CERN/LHCC-2024-009
LHCC-158
May 2024

LARGE HADRON COLLIDER COMMITTEE

Minutes of the one-hundred-and-fifty-eighth meeting held on

Wednesday and Thursday, 29-30 May 2024

OPEN SESSION – STATUS REPORTS

1. Status of the Accelerator: Matteo Solfaroli
2. ALICE Status Report: Daiki Sekihata
3. ATLAS Status Report: Adriana Milic
4. WLCG Status Report: Simone Campana
5. CMS Status Report: Anna Benecke
6. LHCb Status report: Giulia Tuci

CLOSED SESSION

Present: F. Alessio, M. Battaglieri, D. Calvo, M. Diefenthaler, I. Efthymiopoulos, X. Espinal, T. Ferber, T. Galatyuk*, I. Gil-Botella, J.J. Hernández-Rey, T. Higuchi*, A. Ianni, M. Kirby, R. Leitner, M. Mangano, G. Merino, J. Mnich, L. Moneta (Scientific Secretary), J. Nagle*, S. Niccolai, B. Panzer-Steindel*, G. Signorelli*, F. Simon (Chairperson), H. Tanaka, A. Weber, P. Wells, C. Young.

All sessions were held at CERN with some participants (*) connected remotely via Zoom.

1. Procedure

The chairperson welcomed the committee members. The minutes of the previous session were already approved by email. The chair welcomed two new members of the committee, Markus Diefenthaler joining the WLCG and ALICE referee teams and Michael Kirby joining the ATLAS and WLCG teams.

The committee continued the discussion of a proposal from the chairperson to improve the layout of the LHCC weeks by moving the Open Session to Monday afternoon starting at 14:00 or 15:00, also allowing a seamless integration with the LHCC poster session happening once per year on Monday evening. The drawback will be that all referees will need to arrive on Monday. The practicality of this change will be explored with all stakeholders.

The chairperson reported on the progress of a new proposed LHC experiment, ALADDIN (An LHC Apparatus for Direct Dipole moment INvestigation), which aims at directly measuring the dipole moments of charmed baryons with bent crystals. The location will be in the IR3, no civil engineering work is needed and, if approved, it is

planned to be installed during LS3. The costs are being estimated, including the required activities from the machine, and an LOI is expected to be submitted in the near future, possibly in time for the next LHCC meeting.

2. **Report from the LHC Programme Co-ordinator**

During March, the LHC successfully completed the recommissioning of the accelerator in view of operations for 2024. The recommissioning fit within the allocated 4 weeks time slot and it included a huge amount of technical tasks, notably including new optics for the Partial Reversed Polarity in IP1. In their final reports, the LHC coordinators acknowledged that up to 3-4 days could have been gained in the recommissioning if nights, week-ends and the Easter holiday week-end were better covered by machine commissioning experts. Such issues may also present themselves in 2025 should the situation of personnel availability in particular during holiday periods is not improved.

Around the same time, a decision on a linear shift of four weeks of the end date of the 2024 data taking period was accepted and a linear removal of four weeks from the start of 2025 data taking period. It was agreed that all four weeks will be dedicated to proton-proton data taking. An optimised schedule for both 2024 and 2025 was drafted and it allows gaining 6 days in the total amount of physics data taking days between both years. The end of data taking in 2024 is therefore currently set to 25 November 2024.

During the months of April and May, efforts were concentrated in finding the optimal configuration for stable data taking while ensuring integrated luminosity targets are in reach for the rest of 2024. Firstly, in order to protect the RF fingers, the current per bunch has been limited to 1.6 x 10^{11} protons per bunch (ppb) at the start of Stable Beams. In addition, the cryogenics margins were measured and limited to ~190 W/half-cell (+/- 5 W/half-cell) with the machine full with ~2350 bunches. It was noted that as of today, the cryogenics system is operating stably at its limits leaving no extra room for more bunches in the machine or more charge per bunch. As a consequence of such limitations, the only allowed train length to maximise the number of bunches and collisions in the LHC was therefore chosen to be 3x36 bunch per injection. The corresponding best optimised filling schemes with such train length allow for 2340 colliding bunches in ATLAS/CMS, 2133 in LHCb and 2004 in ALICE. In addition, the train length of 3x36bpi allows for a sufficiently short Flat Bottom at the SPS and allows for the BCMS (Batch Compression and Merging and Splitting) beam type to be seamlessly produced throughout the whole injector complex and be injected in the LHC. The BCMS type beams have smaller emittances and achieve sufficiently good transmission. The use of BCMS beams was tested in the LHC, providing a net improvement in beam emittance of ~25%. It is currently used regularly in production and, pending final studies of comparison with the standard 25ns beam type, may be used over summer for luminosity production. Better emittance and constant bunch charge results in better beam brightness which increases the ability to level ATLAS and CMS instantaneous luminosity, hence obtaining a net increase of roughly 6-8% in integrated luminosity, depending on the range of emittances observed.

A few issues were observed nevertheless. The two most notable issues are a collimation hierarchy breaking at IR7 when squeezing ATLAS and CMS to 33 and 30 cm β^* and a loss of physics for FASER/SND due to increased background at their detector locations.

A very high background is also observed in the ATLAS AFP. Both issues are understood to be an undesired effect of the new Partial Reversed Polarity optics and both have been tackled with experts and efforts to solve them. The former has been traced to off momentum halo impinging on the secondary collimators and a possible solution has been identified: changing the dispersion settings in the vertical plane. The latter has been tackled with multiple tests with collimators around the IR1 region, notably with various settings for the TCL4, TCL5 and TCL6 collimators. None of the tests improved the situation. As of today, FASER and SND are suffering from more than 2x background compared to 2023 and such background is heavily composed of >1 TeV muons which are degrading the signal-to-background ratio. In addition, higher background implies higher deadtime for their data taking, higher data volume and higher complexity in the emulsion processing. On the topic of FASER/SND emulsion exchanges, the increased background lowered the limit of integrated luminosity before an exchange from ~30fb^{-1} (FASER) and ~20fb^{-1} (SND) to ~10fb^{-1} for both experiments (range is 8fb^{-1} to 13fb^{-1}). Given the new conditions, a strategy for the exchanges of their emulsion was agreed with the experiments, minimising the impact on machine operations while requiring an increased flexibility from the experiments. For 2025, the experiments wish for optics that minimise the background at their detector locations and, independently of the choice of optics, early simulation and possibly an early test in the LHC to validate the choices.

Lastly, in week 20, the luminosity calibration scans (called van der Meer scans) were performed by all experiments. Long fills allowed the program to be completed efficiently and successfully. Some unaccounted overhead and small downtime brought the total time allocated to the scans to ~3.5 days as opposed to the 2 days initially allocated. Given that 2024 is expected to be the highest luminosity production year of the LHC, a precise calibration of the luminosity (<1%) was deemed to be of high importance for the experiment's physics program and therefore the completion of every item in the list was considered critical. However, it is necessary a follow up on the possible estimation for the scan items for the 2025 program, which precedes the Oxygen run hence with a hard deadline, the proton-proton reference run and the Pb ion run, that are both limited in time.

To conclude, given the aforementioned running conditions, the LHC is currently running at ~70% availability with ~50% of time in Stable Beams. The average fill length is ~13 hours and the average turnaround time is ~4 hours. There is therefore a positive outlook that ~110fb^{-1} are within reach in ATLAS/CMS when running at a levelled pileup of 63, ~10fb^{-1} are within reach in LHCb for a levelled luminosity of 1.7 x 10^{33} cm^{-2} s^{-1} and could be even higher in case LHCb can run at the nominal upgraded levelled luminosity of 2.0 x 10^{33} cm^{-2} s^{-1}, and finally ~50pb^{-1} are within reach in ALICE.

- The **LHCC notes** the significant impact of the increased backgrounds due to the Partial Reversed Polarity optics on FASER and SND, and **appreciates** the flexibility implemented in machine operations by guaranteeing an access for emulsion extraction if an integrated luminosity of 13 fb^{-1} is reached before an exchange opportunity due to unplanned downtimes arises.
- The **LHCC recommends** that the LHC experts explore mitigation strategies for the 2025 run with high priority. Alternative machine configurations and/or possible hardware interventions should be considered and their impact on the backgrounds simulated, including, when possible, tests with beam in 2024. The strategy for 2025 must be defined in time for the necessary corrective actions to be implemented during the YETS.

3. Report from the Director of Research and Computing

The Director of Research and Computing (DRC) reported on issues related to the LHC and CERN. The problem with the schedule of LS3 is one of the burning items and it is currently being discussed during the P2UG sessions with the experiments. The DRC announced that an LS3 readiness review will take place in September. The SPC would like to have a risk assessment on the schedule of the upgrade projects from ATLAS and CMS, including how the risks can be mitigated and what the costs are for these mitigations.

Otherwise, the DRC reported that LHC and CERN are operating smoothly. The management has prepared the Medium-Term Financial Plan (MTP) that will be presented at the next Council meeting in June.

4. Test Beams

Following the release of the Injector schedule 2.0, which provides an additional five weeks (35/38 days) of proton beam for PS/SPS and one extra week for PS ion physics, the PS/SPS user schedule 2.0 was recently released. Due to this extension, the beam requests of all reviewed experiments could be scheduled, even when some requests needed to be shortened. Some last uncertainty in the schedule in September/October is linked to user change requests caused by unavailability or delays of experiment setup components. The ion run in November is currently being finalised, with CIMERA (T8, IRRAD), NA60+ (H8) & NA61 (H2) already scheduled.

5. General Comments

The following comments are applicable to more than one project.

The LHCC has heard reports by the ATLAS and CMS P2UGs, with further details given below in the experiment section. The Phase II upgrade projects of ATLAS and CMS continue to make good progress, with projects transitioning to pre-production and production. However, technical problems resulted in additional delays, which have eliminated the contingency on some critical subsystems, such as HGCAL and Tracker for CMS and ITk Strips for ATLAS.

- The **LHCC takes note** of the existing situation of the upgrade projects and **observes** that significant risks still remain, making further loss of schedule likely. The LHCC **recommends** that all stakeholders engage timely in a global discussion that takes into account the situation of the upgrade projects of the experiments and the machine, the importance of delivering the upgrade, and all consequences of a possible change of the LS3 schedule. Such a discussion will require a careful analysis of the remaining risks in the upgrade projects, and an evaluation of possibilities for acceleration and the associated resource implications.

ATLAS has reported some stability issues with the lpGBT ASIC when operated at temperatures below -35 degrees C. This component is widely used across the HL-LHC upgrade projects.

- The **LHCC recommends** that the observed stability issues of the lpGBT at low temperatures are followed up on, engaging all projects using the component and the developing team (CERN EP-ESE).

6. Discussion with ALICE

Scientific output and current activities:

- ALICE continues to make excellent progress on its physics program, with 8 new papers submitted since the last session of the LHCC, bringing the total number of publications to 484. The latest results include the measurement of the Ω_c baryon production and the measurements of its branching fraction $\Omega_c \rightarrow \Omega^+ e^- \nu_e$, an investigation of the strangeness enhancement in pp collisions using angular correlations and the first measurement of the four pion photoproduction in ultra-peripheral PbPb collisions. The collaboration is starting to analyse the new Run 3 data sets of pp and PbPb collisions, and preliminary results, such as measurements of the charged particle multiplicity density in PbPb collisions, have been shown in the Open Session.
- The skimming of the 2023 pp data is progressing; so far, 25% of the original compressed time frames (CTF) have been deleted, and the resulting skimmed CTFs have been archived on tape. The full 2023 PbPb data have been processed. The expected data accumulated in the online disk buffer, which now has a capacity of 150Pb following the lease of 60PB from IT, is tuned to keep sufficient free disk space for data taking. The 2023 data can be deleted to free more space, if required during the early phase of 2024 PbPb data taking, when 2024 pp data for skimming still require significant space. So far, ALICE has recorded 10.2 pb^{-1} of pp collisions at 13.6 TeV, with a data-taking efficiency of 95%, with all detectors in data-taking.

ALICE upgrades:

- The TDRs for the Run-4 upgrades of ITS3 and FoCAL, after having been endorsed at the last LHCC meeting, have been approved by the Research Board.
- The scoping document for ALICE 3 is being prepared with the aim of presenting a detailed assessment of resources and schedule. The document is expected for the September 2024 meeting of the LHCC. Several design studies and beam tests for R&D of the future ALICE 3 detectors have been performed and were presented in the LHCC Open Session.

- The **LHCC commends** the ALICE collaboration on its continued physics output with results of high impact demonstrating a very broad physics program.
- The **LHCC congratulates** the ALICE collaboration on the successful 2023 Pb-Pb data taking and the rapid first results, including measurements of the charged particle multiplicities at the highest PbPb collision energies to date.

- The **LHCC commends** ALICE on the progress in terms of calibrations, alignment and reconstruction and **notes** the importance of having a continuous and careful effort to achieve the full performance metrics and to understand the high occupancy effects before the next PbPb run.
- The **LHCC notes** the approval of the Technical Design Reports of the ITS3 and FoCal at the last LHCC and Research Board meetings and **recommends** a report at the next LHCC meeting on initial order placement, design selections, as well as funding and schedule.
- The **LHCC notes** the good progress on the ALICE 3 design and R&D and **looks forward** to receiving the Scoping Document before the next September LHCC meeting.

7. Discussion with ATLAS

Scientific output and current activities:

- ATLAS continues to make excellent progress on its physics programme, with 1286 papers on collision data submitted to date, of which 340 use the full Run 2 data set and 9 Run 3 data. Since the last session of the LHCC, 35 new papers have been submitted. Recent new results include the first measurement of the W width at the LHC in a combined fit with the W mass, a test of lepton universality in W bosons decays at the 0.5% level, di-Higgs search results which provide constraints on the Higgs trilinear self-coupling and VVHH coupling. All these results are summarised in physics briefing reports available to the general public.
- The 2024 recommissioning has been successfully completed and the experiment is now able to take data efficiently with a maximum pile-up of 62-63 and sustaining an L1 rate of 93kHz with a dead time of ~3.5%. So far, a total of 22.2 fb-1 have been recorded.
- The handover of the Technical Coordinator has taken place with a new TC organisation. The work on the Phase II infrastructure upgrade has continued over the EYETS with some activities being anticipated to flatten the workload in LS3. There are still difficulties in recruiting collaboration members for shifts and a problem of shortage of experts in several systems appears. A slow dump of the magnet occurred due to a chilled water cooling failure.
- All areas related to triggering are progressing well. The legacy triggers are being phased out and the new PhaseI triggers are being finalised and new complex physics triggers are being implemented, providing substantial gain in efficiency.
- Computing and Software are also making excellent progress. A major MC campaign for the 2022 and 2023 data has taken place and the campaign for the 2024 data is well underway.

Phase-II upgrades:

- The Phase II upgrade projects have progressed well, but the schedule remains critical. The report from the P2UG has been presented to the LHCC and is given below.

- The **LHCC congratulates** the ATLAS collaboration on the wealth of very interesting new physics results produced. The **LHCC is pleased** by the release of Physics Reports summarising the main results of Run 2 by topic.

- The **LHCC is pleased** to see that the area of Computing and Software is progressing steadily, thanks to the vast amount of work and dedication invested.
- The **LHCC congratulates** ATLAS for the sound and continuous progress in commissioning the Phase I triggers and **looks forward to** the final decommissioning of the legacy triggers and the full operation of all the Phase I triggers.
- The **LHCC commends** the collaboration for the successful recommissioning of ATLAS and the stability of the data taking, which has reached an integrated efficiency of around 94.5%.
- The LHCC **is pleased** to see the smooth takeover of the new Technical Coordinator and **is satisfied** that a plan to ensure the installation of the infrastructure required for the Phase-II Upgrade is developing fast. The **LHCC encourages** the ATLAS collaboration to continue taking proper heed of the Magnet systems and keep them under close surveillance.
- The **LHCC praises** the appealing program of training opportunities in a wide variety of areas offered by ATLAS, which makes working in the collaboration exceedingly engaging. The **LHCC regrets** that the ATLAS management has to strive to retain, replace and recruit personnel in certain areas. The collaboration should keep in mind that the proper running of the experiment requires an adequate commitment by all sides. The **LHCC fully supports** the approach of the management to the problem of shifter shortages. The institutes and collaboration members should be reminded kindly but firmly of their unavoidable duty to contribute to this essential task.

8. P2UG report on ATLAS

The ATLAS P2UG held a review of the Phase-II upgrade progress from 6th to 8th May 2024. The review was held in a full plenary format over the course of three days, in contrast to previous years, which required more condensed presentations of progress and status but allowed more of the committee to follow all parts of the Upgrade. A preliminary feedback session with the project proponents was held, with remote and in-person participation, on the 9th of May.

There continues to be excellent communication and rapport between the reviewers and the proponents, both specific to this meeting and in communication outside the scheduled meeting times. Presentations are clear and dialogue is candid and transparent, which increases the value both parties get out of these reviews.

Overall, an impressive amount of work has been achieved on all fronts, but the overall schedule has once again very little float remaining. Over the last two years, there has clearly been a noted improvement regarding the usefulness of the schedule as a predictive tool to guide mitigation of delays and technical issues and explore where efficiencies might be exploited. While there are always possibilities for further optimisation, it seems unlikely that there is something of the appropriate scale to regain a significant amount of slack.

Overall, since the Spring meeting, the schedule situation is considerably worse, as the contingency has shrunk, in many cases significantly, for twenty of the twenty-five areas reported. This does not take into account the fact that the current schedule assumes a cavern closure at the end of 2028, rather than the middle of 2029, which is currently foreseen by the LHC schedule.

In the slack distributions, the dominance of the ITk project at low and negative DCP values is prominent, although every subsystem with the exception of HGTD has some activities within a month of the need-by-date. There is some evidence of ATLAS working on the schedule, as the ITk-Pixel contingency has not appreciably changed in the last quarter, meaning the Task Force put in place last year is doing a decent job of mitigating delays from arising issues. But overall, maintaining a viable schedule has proved to be as difficult or perhaps even more difficult than predicted in previous reviews. With one significant technical showstopper, it is clear they cannot meet the current installation plan. What follows is a concise description of the update from all the areas.

Trigger/DAQ. While there has been good progress in many areas, TDAQ has lost a significant amount of float in various subtasks since the last full report in November 2023, and includes, in fact, the two most negative floats in the full ATLAS Phase 2 upgrade project. The New Small Wheel Trigger Processor has been consistently losing schedule contingency over the last year and now has an overall negative 107 days of float, mainly due to lack of effort both in management and technical expertise. It should be noted that a delay in that area does not affect cavern closure, since the processors are located outside. Additionally, there have been further delays in the muon sector logic boards for the TGC/RPCs due to updating the design to switch to halogen-free materials. The L0 Calorimeter trigger is gradually losing float, partially due to resource conflicts with Run 3. Prototype production of the fFex (feature extractor for the forward region) board is in progress. The L0 Global Trigger is testing v3 of the GCM, and no significant issues have been found with the board. A Final Design Review is planned for Q3 of 2024. Generally, the DAQ has proceeded according to plan, with the exception of the DAQ readout board. As a result of the Preliminary Design Review of the DAQ readout board, the prototype design moved to a more performant FPGA, causing a loss of over 100 working days of float since November.

The Tile project continues to make the progress necessary to meet the current schedule demands. Delivery float is still relatively healthy but is slowly being eroded. Two areas in particular, Daughter Board and Cs-Calibration, demand careful monitoring to ensure issues don’t escalate.

The Liquid Argon Calorimetry had an in-depth review. The last year has seen the resolution of all remaining technical questions such that all ASICS are either at, or near, production and all boards are entering the final prototype phase. Small schedule floats are currently reported, but P2UG is confident that in all cases these floats do not represent the actual situation and LAr will recover reasonable schedule float as the project progresses.

The Muon had also an in-depth review. Many significant reviews were completed recently, but there are many more (~20) scheduled in the next 39 weeks. The MDT chamber production has finished and shipping to CERN is expected to be completed in May. The sMDT mezzanine preproduction boards showed poor results in reverse crosstalk tests. A change of the PCB production process between prototype and preproduction is thought to be responsible. The mechanics FDR for the RPCs was completed. The frames, originally foreseen as a Protvino deliverable, have been reassigned to industry, and two commercial vendors were qualified. Gas leaks were discovered in the first production batch of 2x24 gas gaps packages (BIS+BIL) delivered to CERN, after passing all the QC tests at the production site.

The HGTD had an in-depth review. The major components, i.e. ALTIROC, LGAD sensors, their UBM and hybridisation, and module assemblies, which have so far been the limiting factors to the completion of the HGTD now seem to have almost converged.

The project seems almost ready to enter the new phase of module assembly and the focus will shift from components to module production, and then to large scale assemblies. However, there has been further slippage in the schedule. Although this was mostly not unexpected, it means that the overall schedule is not gaining any positive margin. Some gains may be possible during the module assembly phase, but only after experience is gained, so it will continue to be difficult to recover the lost time.

The ITk schedule continues to be the overall critical path driver for the ATLAS Upgrade. As presented at this review, the current float is -80 working days, a deficit of four months, dominated by the ITk-strips sensor cracking issue, which has halted progress on the critical path. As the solution has not yet been found, and the latency between proposal and validation of a solution is of order 3 months, the P2UG expects the situation to get progressively worse in the next six months.

ITk-Common: in response to a previous recommendation, the Production Database project has been bolstered, with 2 external engineers added and it is now heavily used for uploads, registration and reporting requests. There is impressive work on the SR1 assembly hall preparation, with floor replacements in place, and Global structures such as the Outer Cylinder and L3 and L2 shells expected soon. The cooling plant development is also progressing after the agreement on procurement strategy with CERN. Finally, there is significant progress in common electronics, with some of the dedicated interlock systems already integrated into SR1.

ITk-Pixels: the schedule has stabilised, and in the last quarter there has been no significant change to the ITk-Pixel delivery date. However, there is a long list of non trivial tasks lying, which have to be carefully managed, in particular the link between the flex qualification and module production, as the full flex qualification needs data transmission. The hybridisation situation has improved, with re-optimization of the load distribution among the two validated vendors and potential minor changes to detector configuration to mitigate the hybridisation availability problem. A new issue with the ITk-Pixel ASIC has emerged, the so-called “core columns” problem, where the ASIC ceases sending events until a reset is issued. It has been observed in about 5% of ItkPixV1 chips, but only after assembly in a variety of sensor types, and not in wafer testing. It remains to see if the effect is on ItkPixV2, although CMS sees a similar effect on CROC. The issue can be averted by disabling certain core columns, and is under further investigation. This is a large concern, as a resubmission of the ITkPix ASIC would clearly imply a significant schedule delay.

The ITk strips project was presented in a five-hour plenary P2UG session, providing a comprehensive overview of the project. Major achievements were reported in several areas. The sensor production is recovering from the previous delay and has reached 74% of good sensors delivered. The ASICS productions are progressing well and for the modules, 90% of site qualification is achieved. However, several challenges have been reported and discussed in-depth. The most critical issues are the sensor cracking and noise problems observed during cold testing of operational staves and petals. Although mitigation strategies have been proposed and are being explored, the issues persist and have effectively stalled the project just before the commencement of module production. The search for solutions has been hindered by difficulties in supplying sufficient parts for cold noise and sensor cracking tests. The entire ITk strip community is now intensely focused on addressing the problems with the highest priority. Additionally, other project areas have reported delays, but they are now believed to be progressing well. This is particularly true for the services, where delayed completion of tenders and contracts, as

well as delayed supplies from industrial vendors, have impacted on the schedule, bringing the services close to the critical path.

A more extensive report with key project-by-project observations will be provided in the Appendix to a future version of these minutes. The LHCC agrees with the observations and recommendations made by the P2UG. In particular:

- The **LHCC and the P2UG congratulate** ATLAS for the impressive amount of work going on across all fronts, particularly in the face of extremely challenging problems.
- The **LHCC and the P2UG urge** the collaboration to sustain the effort in all areas, and do not compromise the quality control necessary for a detector guaranteed to function throughout Run 4 and Run 5 for the sake of timeliness.
- The **LHCC and the P2UG note** that the schedule once again has very little float remaining, while significant remaining risks make further delays likely.

9. Discussion with CMS

Scientific output and current activities:

- CMS continues to make excellent progress on its physics programme, with 1288 papers on collider data submitted to date, including 30 since the last LHCC. Many new results have been presented to the Winter conferences and advertised with communication briefings. The new results include the measurement of the effective leptonic weak mixing angle with the forward-backward asymmetry in Drell-Yan process, the observation of $\gamma\gamma \rightarrow \tau\tau$ in pp collisions providing limits on g-2 of the τ lepton and the measurement of time-dependent CP violation in $B_s \rightarrow J/\psi\phi(1020)$ decays.
- After commissioning for 2024 data-taking, the detector is operating well with a pile-up of around 63. The data-taking is proceeding smoothly with 92.7% efficiency. The L1 trigger and HLT have included many new features and improvements to extend the CMS physics reach in 2024.
- The ZDC forward calorimeter, required for PbPb collisions, suffered a bake-out incident, caused by the baking out of the LHC sector with the detector still in place. A working group was created to understand and repair the damage.
- A sudden stop of the cooling plant of the pixel CO2 cooling system occurred due to the failure of a communication module, which was quickly recovered, resulting in moderate data loss.
- Progress has been made in the reduction of emissions of gases with high GWP, with reduced consumption and increased recuperation. The R&D on alternative mixtures is also progressing well.
- A new technical coordination upgrade schedule has been prepared, merging detector installation schedules and infrastructure activities, including resource loading.

Phase-II upgrades:

- The Phase II upgrade projects are progressing well, but the overall schedule remains very tight, and the collaboration is trying to get less active groups engaged in the upgrade projects. The P2UG reported the status of the upgrade, and its summary report is given below.

- The **LHCC congratulates** CMS on the wealth of very interesting new physics results produced with both Run 2 and Run 3 data and for the efforts of the communication team in disseminating them.
- The **LHCC commends CMS** for the smooth and successful ongoing data taking and for the planned reduction in data size by implementing a unified Nano AOD format for Run 2 and Run 3 and the potential of reduced AOD usage and disk occupancy.
- The **LHCC is concerned** about the engagement of the CMS Collaboration in service activities, as indicated by the low participation in the survey of the Task Force on Optimization of Time to Publication, and by the difficulty in finding strongly needed person-power for the upgrade projects.
- The **LHCC congratulates** CMS for the new release of open data sets accessible to the global scientific community, and **is impressed** with the success, in terms of publications, of the previous open data releases.
- The **LHCC appreciates** the ongoing effort by technical coordination to develop a detailed resource-loaded schedule for the upgrade taking all external constraints into account. The **LHCC acknowledges** TC's effort to mitigate the effect of a possible ICA termination with JINR, establishing new links with institutes that can provide expert technical teams. The **LHCC is satisfied** with the implementation of new procedures for P5 site protection to prevent thefts. The **LHCC is concerned** about the persisting problems with the construction of the engineering building and the sub-optimal functioning of the HGCAL clean rooms.
- The **LHCC appreciates** the measures introduced to improve the LHC-CMS communication following the ZDC bake-out incident to prevent such an event from happening again.
- The **LHCC is pleased** to see the positive results of the irradiation tests of systems with alternative mixtures/gases, and **strongly encourages** CMS to move to a 5% mixture of CF_4 as soon as possible.

10. P2UG report on CMS

The CMS P2UG met with the upgrade project on 14–16 of May 2024. This P2UG meeting the Tracker and HGCAL projects were reviewed in-depth. Regular status reports were given by BRIL, Muon, BTL, Barrel Calorimeter, L1 Trigger, DAQ and HLT. Clear and steady progress is made in all technical areas of the upgrade and the team is to be congratulated for its continued success in realizing the project. The tracker project carried out a Fraunhofer-like exercise that has provided valuable insight in the production schedule and enabled the project to gain schedule contingency. The schedules for both the HGCAL and Tracker seem to be highly optimized, are in several areas aggressive and success-oriented and are no further compressible. Any further compression would invariably lead to quality losses and the performance of the detector would be compromised.

The project has unfortunately encountered another setback with the HGCROC. At the previous P2UG meeting, it was reported that there was a delay of potentially up to 3 months in the availability of the HGCROC due to a communication issue between IMEC and TSMC. V3b of the HGCROC was received in February of this year and unfortunately

had two major bugs: an error in the transition of the energy measurement from ADC to ToT and sporadic bit outliers. Both bugs have been reproduced in simulation and the strategy for the fix is being worked out. The team is currently forced to consider the use of the HGCROC.v3b chip for pre-production modules and use them in regions of low energy deposit in the final detector, corresponding to at most 5% of the total production. If production and yield are favourable, the fraction of faulty HGCROC.V3b chips in the final detector can be reduced. The team is aggressively addressing these issues. The ECON-T2/D2 chips have just been received but have not been tested yet. Production of the cooling plates for CE-E is about to get started; the bid for the CE-H cooling plates will be extended in July 2024. Progress on the electronics boards is good, though not all boards have been fabricated yet. Module production is going well, but the production of the CuW and carbon fiber baseplates, for which the procurement strategy is still being developed, may become a bottleneck. Sensor production is going well, with 44% of the sensors delivered. Fabrication of the tile section of the calorimeter is also going well. The SiPM contract was awarded at the end of 2023, with the first delivery to CERN this coming month. Good progress is being made with system tests, though not all components were part of the chain test. The project has incurred another three-month delay and has an overall schedule contingency of about one month.

The Tracker project has made excellent use of the Fraunhofer exercise and now has an excellent understanding of their schedule and of all dependencies during production. ASICs and sensor deliveries for both Inner Tracker (IT) and Outer Tracker (OT) is in very good shape. For most of the production, the production speed of the OT is driven by the availability of hybrids and MAPSAs. One vendor for the MAPSA production has delivered parts with serious quality issues. Production of service and mechanical parts are in the shadow of the critical path. The OT has a schedule contingency of about 2 months. The IT project has regained four months of float thanks to the Fraunhofer-like exercise. Hybridisation remains the most critical item and remains the biggest concern for the project. At the moment, only IZM is fully qualified, while two other vendors still need to be fully qualified. IZM, however, is currently unable to cope with both CMS and ATLAS production demands as they are known today. Both the HGCAL and Tracker projects are in urgent need of more personnel.

The BRIL project has concluded a successful validation of the performance of the ASIC for the Fast Beam Conditions Monitor, but significant noise was observed in test beam operation of an FBCM prototype, which requires a redesign of the Front-End module board.

The Barrel Timing Layer is moving into pre-production. LYSO crystal production is well underway, and the first production tray will be assembled at CERN in July of this year. A major milestone was the delivery of the BTL Tracker support tube. The Endcap Timing Layer has validated the performance of an LGAD assembly readout with the ETROC2 chip and obtained a timing resolution of 45ps. Functional and performance tests of the ETROC2 have been completed, including radiation tests. Although TMR has not been fully implemented in the chip, no adverse effects have been observed, which may obviate the need for another version of the chip.

The project has prioritized completion of ME0 in the muon project. The schedule for module production is driven by PCB availability and procurement and the rate of foil production. It is estimated that by the end of 2024, the production will be limited by the availability of foils. Foil production at CERN and KCMS is progressing well. No drop in gain has been observed in the ongoing ME0 longevity studies, which is very

encouraging. The project is on a very tight schedule and its success depends on the frenetic production of multiple components during 2024.

Chamber production of the iRPCs is on track at Ghent and CERN; one endcap has been completed, and the assembly of the second endcap has started. As reported at the previous P2UG meeting, due to procurement delays, the production of the Front-End Board for the iRPC system, FEB V2.3, has seen a delay of 8 months and the full delivery of FEB V2.3 boards is now expected by summer 2024. As a consequence, the schedule for iRPC production is defined by FEB availability, with all FEBs expected by September 2024. The completion of all chambers is planned for November 2024 with only a one-day quality control test per chamber. This is a very aggressive schedule.

The Barrel Calorimeter has observed an unexpected 50% increase in LiteDTU power consumption. The team has a clear plan to address this issue and understand its origin. No significant schedule delay is expected. The Level 1 trigger, high-level trigger and DAQ projects are all making very good progress.

The two biggest sub-projects of the CMS Phase 2 upgrade are the Tracker and the HGCAL, that drive the overall schedule. Currently each project has a schedule float of about one month, based on a schedule that is highly optimized. A more extensive report with key project-by-project observations will be provided in the Appendix to a future version of these minutes.

The LHCC agrees with the observations and recommendations made by the P2UG. In particular:

- The **LHCC and the P2UG congratulate** CMS on the technical progress in all areas of the challenging Phase II upgrade projects, in particular on the successful transition to preproduction for the BTL, the completion of the BTL Tracker Support Tube and the excellent results of the ETL module test obtaining 45ps timing resolution.
- The **LHCC and the P2UG** note that the schedule once again has very little float remaining, while significant remaining risks make further delays likely.
- The **LHCC and the P2UG** urge all invested parties to continue to make the necessary investments in an expeditious way to deliver the project so that it can meet its physics goals in a timely manner.

11. Discussion with LHCb

Scientific output and current activities:

- LHCb continues to deliver high-quality physics results, with a total of 733 publications to date, including 11 new papers since the last session of the LHCC. New results include a new measurement of the Ξ_b^- baryon lifetime, improving the precision of the current world average by a factor of 2, the observation of the rare decay $J/\psi \to \mu^+\mu^-\mu^+\mu^+$ and the search for the lepton-flavor violating decay $B_s^0 \to \phi\mu^+\tau^-$.
- The 2024 data taking is progressing well with about 1.4 fb-1 of data collected. The performance of the detector is good, including the VELO, which has been successfully commissioned with the beam during the intensity rump up. The VELO operates in the closed nominal position with the expected efficiency. A small drift of the VELO C side is observed during running, but it has been mitigated by using a dedicated closure procedure and applying software

corrections. Despite the higher luminosity, the RICH detector shows stable performance comparable to Run 2. The UT detector is currently being commissioned locally and is included in dedicated global runs. The final goal is to include the UT in the trigger decision at the HLT1 reconstruction. Due to the new software trigger deployed in Run 3, the efficiency gains for hadronic final states, especially at low transverse momentum, are a factor between 2 and 4 compared to the previous Run, depending on the specific decay channel.

LHCb upgrades:

- The TDR for LHCb PID Enhancement, endorsed by the LHCC, was approved by the RB. Currently not the full sum of the planned CERN contribution to the CORE cost of the upgrade is foreseen in the outer years of the MTP. The TDR on the DAQ enhancements for Run4 is currently under review, which is expected to be completed at the September LHCC meeting.
- LHCb is expected to present the Phase IIb upgrade Scoping Document at the September 2024 meeting. Currently, three scenarios are being studied in detail, ranging from 123 MCHF to 182 MCHF. Simulation studies are being performed to study the physics performance in the different scenarios, and a preliminary cost contribution matrix is being prepared.

- The **LHCC congratulates** the LHCb collaboration on the excellent physics output with top-quality physics results in the field of flavour physics and several high-precision measurements.
- The **LHCC commends** LHCb on a very successful start of data taking in 2024 and acknowledges the significant improvements in addressing the challenges in data taking since July 2023, resulting now in a detector being fully prepared for the 2024 data taking.
- The **LHCC congratulates** LHCb for the recovery of the VELO after last year accident and **supports** the prompt study of the effect of a residual VELO drift to physics performance after the applied corrections.
- The **LHCC agrees** with continuing the LS3 enhancement projects as planned, despite reduced funding currently foreseen in the MTP, and **encourages** CERN to allocate the full funds required in future MTPs, provided the project demonstrates its ability to spend resources on the expected funding profile.

12. Discussion with WLCG

The WLCG infrastructure continues to perform extremely well. The data taking in 2024 is progressing as planned and the available resources are sufficient to cover the 2024 needs for CPU, disk and tape. The experiments are making progress in adopting reduced data formats for analysis, reducing the data read loads and the storage needs. ATLAS and CMS are also progressing in the software R&D for Phase 2, by porting some of the reconstruction algorithms, such as tracking, to GPU and making use of machine learning for both fast simulation and tracking.

The WLCG community agreed during the recent WLCG/HSF workshop in DESY on a set of questions for experiments to better define the requirements for the Analysis

Facilities and use them in a process to help sites better understand what running this infrastructure entails.

A data challenge occurred in February 2024, and its results have been presented. The goals in data transfers were achieved thanks to the extensive prior work, and valuable insights were gained that will guide future improvements. For the first time, tokens were exercised at scale during DC24 and this resulted to be helpful in uncovering relevant design issues that need to be followed up and addressed.

The status of the HSF activities and the status of the common software projects from SFT were presented. The activity metrics of the ROOT project show an important reduction in the number of open issues while making progress with the program of work. The ROOT/RDataFrame tool is ready to run on various types of Analysis Facilities. Thanks to a joint effort between IT/ST and EP/SFT, large-scale tests of ROOT/RNTuple are being performed, confirming the expected speed improvements. The binary RNtuple format is expected to be frozen by the end of the year.

- The **LHCC congratulates** the WLCG and the experiments on the successful and efficient use of the computing resources. The **LHCC commends** the four experiments for their work on validating their software stack on new computing architectures such as ARM in order to be ready to exploit the potential energy efficiency, performance or cost advantages that those might bring. The **LHCC notes** the ALICE approach to using the ARM validation to prepare generic infrastructure that can be useful in speeding up any other architecture validation in the future.
- The **LHCC commends** the Glasgow Tier2 for their thorough benchmarking efforts and looks forward to the results of power efficiency tests with newer CPU models. The committee encourages WLCG to streamline these findings to provide clear guidance to sites for their hardware purchases.
- The **LHCC congratulates** the entire WLCG community on the successful completion of the DC24 exercise. The **LHCC supports** the community's proposal to establish a Task Force to investigate the discovered issues due to tokens and develop an effective and scalable token usage process. The **LHCC suggests** that this Task Force prepare a comprehensive plan, including a detailed list of necessary actions and timelines, to achieve consistent and scalable token usage across the entire system.
- The **LHCC recommends** that a future LHCC meeting include a presentation on how both SFT and HSF groups are collaborating towards the success of common scientific software.
- The **LHCC welcomes** the kick-off of the new project in SFT for common ML activities aimed to provide service and support to the experiments on common ML issues. The **LHCC is pleased** to see that specific goals have been identified, as well as near and long-term plans. The **LHCC recommends** regular updates with the experiments on requirements for the common software framework to train and optimise machine learning models, ensuring its close integration within the experiment workflows.

13. Report on FASER

New physics results were released by FASER, covering a search for ALPs, with limits covering previously unconstrained parameter space, and the measurement of νe and νμ cross sections in the TeV energy range. The first studies with 2024 data confirm the good performance of the calorimeter readout upgrade, to be further optimized via a final calibration of the energy response. The ASIC production for the pre-shower upgrade has been delayed due to post-processing issues at the foundry. 427 ASICs have eventually arrived in Geneva on May 17, and the first assembly and testing are ongoing. The project remains on track for surface commissioning of the full detector by the Fall and installation during the YETS.

Progress is ongoing to enhance the exploitation of the FASERν output, ranging from improvements to the scanning process to the improvements to the tracking, alignment and neutrino energy reconstruction. The scanning of the 2nd 2022 module of emulsions, which integrated 10 fb^{-1}, has been completed. Scanning of the 3rd module (30 fb–1) is in progress.

As discussed in the Programme Coordination report, FASER is observing in 2024 a much-increased muon background. Contrary to 2023, repositioning of the TCL6 collimator does not appear effective in reducing such background. The doubling of the muon rate is accompanied by a factor 3.5 increase in hit occupancy, due to the high momentum spectrum of this background. FASER strongly requests that the optics for 2025 are optimised to improve this situation, and favours reverting to non-reverse polarity optics as the safest solution to the background issue. This increased background reduces to O(10 fb^{-1}) the maximum integrated luminosity for each emulsion batch. A first match was removed in mid-May, leaving just the tungsten absorber in the FASERν detector. A second batch is planned to be installed during the June technical stop, and to be extracted according to the agreement reached within the LPC (see the LPC report). The final third emulsion batch will be installed during the access in concurrence with the VIP visits of Oct 1.

- The **LHCC congratulates** FASER(ν) for the release of new physics results, the successful deployment of the calorimeter readout upgrade, and for the steady progress towards completion of the pre-shower upgrade project.

14. Report on MoEDAL

The first stack of 60 m2 of MoEDAL emulsions has been removed and is ready for etching and scanning. The preparation of the MAPP detector readout boards is back on track. The delivery of the first two (out of three) boards is expected for May 31. The last scintillator bars and photomultipliers will be installed during the June technical stop, together with the available front-end electronics. The completion of the detector is foreseen for the YETS, but access opportunities during the run could be exploited to begin at least a partial commissioning of the detector, with beam, before the YETS.

Two papers documenting MoEDAL's Run2 searches for magnetic monopoles and highly ionising particles have been accepted for publication by PRL. A new group, from Oujda's University (Morocco) has applied to join the experiment.

- The **LHCC recommends** the urgent completion and deployment of the MAPP detector to guarantee effective data taking in 2025 and to enable a timely review

of the outriggers proposal.

15. Report on SND

The SND detector has been read for 2024 data taking with the third veto layer operational, which allows to increase the neutrino fiducial volume. However, an increase in the muon background rate by a factor of 2, as measured by the SciFi detector, has been observed. Given this increased background, SND decided to replace the emulsions every 10fb-1 of integrated collected luminosity instead of the 20fb-1 planned. In addition, since the muon background rates are higher in the upper part of SND and the expected neutrino events are higher in the lower part, SND decided to install only half of the emulsions in the lower part and leave the upper part as passive target without emulsions. This allows them to have emulsions for 7 further half targets, covering almost all the 2024 running period.

The scanning of the emulsion films of the data collected previously is proceeding well, with 4 microscopes steadily running at CERN in parallel. 20 fb-1 of emulsion data are expected to be analysed for the coming summer conferences.

SND presented at the LHCC the LOI for the SND upgrade for Phase II, the magnetised minimal AdvancedSND. The plan is to excavate the floor of the LHC tunnel to bring AdvSND on the same axis with the beam and to minimise the crossing angle effect foreseen in Run4 and mitigate the switch to a vertical crossing angle in Run 5. The neutrino target will not be based on emulsion as the current SND, but made of tungsten slabs interleaved with silicon detectors. It will be complemented with a magnetised HCAL consisting of 22 layers of silicon detectors or drift tubes. An agreement with CMS has been reached to reuse the CMS silicon tracker modules. The costs of the detector, excluding the civil engineering work, which is expected to be covered by the host lab, are expected to be slightly more than 1 MCHF.

- The **LHCC congratulates** SND for the new neutrino results with 2022-23 data presented at the winter conferences and **looks forward to** the first analysis of the emulsion sample for the summer conferences.
- The **LHCC congratulates** AdvSND for the submission of a Letter of Interest and the subsequent Addendum for operation during run 4.
- For further consideration, the **LHCC requires** AdvSND to
    - quantify systematic uncertainties on the main physics program objectives, namely the neutrino LFV and gluon PDF measurements without an additional near detector, and to compare them to what can be achieved at non-advanced SND.
    - Quantify the detector performance for all planned measurements, including energy resolution of the silicon calorimeter and track efficiency in pessimistic background scenarios.
    - Promptly assess the feasibility of the CE works and present an ECR (to the LMC) in time to be considered at the LHC LS3 readiness review in September.
- The **LHCC notes** that the foreseen CE works can only be completed within the two-month window in Q3-2026 to avoid impacting the HL-LHC installation and

the planned activities near P1 during LS3.

- The **LHCC asks** AdvSND to complete the cost estimate, including those for the CE works and overall infrastructure and installation of the experiment, and clarify the possibility of funding with the host lab.

Comments applicable to both SND and FASER:

- The **LHCC acknowledges** the very challenging background conditions for FASER and SND due to the optics change in the 2024 LHC running. The **LHCC is concerned** about the significant reduction in integrated luminosity for both FASER and SND in 2024.
- The **LHCC appreciates** the high degree of flexibility of SND and FASER in making use of opportunistic access to exchange emulsions and to adapt to the challenging background conditions.

REFEREES

The LHCC referee teams for this session are as follows:
ALICE: M. Battaglieri, D. Calvo, M. Diefenthaler, T. Galatyuk, J. Nagle (Co-ordinator)
ATLAS: J.J. Hernandez-Rey (Co-ordinator), M. Kirby, G. Signorelli, H. Tanaka
CMS: A Ianni, I. Gil-Botella, S. Niccolai (Co-ordinator), A. Weber
LHCb: B. Golob (Co-ordinator), T. Higuchi, R. Leitner
LHCf: M. Mangano (Co-ordinator)
MoEDAL: T. Ferber, M. Mangano (Co-ordinator)
WLCG: M. Diefenthaler, M. Kirby, G. Merino (Co-ordinator), A. Weber
FASER: T. Ferber, M. Mangano (Co-ordinator)
SND: T. Ferber (Co-ordinator), M. Mangano

The LHCC received the following documents:

CERN/LHCC-2024-005	Minutes of the one-hundred-and-fifty-seventh meeting held on Wednesday and Thursday, 28-29 February
CERN/LHCC-2024-006	AdvSND, The Advanced Scattering and NeutrinoDetector at High Lumi LHC Letter of Intent

DATES FOR LHCC MEETINGS

Dates for 2024

11-12 September
20-21 November

Lorenzo Moneta
E–mail: Lorenzo.Moneta@cern.ch
Tel. 74886

LHCC Secretariat: csc.secretariat@cern.ch

Appendix A: **Report of the ATLAS Phase-II Upgrade Project Review (P2UG)**

Appendix B: **Report of the CMS Phase-II Upgrade Project Review (P2UG)**

A. Summary Report of the ATLAS P2UG review

G. Barker, P. Collins, J.-L. Faure, M-C Fouz-Iglesias, G. Hall, R. Mankel, M. Moll, S. Nahn (chair), V. Odell, D. Wood

Introduction

The ATLAS P2UG held a review of the Phase-II upgrade progress from 6th to 8th May 2024. The review was held in a full plenary format over the course of three days, in contrast to previous years, which required more condensed presentations of progress and status but allowed more of the committee to follow all parts of the Upgrade. A preliminary feedback session with the project proponents was held, with remote and in-person participation, on the 9th of May.

There continues to be excellent communication and rapport between the reviewers and the proponents, both specific to this meeting and in communication outside the scheduled meeting times. Presentations are clear and dialogue is candid and transparent, which increases the value both parties get out of these reviews.

General Comments

The general feeling from the committee was succinctly put in the closing remarks of the summary presentation:

> *The ATLAS Upgrade should be commended for the impressive amount of work going on across all fronts, particularly in the face of extremely challenging problems. We urge you to sustain the effort in all areas, and do not compromise the quality control necessary for a detector guaranteed to function throughout Run 4, Run 5 for the sake of timeliness.*
>
> *That said, it is clear that for the detector as a whole, the current schedule will not be met. Over the last two years, there has clearly been a noted improvement regarding the usefulness of the schedule as a predictive tool to guide mitigation of delays and technical issues and explore where efficiencies might be exploited. This panel as well as others have examined to see if there is more to be done, and while there are always a few optimisations, there is nothing of the appropriate scale to regain the amount of slack needed to meet the schedule.*

There are two important clarifications to make

1. The LS-3 schedule utilized by ATLAS in this review has the end of LS3 at the end of 2028, rather than the middle of 2029, which is currently foreseen by the LHC. Thus ATLAS believes they can recover about five months of slack, but likely not enough, by making this change.
2. While it is true that the ITk-Strips "sensor-cracking" problem draws much of the schedule focus, it should not be seen as the cause of the delay, but rather the proverbial "straw that broke the camel's back". As discussed below, many different areas are losing valuable slack across the whole project.

Progress against schedule

Overall, since the Spring meeting, the schedule situation is considerably worse, as the Distance to Critical path has shrunk, in many cases significantly, for twenty of the twenty-five areas listed above. The schedule is discussed in detail in the technical sections, and

some of the loss is an artefact of sequential planning where execution will be in parallel, but some notable features are:

- TDAQ L0: there is an ongoing resource crisis associated with the New Small Wheel trigger processor, which continues to have negative float. New resources have recently been brought to bear, but have not yet had an impact.
- TDAQ DAQ: another iteration of the DAQ Readout board design was needed, incorporating a more performant FPGA, at the cost of 100 days of float.
- ITk-Strips: in addition to previous delays, a new technical issue has emerged in pre-production that has caused significant loss of float in the Barrel and Endcaps. This is now the most critical part of the schedule, driving ITk readiness to 4 months (80 working days) past the need-by date from the current LS-3 schedule. Given there isn't yet a solution to the issue, this discrepancy is expected to grow before the next P2UG.
- Muon MDT Front End: excessive crosstalk revealed a layout issue which triggered a second pre-production before PRR.
- Muon TGC Electronics: a single system, the Charge Monitoring System, is responsible for the loss of slack, but is in fact not needed for installation of the rest of the system.

In the slack distributions, the dominance of the ITk project at low and negative DCP values is prominent, although every subsystem with the exception of HGTD has some activities in the lowest positive bin, within a month of the critical path. There is some evidence of ATLAS working on the schedule, as the ITk-Pixel DCP has not appreciably changed in the last quarter, meaning the Task Force put in place last year is doing a decent job of mitigating delays from arising issues. But overall maintaining a viable schedule has proved to be as difficult or perhaps even more difficult than predicted in previous reviews. With one significant technical showstopper, it is clear they cannot meet the current installation plan.

TDAQ (Summary)

While there has been good progress in many areas, TDAQ has lost a significant amount of float in various subtasks since the last full report in November 2023, and includes, in fact, several of the most negative floats in the full ATLAS Phase 2 upgrade project.

The New Small Wheel Trigger Processor has been consistently losing schedule contingency over the last year and now has an overall negative 107 days of float, mainly due to lack of effort both in management and technical expertise. Several new groups (NCSR Demokritos (lead), Athens, University of West Attica, and Thessaloniki), along with a new project coordinator joined the effort in January, and the effort to review the project schedule and requirements is underway. The committee looks forward to looking at the details of this effort at the next in depth review (November, 2024). Engagement of these four new groups in NSW-TP is very encouraging. It is, however, important to assess whether this will bring the project back on schedule, or if further strengthening is needed.

Additionally, there have been further delays in the muon sector logic boards for the TGC/RPCs due to updating the design to switch to halogen-free materials, however, the project feels that there is a potential to regain some of the lost float as the firmware tasks have progressed in parallel with the hardware design and are quite advanced. Nevertheless, the trend is worrying, as there is not much float left. The Final Design Review for the SL boards is scheduled for the first quarter of 2025; it is important to hold that schedule.

At the last review, a recommendation was made to report on the findings of a task force that was set up to develop a mitigation strategy against a possible loss of RPC coverage. The Muon/TDAQ task force assessed possible scenarios of future issues with RPC performance

and are working on a full report, however, the current status of recommendations foresees no changes to L0Muon trigger processors.

In addition to the L0 Muon trigger, the L0 Calorimeter trigger is gradually losing float, partially due to resource conflicts with Run 3. The situation with the feature extractor for the forward region (fFex) needs to be carefully monitored, since as opposed to the other Fex boards, the system is completely new and scheduled to arrive relatively late. Prototype production of the fFex board is in progress. The team has dropped the idea of retaining the use of the Run 3 L1Topo as a trigger legacy system for Run 4 startup. It was felt that the effort is better spent focussing on commissioning the L0 Global trigger on time. The detailed commissioning plan for L0Global is in development.

The L0 Global Trigger is testing v3 of the GCM, and no significant issues have been found with the board. A Final Design Review is planned for Q3 of 2024. However, there is a major concern with the Samtec Firefly 25 Gbps transceiver modules: they failed validation tests at CERN, and the vendor has confirmed the problem. There should be a new version of the module (with a different VCSEL) for testing in May/June; if this module also fails, it is possible that a new transceiver will have to be incorporated into the design, potentially requiring PCB re-work to accommodate new routing (and schedule delay).

Generally, the DAQ has proceeded according to plan, with the except of the DAQ readout board. As a result of the Preliminary Design Review of the DAQ readout board, the prototype design moved to a more performant FPGA, causing a loss of over 100 working days of float since November. The overall schedule contingency of the DAQ readout is -77 days. By optimising board procurements, since most of the early needs of FELIX cards for detector commissioning could be supported by earlier prototypes, the group feels they can speed up the schedule and return the readout to +75 days of float. The final design review is scheduled for early 2025, and the final FPGA selection will then be made. It is important to establish that the substitute FPGA works satisfactorily, and this third prototype can be the final version. This board also uses the Samtec Firefly transceiver; it is monitoring the situation closely.

For the DAQ network, all hardware for the demonstrator has been received and is being commissioned at the CERN testbed, with the final design review scheduled for the end of 2024. At the same time, CERN IT is working on a network procurement market survey. Ideally, the ATLAS DAQ network will be able to take advantage of the CERN IT tender. However, if not, then the ATLAS DAQ team may have to make a separate vendor, resulting in an overall increase in cost. The outcome of the IT market survey should be known in early summer.

Dataflow testing is going well, culminating in a successful large-scale test at P1 during the 23/24 EYETS. A seven-hour test at close to a 1 MHz L0 rate ran smoothly. Further studies and development continue in the testbed lab with upgraded hardware and network.

Event Filter technology studies are ongoing. The first EF Calo and Muon demonstrator phase is just concluding, with the Calo EF using an advanced GPU TOPO clustering demonstrator, and an FPGA version well underway, as well as developing machine learning algorithms. The muon ML algorithms are also progressing, but other high-priority tasks are suffering from a lack of person power. The EF Tracking second demonstrator phase is on track, putting together candidate implementations of each part of the tracking chain. It is on track to meet the milestone of having at least one pipeline of each technology by Q3 of 2024, and is leveraging the significant effort from ATLAS offline computing using GPUs and CPUs. Overall, the technology choice criteria for the EF are currently under review, with a final draft being assessed.

TileCal (Summary)

The recommendations from Nov P2UG have been addressed. Many aspects of the project are now in production, and the levels of project float available for the start of partition assembly are still relatively healthy (~100-200 days).

On-detector electronics is mostly on schedule with the exception of the Daughter Board which continues to raise a concern. The FDR date is reported to have slipped further from Q1 of 2024 to later this summer. Confidence has risen with the production of v6.4, and the FDR is expected to be completed by September of this year. Off-detector electronics FDRs have been delayed to 2024 partly because of the delays in producing v6.4 of the Daughter Board but also to verify the thermal performance, plus some other technical aspects.

In the low voltage system, the use of the original choice of isolation amp (SI8920) is back on the table following new radiation tolerance data from CHARM. If lot/batch testing of the components (already procured) checks out, the current baseline choice becomes the backup and will again allow surface mount rather than the work-around socket connection that was being proposed. The new high-voltage cable and connector design (forced by the closure of the preferred vendor's production line) are now tested and ready to close out the FDR. Production and acceptance testing of PMTs is now running at full capacity over 3 test sites (50 pcs/month). Results are showing devices of excellent quality.

In conclusion, the Tile project continues to make the progress necessary to meet the current schedule demands. Delivery float is still relatively healthy but is slowly being eroded. Two areas in particular, Daughter Board and Cs-Calibration, demand careful monitoring to ensure issues don't escalate.

Liquid Argon Calorimetry (In-depth review)

The last year has seen the resolution of all remaining technical questions such that all ASICS are either at, or near, production and all boards are entering the final prototype phase. Small schedule floats are currently reported for FEB2, Calibration Board, FEC LVPS and LASP but, P2UG is confident that in all cases these floats do not represent the actual situation, and LAr will recover reasonable schedule float as the project progresses. For the FEB2 the schedule has one year between FDR and PRR, which is conservative. The FDR is expected to move into early 2025, but the PRR will not need to move. The current PDR schedule for the Calibration Board (March 2025) is overly pessimistic and is expected to take place in the fall of this year, so recovering ~100 days of float. The FEC LVPS schedule can be re-baselined once the results of the tender are known. For LASP, the prototype board is expected to be ready mid-October so gaining ~40 days of float.

Muon (In-depth review)

Many significant reviews were completed recently, including FDRs for RPC BIS Mechanics, RPC FE ASIC (Part 1), RPC BIL Mechanics, PDR for RPC Gas, and PRRs for RPC Singlets, EIL4 chambers, and TGC Electronics (PS-, JATHUB-). We note that there are many more reviews scheduled in the near future - 20 in the next 39 weeks. It will be challenging for the project to carry out all of them, but it is important to try to keep this schedule because this is an important period with many projects going into production.

A task force was convened to lay out options for L0 muon trigger for various scenarios of failure of legacy RPCs. A report is expected soon.

sMDT chamber production: Chamber production has finished. Shipping to CERN is expected to be completed in May. Re-testing is carried out at CERN after reception, and the quality is good. The gas tightness tests at CERN are systematically better than those at the production site, but this might be related to more precise measurements at CERN.

MDT electronics: The sMDT mezzanine preproduction boards showed poor results in reverse crosstalk tests. A change of PCB production process between prototype and preproduction is thought to be responsible. A new preproduction with a small layout change addresses the problem, but some schedule loss resulted. The delay is being mitigated by carrying out production at two sites. For the TDC chips, 30k were produced, and about 12k were tested so far, with 94% yield. Manual testing at Michigan is supplementing robotic tests at Fermilab to avoid delays. CSM motherboard has mechanical alignment problems with CSM modules. Precision gauges were developed to allow the assembly company to check connector alignment. The mezzanine to CSM cables had been procured earlier and stored. When tested, they were found to be of poor quantity. A new vendor has been found and is producing samples that look good so far. Some progress was also made in integrations, with tests of MDT electronics together with L0MDT, and sMDT chambers with FE electronics.

BI RPCs and FE electronics: The mechanics FDR for the RPCs was completed. The frames, originally foreseen as a Protvino deliverable, have been reassigned to industry, and two commercial vendors were qualified. Production is expected to start in October. RPC production is planned to take place at over six sites to spread expertise, with fallback options for production at GTE. Gas leaks were discovered in the first production batch of 2x24 gas gaps packages (BIS+BIL) delivered to CERN after passing all the QC tests at the production site. Several leaks showed up in tests at BB5, and the sniffer shows that leaks are located in the corners. The glueing of the production frames is suspended until a more complete understanding of the origins of the leaks is achieved. These leaks in the newly built RPC gas gaps are concerning. High priority should be given to understanding the origin of this problem, and implementing and verifying a solution. It is important to continue rigorous QA/QC on the gas tightness.

Following a recommendation from the last P2UG review, ATLAS has created a roadmap to formally update the plans for the BIM/BIR chambers. It is now officially decided that the 120 BIM/BIR chambers will not be built, and instead, 80 BOM/BOR chambers will be installed to cover the same angular range at a larger radius. There was no viable installation option for the BIM/BIR location. This is not considered a change in scope. ATLAS continued to follow the strategy for the FE ASIC that was presented to the P2UG in May 2023: three versions of the ASIC are included in the multi-project runs, where there are sufficient quantities of any of the three to satisfy the production needs. The ASICs from the May 2023 submission were received recently and are being tested. The remaining tests include the final qualification of the communication between FE and DCT with the full chain and realistic cable, and completion of the SEU radiation test. The FDR was split into two parts, and the final choice of ASIC version will be made in the second FDR. Following a previous P2UG recommendation, ATLAS has reserved space in an additional future multiproject run, but this will not be needed if the tests of the current chip are successful.

RPC Trigger & Readout electronics: A second BMBO DCT prototype was produced, but signal connection problems were observed that were not present in the first prototype. This is under investigation. The first BI DCT prototype is expected to be finalised in about one month. The radiation testing of all DCT components is done except for one remaining TID test on a voltage translator. Good progress was made on BI DCT firmware for Manchester decoding.

EIL4 TGC: The PRR for the EIL4 TCG chambers was passed, and construction is underway. Fifty chambers are needed, and currently two are done, and eleven more are in various stages of production. ATLAS plans to install one section in the 2024/25 YETS.

TGC Trigger and Readout Electronics: The PRR was passed for the PS board and JatHUB board. A demonstrator is in place. The charge monitoring board is behind schedule due to a lack of person power. This currently drives the DCP for the TGC Trigger and Readout Electronics, but the charge monitoring board is not critical for operation and can be installed later than the other boards. A BCP is foreseen to correct the misleading small DCP. In searching for a new radiation hard ADC for the TGC electronics, the ADC used in the LAr project could be considered as it is well-tested.

Power System: A prototype system controller was delivered and tested. Some bugs were found and reported to CAEN. Radiation tests at CHARM are expected in July and October, and these are needed before moving ahead with the order.

HGTD (In-depth review)

Once again, the reviewer is grateful to have found slides posted in advance, even if only by a couple of days, as there is a lot of material to try to absorb at relatively short notice. The constructive reactions to previous review recommendations are noted and, once again, steady progress over the six months since the last review is evident, reducing some previous concerns, especially the Peripheral Electronics Board (PEB).

The major components, i.e. ALTIROC, LGAD sensors, their UBM and hybridisation, and module assemblies, which have so far been the limiting factors to the completion of the HGTD, now seem to have almost converged. The project seems almost ready to enter the new phase of module assembly, and the focus will shift from components, to module production, then to large scale assemblies.

However, there has been further slippage in the schedule. Although this was mostly not unexpected by the reviewers since, e.g., for the last ASIC submission the amount of time predicted for recent work needed seemed very tight, it means that the overall schedule is not gaining any positive margin. Some gains may be possible during the module assembly phase, but only after experience is gained, so it will continue to be difficult to recover the lost time. The overview presentation appeared to foresee a significant acceleration soon to match the final delivery and installation schedule. This will probably be difficult and must not come at the expense of detailed evaluation and quality control.

A realistic demonstrator now seems imminent. This will be an important step forward.

Other observations and comments.

This review cannot improve on the project's own evaluation of sensors, whose quality and yield look very promising, including UBM. Sensor QA/QC plans appear to be in good shape. As observed before, the sensor production flow is complex and remains to be proven, so this is the next challenge.

Similarly, recent ASIC progress appears to be much as the P2UG expected, including taking more time than the November 2023 plan foresaw. The time to the PRR is short, since the next ALTIROC version – which ought to be final - will not be delivered until June. A lot of evaluation is expected between June and October, which would be challenging even without the usual slowdowns during the summer vacation period.

An important aspect of the ASIC evaluation is large-scale wafer probing, whose preparations still seem to be taking longer than desirable, and do not yet appear to be firm, or based on

much experience with ALTIROC wafer testing. It is noted that, unavoidably, almost all evaluations so far are based on limited statistics, and small issues can show up when larger numbers of units – ASICs, modules, etc – can be scrutinised in detail. It is crucial that evaluation of final modules does not identify any ASIC-related problems to avoid any further ASIC design changes.

Some changes have been made relatively late in the project, e.g. final ASIC design, organisation of wafer probing, extra bond pads on the ASIC, flex-module design, thicker modules, PEB fine-tuning, etc, and though they seem to be well scrutinised and justified, must continue to be checked.

Evaluation of modules in great detail will be essential. Quite a range of environmental conditions are to be covered, and statistics so far are very limited, and on non-final components. This was noted in some of the presentations, e.g. referring to possible issues of enlarging the module thickness. The justification for this change appeared to be sound but the HGTD team correctly noted the care that must be taken to evaluate any possible unexpected effects. Up to now, it appears to be assumed, inevitably given the availability of parts, that all issues factorise to component level, and this cannot yet be considered proven. Other parts of the upgrade project, such as the ITK, have unfortunately already demonstrated this.

The production database will be crucial, and its usage must be well-tested at an early stage. The HGTD team do seem to be as ready as possible.

There has been excellent progress with PEB manufacture and assembly, which was very encouraging. The PEB itself is essentially passive, but does provide power, monitoring and control. Test results look very good, but are still incomplete, and full validation will require final modules. PEB and module assembly issues, such as assembling the dense arrangement of kapton tails and handling the modules during the assembly process, can soon be explored with final components, as well as making performance checks.

The demonstrator is now imminent, and results will be eagerly awaited, including readiness of DAQ and data analysis, and scrutiny of many details. This will be a big milestone for the project.

Some issues were raised regarding the final large-scale detector assembly. The clean room preparation seems to be well established and almost complete, but there are some uncertainties about the delivery of a mechanical support structure from Dubna, which is due in June. Time did not permit to discuss this in more detail, and clearly there are some special considerations surrounding support from Russia. In both this and previous reviews, questions were raised about the staffing of the clean room – e.g. are enough people available? Can they all be accommodated in the area? The proposal to quickly test assemblies with limited cooling must be well verified to avoid risk.

The report from the timing group was very welcome, including news about the collaborative links with other projects. It seems likely that this issue will grow greatly in importance as the full extent of important technical details becomes evident, and new challenges can be confronted. It was noted that timing resolution alone is not the only issue, as we learned about LHC machine-related variations which have come to light. This indicates the desirability of wide collaboration with others who are concerned with such high precision timing as will be available from the HGTD, which is still unprecedented on such a scale.

Use of the HGTD for physics in ATLAS was extensively discussed in the TDR (2020). It may soon be time to revisit this, to establish if all the methods and assumptions are still valid for a real system, e.g. calibration and synchronisation issues.

ITk Schedule

The ITk schedule continues to be the overall critical path driver for the ATLAS Upgrade. As presented in this review, the current float is -80 working days, a deficit of four months, dominated by the ITk-strips sensor cracking issue discussed below, which has halted progress on the critical path. As the solution has not yet been found, and the latency between proposal and validation of a solution is of order 3 months, the P2UG expects the situation to get progressively worse in the next six months.

That said, it is clear the ITk team is pursuing every avenue to solve this problem and move into production, and there is significant progress on many fronts unaffected by this issue. Despite the "no-show" Fraunhofer follow-up review, there was a substantial effort to look at many potential scenarios for acceleration, and to implement the most promising candidates. It remains the case that there is no "silver bullet"; however, for the first time in several years, the ITk-Pixel schedule seems relatively stable, chiefly due to the constant work per the "6-prong plan" to streamline and consolidate schedule in order to compensate for delays encountered. Some of the principle handles for acceleration involve reductions in Quality Control and/or shift work, each with their own obstacles – QC reductions can only happen after there is sufficient experience to have confidence that sub-standard performance will not escape detection, and additional person-power implies more training from the already overconstrained experts, as well as potentially additional funding. We continue to admonish against the reduction of a robust QC program for the sake of minor schedule gains.

Although not a major critical path driver, the services have also suffered some delivery and QC issues. Continued close attention will be required here, particularly in areas which rely on a single vendor, where additional items for further diagnostics and testing are needed, where long-term support may be required, and where there is still some effort missing, such as Common Electronics, Production Database, logistics(2 persons), FOS readout hardware, and cooling installation. In addition, there is a danger that logistical demands can divert technical/engineering efforts, so care should be taken to match jobs and personnel skills.

ITk-Common (Summary)

In response to a previous recommendation, the Production Database project has been bolstered, with 2 external engineers added, a request for EOS disk space for images, a facility for local copy, and a discussion with supplier for long-term maintenance and yearly payment process. The Database is now heavily used for uploads, registration and reporting requests.

In development of the online software paradigm, two consortia have made detailed and convincing proposals for the organisation of stage 2 of the online software development. It is then hoped that the entire community will rally around a single way forward. This determination may require Management or IB interaction, in order to preserve resource availability focused on the final system.

There is impressive work on the SR1 assembly hall preparation, with floor replacements in place, and Global structures such as the Outer Cylinder and L3 and L2 shells expected soon. The cooling plant development is also progressing after agreement on procurement strategy with CERN, the 1-head plant pressure tests were achieved in April, and the 2-head installation is targeted for September.

Finally, there is significant progress in common electronics, with some of the dedicated interlock systems already integrated into SR1, with the foreseen upgrade of the Main Interlock

Crate (MIC) and Embedded Processing Monitor (EMP), and continual progress on the environmental sensors, radiation monitoring, and Beam Conditions monitoring systems.

ITk-Pixels (Summary)

There has been a successful changing of the guard in the ITk-Pixel project, with Jo Pater seamlessly taking over the ITk-Pixel helm. The P2UG encourages the appointment of a second deputy, to help share the load of leading this non-trivial enterprise.

As mentioned above, the schedule has stabilised over recent BCPs and statusing periods, and in the last quarter there has been no significant change to the ITk-Pixel delivery date. However, there is a long list of non-trivial tasks that have to be carefully managed, in particular the link between the flex qualification and module production, as the full flex qualification needs data transmission.

The hybridisation situation has improved, with re-optimisation of the load distribution among the two validated vendors and potential minor changes to detector configuration to mitigate the hybridisation availability problem. Nevertheless, neither of the other two vendors, including the one providing the bulk of the quads, have passed PRR.

A new issue with the ITk-Pixel ASIC has emerged, the so-called “core columns” problem, where the ASIC ceases sending events until a reset is issued. It has been observed in about 5% of ItkPixV1 chips, but only after assembly in a variety of sensor types, and not in wafer testing. It remains to see if the effect is on ItkPixV2, although CMS sees a similar effect on CROC. The issue can be averted by disabling certain core columns, and is under further investigation. This is a large concern, as a resubmission of the ITkPix ASIC would clearly imply a significant schedule delay.

Finally, due to the hybridisation and “core-columns” problem, there has yet to be a demonstration that the assumed assembly rates in the production model can be realised; thus, there remains substantial uncertainty in the finish date of the ITk-Pixel project, and potential for an unwelcome surprise.

ITk-Strips (In-depth review)

The ITk strips project was presented in a five-hour plenary P2UG session, providing a comprehensive overview of the project. The presentations and discussions covered the schedule status and delays, recent technical progress and challenges, technical and schedule risks, as well as focal points for the coming year.

Major achievements were reported in several areas: The sensor production is recovering from the previous delay and has reached 74% of good sensors delivered. The ASICS projects are progressing well, with HCCStar and AMACStar at >100% and ABCStar at > 50% completed. For the modules, 90% of site qualification is achieved. The End of Substructure (EoS) and EoS-DCDC PRRs have been passed, and for the barrel, the bus tapes production has commenced, while the endcap petal cores are already in production. For global mechanics, the completion of both endcap global structures and the L3 dressing could be reported. However, several challenges have been reported and discussed in-depth. The most critical issues are the sensor cracking and noise problems observed during cold testing of operational staves and petals. Although mitigation strategies have been proposed and are being explored, the issues persist and have effectively stalled the project just before the commencement of module production. The search for solutions has been hindered by difficulties in supplying sufficient parts for cold noise and sensor cracking tests. The entire ITk strip community is

now intensely focused on addressing the problems with the highest priority. Additionally, other project areas have reported delays but are now believed to be progressing well. This is particularly true for the services, where delayed completion of tenders and contracts, as well as delayed supplies from industrial vendors, have impacted on the schedule, bringing the services close to the critical path.

Schedule

The ITk Strips barrel ready date shifted from 13.5.2026 (May 2022) to 20.11.2026 (May 2023) to 13.10.2027 in May 2024. This corresponds to a shift in schedule of 17 months over the last two years, respectively 10 months over the last year. The ITk Strips barrel and endcap need-by date is set to 25.6.2027 (unchanged from May 2022), and thus ITk Strips project holds a negative float of 80 days (4 months) against the need-by date; the ITk ready date is shifted correspondently.

The primary reasons for the dramatic loss in schedule are technical issues, specifically the "cold-noise" and "sensor cracking" problems in strip modules. These issues have stalled the project, preventing the commencement of module production.

Comment on schedule:

The present unavailability of proven solutions for the sensor cracking problem, the challenging situation for solving the cold-noise in the short strip modules and the limited potential for accelerating module production or descoping in the later phase of the project, give at present little prospects of recovering the schedule within the ITK strip project.

Sensors

The production of the 20800 sensors started in August 2021 and 74% of the sensors have been received. The delivery schedule for the Endcap sensors has been fully recovered, while the Barrel sensors are currently about 3 months behind the initial schedule. The overall quality is very good. Some issues were flagged in the QA/QC process and solved in close collaboration with the producer. In total, 5% of the sensors that were delivered were rejected and will be replaced by the producer. An additional sensor order amounting to 15% of the production is about to be ordered to cover additional needs for solving the technical issues around the cold-noise and sensor cracking.

ASICs

The ASICs production, driving the ITK project critical path in 2021/22, is no longer a concern. The HCCStar and AMACStar production is completed, while half of the 311373 needed ABCStar chips passed the final production step. In a fraction of the ABCStar chips minor flaws not compromising the operation were detected and classified as Category B. About 2% the ABCStars in the ITK tracker will be of this category.

Modules and module mounting

On the 4 barrel strip layers, 7168 Long Strip (LS) and 3808 Short Strip (SS) modules are needed, and for the 2 Endcaps, a total of 4608 modules realised in 6 different geometries are needed.

The pre-production of those modules has been heavily impacted since 2022 by a series of technical issues in the barrel and endcap modules, in particular, the "cold-noise" issue and, since 2023, the "sensors-cracking" issue. These issues prevent to-date the module production, and did not allow for gathering comprehensive experience along the module production chain and across all the participating institutions. They forced the project to revise and split up the module PRRs in a way that allowed them to proceed with orders or production of components not impacted by the technical issues. PRRs on first production hybrids and modules have been

passed in 2023, but only under the assumption that no further changes of hybrid and powerboard designs are needed to cope with the technical issues. Barrel LS and Endcap module PRR follow-ups are acting as blocking points on the critical path, awaiting the resolution of the sensor cracking issue.

The origin of the cold noise was traced back to mechanically vibrating capacitors on the power boards that are glued on top of the sensors. Mitigations for the Endcap modules (re-design) and the Barrel long strip modules LS (change of glue type) were found in 2023. For the Barrel short strip modules SS, which have a higher load on the power board compared to the LS modules, some cold-noise remains after the glue change. Further mitigation techniques are thus being investigated (stress reduction, "anchoring of glue edges"), and beam tests to better evaluate the operational limits of modules with different levels of cold-noise are under way.

Further studies on the "early breakdown" behaviour of sensors on staves in cold tests resulted in 2023 in the finding that some sensors fracture when cycled to low temperatures ("sensors-cracking issue"). The sensor cracks are found close to the power boards, only occurred for modules mounted on staves or petals, and amounted up to about 8% of broken sensors on the staves subjected to thermal cycling. An impressive amount of analyses and testing have been performed, and a simulation has been developed that is capable of delivering qualitative predictions. Mitigation techniques were developed in form of using an alternative glue ("Hysol mitigaton"), increasing the separation between hybrids and power boards ("wide gap modules"), using thinner hybrids and, as a more significant change to the module design, placing interposers. So far, no viable solution has been found, although not all tests are concluded.

While the technical issues prevent the start of production, the module sites are approaching full production readiness. Module site qualification has reached a level of 93% of all reviews completed and significant effort has been spent on optimising production ramp-up, part flow for production and evaluating scenarios for descoping in QC and overall production speedup. The latter is difficult due to the limited experience in exercising production-like assembly rates.

Procurement / production management / logistics / databases

The production management team works on providing tools to streamline logistics and optimise shipping processes and costs. In anticipating scenarios of accelerated production ramp-up and/or increased production rates potential bottlenecks in throughput and/or storage needs have been evaluated and mitigated proactively. Additionally, the production and procurement databases have been further improved, and a database backup system has been implemented. The simulation tool for modelling production has also been further refined.

EoS and Local support

The End of Substructure (EoS) project is advancing well. About 85% of the EoS card PCBs are produced, and fully populated EoS cards are starting to enter the QC procedures. Also, the production of the 1900 EoS DC-DC PCBs, required to power the EoS cards, was launched in April after successfully finishing the pre-production of 250 units. A point of attention is the observation of failures of the lpGBT at low temperature testing (-30°C), which is now under investigation at DESY and with the CERN-ASIC team.

The LS tapes and LS cores went through a PRR in July 2022, with the petal tapes and cores passing with recommendations and the stave tapes passing later after follow-up actions. The stave cores passed only conditionally and still must pass a PRR follow-up. The stave local flatness turned out to be challenging, but a revision of the specifications seems to offer a feasible solution to pass the PRR and to proceed towards production.

Comment on LpGBT: A small amount of lpGBT issues have been experienced when operating at low temperature. These issues need a thorough explanation as any problem would have a high impact on the entire HL-LHC.

Services, Power suppies, Global mechanics and Integration

The services WBS covers the connectivity from the stave/petal (PP0) up to the service cavern (PP4) and is charged to produce the service modules and several different types of custom cables and connectors. Progress was achieved in most areas, but further delays also resulted: Type I cables have been produced and delivered and passed the Construction Products Regulation (CPR) test. A fraction of the PP0 (27%) and PP1 (60%) connectors was produced. Unfortunately, for the PP0 connectors, an issue was detected that led to production stop and the need to repair the already produced connectors at the supplier, which is expected to last a few weeks. For the Type III connectorized cables, pre-production was stopped for several weeks due to a production mistake of the PP2 connectors (male instead of female pins) but has now picked up again. Overall, several milestones ("Type II/Type III cables ready for installation", "Type II and optical link PRR", "PP3 FDR") are delayed due to procurement and technical issues and Type II and Type III services are at present left with less than 10 working days against the critical path. They are needed to start integration; however, Type II cables are not expected to arrive before Q4/2024, and thus, a further float reduction is expected within the present schedule.

For the power supplies, the invitation to tender was closed in June 2023, the contract was signed with the vendor of the winning bid, and the order is out with only very minor changes w.r.t. the previous prototypes.

Global Mechanics advanced well. Both Endcap structures have been completed and bulkheads and cooling serves are being added. For the Barrel the Outer Cylinder is ready and about to be shipped from LBL to CERN, the L3 is already at CERN and dressed, while the L2 is being dressed in Oxford. It is expected that the project is ready before arrival of production staves and petals.

The site infrastructure to support integration is largely in place at SR1 at CERN and progressing well in the two Endcap sites. The insertion tooling for the petal insertion at the two EC sites is in place, and the one for the staves is well-advanced. The setups for stave and petal reception are approaching commissioning and are assumed to be ready in one to two months.

Recommendations from the ATLAS P2UG review

Trigger/DAQ

- At the next in-depth review in November, please present in detail the consolidation of the NSW-TP effort, including the FPGA choice and firmware development plan. An interim report on the occasion of the next intermediate meeting would be very much appreciated.
- At the next intermediate meeting, please give an update on the situation with the DAQ readout, in particular the outcome of the BCP.
- At the next in-depth review in November, we would like to see more details about the L0Global startup plan.

TileCal

- Ensure that the Daughter Board group has enough human resources to complete the full evaluation of v6.4 in a timely way. Integrated tests of the full online-to-offline chain should not be delayed further by the readiness of the Daughter Board.
- Continue efforts to identify a group/institute to take over the lead responsibility in delivering the CS-Calibration project and the new sharing of responsibilities should be formalized through an addendum to the MoU.

Liquid Argon Calorimetry

- For the P2UG November meeting, update the FDR schedule across the project with realistic dates.

Muon

- Understand the RPC gas leak problem before proceeding with production. Consider setting up more stringent leak testing at the assembly site. Consider retesting chambers over months to see if a change in leak rate occurs.

HGTD

- Ensure the wafer probing capability is fully ready, including automation and data analysis.
- Confirm readiness across the project to handle the next phase of module assembly and integration activities, which must ramp up rapidly. This includes data acquisition, analysis, and tracking using the production database.
- Confirm the availability of the mechanical integration structure from Dubna, and ensure there will be a backup if necessary.
- Verify the electronic testing procedures on large scale assemblies can function with appropriate cooling without risk.

ITk Common

- The major slippage in the schedule can lead to logistic issues at the start and end of the project timelines, therefore;
    - We encourage the management to actively continue to pursue direct discussions with the funding agencies concerning the need for contract extensions, infrastructure booking and maintenance. Where it is clear that the project delay results in a slow down or delay in a sub activity, this should be reflected in the schedule.
    - The potential "standing army" issues in the project should be actively reviewed, to make sure that no opportunities are missed for project consolidation or cross support or temporary redirection.

- The reinforced effort on the database is very welcome and it is clear that the demands on the database are increasing as the usage ramps up. The committee recommends that every effort be made to encourage interactions between the database engineers and users, in order to ease the introduction of user requirements and enhance visibility of the work. We would like to see a dedicated slide on the status of the database reporting in the next review.
- In areas of the project where there is dependence on one supplier for a very custom component (examples include the surface treatment of the diamond sensors, connector manufacture,…) ensure plentiful spares as there may be no further supply.

ITk-Pixels

- Expend all resources possible to determine the cause of core column problem, including collaboration with RD53 and CMS teams.
- For the Fall P2UG, be prepared to compare the preproduction assembly rate and yield with the assumptions used in the production model once the ASIC and hybridisation issues allow preproduction to ensue.

ITk-Strips

- Keep working with the highest priority on solving the sensor cracking problem. Ensure that sufficient parts are available for cold-noise and sensor cracking tests and take measures to replenish the parts needed for the final production. A sign-off procedure should be established in advance for any acceptance of a solution, stating exactly which criteria have to be satisfied including long-term performance, performance after irradiation, and operational windows.
- Persist in addressing the Cold Noise Issue for Barrel Short Strip modules. Evaluate the impact of cold-noise on operation, i.e. clarify how much cold-noise could be tolerated.
- Make concerted efforts to move the services further away from the critical path, per recommendation IS-10-2 from November 2023.

B. Summary Report of the CMS P2UG review

F. Bedeschi, A. Cardini, M. Demarteau (chair), V.V. Gligorov, A. Kluge, G. Kramberger, M. Martinez-Perez, C. Roda, M. Stanitzki, D. Zerwas.

The CMS P2UG met with the upgrade project on 14–16 of May 2024. This P2UG meeting the Tracker and HGCAL projects were reviewed in-depth. Regular status reports were given by BRIL, Muon, BTL, Barrel Calorimeter, L1 Trigger, DAQ and HLT. Clear and steady progress is made in all technical areas of the upgrade and the team is to be congratulated with its continued success in realizing the project. We note the completion of the BTL Tracker Support Tube (BTST), the successful submission of the CROCv2 for the inner tracker, completion of CE-H raw steel production and the successful irradiation campaign of ME0, reaching the target irradiation dose with no degradation of gain observed.

The tracker project carried out a Fraunhofer-like exercise that has provided valuable insight in the production schedule and enabled the inner tracker project to gain schedule contingency. After these Fraunhofer exercises, the schedules for both the HGCAL and Tracker seem to be highly optimized, are in several areas aggressive and success-oriented and are no further compressible. Any further compression would invariably lead to quality losses and the performance of the detector would be compromised.

The project has unfortunately encountered another setback, this time with the HGCROC chip. At the previous P2UG meeting, it was reported that there was a delay of potentially up to 3 months in the availability of the HGCROC due to a communication issue between IMEC and TSMC. V3b of the HGCROC was received in February of this year and unfortunately had two major bugs. The team is aggressively addressing these issues and a delay is unavoidable. The schedule float for the HGCAL project is only one month, with approximately three years to go before the scheduled installation of the first endcap.

The Tracker project has made excellent use of the Fraunhofer-like exercise and now has an excellent understanding of their schedule and of all dependencies during production. The IT project has regained four months of float due to this exercise The production speed of the OT is driven by the availability of hybrids and MAPSAs. Production of service and mechanical parts, however, are in the shadow of the critical path. The OT has a schedule contingency of about 2 months. Hybridisation remains the most critical item for the Inner Tracker. Both the HGCAL and Tracker projects are in urgent need of more personnel.

The Tracker and the HGCAL are the two biggest sub-projects of the CMS Phase 2 upgrade and drive the overall schedule. Currently, each project has a schedule float of about one month. Given the very limited, effectively zero, schedule contingency that the project now has, based on a schedule that is highly optimised, the current schedule is deemed untenable.

The BRIL project has concluded a successful validation of the performance of the ASIC for the Fast Beam Conditions Monitor, but significant noise was observed in test beam operation of an FBCM prototype, which requires a redesign of the Front-End module board.

The Barrel Timing Layer is moving into pre-production. LYSO crystal production is well underway, and the first production tray will be assembled at CERN in July of this year. A major milestone was the delivery of the BTL Tracker support tube. The Endcap Timing Layer has validated the performance of an LGAD assembly readout with the ETROC2

chip and obtained a timing resolution of 45ps. Functional and performance tests of the ETROC2 have been completed, including radiation tests. Although triple modular redundancy (TMR) has not been fully implemented in the chip, no adverse effects have been observed, which may obviate the need for another version of the chip.

The project has prioritized completion of ME0 in the muon project. The schedule for module production is driven by PCB availability and procurement and the rate of foil production. It is estimated that by the end of 2024, the production will be limited by the availability of foils. Foil production at CERN and KCMS is progressing well. No drop in gain has been observed in the ongoing ME0 longevity studies, which is very encouraging. The project is on a very tight schedule and its success depends on the frenetic production of multiple components during 2024.

Chamber production of the iRPCs is on track at Ghent and CERN; one endcap has been completed and the assembly of the second endcap has started. As reported at the previous P2UG meeting, due to procurement delays, the production of the Front-End Board for the iRPC system, FEB V2.3, has seen a delay of 8 months and the full delivery of FEB V2.3 boards is now expected by summer 2024. As a consequence, the schedule for iRPC production is defined by FEB availability, with all FEBs expected by September 2024. Completion of all chambers is planned for November 2024 with only a one-day quality control test per chamber. This is a very aggressive schedule.

The Barrel Calorimeter has observed an unexpected 50% increase in LiteDTU power consumption. The team has a clear plan to address this issue and understand its origin. No significant schedule delay is expected. The Level 1 trigger, high-level trigger and DAQ projects are all making very good progress.

Project Specific Observations:

The Beam Radiation, Instrumentation and Luminosity (BRIL) team has made very good progress on the three detector systems that are needed on day 1: the Beam Condition Monitor for Losses (BCML), the Fast Beam Conditions Monitor (FBCM) and the BRIL Trigger Board (BTB). The readout ASIC for the FBCM has been received and meets all specifications. The successful performance of the FBCM FE ASIC on the first try is to be commended since it obviates the need for a second version of the chip. However, in testing the beam operation of an FBCM prototype a large pickup noise was observed. This has been attributed to the long traces from the sensor to the ASIC and a redesign of the Front-End module board is required. The BRIL Trigger Board (BTB) team has established an effective cooperation with the inner and outer tracker teams for firmware and hardware development and are participating in common testing stations. Progress is good, but overall, the BRIL program is marginal in person power.

The BTL is moving into module production. Production of the LYSO crystals is well underway. Delivery of the SiPMs for the readout has been delayed by 3 weeks but is imminent. The packaging of the TOFHIR chip has also seen some delay, but the acceptance rate is good. The first production tray will be assembled at CERN in July. A major accomplishment has been the completion of the BTL Tracker Support Tube. Rail installation for the BTL is scheduled to start in mid-July. The schedule for the BTL is adequate with a float of 3 months with several items very close in driving the critical path.

The Endcap Timing Layer has validated the performance of an LGAD sensor assembly read out with the ETROC2 chip in a DESY test beam and obtained a timing resolution of

45ps. Functional and performance tests of the ETROC2 have been completed, including radiation tests. Although triple modular redundancy (TMR) has not been fully implemented in the chip, no adverse effects have been observed, which may obviate the need for another version of the chip. Vendors for both the high and low fluence regions have been identified in the market survey: three for the low fluence region (80% of sensors) and two for high fluence region. A procurement readiness review has been scheduled for the LGAD sensors for July 3rd, 2024. A market survey with three vendors for hybridization, including wafer post-processing for LGADs (thinning, UBM, dicing) and ETROC and bump-bonding, is in progress. With the confirmation of the technology in the test beam (σ_t ~ 45 ps/hit), procurement should not be delayed. Hybridization remains a point where more experience is needed; it is the critical path of the ETL. There is a tight constraint on the overall envelope of the ETL package and the final layout of the ETL has not yet been decided. This layout is independent of the module design.

The project has prioritized completion of ME0 in the muon project. The schedule for module production is driven by PCB availability and procurement and the rate of foil production. It is estimated that by the end of 2024, the production will be limited by the availability of foils. The project is on a very tight schedule and its success depends on the frenetic production of multiple components during 2024. The project has lost one month of float since the last P2UG meeting.

Good progress has been reported in all areas for the muon drift tubes (DT). An engineering system review (ESR) follow-up was completed in April 2024. Work on the mechanics and cabling is progressing satisfactorily. The OBDT-phi modules are ready for production and the OBDT-theta are being tested. A slice test has been carried out that pointed to sensitivity to possible overvoltage affecting the boards. Protection measures have been implemented. Longevity studies at GIF++ have been concluded. The reduced efficiency of the DT observed in 2017 is understood in terms of carbon deposits on the wires. This is not present in newly tested modules. Although the origin is uncertain, the project is confident that this will not repeat itself in the future. Excellent results were presented on the reduced dose in the wheels thanks to new shielding in P5. A rate reduction up to a factor of two in MB4 has been achieved. The project is to be commended for the excellent progress.

Longevity studies of the Cathode Strip Chambers (CSC) have been carried out with nominal gas and with 10% CF4 for the expected HL-LHC total dose. No signal degradation has been observed in any of the irradiation runs. Irradiation tests at GIF++, with 5% CF4 in the chambers, will continue until mid-2024. Readout board production is going well, with full production of the boards expected to start by the end of the year.

iRPC chamber production is on track; one endcap has been completed and assembly of the second endcap has started. The schedule for the iRPC production is defined by FEB availability, which has experienced significant delays. All FEBs are expected to be delivered by September 2024 with completion of all chambers in November 2024. The team considers a 1 day/chamber quality control test sufficient.

Extensive longevity studies are ongoing. As noted above, the GEM-based ME0 chambers reached the target irradiation level with no observation of degradation of the gas gain. The CSC ME2/1 with 5% CF4 have been exposed to 435 mC/cm, which corresponds to >3x the dose expected during the HL-LHC running, with no degradation observed. RPC longevity studies with eco-friendly gas mixture studies continue with the selected candidate gas ECO2: 35% HFOze + 60% CO2 + 1% SF6 + 4% iC4H10. A mixture of 30% CO2 + 0.5% SF6 is also being studied. No degradation of RPC performance has been observed.

The barrel calorimeter project has seen an unexpected 50% increase in LiteDTU power consumption. A three-step plan has been presented to understand this effect including mitigation techniques on the system side. The team plans to validate the existing chips for use in the detector; identify the origin of the increased power consumption; and procure more v3.0 chips. Given that it is possible that the additional power consumption is due to process variations it is proposed to produce an additional 4 wafers (40% of full quantity) to verify if it was a one-time problem. The team is to be commended for having developed a well-thought-out plan to address this issue, even though it introduces some delay. Significant effort is underway on the FW design for the Barrel Calorimeter Processing board.

The HGCAL project has made great progress over the last six months. Ever so important system tests continue at an excellent rate. Many, but not all, components are being finalised and there remain important gaps. Unfortunately, the project continues to incur delays; over the last six months a delay of three months has been realised and the current schedule contingency is only about one month, insufficient to meet the baseline schedule. Furthermore, significant vulnerabilities remain, both internal and external to the project. Although the start of production is in sight, it is the P2UG's opinion that the project underestimates the time it will take to get to production.

The project has unfortunately encountered yet another setback with the HGCROC. At the previous P2UG meeting, it was reported that there was a delay of potentially up to 3 months in the availability of the HGCROC due to a communication issue between IMEC and TSMC. V3b of the HGCROC was received in February of this year and unfortunately had two major bugs: an error in the transition of the energy measurement from ADC to ToT and sporadic bit outliers. Both bugs have been reproduced in simulation and the strategy for the fix is being worked out. The team is currently forced to consider the use of the HGCROC.v3b chip for pre-production modules and use them in regions of low energy deposit in the final detector corresponding to at most 5% of the total production. If production and yield are favourable, the fraction of faulty HGCROC.V3b chips in the final detector can be reduced. The team is aggressively addressing these issues. The P2UG agrees with this strategy given that in at most 5% of the detector these faulty chips will be deployed where the risk is further mitigated through strategic placement of the modules where the energy deposition is lowest. The ECON-T2/D2 chips have just been received but not been tested yet.

Sensor production is going very well; 44% of full silicon sensors have been delivered. Board design is also proceeding well, though not all boards have been designed yet. High-density (HD) engines and wagons still need to be designed. Five module assembly centres (MACs) have each produced 10 to 25 modules, which is an important achievement. Not all flavours of the silicon modules have been made yet. The current assembly yield is about 70% (95% required during production), which is being addressed.

A significant increase in the production of the stainless-steel absorber plates for the CE-H has been realised. The market survey for the absorber plates for the CE-E has been closed. The procurement of the CE-E copper cooling plates, a responsibility of LLR, is to be launched in two weeks. The process for the CE-H cooling plates is just starting, with the bidding to start in June 2024. The nominal production rate is expected to be reached in June 2025 for the CE-H. Cassette production is to proceed at a rate of 6 cassettes/2wk for CE-E and 21 cassettes/2wk for CE-H. It is noted that the pre-series exercise has not tested the CAF extensively.

The production of the CuW and carbon fiber baseplates, for which the procurement strategy is still being developed, may become a bottleneck and could start driving the critical path. Close attention is required.

The HGCAL team is to be commended for carrying out many system tests, both on the bench and in test beams. Good progress was reported with the V3 system tests, but not all components were part of the chain; for example, no HD wagons were included. Studies were carried out in a beam with 2 LD V3 systems equipped with 1 LD module. Further beam tests are being planned for August and September that will include 1-2 layers with 3 LD silicon modules, 1 layer with HD modules and 1-2 layers with tile modules. ECON-D chips will be included in this test. An LD test with 6 X_0 in front is planned, as well as a test in a magnetic field. The P2UG very much welcomes these extensive tests.

The tracker has made excellent use of the Fraunhofer-like exercise and is moving towards production. The inner tracker has regained four months of float thanks to this exercise. The subproject now has full command and an excellent understanding of its schedule and of all its dependencies during production. ASIC and sensor deliveries for both IT and OT are in very good shape. The project organisation is gearing up for production with the creation of a Tracker Integration Coordination Area at CERN and an Outer Tracker Module Production Office.

The production speed of the OT is driven by the availability of hybrids and MAPSAs. Despite best efforts, the project has lost around three months since the last meeting. In particular, the OT schedule float is getting close to zero. It is noted that in several places some additional manpower could have a big impact and make the schedule significantly more robust. A tendency that has to be closely watched is the progress with general services, such as cable procurement and mechanics, that are trending to getting closer and closer to the critical path.

The hybrid kick-off batch for the OT has been completed with a yield that has been below the targeted 95%; this issue is being addressed but needs to be monitored. The second batch of the pre-series is now in production. Pre-production of the MaPSAs is ongoing with excellent yield at HPK (94%); there are, however, issues at QPT. All QPT pre-production MaPSAs had bad IV behaviour with two different failure modes (early channel breakdown or continuous current rise). Another setback has been that the dicing vendor did not use wafer maps for thinning and dicing of the MPAs for batch 2 (worth > 200 MaPSAs or several thousand MPA), which implies that bad MPA chips could be used in the MaPSA production.

Module production is proceeding, but the number of modules built is still limited by the number of available hybrids. Final design changes have been implemented (such as a “ground balancer”), which are great steps to consolidate (and improve) the design. Preparation of the 2nd TB2S ladder assembly centre at CERN is ongoing to alleviate schedule pressures, which is a great move. Preproduction for all major components of the TBPS, TB2S and TEDD has started. We note that the end of the Tracked Endcap Double Disk (TEDD) Dee production is forecast to be November 2026, the same as the PS modules, which puts it very close to the critical path.

For the Inner Tracker, the hybridisation remains the most critical item and remains the biggest concern for the project. At the moment only IZM is fully qualified while Advafab and Micross still need to be fully qualified. Contracts have been signed with IZM and Advafab, with pre-production orders placed. Module assembly sites and testing sites are well advanced, and the large overcapacity offers great redundancy. First pre-production

at pilot sites is expected for August 2024. We note that some failures have been observed for extended temperature range thermal cycling of inner tracker modules.

As far as the readout is concerned, a “core column issue” has been observed on some multi-chip CROCv1 modules, which manifests itself as non-performing (<=8) pixel columns. The hypothesis is that this is due to dust contamination during dicing. The origin, however, is not clear, but the impact at this stage seems limited to <<1% of the total, which is acceptable. There has been great progress on the systems tests of the TBPX, TFPX, and TEPX. The tests with serial power chains are achieving their targeted values. The team is to be commended for launching an ambitious but strongly supported integration activity of a pre-production slice, planned for 2025, equivalent to 1/9 of the final system. Ensuring adequate on-site person-power from participating institutes will be crucial to support activities at the Tracker Integration Facility.

There has been good progress in the area of the L1 Trigger, HLT, DAQ and TCDS system. The DAQ, HLT and TCDS teams are relatively small team but effective.

The DAQ team is facing a shift in the availability of the new control in the building 904 Integration Center. This results in an unfortunate delay of 5 months with respect to the original schedule. The cause is partially due to the longer time to complete the infrastructure (electrical installations). The team is commended for trying to find ways to gain back time by taking actions in parallel where possible.

The HLT is making very good progress. The person-power shortages are concerning. It is important to maintain person-power at the required levels to be able to follow external technology evolution over the next years and tune the software to best exploit it. Furthermore, the team is strongly encouraged to retain a good fraction of the people who built the system (and in particular the performance/heterogeneous computing specialists) for commissioning and maintenance in particular during the early phase of Run 4.

The full report by the CMS P2UG is available as a separate report.

Recommendations from the CMS P2UG review

BRIL: Fast Beam Condition Monitor

- Work on the redesign of the frontend hybrid to address the observed noise issue.

BRIL: Beam Condition Monitor for Losses

- None.

BRIL: Trigger Board

- Continue and expand cooperation with IT and OT teams in the development of BTB for the cluster/stub testing.

BRIL: Neutron and Radiation Monitoring

- None.

Muons: Drift Tubes

- None.

Muons: Cathode Strip Chambers

- None.

Muons: GE2/1

- None.

Muons: ME0

- Carefully monitoring the PCB quality and the production schedule of all the electronics boards.
- The project is under a very tight schedule. We recommend again to scrub the overall schedule to look for opportunities to accelerate the production schedule without compromising quality assurance.

Muons: iRPC and iRPC Link

- Carefully monitor all steps of board production/testing activities to ensure sub-project stays on schedule.
- We recommend the team to re-evaluate that 1 day/chamber QC test duration will meet the requirements to guarantee high Q of the chambers in the long term.

Muons: Eco-friendly gases and Longevity Studies

- Continue the important work towards meeting Eco goals and restrictions.

MTD-BTL

- None.

MTD-ETL

- Proceed as quickly as possible qualifying hybridisation vendor(s).
- Fully validate the ETROC2 chip, including irradiation studies with adequate safety margin. If no errors are observed, move directly to production and skip the ETROC3 phase.
- Finalise the layout as soon as possible.

HGCAL

- Prioritise detailed HGROCv3c and ECON-T/D verification with respect to early submission.
- Ensure that all fixes for the known issues are implemented for the HGCROCv3c and are fully tested (including simulation); ensure all observed issues are addressed before submission.
- Engage the CERN ESE team to ensure that the HGCROCv3c is okay before submission.
- Extensively test the ECON-T/D chips with help from ESE before the go-ahead for production.
- Finish the design of all boards with priority.
- Complete thorough testing of all boards with necessary system tests before production.
- Understand the failure modes in the assembly process and see if rework is possible.
- Qualify the 6^{th} Module Assembly Center.
- Prepare the sites for staffing to have “the right technical level at the right time” and demonstrate that the needed production rate can be reached at all sites.
- Define the CuW baseplates procurement strategy expeditiously.
- Continue to aggressively pursue the pre-cabling of YE1.
- Continue to anticipate the resource needs as function of time to have a smooth production.
- Evaluate the impact of a potential absence of a contribution of currently available personnel (especially for the expert knowledge potentially no longer available) and develop a mitigation plan.

- Map the identified needed resources (person-power, ...) to available resources and make sure that it is an institutional commitment.
- Develop a tight collaboration with the system test group, which is necessary to ensure that the detector is working after assembly.
- Ensure that the cassette assembly is fully qualified as soon as possible, incorporating its readout as part of the system tests.
- Given the current state of the schedule, the HGCAL project should study how much additional float would be necessary to comfortably deliver the project.
- Continue the system tests, increase the staff involved in the system tests, and increase the scope whenever possible, including planning with the CAFs.
- Consider the possibility to demonstrate the use of the HGCROC in the full energy range, especially in the ADC to ToT transition with real showers by inserting a variety of X_0 using the existing setup to mimic the shower development on a single module.
- Assess the possibility, when HW is available, to test in the test beam systems that have been extended longitudinally.

Tracker

- Involve all project institutes to make sure, that the tracker integration effort is adequately staffed.
- Conclude the procurement process for both power supplies and cables as quickly as possible.
- Given the current state of the schedule, the Tracker project should study how much additional float would be necessary to comfortably deliver the project
- The Tracker project should identify and pursue ways to increase the overall OT module throughput, whether that be by increasing capabilities at existing sites, leveraging capabilities of IT module sites, or even by bringing in a new wire bonding site.
- Remain very vigilant regarding hybrid production and act proactively.
- For the next P2UG we'd like to see the hybrid yields using both the old and new criteria.
- Establish testing capacity for the MAPSA in Europe at DESY and Padua.
- It is crucial for the project to decide on how to proceed with OT MaPSA production considering the difficulties at QPT. Initial steps for a potential vendor transition should be taken now given that the process for doing so is likely to be involved.
- Reassess the ASIC stock for the MAPSA to compensate for dicing problems.
- Make the best use of the pre-series production to train all module centres.
- Start with the site visits as soon as possible and keep monitoring the production logistics.
- Increase statistics of module testing and test also larger objects (ladders) with enough modules to spot possible surprises timely.
- Complete the 2nd ladder assembly centre at CERN as soon as possible and ensure appropriate staffing.
- Closely monitor the TEDD production and identify schedule mitigations.
- Ensure that steps in the qualification of hybridisation are properly evaluated with sufficient statistics for all variants of the modules.
- IZM is currently unable to cope with both CMS and ATLAS production demands as they are known today. This limitation might be overcome by improving communication between the CMS IT and ATLAS ITk pixel projects

and through direct engagement with Fraunhofer from CMS, ATLAS, and CERN management.
- Check for the origin “core column issue” carefully in CROCv1 modules, particularly if they show also in CROCv2 modules with slightly different dicing.
- Ensure full phase space of system is adequately explored before moving to production.
- Continue as fast as possible with the prototyping and strive to reach final designs as soon as possible.
- Keep monitoring the schedule and avoid schedule slippage.
- Strongly engage with the tracker institutes to ensure adequate staffing both at the assembly sites and on-site at CERN so that the systems get fully exercised since testing and feedback at assembly sites is crucial.

Barrel Calorimeter

- Study the cause of the increase in LiteDTU power consumption before taking the decision to modify the readout (PCB) system configuration.
- Produce the full complement of LiteDTU chips (100%) compared to the proposed 40%.

Level 1 Trigger

- Continue to evaluate the issue with the Firefly x12 25G optics.

DAQ

- Fully evaluate the implications of the delay in the availability of the control room and continue efforts to parallelise tasks.

HLT

- Explore how the physics validation machinery can be used to pre-emptively evaluate reconstruction robustness against hit efficiency drops in detectors.
- Review the TDR schedule for reconstruction components based on the experience over the past years (for example the Patatrack integration reported at this meeting, but also other relevant items), and consider whether there are enough people available with the required expertise to achieve the TDR goals on the timescale foreseen. Present this review and potential mitigations/actions (if any) at the next P2UG.
- At the next P2UG, report on the evolution of computational performance, memory usage, and physics performance using the new automated testing machinery and compare it to TDR expectations at this stage of development.