EDMS No 466 475

April 30, 2004

Beetle Production Readiness Review

Date:	April 20, 2004, 9:00 – 15:30
Place:	Kirchhoff-Institut für Physik, Heidelberg
Agenda:	`http://agenda.cern.ch/fullAgenda.php?ida=a041377#2004-04-20`
Reviewers:	Christian Bauer, Achim Vollhardt (Silicon Tracker)
	Jan Buytaert, Massimiliano Ferro-Luzzi (VELO)
	Martin van Beuzekom, Eddy Jans (VETO)
	Jorgen Christiansen (LHCb management)
	Michael Schmelling (Chair)

All presentations shown at the meeting are accessible through the agenda Web-page.

1 Executive Summary

The reviewers would like to congratulate the designers for the excellent job done in the development of a readout chip that satisfies the requirements of the LHCb silicon strip detectors. The existing chip, Beetle1.3 is suitable for VELO and ST and should be produced without further delay together with a new version, Beetle1.4, which contains a new comparator design for the pileup VETO plus the low risk fixes for the known minor problems. The submission may include a third design, Beetle1.5, if a deeper understanding for the up to now not-understood problems should emerge and a fix can be implemented without affecting the projected submission date.

2 Qualification of the Beetle-Chip

All groups planning to use the Beetle chip did a careful qualification of the Beetle 1.3.

2.1 Lab Test Results

The chip was subjected to detailed lab-test, as presented at the Beetle Users Meetings November 2003 and February 2004. The tests covered

- the voltage range for stable operation of the chip
- measurements of the total power consumption
- startup and operation at different temperatures
- over-clocking
- random trigger test
- measurements of front-end performance as function of capacitive load
  - rise time, peaking time and remainder
  - ENC
  - temperature dependence
- Problems were identified in
  - different baselines between consecutive and non-consecutive readouts
  - “forward” cross talk into the successor channel
  - few % cross talk with even-odd pattern between readout lines of the pipeline
  - wrongly encoded parity bit for pipeline column number
  - marginal performance of the comparator
  - non-working daisy-chain mode
  - reduced length for last readout sample if `Rclk<Sclk`

The last two points do not affect the LHCb operation. The comparator issue is discussed in greater detail below. Forward cross talk is seen both from the last header bit to the first data sample and between subsequent data samples. Design and operation of the Beetle1.3 are described in detail in the Beetle1.3 Reference Manual. In addition an extensive review of the Beetle architecture and some considerations defining the Beetle1.3 are available.

New results were presented at the meeting from an SEU-test including a total dose irradiation at the low energy proton irradiation facility at PSI, and a Focused Ion Beam patch verifying that the pipeline column number parity bit can be fixed by a simple swap of two signal lines. The conclusion from the SEU test was that at nominal operation conditions one will have at most one SEU per 25 minutes in the VELO. Most of those will be corrected by the triple redundant layout of all crucial flip-flops. The non-corrected ones will only affect the pileup veto at a rate of about 1/day. These limits are derived from the observation of 4 SEUs within a window of 0.923 Mrad from a total of 7.9 Mrad. The reason for the clustering is not understood. Assuming only poissonian statistics at otherwise unchanged conditions, the probability to have a burst of 4 or more SEUs in such a window is only 0.2%. No degradation of the analog performance of the front-end has been observed after 7.9 Mrad proton irradiation. A radiation induced 10% variation of the current through the chip was observed, but is so far not understood. After a gradual rise at the beginning, the current dropped quickly after changing the beam energy and continued to approach the default values even after going back to the initial energy. For LHCb it is believed not to be a problem, since a flux in excess of 10^9 hadrons/cm^{2}s is not expected for extended periods of time.

2.2 Silicon Tracker

The main issue addressed in tests performed by the Silicon Tracker group concerned the high rate behaviour of the chip in LHCb 4-port readout mode and whether the common mode subtraction does compensate baseline variations. Fixed trigger rates, varied between 1 kHz and 1 MHz, showed no impact of the trigger rate on the noise performance of the chip. Running in 4-port mode and using the linear common mode subtraction scheme foreseen in the TELL1 board gave practically identical performance as a 3rd order polynomial in single port mode. It was found that the baseline had quite different shapes between chips, but the variations were always linear and could therefore be well corrected for by the linear common mode subtraction. The variations were of the same order of magnitude as detector induced noise. The results are independent of whether or not the readout is consecutive, or whether pedestals were determined in consecutive or non-consecutive mode. The conclusion was that the Beetle1.3 satisfies all requirements of the Silicon Tracker.

2.3 VELO

Detailed studies by the VELO group, using bursts of events with variable delay between the bursts read out in 4-port LHCb mode confirmed the results reported by the Silicon Tracker. The measurements were done with an ATLAS "baby sensor" and will be repeated with a full VELO hybrid. A regular "hairy" structure contributing less than 10% to the noise when not biasing the detector was understood to originate from feedback through the bonded strips. A lower than expected S/N performance from last year's testbeam data is not yet understood and needs to be followed up, but is not believed to be caused by the Beetle chip. First tests at the system level gave no indication for problems with the full hybrid. Cross talk and pulse shape measurements are planned for the testbeam beginning of June. In conclusion, the VELO group shares the opinion that the Beetle 1.3 will work properly in the final detector.

2.4 Pile-Up VETO

The VETO group showed detailed studies of the performance of the comparator for the pileup-veto. Apart from a non monotonic behaviour in the DACs for the comparator threshold, which reduces the effective range to 27 rather than 32 values, feedback in the comparator leads to oscillatory behaviour for positive signals when the threshold setting is too low. The situation is stable for negative signals, i.e. for the case of using n-on-n detectors. Still, also here the minimum possible threshold setting is so large that it would significantly compromise the efficiency of the pileup veto. In the simulation at least part of the observed feedback could be traced to switching noise in one of the inverters in the digital signal path. A redesign exists which fits into the existing space of the comparator and reduces this noise by one order of magnitude. The new design also provides a linear DAC characteristic for the threshold. In addition, sensitivity to variations in the power supply is reduced by new dimensions of the input buffer. The VETO group confirmed their conclusion that they definitely need the Beetle1.4 for the pileup detector. In case the performance of the Beetle1.4 comparator should still fall short of the specifications for the VETO, the required number of chips could still be produced in an MPW run without affecting the schedule for the other users.

3 Modifications with Respect to the Beetle1.3

A list of modifications with respect to the Beetle1.3 was presented to fix some of the known problems and to improve the handling of the chip.

- exchange two signal lines to correct the parity bit of the pipeline column number
- introduce a chip-ID by hard-wiring the unused bits of the comparator threshold registers
- realize a visual identification as optical alignment marker for automatic chip testing and bonding, and to allow optical discrimination between different Beetle-versions
- increase the spacing between the readout lines in the pipeline to reduce the coupling responsible for the even-odd pattern observed in cross-talk measurements as function of the channel number
- put the pipeline cells into an n-well structure to increase the linear range of the chip, enhance slightly the signal gain and to shield the pipeline against cross talk via the substrate.

Except for the last item all modifications are considered very low risk. Although desirable from the design point-of-view, the benefits of the n-well for the actual operation are marginal, since the small gain in linearity for large signals is not relevant for the application in VELO and ST. Understanding is still lacking for the forward cross talk from the header to the data samples and between subsequent data samples, and for the baseline variation between consecutive and non-consecutive readouts.

4 Production Planning

The production planning was presented for an “optimistic” and a “late” scenario, where the optimistic scheme should be compared to the chip requests by the different groups. In that scenario, with submission of the engineering run on May 15, first untested chips would be available beginning of September 2004. Large tested quantities can be expected after beginning of October. Depending on whether two or three designs are submitted to the engineering run and assuming a yield of 85%, 1200 or 1800 chips of each kind can be expected from 6 wafers. Ordering of the production run then is expected in October 2004, the final delivery would be between 18 and 26 weeks later. For a yield of 85% this would result in a total of 32400 working chips shared evenly between all submitted designs.

Currently there are still 170 Beetle1.3 chips on stock, which can be used for pre-production and prototyping, and have to be shared between VELO, VETO and ST. Although this appears rather tight, the representatives of the different subgroups were confident that the supply would be sufficient until the chips from the engineering run become available.

5 Decisions and Conclusions

No arguments or pending tests were identified which would justify a delay in the production of the chip. There was consensus that the submission should include the Beetle1.3 and the Beetle1.4:

- Beetle1.3 – unchanged with respect to the current design
- Beetle1.4 – with the following modifications compared to the Beetle1.3
  - modified comparator as presented at the meeting
  - fix the parity bit of the pipeline column number
  - fix the cross talk in the pipeline
  - Beetle version number coded into unused bits of the comparator threshold register
  - optical alignment marks serving also for visual identification of the chip version

A third design containing higher risk modifications, the Beetle1.5 as discussed below, will be included if a significant improvement of the chip's performance can be expected. The reticle will be equally shared between the two or three submitted designs, which in both cases will provide a sufficient number of chips to equip the respective subsystems of LHCb with either of the versions. This includes a future upgrade of the VELO. Only in the event of a dramatic loss in yield another production run may be needed.

The Beetle1.4 shall be finalized with high priority as soon as possible. If before the final submission date, targeted for May 15, 2004, a substantial understanding should emerge about the origin of the forward cross-talk and the baseline variation between consecutive and non-consecutive readouts, and if a simple solution can be implemented for either of the two, then a Beetle1.5 will be realized which also includes the n-well for the pipeline. Design work on the Beetle1.5 must not delay the submission. Apart from the academic interest in understanding the not yet understood features of the chip, the potential gain in chip performance make it appear worthwhile to produce the additional design and to qualify it for the construction of the VELO, VETO and ST.

It was pointed out that if possible for the production, all chips on the reticle should have the same metal filling in order to achieve identical parasitics independent of the location of the chip on the reticle. This is, however, not considered critical if it cannot be achieved.

The next steps in the chip production concern the definition of the testing procedure and quality control. The discussion, with input from the user groups, should start immediately after the submission. Production tests of chips at the wafer level must be prepared to be capable of delivering tested chips to the users during fall 2004, at a pace defined by the production schedule of the silicon detector modules. In addition, the ST, VELO and VETO groups must be ready to do quick tests of the Beetle1.3 and 1.4 when they return from the foundry, to be capable of making the step to the final production run ASAP.

Until the final chips are available the groups should negotiate about the distribution of the remaining 170 Beetle1.3 dies for test and pre-production issues. The general feeling was that available supply is tight but still sufficient.