gpixel for a similar design with uniformly 20 μm pixels), with consequently lower power dissipation. However, merged hits on any layer make it useless for precision measurement of the affected tracks, so such effects do need to be kept to a low level.

The estimated power dissipation of 600 W during readout can be reduced to 30 W, easily handled by gas cooling, by operating at a 5% duty cycle. Since the ILC bunch train timing allows in principle 0.5% duty cycle, they have a considerable safety factor, some of which will be needed for settling time of the electronics on power-up, for which they do not yet have estimates. This will become apparent with their first sensors having a full system of integrated electronics.

At the time of the review, this collaboration estimated that they would have complete ILC-compatible ladders operational by 2010, ahead of most other options. While we are informed that this will not now be possible, a very important milestone will be the assembly of the STAR HFT, which will represent the first use of MIMOSA sensors in a physics experiment. However, for the very challenging ILC application, they recognise a number of possible showstoppers, such as performance limitations due to the limited space available for integrated electronics. Should

such problems prove to be insuperable, they see the way forward in 3-D sensors, for which these space constraints would essentially disappear. However, this approach has its own considerable challenges, as discussed in Section 5.7.

5.2.   Deep n-well sensors (Valerio Re et al)

This approach takes advantage of deep submicron processing (130 nm and below) to provide in-pixel data sparsification and time stamping. Readout can then take place in the quiet periods between bunch trains, when problems of EMI will be minimised. We were impressed with the writeup and presentations by this group, particularly the extent to which they have simulated critical aspects of the device performance, notably as regards noise. The deep n-well is used in modern triple-well CMOS technologies to shield NMOS devices from noise coupled through the substrate. This group takes advantage of this feature and uses it for sensing the signal charge generated in the epitaxial substrate. In their case, using a 0.13 μm process, they follow the standard triple-well approach of integrating their NMOS transistors used for analog processing in a p-well inside this signal-sensing deep n-well. However, their digital processing requires additional NMOS transistors (which can conveniently be integrated in the p-epi material), and also PMOS transistors which require their own little n-wells, since they would disturb the signal sensing if located in the deep n-well. These additional n-wells unavoidably compete for signal charge collection with the deep n-well, so should cover as small a fraction of the pixel area as possible. In their current prototype, they achieve a 2:1 ratio between the wanted and unwanted n-well areas, which may or may not suffice for efficient min-I detection. Their first large prototype SDR1, 256x256 pixels, will be tested in a beam in 2008, at which time this issue will become clear. In the meantime, they showed results of simulations in their presentation (not in their written report) based on the 3-D package ISE-TCAD which allows them to explore the charge collection irregularities in their structure. These simulations are laborious, but since their charge collection is mainly by diffusion from undepleted silicon, it has been possible to develop a simpler approach in collaboration with Dave Christian. It will be very interesting to see quantitative results on the irregularity in charge-collection efficiency over the pixel area, and hence of possible hit inefficiencies taking account of realistic threshold settings. Intuitively, there might be significant problems.

They take advantage of the very low occupancy during a full ILC bunch train. There is no need to include a pipeline for storing more than one hit in any pixel, so their in-pixel logic consists firstly of a charge-sensing amplifier (which is designed to achieve low noise in spite of the relatively high capacitance of the large area deep n-well (~113 fF)). The amplifier is followed by a discriminator to record the hit, and a 5-bit register to store the hit time with a precision of ~30 μs (1/32 of the bunch train). (The encoded time within each bunch train is distributed over the array from edge logic.)

In order to achieve the desired few micron spatial resolution in this system with binary readout, the pixel size should not exceed ~15 μm square. Their first prototypes will not quite meet this goal, but they have several good ideas to reduce pixel size in the future (by progressing to a 90 nm CMOS process and implementing a more compact analogue time-stamp). If there are efficiency problems for min-I particles with their first prototypes, these developments could improve the situation by increasing the pixel fraction covered by the deep n-well.

Despite the significant amount of in-pixel logic (~50 transistors) they profit from the relatively relaxed timing requirements and are able to keep the power dissipation per pixel to 4.7 μW. By power cycling the signal-sensing

electronics to be off between bunch trains when they are reading the stored data, they expect to keep the average power below 20 W, appropriate for a gas-cooled system.

From tests on a related chip (Apsel) and from other experience with DSM devices, they are confident of meeting the ILC radiation tolerance requirements.

We wonder about the feasibility of producing large area stitched devices with acceptable yield, using the 130 nm process or below. It seems that they might encounter the same limitations as described by the Strasbourg group, with ladders needing to be made of ~10 reticle-sized devices. While this is undesirable, it could provide an acceptable compromise. Much more R&D will be needed to clarify all the associated system aspects.

5.3. CAP devices (Gary Varner et al)

While primarily directed at developing a vertex detector for an upgraded *B* Factory, where background rates will make fast-readout pixel-based vertex detectors obligatory, this collaboration is also making important contributions to the requirements for ILC. They are following two main approaches, MAPS-based and SOI-based.

For their MAPS-based devices, the main difference compared with the MIMOSA family of devices is the use of in-pixel analogue signal storage on a string of capacitors (10 or more storage elements per pixel). In the ILC context, this would allow time slicing of the bunch train, with consequent reduction in the density of background hits.

In moving from small test devices, they have encountered two general types of noise problem. Firstly, and this they suggest applies generally to MAPS devices, the system noise increases with readout speed (which is obvious) and also with device *length,* in column readout architectures, is to be expected due to increased off-pixel capacitances, but they also see a dependence on the row length (device *width*). This reflects the complexity of these devices, for example the fact that increased width means more tenuous connections to voltage supplies, increased parasitic capacitance, and other effects. Their analysis is not sufficient to separate all these effects, but their results flag up a warning for groups which have so far focused on small prototypes. Reading out their device at 10-20 MHz, they find ~50 e$^-$ rms noise for a 1 Mpixel device. A typical ILC sensor might comprise 6 Mpixels, so there is cause for concern. The second problem is specific to the pipeline CAP approach, and relates to the noise associated with each in-pixel storage cell. To increase the number of cells for a given pixel size (say 20 μm square) the capacitance of the cell must be driven down. While this may be possible by moving to a deeper submicron process, the noise associated with very small storage capacitance increases rapidly. There are also concerns about EMI sensitivity of CAP devices under ILC conditions during the bunch train, which will need to be studied once full-scale ladder prototypes become available.

These noise problems could be reduced by increasing the signal above that typical of standard MAPS devices, where the epi layer is typically only a few microns thick. For this reason, they are actively pursuing the SOI process, which is discussed in Section 5.7. Like the Fermilab and LBNL groups, they are working with OKI Semiconductors in Japan, and have begun work with a 0.2 micron process.

Whether this group will end up developing large-scale CMOS sensors with in-pixel storage, or SOI devices with fast readout and/or in-pixel storage, remains an open question, to be determined by their ongoing R&D work. They point out that moving to large-scale devices is expensive, so one wants to learn as much as possible by simulation and

production of small prototypes before making that commitment. Overall, they are making excellent progress and pushing the CMOS approach in novel directions which might well prove to have unique advantages for *B* factories or ILC or both.

5.4.   DEPFET Collaboration (Laci Andricek et al)

This powerful collaboration has a balanced programme, having made major contributions over many years to all aspects of the ILC R&D connected to the vertex detector. The DEPFET technology they invented (Kemmer and Lutz, 1987) has a distinguished track record, and is configurable for a wide variety of scientific applications. Their ILC work is aimed at one of the most challenging of these applications, since it requires the largest arrays of the smallest pixels, the most severe constraints on material budget, and the fastest readout, together with minimal power dissipation. The only respect in which it is less demanding than some other DEPFET applications is in terms of the radiation tolerance, where the requirements are modest.

Their R&D programme is divided into sensor development, associated electronics development, mechanical R&D, test beam studies and detector simulations. All this work is of an extremely high standard, demonstrating that they are on track to achieve their ambitious goals.

For the sensor design, the goal of 25 μm square pixels is largely established. The readout mode is 'rolling shutter', with a frame readout time of 50 μs, which could be achieved with a 40 MHz row readout rate. That might be beyond the DEPFET capability, given the other constraints, and an attractive option is to read row pairs synchronously, reducing the clocking rate to 20 MHz. This doubles the number of readout channels, but that appears to be feasible. The noise performance with small sensors is excellent, but in order to scale up, it would be advantageous to increase the gain of the integrated amplifier beyond the present value of 500 pA/electron. The gain could be doubled by reducing the device gate length from 4 μm to about 3 μm, which again seems possible.

There have been problems achieving uniform signal charge collection, due to the fact that the 'clear' electrode acts as a sink for signal electrons generated by particles traversing this region of the sensor. This can be solved by introducing a deeply implanted p-layer to shield this region, inducing most of the charge to drift to the internal gate.

The gate oxide used in the DEPFET is relatively thick, so radiation effects are significant. However, given the benign ILC radiation environment, this is not a problem.

The evolution in sensor design is impressive, addressing the remaining performance problems, and moving forward to full-scale devices. However, there is a potential noise problem in large devices due to the capacitance of the drain lines which run the full length of the columns. This is exacerbated by the need to move to thin active layers (not greater than 75 μm) and hence smaller signals, in order to achieve the needed precision in measuring oblique tracks. It may be that advances in readout chip design will permit acceptable S/N performance to be preserved. The collaboration has in addition devised a strategy of 'column segmentation' which would reduce the capacitance seen by the amplifier input.

At the time of the review, they expected have 5x1.2 cm^2 devices, proof of principle for layer 1, by the end of 2007, thin sensors (which involves an SOI process discussed in relation to the mechanical design) by end of 2009, and complete ladders by 2010/2011. This schedule would ensure that the DEPFET approach would be proven easily in time for consideration by experiment collaborations as they form.

The associated electronics (steering chips and readout chips) are also challenging, and form a substantial part of the R&D programme. The steering chips have two functions, to control the row-enable during readout, and to provide the clear pulse, so generating the sequence *read/clear/read* needed for true CDS. The clear pulse needs to be fast, to satisfy the rate requirement, and of large amplitude (~10 V), to completely remove the signal charge from the internal gate. If clearing is incomplete, this introduces a type of reset noise into the system. There are ideas to modify the DEPFET design so as to reduce the amplitude required for the clear pulse, thereby diminishing the problem. While chips capable of switching such large amplitude pulses are expected to still be available into the long-term future, for applications in industrial and automotive systems, they necessarily have thick gate oxide and hence may not be sufficiently rad hard. So the combined optimisation of sensor and switcher design is a subtle process in which the collaboration is making very good progress. As well as the chip design, there is a further non-trivial requirement associated with the fact that these chips sit in the active volume for precision vertexing, so material budget is important. It is expected to be able to thin them to approximately 50 μm. While this in itself is established technology, the further requirement to bump-bond these chips to the sensors may be challenging. At present, they use in-house gold stud bumps, but consider going to a smaller scale, less mechanically forceful process when using the thinned chips. Advanced interconnect technology is an area in which industry is making rapid progress, so this is likely to be a soluble problem.

The readout chips at the ladder ends are beyond the active detector volume, so do not have the same thickness constraints, but the performance requirements are challenging. The original CURO (current readout) family of chips work beautifully with small sensors, but performance naturally degrades as the readout rate is increased, and also as the sensor size (length) increases, due to the greater input capacitance. Excellent progress was reported in the development of a next-generation readout chip, the DCD (drain current digitiser) which may achieve the required noise performance, particularly given the reduced bandwidth associated with the 'double pixel' sensor design. There are some very talented chip designers working on this, and the problems are likely to be solved within the next few iterations.

For the mechanical design and prototyping of ladders, this collaboration is following a truly novel and promising approach. The idea is to eliminate thermal expansion mismatches by an all-silicon ladder structure, combining the thin sensors with thick 'window frames' to provide adequate mechanical rigidity. The ladders can be engineered by an SOI-inspired approach, where the thick handle wafer is selectively etched to leave the thin sensor chips exposed over the full area other than the support frame. The frame also provides a robust base for mounting the steering chips along the long edges, and the readout chips at the ladder ends. This work has already reached a high level of development, mostly using dummy sensors, and the extension to full-sized working ladder assemblies is beginning to come into sight. There is a good chance that the overall objective of 0.1% X_0 average thickness over the active length of the ladders will be achievable with sufficient stiffness that, when combined with a well-designed support shell, the overall mechanical stability will be adequate.

The simulation by this collaboration of a DEPFET-based vertex detector has reached an advanced stage of development, more so than for any other sensor type. The charge generation in the sensitive layer is simulated as function of track angle of incidence. This is particularly important since their overall detector design, following the approach in the TESLA TDR, is for long barrels giving excellent polar angle coverage, without endcap disks for vertexing. The collaboration has developed these studies to investigate the track reconstruction capability of their 5-layer detector, both with and without the dominant pair electron background. Their studies have revealed that such a

detector system will have excellent performance (high efficiency and low fake rate) for track momenta down to ~200 MeV/c. It may be possible to reduce the low momentum limit by using information from the forward disks of the main tracking system, but this depends on design details of that system, which are not yet settled. The results from the simulations by this group are world-leading, highly significant and deserving of long term support, as the interplay between detector design and physics studies advances.

The collaboration has also undertaken an impressive test beam programme, culminating in a 5-layer telescope in a CERN test beam. These studies have demonstrated excellent (few micron) point measurement precision, and have also provided data on cluster shapes and cluster signal charge as function of track angle, which have been compared with the simulations mentioned above. Excellent agreement has been demonstrated, proving the reliability of the simulations of these fundamental quantities, and hence giving confidence in the higher level studies they have undertaken.

Overall, the DEPFET collaboration is strong and composed of extremely high quality people in all areas of their comprehensive R&D programme. They may take a little longer to achieve their ambitious goals than they would like, but in all circumstances they are extremely well placed to have full-scale prototype ladders, satisfying all the ILC requirements, operating in a test beam by the 'official' target date of 2012. They also have interesting ideas for upgrading the technology, still using DEPFETs but benefiting from the emerging 3-D technologies, discussed in Section 5.7.

DEPFET sensors are already produced on 150 mm wafers, and previous work by this collaboration has demonstrated the feasibility of full-scale ILC sensors, which puts them in a strong position in relation to some alternatives.

5.5.   LBNL Group (Marco Battaglia et al)

This excellent international collaboration (including groups from INFN Padova and Torino) has developed a strong and balanced programme, making major contributions in physics simulations, sensor development (with two technologies), provision of telescopes in test beams, R&D on mechanical structures as well as on power distribution.

For the simulations, Battaglia has for many years been at the forefront of defining reference reactions that provide a balanced test of all detector features. The collaboration has now built up a full chain of simulations from the pixel response (for the case of monolithic CMOS devices) through to physics capability. They are now able to make quantitative studies of the impact of changes in detector parameters (eg layer thickness) on the physics. They will also make comparisons between their simulation code (eg for flavour tagging) with LCFIVertex and other packages as they become available. Their simulations need to include the predominant electron pair background, and here (following the lead of the FPCCD group) they will explore possibilities for cluster shape filtering, though this is likely to be less helpful due to their larger pixels. Their suggestion of cross-checking these background simulations with test beam data (they are working to cover the range 50 MeV to 1 GeV electrons at LBNL) will be of value to the whole community. Their simulation work will evolve over the next few years, as detector R&D becomes better able to define the realistic material budget associated with stably supported sensors and associated services.

Their 'baseline' sensor technology is CMOS pixels similar in architecture to the MIMOSA group. However, they are pursuing an independent design path. Given the complexity of such devices, experience (for example at LHC) has

shown the value of separate R&D approaches by independent collaborations which keep in touch with each other. Their approach involves rolling shutter readout with a frame period of 25-50 μs, in-pixel CDS (with limited functionality due to the fact that successive samples are made at the frame-rate), on-chip digitisation and sparsification. Test devices made so far (using the same AMS 0.35 μm process as the MIMOSA group) represent an orderly buildup to these complete goals. One major milestone is the successful design and operation of 5-bit ADCs on a pitch of 20 μm and length only 1 mm. They also follow the 'row-parallel' readout approach, so their ADCs create an inactive band about 1 mm wide at the edge of the sensor, along the length of the ladder. They have already demonstrated capability for reading out at 25 MHz, giving a readout time for a 10 mm wide sensor of about 25 μs. Their goal is to demonstrate reticle-size devices (2x2 cm^2) with full functionality by about 2011. They will then decide which technology to push for full-scale ladders, and whether to populate these ladders with mosaics of such devices, or develop stitched devices of half or full ladder length. During the next few years, they will also decide whether to retain their present process, or to move to a more advanced process (possibly 0.18 μm design rules, if this becomes available with sufficient epi layer thickness).

Their 'alternative' technology is SOI-based, following the availability of commercial processing of such devices via KEK, from OKI using their 0.15 μm FD-SOI process, which has the advantage of permitting full CMOS functionality in the pixels. Combined with the finer feature size, this allows greater functionality and/or smaller pixels. Their goal is to develop 10 μm pixels with single bunch time-stamping capability. The precision requirements can then be met by binary readout, reducing the need for signal processing at the chip edge. Fortunately, such devices are of great interest for other applications, so the development cost does not have to be borne entirely by their ILC budget. They are among the first (probably *the* first) group to be operating SOI devices in a test beam. Not surprisingly, they have encountered problems. They suffer from the back gate effect discussed in Section 5.7. As a consequence, the output signal as function of detector bias behaves normally up to 9 V, but then saturates and even decreases as the bias is raised. This is the beginning of a long and interesting R&D programme.

Their mechanical R&D involves working with the Aptek company to back-thin wafers. They have made extensive tests of devices thinned to 40 μm and find the performance to be entirely unchanged, as it should be. Their ladder design starts from the STAR HFT prototypes, which should establish a new record for layer thickness of 0.3% X$_0$ They are exploring the possibility of 'cool-core' ladders, in which sensors are mounted on both sides (hence balancing any associated forces) and the space between is filled with RVC or graphite carbon foam. Such an assembly is relatively stiff, and by machining channels through the foam, it is possible to provide 'internal' cooling with gas flowing through these channels.

Power distribution options, with the implications for cabling, could be decisive in selecting candidate sensors for ILC. Their collaboration is extremely well-placed to work on this, since much of the international expertise is at LBNL, in the form of people working on this problem for sLHC.

This group has an excellent track record for building telescopes of detectors and operating them in low energy electron beams at the ALS facility at LBNL, and (more significantly) in the 120 GeV MTEST beam at Fermilab. If work proceeds as they hope, they will be in a position to choose which sensor technology they prefer, and move to the construction of an ILC 'demonstrator' system, consisting of maybe 5 fully equipped ladders operating in a test beam.

Overall, this group is doing excellent work in all areas, and is ensuring that the lessons learned from other advanced vertex detector programmes at LBNL are being communicated efficiently within the ILC vertexing community.

5.6.   SiD Collaboration (Su Dong et al)

The work of this collaboration is closely linked to that of the Fermilab group considered in the next section, where the mechanical details of their design are discussed. Their presentations here focused on an overview of their design philosophy, backed up by important studies of track-finding performance. Their approach is to a great extent technology-independent, in that they will choose the sensor technology which performs best when fully functional ladder prototypes have been studied in test beams.

By operating with a 5 Tesla solenoid field, they are able to locate their layer 1 barrel of length 10 cm on a radius of 1.4 cm, giving excellent coverage to $\cos(\theta) = 0.96$. They choose to keep their 4 outer barrels to the same length, obtaining 4 additional hits in the forward region by using disks of endcap sensors at z values of approximately 8, 10, 12 and 18 cm, enabling them to extend the polar angle coverage to $\cos(\theta)$ of 0.98. This enhanced coverage is one advantage over the long-barrel concept, and there are others. For tracks at low polar angles, it is easier to control the geometry of endcaps compared to barrels in which these tracks are highly oblique, so small uncertainties or instabilities in radial position can have serious effects. Also, charge depositions in the disks are free of Lorentz angle effects.

However, there is one particular disadvantage of disks, which is that the tracks they see have been obliged to traverse the mechanical supports and services at the ends of the barrel. Their studies are based on an assumption about the amount of extra material involved (equivalent to 2 mm thick pads of G10) which may eventually prove optimistic or pessimistic - the sensor R&D groups are at present far from being able to evaluate this. This problem is exacerbated by the use of pulsed power, which may induce significant mechanical shock to the structure, in view of the Lorentz forces on high-current cables servicing the detector. However, on the assumption that the material for supports and services can be held to these suggested limits, their simulations show the inevitable degradation in performance as function of polar angle, due at least to the increasing distance of the first measured point from the IP. By $\cos(\theta) = 0.94$, the impact parameter resolution on 5 GeV tracks (typical of particles in multi-jet events) has degraded by about a factor 2 to ~10 μm, by which point it becomes questionable whether the endcap sensors need to be of 'vertexing quality' or might be relaxed to more typical pixel tracker precision of around 10 μm. However, none of these detailed questions can be addressed until we have a better idea of what actual working ladders, in the different technologies, will really look like.

Their general layout, with the beamipipe cleanly supported over a length of about 1.8 m, gives confidence in the engineering realism of their thinking. Everyone assumes a thin-walled cylindrical section of beryllium pipe in the innermost region, but not all designs have considered the issues of stresses on braze joints during installation and removal. Their layout appears to satisfy these crucial requirements., even for the aggressive beampipe wall thickness of 0.4 mm, plus 25 μm titanium liner within the cylindrical section.

This group presented a new idea in the review, namely the possibility that one consider a different technology for layer 1. The argument is that the pair backgrounds are so much more severe for this layer, that it might pay to use a technology with a shorter sensitive window (possibly even insisting on single bunch timing) tolerating more massive services and/or higher power dissipation than would be acceptable for the entire vertex detector. There are special

factors associated with layer 1 (possibility of routing services and/or liquid cooling at small angles below the active volume for tracking) which make this a potentially attractive option. However, there are the usual good arguments for avoiding mixed technologies, so the need would have to be seriously demonstrated.

The updated studies of Takashi Maruyama in this collaboration provide a clear picture of the background to be faced, which is dominated by pair electrons outside the main envelope. There is now some confidence that other sources such as backscattering of electrons and multiple bounce shine from SR photons are relatively minor. However, the pair background is strongly dependent on the final focus design, being as high as 200 hits/mm^2/train for some unfavourable options to which the machine people may be driven if (as will surely happen in the early phase of running) luminosity falls short of the goals. Several technologies intend to divide the train into 20 time windows, giving a hit density of 10 per mm^2/train for this worst-case scenario. This group believes that the upper limit for comfort, based on the track reconstruction studies of Hawkings and Sinev, is only about 1/10th of this value. However, in these studies, layer 1 is used democratically with the others. Consider instead that the track finding, and linking to the outer tracker, is all done with VXD layers 2-5, for which the efficiency should still be excellent. Layer 1 is now used purely to refine the extrapolation of the fitted track to the IP. The 2-sigma extrapolation ellipse from layer 2 to layer 1, will generally home in unambiguously on the real hit. For very low momentum tracks (below 200 MeV/c) a few percent of measurements will be spoiled by a background hit in the search ellipse, but in such cases, even if there were no background, the error in the extrapolation to the IP, due to the material in layer 1 and the beampipe plus liner, would make such tracks relatively less useful for vertex finding etc. Before deciding whether or not the '20-slice technologies' are acceptable or not for layer 1, more detailed simulations will be needed, in which one can evaluate the degradation in critical physics measurements (such as processes needing knowledge of the vertex charge) as the background is raised to the most pessimistic levels. Such studies should shed light on this important new question of mixed technologies or not for the VXD barrel section. Further information on this issue can be found in Section 5.8.

While being open to consider all technologies under development, this group also presented a new and interesting idea for a short-column CCD (SCCCD), which could provide single-bunch timing capability while using a CCD technology. The idea is firstly to orient the columns across the width of the ladders rather than along the length, and secondly to use small enough pixels (~10 μm) that a cluster is virtually always split between at least two columns. Finally, by reading the columns in alternating directions (to left and to right across the ladder width) one ensures that a good match between the cluster fragments will be obtained only for hits generated by the appropriate bunch crossing. Out-of-time background hits can be rejected since their cluster fragments will be displaced from one another, when seen from the event-related bunch crossing. This idea has been briefly discussed with CCD manufacturers, who confirmed that there is nothing unfeasible in this concept.

This collaboration has taken up the challenge of investigating beam-related electromagnetic interference (EMI) using VXD readout electronics from SLD, operating in proximity to a SLAC test beam, with various leakage paths allowed for escape of RF radiation. They established that beam-induced EMI can disrupt this electronics through small holes in the electrical screening, but that complete screening by metal shielding reduces EMI effects below threshold for that particular system. RF leakage is found to be proportional to the bunch charge, but independent of the bunch length in the frequency range considered (below about 20 GHz). For much higher frequencies, effects would be expected to increase as the bunch length is reduced below 1 mm. RF connectors penetrating the beampipe provide significant

leakage paths, whether or not the connectors are fitted with terminating resistors. Complete metal screening of these assemblies again provides effective attenuation. This group has more work to do (for example, investigating leakage levels through coax cable) but already important conclusions can be drawn. Given that the electrical integrity of the ILC beampipe near the IP may be compromised by feedthroughs carrying BPM signals etc, and given the possibility that RF escaping along coax cables could be radiated into the environment at the remote end of those cables, it will be important to search for possible sources of RF leakage during ILC commissioning, before the detector is in place. Wide angle antennas with high bandwidth, able to precisely determine the source time/distance relationship, will rapidly pin down any beam-related sources. Identifying these and taking corrective action before the first detector is installed, will be much easier than trying to dig them out later. In the case of SLD, such investigations proved fruitless, not surprisingly since electromagnetic radiation in a crowded experimental environment, with numerous metal surfaces, can suffer multiple reflections and appear to come from many directions. Given the susceptibility of pixel detectors (having thin sensitive silicon layers hence producing very small signals) to all sorts of spurious noise sources, including from other detector systems active at the same time, careful monitoring of all potential noise sources will be important. If in the end a vertex detector that has proven to be robust as regards EMI when operating in its own detector system, is offered up to ILC after the interaction region has been cleansed of any unacceptably large beam-related or other noise sources, one will be confident of the proper functioning of the VXD electronics during ILC running. As part of the ILC vertex detector R&D programme, it will be important to establish the susceptibility of the different technology options to RF interference, by testing fully instrumented ladders under common conditions, preferably in the same test facility.

5.7.   Fermilab Group (Ron Lipton et al)

This very strong collaboration presented their work on mechanical design, sensors and electronics, and simulation software. They are relatively speaking newcomers to the ILC R&D community, and are providing a stream of new ideas and energetic pursuit of critical issues, partly building on their leading work over many years in detectors for hadron physics experiments.

One of their very useful initiatives has been to suggest world-wide coordination of the mechanical R&D work through joint meetings. Participants currently include those working in this area in the LCFI, DEPFET and FPCCD collaborations. Participation is informal, and is likely to increase as more people join this challenging work area, since most of the problems are independent of the particular sensor technology.

Regarding their own mechanical design work, they have provided a real engineering outline for the small-radius region of the SiD detector concept, starting with a realistic mechanical support system for the delicate beampipe in the form of an exoskeleton made of CF laminate. The innermost section of beampipe made of thin-walled (0.4 mm) beryllium of length 12.8 cm is particularly delicate. It needs carefully planned procedures for assembly, installation and removal, which they are beginning to provide. With their 5 Tesla solenoid field, they are able to push to a 12 mm beampipe radius, providing a reasonable stay-clear wrt the envelope of pair background for the 500 GeV machine. This places their inner layer sensors on a radius of 14 mm, which has significant advantages for physics.

They are investigating three approaches to ladder fabrication for the barrel region, sensors mounted on CF laminate supports, all-silicon structures in which the sensors are themselves the supports, and silicon-foam-silicon structures, as

preferred by the FPCCD collaboration. They point out that fabrication techniques and tooling are the key to establishing what is achievable for each, particularly regarding mechanical stability vs thickness. They are using their considerable experience to advance these studies with prototype structures including complete half-barrels, with already encouraging results.

In order to minimise the number of prototypes to be laboriously assembled in laboratories, they are able to explore and narrow down options with the help of FEA calculations. These are demonstrating that for operating temperatures around -10 C or below, which may be preferred due to considerations of dark current, thermal effects dominate over gravitational distortions. Lorentz forces related to cabling in the solenoid field are likely to be even more serious, particualrly since the 'pulsed power' approach may deliver small 'hammer blows' to the support structure. Such effects will depend critically on the choice of sensor technology, and these systems aspects are only now beginning to emerge, including the question of the routing of the VXD power cables out of the detector. While this is a serious challenge, given the universal plan to read out detector data at leisure between bunch trains, the data cables will be reduced to one or two optical fibres each end, so they should be a minor issue. Bill Cooper's group is doing a great job of pushing the sensor technology groups to address the critical cabling questions.

For the forward disks, a CF-foam-CF base is considered, to which sensors of half-reticle size are attached to form mosaics. Some instability in the z direction is likely to be encountered, but this is less serious, given that the tracks concerned will be oriented approximately normal to these disks.

Their plans for vertex detector access and servicing benefit from the push-pull operation of the ILC detectors. Major maintenance and other work would be done off beamline, in a garage location that allows for 3 m door opening. They will there be able to roll the main tracker forward by 3 m along the beam direction, exposing the vertex detector assembly inside its exoskeleton. This assembly can then be moved to a clean room, where major servicing or replacement of the vertex detector can be carried out.

Their work on sensor technology is guided by the perception that the conventional CMOS MAPS approach may run into major difficulties in two main areas. Firstly, the need for small (~20 μm) pixels restricts the amount of logic that can be accommodated in each pixel, and secondly the need for efficient charge collection over the pixel area inhibits the use of PMOS transistors within the pixel. The required functionality can be provided by edge logic, but this collaboration is pursuing the interesting alternative of the SOI or vertically integrated (3-D) approach, in which the sensor layer has additional tiers (each 5-10 μm thick) of analog and digital CMOS stacked above it, with special through-wafer via formation and metallization providing interconnections between them, for example 3 levels (tiers) of transistors, 11 levels of metal, in a total vertical height of only 22 μm.

The ILC VXD community had for years been following with great interest and enthusiasm the SOI approach, through the work of the Polish collaboration, Lukasz Maczewski, Halina Niemic and others. They were unable to attend this review, since their funding has run out (we hope temporarily). Furthermore, their SOI resources were restricted to in-house R&D facilities. Now that SOI and 3-D technologies are becoming commercially available and affordable, it is most welcome and timely that the Fermilab group is able to take a lead in these approaches for ILC VXD development, within their broad range of detector development. In this review, they presented their work on a variety of topics, some of which are guided by other applications, but all are relevant to the VXD challenge.

At least for the inner layer, and to some extent overall, a major challenge for conventional pixel sensors (CCD, MAPS) is the readout speed required to achieve acceptable limits to background hit density. With a background hit rate of ~200/mm^2/train, and a requirement of possibly 1 hit/mm^2 on layer 1 for efficient track reconstruction, they would like to offer a SW (sensitive window) of better than 15 bunches, possibly going all the way to single bunch timing (300 ns). The required SW is discussed in Sections 5.6 and 5.8, and it may be that this ultimate performance will be needed for physics. If so, they argue convincingly that the best approach may be to adopt the emerging SOI and vertical integration (3-D) technologies.

Most of the work they presented relates to chips and test circuits motivated by non-VXD applications, but which are proving extremely useful in learning the 'rules of the game' for these novel technologies. For example, in developing their MAMBO chip (a photon-counting x-ray sensor, also applicable to electron microscopy), they have found that the obvious assumption of one sense diode per pixel, with contact from the high resistivity 'handle wafer' to the SOI transistor level, raises serious difficulties. Unlike the 'conventional' SOI technology, in which the handle wafer is inert undepleted silicon, the depleted sensor layer introduces a back gate as seen by the SOI transistors, shifting thresholds in a way that depends on their postion relative to the sense diodes, as well as on the reverse bias voltage applied to the sensor layer. These problems have shown up on chips produced by the OKI company through a multiproject wafer run hosted by KEK. The collaboration is making changes to the design (including 4 sense diodes per pixel) so as to minimise the problem. Another approach, offered by American Semiconductor Industries (ASI) is a dual gate transistor (Flexfet) that employs a bottom gate to shield the transitsor channel from the voltage below the buried oxide (BOX) layer. However, this may drastically increase the gate capacitance of all those transistors. ASI and MIT-LL also offer a pinning implant to stabilise the voltage at the surface of the sensor wafer, and so control the back-gate effect as well as reducing crosstalk between digital transistors and the analogue input. However, this also has disadvantages, such as increasing the effective sense node capacitance, as well as trapping signal charge.

The SOI technology clearly has great potential when compared with the standard MAPS approach, but even more may be needed to achieve the full functionality of a VXD chip with 20 μm pixels, and this is being explored with the 3-D technology available from Lincoln Labs (MIT-LL). They are starting with a 64x64 pixel prototype, with the intention of eventually producing reticle-size 1kx1k devices to be assembled into ladders. The architecture now employed is an analogue front-end with storage of two samples, one before the bunch train, and the second corrseponding to the hit sensed by a discriminator. With ~100 ns shaping time, the signal can provide single-bunch timing, with in-pixel time-stamp logic, for which there are several options. Since the signal is the difference between the two samples, CDS in the sense of reset noise suppression is achieved. But in view of the delay of up to 1 ms between the samples, there is the danger of low frequency pickup shifting the baseline enough to disrupt the true threshold. In principle, more sophisticated noise filtering (closely coupled sampling) could of course be considered, but having to implment anything more at the pixel level could raise the power excessively. At present, the main contributor to the power is the front-end preamp, for which a node capacitance as low as 10 fF is considered a realistic goal (provided a clean solution is found to the back gate problems). In these circumstances, they calculate that good noise performance will be achievable with 0.5 μA drain current, delivering a risetime of ~100 ns.

On the question of power distribution, this collaboration is leading the way in the ILC VXD community, by investigating options such as serial powering that are active areas of R&D for sLHC. Eventually, once this has been

worked through for all sensor options, these considerations, leading to realistic material budget estimates for the necessary services (electrical components, cables and cooling) will provide important input for narrowing down the acceptable technologies for ILC.

Given that there is not much likelihood of such complex devices (SOI or 3-D) being amenable to stitching to the scale of full- or even half-ladders, this collaboration is focusing on reticle-scale devices daisy-chained together by wire bonds. To avoid unacceptable gaps in z, they are exploring the production of edgeless sensor asemblies. Whether edgeless or not, the handle wafer needs to be thinned to ~50 μm in order to meet the goal on material budget. Fortunately for our community, thinning technologies are now widely available commercially. However, this collaboration has found it difficult to engage one of the companies with whom they are working to perform the laser annealing needed after implantation of the new back surface, so they are developing this capability themselves.

Overall, this is a very ambitious annd impressive programme, which fortunately is supported by several possible applications, by no means confined to ILC vertex detectors. The whole ILC detector community is watching these developments with great interest and enthusiasm.

The 'SLAC+collaborators' part of the ILC community has made important contributions to simulation and reconstruction software over many years. In this review, Rob Kutschke informed us of the developments he is leading at Fermilab to make further progress. It was very encouraging to learn of this increased manpower, since until now, some of the issues relating to detector performance, such as the tolerable rate of background hits on layer 1, may to some extent have reflected limitations in algorithms as opposed to fundamental limitations of the detector. Increased manpower studying these issues, that we heard will be coming from the 'Fermilab+collaboators' team, is most welcome. The combined effort of all these developers, bringing experience from different HEP backgrounds to bear on these problems, are sure to bear fruit. The entire silicon tracking system of SiD, for which there is still considerable flexibility as to where the pixels end and the strips begin, provides an opportunity for simulation and reconstruction software development, which is essential in order to truly optimise the performance, particularly as regards efficient tracking over the full solid angle and minimising the material budget in front of the calorimetry.

5.8.   FPCCD Collaboration (Yasuhiro Sugimoto et al)

This collaboration is following a unique and interesting approach to the problems of ILC vertexing. Most optons aim to overcome the high background density in layer 1 by multiple readout during the train, or by time stamping the hits, in other words by reducing the sensitive window by a factor 20 or more. The FPCCD approach is to increase the pixel density by a similar factor (CCDs with 5 micron square pixels are suggested) and hence resolve the signal and background hits with high efficiency. In order to keep the clusters as compact as possible, they will operate with fully depleted sensor material (20 μm thick epi), following common practice for back-illuminated CCDs.

Their approach has several advantages. It builds on the track record for CCDs with their excellent noise filtering capability, which have proven to be viable in collider environments, notably SLD. By reading out in the long inter-train period, the readout involves only modest clock speeds, where excellent noise performance is guaranteed, and power dissipation is minimal. Pulsed power is neither required nor desired - so they avoid all the challenges of huge peak

currents requiring thick conductors, blasts of heat during each train, and significant impulsive Lorentz forces associated with the high currents.

In order to be quite safe as regards radiation damage effects, they propose to divide the sensitive area into 16 sections, each read out by a long serial register, so the length of the parallel register (where most of the CTE loss due to radiation damage would take place) is only 128 pixels. The 16 readout registers also act as storage elements for signal charge during the train, so there is no loss of efficiency from interleaving them within the imaging area.

Will this factor 20 in granularity be sufficient? This must be answered by simulations, and here they have some interesting new results. Most approaches to VXD pattern recognition have used all layers democratically, and find that with excessive hit density layer 1 becomes useless, initiating an unacceptable rate of fake tracks. This group follows a different approach, performing the track finding from the outside, and using layer 1 only to refine the impact parameter precision for the track extrapolated into the region of the IP. Results are encouraging; with the nominal background rate, the algorithm assigns the correct hit approximately 96% of the time. Furthermore, they have more to gain, since many background hits can be suppressed on the basis of the cluster shape, given that their tracks are generally spiralling with oblique angles to the sensors. This approach works best near the ends of the ladders, with a 'suppression factor' increasing from ~2 to 20 as one moves from z=0 outwards. Furthermore, the sensors are paired on each ladder (one inward facing and one outward), so consistency checks between these can further suppress background.

As well as the sensor development, they are engaged in the design of multi-channel ASICs incorporating all the off-sensor electronics, namely preamp, filter, CDS, ADC and (eventually) data sparsification.

They developed this concept soon after the ITRP chose the cold technology for the machine in 2004, but only in 2006 were they awarded funds to start the R&D work with HPK. They envisage 3 cycles of design/manufacture/test with small-scale devices, followed by another 3 cycles with full-scale devices, culminating in complete ladders by about 2012. In order to meet this schedule, a considerable increase in resources needs to be allocated to this collaboration. Given their limited resources, they are understandably enthusiastic for pooling of engineering effort in the ILC VXD community to solve common problems such as wafer thinning, ladder assembly procedures, mechanical supports and cooling.

5.9. LCFI Collaboration (Tim Greenshaw et al)

This strong collaboration has been playing a major part in ILC vertex detector developments over the past 10 years, integrating to a considerable investment in money and manpower. Their two sensor options (CPCCDs and ISIS devices) were conceived even before the cold option for the machine was selected. The work they presented was divided into four main work package areas, physics studies, sensor developments, sensor-related electronics, and mechanical R&D.

For the physics studies, their early work was done using the fast MC simulation code SGV. Fast MC effort has focused on studies of 'vertex charge' reconstruction, permitting discrimination between leading quarks and anti-quarks in heavy flavour jets. The algorithm to reconstruct vertex charge builds on pioneering work in the SLD collaboration, and has demonstrated that the superior detector performance achievable at ILC will permit an unprecedented quality of vertex charge measurement. Measurement of *B* vertex charge is more challenging than for charm, and their work shows

that this physics tool will push the limits of detector design and performance, particularly as regards the beampipe radius and material budget. Indeed, it is mainly these studies which motivated the goal for layer thickness of at most 0.1% X_0. Given the delicacy of the measurement of vertex charge, where a single low momentum track that suffers excessive multiple scattering in the beampipe/layer 1 can completely invalidate the measurement, it is not sufficient to make studies with the fast MC programs that have been used to date. Since full Geant 4 simulations are obligatory, the collaboration has developed and is maintaining the LCFIVertex pakage of code (~20k lines) for vertex finding, heavy flavour ID and vertex charge reconstruction. This package currently provides the algorithms that originated at SLD and were developed further in later ILC studies, newly implemented in C++ in a user friendly way, permitting running conditions to be easily changed and including exhaustive documentation. This code is fully interfaced to the European software framework MarlinReco. With I/O being based on the internationally agreed common ILC data format LCIO, the code can also be run in conjunction with packages from other ILC software frameworks. It is being used for optimisation of detector concepts in all three regions. The small group responsible for delivering these tools to the community is working effectively with international partners to be sure that the code is bug-free, efficient in execution and fully compatible with the overall simulation/reconstruction code, as regards such common features as steering parameters for the tracking code. They have also taken on responsibilities to maintain and develop the code in new directions, not previously explored due to lack of time. These include features such as the correct handling of K-short and lambda decays, as well as photon conversions and secondary interactions in the detector material, since all these processes can confuse the flavour ID procedures if not handled carefully. Furthermore, they intend to explore ways to improve the current flavour tag and vertex charge determination, eg by using additional information and/or applying novel data mining techniques. As these new tools are developed in conjunction with Geant 4, they will provide the appropriate basis for making critical decisions regarding the detector design, such as the physics impact of pair background particularly in layer 1, the relative advantages of long barrels vs short barrels plus forward disks, and other important detector parameters. Continuing tool development is vital not only for optimising the design of the vertex detector, but also for optimising the design of the overall detector (notably the tracking system) and for benchmarking with physics studies. Much work is still needed, and the collaboration has a long-term responsibility to the world-wide community to maintain and develop this code in the light of evolving ideas for detector design and physics analysis.

For the sensor development, the CPCCD approach is forging ahead towards its goal of 50 MHz column parallel operation, set by the occupancy requirements in layer 1. This work is opening new doors for high speed imaging. If these goals are achieved, this will offer a secure route to a detector using full length sensors on ladders. Even if the full clock rate is unachievable, CPCCDs could offer a robust solution for layers 2 and beyond, in conjunction with an alternative (faster) solution for layer 1.

The LCFI collaboration has established bump-bonding between CPCCDs and readout chips, with electrical connections at a pitch of 20 μm, opening the door for numerous other applications needing a factor 100-1000 faster readout than is achievable with standard CCDs. They have developed readout chips with front-end amplifiers feeding 5-bit ADCs, pixel-rate CDS and sparsification logic, all on the same 20 μm pitch as the CCDs. While they have encountered minor deficiencies in these chips, which the next generation is likely to correct, their current status represents proof-of principle that sophisticated bump-bonded CCD/CMOS hybrids are within reach, a message that is

relevant to other application areas such as high speed imaging systems for soft x-rays, for which a CPCCD made on high resistivity epitaxial material should be efficient and (by virtue of back illumination) highly radiation resistant.

In some respects, delivering the precisely controlled high current drive pulses to the CPCCD is as challenging as dealing with the outpouring of data. The LCFI collaboration has made excellent progress in developing extremely compact high current driver chips, which will be mounted at the ladder ends, also bump-bonded to achieve the very low inductance connections needed for high speed clocking. Both in terms of the clocking and the readout, this R&D programme is opening new windows of opportunity into the world of CCD imagers, for which the device manufacturer (the e2V Company in the UK) has identified markets beyond particle physics.

The other approach that LCFI is pursuing for sensors, namely In-Situ Image Sensors (ISIS) devices, is newer, more advanced and consequently less developed, but the potential for science in general is at least as great. Since it requires both CCD and CMOS technology to be integrated in a single chip, there are few potential manufacturers. However, the collaboration has identified Jazz Semiconductors as having all the required capabilities, and is starting to work with them. In addition, e2V have made a small proof-of-principle device with large pixels, using a CCD manufacturing process, which performs well in lab tests. This approach relies on developing a deep p-implant to shield the storage register and active circuitry from the charge-collection layer below. If this can be made to work, it will provide the 'holy grail' for a number of astronomical applications such as star trackers for which high speed devices with 100% fill factor are needed.

The other major work area for LCFI comprises the mechanical studies, tackling the challenge of 0.1% X_0 ladders, respecting the requirement for a precision vertex detector of few micron stability. The collaboration started with a daring unsupported (stretched) silicon approach. While this did not provide sufficient stability, it taught the community many useful lessons, and provided the springboard for other approaches. Mounting to carbon fibre composite (CFC) support assemblies (ladder substrates or support barrels), is an interesting approach, but for detectors operating at reduced temperature (-10 C or below) one needs to be very careful of differential contraction producing compressive forces which can lead to shape instabilities (buckling) in the silicon sensors. A new approach is based on low density materials, such as silicon carbide foam, having a cte very closely matched to silicon. This material is machinable without great difficulty, and could be used not only for the ladder substrates, but also as the structural material of the entire vertex detector assembly, and assisting with the task of stabilising the very thin and delicate innermost section of beampipe. Manufacturers are working to reduce the foam density to well below 10%. Here again, there are synergies with other areas of science using silicon sensors where the lowest possible material budget is required.

Assuming financial support continuing at current levels, this collaboration is well-placed to complete their R&D programme by the target date of 2012.

5.10. Chronopixel Collaboration (Charlie Baltay et al)

This important collaboration follows the CMOS MAPS approach, but instead of the fast rolling shutter of the MIMOSA collaboration, aims to achieve in-pixel sparsification and single bunch time stamping. Their 'special features' which make this possible are the use of a deep p+ implant to exclude the PMOS transistors from stealing signal charge, and the planned evolution to 45 nm technology and 10 μm pixels. With such small pixels, binary readout should

provide adequate precision, and with the small feature size they can fit in the necessary logic without resorting to the 3-D approach. Indeed, they propose to fabricate full-size sensors of 125 x 22 mm^2 by stitching between reticles for the device processing.

Signal charge is collected from a 15 μm thick epi region which is fully depleted. There are concerns (only short-term) about obtaining high resistivity material that can be fully depleted by the limited bias voltage associated with the 45 nm process. There are prospects of a thicker oxide layer, hence increased voltage limit, for the analogue parts of the circuit in each pixel. They present new evidence (from simulations) for collection of signal charge being helped by a weak electric field in the undepleted epi material, which at first puzzled our committee. However, we have been greatly helped by subsequent e-mail exchanges with Nick Sinev, who referred us to his recent studies with the Synopsys TCAD package. He appears to have uncovered an effect of potential importance for many LC vertex detector options, namely a weak E-field within the undepleted epi material which aids collection of the signal charge. As we understand it, this field is created by the flow of dark current through the high resistivity material, and hence is much reduced at low temperatures, so the effect would not have been observed previously in typical CCD-based scientific imaging systems. Since their simulations indicate more compact and efficient signal charge collection as a result of this field, the effect may be of great interest for other VXD options. If confirmed, it will encourage operation at somewhat higher temperatures, and with the highest possible resistivity for the epitaxial layer.

The deep p-implant is essential to prevent spurious charge collection by the n-wells associated with the PMOS transistors. They want to collect at least 50% of signal charge within one pixel, and given the benefits of the newly-discovered E-field, this appears in simulation to be possible.

The signal processing through the bunch train is not very clearly described. Evidently, they use a high bandwidth preamp (risetime <100 ns) and sample the signal after each bunch crossing. So they need to distribute a clock to all pixels. They assured us that the peak current and power associated with this has been included in their overall estimates, and is not the dominant contributor. After some e-mail exchanges, the committee remains puzzled by their plans for noise suppression (bunch-to-bunch CDS or something weaker)? Our impression is that their strategy for noise suppression is in a rather early stage of development. Their in-pixel processing is able to store times associated with up to 4 hits per train in each pixel.

They mentioned 2009 as the timescale for 45 nm processing, but have explained that this refers to the technology roadmap. They do not expect to be using this technology themselves, on that timescale. During 2008, they will produce first prototypes using the TSMC 180 nm process. However, this route is limited in terms of epi layer resistivity (10 Ωcm whereas 10 kΩcm is desired) and furthermore the deep p-well is not available. For these reasons, these prototypes will be more useful to test the electronics than the sensor functionality. It is also not clear from the writeup whether the CDS functionality for bunch train operation will be included in this prototype, probably not. It will clearly be most interesting to move as quickly as possible to realistic prototypes. The review committee wondered whether the plan for two prototype phases with TSMC would be justified, or whether they should proceed next to a 45 nm process with full capability for all their requirements (deep p+ implant in high resistivity epi). Given the cost of a 45 nm mask set, do they (via Sarnoff) have access to a multi-project arrangement with some foundry to make this affordable?

In general, the committee was a little concerned by the approach of leaving so much of the sensor design in the hands of a company which is not participating in the science goals of the project. For example, what would happen if the

company decided to stop work on the project - how much of the costly knowledge base would be useable by the collaboration? Having an intimate involvement of one or two collaboration members in the detailed pixel design would appear to us to be desirable.

6. CONCLUSIONS AND RECOMMENDATIONS

6.1. Environmental Issues

The most important external constraint which influences the potential VXD performance is of course the beampipe radius, which follows directly from the machine backgrounds. At SLD, the limit came from an unexpectedly high level of background originating from SR produced by unpredicted non-Gaussian tails on the phase space population of the beams entering the IR. There is almost 100% confidence from the machine people that such backgrounds will be negligible at ILC, but we could again have some nasty surprises. If all is well, the limit to the beampipe radius will this time come from the hard edge to the envelope of pair electrons generated as result of the beamstrahlung process of the colliding beams at the IR. This has been calculated with a high degree of confidence, and it seems reasonable to assume that this does represent the 'worst case', without adding safety factors which would degrade the physics performance. If there are again nasty surprises, one may need to modify the detector after the first push-pull, removing the inner layer or whatever. This could most efficiently be done on the second detector while the first was still running. Meanwhile, there is widespread agreement on the shape (radius and length of inner cylindrical section) of beampipe that can be accommodated with reasonable stay-clear to account for misalignments between the IP and the beampipe centre. However, this radius is obviously a function of the assumed solenoid field, which is likely to be in the range 3-5 Tesla. The physics benefit of the higher B field as regards VXD-related physics has not yet been evaluated; when this becomes possible, it will be one of several factors which will go into the eventual choice of solenoid field for the real detectors.

In addition, there is the uncertainty as to which final focus option the experiments will wish to use. One could well imagine that at startup, being far below the design luminosity, they will opt for the 'low P' option, with much greater beam-beam disruption. But if the beampipe is sized for that, there will be a significant loss of physics in the future, when the nominal luminosity is achieved and good vertex detector performance is required. Tradeoffs should be studied carefully before taking firm decisions.

Almost as important as the beampipe radius is its wall thickness, which is generally assumed to be 0.4 mm. This parameter isn't quite as important, since the impact parameter resolution scales linearly with the radius, but only as the square root of the thickness. 0.4 mm is certainly overkill in terms of supporting the vacuum, but the braze joint between the cylindrical and conical sections will be delicate. Schemes have been devised to minimise stresses (mechanical help from the VXD support shell - ideas from the TESLA TDR era, and a large exoskeleton - ideas from the preliminary SiD engineering design). These studies should reach a high level of confidence during the forthcoming TDR phase.

Despite the confidence about SR-related background (typically fluorescence x-rays from surfaces visible to the detector, which may be one or more 'bounces' away from surfaces that see direct SR), it seems prudent to assume a high-Z liner inside the beryllium beampipe; 25 μm of titanium foil is typically assumed. This imposes a significant multiple scattering penalty, so it is important to study this in detail, to determine what is really needed. Here again, if one gets it wrong, it could easily be corrected at the first push-pull, by modifying the other beampipe assembly.

The early commissioning of the machine, with diagnostic equipment in the IR in place of either detector, will provide a golden opportunity to track down and fix any problems to do with leakage of beam-related RF, due either to imperfections in shielding along the beamline, or to RF leaking from cables coming from BPMs, vacuum pumps and gauges etc, which are exposed to RF associated with the beams and their image currents. From the excellent studies at SLAC ESA, we can be confident that appropriate wide angle antennas with fast edge timing, will be able to pin down any sources of RF leaking into the experiment pit. Similarly, the commissioning of the detector system at ground level will provide the opportunity to search for other RF sources that may disrupt the delicate VXD system. Despite the wonders of CDS and ERF (Section 6.4) it is possible for external noise sources to cause havoc with a 1 Gpixel system, totally overwhelming the sparsification capabilities of any DAQ system.

As well as machine backgrounds causing extra hits in the VXD system, one has also to be concerned about radiation damage effects. All the R&D groups have followed designs which they expect to be robust wrt the effects (ionisation and bulk damage) in their sensors and electronics, related to the $e+e-$ pair background, though this has so far only been demonstrated by the MIMOSA collaboraton. However, there may be problems with neutrons from the beam dump. Since, for reasons of saving money, the beamstrahlung dump and main beam dump have been combined into one, this dump has by design a line-of-sight view of the vertex detector layer 1. Neutrons originating from that dump could produce dangerously high fluxes in this detector. Progress was made last year with ideas to mitigate these effects, and maybe this is a solved problem, but it is certainly important for the VXD community to keep in touch with the MDI group about this.

Another machine-related issue of great importance for the VXD performance is the time stability of the beam spots. Clearly their intrinsic size is negligible on the scale of the required impact parameter precision, and in order to make luminosity the beams have to track each other to this level of precision. But in principle, given the capabilities of the FONT system, they could wander slowly around in unison at the level of many microns, significantly degrading the VXD performance. As at SLD, the way to define the IP is to average stiff tracks over some tens of events, which will again work beautifully if the beam positions are stable at the micron level over a period of some minutes. Recent discussions with MDI experts suggest that the optical anchor will be able to provide the required absolute stability, but this implies some significant work and will only happen as result of a formal request from the vertexing community, most logically after in-depth discussions by the suggested Vertexing Coordination Group.

6.2.   Long vs short barrels

It is extremely fortunate that two of the LOI collaborations are following different approaches on the question of long barrels (with coverage to $\cos(\theta) \sim 0.96$) and short barrels plus disks, with somewhat greater coverage. If the material at the ends of the barrel ladders can be reduced sufficiently, the short-barrel option will be preferred. Each of the nine sensor options has its own collection of problems and requirements regarding off-sensor local mechanical and electrical services, and it will not be until groups have complete serviced ladders operating in a test beam that one will be able to take a decision on this. It is significant that with the tracking systems for the LHC general purpose detectors, the material budget has continued to grow long after the end of the official R&D phase, reflecting the difficulty of getting a complete picture even after extensive test beam studies of smaller scale systems. It is clear for all ILC detector

concepts that forward tracking disks will be essential, and it is pretty obvious that these should be pixel-based, and that this technology may usefully be extended to the rest of the tracker. However, 50 μm pixels with binary readout should suffice, which will be much simpler and consume much less power than few-micron-precision pixels, as would be needed for vertexing. The open question is whether or not the VXD barrel services will degrade the performance so that higher precision disks are insufficiently useful for physics. It could even happen that long barrels are preferred for startup, with a superior VXD technology coming along later and tipping the balance the other way.

6.3.   Possible technology differences between layers

As discussed in Sections 5.6 and 5.8, there is uncertainty at present as to the target hit density that one needs to reach on layer 1 in order to do the physics. If 1 hit/mm^2 is really important, one can forget about all except the 3-D and chronopixels for layer 1, assuming one of them can be made to work. However, it could be that the limit is 5 times higher, and that other considerations (such as mechanical stability or being ready on time) favour a different option. Again, we can be thankful to have talented groups exploring all these options. It is an excellent outcome of this review that Su Dong opened our minds to the possibility of a mixed technology VXD system, since this may provide an optimal strategy, given that conditions for layer 1 are in many respects different - there are both challenges and opportunities, as discussed in Section 5.6. If a different technology is needed for layer 1, there might be a case for extending this to layer 2, depending on the outcome of the detailed pattern recognition studies. However, layer 2 does not share the special advantages of being adjacent to the beampipe, so it is largely subject to the power dissipation restrictions of the outer layers. As well as different technologies between layers, there is the less dramatic question of possibly increasing pixel size with layer number, suggested by the MIMOSA collaboration. All these open questions reflect the immaturity of the simulations related to the vertex detector, despite many years of hard work. However, the essential software tools are now in place, and these questions should all receive clear answers over the next few years.

6.4.   Correlated Double Sampling and Extended Row Filter

The heart of a collider experiment is intrinsically a somewhat noisy environment electrically. Power supplies, vacuum pumps, environmental monitors, and readout systems generate noise which can be significant for monolithic pixel systems where the signal thresholds need to be set in the region of hundreds of electrons. Problems are exacerbated by the need to minimise electrical screening throughout the tracking volume, in order to keep material budgets under control. The concept of each subdetector residing in a hermetic Faraday cage is unrealistic. The most powerful handle in terms of noise suppression is the fact that the deposited signal charge is 'permanently' available (till the sense node is reset) whereas capacitively coupled electrical noise is intrinsically transitory. This fundamental difference was the basis for the noise suppression system developed for SLD, where it proved highly effective. The procedure of correlated double sampling (CDS) combined with extended row filter (ERF) can in principle be used to varying degrees by all pixel systems envisaged for ILC. Most of the VXD R&D collaborations have said that they will use CDS for noise reduction, but it appears to mean different things to different people in this community, so we aim here to provide some clarification and suggestions from the previous experience.

CDS was initiated as a means of eliminating reset noise in CCD readout systems, primarily for astronomy where optimal noise performance has been developed over several decades. For their applications, the readout sequence for each pixel typically consists of:

node reset, followed by shaping/sample A, followed by signal charge transfer, followed by shaping/sample B.

The analogue or digital difference *B-A* provides the optimal measure of the signal free, of reset noise. The time difference between samples should not exceed a few times the signal shaping time. If the samples are less correlated in time, the system becomes susceptible to an ever increasing noise bandwidth, up to the cutoff set by the shaping time. For this reason 'frame-rate CDS' during the bunch train, which is the only option available to some of the technologies, while it does indeed suppress the reset noise and other fixed-pattern noise, has limited capability for pickup suppression, as can be inferred from the difference between the shaping time (necessarily at most ~100 ns for readout during the train where the bunch crossing period is ~300 ns) and the frame readout time (for example 25 μs, 250 times longer).

For vertex detectors with very sparse data, it was realised that one could speed up the readout compared with the astronomy sequence by applying the reset only once per row, so the sequence became:

sample A, charge transfer, sample B, charge transfer, sample C, charge transfer, etc.

While this is effective at eliminating the reset noise, which would otherwise dominate, it is still susceptible to external pickup spikes which can disrupt either of the samples *A* or *B*. It was found at SLD that the noise rate at the required pixel threshold (around 200 e⁻) was dominated by such pickup spikes. A clean solution was to implement a simple algorithm by which the signal in pixel *C* for example would not be determined simply by *C-B*, but by the lesser of *C-B* and *D-A*. The probability of two samples being hit by pickup was negligible, and the noise rate was accordingly reduced by a factor ~100. This procedure, called extended row filtering (ERF) will probably be needed at ILC, to avoid a data deluge, given the unprecedented scale of the system, ~10^9 pixels. This refinement is available to all options, at some expense in readout rate and power dissipation for the additional logic. While different pixel systems have different output circuit structures, the principle of sampling twice before the signal charge is collected by the node, and twice afterwards, at an interval determined by the signal shaping time, is always available. (For the DEPFET, the time sequence is reversed, with the first samples being taken in presence of the signal charge, followed by further samples after the clear.)

We comment in Section 6.8 on our understanding of the applicability of CDS to the ILC candidate technologies, and some possible implications for each.

It could of course be argued that such noise suppression may not be needed, if adequate electrical screening of the detector is implemented. However, as Marvin Johnson has pointed out, most Faraday cages in HEP detectors are usually 'little better than dust covers'. Much higher quality screening would of course be possible, but at the expense of considerably increased material budget. Since the pressure at ILC is to greatly reduce the material in the tracking systems compared with anything that has gone before, it seems imprudent to assume that less robust noise filtering will be acceptable.

6.5.   Mechanics and Alignment

The discussion about whether to make ladders from large sensors, or as mosaics of reticle-scale devices, will clearly continue for years. To a great extent, it may be decided for each option by practical considerations. For CCDs which can be made in large area with high yield, there is every reason to follow the plan for full-ladder or half-length sensors. For 3-D devices, reticle size is presumably going to be the upper limit. For monolithic CMOS options, it isn't so clear. One argument for limiting to reticle scale, made in the review, was based on yield considerations, which might currently be 50% for some process. However, Dave Christian made the important point that even in this case, large devices might be feasible. The point is that most of the defective devices tend to come from the edge of the wafer, so that the yield for a single large device, placed in the middle of the wafer, might be much higher than one would infer from the average yield of smaller devices. This is something not to be forgotten as collaborations move towards full-scale ladder prototypes.

A great deal of effort is going into the development of ladders that meet the design goal of 0.1% X in thickness over their *active length.* Minimising the material at the *ladder ends* is also extremely important as mentioned in Section 6.2, but here work is only just beginning. Once optimised mechanical designs for ladders begin to emerge, the question will be how best to assemble them into concentric barrels. There has of course been considerable experience in this, but modern equipment provides new opportunities. Two main approaches for support structures are complete half-shells of CF or other material, to which the thinned sensors are attached, or separate support strips which lead to ladders that have to be supported and stabilised by some arrangement of end rings. There are advantages and disadvantages to these contrasting approaches, which are partly linked to the chosen sensor technology. One guiding principle is that nothing needs to be rigid; the structure needs only to be mechanically stable when installed and assembled. Thus, two half-cylinders might be torsionally weak, but gain stability when clamped together round the beampipe. What has absolutely to be avoided is instability such as vibration or creep while operating, due for example to stresses from cables that impose excessive external forces.

Individual sensors, thinned to about 50 μm, will be pretty flexible, more so if they are of ladder size than reticle size. Either way, they should gain their mechanical stability by attachment to their substrate. There is plenty of experience doing this with CCD sensors, but it is well known that chips with multiple metal layers (CMOS, even more extreme with 3-D) tend to be significantly bowed after thinning, and require correspondingly greater force to flatten them. Of course, perfect flatness is not obligatory, but what is essential is that the assemblies of these lively chips to their minimal-mass substrates should be stable over time. Deviations from planarity can be measured precisely with modern optical survey equipment, but significant movement after assembly into barrels could not be tolerated. This also bears on the question as to whether sensors should be inward- or outward-facing on their supports. At first sight, one would prefer them to be outward facing, and survey each layer in turn through the buildup, starting with layer 1. But there are arguments in favour of surveying the inner surfaces, working with half-shells starting with layer 5 in its support shell, and proceeding inwards. Again, much depends on the sensor technology, mode of attaching cables and other factors.

What is clear is that, if correctly designed, one can expect to build the VXD up while surveying it layer by layer with modern optical measuring equipment, then to assemble it round the beampipe and to know, with micron precision, the shapes and positions of all the sensors in the system. If disks are included, their survey is somewhat easier than the

barrel of nested cylinder. It bodes well that Bill Cooper and Joel Goldstein are running open meetings of all ILC VXD participants on the subject of mechanical R&D. Steady progress is being made, and the availability of new materials such as silicon carbide foam and of new survey equipment having adequate precision in all 3 coordinates, gives confidence that the problems will be solved with minimal overhead in terms of material for mechanical stabilisation at the ladder ends, despite the intrinsic flexibility of each individual ladder.

While there are many complications wrt SLD, mostly related to the required readout rate, some things have become much simpler. All technologies may profit from operating at reduced temperature (~ -10 C) and some may need to run much colder. In SLD, operating at ~180K caused considerable complications related to differential contraction. For ILC, it seems likely that no technology will need to run that cold. Furthermore, the availability of materials with excellent mechanical properties, and cte closely matched to silicon (such as silicon carbide foam) promise mechanical structures that will be entirely liberated from these problems. This is a major reason for confidence that the VXD will, if designed correctly, be capable of a micron-precision optical survey during assembly, and that its shape will remain stable with that level of precision throughout the life of the detector.

There have been some misunderstandings about the implications of push-pull operation on the mechanics of tracking and vertex detectors. It is clear that the VXD system mounted on the beampipe will, after each push-pull cycle, end up in a different position and orientation wrt the time before. However, by appropriate mechanical design, there is no reason for there to be any significant external forces other than gravity on the VXD or main tracker, so the shape stability of each should be guaranteed. The relative alignment between these two systems can of course be determined by a track-based procedure. The rate of mu-pairs in the ILC energy range is adequate, and they will provide close-to-ideal data samples for alignment purposes. They aren't quite perfect, since the energy spread at ILC will make them deviate significantly from back-to-back. These issues need to be studied in detail, but it appears that by appropriate design (for example, enough ϕ-overlap to link each ladder to its neighbour by tracking) micron-level precision should be achievable, well matched to the intrinsic precision of the detectors. There is good reason to expect that Z-pole running will not be needed in order to maintain full control of the alignment between the VXD and its associated outer tracking system. It will be much better to design and build these subdetctors carefully than to later need to impose time-wasting inefficiency in the operation of this extremely costly accelerator facility.

6.6.   Power and Cooling

This could be the Achilles heel of some of the currently considered technologies. It has long been apparent that for the cold machine, the preferred strategy would be pulsed power, taking advantage of the 200 ms off period to let things cool down. With an average power below 100 W, maybe as little as 30 W for some options, this will be compatible with gas cooling (either dry air or nitrogen, with arguments both ways). The main concern is the peak power during the bunch train, and how to avoid massively thick cables for the delivery. An ideal solution could be local charge storage in super-capacitors at the ends of the ladders, with multiple bump-bond connections to avoid excessive voltage drop as the current of possibly 10-20 is amps is delivered to the sensor. Depending on the technology, there may or may not be super-capacitors adequate for the job, without expanding the size of the ladder ends to ridiculous dimensions.

The Fermilab group is doing pioneering work in the field of serial powering and related approaches to minimising the external cable plant, which will be of vital importance to all the VXD technologies. Since the VXD barrel is the innermost part of the innermost detector, the issues here are particularly important, not only for the VXD endcaps (if any) but for the overall tracking and PFA or compensating calorimetry systems. Photon conversions and hadronic interactions in the VXD material are the most damaging to the overall system, since charged interaction products will in general be bent out of their parent jet. Recovery by software is not excluded, but it will be much better to minimise these interactions. Once we have one or two realistic designs, it will be important to determine the physics consequences of these effects, and to know how hard we should work to make further improvements.

As well as the consequences regarding the material budget, high-current cables threaten to impose large Lorentz forces on the delicate VXD structure. Very careful cable routing and other measures will be needed to keep things under control. Test data of fully serviced ladders in the appropriate magnetic field will be essential.

6.7.   Installation and Access

Since some years ago, everyone has agreed to follow the approach pioneered by SLD for installation and access, which has become easier since the push-pull decision. In the 'garage' position, detector end doors can be opened by as much as 3 m, leaving plenty of room to roll the main tracker along the beam direction, exposing the inner detector system. As is clear from this review, easy access is essential for detector modifications (eg if backgrounds proved worse than anticipated), and to allow one or two detector upgrades during the life of the project. For such major work, the inner detector would obviously be moved to a clean room, and this needs to be taken into account in thinking about avoiding excessive stress to the braze joints of the innermost section of beampipe. Another issue which needs to be agreed once and for all with the machine people is that of bakeout. From SLD experience, it seems that it should suffice to have an initial bakeout in the absence of the VXD system, after which the inner section of beampipe can remain forever at or below room temperature, being held either under vacuum or filled with dry nitrogen, during the long period of VXD buildup, optical survey, etc. Close continuing liaison with the MDI group on such matters will be needed to avoid misunderstandings. People come and go, and clear lessons from the past can easily be forgotten.

This easy access to the small-radius equipment is essential not only for the VXD, but for other delicate equipment such as parts of the forward calorimetry and BPMs.

6.8.   Choosing technologies

Given the uncertainties surrounding large-scale pixel systems in the ILC environment, it would seem prudent to use different technologies for the two large detectors. Since ILC will support two detectors operating in push-pull, and since there may be a case for layer 1 to employ its own special technology, there may be up to four pixel technologies used in vertex detectors at the startup of ILC. In addition, promising technologies that couldn't reach maturity in time might be strong candidates for future upgrades. All designs that might satisfy the ILC conditions are adventurous, and cannot be guaranteed. From the comments below, one can infer possible showstoppers for each of them, bearing in mind that the real showstoppers are likely to be things that none of us has yet thought of. The 'example showstoppers' are listed to emphasise the importance of avoiding the temptation to organise a premature shoot-out to downselect the

'winners'. Our review came to the conclusion that all current options should continue be supported, including if possible the SCCCD (see Section 5.6) which we omit from the list below since resources haven't yet been found to work on it. Narrowing the field down will be appropriate on the basis of results from full-scale fully-serviced ladders operating in a test beam, on the timescale of around 2012, or later for options considered candidates for upgrades.

For further details on all the technologies, we refer readers to the slides of the excellent presentations in the review, which can be found at the workshop website, http://ilc.fnal.gov/conf/alcpg07/

6.8.1. CMOS MAPS - MIMOSA approach

Brief description: Uses 'standard CMOS' sensors, hence limited to NMOS transistors in their pixels. Rolling shutter readout - row-parallel as opposed to column-parallel in order to achieve the required frame rate. Modest penalty in material budget (~2 mm of inactive silicon along the long edge of the ladder).

Mechanical issues: Each ladder probably composed of 10-20 sensors, since stitching may be unacceptable due to yield considerations.

CDS: Frame-rate, ie 25 μs on layer 1, longer on outer layers. This form of CDS, while it suppresses reset noise and other fixed-pattern noise, is not robust against baseline drift and low frequency pickup during the signal integration period.

Example of a possible showstopper: Insufficient pickup immunity due to frame-rate CDS.

6.8.2. Deep n-well

Brief description: Full CMOS in pixel, collecting signal charge on the deep n-well that houses the NMOS transistors for the analogue circuitry (CMOS triple-well process). Allows in-pixel data sparsification and time stamping with ~30 μs precision. Disadvantage that n-wells of in-pixel PMOS transistors collect some of the signal charge. Goal is ~15 μm pixels so binary readout OK. Prototyping with 130 nm process, plan to move to 90 nm later.

Mechanical issues: Similar to MIMOSA?

CDS: This function is satisfied by means of an in-pixel time-invariant signal processing architecture.

Example of a possible showstopper: Fall short of 100% min-I efficiency due to unavoidable charge collection to additional in-pixel n-wells.

6.8.3. CAP

Brief description: CMOS MAPS, differs from standard approach by signal storage (after charge-to-voltage conversion) on in-pixel capacitors. SW <100 μs for >10 storage elements. Study noise degradation with device size and with reduced capacitance of storage cells, with aim of increasing their number.

Mechanical issues: Similar to MIMOSA?

CDS: Need fast risetime to ensure good fidelity for signal from last BX before each sample. Signals are related to a fixed baseline established at start of train, so there is an exposure to baseline drift during this period.

Example of a possible showstopper: Inadequate pickup immunity due to charge-to-voltage conversion during noisy bunch train.

6.8.4. DEPFET

Brief description: Unique in-house technology in which the signal charge is stored in the bulk, forming an internal gate. Complete clearing of the charge results in precise time-correlated determination of the signal. Rolling shutter readout; column parallel, reading row pairs synchronously in order to achieve the required frame rate.

Mechanical issues: Full-scale devices envisaged, with shorter steering chips strung along the long edge of ladder. Unique mechanical support system (SOI-inspired) based on selective thinning of the handle wafer.

CDS: Despite the rolling shutter readout, pixel-rate CDS is achieved by the sequence; *integrate, sample, clear, sample,* with the two samples closely spaced in time (enabled by the fast clear).

Example of a possible showstopper: Failure to reach required readout rate due to complexities of overall system of 3 chip types working together.

6.8.5. SOI and 3-D

Brief description: An impressive strategy to be liberated from the constraints of CMOS sensors by developing tiered systems. Potential for data driven sparsification with single bunch time stamping, the 'physicists dream'. Current plan is for very small pixels with binary readout.

Mechanical issues: Assume mosaics of reticle-size devices. Will the internal stresses in individual sensors be compatible with the required mechanical stability and material budget of delicate barrel assemblies?

CDS: Robust bunch-crossing rate CDS is of course possible in principle.

Example of a possible showstopper: 4 Gpixels, when full functionality is designed in, may exceed the power limits for gas cooling.

6.8.6. FPCCD

Brief description: Concept is based on the belief (shared by others, and for which there is evidence) that reduction by a facto 20 in effective background level, should suffice. This they achieve by a 20 times increase in pixel granularity, combined with the use of cluster shapes to suppress background.

Mechanical issues: Aim to develop ladder-size CCDs, and use developments from others with similar goals to investigate suitable support structures.

CDS: Robust pixel-rate CDS between bunch trains - no need to break new ground in this area.

Example of a possible showstopper: Insufficient rejection of pair background while retaining full min-I efficiency (possible discrepancies between simulations and real data).

6.8.7. CPCCD

Brief description: Fast readout of standard CCD aiming to reach 50 µs fame rate, thereby cleanly subdividing the bunch train into 20 frames of data. Development with minimal changes from the successful SLD vertex detctor.

Mechanical issues: Most promising approach is probably to mount the ladder-size sensors on silicon carbide foam substrates, having good stiffness, excellent cte match to silicon, easily machinable and generally sound mechanical properties (like beryllium, without the disadvantages).

CDS: Robust pixel-rate CDS standard, but this is going on during the train, so relies on a clean experimental environment, as regards beam-related pickup. Probably that is now a safe assumption. However, there may be significant noise generated by other detector systems during the bunch train.

Example of a possible showstopper: Unacceptable bulk of service electronics (supercapacitors etc) at ladder ends.

6.8.8. ISIS

Brief description: Storage of signal charge in a 20-element in-pixel serial register (hence SW = 50 µs) for leisurely readout between bunch trains.

Mechanical issues: As for CPCCD, but far less services needed at ladder ends.

CDS: Robust pixel-rate CDS between bunch trains - no need to break new ground.

Example of a possible showstopper: Tails from deep p-well could disable charge transfer in CCD register.

6.8.9. Chronopixels

Brief description: Goal to obtain single-bunch timing by implementing pixel functionality that might be fitted into a 10 μm pixel with 45 nm design rules. Such a compact pixel will permit binary readout with sufficient precision. Needs deep p-well to shield signal charge from in-pixel PMOS transistors, which would otherwise collect part of the signal.

Mechanical issues: Probably similar to other ladders built with reticle-size CMOS sensors.

CDS: Robust bunch-crossing rate CDS is of course possible in principle.

Example of a possible showstopper: Same as 3-D - possibly unacceptable power requirements when full functionality is designed in.

6.9.   World-Wide Coordination

To a good approximation, the ILC vertex detector community keeps in touch with what the others are doing, and there is always an open sharing of new ideas and discussion of current problems. As well, most but not all are associated with LOI groups and will participate in the LOI/TD activities. However, most of the work is common to all detector concepts, so the preservation of the ILC VXD community as a collaborating body is really important for at least the next 5 years.

Firstly, there is the long discussed need for a high energy (~100 GeV) test beam with the appropriate (coarse) time structure of the ILC bunch trains. This could be provided by Fermilab or CERN or both. There are many good arguments for having it in one place, which would then become the centre for other shared facilities needed for evaluation of the technologies. These facilities could include a modest-size 3-5 T solenoid, with appropriate apertures for taking the beam through at various angles, including normal to the stack of ladders. They could include an anechoic chamber with calibrated equipment for generating electronic noise, so as to obtain truly comparable data on the noise immunity of different technology options, with possibly evolving designs of noise suppression circuitry, as discussed in Section 6.4. They could include optical equipment for measuring the mechanical stability of ladders when supported as in the intended support system, particularly their freedom from vibration when subjected to pulsed power in the field of the test solenoid. As mentioned in Section 6.1, there are various issues in which it will be important to maintain and enhance the communication path between the vertexing community and the ILC MDI group.

This is only the beginning of a list which effectively calls for some coordination mechanism. Within the vertexing community, there was almost unanimous agreement that this would be desirable, when we discussed it during the review. The only dissenting voice was concerned about possible disagreements within this group, due to conflicts of

interest. He agreed that if the function of this group were primarily to organise facilities and gather data, this should be beneficial. The idea would be for the Vertexing Coordination Group (VCG) to comprise one representative per R&D group, and possibly a few representatives from other coordination groups, for the ILC test beam, DAQ system, etc. This group would be self-organising. They might appoint one member as point of contact with the MDI group, another as point of contact with the test beam group, etc. One advantage would be that the VCG members would develop a detailed understanding of all the technology options, and might be best able to quietly over the coming years self-organise changes of direction when the R&D pointed in particular directions. Such a world-wide coordination at the working level would be much preferred to some high level committee coming in and organising shoot-outs, as have sometimes caused considerable damage to morale and to scientific programmes in the past. Doing this through the LOI groups would not be as effective, since some of the R&D groups are members of one, some of another, and several have by choice and for good reasons no such allegiance at this time. This review committee is speaking for the VXD community we reviewed in recommending to the ILC Detector Directorate that they consider setting up a Vertexing Coordination Group, just as our previous reviews last year suggested a similar structure for tracking and calorimetry.

Acknowledgements