B-factory Programme Advisory Committee
Full report of
the Focused Review Meeting

30 June and 1 July 2025, Remote meeting

G. Corti (CERN), M. Demarteau (ORNL), R. Forty (CERN),
S. Gori (UCSC), G. Isidori+ (Zurich), P. Mato* (CERN),
F. Meijers (CERN), N. Neufeld (CERN), A. Petrov* (USC),
B. Ratcliff* (SLAC), M. Sullivan (SLAC), H. Tajima (Nagoya),
O. Tajima (Kyoto), M. Titov* (Saclay)
and chaired by T. Nakada (EPFL)

+ Partly absent
* Expert member.

17 September 2025

1 Short summary

A focused review meeting of the B-Factory Programme Advisory Committee (BPAC) was held remotely on 30 June and 1 July 2025. Presentations were delivered by members of the KEK accelerator groups and the Belle II collaboration on the status of the accelerator complex and the Belle II detector. They are currently in preparation for the upcoming data-taking period in autumn. Progress on the upgrade plan was also discussed.

The BPAC is impressed by the work being carried out and acknowledges with appreciation the high quality of the presentations. The committee wishes to express its appreciation for the continued dedication of the KEK accelerator groups and Belle II collaboration, and looks forward to hearing the progress during the next meeting.

In this section, the committee provides feedback on the two specific issues raised by the management of the Institute of Particle and Nuclear Studies. A more detailed report is provided in the following sections.

Assess the progress of the preparation for the coming 7-month long operation and identify areas for further focus.

The upcoming run, starting in autumn 2025, marks a critical phase for the KEK B-Factory programme. For SuperKEKB, it is essential to clearly demonstrate the potential to reach peak luminosities well above 10^{35} cm^{-2}s^{-1}. This will require further increases in beam currents and reductions in beam spot sizes at the interaction point, along with improved efficiency and stability for the beam injection. These enhancements must be implemented while maintaining safe machine background levels for the detector components.

The committee finds that the progress made and the planned work for the SuperKEKB rings and injection chain are adequate to meet these challenges. Once the rate of Sudden Beam Loss (SBL) events is significantly reduced, through the cleaning of the beam pipes from Vacseal contamination, systematic machine studies can be carried out to better understand the accelerator's behaviour under various parameter settings. While the proposed hardware changes in the injection chain may not address all existing issues, they are expected to bring notable improvements to operational performance. The committee also appreciates the continued efforts to reduce machine background and to protect the Belle II detector from uncontrolled beam losses, especially through fast beam abort systems. The strong collaborative framework between the accelerator teams and the Belle II collaboration should be maintained during the current shutdown and throughout the upcoming run period. The ongoing simulation effort to reproduce observed beam backgrounds also remains a critical activity.

On the detector side, the ongoing hardware activities are addressing issues identified during the previous run. Upon completion of these efforts, the committee believes that the Belle II detector will be ready for the upcoming seven-month run. In addition, the improvements presented in the data acquisition system, trigger, and data processing will enable Belle II to collect data more efficiently and support prompt physics analysis. Given the large number of concurrent activities, the committee advises the Belle II management to carefully monitor the allocation of resources, particularly with regard to qualified personnel. The operation of CDC at high beam currents remains a major concern for the coming run and warrants close attention.

The current plan for the coming run prioritises achieving an integrated luminosity of at least 425 fb^{-1}. To this end, two possible scenarios have been presented:

A Baseline plan targeting a peak luminosity of 10^{35} $\text{cm}^{-2}\text{s}^{-1}$ with an accelerator efficiency exceeding 67%.

B Alternative plan, targeting a peak luminosity of 6×10^{34} $\text{cm}^{-2}\text{s}^{-1}$ with an accelerator efficiency exceeding 85%.

In both cases, 80% of beam time is allocated for physics data taking. The committee recognises the importance of the integrated luminosity goal set by the Belle II collaboration to deliver significant physics results surpassing those of Belle and BaBar, particularly in core B physics measurements. However, achieving or approaching a peak luminosity of 10^{35} $\text{cm}^{-2}\text{s}^{-1}$ is an essential step for the final goal of the project. The committee therefore urges the accelerator and Belle II teams to engage in further discussions to refine the run plan and develop a shared strategy that accounts for the two objectives: peak and integrated luminosities. In this context, a stepwise approach, with agreed-upon intermediate targets for peak and integrated luminosities, could be a practical solution. Above all, all stakeholders must remain flexible and responsive to evolving operational conditions.

The upgrade options need to converge in one year to determine a baseline for the upgrade and to summarise it in TDR in 2027. Review the studies planned for the year and the steps toward TDR 2027, and advise on those plans.

Flavour physics remains one of the least understood areas of the Standard Model. In recent discussions at the Open Symposium for the update of the European Strategy for Particle Physics, flavour physics was identified as a key subject in evaluating the next flagship project at CERN. In this context, a credible demonstration that SuperKEKB can collect 50 ab^{-1} of data by early 2043 would be highly impactful. Furthermore, SuperKEKB's performance will serve as a benchmark for evaluating the feasibility of future high-luminosity circular colliders. Achieving the peak luminosity target of 6×10^{35} $cm^{-2}s^{-1}$ would be a major milestone for the global particle physics community.

The committee is pleased to see steady progress in defining the LS2 upgrade plan. Establishing a baseline scenario for the new final focusing magnet (QCS), including agreement on its dimensions, is an important step. Concentrating upgrade efforts on the vertex and tracking detectors, while assessing the continued viability of other subsystems for post-LS2 operation, is a sensible approach.

That said, the proposed timeline, i.e. completion of a detector Technical Design Report (TDR) by autumn 2027, construction starting in 2029, and installation beginning in 2032, is highly ambitious. While some benchmark studies have begun, a substantial amount of work remains to reach the level of detail required for a full TDR. A dedicated and sufficiently staffed team will be essential to this process. The committee also notes that a comprehensive and convincing overall project management plan was not presented and encourages the preparation of such a plan as a matter of urgency.

The committee also emphasises the need for synchronisation between the detector and accelerator upgrade planning. The mechanical design of the vertex and tracking detectors depends critically on the confirmed position and design of the QCS and other constraints from the machine around the interaction point. A joint accelerator–detector upgrade plan, including a resource loaded schedule is essential for informed evaluation and planning. Establishing an upgrade project office should be considered.

Finally, achieving the peak luminosity goal will require a comprehensive performance review of the entire accelerator complex, including the injector system. Necessary upgrades should be identified and planned accordingly.

2 Accelerator preparation

2.1 Status

Compelling evidence from last fall indicates that residual VacSeal that had gotten into the vacuum chambers is responsible for at least some and perhaps many of the Sudden Beam Loss (SBL) events. A particular flange type (MO-flange) often had VacSeal applied to it when the flange was hooked up. These flanges are used in the IR area and in the LER wiggler sections. All of the MO-type flanges in the IR straight section have been opened, inspected and cleaned. Many were found to have black stains from VacSeal residue. The three wiggler sections of the LER (D10, D11, D04) also contain many MO-type flanges and all of the flanges in D04 and D10 have been inspected and cleaned. Work is still ongoing to inspect and clean all of the MO flanges in D11. The HER also has some MO-type flanges, and these are also being inspected and (if needed) cleaned where VacSeal was applied.

In addition, the damaged collimator heads in the LER have been replaced. The top jaw of the Non-Linear Collimator (NLC) which was made of Ta has been replaced with a Ti jaw. This is a test case of comparing the durability of a Ti jaw verses a Ta jaw. A collimator located in the arc before the Oho straight section (wiggler and RF) has been moved to the arc before the Tsukuba straight section (IR and Belle II) to improve the protection of Belle II and the super-conducting final focus (QCS) magnets. A collimator in the HER has been making local pressure bursts either through Higher-Order Mode (HOM) heating or arcing and has been replaced with a water-cooled version.

The radiation shielding for the Oho straight section which contains the NLC has been improved, and the closed radiation zone has been enlarged. This will allow the full use of the NLC which should improve the background levels in Belle II without affecting the ring impedance.

The linac improvements include a new RF gun with a lower field at the cathode plug eliminating the arcing seen in the current gun. Unfortunately, the new gun experienced a brazing misstep resulting in a frequency change of 10 MHz. Tuning may fix this issue, and a backup plan is to use a new IrCe cathode with the present gun which should work through run 2026b.

Other improvements include installing an Energy Compression System (ECS) in the electron transfer line, strengthening the support structure for the Q-magnet in the HER where high beam currents cause magnet displacement, installing higher precision displacement sensors for the vertical collimators, replacing the poles of 11 bend magnets in the positron beam transfer line, and checking and realigning the LER injection point beam pipes.

The background group continues to monitor and improve background simulations. The group found an unexpected increase in background rates for the 2024c run compared to the 2024ab runs and simulation. An increase in beam-gas events could be resulting from the fact that several chambers in the LER wiggler section had been flipped to see if dust coming from the clearing electrodes was a source of SBL events. Flipping the chambers means that the SR fans in this region have a new beam pipe surface and this initiated excessive outgassing until the surface had been scrubbed by the SR. This had not been properly modelled in the simulation. Several other differences between the actual machine and the modelled machine (including time-dependent changes, i.e. tune changes along the bunch train as a function of beam current) are not included in the simulation and are being worked on.

2.2 Concerns

- The effort to inspect and clean as many flange connections as possible is a very good step toward reducing the SBL events. However, venting this much of the storage ring will require some initial scrubbing time as was mentioned in the presentations.

- Damage to collimator jaws continues to be an issue and testing new jaw materials for robustness is important.
- The extra radiation shielding for the NLC should help in gaining the full use of this collimator but as the beam currents increase it may be necessary to further increase the shielding.
- The improvements mentioned for the linac should improve the injection efficiency for both beams. The injection and linac teams will need some dedicated time to fully exploit these changes and also to look ahead to further improve the linac complex.

2.3 Recommendations

- Continue to monitor SBL events especially as beam currents increase. New sources of SBL events may arise.
- Continue attempts to find early detection signals for SBL events and to abort the beams as soon as possible. Faster improvements in beam aborts should reduce the risk of damage to the PXD and the QCS.
- The committee supports strongly the effort by the background team tackling more difficult issues of simulation.
- Be sure to give the linac and beam transport teams time to fully deploy the improvements implemented during this downtime.
- Monitor the radiation backgrounds in the Oho Hall as beam currents increase and with the full utilization of the NLC.

3 Plan for the coming data taking

3.1 Status

Achieving an integrated luminosity exceeding 425fb^{-1} is the top priority for the run period from November 5th, 2025, to June 1st, 2026 (2025c-2026b runs). Excluding 13 days of vacuum scrubbing and 15 days of winter shutdown, approximately 180 days of collision runs are expected. In general, 4 days of machine-related activities are allocated every 3 weeks, resulting in 150 days dedicated to physics runs.

In the baseline plan, the target integrated luminosity can be achieved by reaching a peak luminosity of 1.0×10^{35} cm^{-2}s^{-1}. This requires $\beta_y^* < 0.9$ mm and 1.8 A in HER and 2.6 A in LER. Additionally, 2-bunch injection is necessary in both HER and LER to maintain such substantial beam currents.

In the alternative plan, the target integrated luminosity can be achieved by reaching a peak luminosity of 0.6×10^{35} cm^{-2}s^{-1} while ensuring efficient and stable accelerator operations. Notably, 2-bunch injection is not necessary for HER in this scenario.

Essentially, these plans present trade-offs between a higher luminosity and more stable operations.

3.2 Concerns

- While setting a short-term goal for the integrated luminosity is a good practice, it is unclear how this goal aligns with the long-term strategy to reach the final integrated luminosity.
- It is unclear whether the allocated time for machine-related activities is sufficient.

3.3 Recommendations

- A careful monitor of real-time performance should be initiated by a close collaboration of both the accelerator and Belle II teams, and adjust the run plan as more about peak and integrated luminosities, as well as beam backgrounds become known.
- A stepwise approach, with agreed-upon intermediate targets for peak and integrated luminosities, could be considered.
- The committee suggests allocating as much time as possible to machine-related activities once the target integrated luminosity for this run period is reached.

4 Belle II detector

4.1 Vertex detector (VXD)

4.1.1 Pixel detector (PXD)

4.1.1.1 Status

To avoid further damage to the detector due to the sudden beam loss events, the PXD was turned off on May 7, 2024 for the remainder of run 2024ab and remained off for the entire 2024c run that started on October 9, 2024 and ended on December 27, 2024. The PXD system has been kept cold during run 2024c.

One of the current focuses of the team is on the development of a fast shutdown of the detector. Currently, a regular shutdown takes about 600 μs. A fast shutdown board that rapidly discharges capacitors on power lines through resistors by switching MOSFETs is being developed that could turn the PXD detector off within 30 μs. The signals from the clearing electrodes, that are indicative of an SBL, are the input to trigger a shutdown of the PXD. Board production is planned for this summer, while the optimization of the resistor value, including simulation studies, will continue to determine the optimal settings and avoid dangerous, fast voltage crossings that might damage the switcher chips. A shutdown of the PXD within 10 μs can cause the order of several critical voltages to be flipped resulting in gate damage. Irradiation tests with the new discharge board are tentatively scheduled for October. The preliminary schedule calls for installation of the new boards by the summer of 2026.

The Data Handling Hybrid (DHH) system, composed mainly of an FPGA and a high-speed optical link, is located at the top of the Belle detector and its upgrade has been successfully completed. This improvement doubles the data bandwidth and is expected to reduce exposure to radiation as well as minimize SEUs. Test are currently ongoing and the results are promising.

The team expressed a concern over the departure of the DHH expert from the PXD team. A person to take over the DHH system hasn't been identified yet and a smooth knowledge transfer is essential.

The PXD team is actively preparing for the next run. However, the conditions for when the PXD can be turned on again are difficult to quantify but are actively being discussed within the collaboration. This is a multi-dimensional problem driven by the performance of the accelerator and will require a risk versus benefit analysis regarding physics impact. A PXD expert will be resident at KEK when the collaboration has decided to turn the PXD on again.

At the last review it was recommended to work with the accelerator group to mitigate the IR losses during SBLs. The team has proactively engaged in this exercise and has identified that a rearrangement of the collimators – relocate the LER collimator at D06V2 to D03V4 – could prevent large beam losses more effectively. The team is also installing additional CLAWS beam abort sensors at collimators D03 in the LER and D12 in the HER regions to achieve a reduction in signal propagation delay by about 10 μs. We hope this effort to identify mitigation strategies will be maintained.

4.1.1.2 Concerns

- As noted in the previous report, although there is good progress on the investigation of the sudden beam loss events, their origin is at the moment not confirmed. There is a non-negligible probability that the SBLs may never be completely eliminated. Turning the PXD on thus carries an inherent risk that further damage to the detector will be incurred.

- Although the team has strengthened the management of the PXD project, the personpower to operate, maintain and analyse the data remains marginal. The ability of the collaboration to ensure adequate strength of the PXD team remains a concern.

4.1.1.3 Recommendations

- Continue the development of the mitigation strategy to quickly ramp down the HV to put the detector in a safe state with a target date of implementation before the next run.

- Continue to develop a strategy to turn the PXD on again through a risk-benefit analysis based on physics impact.

4.1.2 Silicon strip vertex detector (SVD)

4.1.2.1 Status

The SVD ran very stably during the 2024c run, also in high background conditions. No occupancy issues were observed and the hit efficiency was excellent with all layers having a hit efficiency exceeding 98.5%. Mainly due to SEU events, the SVD contributed 7.2% to the total number of accidental stops in the data acquisition. The SEU events are due to the worse injection background conditions compared to previous runs, most likely related to the 2-bunch injection and the relaxed diamond threshold during the injection. A firmware-based automatic SEU recovery without stopping the run within 20 ms of an SEU has been implemented and tested. It will be implemented in the next run, run 2025c.

During the 2024c run a picoscope was procured to check real-time injection background on both the aborting and monitoring diamonds. It is being explored if this tool can provide direct feedback to the machine for real-time injection monitoring. The team is preparing an epics Input Output Controller and monitor panel that can be shared with other groups and the accelerator. This should provide useful information for real-time monitoring of the background.

4.1.2.2 Concern

- No particular concern is raised by the performance of the SVD.

4.1.2.3 Recommendation

- The SVD has significant strengths and potential in monitoring the accelerator performance. Continuing the interaction with the accelerator group is recommended.

4.2 Central drift chamber (CDC)

4.2.1 Status

Several studies to better understand the current CDC performance and potential mitigation strategies to improve chamber longevity are under evaluation in the CDC Task Force. Among them are optimization of the amount of oxygen and water content during collision runs, operation at lower gas gain due to the decrease of the high voltage in the innermost layers, increase of the fresh gas flow rate, introduction of additional additives, e.g. alcohol, modification of the gas mixture composition after verification with a test chamber, and ageing studies with the CDC test chamber. The BPAC committee strongly supports these R&D efforts and encourages to continue these studies in a systematic fashion.

During the 2024c run, the gas gain dependence from the chamber current and gas conditions (e.g. oxygen, water and hydrogen content) was observed. Run-gain analysis using collision data was developed to disentangle potential contributions to the gas gain behaviour from different components. To improve precision of separation between several sources, the cosmic run campaign with different H_2O/O_2 content is currently being held during the shutdown. Additional cosmic data runs will further help to constrain different parameters and can facilitate evaluation of already existing ageing effects (gain drop) in the CDC chamber.

A possibility to operate CDC with a decreased high voltage to mitigate the risk of Malter-effect appearance and to minimize the total chamber accumulated charge is also under study. The operation at a lower gas gain will also allow to extend the CDC chamber lifetime. Special run with a decreased high voltage—by 100 V leading to 50 % of gain reduction in the innermost 14 layers—was taken during run 2024c. The prompt online checks have shown that 2D track trigger efficiency remained unchanged (efficiency drop was compensated by decrease of electronics threshold), while 3D efficiency dropped by 2%. CDC performance studies with lower operation voltages—in the innermost 14 layers (same as for collision run in run2024c) and for all CDC layers—will be carried out with cosmic data during the ongoing shutdown. These tests will be complemented by studies with collision data in run2025c before the final decision will be made about the CDC operating voltage.

Modification of the gas recirculation system is foreseen in order to add a controlled amount of oxygen to the gas, whenever necessary. Discussion is currently ongoing, based on a drawing provided by external company with KEK mechanical experts involved. The goal is to implement necessary modifications to the gas system in the fall of 2025, before the start of the new physics run.

An increase of the flow rate of fresh gas is an effective measure to mitigate ageing effects and to remove free radicals produced in the avalanches. The standard CDC flow rate corresponds to one volume exchange per day using a recirculation system, and to the replacement of one CDC volume with a fresh gas on a monthly basis. This is much lower compared to other similar-scale experiments. There is an ongoing effort to increase the fresh gas flow rate by factor of 10 for the run2025c. An additional fresh gas flow supply line has been already installed and needs to be tested. An increase of the gas cost and corresponding funding allocation, procurement of gas supply and contract procedures are currently being investigated.

Due to the ongoing repairs of the X-ray source at KEK, several facilities outside of KEK are being investigated for the CDC ageing studies with a test chamber. The first benchmark will be an integrated charge of 1 C/cm per wire, accumulated in a timely manner. An ultimate goal is to achieve the dose of 5–10 C/cm per wire. A collaboration with Tohoku University, IHEP Beijing and a possibility to advance interactions with DRD1 through internal discussions and informal meetings would be further exploited.

4.2.2 Concerns

- There might be already some permanent gain loss in several CDC layers, in particular the innermost layers, considering the total accumulated charge of up to 150 mC/cm per wire.

- The CDC operating current was very close to the maximal acceptable limit during the run 2024c. Therefore, the operation of CDC at high beam currents in run2025c at peak luminosities above 5×10^{34} cm^{-2}s^{-1} remains a major threat if there will be no significant reduction in background conditions. Potential onset of Malter-like currents represents a major risk for the CDC future.

- For the alternative gas studies with CF_4 admixture, the potential impact of the HF presence on all materials in the CDC and the gas system itself needs to be understood.

4.2.3 Recommendations

- There is significant progress in quantitative separation of potential factors contributing to the gas gain variation during collision data. The ongoing cosmic run campaign could further facilitate separation between different sources and help to evaluate already existing ageing effects in the CDC. The continuation of the advanced run-gain analysis is very much encouraged.

- All possible efforts should be made to optimise CDC operating conditions, such as the oxygen and water content, operation at lower gas gain, increase of the fresh gas flow rate, introduction of additional additives, modification of the gas mixture composition, etc., in order to increase the chamber lifetime.

- Ageing studies with a CDC test chamber needs to be done in a timely manner, at an irradiation facility outside KEK. The optimal set of operating conditions during the test chamber irradiation should be discussed and finalised. Even if such a study could not fully resemble CDC operating conditions, on a longer term, it will facilitate development and test of different remediation strategies for gain recovery, if anode/cathode ageing effects are observed.

- The study with alternative gas mixtures needs to be extended from simulation activities to the measurements, e.g. HV plateau and discharge behaviour, wire hit efficiency, position and dE/dx resolution, using a test-chamber that resembles the CDC cell structure;

4.3 Particle identification system

4.3.1 Barrel system (TOP)

4.3.1.1 Status

During 2024, the TOP group successfully operated the detector for data taking even as backgrounds increased substantially. Frequent *b2link* errors occurred in the 2024c run, especially under bad injection background conditions, which are not yet well understood. They are not caused by SEUs as they are not seen in null runs during the same accelerator conditions. Studies simulating the underlying mechanisms continue. If the issues remain unresolved before run 2025c begins, the TOP internal injection veto will be used to suppress errors.

Relocation of feature extraction from the frontend processing system to the PCle40 computers in order to avoid SEU problems induced by the high radiation environment at the frontend boards is ongoing. The new system should be ready before the fall data taking period. The TOP group continues to investigate aspects of performance in an attempt to improve the agreement between data and the run-dependent MC. Channel-by-channel photon-hit efficiency calibration was implemented, which has reduced the difference, but the PID performance seen in the data continues to lag behind the MC predictions. Upcoming studies will address possible issues with track extrapolation and background hits. In particular, it appears that the MC may underestimate the delta ray contribution by around 30%—an interesting observation which needs further study.

The PMT quantum efficiency (QE) measurement programme at Nagoya University has continued to use conventional PMTs removed from the detector. Jigs are now being built to support tests of the combined effects of the B-field and high temperatures on QE.

For possible QE ageing mitigation in the longer term, further testing is underway to measure the lifetime with a different HV ratio for the divider chain, to reduce outgassing and ions from the first MCP layer that can strike the photocathode. If changing the relative gain is found to be beneficial for the MCP-PMT lifetime, mass production of new dividers could be planned to begin around the end of 2026JFY so they could be retrofitted during an upcoming shutdown.

4.3.1.2 Concern

- The ageing of PMTs is still not well understood.

4.3.1.3 Recommendations

- Efforts to better understand the loss of PMT QE should continue.
- Contributions from the TOP group to characterise and model the performance at the detector level are very useful and should continue as far as possible, including the effect of delta rays.

4.3.2 Endcap system (ARICH)

4.3.2.1 Status

ARICH is a stable detector that has generally run well during all data runs until 2024c when a problem arose with the cooling water as described at the March 2025 meeting. The group has thoroughly investigated this issue by opening the Belle II end-yoke on June 16, so that the pipes could be accessed directly. It was quickly observed that couplers between the supply pipes were clogged, and that all 6 sectors have no flow problems nor blockage inside, as seen through a borescope. The couplers were blocked by swollen O-rings, not by extraneous material. It is not yet clear why the couplers' O-rings blocked flow, but tests will continue. Meanwhile, since the couplers are unnecessary, they will be replaced with direct connections and valves during the period before the end-yoke is closed at the end of September. Additionally, filters and additional water monitors will be added and the chiller will be replaced so that all piping material is compatible for future operations.

4.3.2.2 Concerns

- Neutrons are a future concern as backgrounds increase, as they can lead to HAPD deterioration and increased SEUs.
- Appropriate replacement photon detectors may not be available if needed in the future.

4.3.2.3 Recommendations

- Continue to study alternative photon detectors to the HAPDs to be ready for a future upgrade, in case that becomes necessary.

4.4 Electromagnetic calorimeter (ECL)

4.4.1 Status

The ECL operation has been stable and continuous efforts have been made to enlarge the group. New firmware for improved waveform fitting has been developed to cope with the injection background. The group also plans an upgrade of the front-end electronics, new ShaperDSPs, and to produce its prototype in early 2026. Impacts of the beam background have been studied with the Performance Group. Using the 2024c data set, the group confirmed π^0 mass width is still acceptable, and good agreement with simulation.

4.4.2 Concern

- The background increase at higher luminosities leads to a degradation of the photon reconstruction, as shown by the degradation of π^0 yields at low momentum (below 1 GeV) in the Run II data.

4.4.3 Recommendations

- Performance studies in a higher background condition should continue.
- Performance evaluation with the new firmware should be done. This may provide useful inputs for the new ShaperDSP development.

4.5 K-long Muon Detector (KLM)

4.5.1 Status

The HV Sudden Turn-Off (STO) events without any error messages has been persisting since run2024ab with new-type CAEN SY4527 crates. Investigation of this issue with CAEN engineers is ongoing. Pinpointing the nature of the problem is not straightforward because its frequency varies for unknown reasons and it occasionally re-appears after long time intervals. Recently, it took nearly 80 days. Testing with a new CPU board firmware that records more detailed information at the sudden shutdown is ongoing. Several issues with the interlock system setup have been identified, which might cause instability in the interlock behaviour and require further investigation. CAEN engineers have suggested to use the global interlock input on the crate, instead of the individual interlock inputs on the HV modules. This approach eliminates the need for multiple cable connections and interlock fan-out modules.

No sudden HV shutdown have been observed in run2024c with an old-type crate, CAEN SY1527. However, the old-type crate is no longer produced and no spares are available on the market. Only one spare exists in the TOP group.

4.5.2 Concern

- The lack of spare old-type CAEN crates is a threat for the stable KLM operation and will remain until solution to the problem with a new-type CAEN SY4527 crate is found.

4.5.3 Recommendations

- Ongoing investigation, including diagnosis and prevention of new-type CAEN SY4527 crate module failures, should remain the highest KLM priority.
- The RPC gas system dates from the Belle era, and the ageing of gas system components is a risk. Therefore, regular inspections and continuous monitoring will have to be strengthened during the upcoming run 2025c.
- The ongoing R&D studies aiming to switch Belle-type RPC operation from streamer to proportional mode needs to be advanced. Solid justification that such a change will not lead to KLM performance degradation has to be presented (e.g. using RPC test demonstrators) by the time of the upgrade TDR submission in 2027. Back-up options for the RPC chambers upgrade should be considered and developed.

4.6 Trigger and data acquisition system (DAQ)

4.6.1 Status

The central data acquisition and trigger work well. The reported issues with data acquisition with high backgrounds are all related to limitations or weaknesses in the front-end design and can only be mitigated by (automatic) recovery procedures at the central DAQ level. The central DAQ team has implemented a number of measures to increase overall efficiency by implementing automatic restarts ("SALS") and making the restarts faster.

The HLT throughput has been studied and after increasing from 12 to 14 units there appears to be sufficient headroom. A HLT pre-filter could increase the throughput further by rejecting time-consuming events in the active veto window. In addition to the HLT, there is further possibility to reduce the trigger rate, for example by enabling the new 3D neural network algorithm, among others.

4.6.2 Concern

No specific concern have been identified at this time.

4.6.3 Recommendations

- The collection of failures on restarts, broken down by subsystem, including the reason, should be continued, as it can provide a basis for further optimisation strategies.
- The studies on a GPU trigger should be continued, as this is a good investment into future upgrades of the HLT.
- The DAQ group is performing well. An important factor to this is the support of three FTEs by the collaboration common fund. This support should be maintained.

4.7 Data processing

4.7.1 Status

The committee received an update on the status of calibration and data processing. The *b2luca* tool, developed with support from the *b2luigi* team, is progressing rapidly and is on track to be ready for the next data-taking period. It now covers nearly all steps of the calibration workflow and has successfully passed a dry-run test without any show-stopper.

Full-scale calibration tests on an entire bucket, for both Run1 and Run2, were underway using the *rel-09* software. This version is intended to be used for the 2025c run, therefore testing with the *rel-10* software is less urgent. A full dress rehearsal for training and interaction of sub-detectors' experts is planned for Autumn. In summary, all calibration steps are now fully automated, with the exception of producing the dedicated MCrd required by two specific calibrations. Options to automate this final step are already under consideration.

Reprocessing of the full Run 2 dataset (~150/fb) with a patched release to fix a tracking issue has been completed in about two months. This reprocessing includes Data, MCrd production, and the associated skims. MC signal production was progressing well, with nearly all requested signals already produced; those produced locally were in the process of being registered on the grid. MC liaisons are now in charge of producing signal samples locally at KEKCC using the tools provided by the Data Production team. This not only helps lessen the workload of the Data Production team but also acts as preliminary training for prospective future team members.

The skim production campaign for the entire Run 1, using flags to group skims with large overlaps and thereby dramatically reduce the number of jobs, had been fully completed. The current plan for 2025c is to produce the same set of skims as in 2024 for consistency in the analysis. Additional skims can be accepted if defined before the start of data taking. The possibility of adding new or updated skims afterward remains open, but they will be treated as off-campaign.

In summary, the reprocessing of Data and production of MCrd with *rel-08*, including skims for Run 1 and Run 2, is complete and available for physics analysis. The committee congratulates the production team for this achievement.

4.7.1.1 Concern

- No particular concern.

4.7.1.2 Recommendation

- The committee encourages the team to finalize the improvements and missing steps in *b2luca*, and to automate the production of the specific MCrd required for some calibrations, in order to achieve complete automation of the calibration workflow.

4.8 Performance

4.8.1 Status

The Belle II Performance Group provides corrections for the physics analyses, fast feedback on data quality, develops and releases new tools, and publishes the performance, to serve as a reference. The coordinators recently changed. *Release-8* data and simulation samples are ready, and highlights were shown of the improved Particle ID. The CDC calibration has been improved with corrections on time-since-last-injection (TSI), and significant improvement made for TOP reconstruction. Run 2 data shows better performance than Run 1, mostly due to the CDC improvements. For tagging via Full Event Interpretation (FEI), a significant (30%–50%) improvement is seen for the overall tagging efficiency. It is much better than in *Release-6*, while there is a slight degradation of performance for Run 2.

The Performance Group aim for their recommendations to follow the following guidelines: Flexibility (allow for multiple working points to enable optimization), Orthogonality (they should not influence each other), Optimality (should not produce large spread of weights e.g. for PID fakes which increases MC statistical uncertainty), Coverage (should cover all objects used by the analyses and all aspects of uncertainties), Universality (should be similar for various analyses), Simplicity (should be possible for non-expert to apply), and Traceability (should allow to trace correlations even when analyses decide on different working points). With further development the Systematics Framework (ntuples and tools for calibration modes) can address many of these points. Examples were shown of improvements in progress:

- **Photon/π^0 efficiency correction:** the current prescription involves separate correction for photons and π^0; this will be unified, as the π^0 efficiency is a product of photon efficiencies. The photon efficiency is determined using radiative dimuons, and extended to low momentum using asymmetric π^0 decays. It can be validated using channels previously used for the π^0 efficiency calculation.

- **(Anti)neutron calibration:** (anti)neutrons are reconstructed as an inclusive recoil to a charged system with an (anti)proton and up to seven pions or five pions + two kaons. Clear signal separation is seen for the recoil mass, and the recoil momentum covers the relevant range. The method may be extended to K_L.

- **Improving FEI recommendations:** standard FEI correction factors are provided for "good" (i.e. correctly reconstructed) tagged B mesons, using e.g. $B \to D^{(*)}\pi$ decays for the probe side. There is a significant fraction of "bad" tags, increasing towards low signal probability. The fraction may not be perfectly modelled by the simulation; it is also more likely to be probe-side dependent: dedicated corrections will be provided.

- **Preparing for higher background:** increase of background at higher luminosity leads to degradation of the CDC tracking efficiency and photon reconstruction in ECL. This is studied using data from Run 2. Losses are seen of K_S in the CDC acceptance and π^0 at low momentum. To mitigate those losses on a short term, a local track finding algorithm is considered (with additional filters against extra fake tracks and checking extra CPU). On the longer term, graph neural network finding has been shown to perform better for displaced tracks. The efficiency loss from background is increased by CDC electronic cross-talk which may be helped by the new ASIC.

Corrections for physics analyses using *Release-8* are being prepared, with several improvements for the corrections underway. The group is preparing for the next data taking, identifying the main weak points. Performance analyses are prepared by the group for publication: a Particle ID paper from *Release-6* is available on arXiv.

4.8.2 Concern

- No particular concern, the group appears to be well motivated and focused on improving the performance for physics analyses.

4.8.3 Recommendation

- Concerning the neutron calibration, different Geant4 physics lists might be investigated to better control the calibration; it was pointed out that the data/Monte Carlo discrepancy may also be affected by the generators, and further study is required.

5 Physics

5.1 Status

Since the last remote BPAC meeting in September, the Belle II collaboration has made significant progress in generating compelling results. Eleven new analyses have been completed. Several results were presented at various conferences, including five at FPCP. An additional thirty analyses are in the final stages of review to be presented at the EPS High Energy Physics Conference, the Lepton-Photon symposium, and the 13th International Workshop on the CKM Unitarity Triangle (CKM 2025).

The new results include non-leptonic B-decays into baryons, studies of V_{cb} from untagged semileptonic decays $B \to D\ell\bar{\nu}$, CP-violation studies in D-decays into pion and kaon states, as well as direct bottomonia production in e^+e^- collisions. Physics highlights include the new inclusive studies of V_{ub}, updates of the inclusive tag analysis of $B \to K\nu\bar{\nu}$, searches for the flavor-violating radiative tau decays and searches for dark photons in the $A' \to \mu^+\mu^-$ decays. The collaboration has also anticipated the release in the summer of results on a number of electroweak penguin decay modes of the type $B \to X\nu\bar{\nu}$ (with $X = K$ or the inclusive state), which would be very useful in constraining extensions of the Standard Model.

5.2 Concern

- No specific concerns have been identified at this time.

5.3 Recommendation

- The committee is pleased with the progress on the core flavour physics program and the dark sector analyses. As the dataset continues to grow, the committee encourages maintaining steady progress on these analyses and also paying attention to other high-stakes analyses, such as measurements of $e^+e^- \to \pi^+\pi^-$ spectrum and precision studies with τ decays.

6 Long Shutdown 2 (LS2) upgrade

6.1 Status

To reach a total integrated luminosity of 50 ab^{-1} of data by early 2043 remains the end goal of the SuperKEKB physics program. This is scientifically well motivated, given that flavour physics remains one of the least understood areas of the Standard Model. To achieve that ambitious goal both the accelerator and detector need to be upgraded. A strategy was presented towards a technical design report for an upgrade of the Belle-II detector during LS2. Eight steps with a tentative timeline were identified for this process:

- 2025: Definition of a working plan with milestones
- 2026: Definition of the baseline detector configuration
- 2027 Spring: Realization of technology demonstrators
- 2027 Autumn: Submission of the technical design report
- 2028 Spring: Submission of the engineering design report
- 2028 Autumn: MEXT review under consideration
- 2029: Start of Production
- 2032: Start of integration of new detector components

The committee is pleased to see progress in defining the LS2 upgrade plan, but finds the plan severely lacking in detail. Given the amount of work required, the effort by the Belle II and KEK managements would be needed to attract further interest of institutes within and outside of the collaboration for the upgrade activities, both in the detector and the accelerator, by presenting SuperKEKB as a pathway towards the next future e^+e^- collider.

6.2 Concerns

- The performance degradation of the current detector until LS2 is not fully understood nor simulated. The background model used for the simulations is from 2021 and is outdated. The performance of the detector beyond LS2, for those detector elements that need to be retained, has not been evaluated. The scope of the detector upgrade is not yet well-defined at this point.
- Although no details were presented, it was great to hear that progress has been made on establishing a baseline scenario for the new final focusing magnet (QCS) and that a new envelope has been defined that provides the boundary conditions for the size of the new vertex detector. Nevertheless, there is a clear need for synchronisation between detector and accelerator upgrade planning given the strong interplay between the two.
- Achieving the peak and integrated luminosity goals will require a comprehensive performance review of the entire accelerator complex, including the injector system. Necessary upgrades should be identified and its effect on the background and impact on the design of the upgraded detector evaluated through simulations.

The committee felt that the detector upgrade plan was presented almost independent from the accelerator upgrade, which is of great concern. Both upgrades are intimately connected and should be developed in parallel, with one informing the other. A joint accelerator–detector upgrade plan will be essential moving forward.

- The proposed timeline with a complete Technical Design Report by autumn 2027, construction starting in 2029, and installation beginning in 2032, is highly ambitious. While some benchmark studies have been completed and presented to the committee before, a substantial amount of work remains to be done to reach the level of detail required for a full TDR. A dedicated and sufficiently staffed team will be essential to this process.

- The committee appreciates that the team is concentrating the upgrade efforts on the vertex and tracking detectors, while assessing the continued viability of other subsystems for post-LS2 operation and considers this a very sensible approach. However, a concern is that the scale of the upgrade, even if it is "just" limited to replacing the VXD and CDC, together with the associated background instrumentation, readout, data acquisition system and trigger, seems to be underestimated. A project of this scale would benefit greatly from a project management office. A resource loaded schedule for example, even if very rudimentary, was not presented, nor was a comprehensive overall project management plan. Such a plan would be very beneficial to guide the overall approach to defining and implementing the upgrade.

6.3 Recommendations

- Working with the accelerator team, define as soon as is reasonably possible a baseline scenario for the new final focusing magnet (QCS), including agreement on its position and dimensions, which will define the available space for the new inner detector.

- Evaluate the evolution of the detector performance until LS2 and post-LS2 (using the new QCS design) with simulations to define the scope of the upgrade. Specifically, define the radiation tolerance in terms of occupancies, radiation dose and neutron fluence for key detector components, in addition to the minimum requirements for physics performance.

- Continue simulations of benchmark physics processes, with representative backgrounds, to start defining the detector design.

- Establish a joint accelerator-detector upgrade team to develop a resource loaded schedule for the full upgrade.

- Establish a project management office for the detector upgrade as soon as possible to move the upgrade forward expeditiously.

- Achieving stable running conditions with high luminosity during run2025c is central to defining the scope of the LS2 CDC upgrade. Best practice developed for the construction and operation of radiation-hard gaseous detectors at CERN needs to be studied and followed for the design and construction of the new drift chamber. Collaboration with the DRD1 community is encouraged.

- While detailed comments on the DAQ and trigger cannot be provided at this stage, a hardware trigger along the lines of the present Belle II experiment followed by a flat single stage data acquisition and event-building followed by a software trigger is recommended as a general strategy. In the design, the number of types for custom components (custom-made electronics boards) should be kept to a strict minimum. The design read-out rate after the trigger should be high enough such that unexpected fluctuations in the trigger rate or background can be absorbed in the DAQ and dealt with in software. The software-trigger-on-GPU project presented should be continued as it promises the best cost-performance ratio, at the time, and is widely used in other HEP experiments.