PROCEEDINGS OF SCIENCE

The silicon pixel tracker – beginning of a revolution?

CJS Damerell[1]
Rutherford Appleton Laboratory
Chilton, Didcot, OX110QX, UK
E-mail: c.damerell@rl.ac.uk

KD Stefanov
Sentec Ltd
5 The Westbrook Centre, Milton Rd, Cambridge CB4 1YG, UK
E-mail: kstefanov@sentec.co.uk

Despite the 'hermetic' nature of modern detector systems at the energy frontier, it is well-known that tracking has always degraded badly in the forward region of these systems. In an attempt to overcome this weakness, consideration has been given to a new approach – the Silicon Pixel Tracker (SPT). This exploits the fact that, while one clearly needs to know the precise time associated with each track (so that it can be unambiguously associated with the correct bunch crossing) one does not need this information for every point on the track. Thus one can conceive of a silicon tracker which has extremely thin time-integrating layers over most of the tracking volume, where the term 'time-integrating' depends on the application. This approach could reduce the material in tracking systems for ILC/CLIC and HL-LHC by a factor ten or more, compared to systems with silicon microstrips which, due to their relatively coarse granularity, require fast timing on every strip.

The 20th Anniversary International Workshop on Vertex Detectors – Vertex 2011
Rust, Lake Neusiedl, Austria (June 19 – 24, 2011)

[1] Speaker

1.Introduction

The ILC Detector R&D Panel review of tracking in 2007 [1] raised concerns about the likely material budget associated with the tracking systems under consideration. While the R&D groups optimistically predicted a total material budget of ~0.1 X_0 at all polar angles, which would satisfy the physics requirements, the panel was doubtful that this could be achieved, particularly in the forward region down to the minimal polar angle for tracking of ~7^o.

These concerns inspired efforts to find a new approach, and a year later the Silicon Pixel Tracker (SPT) concept was presented at the ILC workshop at Tohoku U in Sendai [2]. Around this time, the UK formally withdrew from ILC but the work has continued at a lower level, with welcome support from the SiLC collaboration [3], and has been extended to a possible application for LHC general purpose detectors (ATLAS and CMS), where current tracking systems also suffer from excessive material and hence degraded performance in the forward region [4]. This situation is likely to deteriorate further with the extension to HL-LHC over the next decade, due to the increased number of pileup interactions in each bunch crossing.

The SPT concept is based on two instances of what we call 'separated function' design, by analogy with the invention of strong focusing accelerator systems in the 1950s. Firstly, we acknowledge the need for single-bunch timing on every track, but do not insist on this for every point on the track. So we suggest an appropriate mixture of *timing layers* which provide points with precise single-bunch timing, and *tracking layers* which allow some degree of time integration over a number of bunches.

The second 'separated function' feature is in the pixel design. In order to achieve the large area coverage required for tracking systems, we propose the simplest approach (monolithic CMOS pixels), with pixel size ~50 μm diameter, matched to the required momentum measurement. Despite these relatively large pixels, we suggest very compact sense-circuits with extremely small front-end (FE) transistors. We are able to evade the capacitance-matching theorem by using 4T pixels, also known as charge-coupled pixels, in which the signal charge is collected into a buried channel then transferred by the inbuilt electric field to the collection node of the FE transistor. These important points are discussed in Section 4.

The power dissipation in the tracking layers can be reduced to the level at which liquid or evaporative cooling is no longer needed – one requires only a gentle flow of coolant gas, either nitrogen or air. By eliminating cooling pipes and liquid coolant, the material in the tracking layers can be greatly reduced. This results in a 'virtuous circle'; as less detector material needs to supported, the support structure itself can be made more lightweight. This approach was used to good effect in the SLD vertex detector [5], resulting in a layer thickness of 0.4% X_0, an achievement that remains unmatched after 30 years.

This paper is divided into the following sections: Design overview, Mechanical design, Pixel design, Possible use at LHC, and Next steps.

2.Design overview

We first consider the case of ILC/CLIC, since this provides the simplest application for this approach. The bunch trains determine the 'natural' integration time for the tracking layers. Thus for ILC we have trains of ~3000 bunches spaced ~300 ns apart, hence a train duration of ~1 ms, and a train frequency of ~5 Hz. For CLIC, we have trains of ~300 bunches spaced 0.5 ns apart, hence of duration ~150 ns, and a train frequency of ~50 Hz.

In both cases, the baseline already includes a pixel-based vertex detector of 3-5 layers (barrel plus endcaps) which delivers hits with single-bunch time resolution (or close to it, ~ few ns in the case of CLIC). The vertex detector comprises the first system of timing layers envisaged in this concept. The second timing layer system comprises an external envelope also of ~3 or 4 layers (barrel plus endcaps) in the form of an external shell or envelope within ~10 cm of the inner wall of the Electromagnetic Calorimeter (ECAL). Since the timing layer systems deliver a precise time measurement with every hit, their pixels must be equipped with continuously active FE circuits, which implies active cooling (presumably evaporative CO_2 unless some better alternative is found within the next few years). In the case of the vertex detector, the implied material budget is undesirable but inevitable. In the case of the outer timing system, the physics penalties are much less serious. This shell can be viewed as a sort of pre-shower detector, since the products of photon conversions and hadron interactions proceed immediately into the calorimeter system close to their production points.

Between these two timing layer systems, we have a number (~5) of tracking layers, as shown in Fig. 1, each of which is populated with pixels which deliver time integration through the bunch train (1 ms and 150 ns for ILC and CLIC respectively). During the bunch train, the in-pixel active circuitry is switched off, and the signal charges are simply drifted to the input nodes by the inbuilt electric fields. During the inter-train period (200 ms and 5 ms for ILC and CLIC respectively) the pixels of each device are read out row by row. Those which have been hit generate a binary signal, so by the end of this period, all hit addresses have been transmitted to the event builder. Since these data comprise the total accumulated hits from their bunch train, the hit density is relatively high, but the occupancy is still tiny due to the high granularity of the pixel system (there are ~30 Gpixels in the overall tracking layer system).

Offline, the track reconstruction proceeds as follows. The first step is to find the on-time track stubs or mini-vectors in the two sets of timing layer systems (see Fig 2). Standalone track-finding in pixel systems is straightforward and can be made extremely efficient, due to the advantage that such detectors provide unambiguous space-points in contrast to the extended regions associated with strip detectors. The next step is to link the inner and outer track stubs, which again can be made clean and efficient given that one is working in 3-D, and that the linked tracks must line up precisely with the stubs in the RZ view. The final step is to 'drill inwards' from the outer stub, looking for matching hits in the tracking layers. One needs to allow a search region of say 3-sigma, where the search ellipse is determined by the multiple scattering for that particular track. For some steps (for example when extrapolating the track from layer 4 to layer 3) one may find ambiguous hits in the search ellipse. Usually, when these candidates are propagated to the next layer (layer 2 in this example) the ambiguity will be

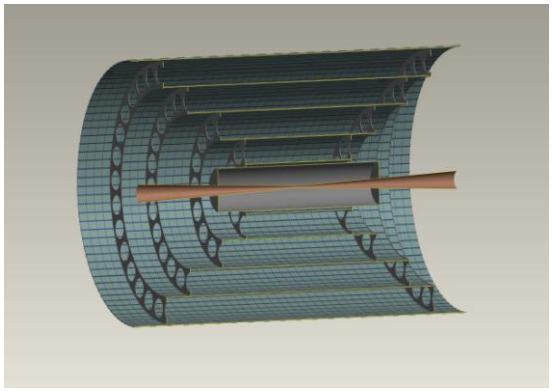

Fig. 1 Arrangement of tracking layers (5 barrels, 5 pairs of endcaps of which only one is shown). Overall length ~ 3.5 m

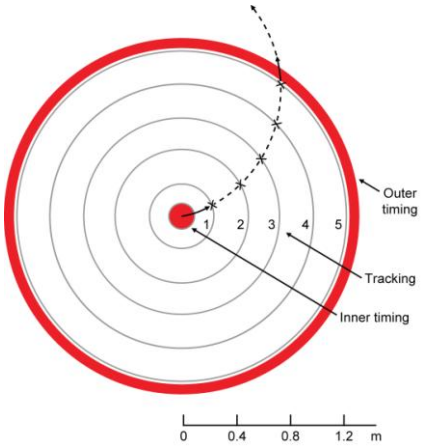

Fig. 2 rϕ view of barrel system, with a single track for which are indicated the track stubs in the timing layers, and the hits in the tracking layers

resolved. If for the fully-defined track, there remain ambiguous points on one or more layers, the best approach is to omit them from the track fit, at little cost to the fit quality.

It is not possible in this short paper to describe the special procedures to handle long-lived *B* decays, photon conversions and hadronic interactions in the material of the vertex detector and in the tracking layers. However, back-of-envelope calculations [6] indicate that this track-finding and fitting procedure can be made close to 100% efficient. Of course, the only way to establish the performance precisely will be by a full simulation of the suggested procedure. It might then turn out, for example in the forward region, that the link between the inner and outer track stubs is too tenuous to be bridged without further help. In that case, one could consider adding tracking layers (this would be the first choice since they are so thin), and as second choice adding one or more intermediate timing layers. This would be disfavoured due to the addition of material where it is least wanted. It will be interesting to explore what is the optimal balance between minimal material and adequate tracking efficiency over the full angular and momentum range.

3.Mechanical design

The outer shell of timing layers can be constructed with ‘conventional’ technology, similar to large microstrip detector sytems. Each layer with its cooling system needs to be quite heavy, so the support structure cannot be made extremely lightweight. The main challenge (and the opportunity for a new approach) comes with the tracking layers, where reduction of material is at a premium. From work done in the LCFI collaboration and beyond [7], we suggest that the mechanical supports be constructed from silicon carbide foam, which can be produced with as little as 5% of the solid density. Adhesive-bonded to this on the inner surface (see Fig. 3) are the silicon sensors. Each sensor is of size approximately 8x8 cm^2 and they are arranged side by side along the ladder, with a gap in active area of ~200 μm between neighbours. The

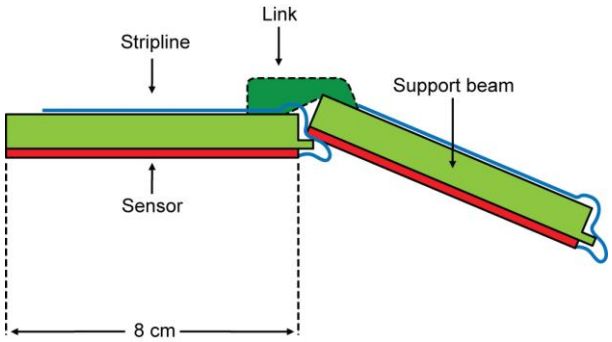

Fig.3 End view of two ladders in a barrel, linked at ~50 cm intervals along their length. Full azimuthal coverage by overlap between sensors.

orientation of the devices is with the columns running across the ladder, in the azimuthal direction, and the electrical connection and readout are at the edge adjacent to the bottom of the columns (details in Section 4). Wire bonds connect from the device to a copper-kapton tab which is folded to pass between adjacent ladders. On the outer surface of the ladder, these tabs connect to a wide single-layer stripline which runs the length of the ladder, to electrical connections at the end of the barrel. A similar structure is used for the endcap detectors.

The ladders of each barrel are arranged in a spiral geometry. We assume a thickness of SiC of ~10 mm for the support structure. This material has the important advantage of extremely homogeneous thermal behaviour, and a very small expansion coefficient which closely matches silicon. For this reason, the sensors, assumed to be thinned to ~50 μm, can be adhesive-bonded to the substrates with silicone elastomer, providing a stress-free joint regardless of the operationg temperature, hence avoiding the requirement of precise temperature equality within the tracker volume.

Given the very low mass of the sensors and their services, the main mechanical load on the support structure is the gravitational force due to its own mass. If it were made as a complete foam cylinder, the grvitational sag would be negligible, ~10 μm even for the largest barrel in the system. However, if one simply slices such a cylinder into ladders and adds the detector-related weight, one finds prodigous sag (~20 mm) for the upper and lower ladders in the system. This can be almost entirely avoided by adding small foam links between ladders at intervals of about 40 cm along their length, as indicated in Fig 3, which reduces the sag in the worst case ladder to ~ 20 μm, almost the same as for a complete cylinder.

It is suggested to assemble the detector in half-shells. These would be assembled sequentially onto the beampipe/vertex detector assembly, working outwards from layer 1 (barrel plus two endcaps), using adhesive bonds to connect together each pair of half-shells. This approach has at first sight the disadvantage of difficulty of servicing. However, one has to consider the great advantage of truly minimal material. The SLD vertex detector (307 Mpixels) never required any servicing, and this is also true of much larger space-based systems of monolithic pixels currently in service. By minimising electrical connections, and being careful with assembly procedures and associated quality control, there is every reason to expect such a system to perform reliably. If necessary, one could cut through the adhesive bonds for

disassembly if required, and the work of doing so would be far less than the necessary preliminary of moving the detector into the parking area and extracting the tracker from the heart of the detector.

The aim of the pixel design is to achieve a total power dissipation in the tracking layer system of at most 1 kW, which is about 20 times greater than for the very compact SLD vertex detector, which pioneered gas cooling of silicon tracking systems. Given that the tracker volume is ~500 times greater than the SLD cryostat, it is clear that with a set of foam-insulated carbon fibre composite pipes of very low mass, one can easily deliver and extract the gas needed to cool such a system (~20 g per s) with no risk of significant vibrations, pressure excursions or other undesirable features.

The overall material budget for such a system will be ~0.6% X_0 per tracking layer. One needs to allow some extra for end-services (LVDS cables, a few EO converters), and there is the inevitable obliquity factor, but it is clear that a 5-layer system made in this way can present a material budget considerably below 10% X_0, maybe only about half that.

4.Pixel design

R&D for particle physics pixel detectors is an active topic – with concepts of varying degrees of complexity ranging from the simplest (monolithic devices such as CCDs and CMOS pixels) to hybrid pixels and vertically integrated designs of considerable structural complexity. For a large-scale pixel tracker to be a serious contender for use in experiments within the next 10 years, it is important to work with a simple construction concept, and fortunately the design requirements are in our case best met by a monolithic pixel architecture. The most challenging aspect is to design the tracking layer pixels, given the need for at most 1 kW from ~30 Gpixels, about 33 nW per pixel. For the outer timing layers, a similar architecture is envisaged, but much faster sampling is needed to provide the necessary timing precision. This is in fact less challenging given the acceptability of evaporative cooling in the shell of timing layers close to the ECAL.

In view of the space constraints in this short paper, we discuss only the design of the tracking pixels, for which the basic structure is shown in Fig. 4. The main characteristics are often ambiguously referred to as 4T (for 4-transistor) but are best described as charge-coupled CMOS pixels, the term used by Jim Janesick, who pioneered their use for scientific imaging [8]. In our case, the essential requirements are:

- an active thickness of ~30 μm, sufficient to provide a substantial min-I signal (on average 2400 e^-)
- a depleted structure which achieves charge collection by drift. This involves depletion of the epi layer by reverse bias with respect to an n-channel over most of the pixel area, followed by patterning of this n-channel to establish a potential gradient towards the pixel centre, followed by a circular transfer gate of variable width which drifts the charge to the region of the output node, and the node itself which comprises a reverse-biased diode plus polysilicon contact which includes the gate of the FE transistor. The thick arrows in Fig 4 depict the sequential drift

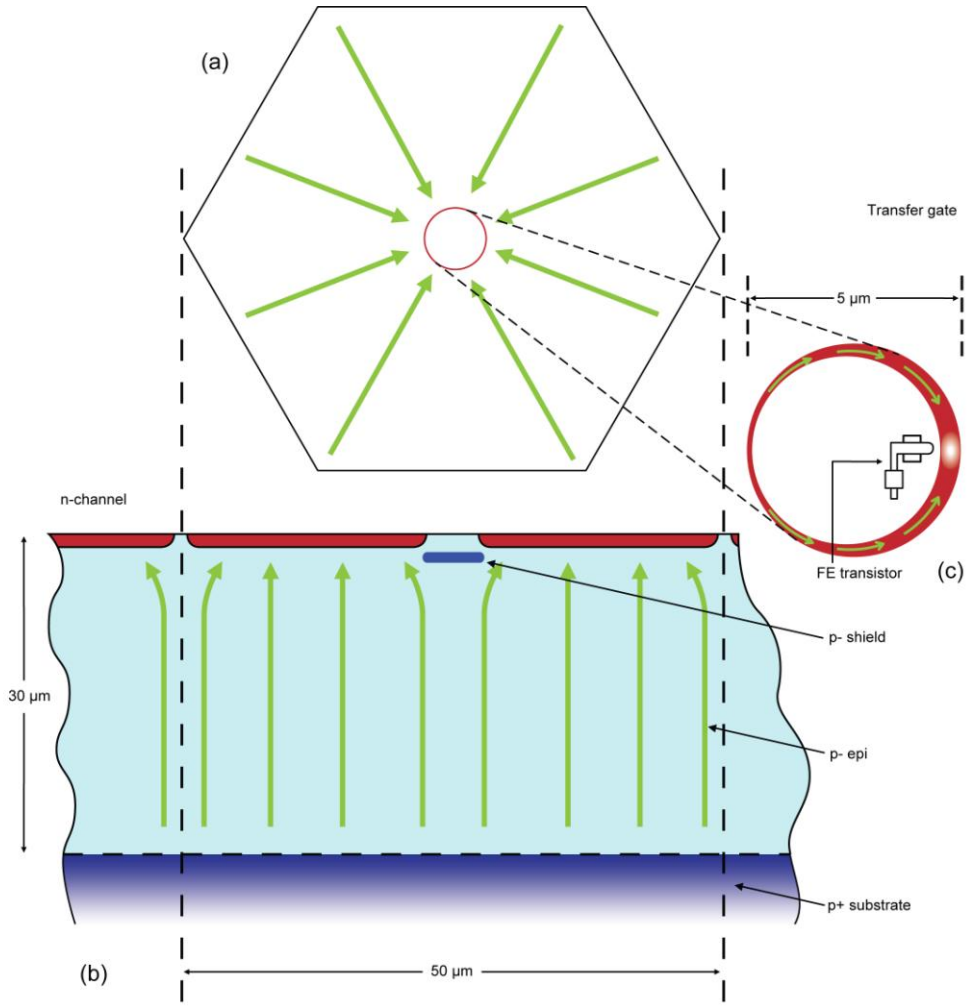

Fig.4 Pixel layout. (a) shows plan view, (b) the cross-section,and (c) an expanded view of the central region which incorporates all in-pixel CMOS electronics, of which only the front-end and reset transistors are sketched.

of signal charge in the bulk (b), in the buried channel (a) and beneath the transfer gate (c), for transfer to the gate of the FE transistor.

- The FE transistor and additional in-pixel CMOS logic (not shown in Fig 4 (c), but seen in Fig 5) is protected from parasitic collection of signal charge by a depleted p^+ implant, the p-shield indicated in Fig 4 (b).

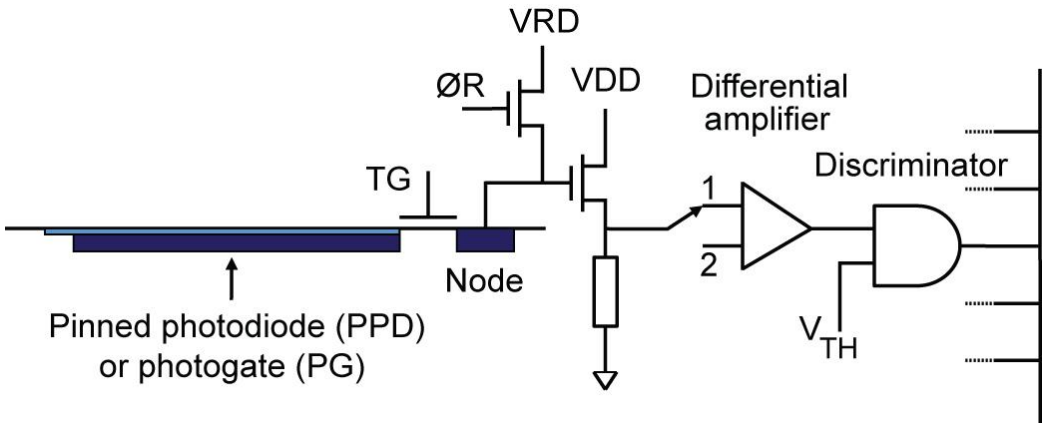

Fig.5 Pixel logic.

This structure achieves fast, efficient collection of signal charge from all areas of the large pixel (typically 90% charge collection within 5 ns), and concentrates it on the very small FE input gate. Without this feature, one would be limited by the capacitance-matching theorem [9], and with such large pixels, the noise performance would be seriously degraded. These charge-steering structures have been developed over many years of CCD design, with recent progress by Goji Etoh [10] and the e2V imaging team [11].

The main structure can be formed either as a photogate (PG, with thick field oxide and biased polysilicon above the n-channel) or as a pinned photodiode (PPD). The second option is probably preferred, since it avoids the need for dual gate thickness in the device processing.

In either case, the in-pixel circuitry will use very thin dielectric associated with a small-scale deep submicron process. Minimal power will be achievable with smaller feature size. This is considered quantitatively for the most challenging application (an SPT for HL-LHC) which is discussed in Section 5.

The in-pixel functionality for ILC/CLIC works as follows. The outer timing pixels, like the vertex detector, are read out for each bunch crossing in the case of ILC, and every 5-10 bunch crossings for CLIC. As regards the tracking pixels, these are are biased but static through the bunch train, with VDD switched off, hence zero pwer dissipation. After the train, the devices are read out slowly, through the entire inter-train period of 20 ms or 200 ms for CLIC and ILC respectively. This readout takes advantage of the full noise suppression capability of correlated double sampling (CDS). The devices (8x8 cm^2, 1600x1600 pixels) are read by a rolling-shutter variant, in which each row is powered in turn. The signal level from the FE source follower is sampled and held on one input of a differential amplifier (Fig 5), where the other input holds the voltage from the previous train on a storage capacitor. The difference output is compared with a reference voltage, and if it exceeds the threshold V_{TH}, the discriminator output flags the hit on edge logic at the periphery of the device. Occasionally, the node potentials are reset to standard values (V_{RS}) to avoid the input being pushed beyond the linear range by an accumulation of hits.

This simple approach is capable of refinement if tests show the need of (for example) more closely coupled CDS. Signal charges can be held under the transfer gate, allowing the

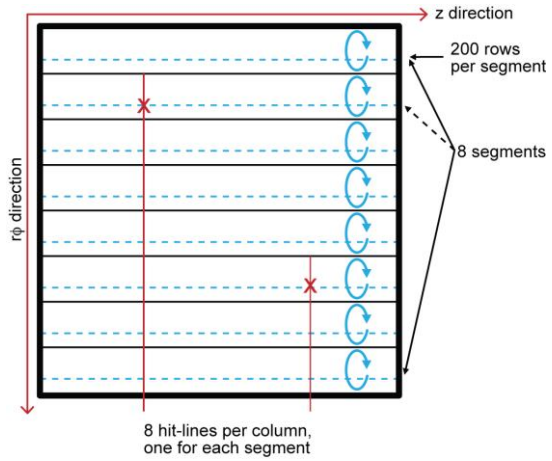

Fig.6 Segmented rolling shutter readout. To avoid ambiguity, the column line (Fig 5) is duplicated 8 times, one for each segment.

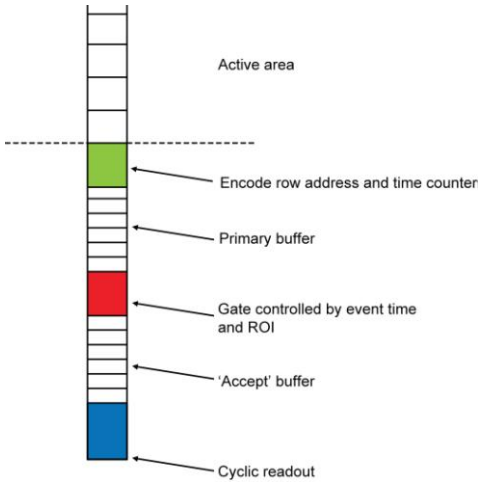

Fig.7. Column edge logic.

rolling shutter to proceed for each row in a 2-step sequence, firstly to sample the pre-transfer level, and immediately afterwards the level with the latest signal charge added.

5.Possible use at LHC

In LHC, there is no interruption to the stream of bunch crossings, so the concept of readout during inter-train periods does not apply. Nevertheless, one can still consider the 'separated function' tracker concept. The timing layer pixels, like those already used in the ATLAS or CMS vertex detectors, are read at the bunch-crossing interval of 25 ns. For the outer timing layers, the requirements in terms of radiation resistance are far less stringent, so the architecture of monolithic pixels is likely to work well, greatly simplifying the construction of these layers compared to the inner timing layers (vertex detector), where hybrid pixels will probably still be obligatory.

For the tracking layers, if we are to retain the advantages of low power, hence gas cooling, hence much reduced material, we need some amount of time integration just as at ILC/CLIC. As an example, we have considered the case where this integration time is set to 10 μs (400 BX). The first question is whether the hit density accumulated during this time is or is not acceptable, in terms of clean track reconstruction. We have of course to consider the hit rates for the high luminosity upgrade HL-LHC, since this is the timescale on which it is intended to replace the current ATLAS and CMS tracking systems. Using figures supplied by ATLAS colleagues [12] we find firstly that the occupancies in the tracking layers are entirely acceptable, thanks to the fine granularity of the pixel devices (1600 times finer than with microstrips). So at first sight, one might imagine implementing a segmented rolling-shutter readout (see Fig 6) and transferring the accumulated data off the detector at 10 μs intervals for offline reconstruction, by analogy with ILC/CLIC. However, due to the high rate of pileup events at HL-LHC, the data volumes with this approach would be completely excessive (~50 MB per event).

For this reason, just as with other LHC detectors, we suggest localised readout of Regions of Interest (ROIs), and this looks extremely promising. This strategy involves firstly reading the data from the timing layers for each Level-1 trigger, and using this to perform fast track

reconstruction in a matter of microseconds. This should be relatively straightforward, since track-finding with pixel-based detectors is much simpler than with microstrips. As a first step, we quickly find the track stubs associated with the vertex detector and outer timing layers (Fig 2). The more problematical next step is to link these unambiguously. The measurement precision on the track stubs is particularly helpful for linking in the RZ plane, where the track projections are accurately colinear. The precision deteriorates due to multiple scattering as momentum is reduced, but the low momentum tracks are well separated from the dense jet cores, and multiple scattering in the SPT is much smaller than in the current tracking systems, due to the favourable material budget. Back-of-envelope calculations suggest that this step can be made cleanly, but simulations are of course needed to study this properly. Even if it turns out that the linking is ambiguous for a small precentage of tracks, it is important to remember that while the current trackers work well for muons, they are in some trouble with pions and electrons [4], so we do not have a hugely challenging target to beat. A small degree of confusion in found tracks could be more than compensated by the greatly reduced rate of spoilage due to secondary interactions and photon conversions. The next step is to accept as a 'done deal' those tracks found in this way which have momentum below some threshold P_{THR}. The tracking layer data are of course needed for precise measurement of high momentum tracks, but below some threshold, the linked stubs from the timing layers will suffice. It is suggested that P_{THR} should be set to ~3 GeV/c. For above-threshold tracks, the found trajectory is used to define a broad road (~1x1 cm^2 in rϕ and z) through the tracking layers. The ROI logic opens gates (Fig 7) which transmit selected data from the primary buffer to the accept buffer, to be read out by the cyclic readout (LVDS to electro-optic converters, one at each end of the tracking system). Data stored in the primary buffer which are not selected by the ROI are simply overwritten. This strategy can reduce the data volume from the tracking layers to some tens of kB per event – very easily handled by a single optic fibre at each end of the tracker.

The most critical question regarding the feasibility of the SPT concept is the power dissipation in the tracking layers, which is of course most challenging for the LHC application. We have carried out a preliminary pixel design for the functionality of Fig 5, which has the following key features:

- ~14 transistors per pixel
- FE source follower operating in subthreshold region
- FE signals stored on 0.1 pF capacitors at input to differential amplifier (noise performance is entirely adequate)
- Differential amplifier gain x4
- Digital output with sufficient drive capability for fast signals on the column line to chip edge

Since the hit density is low, the power dissipation is dominated by the circuitry before the discriminator output.

For a 250 nm process (with which we have considerable experience) we estimate a power dissipation during the pixel-on time of 5 μW. The response of the circuit (simulated with SPICE) is sufficiently fast that a total on-time of 50 ns is sufficient, so the average dissipation

for the 10 μs cycle time is 25 nW. We repeated the simulation for a 180 nm process (where we also have good experience) and the dissipation is reduced to 18 nW. Given the limit for gas cooling of ~ 33 nW (Section 4), this is a comfortable situation, leaving open the possibility of shorter integration time if necessary. Furthermore, there would be a power advantage of reducing the feature size still further, but the implications for these analogue circuits would need detailed study. These encouraging results elevate this aspect of the SPT concept from back-of-envelope calculations to a high level of confidence.

What about radiation effects in these monolithic pixels? The innermost tracking layer will be at a radius of ~30 cm for the barrel (closer for the endcaps), so radiation effects are much reduced compared to the inferno of the vertex detector. Charge-coupled CMOS pixel devices, being composed of relatively low resistivity silicon, are more radiation resistant than microstrip detectors, other things being equal. However, effects of hadronic radiation on silicon devices are sufficiently complex that one cannot make any assumptions. What is needed is a comprehensive set of measurements. Fortunately, devices already existing for astronomical and other scientific applications, are sufficiently similar to our needs that such measurements could be made with existing devices.

6. Synergies, next steps and conclusions

The scale of this tracking system (~30 Gpixels for the tracking layers) certainly sets a new scale for scientific pixel systems. On the other hand, such systems have on average been advancing by a factor ten in number of pixels every nine years for decades. The current record is set by the LSST focal plane array (3.2 Gpixels) so our suggested system will be 'on the line' by 2020 [13]. Furthermore, the 307 Mpixel SLD vertex detector was in 1993 ten years ahead of Luppino's version of Moore's Law, so one can always push a bit harder than the average pace.

One concern raised in the LC community was the likely cost of such a tracker, compared with microstrips. Projections 10 years ahead are of course difficult, given the pace of progress with such devices for other science, but some of our commercial partners have provided useful guidance. By following the same procedure as for SLD (purchasing devices that passed only a basic DC test, so the bulk of testing rested with the customer), it is estimated that the total system cost would be in the region of $40M. While this is more expensive than microstrips, the advantages may greatly outweigh the cost difference. If there is a clear performance advantage (and microstrips may be in serious trouble due to hit densities from pileup events at HL-LHC), the pixel option may rapidly pay for itself in terms of 'luminosity factor', the associated reduction in integrated luminosity needed to achieve a specific physics goal. This is likely to be particularly true of any multi-jet physics, for which there will always be one or more jets in the forward region, for which the current trackers are working badly.

The development of charge-coupled CMOS imaging systems is advancing vigorously for numerous scientific applications, and there are synergies to be exploited. For example, Jim Janesick's 'sandbox' processing runs [14] provide an opportunity for developing sophisticated prototypes in an inexpensive multi-project environment. Now that these devices have made an entry into the scientific market with major commercial backing (Fairchild/Andor/PCI [15], and

Hamamatsu separately), we can expect this technology to be pervasive on the timescale of the next major construction projects for particle physics at the energy frontier.

References

[1] *Report of ILC Detector R&D Panel review of tracking systems*, Feb 2007: http://www.linearcollider.org/wiki/doku.php?id=drdp:drdp_home

[2] *Silicon Pixel Tracker for the ILC*, K Stefanov, ILC Workshop, Sendai, March 2008: http://ilcagenda.linearcollider.org/conferenceDisplay.py?confId=2432

[3] SiLC Collaboration: http://lpnhe-lc.in2p3.fr/

[4] *Performance of the ATLAS silicon tracking detectors*, M Elsing (report at this workshop)

[5] *Design and performance of the SLD vertex detector: a 307 Mpixel tracking system*, K Abe et al, NIM A 400 (1997) 287

[6] *Performance estimates for Silicon Pixel Tracker*, C Damerell, ILC Workshop, Warsaw, June 2008: http://ecfa2008.fuw.edu.pl/

[7] R Page, PhD thesis Bristol U 2011 (to be published shortly)

[8] *Fundamental performance differences of CMOS and CCD imagers*, J Janesick, Part I, Proc SPIE Vol 6276-7 (2006), Part II, Proc SPIE Vol 6690-2 (2007), Part III Proc SPIE Vol 7439A-6 (2009), Part IV Proc SPIE Vol 7742-11 (2010)

[9] *Design of analog integrated circuits under low-noise constraints*, PF Manfredi, IEEE-SSC lecture notes (2007)

[10] *Toward 100 Mega-Frames per second: design of an ultimate ultrahigh-speed image sensor*, Dao Vu Truong Son et al, Sensors 10(1) 2010, 16

[11] D Burt, supplementary channels etc (private communication). For details, contact David Burt at david.burt@e2vtechnologies.com

[12] N Konstantinidis and G Crone (UCL, London) (private communication)

[13] *High energy, optical and infrared detectors for astronomy*, G Luppino, Proc SPIE 6276 (2006)

[14] For details, contact Jim Janesick at CMOSCCD@aol.com

[15] *The online home of scientific CMOS technology*: http://www.scmos.com/