Alice ITS Upgrade

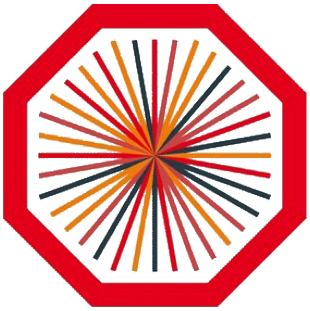

Production Readiness Review

Answers to the reviewers questions and comments

15 May 2018

Summary

This document provides the responses to the questions and remarks the reviewers arose at the ALICE ITS Readout Electronic Production Readiness Review.

1) The pre-production board should be carefully qualified. Although it was reported that it has mostly the same layout compared to the previous version, given that it has a different form factor and that some features (e.g. compatibility with FEAST) where added, it is not excluded that some mistakes were introduced in the layout.

A first, pre-production batch of boards (4 with the option for 5 more) has been ordered, and it is now in production, with the delivery expected at the beginning of June.

2) Verify the optical power margin for the CTP to read-out unit link and possibly provide even measurements of the optical margin with attenuators. Quoted numbers refer to maximum output instead to min.

The optical margin issue has likely been underestimated as we did successfully test with even 1:32 splitting: while we installed enough fibers to go with a lower splitting ratio (down to 1:1), this will require some firmware work on the LTUs to provide standalone synchronization between them (not an issue for "standard" data taking operation). Therefore, we decided to proceed with more accurate optical power measurements together with the CTP-LTU team to assess which is the most effective solution.

3) Tremendous amount of work was presented. Are BER values with real life length cable available? Perform integration tests with the final cable and CRU/CTP. Are eye diagrams to detector and CRU available?

Final cables prototypes just arrived during the second week of May, and BER testing are ongoing before certifying the solution for the final production. Tests conducted so far show no problem for an 8m long cable with an intermediate connection point (to accommodate a patch-panel in the cable path). The GBT team has tested the GBT communication link with various configurations and cable lengths (reference?) and therefore WP10 obtained no eye-diagrams of the GBT link on the CRU or CTP. We did perform BER rate tests over a 300m optical fiber with the link running for several days while sending PRBS patterns in both directions and did not observe any errors.

4) Provide temperature of FPGA and how much power FPGA is using. Are thermal images available?

Not yet, also because to operate safely the FPGA need cooling, either air of liquid, making it difficult to take operational equilibrium heat images. As soon as we will receive prototype cold-plates (already in production), we will do temperature measurements using the internal temperature sensor while operating the board in realistic condition (enclosed in crate with no ventilation); we will also try thermal imaging of the cold plate to verify thermal transfer distribution and thermal interface efficiency.

We used the SCA to read several on-board currents and the two temperature sensors on RUv1, one close to the FPGA, one close to the DC-DC converters. These measurements were taken with one GBT bi-directional link active, the firmware in continuous trigger mode, and the read out of an inner barrel module:

I_MGT (1V):	*0.695 A*
I_INT (0.95V):	*0.726 A*
I_1V2 (1.2V):	*0.504 A*
I_1V5 (1.5V):	*1.311 A*
I_1V8 (1.8V):	*0.705 A*
I_2V5 (2.5V):	*0.658 A*
I_3V3 (3.3V):	*0.519 A*
I_IN (total board, 4.5V):	*2.338A*
T1 :	*42.8C*
T2:	*39.6C*

5) Will optical transceivers be cooled? Is there airflow?

At the moment there is no cooling foreseen for the optical transceiver. No air-flow is foreseen within the crates, as they are installed in a magnetic field. In all our testing no problem has been observed, yet the boards never operated in a "crowded" installation like in the final crates. We plan to proceed with thermal imaging of the transceiver in operation to verify the operational temperature margin we have in the present configuration.

6) As the layout is changed from V1 to V2, why is it not considered to make a clock path to the detector without going through FPGA in order to bypass the radiation induced increase of jitter?

The jitter increase is actually due to the scrubbing activity, and no jitter increase has been observed even in FPGA irradiated with more than 100 kRad (those used to firmware qualification), way in excess the target radiation level of 10 kRad. In any case, the jitter issue was solved by not using the MMCM tile for clocks to the sensors and GBTx, relying only on simple combinatorial logic and clock dividers, which have been tested to not be affected by the scrubbing activity.

7) It was noted that at that stage of the project still a number of different options are kept open. The review panel would appreciate if these options would be decided quickly (FeastMP, jitter cleaner). Is the jitter cleaner needed?

The FeastMP option has been dropped. The jitter cleaner was there because the US transceiver jitter specifications are not met by the GBTx phase adjustable clocks. It should be noted, however, that these specifications are likely for the highest transceivers speed, not the quite "slow" 1.2 Gb/s we use on RU. The latest measurements of the eye diagram from the transceiver anyway could still hint at the GBT clock affecting the transceiver performance, even if only in a subtle way. As the jitter cleaner is a non-critical component and the fact that we are still not able to rule out some influence of the GBTx clock jitter on the transceiver performance, it has been decided to retain it.

8) Power board seems to lag behind the read-out board in design maturity and is possibly less production ready. Please comment.

The power board is fully advanced to the state of production readiness. As was shown in the review, the production prototype board has been produced and fully tested and characterized. Test results show that it meets the requirements in all aspects and all of the components used have been radiation tested and qualified. The existing production design (same components and circuit topologies but in a different board form factor) have been in use at multiple ALICE ITS upgrade sites for testing modules and staves for 9 months and have been performing as expected and meeting the requirements.

9) Present an updated production/test plan and strategy.

The production plan foresees the production of the RUs by autumn this year, compatible with the present components availability problem in the semiconductor market. While key components (FPGAs, actives, ICs, etc.) have already been purchased/secured, it could be that lead times of quite standard passives like capacitors could negatively affect the production schedule. We are currently waiting for market searches results to better evaluate component availability: three scenario have been identified:

1) No major shortage of components: we go as planned, i.e. the mounting company procures the components. Key components have already been secured.

2) Few components have long (>> weeks) lead times: we pre-buy them asap, so to have them ready for production, and then ship them to the mounting company awarded the contract when production will actually start.

3) Few components (likely capacitors) have impossibly long lead times (>> months): we will replace these components with available ones, which will require mounting 2-4 additional prototypes to verify compatibility before launching production.

We plan to mix plans for scenarios 2) and 3) according to the market search outcome.

The testing, following the suggestion of the reviewers, is being re-organized, delegating to the manufacturer only the industry-standard verifications (bed-of-needles, traces electrical integrity, x-rays of BGAs, etc…) and moving the board functional testing to the collaboration. Acceptance testing for all ITS boards will happen at Nikhef, where on-purpose manpower has been allocated. Acceptance testing will include power rails and basic functionalities verification. Full board features testing will be distributed to sites equipped with back-ends and sensors to verify the entire IOs and processing capability of the boards.

To the purpose of both acceptance and in-depth testing, a specific loop-back and test board is currently under development at Austin and Nikhef, with first prototypes of the loop-back board already in use. The test board is based on a commercial FPGA dev-kit with an adapter board to interface with the RU. The test firmware will be an evolution of what currently used to test RUv1.x boards.

10) Have the boards been tested in magnetic field?

Not the whole assembled board, only the DC-DC have been tested. The possibility of magnetic field full board testing is under discussion (to assess manpower necessary for an effective testing), especially since the DC-DC are considered the only sensitive part, and these have already been extensively tested, as reported in the review.

11) Will during the production and before acceptance thermal cycles sequences be conducted?

Executing accelerated lifetime tests and thermal cycle tests before acceptance or as pre-series validation has been considered by the project. The final procedure will depend on availability of specific testing resources.

12) The reviewers suggest to reconsider the choice of the Microsemi ProASIC3 FPGA device. Since the current firmware already takes 1/3 of the resources of the device without TMR on all blocks, maybe it would be wiser to consider using the pin-compatible larger device, even if more expensive. Usually the requests on FPGA resources increase during the life of an experiment, so starting with an already fully exploited device should be avoided. If it is considered impractical to use the large device a possible approach would be to implement and test the full functionality with TMR and to reassess the resource utilization before launching the production.

As the only task of the Microsemi FPGA is scrubbing, it was determined that the A3PE600L device is up to the job, and offers enough room for triplication of the vulnerable parts of the design. Further considerations supporting this decision:

- *In all the radiation tests so far the system failures due to the Microsemi experiencing an upset are not significant compared to other failure modes, and can be recovered by simply resetting the Microsemi device, without affecting system data taking.*

- *It has been verified that there is room for protecting almost all firmware blocks with the A3PE600L resources (about 52% utilization with TMR), and the firmware is being further optimized after the first development phase.*
- *The main FPGA (XCKU060) offers plenty of resources for system upgrades in the field, which is indeed foreseen during the lifetime of the experiment.*

Considering all of the points above, the cost penalty of going with the bigger A3PE3000L (1000 USD vs 250 USD per device) was deemed not worth the potential performance gain.

13) Please clarify some details of the readout system test under radiation planned at CHARM? Is it planned to use a fully-equipped VME crate or just one RU board?

The planned CHARM test beam will use a complete system mock-up, including:

- *Readout Unit*
- *Power Board*
- *Actual sensor modules/staves*

The three items above will be in the radiation area, and the connection with the back-end will be through the fiber-optic connections, exactly in the same configuration as in the experiment. We plan to only use one RU for this test (one full readout chain) as it is destructive.

14) It would be good to perform integration tests with the CTP and CRU

The CTP team recently tested a LTU prototype equipped with 10 SFP, which will also be available for our trigger implementation tests. The CRU connection has been tested using an Altera development board with the same FPGA as the CRU (CRU emulator) implementing the same firmware and optical communication (GBT) as the final CRU; therefore we do not expect to see any difficulties when migrating to the final CRU.

What could be a challenge is the clock distribution management, as we have two different GBT links providing a reference clock (from CTP/LTU and from CRU) and we will rely on the CRU clock as our master clock. Therefore, we are planning a test involving the CTP/LTU and the CRU both communicating optically with the RU, in order to verify that we can correctly recover trigger tokens into the CRU-clock domain.