Upgrade Cost Group Review of the EF amendment to the Phase II upgrade of the ATLAS TDAQ[1]

C. Biscarat, R. Calabrese, R. Mankel, D. Newbold,
AJS Smith (chair), P. Vande Vyvre, W. Wisniewski

Process

In response to fundamental changes to their initial TDR for the ATLAS Phase II TDAQ upgrade, ATLAS submitted an amendment to the TDR on October 14, 2021. The scientific-technical review for the LHCC and the UCG review of cost, schedule, resources, management and risks were then conducted in parallel by the above panel. All review sessions took place remotely because of the pandemic. The TDR and supplemental UCG Cost Package were presented to the panel at a "Kickoff meeting" on October 21, which was followed by a round of questions and answers that were discussed at an "iterative" meeting on December 7. The final review took place January 20, 2022[2], with presentations by the EF project coordinators and several of the leaders developing the project, including three younger scientists. We also met with ATLAS and TDAQ project management to discuss funding and resources, and reviewed the confidential money matrix.

Note: *The review took place a few weeks before the new HL-LHC schedule was approved. We did not update the report on the review, because it already takes into account the possible impact of future schedule changes, and because there could be additional schedule changes over the course of the HL-LHC program.*

Overview

The ATLAS Phase-II upgrade of the TDAQ, combined with the Phase-I upgrade, will lead to the almost full replacement of the original ATLAS trigger, readout and dataflow systems. In the TDAQ-EF Amendment we are evaluating, the originally-proposed hardware tracking trigger system (HTT) will be replaced by a software-based Event Finder that exploits advances in computing technology and tracking algorithms to retain or even improve the performance of the original system, significantly reducing cost, need for firmware, etc. The new system will also allow greater flexibility in responding to new technologies.

General Observations

The TDR for the Phase II upgrade of the ATLAS TDAQ – EF amendment represents a very positive development. The elements for reconsidering the original baseline of the HTT are solid and convincing: they cancel the option to evolve to a L0/L1 system, improve software tracking by a factor 8 thanks to some updates of the ITK geometry and performance improvements, ease interactions with off-line computing, and show potential for significant cost reduction.

The new baseline of the Event Filter will consist of a commercial system which allows the inclusion of accelerators, where currently FPGAs and GPU accelerator cards are being considered. In the future it could evolve to hybrid computing units with multiple types of CPU/GPU/FPGA cores implemented in the same chip and make use of a shared high-speed memory. The EF tracking code will be based on the new software framework ACTS developed mainly by ATLAS and other experiments or projects. The flexibility of the new baseline makes it better prepared to take advantage of advancements in technology and a rapidly evolving market.

[1] CERN-LHCC-2022-004 ; ATLAS-TDR-029-ADD-1
[2] https://indico.cern.ch/event/1116341/

The project includes a large multi-platform development of algorithms. This is a major undertaking in general, but even more with the introduction of the heterogeneity with FPGAs and GPUs. Ultimately, given the project timespan and the possible evolution over time a single code base usable from different hardware platforms is a must. It is reassuring that this has been stated as an explicit goal of the project.

Organisation and Management

The EF project is well organised, and the senior management team is strong and well integrated into the overall TDAQ upgrade project. The leaders of the subsystems are all in place, and positions are filled down to level 4. The strength and coherence of the team was clear from the excellent presentations during the review. We see no issues of concern.

Cost Situation

The total estimated CORE value of the event filter system (EF) is 8.8 MCHF, approximately 12 MCHF less than that allocated to the EF in the original TDAQ TDR. The baseline system consists almost entirely of commercial commodity CPU servers, with an additional 3% to cover various compute services. If ongoing studies conclude that a heterogeneous system with accelerators turns out to be a better choice, its cost will fit within the 8.8MCHF envelope. Finally, the estimate assumes that M&O funds will be available to replace the servers of the existing Run 3 EF farm.

The main uncertainty in the cost lies in the extrapolation of the cost of CPUs over the next decade: the purchase for run 4 will be made in 2027; for run 5 in 2032. To produce their estimates ATLAS has used the "standard" CERN cost-reduction factor of 15%/year and has set the quality factors at QF 3 to take into account the uncertainties in this factor.

The preliminary money matrix indicates a good matching in general with the required funding, although further negotiations are needed to secure the full CORE budget needed. Where possible the funds saved by going to the new EF design will be assigned to other ATLAS upgrade projects.

Schedule

Detailed schedules have been made available for all the TDAQ subsystems down to Level 6 and reasonable contingencies of the order of one year have been allocated for most deliverables. In the original schedule the HTT had the shortest float and the largest number of risks with possible schedule impact. The new plan helps very much to reduce this risk, and frequent milestones have been defined to help ATLAS and the P2UG to effectively track progress.

The schedule for installation, integration and commissioning will be particularly critical given the short time available and the many dependencies with and within TDAQ system. We encourage the TDAQ team to proceed with the early planning activities already started.

Resources

Time profiles of person-power needed for each professional category have been evaluated with a bottom-up approach and it seems overall reasonably consistent w.r.t. expected person-power available. The situation has been greatly eased by the move away from specialised hardware-based tracking. We continue to support the approach taken by TDAQ to clearly identify tasks and non-CORE deliverables related to software and firmware developments and to associate them to specific Institutional responsibilities to be defined through the on-going MoU negotiations. To maintain the necessary level of redundancy and continuity throughout the project lifetime it is especially important that sufficient technical experts are well identified and secured in this process.

Risks

Schedule: The move to commercial commodity CPUs has significantly lowered the overall schedule risk to the project. Dependence on specialised firmware is greatly reduced, and the flexible framework creates opportunities to exploit new technologies if appropriate. The major remaining risks include delays in deliverables from other subsystems, including simulations and off-line software; and the inevitable loss of experts along the way. ATLAS is mitigating the latter by forming and training teams so that younger people can step up if senior people leave.

Cost: Uncertainties in the extrapolation of CPU costs over time represent the major cost risk (and possible opportunity). The 15%/yr factor is reasonable but subject to factors like external demand and supply-chain problems. There are also technical issues that could increase the cost such as excess power demands, lower performance of the software, etc. Should these risks materialise, the cost increases are relatively modest and should be manageable.

Opportunities to reduce cost:

1. Significant delays in the HL-LHC schedule represent a possible or likely opportunity to reduce the cost of CPUs and other commodity components. (**Note:** *such a delay indeed took place after the review –accordingly, a new cost estimate, and/or a reduction in risk might be useful at this point.*)
2. Better than anticipated performance of, and/or improvement in tracking algorithms could reduce the number of CPUs required.
3. Emergence of new technologies could lead to less expensive systems.

Conclusion:

In summary, ATLAS has carefully studied and evaluated the significant risks to the project, and has produced credible mitigations. We therefore recommend Step 2 approval by the RB and RRB to allow resources to become available and MOU's to be signed.