CERN –LHCC-2018-008; UCG-029 April 17, 2018

UCG Report on the Phase-II Upgrade of the ATLAS TDAQ System[1]

V. Beckmann, U. Behrens, A. Boehnlein, M. Convery, F. Forti, V. Gligorov, M. Huffer, T. Kuhr, R. Mankel, L. Moneta, M. Morandin, D. Newbold, F. Simon, A.J.S. Smith, P. Vande Vyvre

Process

In accordance with the Phase-II TDR approval procedure, on 27 January 2018 the ATLAS TDAQ team made available the specific documentation set required for the review of the cost, schedule and management aspects of the TDR ("UCG package"). The material provided was presented to the UCG during the "Kickoff meeting" held on 2 February, soon after the scientific-technical review of the TDAQ TDR had taken place.

Over the following two weeks the UCG examined the documentation provided by ATLAS, leading to the formulation of a set of questions that were submitted to ATLAS on 13 March. The answers, delivered in written form on 19 March, together with an update of the risk registry, were discussed during a Vidyo meeting held on 22 March. Further questions were submitted to ATLAS on 5 April and the final version of the UCG package was released by the TDAQ team the following day.

On 12 April the UCG review took place at CERN with a four-hour long plenary session, followed by three parallel session tracks that gave the UCG sufficient time to examine and discuss all the relevant aspects of the project with the L2 and L3 managers. In the same day the core UCG team met with ATLAS and project management to discuss the situation of the money matrix. The revisions of the TDR were considered and accepted by the LHCC panel, finalizing the Scientific/Technical approval process.

Project Overview

The Trigger and Data Acquisition System (TDAQ) Upgrade is an essential part of the ATLAS upgrade Phase-II program. Combined with the Phase-I upgrade, it will lead to the almost full replacement of the original ATLAS trigger, readout and dataflow systems and to the introduction of a new hardware tracking trigger system (HTT) embedded in the Event Filter system.

It's a huge and complex project, articulated in three systems (Level-0 Trigger System, Data Acquisition System and Event Filter), twelve L3 subsystems and 39 L4 "deliverables" that involves more than 50 Institutions for a total estimated manpower effort of nearly 700 FTE-years.

A very large amount of development work will have to be carried out, starting with the very first internal specifications reviews in 2018. However, in many cases, the new system can be considered as an evolution of the current system and of developments being pursued in the context of the Phase-I upgrade. This holds true in particular for the HTT system, that represents the most challenging part of the upgrade, where the experience with the current FTK system has turned to be of paramount importance, in particular for the many lessons learned so far.

[1] CERN-LHCC-2017-020 ; ATLAS-TDR-029

General Comments

The UCG was impressed with the details and clarity of the presentations and has much appreciated the very good quality of the documentation and the quick and accurate responses given by the team to reviewer's questions.

All issues raised had been satisfactorily addressed with a significant amount of additional information submitted or presented together with the several changes introduced meanwhile in the TDR.

Cost Situation

The total estimated CORE cost of the upgrade amounts to 44.9 MCHF, in line with the estimate in the scoping document (43.3 MCHF), despite the important evolution of the project that has taken place meanwhile.

The largest contribution to the costs comes from commodity components like FPGA (more than 30%), servers (~20%), networking devices, etc., that will have to be purchased in 5 or more years from now, and most cost estimates have been ranked with QF3. In fact, several external factors contribute to the uncertainty of the estimates, like possible variations in the technology evolution trends, in market conditions and in currency exchange ratios and therefore significant cost deviations, in the TDAQ more than in other upgrade projects, may occur at the time of purchase w.r.t. the CORE estimates.

We note in particular that for FPGA procurement, reduced cost available only through centralized ordering have been assumed in some cases and that there is a significant risk that, in the event of a less than optimal outcome to tendering, actual costs may eventually turn out much higher.

Three quarters of the expenditures will take place in only two years (2023-24) and therefore good planning and management support will be needed to identify potential risks associated with such a compressed schedule, develop a mitigation plan and successfully execute the procurement.

For what concerns the comparison with institutional aspirations, there is a good matching in general with the required funding, although further negotiations are needed to secure the full CORE budget needed.

Schedule

Detailed schedules have been made available for all the subsystems down to Level 6 and reasonable contingencies of the order of one year have been allocated for most deliverables; however no critical path analysis has been performed yet. The system with the shortest float and the largest number of risks with possible schedule impact is the HTT so any opportunity to recover development time should be proactively pursued in the coming years.

Frequent milestones have been defined to effectively track the progress of the activities and we commend the decision by ATLAS of including the firmware and software components in the review process of the electronic boards and modules.

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 profile of manpower needed for each professional category has been evaluated with a bottom-up approach and it seems overall reasonably consistent w.r.t. expected manpower available. However the mismatch between skills required and the professional categories used in the Institutional survey makes it difficult in some cases to verify that there is a substantial correspondence between needs and availabilities. Remaining shortfalls seem to affect in particular the HTT system and require therefore to be proactively addressed.

The UCG believes that the single most critical element for the success of the project is related to the availability of adequate qualified technical expertise for firmware and software developments. The UCG therefore strongly supports 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.

It's especially important that sufficient resources providing key technical expertise with the necessary level of redundancy and continuity throughout the project lifetime, are well identified and secured in this process.

In addition, the UCG encourages the TDAQ team to consider the design approach to firmware development shown by the Global trigger as a model for any TDAQ program which has a dominant firmware element. Such approach deals directly with the issue of developing and maintaining a system containing a large firmware component and which addresses the further issue of leveraging a larger pool of manpower to aid in its development.

Risks

The TDAQ team has developed an adequate Risk Management Plan and a comprehensive set of risks have been included in the risk registry.

Very preliminary risk analysis outcomes were reported during the review regarding the global cost impact but no similar evaluation was performed on the schedule impact due to lack of adequate tools to correctly handle risk correlations.

We encourage ATLAS to address the problem so that proper assessment of risk impacts can be carried out and tracked through the project life-time.

L0 Trigger System

The Level-0 Trigger system is made up of the Central Trigger, L0 Calorimeter Trigger, L0 Muon Trigger and Global Trigger.

In general, the development program envisioned seems to be well characterized and understood. The members of the teams, besides being highly experienced, appear to be well integrated, highly organized and exhibiting a very good understanding of their external interfaces.

The costs of the system are adequately enumerated and seem reasonable for a project of this complexity. The schedule produced is credible and exhibits a good understanding of all schedule dependencies.

The UCG noted that a well developed and focused simulation program was presented and already underway.

Central Trigger

The Central Trigger System consists of the Central Trigger Processor (CTP), the Muon-to-CTP-Interface (MUCTPI), and the distribution of Trigger, Timing, and Controls signals (TTC) to the sub-detector readout systems, including the Local Trigger Interface modules (LTI).

A very clear and convincing presentation of the project was given including the resources needed and the schedule. The management structure is simpler than other parts of the TDAQ upgrade because it is under the responsibility of a single institute (CERN) and involves two groups which were already in charge of the CTP and TTC in ATLAS since Run 1.

The project follows an evolutionary approach by replacing some components of the system at each upgrade. The UCG seconds the decision by the team of advancing a key milestone of the project involving both trigger and DAQ components to Q1/2020, a seminal step for demonstrating early enough that the DAQ hardware can sustain the peak performance needed for the Phase II upgrade.

The schedule includes the availability by the end of 2022 (in time for Run 4) of key equipment needed for the detector tests in the lab or in test beams. The project should be in close contact with detector groups to be informed of the detailed need for each of them and we suggest that the relevant equipment delivery times are included in the project schedule.

The software will be developed by the two groups involved. We deem important that the software development schedule is decomposed into its different packages and their delivery schedule properly monitored by the project.

The team is well organized and has the adequate manpower and competences. The budget is based on well-established estimates and the schedule presented is credible. There is no risk identified which should represent a serious threat for the successful completion of the project.

Calorimeter Trigger

The architecture of the Calorimeter trigger is established during Phase I upgrades, LS-2 and Run-3. Its architecture emphasizes modularity, allowing new functionality to be added incrementally. For the high-luminosity upgrade all pre-existing modules are reused and simply, only a new module is incorporated called the Forward FEX (fFex). From a hardware perspective the system implementation is heavily dependent on ATCA and fiber-optics running serial protocols at high-speed.

As its architecture is being currently implemented, experience gained before the Hi-Lumi epoch will prove out the majority of the design. Consequently, much of the development risk will be "burned down" by the time of its incorporation into the new ATLAS trigger infrastructure. This is judged to be a prudent and conservative approach to the system's implementation.

The fFex is itself based on a pre-existing ATCA board. Consequently, the major engineering deployment will be in the development of its firmware. Again, this is a prudent approach as it minimizes the amount of hardware to be produced and tested.

The team has considerable ATCA experience and is adequately staffed with the engineering and physics experience to develop the necessary hardware and firmware.

The amount of platform software required is minimal and well understood by the team. It exists mainly in the DAQ/Felix interface.

Muon Trigger

The Muon Trigger system is in solid shape for this point in the project's lifecycle. The key points are that the project is significantly over-resourced in every category of personnel, and that it is building on a solid base of expertise with the current Muon Trigger system. The management structures and the project planning seem solid and appropriate to the task.

The system is designed to deliver the required physics performance improvements to allow ATLAS to trigger on muons with acceptable rates in the HL-LHC period. It is worth noting that these improvements come not only from improvements to the Muon Trigger but also from the installation of several new muon detectors during both LS2 and LS3, so there are shared risks which should be kept in mind throughout the development. In particular, there are worries related to the NSW processor, namely that the problems with the Phase-I NSW development might lead to knock-on delays to the Phase-II developments. A specific risk is if certain people who are expected to transition from the Phase-I to the Phase-II NSW work are held back by developing or commissioning the Phase-I NSW system. This is however mitigated by the fact that the NSW is not the core part of Muon Trigger and the fact that the Phase-II NSW developments will benefit from experience gained with Phase-I.

As with many other parts of the system, the cost driver are the FPGAs. It is reassuring that the TDR presents a clear R&D programme to address the remaining design choices before the PDR/FDR, but it is clearly important to keep an eye on how the FPGA resource usage and latency estimates evolve during this process.

Global trigger

The design goal of the global trigger is to bring "Event Filter" like capabilities to a Level-0 trigger system. The design approach is modern, and perhaps even unique within our community, since it emphasizes a firmware rather then hardware-centric architecture. As such the team will build a single common hardware module to support the bulk of the engineering which is now reduced to firmware, thus leaving to the amount and complexity of the necessary firmware its significant design and implementation challenge.

To isolate the algorithmic code from the mechanics of input and output flow the team proposes to provide a "framework" into which all algorithmic code would execute. The advantages are two-fold: First, this makes the bulk of the engineering much more accessible to a wider pool of implementors who no longer require all the specialized knowledge and experience typically required to program an FPGA. Second, this allows their system to evolve and grow both incrementally and gracefully in a modular fashion.

To facilitate the rapid implementation of firmware the team proposes to build an intermediate hardware module called the firmware Deployment Modules (or PFM).

From a hardware perspective the system implementation is heavily dependent on ATCA and fiber-optics running serial protocols at high-speed. Some of the link speeds envisioned (> 25 gigabits/second) are at the current state of the art.

The nature of the Global Trigger hardware architecture is such that it exhibits very strong scaling properties, thus providing significant design margin for meeting their performance requirements as well as providing a system which may naturally evolve as requirements change. This again is an excellent approach.

The team has considerable ATCA experience and is adequately staffed with the engineering and physics experience to develop the necessary hardware and firmware.

The amount of platform software required is minimal and well understood by the team. It exists mainly in the DAQ/Felix interface.

DAQ

Management structure of these projects is well in place with key positions identified and filled. The team has solid experience in the relevant aspects from previous runs. Risks in these subsystems are well identified and mitigation strategies are in place. The fact that the work is done in a fast moving field of computing hardware evolution is taken into account by on-going market surveys and testing of viable cost-effective solutions. These on-going activities are accompanied by formal design reviews.

For the Readout subsystem a significant risk is related to the possibility that the chosen FPGA cannot achieve the required performance and this could indeed result in substantial cost increase.

The Front-End LInk eXchange (FELIX) is a crucial element of the ATLAS TDAQ, because many subsystems are connected to it and rely on its timely installation. However, the FELIX is a modular system and can accommodate rather easily different requirements. In addition, the prototype Phase-1 FELIX can be used in demonstrator tests and in case the Phase-2 FELIX is delayed.

Event Filter

The upgraded Event Filter comprises a farm of software processors running full offline algorithms and the new hardware-based HTT subsystem to whom the most demanding pattern recognition and track fitting tasks are delegated.

The farm is a low-risk subsystem since it is based on scalable commodity hardware and its installation is not dependent on time critical developments. 70 FTE-years of effort is foreseen for the development of the software that is planned to be carried out mainly by scientists with a very limited contributions of software engineers. The UCG noted that 38 institutes are involved in the project, with the manpower rather evenly split over the partners. In this context it is essential to have in each working group identified experts with long-term or permanent contracts. The project needs to clearly identify the key personnel for covering the software development work and map it to the tasks in the Work Breakdown Structure (WBS), in accordance with Institutional commitments to be negotiated with the Funding Agencies.

The project will decide on the Processing Unit (PU) technology rather late in the schedule at the beginning of 2025. This does not present a major risk, because the possible use of accelerators is seen as optional. However, the review panel sees opportunities by getting specialized engineers involved in the process, in order not to miss potential optimization and acceleration of EF code and cost saving in this context.

Hardware Tracking Trigger

The HTT is a novel and extremely technically challenging project, which is nevertheless vital in obtaining the necessary cost / performance ratio for the ATLAS trigger system. The existence of alternatives to the HTT development was fully explored during discussions, and strong evidence presented that, despite some risks, the HTT was the most effective option. However, the collaboration should keep open, as a risk mitigation strategy, the option to fall back to a more conventional software-based high-level tracking trigger, in the case of delays to HTT delivery or its commissioning in 2026. This implies the existence of at least a minimal parallel project to produce and optimize track reconstruction software suitable for use in the EF with the new ATLAS tracking detectors. There does not appear to be a cost-effective alternative to the HTT in the longer term.

The HTT project is complex, comprising a large number of sub-projects and the efforts of almost 40 institutes. There are many interdependencies between deliverables. The schedule is very tight, with only around six months of float overall. The collaboration has taken steps before and during the UCG review period to introduce into the schedule sufficient integration exercises and pre-production tests (of both ASICs and electronics modules) to minimize the risk of schedule slip or technical failure. Nonetheless, a comprehensive and proactive project management regime is mandatory to ensure timely delivery.

The HTT CORE cost has been broken down sufficiently to verify the appropriateness of the individual items. The cost estimates are well-founded, though as in other components of the project, the cost is dominated by FPGA prices, which are assessed at QF3.

The basic concept, and some components, of the HTT are directly evolved from the current ATLAS FTK project. The FTK is extremely late compared to the originally proposed schedule, with the collaboration having identified delays in firmware development and commissioning as the main issue, in some cases due to loss of expert personnel. The lessons from the FTK appear to have been fully absorbed by the collaboration, the HTT architecture significantly simplified, and the estimate of the required expert resources increased. The overlap between institutes responsible for FTK and HTT is not large. However, this reinforces the need for a strong central engineering management organization, with maximum commonality in hardware, firmware and software components across the TDAQ project. The collaboration responded positively to this observation, though in the UCG materials, the necessary management positions have not been identified.

Although the assessment of expert human resources for HTT appears reasonable, it is not clear that all the necessary resources have been identified within the collaboration. This is particularly noticeable in the area of software development, where it was difficult to obtain a clear estimate of the need for professional software engineering effort, or its availability. This is a critical issue for the HTT project, that must be addressed proactively by management in the immediate future.

The HTT project will be challenging for the collaboration. We suggest that further, incremental, steps in project organization and planning should take place before the MoU phase. In particular:

- The need for strong and dedicated HTT project management team, including necessary engineering management, cannot be overemphasized. The collaboration should seek to ensure that the strongest possible management is in place, with all necessary resources available.
- The ATLAS management should proactively seek to renegotiate contributions to the TDAQ project, such that any lack of expert effort within the HTT is addressed.

Conclusions

We recommend Step 2 approval by the RB and RRB to allow resources to become available and MOU's to be signed.

Vigorous oversight of project, including external reviews, is essential going forward.