NSF Review of ATLAS High Luminosity Upgrade Construction Progress

July 23, 2024

Panel Report

1. **Construction progress:**
   a. Review Rebaseline panel recommendations and comment to NSF on whether they have been satisfactorily addressed.
   b. Examine the status of production or pre-production for all major MREFC funded subsystems.

   a) Comments: All Rebaseline panel recommendations are specific to detector components and will be addressed in the corresponding sections. In general, we note that the various L2 manager presentations have included clear summaries of the action taken in response to the recommendations. A strong effort to address the concerns about preservation of technical expertise, as some key technical staff are expected to retire, led to several new hires in the relevant areas. In addition, comprehensive documentation on the hardware and firmware have been developed, and the needed attention to QA/QC procedures has been shown.
   b) Comments: Significant progress has been achieved in all the major components within the scope of the NSF MREFC: more than one half or the PRRs needed for the various subsystems will be completed within the CY 2024. The remaining ones are scheduled for CY2025, with the exception of the EF Tracking Trigger, scheduled for Q3 2027.

2. **Project Management and Project Governance:** Examine and comment to NSF on:
   a. Project management functions, as outlined in the latest PEP, including project controls and financial reporting and subawardee management and oversight, representative change control actions, and recent Earned Value Management Reports and advise NSF on their use as effective management tools.

   The project established formal project management processes years ago and since their definition it has been successfully exercising these processes. Besides the well-developed definition of project management, captured in released documentation (PEP, CEP, Risk and Contingency Management Plans, Configuration Management Plan, Quality Assurance Plan, among others), corresponding practice is also widely accepted by the project team.

   The US Atlas HL-LHC Upgrade Project office includes experienced members from Columbia University and, most recently, a faculty from MSU to monitor risk

management. All the deliverables within the scope of the NSF MREFC are led by experienced L2 managers from institutions supported by NSF. The project control is exercised by an integrated NSF/DOE team.

ATLAS is a very well-managed project which employs the best-practice tools necessary to successfully complete its NSF award. This was demonstrated from the outset in the PI's talk which provided an excellent overview of the accomplishments, challenges, opportunities and plans to complete the project. The PI's clear and well-supported overview served also as the template for each subsequent talk making the telling of the ATLAS story consistent and easy to understand. As a result, the panel was provided ample evidence of the excellent way ATLAS is managed and gained confidence in the project's leveraged use of the key functions and systems essential to successfully execute an NSF MREFC award.

b. Consider the effectiveness of the management and oversight tools and processes that the project has in place. Advise NSF on whether they are verifying progress, identifying issues, communicating effectively, and informing corrective actions.

Project Management is clearly in control of the 34 subawardees' activities. The provided examples indicate the effectiveness of change control, EVM reporting and analysis, and risk management.

The EVM system is in compliance with NSF's implementation of federal government standards. The project is expertly using these tools to measure status, advise decisions, formalize and document necessary corrective actions. The panel notes that an EVM system surveillance review of the project, performed by NSF in 2023, found no issues. Further, the Brookhaven EVM system is subject to ongoing oversight by the Department of Energy, providing confidence in the reliability of EVM reports.

c. Risk management process, including completeness of current projections of risk and potential mitigation costs, and the adequacy of the contingency budget and schedule to mitigate future risks. Advise NSF on the completeness of the risk register in identifying currently foreseen threats and opportunities with appropriate probabilities and estimated cost and schedule.

The risk register and estimation uncertainties are properly integrated with the Resource Loaded Schedule, enabling Monte Carlo analyses of cost and schedule, including contingency.

The risk register is regularly (at least quarterly) reviewed and updated as needed.

d. Adherence to the QA/QC processes for all major components. Advise NSF on whether they are being implemented as intended.

Scrutiny of the QA and QC is integral part of the international ATLAS review process that the various components need to pass in the course of the major reviews: QA of all the components of a given deliverable needs to be demonstrated in the course of FDR reviews, QC procedures are scrutinized in PRR reviews. More details will be addressed in the discussion of specific subsystems.

3. Plans for FY 2025

a. Examine and comment to NSF on the realism of plans for the coming project year (10/1/24 – 9/30/25) based on the lessons learned from construction and testing so far.

b. Examine the forecast risk-adjusted expenditure and obligation profile for the coming fiscal year and advise NSF on the adequacy of its substantiation. Make recommendations for modification, if appropriate.

The $16.8M requested for FY25, including $5.3M contingency is informed by prior history and substantiated with the technical, cost, schedule and risk detail necessary to give the panel confidence that they were realistic and successfully executable.

Currently available contingency for the remainder of the project is $17.2M, in good agreement with the estimated contingency of $17.5M at 90% confidence level (CL). It is ~50% of the remaining work. This seems to be ample, and provides an opportunity to aggressively address issues as they arise.

c. Does the project show appropriate progress since rebaselining?

The project showed significant progress in several key components since rebaselining. While some milestones are noticeably lagging, resulting in $3.8M negative schedule variance at project level, in general this does not impact required subsystem floats. The only noteworthy exception is LAr, which has an ~6 months of float deficiency.

During the LAr talk, there were good discussions of ways to possibly mitigate the LAr schedule risk and a possible offset due to a currently unknown but potential slip in CERN's schedule. It was unclear to the panel whether or not ATLAS plans to complete its LAr deliverables as scheduled independent of a CERN schedule slip.

We appreciate the documentation provided on the various DEI initiatives woven in the High Luminosity Construction project that shows a firm commitment to leverage on this excellent effort to disseminate science among underprivileged national and international communities.

4. **Technical Progress**

a. Comment on the adequacy of progress and planning across all Level 3 Work Breakdown Structure (WBS) elements. Identify any schedule or budget related impacts to technical progress and comment on the adequacy of plans and efforts by ATLAS to mitigate their impacts.
b. Describe the status of remaining major risks to the project and their mitigation strategies. For example, the FEB2 schedule, L0MDT effort, and firmware engineering.
c. Are the LAr electronics projects progressing satisfactorily towards FDR and PRR? Are the risks to budget and schedule for the start of production fully understood.
d. Examine the schedule and plans for the production of Tile LVPS.
e. Assess the results from initial pre-production of muon subsystems funded by the MREFC. Is there satisfactory progress with firmware development?
f. Examine the status of construction and firmware development for the Trigger systems covered by NSF scope. Are there any new risks from the change in scope to the tracking trigger?

LAr

Findings

- NSF scope for LAr consists of three areas: 6.4.1 FE Electronics; 6.4.2 FEB2; 6.4.3 BE Electronics
- 6.4.1 FE Electronics
    - The COLUTAv4 ADC ASIC met all requirements and is the final production version. All 50 wafers have been ordered and received. A robotic QC setup at UT Austin is in validation to test the required chips.
    - The required lpGBT ASICs for FEB2 have also been received, and while deliveries of the VTRx+ optical transceivers from CERN are delayed 3 months, this should not affect completion of the FEB2.
- 6.4.2 FEB2 and LSB2
    - The on-board powering issues of the FEB2 have been resolved by design of an additional prototype(v2). ATLAS held a revised PDR in January 2024, and a v2.5 prototype is underway. Tests are expected to fully validate the new design leading to ATLAS FDR in March 2025 and pre-production in FY25.

    - Prototype LSB2 have been designed, fabricated and used to test the interface to FEB2 and the ALFE2 preamplifier/shaper ASIC.
    - Front End Crate System Test at BNL is a key milestone prior to pre-production of the frontend deliverables, LSB2 and FEB2.
    - For FEB2, the need for schedule contingency is evaluated at 213 days, while only 86 days of float are projected in the working file.
- 6.4.3 BE Electronics
    - The v2 prototype of the SRTM was tested with the SRTM tester, but it remains to test integration with the LASP prototype before the ATLAS FDR in February of 2025.
    - Delays arising from non-US scope (LASP FPGA change) have been mitigated with BCPs to speed up QC testing of the SRTM and add firmware engineering for integration and the FEC system test.

Comments

- The FE ADC ASICs have been produced and QC plans are advanced with a PRR planned for September for testing of 40k ASICs at UT Austin.
- Procurement of VTRx+ at CERN is delayed contributing to the negative schedule variance. Recovery is expected as these are delivered to the project, in advance of needs.
- The project responded well to the sole recommendation for LAr Electronics from the rebaselining review, namely, to examine the FEB2 risk from component availability and powering scheme, addressing appropriate mitigations.
- The project team is strongly commended for resolving the powering challenge for the FEB2 by undertaking an additional prototype to validate revised design using CERN bPOL48V, bPOL12V and bPOL2V5 ASICs.
- The remaining float for the FEB2 is short of the estimated need by six months. Substantial risk will be retired after successful tests of the v2.5 prototype, but the risk analysis indicates mitigations are required at present to maintain schedule. A viable plan to increase production and testing throughput by running double shifts was presented, but not yet implemented.
- The front end crate (FEC) system test is a key milestone set by the project to validate the design of the LAr Electronics prior to production.

Recommendations

- The management team should prepare a BCP to address the delays in FEB2 and restore an acceptable amount of float in the baseline schedule prior to November.
- The management team should revise the risk register by November to reflect the residual risk for FEB2 prototyping, for example, a possible need for further revision.

Tile

Findings

- NSF scope of Tile Calorimeter upgrade consists of
    - 6.5.1 Main Board (Chicago) - all 896 Main Boards;
    - 6.5.3 ELMB 2 Motherboard (Michigan State University) - 6.5.3: all 256 ELMB motherboards;
    - 6.5.4 Low Voltage Power Supply (NIU, UT Arlington) - 6.5.4: 128 boards. 50% UTA, shared with WITS (South Africa) and Prague.
- 6.5.1 Main Board: Production Completed on budget (Chicago). Boards completed and stored at CERN 2+ years in advance of CERN need-by date.
- 6.5.3 ELMB-MB Production Completed under budget (MSU). Boards completed and are being shipped to CERN 2 years in advance of CERN need-by date.
- 6.5.4 LVPS: LVPS bricks and LV box - UTA making 50%. Pre-production completed.
    - LVPS - CERN requested a combined UTA+South Africa LVPS PRR. PRR was delayed until South Africa completed pre-production. PRR now being scheduled. Production LV bricks on track to start Q4 CY24. Production of LV boxes on track to start Q1 CY25. All production components are in hand or are on order. LVPS deliverable scheduled to complete with 220 days of float remaining.
    - Safran analysis shows a 90% CL schedule float to need-by date of 70 days. Current baseline schedule has 220 days of float remaining.
- There are no remaining external dependencies for 6.5.4.1 LV brick. External dependencies for LV boxes have been lowered now that ELMB2 is obtained.
- Overall TileCal SPI/CPI is 1.05/1.05.
- Remaining risks are related to components not meeting radiation tolerances and shipping risks. Risks will be retired as components are procured.
- NIU LVPS boxes scope moved from NIU to UT Arlington due to illness of NIU lead. This had a small associated cost/schedule impact (BCP 1080).
- TileCal boards Main Boards and ELMBs will be stored at CERN for an extended period of time. LVPS production is scheduled to complete approximately 1 year before CERN need-by date.
- Recommendations from prior reviews have been addressed.

Comments

- Significant progress has been made in the NSF scope for the Tile Calorimeter. The TileCal team should be commended for completing production of the 6.5.1 Main Board on budget and the 6.5.3 ELMB-MB under budget. Good progress has been made on 6.5.4 LVPS, which is

moving towards its PRR and production. All components are in hand or on order, allowing the team to be able to move quickly into production.
- The TileCal team was able to efficiently move LVPS scope from NIU to UTA with minimal cost and negligible schedule impact.
- While LVPS PRR scheduling has slowed having to wait for South Africa to complete their pre-production testing, there is sufficient float in the schedule to complete production well in advance of the CERN need-by date in Q1FY28.
- TileCal L2 mentioned a concern that deliverables will sit at CERN for an extended period of time. To mitigate environmental impacts from long term storage, the team plans to follow industry standards. The CERN TileCal building (Bat 175) will be used for storage. This building is temperature controlled. Electronics will be stored in sealed anti-static bags with desiccant. The team is developing a plan to randomly conduct charge injection tests on boards and compare to prior QC test data from production. This storage and retesting plan sounds reasonable and should be further developed.

Recommendations

- None

Muon

Findings

The NSF scope in the Muon system consists of

- 6.6.1 sMDT
- 6.6.3 TDC ASIC
- 6.6.4 Chamber Service Module (CSM)
- 6.6.5 L0 MDT processor

Note that in international CMS, the L0MDT falls under the TDAQ system while the other elements fall under the Muon system

The sMDT production is complete with all chambers delivered to CERN, and the test results are good.

The TDC production is complete and enough good chips have been tested for the full needs of the project.

The CSM experienced some delays due to issues with PCB production, but these have been corrected and the new pre-production modules are ready for PRR.

The pre-production phase of L0MDT is expected to continue into 2025, where the PRR may be delayed until Q1FY26. The L0MDT schedule is partly driven by delays in CERN procurement.

The critical path in the Muon scope is driven by the L0MDT project.

Comments:

On the whole, the results from pre-production in the Muon scope are very good, with some projects nearly complete (sMDT and TDC) and some approaching PRR (CSM).

The recoveries from technical delays in the TDC and the CSM have been handled, and these do not appear as schedule risks.

Many of the delays experienced in the MREFC Muon scope are due to external factors, such as delays in the mezzanine board and procurement delays at CERN for L0MDT.

A recommendation of the baselining review was to track and report on the firmware progress for L0MDT, but it was not easy for the committee to assess this progress. At a low level, the development tracking in Gitlab is an excellent tool, but from what was presented, we were not able to see if the progress is meeting the expected schedule. At the high level, there are milestones associated with completing the production version of the firmware, which come too late to guide any formative recommendations. To assess the firmware effort, it would be useful to see a tracking of the lower-level firmware milestones to compare the progress with the anticipated schedule.

The plans for FY25 look appropriate. There are large procurements planned for L0MDT, and completing these will improve the SPI.

Recommendations

The committee recommends exercising the scope opportunity of carrying out a vertical slice test of the sMDTs with production electronics at CERN (already identified as N-SO-6-06-003) by Oct 2024.

At the next review, the tracking of the firmware development in L0MDT should be presented in a way that illustrates more clearly if the progress is following the anticipated schedule or falling behind.

Trigger

Findings

At the rebaseline review the NSF trigger upgrade made a significant change, dropping a specialized hardware based tracking trigger to focus on co-processors. The ATLAS collaboration faces a choice between an FPGA based co-processor design, or the use of GPUs. Much of the US effort is

applies to either scenario. That said, there is significant effort focused on implementing algorithms in FPGAs. In response to a question the L2 manager made clear that the engineering effort working on FPGAs would quickly reorient to GPUs if that was the ultimate choice.

The selection of FPG or GPU represents an opportunity that has been noted in the risk register.

Other elements of the trigger upgrade include fiber optic cable remapping and firmware for the global event processor. In both cases significant progress was shown.

The SPI for the projects showing firmware development were below one. This was reflective of both the complexity of the firmware being more than originally anticipated (particularly for the event processor) and the challenge of retaining firmware engineering talent.

Plans for the coming year, through the end of the project were shown.

Comments

Excellent technical progress has been made across the trigger project. Risks appear to be well understood and in control. The largest risk is firmware engineering, as is appropriately recognized by the group.

The group has carefully thought through how to respond to the choice between GPU and FPGA for the tracking trigger.

The plans for the trigger project in the coming year, and subsequent years looked solid based upon the progress that has been made so far.

Given the ongoing risk of loss of firmware engineering, it might be prudent (if it has not already been done) to create a mapping of how existing talent across the US ATLAS projects could be reassigned to fill a gap if it appears. (For example, are there other firmware engineers with US ATLAS that could step in and fill a gap if someone were to leave the global event processor project).

Recommendations

None.