Report of LHCb VeLo Electronics Architecture Review

7th November, 2013 at CERN.

https://indico.cern.ch/conferenceDisplay.py?confId=281473

Review panel:

Federico Alessio (CERN), Alex Kluge (CERN), Ruud Kluit (NIKHEF), Ken Wyllie (CERN).

The reviewers thank the VeLo group for the clear presentations and the constructive discussion during the review meeting.

1. General Comments

The proposed VeLo FE architecture is compatible with the global upgrade architecture. Solutions were presented covering all the aspects of the upgrade. The proposed implementations all seem feasible but many parts of the design are still very preliminary. The VeLo group should now accelerate the system aspects and the development and prototyping of individual components.

No potential "shows-stoppers" were observed, but there are some items whose development and verification is critical for the schedule (see below).

The system performance should be simulated together with the TELL40 code and the TFC commands. The throughput of the DAQ system, including the required time information (see BXID discussion & time-ordering) requires more detailed analysis.

System integration aspects always impact on the performance of the front-end circuit. Given that the VeloPix is currently under design and hence still open for optimisation, system issues should be investigated immediately with Timepix3 and fed-back into the VeloPix design.

Although the radiation levels expected in the hottest parts of the detector appear to be well known, the levels in the area of the OPB are based on extrapolation from existing simulations. These estimates can be critical in the choice and evaluation of components on the OPB. The VeLo group should contact Matthias Karacson (Gloria Corti) to request simulation data specific to the location of the OPBs.

This will be done.

No documentation was prepared in advance for the review. This will be essential for the design of the new components and the overall system so should be started immediately. This is particularly relevant for the design of VeloPix. In general, it is recommended to provide a full list of requirements for the system and the chip. This can then be used as a checklist to evaluate the completeness of the designs, allows easy communication between designers and can help in making certain design choices.

A first version of a VELOpix specification document has been written and will be ready for the ASIC review early 2014. The overall system has been described in the VELO upgrade TDR. Separate and

more detailed specification documents will be written for parts of the system, such as hybrid, readout links, OPB boards and Tell40 firmware.

Thermal runaway was mentioned a few times during the review. How severe is this risk? This should be assessed more carefully and the preventative actions implemented in the system design.

The temperatures at which thermal runaway can occur have been simulated with ANSYS for the strip option. The pixel option with microchannel cooling is safer since the sensor tip can be kept more easily at even lower temperatures.

The VeLo group is asked to converge on the final quantities of common components (GBTX, VTTx/VTRx, GBT-SCA, DC-DC convertor, ST linear regulator) by the end of this year. Production of many of these is planned to start in 2014.

These quantities have been decided and are documented in the VELO upgrade TDR. Meanwhile a further optimisation of the number of readout links has reduced the total number of VTTx.

2. VeloPix, hybrid and links

1. The VeloPix design is advancing well, and will be the subject of a design review early 2014. However, a complete set of specifications should be compiled immediately to ensure that the design goals are clear and all requirements are satisfied.

   A (preliminary) written specification exists, but this was not handed out to the review committee because the VeloPix was not officially being reviewed (in hindsight: we should have distributed the specification). The current write-up will be updated, and the recommendations from this review report will be included.

2. Timepix3 will provide valuable information for the VeloPix design. In particular, tests on radiation, low-temperature operation and full-bandwidth readout should be completed as soon as possible in preparation for VeloPix.

   Agreed.
   Radiation tests: SEU test on the medipix3 (same HD library) have been performed recently, and X-ray irradiation test (also non-homogeneously) are planned.
   Low-temperature operation: A set-up is being constructed (by collaborators from UFRJ) to study the thermal performance of the Timepix3
   High speed readout: The SPIDR readout system has been designed to cope with the maximum bandwidth of Timepix3. Besides bench-tests also high rate beam tests are planned.

3. VeloPix will be exposed to a very non-uniform radiation dose (almost a factor of 10 between maximum and minimum). The effects of this should be understood in terms of long-term operation, for example how often the thresholds will require calibration. This should also be understood in terms of sensor degradation and the subsequent reduction in signal-size.

   Agreed.
   (non-uniform) Xray irradiation of bare timepix3 ASICs will give us an estimate of how frequently the thresholds require recalibration. Tests with irradiated sensors are planned, but will be late due to the delivery time of the sensors. How can we speed up the irradiated sensor testing such that we can check that there are no surprises.

4. Can the temperature gradient through the ASIC create any problems?

   Thermal simulations show that the temperature difference across the ASIC is smaller than 5 ‘C. This is not considered to be a large thread to the ASIC. Thermal tests will show if the ΔT of 5’C is indeed met.

5. The data-packet latency has been carefully simulated to optimize the number of BXID bits to transmit. From simulation the peak latency is at 64 clock cycles out of a maximum range of 512 (9 bits). Although this appears sufficient margin, it is recommended to implement a safe and reliable way to monitor the statistics of the data latency.

The most appropriate place/way to monitor the BXID latency is to create a histogram of the latency in the TELL40. Additionally one could make a time window around the current BXID in which the BXID of incoming data packets should fall. An excess of out-of-window data packets could be easily signaled this way.

6. The GWT is a crucial part of the VeloPix design and is on the critical path. All effort must be made to ensure its success and also to minimise the evaluation time of the test chip. The submission, preparation and testing should be carefully planned and scheduled, and manpower reserved in advance. The goal of the test-chip should be to verify a complete solution that can be incorporated directly into the VeloPix design. Hence, all aspects of the GWT and its accompanying blocks (eg PLL) should be included.

   Agreed. Submission of the GWT is on schedule, and the chip will be submitted on Feb 18th. Details of the tests and required infrastructure (PCB, equipment, FPGA firmware) are being discussed and this work must be done in the shadow of the submission such that testing can commence immediately when the chip return from the foundry.

7. The contents of the data packet corresponding to the 'Synch' command should be decided.

   This will be added to the specification document.

8. A reset pulse for VeloPix is required. It is not yet clear where this can be generated. This should be decided as soon as possible. The functions of the different resets sent to VeloPix should be rationalised.
   As was the case with the TIMEpix3, a careful discussion will take place on the reset and power-up requirements of the VELOpix. The optimal place in the ECS to generate these signals depends on the number of signals: a local SCA or an SCA on the OPB? All will be documented in the specification document.

9. The TFC interface to VeloPix is compatible with specifications. The proposed interface will use serial data at 320 Mbps, one link per VeloPix. Care must be taken in making sure the decoding of the TFC bit stream is correctly aligned in time with the 40 MHz clock. This will probably require a configurable delay based on the 320 MHz clock. Five TFC commands (5 bits) are sufficient. The remaining 3 bits can be utilized for other synchronous commands outside the normal TFC specifications. The allocation of the TFC bits must be specified when ready. It is not recommended to transmit the VeloPix reset in this bit stream.

   Tying VeloPix specific functions to the 3 'free' bits in the TFC stream was not considered thus far, but is certainly an interesting feature.

   The (hard) reset of the VeloPix will not run via the TFC interface to the VeloPix. Controlling the reset via the SCA seems a safe option.

10. The VeloPix must have the ability to individually delay each of the TFC bits in units of 25ns. This is foreseen to allow full flexibility in time-alignment. A range of 16 clock cycles is the minimum requirement as defined in the global specifications.

    This requirement has been added to the specification document.

11. The custom SPI implementation for slow-control was presented and meets the specifications from the point-of-view of the overall ECS. This requires four differential lines and hence two Eports. The VeLo group is encouraged to investigate a protocol more suited to a single Eport, for example the frame-based protocol used to communicate between GBTX and GBT-SCA. The reviewers, however, acknowledge that the SPI-like interface may fit more naturally with the internal structure of the VeloPix configuration logic.

    Investigations have been done, and a single e-port implementation seems feasible.

12. An up-to-date estimate of how much data is needed for configuration should be provided for reference for the ECS group. This should be used together with the ECS protocol to calculate the time required to configure the VeLo and for calibration.

    Initial numbers about the amount of configuration data (and the size of the chunks of configuration data) have been communicated to the ECS group. For the configuration data, the time will not be determined by the VeloPix.

    The time required to do a full chip calibration will be calculated and communicated to the ECS group.

    At this moment, an option of self-calibration is still under development. This would speedup the calibration enormously and would also alleviate the required bandwidth of the ECS.

13. There was some discussion on including some 'GBT-SCA functionality' within VeloPix, for example the ADC. This would reduce the complexity of the hybrid. This should be assessed in terms of the VeloPix design and schedule.
    It is under discussion with the designers. If an ADC 'block' is available, it is likely to be integrated in the VELOpix. This will monitor internal DAC's and external supply voltages.

14. At the time when the first VeloPix chip is available, a fully functional test and readout system must be available. This implies a system capable of handling almost 20 Gbit/s. This is not trivial and requires significant preparation time. This will include full ECS/TFC functionality and the GWT decoding.

    Agreed. The expected date for first VELOpix asics is early 2015. Preparations for a readout system will start early 2014 and will be based on prototype Opto-Power-boards and AMC40 hardware and firmware.

15. The VeLo group should carefully study and document the clock distribution on the hybrids. Multiple clocks (40 MHz, 320 MHz and others) will be required. The Master-GBTX is well

positioned on the hybrid for this purpose. Clock management should be done using the GBTX functionalities which allows for phase adjustment and clock "driving". However, the GBTX does not allow the phasing of Eport clocks. This should be considered carefully for cases where data crosses between clock domains.
The GBTx on the hybrid will provide an individual 40 MHz clock to each VELOpix (6 in total).

16. There is no error correction on the FE links. An estimate should be made of the rate of bit flips induced by SEUs. This should be one of the goals of the GWT testing. With this estimate, the robustness (efficiency and down-time) of the complete link should be evaluated. The ability to handle errors should be implemented in the FPGA receiver according to the predicted error rate.
The SEU sensitivity of the GWT will be measured. The sensitivity of the VTTx is known. From these numbers the link sensitivity can be estimated. The 4 parity bits that are transmitted with each GWT frame allow detecting data corruption in each of the 4 30bit SPP (super-pixel packet). This will be detected and handled by the TELL40 GWT decoding. The VELOPix will also have an internal PRBS generator that can help evaluate or diagnose the link quality or the SEU upset rate in a radiation environment.

17. The first ideas for the hybrid design were presented. A prototype based on Timepix3 is proposed, and the VeLo group is encouraged to start this immediately without delay. The goals of this should be to investigate mechanical integration issues like gluing, wire-bonding and cooling connections.

    The TIMEpix3 hybrid design has started in Liverpool and will be used for integration studies done at Manchester and Nikhef.

18. There was some discussion on using packaged chips on the hybrid (GBTX and GBT-SCA). Using these chips in unpackaged format will be difficult and would imply mounting untested chips on the hybrid. The VeLo group should therefore investigate the hybrid design assuming packaged chips. A preliminary pin-out and package description of the GBTX are available and can be used. The GBT-SCA project leader (Kostas Kloukinas) should be contacted for the details of the GBT-SCA.
    We have no other choice but to use packaged GBTx on the hybrid. But we will investigate how to remove the unnecessary copper cooling plate on top.

19. The correct operation of all components together on the hybrid should be verified to work at -40^{o}C. Agreed.

20. The VeLo group plans to use GBLD chips in a number of places. They should contact the Versatile Link project leader to discuss procuring single GBLD packaged chips (not encapsulated within a VTTx or VTRx). After discussion with J. Troska, a quantity of 300 GBLD is 1% of the total production and that should not pose a problem, provided the number is announced soon.

21. There was some uncertainty on the voltages required on the hybrid. 1.5V is certain, but 2.5V may be required for the GBLD. This should be clarified.

According to the designer (Gianni Mazza) it seems possible, but a test must be done to validate this particular voltage supply scenario. He has not simulated this case.

22. It was pointed out that VeloPix could be designed with 'power-pulsing' functionality and that this could be useful for lowering the current if a high-temperature interlock is activated. If the VeLo group desires this feature then it should be specified in the VeloPix requirements.
    The suggestion will be discussed. But the preferred interlock is by switching off the DC-DC converters on the OPB boards. The power pulsing is an additional non-essential design complication.

23. Because of clock-gating circuitry, the digital power consumed by VeloPix will vary with occupancy. Attempts should be made to simulate this to allow optimisation of the power distribution on the hybrid.
    VELOpix simulations will be done to estimate the transient in power consumption induced by varying pixel hit rates, for example caused by the abort gap or partially filled bunch schemes. If feasible, idle data will be generated to minimise these transients.

24. It was discussed if all data outputs of all VeloPix's should be connected on the hybrid and through the vacuum tank. Only the enabled outputs would then be equipped with transmitters on the OPB. If needed, this would allow these outputs to be used later without requiring intervention inside the vacuum tank.
    Meanwhile we have re-optimised the number of links required per module from 32 down to 20. We will nevertheless install 32 links inside the vacuum. The OPB board will be designed for 20 active links and we will study how to best cope with a broken link inside the vacuum. All downstream optical components will be dimensioned for 20 links per module.

25. Measurements on a prototype flex cable were presented. Although these are encouraging, the tests should be expanded to include the other parts of the link (drivers, receivers, patch-panel connectors, optics, fibres) and for the real distance required in the system. The eye diagrams showed degradation using clock-recovery at the receiver. This should be carefully evaluated with the full link. A bit-error-rate test must be carried out with a realistic receiver (eg AMC40). Results from such a test should be fed-back into the design of the GWT transmitter and its pre-emphasis function.
    We will design new prototypes of the 4 segments of the data link (hybrid part, long cable, vacuum feed through and OPB board). With the GWT prototypes as data sources, we will study the signal integrity of a complete link.

26. Testing of the links requires specialist equipment, and the relevant institutes should ensure they are properly equipped.
    Glasgow University is at the point of purchasing VNA or a high bandwidth oscilloscope.

27. An industrial supplier of the flex cable should be identified.
    Prototypes will be developed at CERN pcb workshop. We try to keep the technology compatible to industry standard manufacturing practice.

28. The data and control links require different combinations of drivers and receivers. For example, a GBLD will drive signals to another GBLD (within a VTRx), and a GBLD will drive a GBTX. These combinations should be studied, checking that the signals levels and swings are compatible between the chips.
    This is part of the planned prototype program.

29. The impact of the wire-bonds on the signal quality should be studied and compared with studies from other projects, for example the NA62 GigaTracker.
    The effect of wire bonds was studied. The insertion loss was less than 0.1dB.

30. The grounding scheme of the hybrid and the OPB should be studied, including the powering scheme using the DC-DC convertors. The powering/grounding of the sensor should be included. An estimate of ground bounce should be used to check the power-supply-rejection-ratio of VeloPix. Tests of the signal-to-noise ratio of Timepix3 should be made with DC-DC convertors and the bulk power supplies.
    A module with Timepix3 asics will be built to assess the effect of DC-DC powering and the grounding scheme on the measured pixel noise and threshold stability.

31. The segmentation of the powering on the hybrid should be investigated. In the case of a catastrophic failure of a component, how will the others be affected?
    We cannot further segment the power supplies. 3 ASICs will share a common power supply. Failure scenario's will be investigated during the VELOpix design to ensure that a single failing asic will not affect correct operation of asic powered by the same supply.

32. The positioning of temperature sensors should be considered carefully to give the maximum information in the shortest time to the interlock system.
    We consider a temperature measurement in the VELOpix and a sensor on the cooling substrate. The latter will detect only the absence of cooling liquid circulation and be a permanent monitor of module temperature. It will also activate the interlock system. The internal velopix temperature is the closest to the sensor temperature.

33. The design of VeloPix, hybrid and the mechanics/cooling should now progress together hand-in-hand with good communication between designers. In particular, the pad layout of VeloPix should be carefully optimised at an early stage by the chip and hybrid designers.
    A first proposal of pad layout has been discussed between ASIC and hybrid designers in view of an optimal asic and hybrid layout.

3. OPB, TELL40 and power distribution

1. The use of commercial-off-the-shelf components in the OPB modules should be assessed together with the expected radiation levels in the upgrade environment. Will radiation tests be required? These can take a long time and should be planned immediately.
   This concerns the local interlock circuit on the OPB board. Further studies are planned to consider alternatives avoiding the use of COTS on the OPB board.

2. The VeLo group should carefully assess the need for a metal cage around the VTTx/VTRx modules. This question was already asked to the sub-detectors in the past, and there was no objection to the plastic cage at that time.
   The pcb is not compatible with the standard SFP metal cage. Development of a custom metal cage is too complicated. A different cooling solution is envisaged using a cooling bar contacting the GBLD device (highest heat dissipation).

3. The proposal for the TELL40 processing is sound. However, the VeLo group should concentrate on two urgent tasks. Firstly, the GBT-FPGA receiver code should be adapted to handle the GWT frame structure. The receiver relies on embedded 'hard-IPs' within the FPGA and it must be verified that these are compatible with the GWT. For example, the 128-bit frame of the GWT is not compatible with the 20-bit deserialised frame output from the receivers currently implemented in the GBT-FPGA code. Any special requirements to satisfy the GWT format must be transmitted to the TELL40 designers. Secondly, the VeLo group should investigate the implementation of their specific processing tasks in the TELL40 firmware. This is a critical issue as the required FPGA resources can limit the number of links handled by a TELL40. For both of these tasks, the VeLo group should collaborate closely with the Annecy and Marseille groups.

   The VELO is aware of the urgency of both issues of VELO specific TELL40 decoding and processing. A lot of progress has been made since the review. It has been estimated that the required resource would comfortably fit in the proposed AMC40 fpga. A team of 3 people are developing the code at this moment.

4. The VeLo-specific TELL40 tasks were defined as the time-re-ordering of data packets and clustering. There is a high probability that there are not enough FPGA resources for both of these tasks, and priority should definitely be given to the time-re-ordering. It was agreed that abandoning the clustering would not have a significant effect on subsequent event processing.
   The time-reordering has been implemented with a 'packet router' that can handle the data rate coming from a module with adequate margin. This router requires no more than 5% of the FPGA resources in ALU and memory. No work has yet started on clustering in the FPGA , as more understanding is needed to what can usefully be done to speed-up the pattern recognition.

5. The mechanism in the TELL40 to add the extra 3 bits to the BXID should be studied and documented.
   The complete 12bit BXID will be added to the data as a header information in the bxid event buffer after the router. This will be simulated and documented in the firmware documentation.

6. The proposed GWT data format should be presented and discussed with the TELL40 firmware group.
   Contact has been made with the GBT firmware support group at CERN and presentations have been made since in the AMC40 working group.

7. The VeLo group should rationalise their hardware needs for verifying the TELL40 code (for example, the AMC40 or FPGA evaluation boards).
   The VELO group as ordered one AMC40 board to validate the firmware in real-time.

8. The existing VeLo power supplies appear adequate for the upgrade. However, the designers of the upgrade system should avoid being tied to specific features of the existing supplies in case these need to be replaced in the future.
   The current CAEN LV system is adequate in terms of functionality, segmentation and power rating for operating the upgraded VELO system until they reach end of life or lack of maintenance support. The VELO will anticipate and look for a replacement system with similar specifications when that becomes necessary.

9. Will the power parameters change through the lifetime of the detector? For example, will adjustments for radiation effects change the power consumption? This should be estimated and included in the calculation of the limits of the power system.
   This applies mostly to the HV system when higher voltages above 700V are needed. Also the maintenance of these ISEG HV system may become an issue. The VELO is aware of this and will try to anticipate.

4. Manpower and Scheduling

The schedule was presented. This meets the experiment schedule but is aggressive and depends very much on the VeloPix. It is recommended that resources are optimized to speed up the development of the VeloPix as much as possible. More detail should be added to the schedule to include the submission and tests of the GWT, and decision points related to this. It was not clear how the schedule would be affected if the GWT fails.

An updated time and resource schedule has meanwhile been published in the VELO upgrade TDR.

The test procedures to verify the GWT functionality should be clarified (eg laboratory tests, irradiations) and planned well in advance. This includes the design of electronics required specifically for testing. Manpower for the testing should be identified soon. The same should be done for the VeloPix.

The division of responsibilities was presented and has been decided at the level of host country. The detailed responsibilities should be fixed at the institute level as soon as possible.

A list of responsibilities of each institute has been published in the VELO upgrade TDR.