Report on the Research Programme of the Linear Collider Flavour Identification Collaboration

June 2006

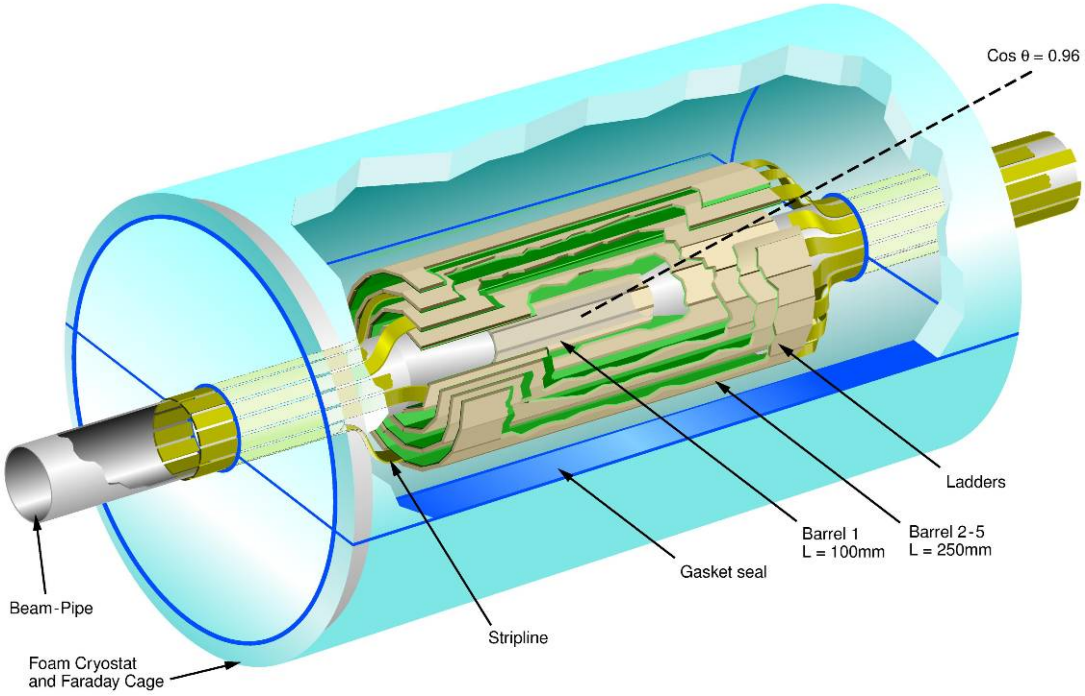

The Universities of Bristol, Edinburgh, Glasgow, Lancaster, Liverpool, Nijmegen and Oxford and the Rutherford Appleton Laboratory

Table of Contents

Introduction

This document describes the progress made by the Linear Collider Flavour Identification Collaboration in the half year since the Oversight Committee meeting in December 2005. The advances made in our research programme are the subject of the majority of this report. Brief mention of developments on other fronts is made here and in the Appendix.

We are pleased to be able to report that the LCFI funding situation has been resolved, firstly through the generous provision of bridging funds by PPARC to cover the period from January to June of 2006 and secondly through the final approval of our funding through the Council of the PPARC: the process that began with the submission of our proposal in April 2005 has thus been brought to a successful conclusion.

LCFI continues to be well represented at national and international workshops and conferences and on various bodies leading efforts towards the construction of the detectors for the International Linear Collider. For example, LCFI members organised a Cosener's Forum on ILC Physics and Detectors, made presentations at recent conferences and workshops, have contributed to the Detector Outline Documents of the Silicon and Large Detector Concepts and one of the editors of the Vertex Detector White Paper is a member of LCFI. A list of the conference presentations and publications given by the Collaboration is provided as part of this report.

The Collaboration is also continuing to strengthen its collaborative links with other groups interested in designing and constructing a vertex detector for the ILC. The focus of our efforts to involve the Valencia group in our activities has changed. As the new Valencia PhD student has an engineering rather than physics background, it is likely that the group will not participate in our physics studies as originally foreseen, but will look at the design of carbon fibre support structures, an area in which they have significant expertise. This collaboration will be aided by our strengthening contact with the Fermilab group investigating carbon fibre support structures for the SiD tracking system.

Collaboration with the DEPFET and Strasbourg groups is also centred on mechanical issues. Both these groups are interested in the problems associated with attaching their sensors to support structures to build ladders. They have agreed to supply us with thinned sensors, giving us the opportunity to gain additional experience with handling processed thinned silicon.

LCFI have also recently been approached by a group from Warsaw University who are interested in taking part in both physics studies and investigations of test CCD structures.

Last, but not least, a group from Lancaster has been successful in obtaining University funding which will allow them to participate in studies of the radiation tolerance of our sensors and to investigate and develop simulations and mathematical models of the performance of CCDs.

WP1 – Physics

Introduction

In the physics area the highest priority near-term goal of the collaboration is the completion and release of the LCFI Vertex Package. This C++ based code, which will provide the vertex finder ZVTOP, flavour tagging tools and vertex charge determination – for charged hadrons closely linked to the sign of the quark charge – will boost the capabilities of the ILC software.

So far, studies requiring vertex detector information are constrained to using the old BRAHMS Monte Carlo (MC) generator and reconstruction with its known limitations, e.g. regarding track-linking, and to fast Monte Carlo programs such as SIMDET or SGV. Since BRAHMS is being phased out and is no longer maintained, such studies are decoupled from any new developments e.g. in the area of particle flow, which makes their usefulness for detector optimisation rather limited.

By providing standard reconstruction tools, the LCFI Vertex Package will enable a wide range of physics studies relying on vertex detector information within the new object oriented software frameworks and hence allow LCFI and other vertex detector groups to optimise their detector designs making full use of all the new developments in other areas. Within the wider ILC community, use of these tools will permit optimisation of various aspects of the global detector design, for instance the design of the FTD, evaluation of the need for an intermediate tracking device between the vertex detector and the main tracker, the design of the main tracker, etc. As well as the detector design, the LCFI package is likely to influence the development of other parts of the ILC software, such as the track reconstruction code. The release of our software is hence eagerly anticipated far beyond the vertex detector community.

The following section gives a more detailed description of the LCFI Vertex Package and outlines our approach to its validation and the current status of its development. Results from the validation of one of the two branches of ZVTOP are presented in a dedicated section. This is followed by a short section on LCFI physics work not directly related to the Vertex Package. The WP1 report concludes with an overview of our plans for the future.

Scope and status of the LCFI Vertex Package

Figure 1 shows a schematic overview of what the LCFI Vertex Package will comprise at the time of the first release. Input to and output from the package will be in LCIO[1] format, providing a set of storage classes on which the ILC community has agreed in order to ensure that data can be exchanged between the three different software frameworks that exist in the three regions involved in the ILC. Input will either be provided by SGV[2], which has been used within LCFI so far and thus allows detailed comparisons and cross checks with previous results, or from MarlinReco[3], the software framework under development in Europe, to which we will provide an interface.

The central part of the software will be the vertex finder ZVTOP, consisting of the two branches ZVRES and ZVKIN. While ZVRES is a very general algorithm coping

with arbitrary multi-prong decay topologies, ZVKIN, also known as the ghost track algorithm, is more specialised. Using the fact that, due to the boost of the B hadron, the interaction point (IP), the B hadron decay vertex and the subsequent D hadron decay vertex lie approximately on a straight line, the algorithm defines a "ghost track" approximating the direction of flight of the decay chain and includes it in the search for vertices. Using this additional information, the ghost track algorithm is able to correctly assign tracks to vertices in decay topologies with one or more 1-prong vertices, thus extending coverage and improving the flavour tagging capabilities as shown for SLD[4], where this code was first developed. For ILC studies, only the ZVRES branch has been available so far, so the LCFI code will extend the scope of ILC software and should improve its performance.

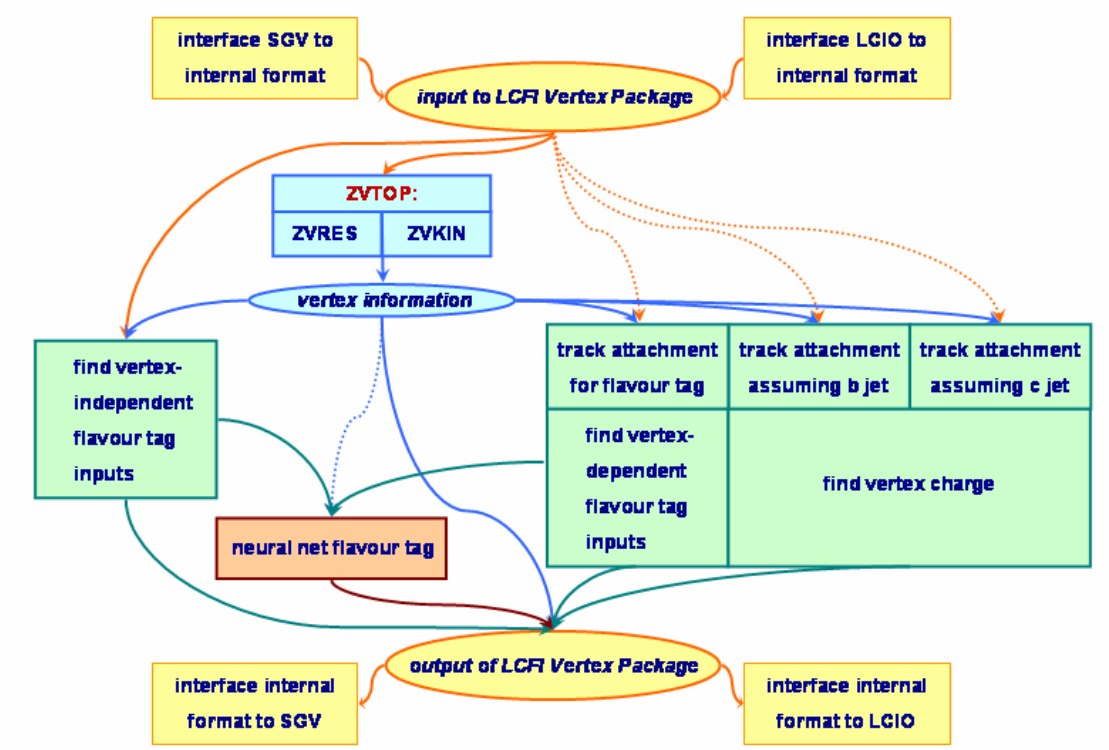

Figure 1 Structure of the LCFI Vertex Package and flow of information between the different parts of the software.

The flavour tag procedure we will implement as the default option was developed for the ILC by Richard Hawkings[5] and is based on nine neural networks, three each for the cases that one, two, or more vertices are found by ZVTOP. In each of these cases, there are separate networks to identify b-jets and c-jets for arbitrary background. Since for some physics processes the background only consists of b-jets, and discerning c-jets from b-jets is simpler than general c-jet identification, networks are trained for this case as well.

If only the IP is found by ZVTOP, the flavour tag is based on information from the two tracks in the jet that have the highest impact parameter significance, plus the "joint probability", i.e. the probability for all tracks to come from the IP. If secondary vertices are found, the neural net inputs are calculated from a set of tracks which contains one or two vertices from ZVTOP and in some cases further tracks assigned by a specific procedure devised to recover some 1-prong topologies for b-jets when running ZVRES. This "track attachment" step, in which additional tracks are added to

the set of tracks contained in secondary ZVTOP vertices, may need to be tuned depending on whether the tracks are to be used for obtaining the flavour tag inputs or the vertex charge, and for vertex charge calculation may differ depending on whether the jet is assumed to arise from a b- or from a c-quark. Track attachment for b-jets has been carefully optimised by LCFI in 2004 using the SGV fast Monte Carlo (MC) simulation. For the first release, we will implement that procedure also for c-jets, but will design the software such that it can later be replaced by an improved algorithm.

For training the neural networks and for obtaining the output from pre-trained neural nets we will use C++ based neural network software developed earlier within the Bristol LCFI group. This allows flexible definition of the network architecture (number of layers and nodes), transfer function and training algorithm.

Development of the LCFI Vertex Package is progressing well. Work on writing new code and documenting and testing it is performed in parallel. For validation of the ZVRES code, it has been interfaced to SGV, allowing detailed performance comparisons between FORTRAN and C++ using identical input events.

The main focus of the work over the past months has been on testing the more conventional ZVRES branch of the code, which is now converging, as described in the following section. This branch of ZVTOP is conceptually more complex than the ghost track algorithm. Also, the ghost track algorithm will to a large extent build on C++ code already developed and tested for ZVRES, such as the vertex fitter, which will save time in the development of that branch of the code. The first stage of the ghost track algorithm, which finds the direction of the ghost track, has already been coded and results from initial tests are encouraging.

To arrive at a coherent software package, we are currently designing the class structure for the parts other than ZVTOP. These classes fulfil several purposes at the same time: they will form the interface to SGV and MarlinReco (see above) as well as provide the “working classes” needed for flavour tag and vertex charge, e.g. for track attachment and calculation of the flavour tag inputs. In addition, they have to ensure that information can be passed around efficiently between the different parts of the code. Since we will use LCIO for both input and output of our software, these working classes will be LCIO based, inheriting from the LCIO storage classes and extending their capabilities for the purposes of our code in a way that users will not have to know about. This encapsulation of functionality is encouraged by the developers of the ILC software frameworks to ensure that non-experts can use the software in a straightforward way.

For the output of part of our code, we have suggested that a dedicated Vertex class be added to LCIO. Design of that storage class has profited from advice from Frank Gaede, one of the authors of LCIO, who has also provided useful information on MarlinReco. Following advice from the LCIO developers group, we have published the suggestion on the ILC software forum[6] for further discussion within the ILC community. The LCIO developers group are supporting our suggestion, which they will discuss in a dedicated meeting in the near future.

ZVRES validation

Validation of the C++ ZVRES code is far advanced, showing overall excellent agreement with the FORTRAN version. Starting from simple checks at the single track level, increasingly complex aspects of the code have been tested, as detailed below, leading up to a test of the full ZVRES branch within the SGV framework.

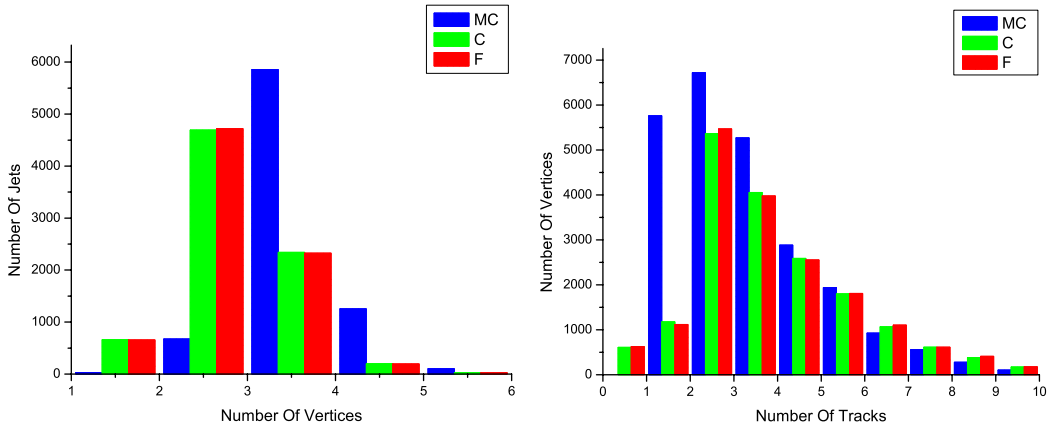

Figure 2 Distributions of number of vertices (left) and of track per vertex (right) for C++ (C) and FORTRAN (F) ZVRES and for the MC, obtained from a sample of 100 GeV b-jets.

Initial tests focused on identifying and correcting various errors in track swimming and vertex fitting. To this end, the IP and one track were input to the code, with errors fixed by hand to a certain value, to allow checks of whether the track is identified as part of the IP up to the expected threshold value of the distance between the two beyond which the two objects are resolved from each other. Similar tests were later performed feeding two tracks plus the IP into the code. For these conditions, agreement between the two versions is now at the level of 99.8%, the remaining 0.2% most likely being due to differences in the implementation of vertex fitting in the two versions.

C++ ZVRES							
Monte Carlo track origin	**Two vertex case**			**Three vertex case**			
	pri	**sec**	**iso**	**Pri**	**Sec**	**ter**	**iso**
Primary	98.1	0.7	1.2	96.9	2.2	0.1	0.9
B decay	8.3	75.6	16.1	2.3	89.1	4.7	3.9
D decay	2.3	79.4	18.3	0.6	17.5	77.5	4.5

FORTRAN ZVRES							
Monte Carlo track origin	**Two vertex case**			**Three vertex case**			
	pri	**sec**	**iso**	**pri**	**sec**	**ter**	**iso**
Primary	98.0	0.6	1.4	97.3	1.5	0.0	1.2
B decay	9.3	74.8	15.9	2.6	90.6	3.9	2.9
D decay	2.5	81.1	16.4	0.6	17.2	79.0	3.2

Table 1 Purity of reconstructed track-vertex association (%) for the C++ and the FORTRAN ZVRES branch of ZVTOP for 100 GeV b-jets. Abbreviations for reconstruction level track assignment are “pri” (primary), “sec” (secondary), “ter” (tertiary) and “iso” (isolated).

Using a sample of 100 GeV b-jets, Figure 2 shows the agreement between the C++ and FORTRAN versions of ZVRES in the number of vertices found and in the number of tracks per vertex. Differences between reconstructed number of vertices and the MC distribution result from the complexity of the decay chain, in which

distances between B hadron decay and IP or between D and B hadron decay are frequently too short to be resolved by the detector.

Track-to-vertex association, in the form shown in Table 1 of the original ZVTOP paper[7], was studied for jet energies ranging from 25 to 250 GeV. Table 1 here gives an example for 100 GeV jets. The first row of this Table lists the percentages of tracks coming from the IP at MC-truth level which are assigned to the primary, secondary and – if three vertices are found – tertiary vertex, or left over as isolated tracks. The following rows show the equivalent numbers for the secondary and tertiary MC vertex. Excellent agreement of the values between the two versions of the code is seen, which is similarly good over the full energy range. The fraction of jets with exact agreement in number of vertices found and of all tracks assigned to them is between 70 and 75%, depending on jet energy.

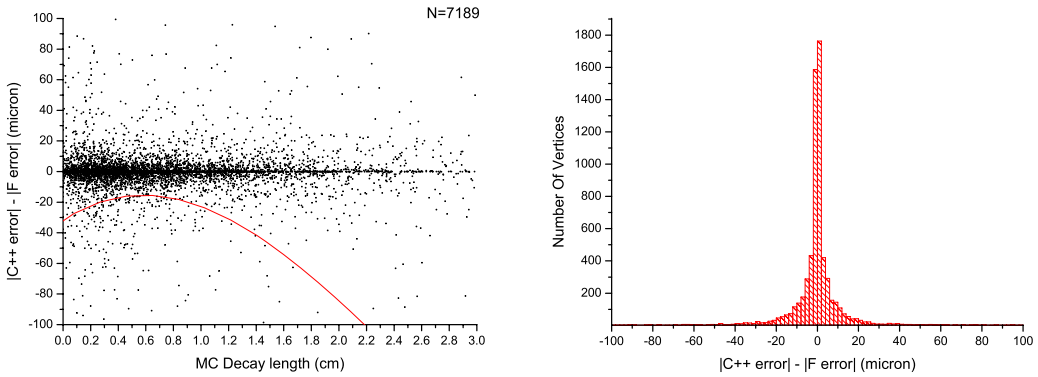

(a) (b)

Figure 3 Difference between the C++ and the FORTRAN deviation from MC in decay length, plotted as a function of decay length (a) and projected onto the y-axis (b). The red line in (a) shows a second order polynomial fit to the distribution.

Physics performance not only depends on the track content of the decay vertices but also on how well their position is obtained from the fit to that set of tracks. This can be measured by calculating the difference between the reconstructed decay length and the MC value. In Figure 3(a) the difference found from the FORTRAN version of ZVRES is subtracted from the difference found for the C++ version and plotted as function of the MC decay length. Thus values smaller than zero correspond to cases in which the C++ comes closer to the MC truth, that is, performs better than the FORTRAN. For most vertices, agreement is very good. In some cases vertices are only found by the C++ version. Figure 3(b) gives a projection of the plot shown in (a) onto the y-axis, with the peak around zero corresponding to small differences arising from the vertex fit and the tails being due to differences in vertex finding. The red curve in part (a) of the Figure shows a second order polynomial fit to all the points, indicating that, on average, the C++ version ZVRES now performs better than the FORTRAN version.

Other work

The automated software that was developed last year for the study of the effect of the beam pipe radius on physics potential of the detector has allowed us to investigate further effects of varying the detector design with minimal investment of effort. The

setup calculates the leakage rate λ_0, i.e. the probability of reconstructing a neutral B hadron as charged, as a function of polar angle for a range of jet energies from 25 to 250 GeV.

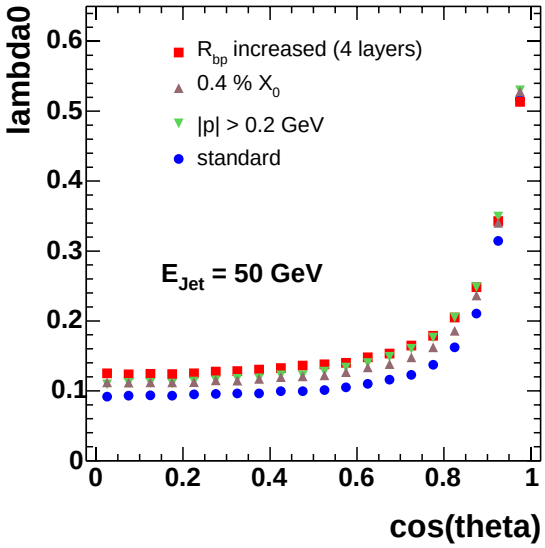

Figure 4 Leakage rate as a function of polar angle for different detectors and reconstruction procedures.

Figure 4 shows a comparison of three effects: last year's result on increasing the beam pipe radius to from 15 to 25 mm, increasing the thickness of the barrel staves from 0.1% to 0.4% of a radiation length per layer, and increasing the minimum momentum requested for tracks to be passed to ZVTOP from 0.1 to 0.2 GeV. While increase of the beam pipe radius has the largest effect, increase of layer thickness and raising the momentum cut-off also lead to sizable degradations. The result emphasises the need to push all these parameters to their limits, since in the real detector all these effects will add up.

Future plans

For the next few months, all effort in the physics area will be focused on completing the LCFI Vertex Package, with a clear plan in place for who will work on which of the remaining tasks. In terms of coding, the main areas needing further work are the completion of the overall class structure and interface for the package (see the section on scope and status for further details), writing of the second half of the ghost track algorithm and translating the calculation of the flavour tag inputs to C++.

Validation of ZVKIN will profit from the work already done in connection with testing ZVRES, since in the C++ version, both algorithms have a number of central classes in common.

We are currently beginning test runs of the full chain of C++ ZVRES, interfaced to SGV, feeding the output into the FORTRAN flavour tag and vertex charge procedure, using the automated setup of job submission and analysis scripts mentioned in the previous section. These tests will be repeated when the C++ versions of the flavour tag and vertex charge determination are complete. Tests of the entire code will be followed by tests within the MarlinReco framework.

Current tests show that some work will be needed to reduce the run time of ZVRES, which is of the order of 3 seconds per jet at the moment in contrast to about 30 ms with the FORTRAN version. For this purpose, we will use a software tool allowing us to identify which fractions of CPU time are spent in which parts of the code. Initial tests with this tool indicate that one of the areas for improvement is the vertex fitter, and we already have some ideas on how a reduction of run time can be achieved.

We aim for a release of the code to the ILC community as soon as possible. Using optimistic time estimates, and assuming that no unexpected performance issues arise in the system test, indicating the need for further work, this should be possible at the end of September, as stated in the previous report.

Beyond the release of the Vertex Package, work on improvement of these tools will continue. While tests of the code within MarlinReco before the release will be constrained to ensuring the software works and yields reasonable performance, the release will be followed by more detailed comparisons of results from MarlinReco and SGV.

Starting with the simplified "cheater" algorithms implemented in MarlinReco, in which tracks are fitted from a known set of hits based on MC information, similar to what is done in SGV, comparisons with results from the full track finding procedure in MarlinReco will be made to evaluate how performance changes. This is currently an area of intense code development by other groups, with tracking based purely on vertex detector information having been published recently[8] and combination with the main tracker and the forward tracking system planned in the near future. Comparisons of performance of our code running with the available tracking options will allow us to give useful feedback to these developments.

Another obvious area for further study is the full exploration of the improvements that can be achieved by the use of the ghost track algorithm at the ILC. These improvements will affect both the flavour tag and the vertex charge determination. The ZVKIN algorithm is the basis for the charge dipole procedure, from which the quark charge can be obtained for short-lived B_d^0 decays.

In parallel to further development of the software tools, we will take up the study of physics channels sensitive to vertex detector performance.

WP2 – Sensor Design and Production

Introduction

The first successful tests of the low-speed version of CPC2-10 have allowed us to proceed with the manufacture of the busline-free double-level metal devices. Presently, the final processing steps of 4 wafers are being completed at e2V. An error during the manufacture of the ISIS1 with the p-well was discovered and currently we are discussing with e2V the possible re-run of 4 wafers to rectify this.

Most of the effort during the last 6-month period was dedicated to the design of the test CCDs targeting a reduction of the clock amplitudes and gate capacitance. Several new ideas for capacitance reduction have emerged, which promise significant improvements in the power dissipation and the integration of the CPD driver into the detector. The work on the evaluation of new technologies for future sensors has continued as well.

Sensor manufacture

The batch of 20 CPC2/ISIS1 wafers at e2V was split into 2 halves to minimise the risk of failure and processing completed to the steps before metallisation. So far 6 wafers (numbers 1, 2, 3, 17, 21 and 22) from the first half have been finished with single level metallisation. Three wafers based on 100 Ω.cm 25 μm thick epitaxial layers (numbers 1, 2 and 3) and one made on a 1.5 kΩ.cm 50 μm epitaxial layer (number 17) have been delivered to LCFI. Wafers 2 and 3 were sent undiced to VTT for bump bonding. Wafers 1 and 17 were diced and delivered to RAL for tests.

The remaining two wafers were processed with a deep p-well for the ISIS1 and do not provide usable CPC2 chips. Only one of the ISIS1 wafers was diced and delivered, for reasons that are explained in the WP5 report.

After observing the first signals from the CPC2 devices, a decision was made to proceed with the two-level metallisation for the remaining 4 wafers and produce busline-free CCDs. The wafers are processed with first level metal and an insulating polyimide coating before being sent to a vendor outside e2V (Innos, Southampton) for second level metallisation. A sputtering process is used for the metal deposition, which removes the native aluminium oxide *in situ* and ensures a high quality of interconnect between the two metal layers. This procedure and vendor have been used before for two-level metal development at e2V, with very high yield.

At the moment, the 4 wafers are receiving second level metal layer, after which they will be passivated at e2V, probed and diced.

During the tests of ISIS1, it was found that the doping of the p-well does not match the design value, which caused very high transistor thresholds and low channel potential. SIMS measurements showed that the p-well had received 2.5 times the desired dopant concentration due to a manufacturing error. We are currently discussing with e2V a re-run of the ISIS1 with p-well, using the design doping on 4 wafers.

Design of test CCDs (CCD-T)

During the last months, significant progress was made in the design of the test CCDs aimed at reducing the gate capacitance and the clock voltages for the next generation CPCCD. Meetings with e2V to discuss this development were held on a monthly basis.

Two devices for clock amplitude studies with rectangular and profiled gates are under consideration. They will be 2-phase devices based on the CPC2 design, but the 4 polysilicon gates will have separate connections. This will allow us to study the fundamental limitations on the amplitudes needed to reliably transport the charge at high frequencies. External offset voltages applied between a pair of gates can be used to create an inter-gate barrier of arbitrary height. Some devices will have low or no p-type inter-gate implant specifically for these studies.

Another process variation will involve making the p-type inter-gate implantation before the oxidation of the polysilicon gates. Simulations show that this creates a much lower potential pocket in the inter-gate gap and can therefore allow lower clock voltages to be achieved. In the third process modification, the inter-gate barrier will be made by using silicon nitride of different thicknesses, which has the advantage of eliminating the need for implantation of very small dopant levels, which are difficult to control accurately.

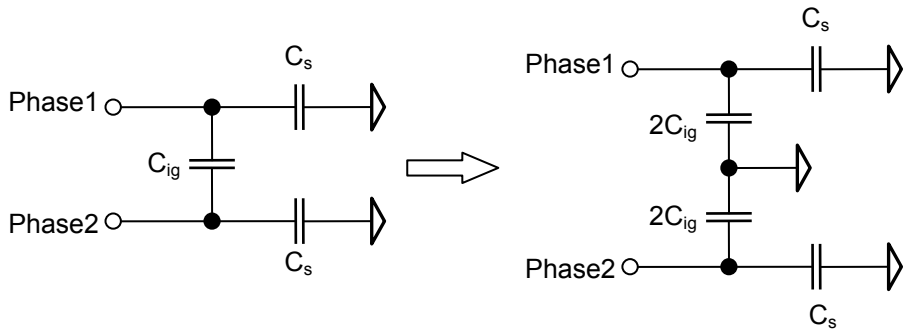

Figure 5 Equivalent capacitance circuit for a 2-phase CCD

The test chips will share a common design and size. They will have 10 columns and approximately 500 rows of 20×20 μm^2 pixels. Four columns in the centre of the array will be equipped with 2-stage source followers, allowing operation at frequencies of up to 50 MHz with an external load. The phase capacitance is estimated to be 150 pF, which can easily be driven at the highest design frequency. Because of the low gate capacitance, the on-chip clock distribution buslines are not expected to be challenging.

Reducing the CCD capacitance could offer numerous advantages for the clock driver system and the overall power dissipation. In the last months LCFI has formulated several ideas which have the potential to reduce the gate capacitance by up to a factor of 4. Around 70% of the CCD capacitance is caused by the inter-gate component C_{ig} and our efforts are concentrated first on reducing this. As shown in Figure 5, the equivalent gate capacitance to ground for a 2-phase CCD is dominated by C_{ig}, because of the clock symmetry. Simulations have shown that C_{ig} does not depend strongly on the actual shape of the gate overlap (obtained from SEM images) or on the gate thickness, as we originally thought. In fact, C_{ig} is dominated almost entirely by the

geometrical capacitance between the gate plates. This important conclusion opened up two new directions for capacitance reduction.

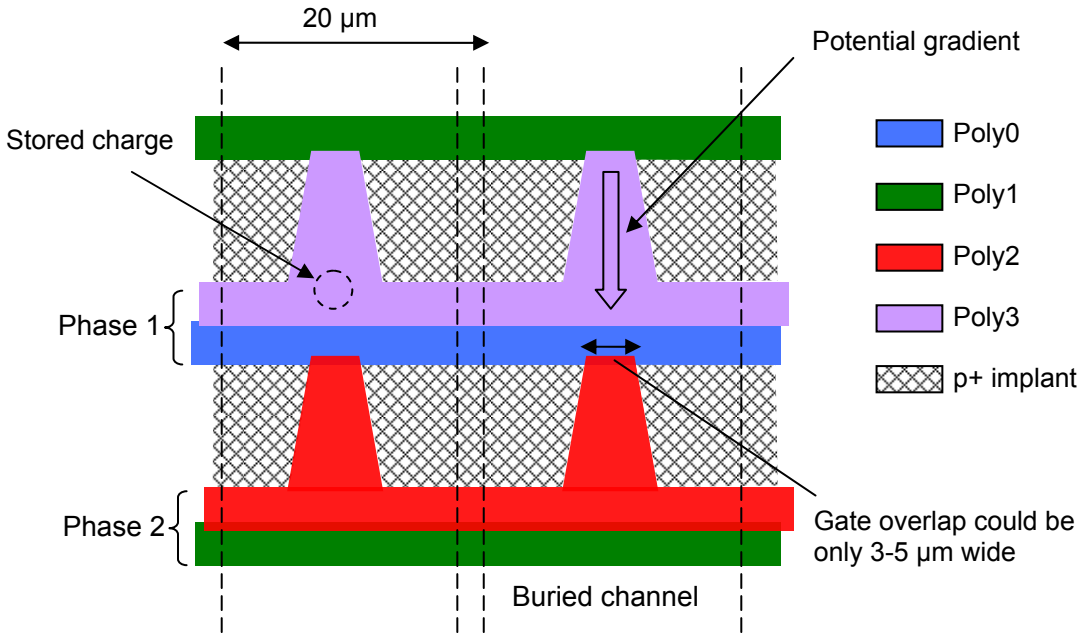

Figure 6 Schematic design of open phase CCD

The first is to keep the same shape of the gate overlap, but to reduce the gate area. Devices with such architecture have been made by e2V for astronomy applications, where capacitance reduction is not pursued but comes as a by-product. Examination of existing capacitance measurements made on such open-phase devices (CCD22) shows that a factor of 2 reduction in C_{ig} is possible. The device uses profiled polysilicon gates to create a potential gradient in-pixel in the direction of the transfer. Self-aligned shallow p+ implants are used to define the potential of the open silicon surface and to confine the charge to the buried channel under the profiled gates. Such a structure is illustrated in Figure 6.

The gate overlap between adjacent phases could be only 3...5 µm wide and one would naively expect a reduction of C_{ig} by a factor of about 4 to 7. In practice, the reduction is around 2, because the p+ areas are partially cut off from the substrate and serve as conductive planes between the two phases, adding to C_{ig}. The signal charges in the open-phase CCD are stored in a small volume at the base of the trapezoidal gates, which improves the radiation hardness with respect to bulk damage. This could be further strengthened by adding a narrow supplementary (or notch) channel. The open-phase CCD uses 3D potential effects to create the inter-gate barrier instead of implants, which inherently offers low clock amplitude drive.

The second route to low C_{ig} is to increase the distance between the two phase gates, thus reducing the mutual geometrical capacitance. Since a gap between the gates will result in poor charge transfer, another “pedestal” gate is added, biased at a fixed voltage, as shown in Figure 7. The pedestal gate is sufficiently wide to greatly reduce the capacitive coupling between Phase1 and Phase2. The equivalent capacitance of the 2 phase gates to ground is $2C_{ig}$* instead of $4C_{ig}$ because the pedestal gate is at AC ground, unlike the alternating phases of the standard 2-phase CCD. Further

improvement could be achieved by raising the pedestal gate above the gate dielectric of the adjacent phase gates. Since $C_{ig}^* \leq C_{ig}$, the capacitance reduction is at least a factor of 2, and calculations, shown in Figure 8, reveal that the total reduction could reach a factor of about 4.

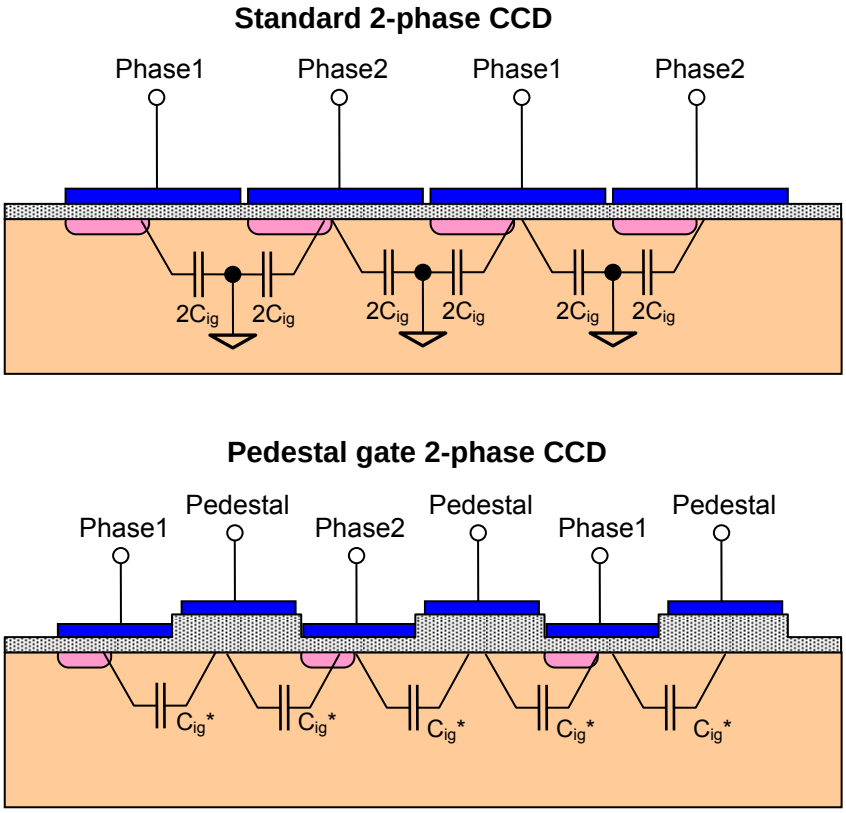

Figure 7 Equivalent capacitance of the pedestal gate CCD. The capacitance to substrate is not taken into account.

For this structure to work, a potential gradient is needed to achieve unidirectional charge transfer under the pedestal gate. This could be provided by profiled doping similar to the open-phase CCD. In fact, the same approach could be used for the phase gates as well, resulting in a buried channel consisting of stacked trapezoidal sections. This structure, dubbed the "Christmas tree CCD", could be easier to manufacture with the existing technology. Different gate thicknesses in the pedestal gate CCD could create potential pockets at the gate edges, which trap signal charge. This effect has been studied in simulation in order to eliminate it. Another potential problem is the higher rate of accumulation of radiation-induced trapped charge under the pedestal gate because of the greater oxide thickness. This could be compensated by adjustment of the pedestal bias and reduced by using a silicon nitride dielectric. Extensive simulations on the pedestal gate structure show that this is a viable option with great promise.

Another idea, this time coming from e2V, is to eliminate the gate overlap over the channel stops of the standard 2-phase CCD. The buried channel is narrow in order to minimize the overlap, and to reduce the capacitance to substrate the oxide is much thicker over the channel stops. The benefit of this approach has yet to be investigated.

The scope of this work has expanded significantly over that originally foreseen, to the point at which a dedicated wafer run will be needed to explore all the options for clock amplitude and capacitance reduction. A much less demanding "passenger" CCD batch alongside the main production at e2V was envisaged in the LCFI Proposal.

We have received a provisional quotation from e2V for the design and production of CPC-T. There will be 2 dedicated batches, each with 6 wafers. The 2 batches will be made with different process modifications and it would even be possible to have process variations within one batch, for example 3 groups of 2 wafers with different parameters. The masks will include about 8 different devices with the common design described above. Due to their interest in applying these ideas to other devices, e2V have kept the cost for the 2 batches to that in the Proposal, i.e. £130k. If the second batch requires any mask changes, the price could increase by up to £70k. The mask will be repeated as many times as the area on the wafers allows, producing hundreds of CPC-T chips.

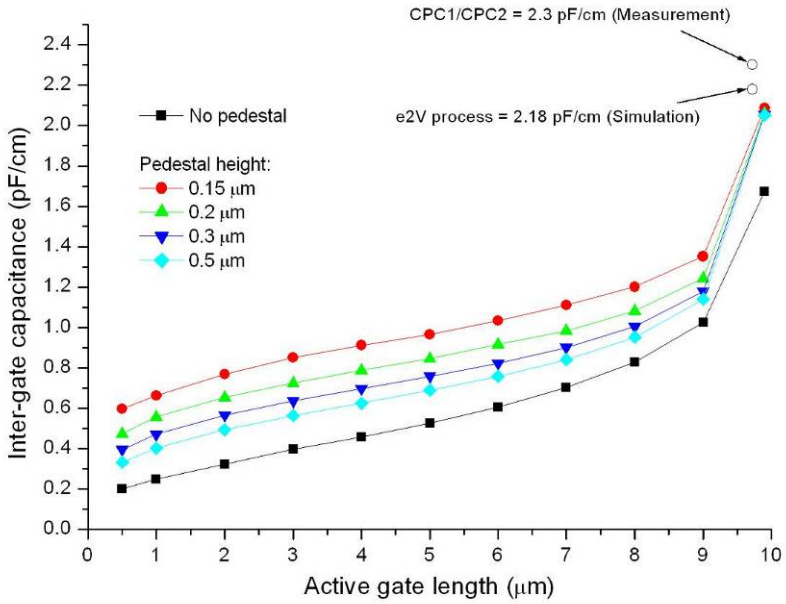

Figure 8 The inter-gate capacitance as a function of the length and height of the pedestal gate. The pixel pitch is 20 μm.

Evaluation of new technologies

Another area of interest for the next generations of CPC and ISIS chips is the evaluation of new commercially available technology from leading manufacturers. Of particular importance are developments which could result in the manufacture of very small CCD pixels for the ISIS2 or devices with low inter-gate capacitance for CPC3. LCFI received five FTF2416 CCDs from DALSA. These are 4 Mpixel image sensors with 9 μm square pixels. The devices are built with non-overlapping gates, which is typical of small pixel CCD and CMOS processes. This technology could produce 3-phase pixels as small as 2.5 μm, but its surface radiation hardness is unknown. The possible problem is that the trapped charge in the thick layers of SiO_2 above the

exposed inter-gate gap could create potential pockets in the buried channel. This would be particularly visible at the low clock voltages used in the CPC and ISIS.

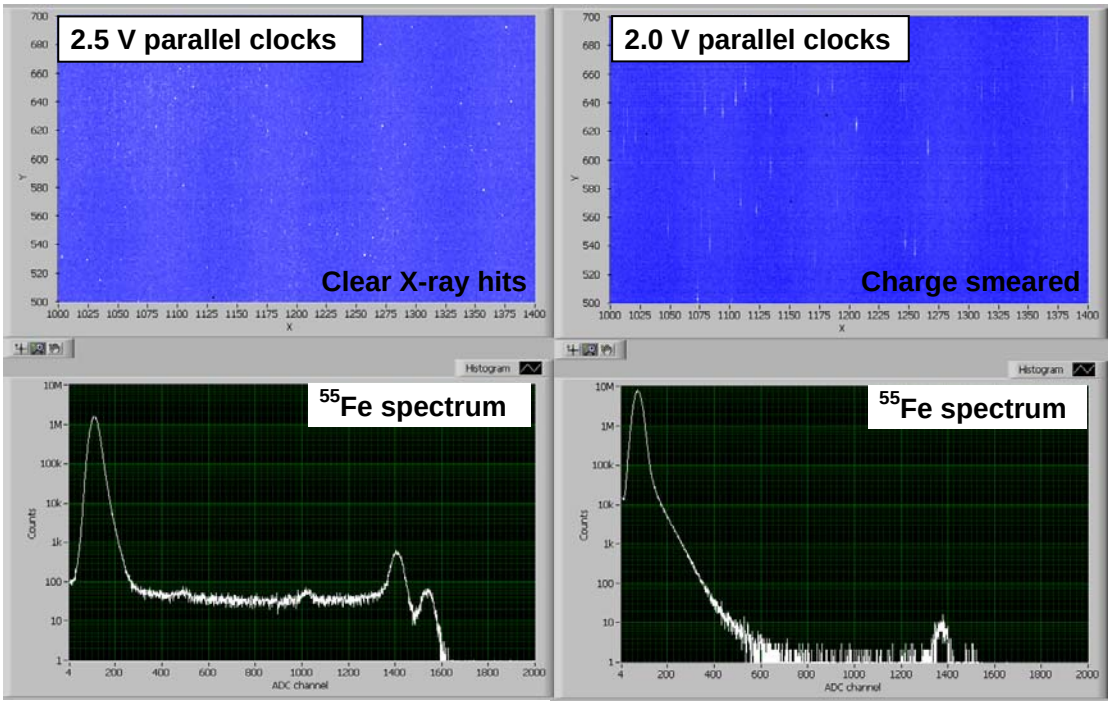

Figure 9 Performance of DALSA FTF2416 at very low parallel clock amplitudes. Even at 2.5 V, the Mn K_α and K_β lines are clearly resolved due to the very low readout noise. The serial clock frequency is 20 MHz.

We initiated a study of this issue, the first step of which was to verify that the FTF2416 can work at very low clock voltages. The supplied hardware was placed in a cryostat and modified to reduce the image clocks below the nominal level of 8 V. The CCD performed extremely well, showing the excellent noise performance of 20 e^- ENC and high quality X-ray response using a ^{55}Fe source, as shown in Figure 9. Charge smear due to insufficient clock amplitude is clearly seen at 2 V. The next step will be to induce flat-band voltage shifts by irradiating these devices using the strong ^{90}Sr source available at RAL.

Future plans

The first busline-free CPC2 chips are expected to arrive in July 2006. After their evaluation by the end of the year, we will decide the number of single- and double-level metal devices to be manufactured from the remaining 10 wafers. The final chip delivery is expected in March 2007.

Provided that a re-run of the faulty ISIS1 with p-well is possible, we will receive 4 new wafers at a time still to be negotiated with e2V.

Significant effort in WP2 is now focused on the design and simulations for the test CCDs. It is important that we finish the design, together with e2V, by October 2006. The large range of CPC-T devices will provide us with valuable insight on the different ways to improve the CPCCD technology and advance it towards the future CPC3. The manufacturing time could be longer than usual due to the complexity of

the processing options and it is expected that it will take 3 months for the first devices to appear.

At the same time, LCFI will continue the work towards the second generation ISIS. At present little effort is available for this, because the new RA at RAL has not yet been recruited. The RA will be charged with most of the device testing in WP5 and will contribute with simulations for new designs, thus enabling key people to spend more time on WP2.

The evaluation of the radiation hardness of the DALSA process is expected to finish by September. After that, we will start detailed discussions with semiconductor vendors, including DALSA, on the first steps in the design of ISIS2.

WP3 – Read-out and Drive Electronics

Introduction

The priority for the last six months has been the design, simulation and layout of the CCD Clock Driver ASIC, CPD1. The Driver will provide the high currents necessary to charge and discharge the large CCD gate capacitances, at frequencies up to 50MHz. The Driver outputs have controlled current limiting, so it is possible to reduce the peak currents when the CCD is operating at lower frequencies.

The drive strengths for positive and negative clock edges are independently adjustable. This has the benefit that the positive edge of the Phase 1 clock can be matched to the negative edge of Phase 2, giving balanced current flows in the bond wires. This minimises the effect of bond wire inductance, thereby improving the transient response.

Summary of drive requirements

1. CCD load: 127 nF to ground per phase, 0.1 Ω distributed gate resistance.

2. Clock waveform: 3.3 V peak-to peak at 25 MHz.

3. Peak current: about 21 A per output (to charge 127 nF with a 1.65 V step in 10 ns).

4. Continuous clock time: 1 ms, with a duty cycle of 0.5%.

5. Higher speed operation with reduced capacitance (50 MHz, 40 nF).

Detailed design

Mode control

The CPD1 has three modes of operation, suitable for different clock frequencies:

a) The highest speed is achieved by driving the output transistors in Inverter Mode, where the gates are connected directly to Vdd/Gnd by switch transistors. This results in the largest swing of gate voltage, with short time constants. However, there is no control over the output slew rate. This mode is best suited to 50 MHz CCD clocking.

b) Frequencies from ~ 5 to ~ 25 MHz can be generated in Fast Mode, where the output gates are driven from source followers with variable input voltage. The NMOS output transistor gates are connected to PMOS source followers, which gives a range of voltage from ~ 1 V to ~ Vdd. The lower gate voltage reduces the overdrive of the NMOS devices, limiting the current flow and creating a ramp waveform.

c) Frequencies from ~ 1 MHz to ~ 10 MHz are generated in Slow Mode, which again features source followers. In this case, the NMOS output transistors are connected to NMOS followers, with a voltage range of ~ 2 V to ~Gnd. This gives accurate control over the slew rate for the lower range of frequencies.

Output enabling

The CPD1 layout is segmented into 8 blocks, each with 4 pads for Phase 1 and 4 pads for Phase 2. The segments are enabled when the corresponding bits in the control register are set high. The enable function will be useful for testing the Driver over a wide range of load capacitors and bonding configurations.

Test channels

Two Driver outputs are permanently connected to an on-chip capacitor of 2 nF. The connection is made through wide metal tracks, which gives much lower inductance than is achievable with an off-chip capacitor. The Driver performance for these channels will not be limited by the external bond wire inductance.

Control circuitry

The Driver includes two DACs to provide references for slew rate control - one for the positive edges and one for the negative edges. The DAC input bits are stored in a serial register which also controls the Mode selection and Output Enable function. This register is loaded in parallel from a second register which is loaded sequentially from the external data input line. The parallel load function ensures that all of the internal control bits change simultaneously, so the Driver is always in a well-defined state. The Driver has a Power Down feature which reduces the static current to zero when the clocks are inactive.

Die layout

The die size is 8×3 mm^2, with the main power supply pads along the left edge, and clock outputs on the right. These sets of pads are on a 100 micron pitch, as is illustrated in Figure 10. The control pads, also shown in the Figure, are along the bottom edge on a 200 micron pitch.

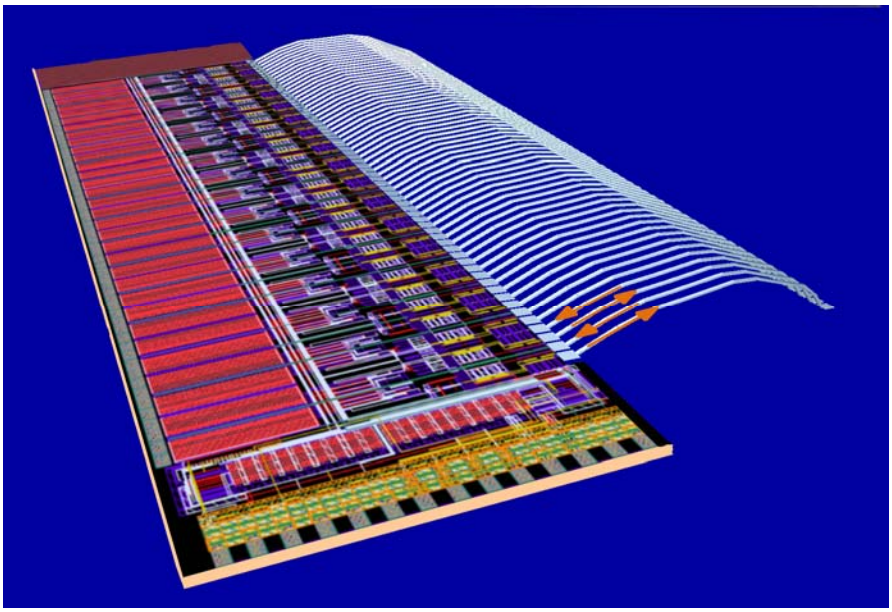

Figure 10 CPD1 layout with output bond wires

The Driver pads alternate Vdd with Gnd, and Phase 1 with Phase 2, in order to keep the bond wire inductances low: when current flows out of the Phase 1 pads, equal and opposite current flows into the Phase 2 pads, the balancing of the currents resulting in the partial cancellation of local magnetic fields, reducing inductance.

Example waveforms

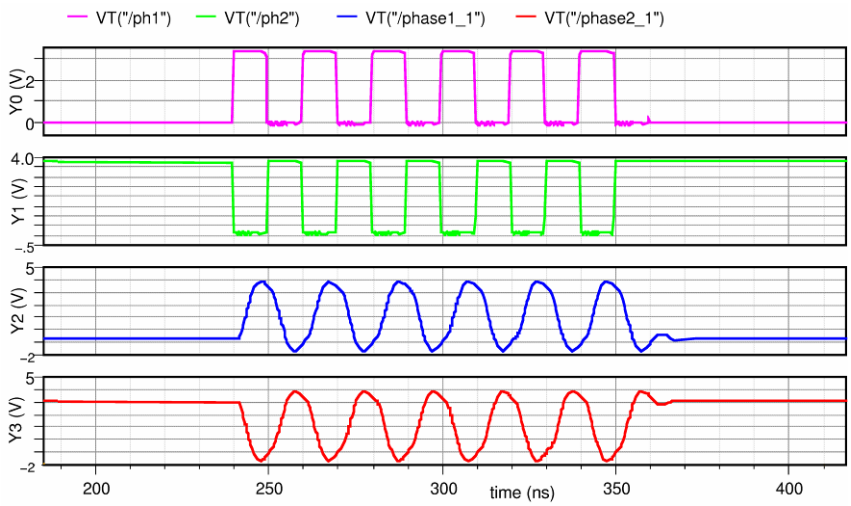

Figure 11 CPD waveforms in Inverter Mode at 50 MHz.

Figure 11 shows the typical simulated outputs at 50 MHz, with the full load capacitance and bond wire inductance included. The signals phase1_1 and phase2_1 are monitored after the bond wires, and show the effects of the resonance of the LC network. The increase of voltage due to resonance is an undesirable condition, as the output then becomes strongly dependent on parasitic effects. Resonance will not occur when the Driver is bump-bonded, as the inductances will be much smaller.

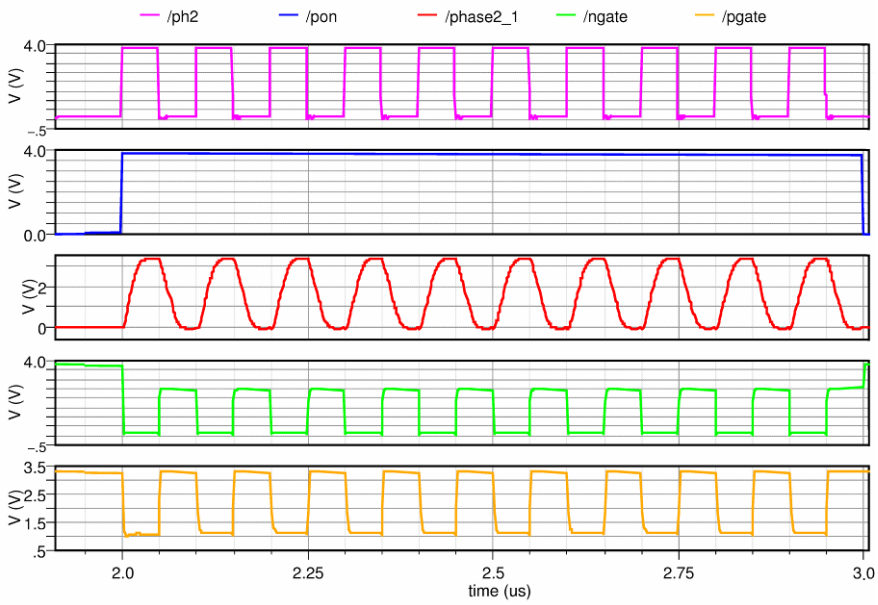

Figure 12 CPD waveforms in Fast Mode at 10 MHz

Figure 12 shows the slew limited phase2_1 output for Fast Mode at 10 MHz. The output transistor gate voltages (ngate and pgate) have reduced voltage swing during the clocking period. When the Driver is powered down (pon low), ngate goes to Vdd

which holds the output at Gnd. The pgate voltage is high, so the PMOS output transistors are disabled and there is no output current.

Future Plans

The main goals of this Work Package in the forthcoming 6 month period are the submission of CPD1 for manufacture. This should be done within the next month. Effort will then shift to further testing of both the digital and analogue sections of the CPR2. These tests will feed into the design of the next readout chip, the CPR2A. The design of this chip will be completed before the end of the year.

WP4 – External Electronics

Introduction

The testing of the ISIS and CPCCD sensors described in this report requires the development of electronic systems which provide all the control, biasing, readout and monitoring necessary for their operation and for the operation of the associated readout and drive chips. Further, test systems for the column parallel drive (CPD) chips must be developed and constructed.

The efficient propagation of clock signals into and across the CPCCD becomes increasingly challenging as both the frequency and the sensor size increase, due to the large capacitance of the sensor. Large currents are required to obtain the voltages needed across this capacitance and even small resistive and inductive impedances in the path of that current can have severe consequences. Minimization of the inductance of the clock connections is an important goal which can be achieved, for example, by exploiting the cancellation of the magnetic fields induced by the opposite currents of the two clock phases.

Several techniques for producing the clock signals are being investigated. It is hoped that the custom CPD chip that is being developed will ultimately provide the drive for the full-scale and full speed CPCCDs, but we are also using miniature air-cored planar transformers, implemented in a multi-layer PCB, and a commercial driver chip for test purposes in the interim. Initial tests of these were successful.

Design and production of CPCCD motherboards

The motherboard is the PCB on which the CPCCD and its readout and drive circuitry sit. Development of this board was particularly challenging due to its overall complexity and the space constraints involved. The 10-layer board must provide all the bias voltages for the CCD, high current clock signals, analogue readout for 12 CCD columns, the digital interface for two types of readout chips, temperature monitoring and also a well for the CCD sensor surrounded by fine-pitch bonding pads.

Two new variants of the CPC2 motherboard have been produced since December 2005: MB 4.2 and MB 4.3. MB 4.2 implemented the MAX5057 MOSFET driver chip, optimized for the available single-metal CPC2, and is shown in Figure 13. Several boards have been assembled and used for CPC2 testing. Some noise pickup from the fast drivers was observed and studied at RAL and Oxford.

MB4.2 is the first motherboard using square pulses with a few nanosecond transitions to drive the clock, as opposed to the sinusoidal signals used in the past for CPC1. This explains the higher sensitivity to pickup. An improved layout of the ground planes, with better separation of the analogue and digital sections, was implemented in the revised version of the board, MB 4.3, which also uses CCD clock drivers based on transformers. The transformers are optimized to drive the double-metal CPC2 chips, which are now becoming available.

The next planned revisions of the motherboard include MB 4.4, which will have a further optimised layout and will use MOSFET drivers, and MB 5.0 for which the clock signals will be provided by the custom CPD1 chip.

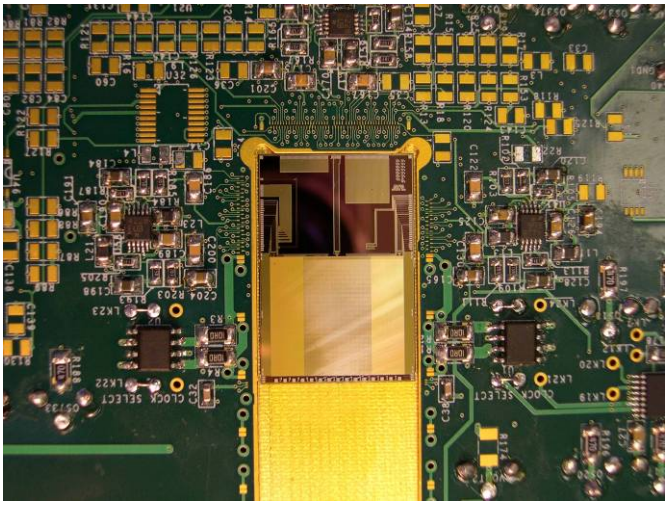

Figure 13 MB 4.2 with CPC2-10 and two MAX5057 clock drivers.

VME module design and construction

A further essential component of the test electronics is the BVM2. This provides a standard, VME based, means of communicating with and controlling all the test boards. A photograph of a BVM2 is shown in Figure 14. The BVM2 daughter cards take care of different specific applications (such as providing programmable delays etc.).

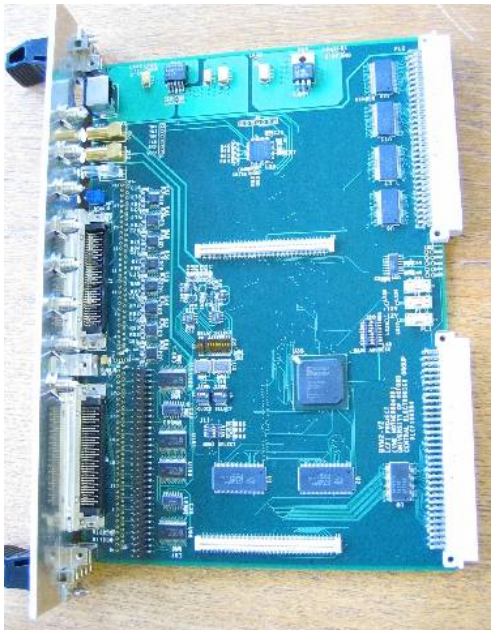

Figure 14 A BVM2 module

Ten BVM2 modules and several daughter cards have been produced and distributed since November 2005 to various LCFI test stands.

Test boards for CPD1 and other applications

The CPD1 custom driver chip will require a new test board, which must be able to exercise the full functionality of the chip. As the design of CPD1 has made good progress and the specifications of the chip are now defined, work has started on the CPD1 test board. The board will be designed jointly by the Bristol and Oxford groups and will use a custom capacitive load with minimal inductance. The power provision for the board should have enough fast decoupling capacitance to allow full exploration of the high current capabilities provided by the CPD1.

A further test board will be needed to provide analogue readout for the e2V test CCDs. The test CCDs will be used to study various possibilities for minimising the sensor capacitance and the clock voltage. This test board will be a simplified version of the full scale sensor motherboard.

Clock distribution

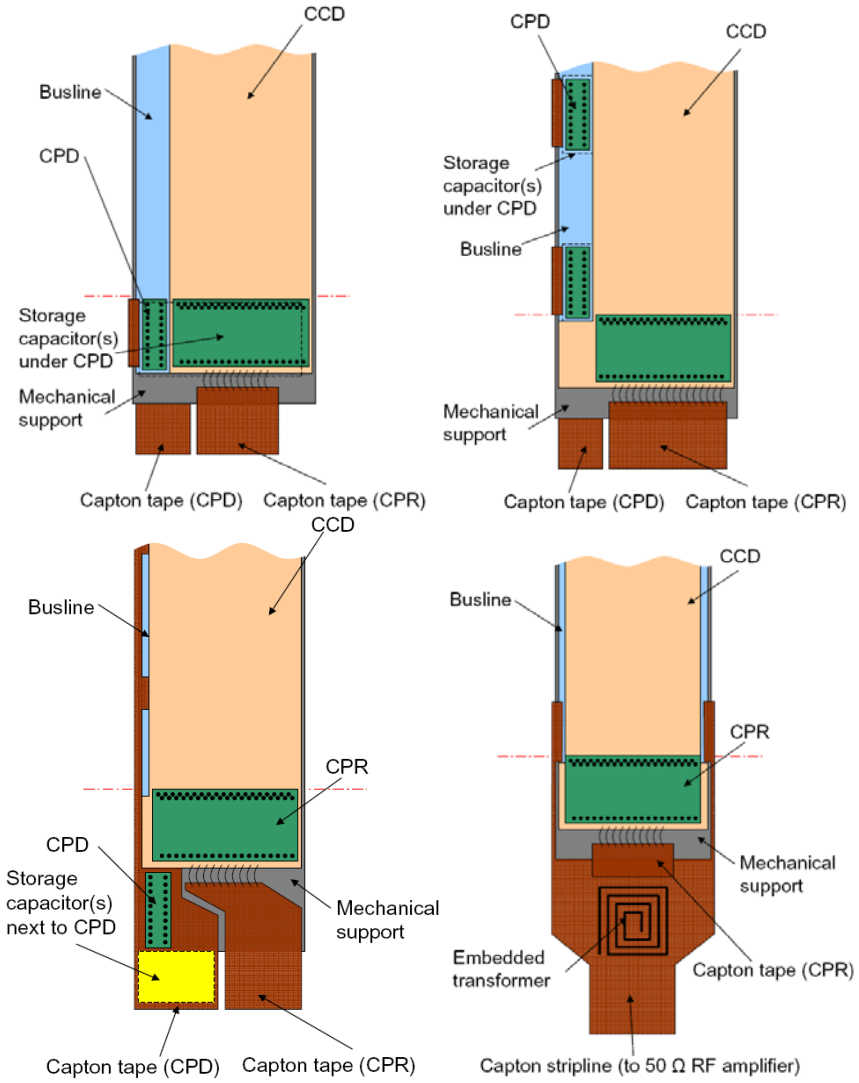

Figure 15 Possible schemes for clock distribution to the CPCCD sensor: single CPD on CPC (top left); multiple CPD on CPC (top right); single CPD outside CPC (bottom left); transformer clock drive (bottom right).

Several concepts, shown in Figure 15, are being considered for the distribution of clock signals to the CPCCD.

1. Single CPD on CPC (top left): one CPD is bump-bonded to the CCD and storage capacitors are located under the CPD chip. This solution provides the minimal material budget, but distribution of the clock to the CCD is most difficult. This solution will be possible only if the CCD capacitance is considerably reduced compared to that of CPC2.

2. Multiple CPD on CPC (top right): several CPDs are bump-bonded to the CCD, possibly with storage capacitors under the CPDs on the back of the ladder. Kapton flat cables provide connections for the CPDs. This approach solves the problem of clock distribution at the expense of extra material along the edge of the ladder.

3. Single CPD outside of CPC (bottom left): one CPD is bump-bonded to a Kapton flat cable outside the sensitive volume of the sensor and the clock is distributed to the CCD via this Kapton cable. Similarly to the previous option, clock distribution is relatively easy, but the amount of material is increased. In contrast to the previous option, the high frequency clock needs to be distributed, as opposed to the DC current for the driver chip. A possible variation of this option has the CPD chip and storage capacitor mounted on the back of the ladder under the CPR, with the clock still being distributed by a Kapton cable.

4. Transformer clock drive (bottom right): Kapton flat cable with embedded transformers connects RF amplifier to the CCD clock buslines. Similarly to the two previous, options the clock distribution is relatively easy, but the amount of material is increased. In this option, every ladder will need a power RF amplifier.

Options 3 and 4 use an “external busline” and make the CCD design easier. The options 2, 3 and 4 all can work with the present CPC2 design, while option 1 depends on reducing the CCD capacitance by a factor of about 2 or more. In all options, there is a low impedance connection between the CCD and a Kapton cable, either to provide the clock itself or the power for the CPD. Considerable effort is needed to determine the feasibility of each option and to investigate the delicate trade-off between mass, capacitance, ease of assembly and mechanical performance.

Future plans

The work in the forthcoming 6 months will concentrate on the design and construction of MB4.4 and MB5.0 and the new types of test board, in particular the CPD1 board. This latter board should be ready in September 2006 when the CPD1 chips become available. The delivery of the test board for the CPC-T devices is planned for December 2006, also in accordance with the CPC-T availability. Studies of clock distribution schemes will also continue.

WP5 – Integration and Testing

Introduction

The first two CPC2-10 devices were successfully tested at low clock frequency with satisfactory results. Some phenomena were observed which are as yet unexplained: these are under continuing investigation. The bump bonding of the CPC2, CPR1 and CPR2 chips is nearing completion and we are preparing the necessary software and firmware for the tests of these assemblies. The phase capacitance of CPC2 was measured using different techniques at RAL and Oxford in order to compare with the simulations done in WP2. This work will continue with measurements of the capacitance overlap structures, expected with the delivery of the busline-free CCDs.

The tests of the ISIS1 chips without p-well have finished and showed that the devices work as expected. It was found that the chips with p-well do not work due to incorrect doping caused by a manufacturing error.

Radiation damage studies have continued with the creation of a Synopsys-TCAD model of the CPC2 in preparation for the planned tests at Liverpool. At the moment, the model is being verified with data taken with very low clock amplitudes, which induces Charge Transfer Inefficiency (CTI). In addition, the simulations of the radiation damage data taken with CCD58 and the building of a reliable model are continuing.

The setting up of the new experimental test stations at Oxford and Liverpool is nearing completion. The bulk of the necessary hardware has been ordered and delivered with some work still required on the cryogenics.

CPC2 tests

Two CPC2-10 devices, one made on a 100 Ω.cm, 25 μm thick epitaxial layer and another on a 1.5 kΩ.cm, 50 μm layer were tested at RAL. These were single level metal chips not optimised to work at high speed, but representative of the design and the process at e2V. Figure 15 shows the performance of CPC2-10 at a 1 MHz column parallel drive frequency and with 5 V square-wave clock signals. These results were obtained using a commercial MOS gate driver chip, which is able to deliver 8 A peak current to each phase. The maximum clock frequency achievable with this configuration is estimated to be about 10 MHz for the single level metal devices and about 5 MHz for the busline-free variants. As described in the WP 4 report, a new CPC2 test board is being produced which will be equipped with transformers and which will allow tests at frequencies of up to 50 MHz.

As is shown in Figure 15, the tests revealed that the gain dispersion between 4 adjacent source followers is higher that that observed in CPC1. During the tests, it was also observed that the clock feed-through to the CPC2 outputs is unstable and vanishes periodically, together with the signal, for time intervals equal to several multiples of the clock period. This behaviour depends on the temperature, the CCD biases and the clocking scheme (e.g. on whether a reset is issued once per line or every pixel). So far, no definite conclusions about the origins of the observed

phenomena have been reached. One hypothesis is that the effects are due to PCB layout problems in the present test board, and many improvements have been implemented in the new design, now in fabrication.

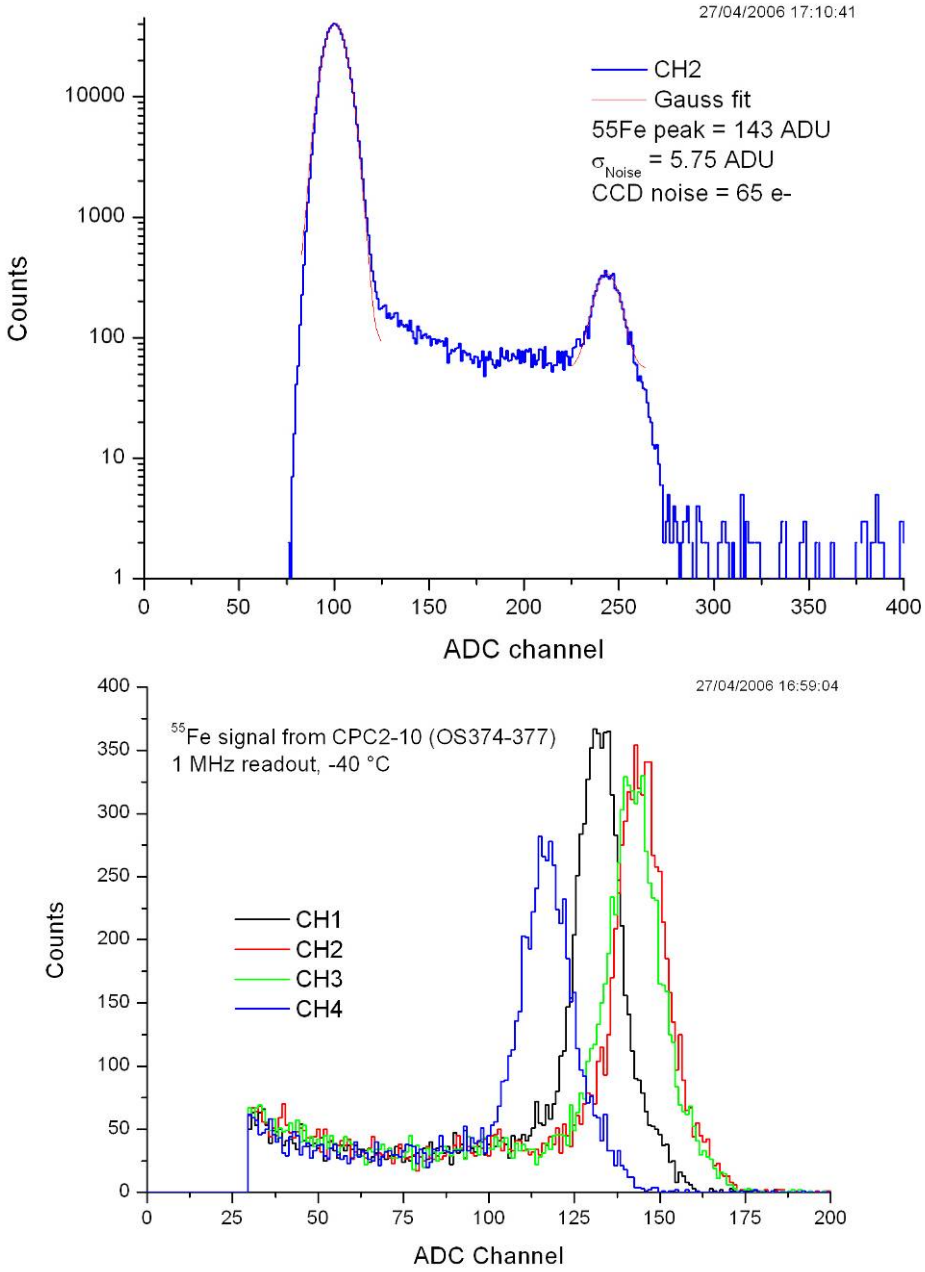

Figure 16 The first ^{55}Fe spectrum obtained from one output of CPC2-10 at 1 MHz clock frequency (top), and the gain dispersion from 4 adjacent outputs (bottom).

Bump bonding

Two CPC2 and two CPR2 wafers were sent to VTT in December 2005 for bump bonding. In order to eliminate possible common sources of error, some CCDs will be bump-bonded only to CPR1 or CPR2 chips, and some to both. Another two wafers will be sent after the successful completion of the first batch. VTT already have CPR1 chips supplied for earlier LCFI bump bonding work.

In March 2006, there was a small fire in the VTT clean room. The smoke from the fire caused some contamination and delayed the LCFI work for several months. Our chips and masks were been affected.

In the last month VTT have processed the chips with under-bump metallisation, bump deposition and protective resist. The wafers are now at the dicing stage, after which the actual bonding should not take long. It was discovered that the wafer-scale GDS file from e2V had not included the top termination field, containing 20 rows of pixels and the IG1, IG2 and IDR pads, as shown in Figure 16. VTT assumed that this area is empty and deposited dummy bumps in order to satisfy the current density conditions for the electroplating. As a result, the IDR pad is partially covered by bumps, which could complicate the wire bonding somewhat. In total, 5 out of 9 chips per wafer were affected by this error, which we do not believe to be critical.

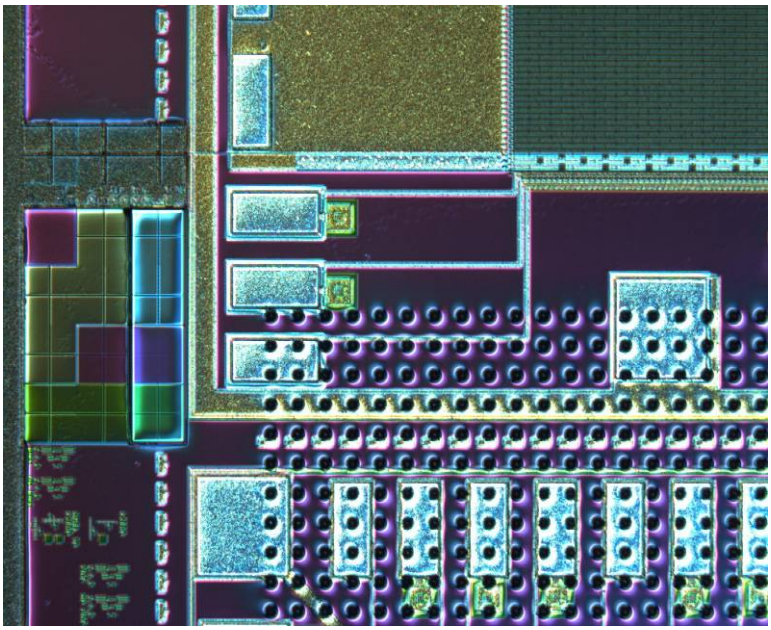

Figure 17 Photograph of a CPC2 with deposited bump bonds and protective resist. The chip is ready for dicing

Capacitance measurements

Accurate measurements of the phase capacitance of the e2V CCDs are important in order to tune the simulation tools used for the study of capacitance and its reduction in WP2. Measurements on CPC2 were carried out at RAL and Oxford using two different techniques.

The first method relies on measuring the time constants for the charge and discharge of the gate capacitance, clocked by square waves from a driver through a resistor. This measurement is done at RAL using a working CCD at low temperature, thus eliminating possible errors due to charge accumulation in the potential wells. This method measures the total gate capacitance, consisting of the inter-gate capacitance C_{ig} and the gate-to-substrate component C_s. The latter is easily calculable from the device parameters and is subtracted to get C_{ig}.

The second technique is more versatile because it uses a precision LCR meter working at many frequency points and able to supply different DC biases. It is

somewhat more difficult to get right, because the CCD may have some stored charge which distorts the measurement. This method is used in Oxford and has the power to determine separately C_{ig} and C_s and does not rely on indirect measurements.

So far both techniques have produced comparable results. The value of C_{ig} shown in Figure 8 was obtained with the first method. These measurements will continue to be used for the study of the dedicated capacitance overlap structures on the CPC2 wafers. These are polysilicon gates designed with an overlap varying from +3 μm to –2 μm (i.e. a gap). Due to the effects of the various processing steps, the real overlap is not a linear function of the design overlap, especially around a value of zero. These structures will provide us with calibration data for the simulation work in WP2. The delivery of these structures is expected to take place at the same time as that of the busline-free CCDs.

ISIS1 tests

The tests of the ISIS1 chip without p-well have been finished, as described in the last report to the POsC. Despite intensive efforts, and the help of e2V, it was not possible to turn on the output source followers of the sensors with a p-well. Parametric measurements on the tests transistors on the same wafer showed that the channel potential is very low (~ 5 V instead of the expected 10 V) and that the transistor thresholds are too high and strongly dependent on the source-to-substrate voltage. The cause of this was identified to be an incorrect doping profile for the p-well, which had received a dose 2.5 times larger than the design value.

After further investigation, it appeared that it may be possible to turn on the transistors by allowing the source to be biased very close to the substrate potential. To achieve this, a new version of the ISIS1 test board was designed to allow the load resistor of the output source followers to connect to negative voltage with respect to the substrate. Despite this, the transistors remained off and no signal was observed from the ISIS1 with p-well.

These measurements were done on chips from the first delivered p-well wafer, number 21. After discovering the error in the doping, it was decided that the chips from wafer 22 were not needed. Some parts of wafer 22 were therefore used for the SIMS study.

Radiation damage studies

After obtaining promising results with the simulation of the radiation damage effects in CCD58, work has continued with first steps towards modelling the effects in CPC2. Presently, the work on radiation damage studies consists mostly of simulations done by the Lancaster team with help from RAL.

The temperature range over which the CTI was measured in CCD58 did not allow good comparison with the simulations because sufficiently low temperature could not be achieved. At high temperatures, the dark current affects the measured CTI because of the long readout time, again making comparisons difficult. Despite all this, the models developed for CCD58 suggest that the radiation-induced CTI should be minimal in the temperature range from –50°C to 0°C, which must be verified

experimentally. If confirmed, this will be an important result with implications for the mechanical support and overall detector design.

Measurements on CPC2 should solve most of the problems encountered with CCD58. They will begin as soon as the Liverpool setup is fully functional. In preparation for the tests, a model of CPC1/CPC2 using realistic doping has been created and is now being verified. The data it is tested against was obtained using very low clock amplitudes to induce “artificial” CTI in CPC1.

Building of new test setups

During the last few months, the LCFI laboratories at Oxford, Liverpool and RAL have seen major upgrades in order to complete the tasks assigned to them. The Oxford setup has been charged with the precision capacitance measurements and in addition it will develop capabilities to test sensors with X-rays in a cryogenic environment. The goal for the Liverpool setup is to measure the radiation hardness of the CPC2 and CPC-T over a wide temperature range. All other measurements are being done at RAL.

The test setups at Oxford and Liverpool are now fully equipped with multi-channel power supplies, fast ADCs, pulse and signal generators, VME crates and oscilloscopes. Most of the equipment has been bought with LCFI funds, but a considerable part of it was procured using separate University funds. The test setup at RAL is fully operational and was recently updated with equipment bought with RAL infrastructure funds.

At the heart of the systems lie one or more BVM2 modules, which have been mass produced and are fully functional. Some work remains to be done for construction of cryogenic enclosures and temperature control. The power RF amplifiers for high speed clocking have been ordered and are expected shortly.

Future plans

The busline-free CPC2 will be evaluated by the end of 2006, shortly after work on the single-level metal chips has finished. Transformer drive and power RF amplifiers will be used to achieve the design speed of 50 MHz. Provided that new ISIS1 chips with p-wells are made, they will also be tested.

The first bump bonded CPC2 assemblies are expected in July. Their testing is significantly more complex than that of the stand-alone CCDs and is expected to end by December 2006. After that, another 2 wafers will be sent to VTT for bump-bonding.

The capacitance of the polysilicon overlap structures on the CPC2 wafers will be measured at Oxford by September 2006 and provide reference data to fine-tune the simulations being done in WP2.

The radiation damage studies on CPC2 are expected to start in September 2006 at Liverpool University and last for 6 months. After that, the work will continue with tests of some of the CPC-T devices, for which the radiation tolerance may be an issue.

After the manufacture and stand-alone evaluation of the first driver chip, CPD1, we expect to test the busline-free CPC2 devices in combination with the driver in March 2006.

WP6 – Mechanical Studies

Introduction

The goal of WP6 is to develop the mechanical technology to build an ultra-low mass silicon vertex detector for the ILC. The mid-term priority is to establish how to stably support the silicon sensors using a total material budget of only 0.1% X_0 per layer within the active volume, but also includes work on other, global aspects such as cooling and thermal issues and ladder mounting schemes.

For the first part of 2006, WP6 work was focused on thin-ladder mechanics (including silicon stress studies) and cooling studies.

Thin ladder mechanical prototypes

Several prototype foam-based ladders were built and investigated towards the end of 2005. These were of two types (Figure 18): a single piece of about 25 μm thick blank silicon wafer attached to a 1.5 mm silicon-carbide foam (Figure 19); and a sandwich of 1.5 mm reticulated vitreous carbon between two thin silicon wafers.

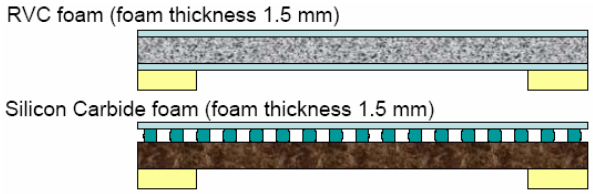

Figure 18 Foam ladder designs.

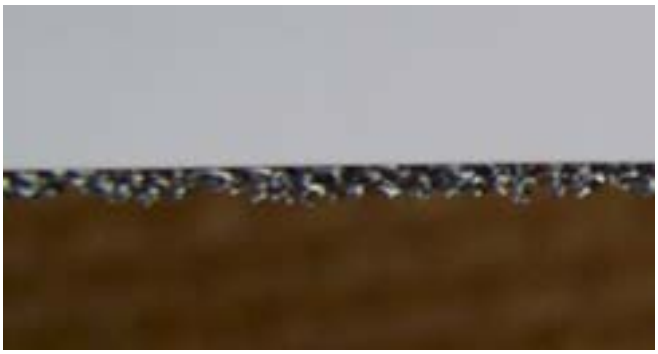

Figure 19 Photograph of silicon/silicon-carbide ladder.

The silicon-carbide prototypes showed promise but were over the required material budget (thickness 0.14% X_0) as we were only able to obtain foam of 8% relative density. Detailed studies of these models revealed the need for a better understanding of the laser survey system, as is described below.

Although the RVC models were within material budget (thickness 0.9% X_0 using 4% relative density RVC), attaching the silicon to the foam proved problematic. A good, rigid bond is required to gain the full benefit of the sandwich structure, but achieving this with a continuous thin layer of adhesive has not proved to be possible. The next

model will be built using the new glue robot (see below) to put down an array of glue pads on the silicon before the foam is brought into contact with this.

No new ladders were built during the January-June 2006 period as there was no thinned silicon available. When funds became available (April 2006) several manufacturers were contacted for quotations to supply this. For financial reasons, it was decided not to use our previous suppliers (Virginia Semiconductor, who out-sourced the actual thinning) but to send some wafers to Aptek Industries for processing to 25 μm. We have recently (June 2006) heard that no wafers survived this process.

Glue robot

Funds from the CCLRC Well Found Laboratory initiative were used to purchase an EFD TT 525 glue dispensing robot, capable of automatic precision adhesive application over a working area of 525×525 mm^2. This has been installed and commissioned in RAL Lab 9 and will be used in subsequent ladder prototyping.

Laser survey system

Although the presence of systematic effects in the laser displacement meter (Keyence LK-031) used for mechanical surveys was known, these effects have only recently begun to limit the studies of ladder prototypes. A detailed study was initiated with the intention of producing a reliable calibration procedure. This study showed not only the expected periodic ± 5 μm "ripple", but also much larger (~ 100 μm) deviations that depended strongly on the angle of the laser head with respect to the measured surface (Figure 20)

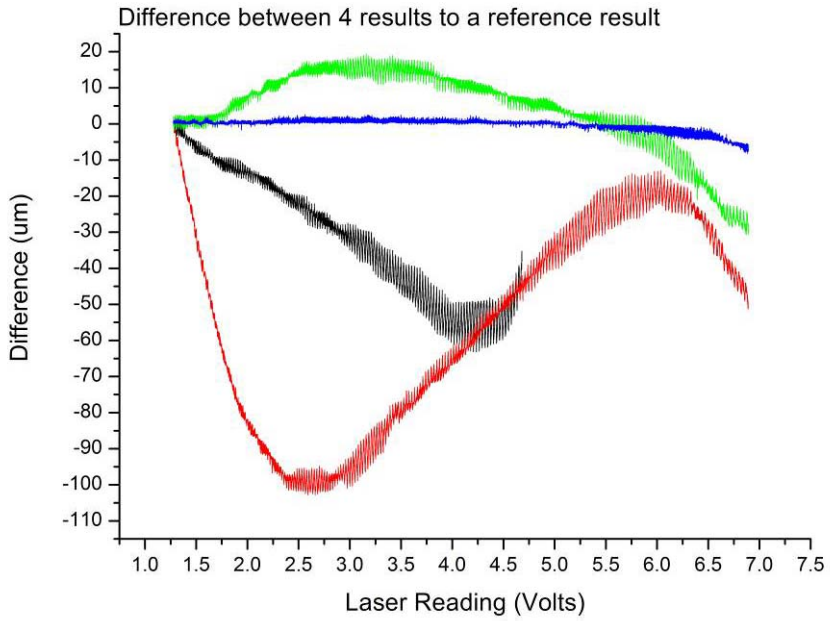

.

Figure 20 LK-031 laser calibrations at various angles.

A new laser displacement meter, the LT-9030M, is now available from Keyence. Not only does this have a much higher precision and linearity specification than the LK-

031, but it also uses confocal optics that should have minimal dependence on the angle of the head. LCFI borrowed an LT-9030M for evaluation and it proved to have negligible angular dependence and small non-linearity. Only a small (< 1 μm) ripple effect was observed (Figure 21), so it was decided to upgrade to the LT. This has now been purchased and installed in Lab 9. The integration of this new equipment into the survey system software is almost finished.

Silicon stress studies

The mechanical properties of thinned, processed silicon sensors were studied by Glenn Christian[9], who discovered that even wafers that had been etched for stress relief showed bowing tendencies; 15 mm wide CCDs thinned to 13 μm had a sagitta of almost 1 mm. A first study was performed using algebraic calculations and FEA to gain an understanding of what may be driving this effect, and to estimate the force required to overcome it.

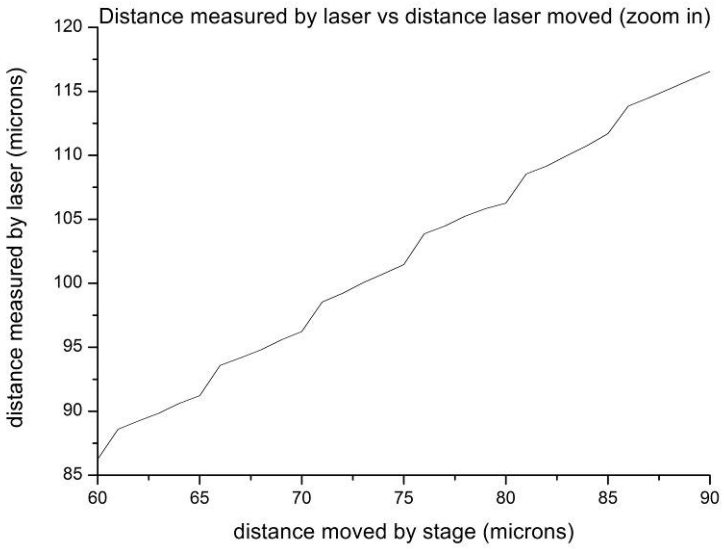

Figure 21 Non-linearity of LT-3010 laser.

The sensor was modelled as a layer of silicon with a uniform layer of a different material on one side. A temperature change was then applied to introduce differential stress. It was found that Glenn's results could be approximately reproduced by using a 1 μm layer of aluminium and heating the composite structure by about 50°C. Scaling this model up to the size of an ILC sensor (25 mm × 250 mm × 25 μm), it was shown than a total force, evenly distributed, of approximately 0.05 N is required to keep the sensor flat.

The model was then used to gain an insight into which steps in the processing could be responsible for the residual stress. Table 2 shows the effective curl produced by a temperature change of about 50°C for layers of various different materials with thicknesses representative of CCD process parameters. Glenn's results correspond to a curl (sagitta) of about −500 μm, so the prime candidate for causing this curl would appear to be silicon dioxide, which is grown at temperatures in excess of 1000°C.

Material	Thickness (μm)	Curl (μm)
Aluminium	1	531
Polyamide	1	78
Silicon Dioxide	1	-49
Silicon Nitride	0.085	3
Polysilicon	0.3	-2

Table 2 Curl caused by different materials attached to silicon, ΔT = 50°C.

Cooling test rig

A test stand consisting of a full-scale model of one quarter of the ILC vertex detector, cooled by nitrogen gas and using thin-film resistors to simulate heat sources, was built and commissioned in RAL Lab 9 in the second half of 2005. During the January-June period a number of measurements were made using this system.

Total power measurements were made at a variety of temperatures and flow rates. Figure 22 shows the total power (in Watts) extracted from the model by nitrogen at various flow rates as a function of the temperature difference between the mechanical structure (averaged over several locations) and the gas inlet. The largest flow rate is comparable to that used for VXD3, the SLD CCD-based vertex detector, and it is clear that even at these relatively low flow rates, 20-30 W can be extracted from the structure.

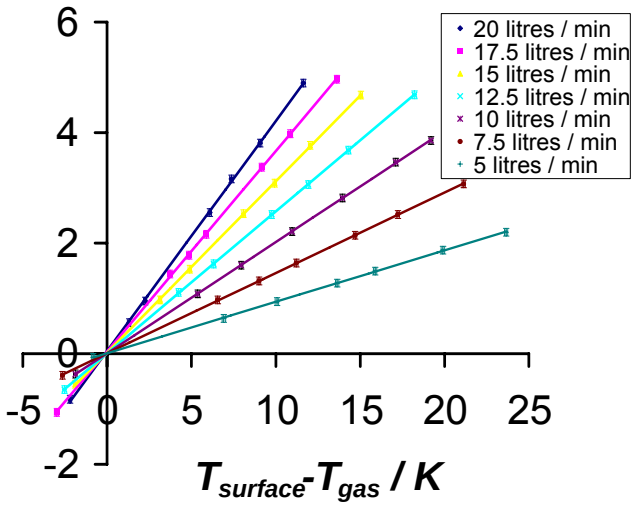

Figure 22 Power extracted (in watts) from vertex detector model by gas as a function of flow rate and temperature difference.

The internal temperature distributions for various configurations of gas inlets and outlets have been tested, with some results for two different flow rates illustrated in Figure 23. It should be noted that the contours represent interpolations of measured ladder temperatures, not actual gas temperatures. Perhaps counter-intuitively, changing the inlet and outlet configuration appears to make little difference, with poorer ladder temperature uniformity resulting for large numbers of outlets.

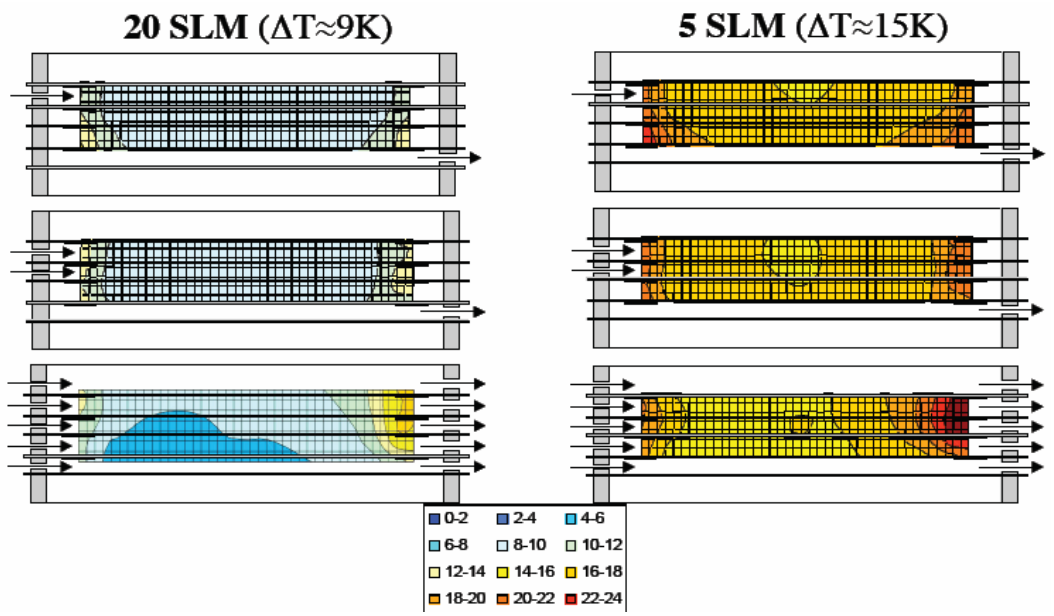

Figure 23 Measured temperature distributions in the thermal vertex detector model for various inlet and outlet configurations.

The current detector model is made from aluminium and steel, so conduction through the barrel and endplates could be affecting the distributions. Since it is not known how thermally conductive the real detector will be, the model is being rebuilt using Perspex components to allow any such effect to be evaluated.

Computational Fluid Dynamics studies

In addition to the measurements described above, a parallel programme of thermal simulations is being performed. Careful tuning of the simulations to ensure they describe the measurements will allow detailed understanding of the results, simplify evaluation of alternate configurations and allow reliable extrapolation to the full vertex detector.

Figure 24 shows the gas temperature distribution in a finite-element calculation of the test rig. This agrees qualitatively with the physical results, and work is ongoing to improve the quantitative match.

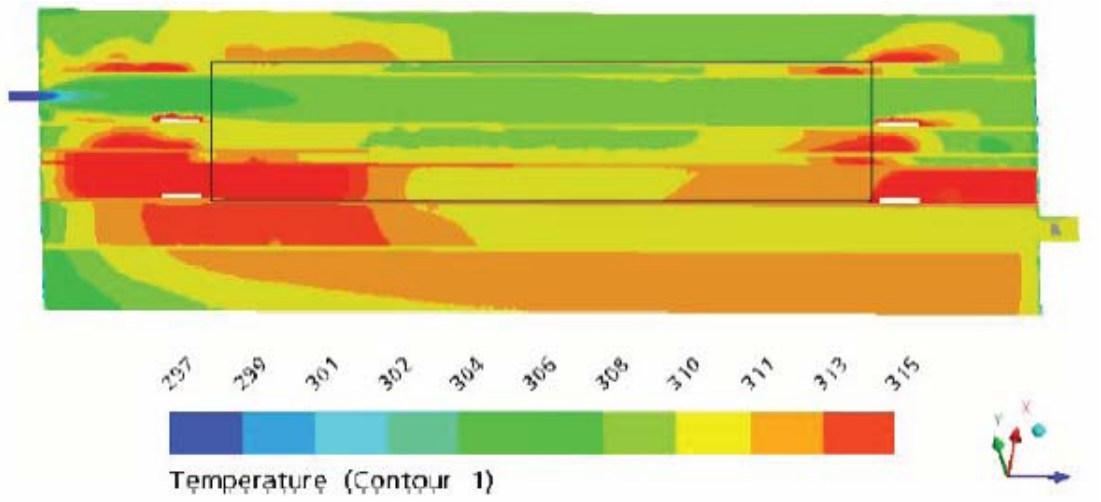

Figure 24 CFD results for thermal model for a gas flow rate of 20 litres/minute.

Other studies

Work has been progressing steadily towards building a cryogenic chamber at Bristol. This will be used for both mechanical and electrical testing, and most of the components and control systems are now in place.

As mentioned in the Introduction, collaboration between LCFI and other ILC vertex groups on mechanical issues is becoming closer, including frequent contact with the Fermilab group working on the SiD tracker and vertex detector design. There have also been recent discussions with the European MAPS and DEPFET groups about the possibility of some level of collaboration. This would involve LCFI receiving thinned mechanical sensors from these other groups for study.

Future Plans

Despite the developing collaborations with the DEPFET and Strasbourg groups, it remains important that we find a reliable and affordable supplier for thinned silicon, allowing large scale ladder prototyping to resume. Once this is resolved, further RVC and silicon carbide structures will be built using the new robot. These will be surveyed using the new laser, completing the basic evaluation of the two technologies.

Other work on support technologies will also continue. The Bristol cryogenic stand should be ready for novel material testing by the end of the summer, and engineering effort at Liverpool will become available within the next few months to look at specific aspects of using carbon fibre.

The simple silicon stress model will be checked against real processed, thinned sensors, obtained both from other vertex groups and from e2v. Some of these samples will also be built into ladders.

The current round of cooling studies should be finished within the next few months. This will involve testing the physical model with insulating endplates and ladders and the tuning of the CFD simulations to match. Once there is confidence in the results, the CFD calculations will be a powerful tool to continue with studies of the full detector and new configurations.

References

[1] http://lcio.desy.de/.
[2] http://berggren.home.cern.ch/berggren/sgv.html.
[3] http://ilcsoft.desy.de/marlin/.
[4] N. deGroot, “From Pixels to Physics”, contribution to the 10th International Workshop on Vertex Detectors, VERTEX 2001, Brunnen, Switzerland.
[5] R. Hawkings, “Vertex detector and flavour tagging studies for the TESLA linear collider”, LC-PHSM-2000-021.
[6] http://forum.linearcollider.org.
[7] D. Jackson, NIM **A 388** (1997) 247.
[8] Alexei Raspereza, “DEPFET Pixel sensors simulation in Marlin”, contribution to ILC VTX workshop at Ringberg Castle, May 2006.
[9] CHRISTIAN, Glenn B, “Aspects of Thin Ladder Mechanical Design and Development for a Linear Collider Vertex Detector”, D Phil, Oxford, 2002.