LHCb Electronics Architecture Review

5th December 2012, https://indico.cern.ch/conferenceDisplay.py?confId=212204

Reviewers: Jorgen Christiansen, Philippe Farthouat, Alex Kluge

Reviewers' comments and response from LHCb:

General comments on architecture:

Referee 1: The proposed architecture for the very challenging readout/DAQ at full 40Mhz rate seems appropriate/optimal and is progressing very well towards the LHCb upgrade.

The extensive use of radiation hard optical links (GBT) and very powerful modern FPGA's is a very good and vital foundation for the proposed architecture.

The availability of the rad hard GBT link is critical for the upgrade.

The use of the existing L0 trigger systems for the muon and calorimeter detectors to make the required rate reduction for initial running could be challenging/problematic. Keeping alive such an old system with significant parts located in the radiation environment could turn out to be problematic. In the new architecture, where all detector data will be available in the back-end (out of radiation environment), it would be advantageous to implement a new trigger (rate reduction) processing in the counting house based on the very powerful TELL40 hardware platform. If this trigger system is only needed for a very limited time period, then using the old system could be appropriate. There is though a large "risk" that this trigger (rate reduction) will be required for a long period. It is encouraged to evaluate how such a new rate control system can be implemented with the TELL40 platform and determine the resources (manpower and hardware modules) needed to implement it.

LHCb response: We acknowledge that this implies maintenance of a legacy system. However, many other systems in LHC experiments will operate for a similar timescale and will face similar problems. The choice of keeping the existing system is essentially based on manpower and cost. If manpower and money become available, then LHCb can indeed investigate a new system if the LLT is still required at that time. Another advantage of maintaining the legacy system is that it minimizes commissioning time.

Schedule wise my major worry is the implementation of all the different front-ends. In particular those based on new ASICs. As these have not been presented this is just a "guts" feeling.

Referee 2: The proposed architecture makes sense and there are no showstoppers. However a few points are to be carefully monitored.

(1) The availability of the GBT chip set and of the versatile link is crucial for the project. The versatile link

production schedule fits perfectly well the LHCb requirement but the GBT chip set is on the critical path. The impact of a possible iteration on the GBTx should be carefully analysed. Same remark for the availability of the SCA.

(2) There is a need for the front-end powering of either a radiation tolerant voltage regulator or of radiation tolerant DC-DC converter modules. The former could be the existing ST regulator if still available. The latter could be a module based on the CERN DC-DC development; the schedule of this development is compatible with the LHCb schedule but not much slack is available.

(3) To take into account the uncertainty concerning the type of switches available at reasonable cost, enough buffering should be installed in the AMC board. This could be (is) very likely at the expense of density and will impact the cost.

(4) Using copper connection at the output of the AMC's to interface to the switch should be considered with very much care and all implications studied. For instance impact on the density of the back-end electronics, need for complex cabling infrastructure allowing full access to each back-end module, etc. should be taken into account when computing the cost of the system.

(5) Although it is assumed that a 25-ns scheme will be used, it might be worth to evaluate the impact of a 50-ns with the same luminosity.

Referee 3: I was positively impressed by the work which has been done on the architecture and the detail of implementation recommendations in the provided documents.

Given the time schedule I think that the proposal is advanced, however implementation for a final installation until 2018 is challenging, especially considering that also detector front-ends needs to be upgraded, partly with ASICs.

I was very impressed by the TELL40 design and the advanced prototyping. I think that the concept of ATCA and mezzanine cards with a large IO bandwidth in both direction is very appealing, especially for a triggerless architecture. I understood the benefit of implementing memory on that card which would make it even more flexible in the system.

I appreciate the concept of one single type of hardware read-out boards, a dedicated group to provide the FPGA firmware code and one group for the DCS software LHCb wide.

Documentation:

Referee: The architecture document is a nice and vital document. In a few places it is not consistent with what was actually presented in the slides, so this should be updated/corrected.

LHCb response: The inconsistencies and missing items pointed out by the reviewers have been corrected.

Referee: It would be good to specify more things in the architecture document, as for example, the use of ATCA, the protocol used for the TTC information distribution in a crate, the recommended ECS protocols used for the front-end configuration (it was mentioned I2C or non daisy-chained SPI in the discussion), the numbering scheme of the BC in the LHC turn.

LHCb response: ATCA and the TFC details are described in detail in the accompanying documents. More details on the ECS protocols have been included in Chapter 10.

Front-end:

Referee: The general architecture of the front-end is quite clear. It is though already known that there will be exceptions (pixel) to this so information on this could be added. More details on how the clocking from the GBT chips should be done can be added.

LHCb response: Agreed, this has been clarified in Chapter 10.

Referee: Details on how the TFC command bits can be mapped into the GBT E-links will be appropriate. It must be decided if the BE module can map TFC control bits into the GBT words according to the best use of E-links on the front-ends (e.g. duplication into multiple E-links). I think some flexibility on this can be very useful for some critical front-end implementations.

LHCb response: Agreed, this has been clarified at the end of Chapter 10.

Referee: Some clear criteria for deciding the size of the front-end buffer should be indicated (e.g. max 0.01% loss at nominal luminosity). Final buffer sizes for each detector must clearly be fixed in separate reviews of each front-end implementation.

LHCb response: A maximum inefficiency has been defined in Chapter 3.

Referee: It is defined that the front-ends must send an event frame with bunch-ID even if there is no data. This could possibly be a significant bandwidth overhead for some detectors. It could be worth considering if it is appropriate to send only BXID when real data has been detected.

LHCb response: We agree with this point, and this will be reviewed on a sub-detector basis.

Referee: It is not really defined how front-ends inserts idles on the link. Especially when the idles may not be aligned with the GBT words. It is not defined if the data packet length (n) is in bits/bytes/words (or fully up to each sub-detector).

LHCb response: This has been clarified in Chapter 10 (use of the txValid signal in the GBTX). The TFC 'synch' command will be used to force the FEs to insert fixed data patterns that will be used for synchronization purposes.

Referee: It could be clarified that "header only" command and "Bunch crossing Veto" from a front-end point of view have same function (at least this is what I have understood)

LHCb response: Agreed, this has been added to the text.

Referee: It could be useful to indicate what should have solid SEU protection (e.g. BX-counter and buffer pointers) and what not (e.g. detector data).

LHCb response: Agreed, more details have been added to the text.

Referee: In the document (page 13) the Status snapshot command is defined in 3.7 which is about running modes. This must be wrong.

LHCb response: Agreed, the text has been moved to the appropriate section.

Referee: Maybe a warning about that the i2C from the GBT-SCA is not standard I2C (voltage levels !) would be a good idea.

LHCb response: Yes, the voltage levels have already been recognized and discussed with the sub-detectors. They are aware of the issue.

Referee: Schedule wise my major worry is the implementation of all the different front-ends. In particular those based on new ASICs. As these have not been presented this is just a "guts" feeling.

Referee: The sequence of commands to be issued and actions to be taken at the beginning and at the end of a run should be specified.

LHCb response: This will be included in the document after detailed discussions.

Referee: The behavior of the FE during a reset is to be clarified and may be relaxed. For instance is it feasible for all FE chip to still send idle packets during the reset sequence?

LHCb response: This will be clarified in the start-up sequence.

Referee: The recommendation of only using ACTEL (ProAsic) programmable devices is very likely safe (although the behavior of the latest devices is to be checked), however progress made by other suppliers and good experience with other suppliers should not be ignored if it brings a substantial gain in performance and/or flexibility.

LHCb response: An Altera SRAM-based FPGA is being investigated in parallel to the ACTEL ProAsic. However, the ProAsic has already been demonstrated as being sufficient for the FE digital tasks and its radiation tolerance has been widely tested.

Referee: The programmation of the FPGAs through the SCA is to checked.

LHCb response: Agreed, this is already under investigation.

Referee: Even though the GBT chip set and the versatile link are the baseline, the use of exotic higher density and home-made packages has been mentioned. Although this cannot be ruled out, LHCb should be very cautious with this kind of solution and make sure that these solutions are fully qualified.

LHCb response: Acknowledged.

Referee: Data format: the correct decoding of the data stream relies on decoding correctly the length field in the header. In case of errors in the length field, data synchronization is lost and an on-line processor must perform data consistency checks, and send hardware re-sync commands. If data

bandwidth allows I would go for a system with an unambiguously identifiable header, either by position or content to increase data robustness.

LHCb response: Acknowledged. We will discuss this issue with each sub-detector together with their buffer implementation and bandwidth estimates. Text has been added to stress that the header must be protected against bit errors.

Referee: bunch counter rollover at 3563. Is this not too early to give this as hardwired value to ASIC designers. Can this still change, especially for LHC upgrade?

LHCb response: Agreed, the text has been updated and the rollover value is now defined as being configurable via ECS.

Referee: will 3 trigger bits be enough?

LHCb response: Three bits are considered sufficient for the different trigger types foreseen (and in use in the existing detector). Note that the 3 bits for 'Calibration Type' also offer similar functionality.

Back-end – TFC – DAQ

Referee: The back-end and the TFC architecture is progressing very well against the final implementation. The use of a common hardware platform for the BE and TFC seems particular promising. The implementation of the complex firmware for the different use of the common hardware module will though be a significant challenge. It is vital that a solid firmware design and support team is made for this and it is probably important to have this as a "common" team to prevent many sub-detectors getting "lost" trying to make their own (experience from the TELL1).

LHCb response: Acknowledged. The core of this design team is already identified and organization of the task is underway.

Referee: It seems attractive (and feasible) to make a BE module that also has some TFC/ECS links to the front-end. Such a module can "easily" be constructed with the flexible TELL40 module and can be very useful for small scale test systems.

LHCb response: This feature will be available on the AMC cards made available for sub-detector tests.

Referee: I would be more specific in the document about the common TELL40. The introduction to chapter 4 is rather vague/generic. A more detailed block diagram of the TELL40 could be included. Alternatively refer to a separate document on TELL40.

LHCb response: Specific details of the TELL40 are all described in the other document that accompanied the review.

Referee: The concept of MEP is not really explained (so if you are not an old LHCb guy like me, it may be a bit hard to understand).

LHCb response: Specific details of the MEP usage are described in the TFC document that accompanied the review.

Referee: The TFC word is given every clock cycle but some fields are only valid in some cases (e.g. when trigger set active). This should be explained a bit better. The MEP destination is a particular case as only a new MEP destination is required for every N triggers. No bit seems to indicate that a trigger is the last event in a MEP. So this is signaled how ? (change of MEP destination ?).

LHCb response: Acknowledged. This will be explained better in the TELL40 and TFC documents.

Referee: The DAQ interface and related re-send, request, buffering, scheduling is clearly the most critical part. If the large date buffers can not be put on a "single width" AMC module then it seems wise to plan/start the development of a "double width" module with sufficient buffer memory.

LHCb response: Implementation of this buffer will be studied by the TELL40 team.

Referee: Deciding/determining if GBT/versatile links can reach to the surface is a delicate point. The fact that LHCb has relatively low (but how low ?) radiation on the transmitter lasers can be a possible plus.

Referee: Having an individual TFC link per ATCA back-end module would cost in density and cabling effort but would bring a lot of flexibility in particular in the definition of the partitions.

LHCb response: Agreed, but partitioning will also be achieved at the level of the ATCA crate. One ATCA crate is treated as a partition, containing all the TELL40s and SOL40s for one particular detector slice.

Referee: The installation of the backend electronics on the surface would have a lot of benefits. In particular the limitations available space, cooling and/or available power in the racks could disappear. There is a clear implication on the LLT latency however this can be easily overcome.

Referee: A cost/performance study with the following parameters would be useful to decide on the further development path.

*) TELL40 with and without memory (less than 4 cards on a carrier/4 cards per carrier)

*) TELL40 with electrical or optical DAQ connections <-> short electrical cables or short/long optical cables.

LHCb response: Agreed, this cost comparison will be done next year.

Infrastructure:

Referee: Check that Wiener can actually still supply Marathon power supply modules with “guaranteed” radiation resistance.

LHCb response: The maintenance/purchase contract with Wiener extends up to 2019. Radiation-critical components that are obsolete have been purchased centrally by CERN to be made available to Wiener when a new order is made. The estimated quantities of new supplies needed for the upgrade are rather small.

Referee: Check/test DC/DC converter with marathon power supplies.

LHCb response: Some groups have already made this test with positive results. All sub-detectors will be required to verify their FE with DC-DC convertors and Maraton supplies.

Referee: Check radiation levels at locations with power supplies when running at high luminosity (and possible different experiment configuration)

LHCb response: This will be checked with the simulated radiation levels that are now available.

Referee: Copper links between the AMC’s and the switch should be considered with extreme care.

LHCb response: The LHCb Online team will consider this with extreme care.