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

3-5 March 2025, Tsukuba, Japan

A. Andreazza* (Milano), P. Collins* (CERN), G. Corti (CERN),
M. Demarteau (ORNL), R. Forty (CERN), S. Gori& (UCSC),
W. Hulsbergen*& (NIKHEF), G. Isidori (Zurich), P. Mato* (CERN),
F. Meijers*& (CERN), N. Neufeld (CERN), R. Patterson (Cornell)
A. Petrov* (Univ. South Carolina), B. Ratcliff* (SLAC), M. Sullivan+ (SLAC),
H. Tajima (Nagoya), O. Tajima (Kyoto), M. Titov* (Saclay)
and chaired by T. Nakada (EPFL)
& Remote participation
+ Partly absent
* Expert member.

25 May 2025

1 Short summary

The annual review meeting of the B-factory Programme Advisory Committee (BPAC) in 2025 took place at KEK from 3rd to 5th of March, with the presentations from the accelerator and Belle II groups on the status of the project as well as the progress in the upgrade plan. The committee is pleased to observe the high level of dedication demonstrated by the machine people and the Belle II collaboration. In this section, the committee gives feedback on the five questions asked by the management of the Institute of Particle and Nuclear Studies. Detailed findings of the committee are given in the following sections.

1. **Are physics analysis plans in this year and toward a few 1/ab convincing?**
   The Belle II collaboration has already been superseding various Belle results with an integrated luminosity half of that of Belle, by exploiting its better detector performance and more advanced analysis techniques. The committee is looking forward to seeing further progress with the substantially increased statistics from the coming data taking. While the recent Belle II results in the dark sector physics are unique and world-leading, the committee considers that *emphasising more on a generic and wider aspect in the analysis could bring further benefit.*

2. **Are all the efforts (computing, data production, performance, software) to publish physics results in a timely manner well organized and sustainable?**
   Belle II data processing has been stable and its functionality has been continuously improved. The tracking software should be *revisited to take the actual performance of the Central Drift Chamber (CDC) with some inefficiencies into account.* If the current level of effort is sustained, the committee thinks that the Belle II computing will be able to cope with the future demands of the timely analysis of the data.

3. **Is the strategy to address the issues of the machine performance, beam background and sudden beam loss in collaboration with SuperKEKB clear?**
   The committee acknowledges the continuous effort to improve the SuperKEKB performance and ongoing hardware improvement activities during the current shutdown period. It is understood that various machine groups as well as the Belle II members are involved in the activities. Since the machine has to achieve high luminosities while keeping the background level safe for the detector, *participation of the Belle II group is crucial.* Given the limited resources available for the machine groups, *involvement of the Belle II group beyond the machine detector interface issues is encouraged.* An effort to *seek support from foreign accelerator laboratories with relevant expertise* is also welcome. In order to profit from the effort coming from those different sources, *introduction of a centralised project structure with a strong leadership to coordinate and prioritise the work is highly recommended.*

   The sudden beam loss (SBL) remains as one of the major obstacles for achieving stable runs. A vacuum sealant, Vacseal, which was heavily used to ensure the vacuum tightness of the beam pipe joints, has recently been identified as a possible main cause for the SBL. While this appears promising, the process how the Vacseal causes the SBL is still unknown. The committee thinks that there are likely to be multiple causes for the SBL. Hence removing the Vacseal could reduce the occurrence of the SBL by an order of magnitude, but other SBL will remain. In order to operate the machine without damaging the detector, *further effort to improve the early detection of the SBL and the fast beam abort is crucial.* This will allow more efficient machine studies to increase the luminosity and Belle II to accumulate the data.

4. **Are all the detector issues found during the 2024 operation understood and are the measures to be taken during this shutdown sufficient?**
   Although the general status of the Belle II operation appears good, there remain several concerns. *PXD2 needs further work to define the beam condition criteria for switching on during the data taking.* Understanding the gain from switching on PXD2 in physics analyses for different machine background conditions could help to weigh the risk against gains. *Long term stability of the CDC operation is a major concern.* The chamber has already accumulated a significant amount of

charge and further increase of the current can hardly be tolerated. Continuous attention to the operation is needed as well as the effort with the test chambers to understand better the characteristics of radiation damage in CDC. Although it is still in the early stage of inquiry, the recently discovered problem with the cooling of the frontend electronics for the ARICH end-cap particle identification system could be serious. The presented plan for the investigation is adequate and the committee is *looking forward to hearing an update during the June meeting.* In the data acquisition, some of the subsystems still require manual intervention for recovery. *This should be all automated.*

5. **Are the procedure and timeline to define the objectives and to make the detailed schedule of SuperKEKB and Belle II upgrades in Long Shutdown 2 clear? Are all critical or limiting points identified?**
   While the new date for the start of Long Shutdown 2 (LS2), around 2032, is more realistic for the upgrade work to be completed than the previous date of 2028, it introduces some concerns whether the subsystems continue to operate till that date with the required performance, in particular for the CDC. Judging from the work plan presented for the new pixel vertex detector (VTX) and final focusing quadrupole magnet (QCS), LS2 could start even later. For this reason, the committee encourages the Belle II collaboration to consider *the decoupling of the CDC upgrade from the VTX-QCS upgrade.* As soon as the layout of the interaction region with the QCS is completed, the design work of a new CDC should start in such a way that the new CDC can be installed in the present Belle II set up as well as with the new QCS-VTX configuration. This will make the upgrade plan more flexible, while maintaining the detector performance until LS2.

NB
The committee took note of the presented computing resource accounting for 2024 and estimates for 2026 to 2029. A dedicated discussion and recommendation will be given by the expert group consisting of G. Carlino, W. Hulsbergen, P. Mcbride and P. Mato, and chaired by the BPAC chair, after receiving the written accounting report and request.

2 Machine related issues

2.1 Operation of SuperKEKB and injection complex

2.1.1 Status

The SuperKEKB accelerator has achieved a peak luminosity of 5.1×10^{34} cm^{-2}s^{-1}. The Belle II detector HV was off when the record was made.

The injection system has been primarily successful in delivering two-bunch injection into the LER. There are higher background levels with the second bunch and there are plans to install fast kickers to correct the orbit of only the second bunch making the two bunches more similar. Two-bunch HER injection studies were curtailed due to an arcing issue with the cathode of the gun. The reason for the arcing has been understood and a new cathode with an improved design will be installed this spring and summer. The positron damping ring presently has a current limit of 35 mA due to radiation concerns and this will also be addressed during the present downtime. Machine Learning has been applied for the Linac operation with the positron beam, and this has improved the total positron charge to 4.2 nC which is higher than the design of 4 nC.

The LER stored beam β function at the injection point will be increased from 100 m to 160 m which will effectively move the injected bunch closer to the centre of the stored beam. This reduces the oscillation amplitude of injected bunches and should improve the injection background. The non-linear collimator has helped to reduce the stored beam backgrounds without increasing the beam impedance and has also helped to lower the injection backgrounds. The full potential of the collimator has not been realized because the radiation levels in the nearby hall have been too high. This issue will be addressed by adding more shielding. There is still a mysterious increase in the HER vertical emittance. This occurred sometime during the initial scrubbing period of the 2024c run. So far, all attempts to understand and reduce the emittance have not been successful. There is a difference in the horizontal tune of the LER bunches along the bunch train. This difference increases as the beam current increases. This is pushing either the first bunch or the last bunch in the train close to a resonance line in the tune plane. Lowering the LER β_x at the IP from 80 mm to 60 mm reduced the LER beam-beam blow-up.

In the previous meeting, it was suspected that dust particles coming from clearing electrodes on the inside top wall of the beam pipe in the LER wiggler sections could be a source of many of the Sudden Beam Loss (SBL) events. Unfortunately, flipping the beam pipe sections with clearing electrodes did not reduce the SBL event rate. On the other hand, a strong case has been made that Vacseal, a sealant for vacuum leaks, is a large contributing factor to the number of the SBL events that have been observed. Pressure bursts have been correlated with some of the SBL events, and a black residue has been found after opening the vacuum system. Removing the residue greatly reduces the rate of subsequent SBLs originated in that region. There is a program now to inspect the regions where Vacseal was used and look for and clean any residue found in the vacuum system. Special attention will be paid to the D04 wiggler region as the SBL events that occur in this area tend to produce high radiation events in the Belle II detector as well as QCS quenches.

2.1.2 Concerns

- Finding that Vacseal residue in the vacuum system is the cause of many of the SBL events could be a big step forward, but there are most likely other sources of SBL events.
- The source of the high vertical emittance in the HER has not yet been found.
- The varying tune shifts along the bunch train will make it difficult to increase the

LER beam current.

2.1.3 Recommendations

- Continue the current efforts to locate the SBL events. New sources may emerge as the beam currents are increased.
- Continue the efforts to shorten the time between detecting an unstable beam and aborting the beam.
- Continue efforts to improve injection efficiency and performance.

2.2 Machine detector interface

2.2.1 Status

The MDI group continues to play a key role in collaboration with the Linac and SuperKEKB teams, as well as the rest of the Belle II detector teams, in understanding and reducing detector backgrounds. They work closely with the SuperKEKB vacuum group to help identify the locations of Sudden Beam Loss (SBL) events. The group is actively working to shorten the time between the detection of an unstable beam and the beam abort, in order to minimize potential hardware damage. Furthermore, the group is investigating the further optimization of the collimator system including possible relocation, to minimize damage from SBLs. In addition, at the request of the MDI team, the VXD diamond system has applied a relaxed beam abort threshold during injection since the 2024c run (details in SVD section), which has reduced unnecessary beam aborts caused by poor injections and contributed to smoother operational startups following maintenance days. Machine-learning-based injection tuning, developed by the MDI team, was extensively tested during the 2024c run and demonstrated the capability to improve injection efficiency without relying on expert manual tuning.

2.2.2 Concerns

- Many efforts are made in parallel and nearly independently within the MDI team, which require regular communication among them.
- This large effort by the MDI team needs support from the laboratory and detector management.

2.2.3 Recommendations

- Further close collaboration of the MDI team with the Linac, SuperKEKB and other detector teams is highly encouraged.
- Their effort on the abort timing reduction, background suppression, SBL finding, etc., are all very important and need to be continued.

3 Belle II detector

3.1 Vertex detector (VXD)

3.1.1 Pixel detector (PXD)

3.1.1.1 Status

The PXD started taking data during run 2024ab with 35 out of the 40 modules operating. Module 2.12.2 was turned off because of excessive noise and both modules in ladders 2.7 and 2.8 were turned off because of the bowing issue. Unfortunately, on April 22 and May 6, there were sudden beam loss (SBL) events that damaged a small fraction of the PXD. After the first beam loss event an increase in the switcher current was observed that resulted in an increase in temperature, especially in the balcony areas, by as much as 5 °C as measured by the FOS sensors. To ensure safe operation, three additional modules were turned off; the second module in ladder 2.12 and both modules in ladder 2.4 were turned off. The ladder glue joints that are temperature dependent and are known to become soft at temperatures above 48 °C, were stable. There is unfortunately little to no margin to provide extra cooling to the switcher chips. Extra runs have been taken to calibrate the FOS sensors and the analysis is ongoing. During the first SBL 84 gates died. The sudden beam loss event on May 6 caused the loss of another 66 gates for a total of 150 dead gates. Overall, approximately 2% of the pixels of the PXD are dead. To avoid further damage the PXD was turned off on May 7 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. During the 2024 run a total luminosity of about 45 fb^{-1} was taken with the PXD and about 146 fb^{-1} without the PXD.

Ladder bending has been studied in some more detail. An analysis of the difference in z position of the measured hit and extrapolated hit, excluding the hit from the track fit, showed displacements up to 60 μm with a strong dependence on beam current that is not understood. The displacement is being corrected for through a time-dependent alignment.

The noisy switcher regions have also been studied in more detail by looking at hits associated with tracks. No significant loss in efficiency is observed and a more performant masking procedure for noisy pixels is being developed. A concerted effort is being undertaken for an early detection of SBL events and it is being investigated if the PXD can be turned off quickly, that is, within one beam cycle of 10 μs. Currently, a regular shutdown takes about 600 μs. A fast shutdown board with active pull-down for all channels is being developed that could turn the PXD detector off within 30 μs. Tests are ongoing to determine the optimal settings.

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 will be upgraded. This is expected to reduce exposure to radiation as well as minimize SEUs. Installation is scheduled to take place during the April and May time frame followed by extensive tests.

The PXD team is preparing for the next run. Software and hardware is being readied and shifters are being trained. An exact copy of the KEK server setup and software has been installed at DESY with three operational modules. This is a great step for which the team is congratulated. This setup will be used to develop and test new software and train experts and shifters for operation. Expert shifters are expected to train regular shifters. 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.

The transition to new management of the PXD has been seamless with two people, assisted by a technical coordinator, now at the helm. The number of available shifters is low, with several leaving the collaboration soon. This is proactively being addressed through the development of a training program with a duplicate data acquisition satellite setup at DESY. A DHH expert has been identified at KIT and is expected to assume that role in May. The overall manpower situation to maintain and operate the PXD remains marginal.

3.1.1.2 Concerns

- Due to the large number of SBLs, PXD could not be switched on during the data taking.
- The origin of the sudden beam loss events is not understood, and could take a while before it is adequately addressed. 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.

3.1.1.3 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.
- Develop a strategy to turn the PXD on again through a risk-benefit analysis based on physics impact.
- Work with the accelerator group to develop mitigation strategies for beam losses targeted towards protecting the PXD.
- Improve the temperature calibration of the FOS sensors to the extent possible.

- Install an additional diamond sensor with a larger dynamic range that does not get saturated in case of a sudden beam loss event.
- Start developing a realistic schedule for the activities to be carried out during LS2 related to the new VXD, including the option for (partial) replacement, in close collaboration with the accelerator group and KEK management (see also comments in the LS2 section).

3.1.2 Silicon strip vertex detector (SVD)

3.1.2.1 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%. Due to SEU events the SVD contributed 7.5% to the total number of occurrences of downtime of the experiment.

The beam abort system based on the diamond detectors has been optimized. New firmware was deployed in the unit that controls the diamond detectors during the 2024c run with two sets of threshold settings for the total integrated dose within a 10 μs window: high beam abort thresholds were used during injection and a set of strict thresholds was used outside the injection windows. A reduction of 57% in beam aborts was observed even under high injection background conditions.

To avoid damage to the diamond preamplifiers, which was experienced during the 2024b run, the high voltage for the sensors monitoring the QCS and beam pipe was lowered from 100V to 5V. 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 background in the SVD has on average been below 1% for all of 2024, though the background during run 2024c was a little higher due to the higher injection backgrounds. The higher injection backgrounds also resulted in a three times higher rate of DAQ issues due to SEUs of the APV25 readout chip. An automatic recovery from SEUs is being implemented for the next run.

The performance of the SVD with integrated radiation dose is as expected. The annealing during LS1 and the lower operating temperature had a positive effect on the overall noise performance at the start of run 2, and a slight increase in noise at the level of 5-10% has been observed so far during run 2.

Maintaining good tracking performance in the presence of higher background is critical. The excellent timing resolution of 3 ns of the SVD can be exploited to reject background hits while maintaining high hit efficiency. A simple hit time selection with $|t| < 50$ ns can achieve a 50% reduction in background while maintaining a hit efficiency of 99%. A more powerful event-by-event time-based cluster classification has been developed. This would allow the SVD to retain excellent tracking performance in an environment with up to 4.7% occupancy in layer 3. These SVD time-based selections are not yet deployed in reconstruction, since background level is still quite low. However, the event-by-event time-based cluster classification is ready to be deployed, even in next run, in case the occupancy in layer 3 gets to 2%–3%. A conservative estimate of the layer 3 occupancy at luminosities of $6 \cdot 10^{35}$ cm^{-2}s^{-1} shows an occupancy of up to 9%. To handle increased occupancies, the use of an average time of clusters-on-track is already implemented and would allow the SVD to retain excellent tracking performance in an environment with up to 6% occupancy in layer 3.

The offline software for the SVD is constantly being improved with an automatic dE/dx calibration being implemented and an automatic production of offline performance plots, called the Validation Interface for Belle II Experiment (VIBE).

The team is to be congratulated on maintaining a very high operational efficiency of the SVD and being proactive in ensuring high performance under future, potentially less favourable background conditions.

3.1.2.2 Concern

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

3.1.2.3 Recommendation

- The SVD has significant strengths and potential in monitoring the accelerator performance. A stronger interface with the accelerator group could be explored.

3.2 Central drift chamber (CDC)

3.2.1 Status

Dark current blow-up, i.e. the Malter-like effect, was observed in the CDC inner layers at the start of the 2024c run, similar to the persistent currents observed in the outer layers in 2018. This was the consequence of the much larger SuperKEKB background, due to injection issues and high LER pressure compared to Run 2024ab, and resulted in the CDC operating current of $\sim$ 300 μA per layer in combination with a relatively low water content of 1200 ppm.

Continuous operation under such conditions carries a high risk that a larger part of the CDC can be affected, possibly causing irreversible damage in the future. Several counter-measures have been immediately implemented to confine the effect. High voltage in the affected CDC layers was decreased by 500 V, water content was increased to 3000 ppm, oxygen removal was stopped, and the limit on the maximum operating current per CDC layer of 250 μA was introduced. All these measures allowed to stabilize CDC operation until the end of the 2024c run, in combination with a better control of machine backgrounds by collimation tuning.

As of today, the total accumulated charge in the CDC innermost layers has already reached 150 mC/cm per wire. This can be compared with 30 mC/cm wire in the BaBar drift chamber during the nine years of operation.

Higher CDC operating currents are expected in the near term at a luminosity of order 10^{35} cm^{-2} s^{-1} with large uncertainties due to the beam-related background conditions.

Therefore, there are serious concerns that Malter-related effects could potentially reappear, as the luminosity increases. The operating current of 300 μA/layer corresponds to a current density of ~20 nA/cm. This means that the requirements for the radiation hardness of the CDC are comparable to those for gas detectors in the LHC experiments.

Several potential measures to mitigate risks are being discussed in the CDC Task Force, such as optimization of the amount of oxygen and water during operation, lowering the gas gain by decreasing the operation voltage in the innermost layers, introducing additional additives, e.g. alcohol, or changing the gas mixture after verification with a test chamber. Meanwhile, ageing studies with the laboratory test chamber did not progress further due to the ongoing repair work of the X-ray facility at KEK.

Several FE-readout boards that led to unstable data links were excluded to avoid DAQ downtime.[1] The net effect of masking the FE-boards was tracking inefficiency in the same phi region of superlayers, since the current tracking algorithm was tuned assuming perfect hardware conditions.

3.2.2 Concerns

- The relation between the decrease of the dE/dx mean, i.e. gain drop, as function of the HER and LER backgrounds, represented by CDC chamber currents, is similar to Run 2024ab. This effect is still not well understood and requires further studies.
- There is a serious concern that there might be already some permanent gain loss in several layers, in particular the innermost CDC layers, taking into account the total accumulated charge of up to 150 mC/cm per wire.
- Appearance of Malter-like currents in several CDC layers during the October physics run represents a major risk for the long-term CDC future.

3.2.3 Recommendations

- Further studies with increased sensitises are needed to pin-down reasons for the gain drop due to the gas composition and the beam backgrounds. High statistics cosmic runs with stable gas conditions are necessary to compare the CDC performance in 2025 with that of earlier data-taking periods.
- As the repair of the X-ray facility at KEK is still ongoing, we encourage the CDC group to construct another chamber and to consider performing ageing test at a different facility. This might facilitate development and test of different remediation strategies for gain recovery, if anode ageing or Malter-type effects are observed.
- In the longer term, all possible efforts should go into the optimization of CDC operating conditions in order to increase the chamber lifetime.

[1] It was understood later that the insufficient voltage of the 1.5 V line was one of the causes for unstable links and 7 out of 16 power supplies were already replaced.

3.3 Particle identification system

3.3.1 Barrel system (TOP)

3.3.1.1 Status

During 2024, the TOP group successfully operated the detector for data taking even as backgrounds increased substantially. The DQM shifter panel had been improved so that TOP data quality issues were detected in the control room display, and bad runs during the high background conditions of run 2024c were easily spotted. Frequent DAQ "b2llost" data link errors occurred which are not yet well understood, but fully vetoing the high occupancy events appeared to work. Studies are underway to understand and mitigate these errors for future operations, including working on injection vetoes. Critical functionality is also being moved from the frontend processing system to the PCle40 computers in order to avoid problems induced by the high radiation environment at the frontend boards. The readout PC software is under construction in Hawaii and will be tested soon, The new firmware will be tested at KEK in March.

The TOP group continues to probe PID performance using the D^* samples from data and comparing with the MC. The timing resolution and the number of photons observed agree well. However, the PID performance seen in the data continues to lag behind the MC predictions. An upcoming study will address possible issues with track extrapolation.

The PMT quantum efficiency (QE) measurement programme at Nagoya University has continued using conventional PMTs removed from the detector. The decrease in quantum efficiency vs output charge is about one order of magnitude steeper in collision data than predicted from the bench tests. Some modest degradation (relative to 25°C) is observed at higher temperatures (50°C) in the bench test but it is several times too small to explain all the difference between the bench and collision rates of QE loss. The group has studied QE degradation in a 1.5T magnetic field with a sizeable effect being measured. However, once again it does not fully explain the much faster QE drop seen in situ, as compared to the bench tests. These studies continue with the next lifetime test addressing the issue of possible enhancement of relative ageing which might occur during low gain operation at high temperature.

For possible mitigation in the longer term, a further test is underway to measure the lifetime with a different HV ratio for the divider chain. This is motivated by the hypothesis that the QE loss is driven by ions primarily released from the first MCP. If changing the relative gain were found to be beneficial for MCP-PMT lifetime, new dividers could be fitted during an upcoming shutdown.

3.3.1.2 Concerns

- Ageing of the PMTs remains to be understood: although no single smoking gun has been found that explains the large differences in the measured PMT ageing rates between collision and bench data, combined effects of high temperatures and the B field might play a significant role, which should be examined.

- Whether the radiation sensitivity of these PMTs is understood sufficiently well to discount it as a possible mechanism for QE loss in situ is not yet certain.

3.3.1.3 Recommendations

- Efforts to better understand the loss of PMT QE should continue.
- Contributions from the TOP group to characterize and model the performance at the detector level are very useful and should continue as much as possible. The committee feels that looking at the performance vs physical and detector variables could help elucidate the overall performance issues.

3.3.2 Endcap system (ARICH)

3.3.2.1 Status

ARICH is a stable detector that has generally run well during all data runs until 2024c (see the discussion of the cooling-water problem below). The ARICH DAQ is stable with the frequency of DAQ errors a few times per week. SEUs in the FEBs are automatically corrected by a scrubber and are not a problem. One concern has been the gradual increase in dead HAPD channels to 6.7% in 2024c. However, this was quite stable during 2024. There were still noisy HAPDs in sector one that sometimes appear and then disappear, without being understood. These have already been reported at the 2024 BPAC. The cause remains unknown.

The ARICH is generally quite insensitive to large beam backgrounds, and no problems have been seen to date with either the DAQ or the PID performance due to present background levels.

The ARICH has developed a significant problem with cooling water flow. Some temperature increase was already seen in run 2024b, and it became substantial during 2024c, with three sectors (2, 3 and 5) having little or no flow by the end of the run. The water flow for each sector is not monitored, so the temperature in each sector is used as a flow proxy. Temperatures greater than 50°C should be avoided to remain within the HAPD operating range. Some sectors were turned off partially or fully as the problem was investigated to keep the temperature within limits. Several results from initial investigations of the problem were presented at the BPAC and further investigations and mitigation measures were discussed. Green murky liquid has been found in the manifolds, which early results from the KEK Environmental Safety office suggest is organic. Further work is ongoing to further isolate and test system components, and develop mitigation strategies.

3.3.2.2 Concerns

- The origin of the noise in sector 1 is undetermined, and though it has a negligible effect on performance at present, it is concerning that it has returned without a clear cause or solution.

- 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.
- The recent cooling problem will lead to significant loss of acceptance due to sectors that are switched off.

3.3.2.3 Recommendation

- Understand the cause of the cooling water problem and develop mitigation strategies. The committee would appreciate a report at the next BPAC.

3.4 Electromagnetic calorimeter (ECL)

3.4.1 Status

The Electromagnetic Calorimeter (ECL) is working stably. There are no dead channels. The organization of data quality monitoring looks to be more robust than at the last BPAC. The online luminosity monitor is in stable operation during 2024c. The online luminosity is consistent with the offline measurement within the statistical uncertainty: online/offline ratio is ~ 0.99. Large pile-up noises from the direction of the LER upstream is observed. ECL can take data under the large pile-up noise, while the energy and time resolution degrade. The group is evaluating the expected performance degradation under future beam background conditions.

During the shutdown in 2025, the group plans to upgrade the FPGA firmware. Implementing a new waveform fitting algorithm eliminates the effects of the injection background. This upgrade is expected to provide better precision of the signal amplitude. The group considers developing a covariance matrix treatment method to save a large amount of memory usage.

Toward LS2, there is no serious concern for radiation damage, i.e., it should be easy to survive until 2032.

3.4.2 Concerns

- The tasks are concentrated in limited institutions.
- There could be an impact of the pile-up noise on physics analysis.
- There is a discrepancy in the time distribution between the data and MC.

3.4.3 Recommendations

- Reinforce the organization; increasing the number of contributing institutions, would be ideal.
- Understand the impact of the pile-up noise on the physics analysis.

- Developing the treatment method to improve the MC/data distribution in time.

3.5 K-long Muon Detector (KLM)

3.5.1 Status

The KLM ran smoothly and with good efficiency in the most recent run, 2024c, thanks to the concerted efforts of the team. However, several subsystems face ongoing challenges. The HV Sudden Turn-Off (STO) events, in which a module trips without triggering alarms, were staved off in run 2024c, but they remain unexplained. Notably, the STO's have been observed with two CAEN SY4527 crates, but not with an older CAEN SY1527 crate. This phenomenon has not been seen by other users or the manufacturer CAEN, and a problematic module sent to CAEN tested OK using their SY4527. While another Belle II group has a spare SY1527 crate, the KLM group has none and they are no longer in production. The KLM therefore must rely on the SY4527 crates, making it essential to diagnose and solve their failures.

The RPC gas system dates from Belle era, and the ageing components are at risk. In 2024ab, a number of plastic tubes were found to have cracks. They were replaced and all other tubes were inspected; a fresh inspection is planned for this shutdown. Mass flow meters are working well, and spares are on hand. A monitoring system is in place and is working well; it flagged the cracked tubes.

In 2022, the efficiency of the BB2 RPCs dropped. Subsequent studies indicated that water contamination was responsible, and that the injection of NH_3 could recover the efficiency; however, by 2024, the efficiency had largely recovered on its own for reasons that are not fully understood, and in Spring 2024, BPAC recommended against the "potentially risky" use of NH_3. The team is currently working to determine the mechanism behind the RPC recovery since lowered efficiency may recur. One possibility is that radiation helped, so they will expose test RPC chambers to radiation and measure the impact using a KEK beamline.

So far backgrounds have not affected KLM performance, and it has an estimated safety margin of 2.2; however, that background prediction is uncertain, and backgrounds could become a problem.

3.5.2 Concerns

- The lack of spare HV crates threatens reliable KLM operation and will continue to be a threat until the problem with SY4527 crates is diagnosed and fixed.
- The gas system is ageing, and regular inspections will be needed going forward.
- Efficiency losses in the RPC chambers could recur, and the mechanism for recovery is not understood.

3.5.3 Recommendations

- Prioritize the diagnosis and prevention of module failures in the SY4527 crates.

- A proposal has been developed to change the operation of RPCs from streamer mode to proportional mode. We consider this a risky solution for the long-term RPC operation and recommend a study of the reliability of this approach in advance of a decision.

3.6 Trigger and data acquisition system (DAQ)

3.6.1 Status

The overall data-taking efficiency has been 85% during Run 2024c. While this falls a bit short of the self-set goal of the 90%, these challenges are explained largely by the difficult background conditions in the accelerator. The losses in the injection veto and due to run-restarts caused by subdetector problems in the front-end electronics, which are in turn caused by increased rates of single event upsets (SEU), outweigh the many measures taken by the DAQ team to improve the efficiency. It must be said however that without these measures the drop in efficiency would certainly have been even larger, and seen from this perspective, these measures deserve to be recognized as a success. Moreover they point the way to future improvements, which will increase the robustness and hopefully recuperate some of the efficiency.

The trigger too was challenged by up to 4 times higher backgrounds, than seen in 2024ab, but it has resisted well and kept the rate comfortably below the DAQ limit with a maximum rate of about 8 kHz.

3.6.2 Concerns

- A drastic additional increase in backgrounds would likely go beyond what the currently envisaged robustness improvements can compensate for.

- The TTD links suffered from a low level of reliability problems since the beginning of Belle-II operation. Some improvement was achieved by properly connecting the ground of the shielded RJ45 copper cables in the connection to the PCIe40 boards Replacement of the copper cable with an optical link was established to reduce the noise problem. Recently, a prototype optical link on the PCIe40 board was tested. Unfortunately, the implementation was not successful due to a mechanical issue. It is understood that Belle II has to live with a low level of failures for the remaining run.

- For more flexible conditions of the HV (high voltage) control, a 'Any Risky Operation (ARO)' flag is being prepared. In this scheme, when SuperKEKB performs a risky operation, each sub-detector can define their own condition independent of the master HV state. The committee does not fully understand how this new scheme would improve the efficiency of physics data collection.

3.6.3 Recommendations

- All proposed and studied improvements, including the more long-term ones, should be brought to production readiness.
- Once the “lower hanging fruits” have been collected, studies should continue to find additional possibilities to increase robustness.
- The committee encourages the continued implementation of the longer-term HLT project to evaluate the use of GPU and/or FPGA acceleration and the integration into the offline software framework.

4 Software, data processing and computing

4.1 Tracking and particle identification software and their performance

The committee received reports on the global performance — with an emphasis on PID — and separately on track reconstruction performance.

4.1.1 Particle Identification and Luminosity

4.1.1.1 Status

The performance group reported on various aspects of Belle II performance, with a focus on PID improvements in *release-08*. Significant gains in lepton identification, in particular for muons, were obtained by exploiting ECL cluster shape information. This improvement is evident in the reprocessed Run 1 data. The performance of the PID detector in the Run 2 data was improved in several ways, among others, due to better control of the CDC gas composition.

Initially developed for K/π separation, the Neural Network-Based PID (PIDNN) has been extended to all six particle hypotheses. It is now available to analysts in data processed with *release-08*, providing improved performance for both hadron and lepton ID.

To harmonise the computation of calibrations and systematic uncertainties among different analyses, the “Systematics Framework” provides a common approach for hadron and lepton PID performance studies. The use of common calibration samples for hadron ID is well-established. A similar implementation for electron and muon ID is progressing. All PID corrections are provided via the Systematic Framework starting from *release-08*. A new common tool (SysVar) has been developed to apply correction tables to analysts’ ntuples, standardising the process and improving error propagation.

A new method has been developed to calibrate the π^0 reconstruction efficiency by exploiting a tag-and-probe method with $\tau \to \pi^0 X$ decays. This has led to a substantial reduction in the systematic uncertainty. Efforts are underway to reduce the energy threshold for photon reconstruction (from 0.2 to 0.1 GeV) and improve K_L PID.

The Run 1 luminosity measurement has reached a total uncertainty of 0.47%, significantly better than previous measurements, considering a combination of ee, $\gamma\gamma$, and

$\mu\mu$ processes. For Run 2 a new Active Injection Veto has been developed. The new veto allows triggering events close to injection, increasing recorded luminosity. However, it results in slightly worse tracking FOMs and PID performance, with a higher track fake rate.

4.1.1.2 Concerns

- The PID performance is highly dependent on stable detector conditions. Variations in detector conditions, such as changes in gas composition for the CDC impact the consistency of PID performance.
- Sizable discrepancies persist in the K/π separation of the TOP counter's data and Monte Carlo simulations.
- The new active injection veto leads to a 10-30% larger rate after HLT1, but only to a 2% improvement in the efficiency of single channels for physics analysis. The higher rate leads to an increase in CPU requirements. Furthermore, the extra data could lead to additional systematics in physics analysis if they are of lesser quality.

4.1.1.3 Recommendations

- Establish a robust monitoring system to detect and correct for any deviations in real-time, ensuring consistent PID performance across different runs.
- Various approaches should be investigated to address the data/MC issue of the TOP counter.
- Assess if the extra efficiency obtained with the new active injection veto is sufficient to motivate the extra expense on data processing, calibrations and performance validation.

4.1.2 Tracking

4.1.2.1 Status

The tracking group reported on the performance of the track reconstruction software with a focus on *release-08* and its impact on Run 1 and Run 2 data processing. With *release-08* significant improvements have been obtained, for instance in the track charge determination and the usage of hit time information in the track finding. This has led to an improvement of about 5% in the D^* reconstruction efficiency. New in *release-08* is also the possibility to access the track time – extracted from SVD hits – for analysis. This helps reduce the contribution from the background tracks by about 50%. It has been shown to be especially effective for analyses involving 3-prong tau decays.

Sudden beam losses in Run 2 led to the decision to turn off the PXD, worsening the single-track impact parameter resolution by about a factor 2. The impact of this on analysis, in particular those most sensitive to the decay time resolution such as TDCPV, has been assessed.

The CDC detector is suffering substantially from failing readout boards. The impact of this on tracking is substantial, as the CDC track finding efficiency is sensitive to missing hits. On the detector side, measures have been taken to reduce the time taken to recover failing boards. At the same time efforts are underway to improve the track reconstruction algorithms to be more robust against missing hits.

In *release-08* an issue with CDC hits being lost in the track fit, in particular for high-momentum tracks in the forward region was discovered with Run 2 data processing. The effect is much less noticeable in Run 1 reprocessing. One cause for this was a selection applied on track quality to reduce the algorithm execution time. The issue was missed in *release-08* validation. Efforts are ongoing both to address the problem and improve the release validation procedure.

4.1.2.2 Concerns

- The Belle-II tracking algorithms appear to struggle with adapting to hardware issues, such as failing readout boards in the CDC. The current lack of robustness could significantly impact data quality and analysis outcomes.
- The sudden beam losses in Run 2, which led to the PXD being turned off, have caused a notable drop in single-track position resolution. As expected this affects Belle II performance in important measurements such as TDCPV. It is unclear to what extent the background conditions will improve. If they do not improve, one may wonder if detector performance would not benefit from removing the PXD to reduce multiple scattering in front of the SVD.
- The example of the issue with the CDC track fit has shown that release validation still needs to be improved.

4.1.2.3 Recommendations

- Continue the development and tuning of more robust tracking algorithms to adapt to the actual performance of the detector.
- Continue to improve the release validation, keeping the physics analysis groups involved at an early stage.
- Develop a strategy for the tracking if the machine background conditions that prevent the usage of the PXD do not improve.

4.2 Data processing and calibration

The committee received an update on the current status of calibration, data processing, and Monte Carlo (MC) production for the 2024 datasets. In addition, a dedicated report was presented outlining the calibration plans for future data taking. The status and plans for future software releases to be used in HLT for the 2025c data taking and in 2026 were also reported.

4.2.1 Data processing

4.2.1.1 Status

Reprocessing with *release-08* of Run 1 data (*proc16*) and Monte Carlo run-dependent productions (*MC16rd*) had been completed at the time of the BPAC. The production of skims for both data and MC was in progress. Procedures are well-established but some refinement is still necessary to shorten the interval between the end of calibration and the start of data processing and MC productions. *MC16rd* samples corresponding to reprocessed `mDSTs` were provided much more quickly compared to previous campaigns. However, while MC samples for the first chunk of data were ready concurrently with the `mDSTs`, the time delay progressively increased for the following chunks. Data collected in 2019 was excluded from reprocessing, due to its small size ($\simeq 2\%$ of the Run 1 data set) and its varied conditions.

Local production of high priority *MC16rd* signal samples had just started at KEKCC at the time of the BPAC meeting. Their fast turnaround (20M events/day) will allow the use of the samples in analysis targeting the upcoming summer conferences. The remaining signal samples will be produced on distributed computing sites with lower priority relative to all other types of jobs. Although different strategies are under discussion, there has been little progress on merging short runs with the same or very similar conditions into single MC multi-runs productions due to lack of person-power. Running short- and medium-productions locally at KEKCC alleviates the problem for MC signal productions.

Skim production for the entire Run 1 started in mid-January and was already more than 90% completed for both the data and the MC at the time of BPAC, well ahead of the May target date. The new flagged skim strategy put in place, which groups skims with large overlaps and adds a flag, shows impressive improvements in reducing load on the production system. The number of productions has been reduced by more than a factor of 3 while the number of output datasets has decreased by a factor of 11 and the total output size by a factor of 2. The committee warmly congratulates the Data Processing group for this achievement that has resulted in a very efficient turnaround in providing skims for physics analysis.

The committee received an update on the prompt processing of Run 2. Data is first calibrated and then processed in *buckets* spanning about 2-3 weeks each. The calibration speed has decreased significantly with respect to what was reported in September for the first *bucket*. Some time was needed to setup and start data processing after calibration, as well as to prepare the payloads for MCrd productions. The experience gained with *2024a/b* data allowed for an almost prompt start of processing for *2024c*. Reprocessing of all Run 2 data and MCrd was in progress with a patched software release to mitigate the tracking problems reported in 4.1.2.

4.2.1.2 Concerns

- No particular concerns.

4.2.1.3 Recommendations

- Continue to strive to keep the delay between processing phases as short as possible.

4.2.2 Calibration

4.2.2.1 Status

Several causes for the delays in prompt calibration of 2024 samples were identified and are being addressed. Better communication and strict collaboration with detectors and operation via the technical board, coupled with improved responsiveness and experience of experts already resulted in a significant reduction in the calibration time for 2024 later *buckets* with respect to the extremely long time of the first. Effort is being put into automating the calibrations that are still run manually. A review of all calibrations to be used in the next data taking took place in January. Several improvements have been identified and are being implemented. A second review and extensive tests are then expected to take place. Problems with the complex Airflow automated job scheduling system required manual intervention by Data Processing coordinators. Following the committee's past recommendations, *b2luca*, a prototype for a new calibration system has been developed based on *b2luigi* tool, already in use by other systems in Belle II. It removes critical AirFlow complexities keeping things as simple as possible to ease operation and maintenance. It is a simple Python package running in the Belle II environment, it runs in a single place at the calibration center without the need of a dedicated server additional layer. Subsystem responsibles interact with the system via GitLab with no need of a dedicate web site. The main components of *b2luca* are in place and tested. A fully working prototype was expected by the end of March with full scale calibration tests planned with *release-09* in April. New additions to integrate MCrd special calibration productions and collection of MCrd specific conditions are expected by the summer to further speed up the overall calibration process. Regular full scale tests are planned before the start of future prompt calibrations to avoid identifying and solving problems during production. The committee congratulates the Belle II collaboration for forming the dynamic small team that is developing *b2luca* in such an efficient manner.

4.2.2.2 Concerns

- Although a new small team has been formed for the development of the new calibration system, person-power to manage the calibration is at a minimum.

4.2.2.3 Recommendations

- Should be ready to intervene quickly in case problem are observed with the new system, both in the testing phase and in prompt calibration.
- Identify additional person-power to manage the calibration.

4.2.3 Core software

4.2.3.1 Status

Release-09 of *basf2* is on track for use on HLT for *2025c*. An extensive validation campaign was carried out to verify that no degradations were introduced and to confirm the expected improvements, e.g. better K_S^0 finding efficiency due to the introduction of CDC FastBDT hit clean up. However, further improvements are needed in the validation of a new software release, as shown by the example of the issue of *release-08* which decreases the CDC tracking efficiency, that also affects *release-09*. In particular, the last round of validation is now based on collision data with the highest beam background. Future validation will also include Run 2 data.

One of the main efforts of the software group continues to be to reduce the CPU footprint of the software and algorithms. In *release-09* a FastBDT[2] has been developed to reduce background hits and speed up CDC track finding. HLT validation on a physics runs shows a large benefit from switching off modules there are not used in the HLT filtering. An additional benefit from switching off PID modules making use of MVA[3] methods is a large memory footprint reduction. Such reductions is not possible for offline processing and MC productions without degradation of PID performance. The software group is investigating ONNX as a possible solution to address the memory footprint in the underlying Python packages.

The freeze of new features for *release-10* is scheduled for July, in line with one major release per year. Many improvements in reconstructions and calibration are foreseen, alongside the support for C++20. Major updates for ROOT (version 6.24 to 6.32) and Pythia (8.215 to 8.3X) are foreseen. *Release-10* is currently planned to be used on HLT for 2026a as well in *proc17* and *MC17* in 2026.

Basf2 makes use of a large number of external software packages. The software group deploys software distributions for EL[4] and Ubuntu and relies on private tests by developers for external packages. The current system that is expected to be sustainable with a small increase in computing resources.

The *b2luigi* workflow manager is now the basis for many Belle II tools: the Systematics Framework, VIBE[5] and the new *b2luca* calibration system.

4.2.3.2 Concerns

- The proliferation in the use of Machine Learning tools is introducing new challenges for the computing performance and maintainability of the software. Memory footprints and API changes in the underlying packages used need to be addressed.

[2]Bosted Decision Tree
[3]Multi Variate Analyisis
[4]Enterprise Linux
[5]Validation Interface for Belle II

4.2.3.3 Recommendations

- Exchange with HEP community on developments for Machine Learning tools sustainability.
- Keep improving validation extending its coverage
- If new person power should become available evaluate the use of LCG distributions, weighting the investment for a complete redesign of how external software is package and built versus long term sustainability.

4.3 Computing

4.3.1 Status

The core and distributed computing project has undergone an intense and productive year. Significant upgrades were carried out on both core computing tools and infrastructure. DIRAC was successfully upgraded to version 8, coinciding with the migration of operating systems from RHEL7/CentOS7 to RHEL9/Alma9 in the distributed computing site. A further upgrade to DIRAC version 9 is planned, which will pave the way for a step-by-step transition to the next-generation system, DiracX, scheduled for deployment toward the end of 2025. Rucio was also updated to its latest long-term support release, and its servers were migrated to RHEL9. In parallel, the metadata catalog system is set to evolve: the current AMGA tool is expected to be replaced by the metadata catalog integrated into Rucio. The KEKCC upgrade was completed as well, albeit with some delays due to the commissioning work. Despite this, computing services remained operational throughout at reduced scale, ensuring continuous operation. The full migration of the Belle II computing system was finalized in time for the physics run and before the retirement of the old servers.

Looking ahead, the committee was presented with an ambitious yet feasible computing plan for 2025. This includes the adoption of new authentication and authorization mechanisms in line with WLCG, the upgrade and deployment of DiracX, although progress on a working prototype has been slower than anticipated, and the transition to a Kubernetes-based deployment model. The latter poses some challenges, especially given that Kubernetes is not currently supported by KEKCC. Resolving this will be crucial to keep the deployment plan on track.

The operation of the computing system over the past year has been very smooth, with several performance improvements resulting from upgrades carried out during LS1. For instance, the transfer of RAW data from the online system to the Grid SEs has been significantly accelerated. The time elapsed between the end of run and the availability of those data on the Grid is now typically under one hour. Similarly, replication of data to the Raw Data Centers is generally completed within two hours. Improvements have also been observed in the HLT skimming of RAW data, where some steps have been automated, although a few still rely on manual intervention.

Within the production system, the introduction of a daily reporting has proven effective in quickly identifying stalled production workflows and recurring issues. This has helped prevent long tails in processing time that were previously common. Another notable improvement relates to the long-standing problem of small files. Since June, a new processing mode, allowing jobs to run without strict run boundaries, has been in production, and this has led to a substantial reduction in the number of small output files. However, some elements of the workflow, such as the production of background overlay files, are still executed on a run-by-run basis, limiting overall efficiency, particularly in the case of *MCrd* production.

Disk space in the RAW Data Centres appears to be under control, while the implementation of automatic unstaging of RAW data put on disk for reprocessing workflows is not yet available. There is also the need to keep the two most recent generations of MC production until the new generation is complete and validated.

One performance concern that persists involves single-core jobs processing large data files. When such jobs take more than 24 hours to complete, they become increasingly vulnerable to computing or network instabilities. To address this, Belle II has begun experimenting with multicore job execution. In particular, *basf2* jobs using 8 parallel processes are now being deployed in 'multicore' queues at the Raw Data Centers. This approach shows promise in significantly reducing execution times and improving overall resilience.

4.3.2 Concerns

- The Kubernetes-based deployment model is becoming the standard across the industry. However, not all the Belle II computing sites, especially KEKCC, are supporting Kubernetes.

4.3.3 Recommendations

- It is recommended to continue and expand the deployment of multicore jobs across all RAW Data Centers, including for MC production, reprocessing, and other workflows.

5 Physics results

5.1 Status

Since the last remote BPAC meeting in September, the Belle II collaboration has made significant strides in producing compelling results. Twelve new analyses have been completed, with at least 10 targeted for presentation at the Moriond Electroweak and QCD conferences. The Belle II collaboration has already surpassed various Belle results with half the integrated luminosity of Belle, by leveraging its superior detector performance and advanced analysis techniques.

The new results include leptonic, rare, and non-leptonic B-decays, including studies of time-dependent CP-violation in $B \to \rho^+\rho^-$, charm and tau decays, as well as QCD studies with quarkonium decays and searches for dark sector particles. The highlights include the first combined analysis of R_D and R_{D^*} performed by Belle II, which has been one of the focuses of attention in winter conferences, as well as a search for the decay $B^0 \to K^{*0}\tau^+\tau^-$ with leptonic tau and hadronic tag where the combination of Belle and Belle II data giving a sensitivity per event twice that of the previous best. Of great relevance are also the evidence for the $B^+ \to \tau\nu$ decay, with the branching ratio of $\mathcal{B}(B^+ \to \tau\nu) = [1.24 \pm 0.41 \pm 0.19] \times 10^{-4}$, and studies of CP-violation in $D^+ \to \pi^+\pi^0$ decays. The collaboration has also anticipated the release in the summer of results on a number of electroweak penguin mode of the type $B \to X\nu\bar{\nu}$ (with $X = K, K\pi, K2\pi, ...$ etc.), which would be very useful in constraining extensions of the Standard Model. On the spectroscopy side, there are tantalizing indications of a possible state in the B^*B^* channel. Furthermore, Belle II completed an interesting search for an axion-like-particle produced from B meson decays and decaying into two photons. This search led to the world-leading bound on the axion parameter space, surpassing the corresponding BaBar search.

The collaboration is achieving a steady state in the number of publications. The committee looks forward to seeing further progress with the significantly increased statistics from the upcoming data collection and analysis of Run 2 data.

The collaboration continues the Physics Week meetings in the new format introduced in 2023, with each meeting focusing on a specific physics topic of interest. Last year's topic was tau and dark sector. This year's theme will be rare decays of beauty and charm hadrons with missing energy.

5.2 Concerns

No particular concern.

5.3 Recommendations

- While the Belle II dark sector program has been highly successful, emphasising more on a generic and wider aspect in the analysis could be beneficial.

- With steadily increasing collected dataset, the committee encourages increasing the commitment toward core B-physics results, such as leptonic and semileptonic decays, the Lepton Flavour Universality ratios, and rare electroweak modes with missing energy ($B \to X\nu\bar{\nu}$). Careful planning in this area is very important given the length and complexity of the analyses.

6 Upgrade

The charge to the committee for this meeting included the question of whether the procedure and timeline to define the objectives and make the detailed schedule of SuperKEKB and Belle II upgrades in Long Shutdown 2 (LS2) are clear. During the meeting a new date for the start of LS2, around 2032, was presented. While this is more realistic for the upgrade work to be completed than the previous date of 2028 considered in the CDR for the Belle II upgrade, it introduces some concerns whether the subsystems will be able to continue to operate until that date with the required performance, in particular for the CDC. At the time of the CDR it was hoped that the CDC would not require replacement as part of the upgrade programme. With the indications of possible ageing and gain loss that have been discussed earlier in this report, it is becoming clear that action will be required to ensure that the CDC continues to operate successfully as the luminosity increases, and that a replacement or upgraded detector should be investigated to be ready on the timescale of LS2 (if not before). Judging from the work plan presented at this meeting for the new pixel vertex detector (VTX) and final-focusing quadrupole magnet (QCS), it may be the case that LS2 could start even later as far as those two systems are concerned. For this reason, the committee encourages the Belle II collaboration to consider the decoupling of the CDC upgrade from the VTX/QCS upgrade. As soon as the layout of the interaction region with the QCS is completed, the design work of a new CDC should start in such a way that the new CDC can be installed in the present Belle II set up as well as with the new QCS/VTX configuration. This will make the upgrade plan more flexible, while maintaining the detector performance until LS2.

6.1 Physics motivation

6.1.1 Status

The physics motivations for reaching $50\,\mathrm{ab}^{-1}$ in the long term are largely unchanged with respect to the original Belle II proposal. The three reference measurements chosen as highlights of this programme, namely measuring the ϕ_2 angle of the CKM Unitarity triangle with 0.6° precision, reaching a below per-mille precision on the direct CP asymmetry in $D^0 \to \pi^0\pi^0$ decays, and the search for the Lepton Flavour Violating $\tau \to \mu\gamma$ decay below the 10^{-8} level, are illustrative of the diversity of such a programme. However, they do not do justice to the broad scope and potential breadth of the programme. Just to mention a few additional examples: i) the leading di-neutrino decay modes $B \to K^{(*)}\nu\bar{\nu}$ should reach the discovery level, providing key constraints on several beyond-the-SM frameworks; ii) improved measurements of exclusive and inclusive $b \to c\ell\nu$ transitions should allow to resolve the long-standing puzzle in the determination of $|V_{cb}|$; iii) a similar statement holds for exclusive and inclusive $b \to u\ell\nu$ transitions and the long-standing $|V_{ub}|$ puzzle; iv) the ratio $B(B \to \mu\nu)/B(B \to \tau\nu)$ would provide a new unique test of Lepton Flavour Universality.

6.1.2 Concern

- With 50 ab^{-1} of data, ensuring systematic errors to be below statistical ones will be a challenge. A real good understanding of the detector performance will be required and simply increasing Monte Carlo statistics would not be a viable solution.

6.1.3 Recommendations

- A strategy must be developed for controlling systematic errors to the necessary level for the key measurements.

6.2 SuperKEKB luminosity upgrade and new final focusing magnet (QCS)

6.2.1 Status

In the current IR, the last vertical focusing magnets for the LER (QC1P- Right and Left) are the quadrupole magnets closest to the IP. These magnets are surrounded by a compensating solenoid for the detector solenoid, but this part of the solenoid compensation is the over-compensation part which is set to cancel the integral of the detector solenoid field between the cryostats. Hence, these final focus magnets for the LER are experiencing a strong solenoidal field of over 2 T. This generates a large chromatic x-y coupling aberration which must be corrected by external rotating sextupoles. This is difficult to do correctly because the optical functions must be known with high precision and there must be good reproducibility of the magnetic fields of the local magnets. The new design of the cryostat in this region is to move the QC1P magnet coils 10 cm closer to the IP. This then leaves room just behind these coils for the corrector windings. In addition, the over-compensating part of the solenoid compensation design is placed in front of the new QC1P coils thereby minimizing the solenoidal field in the QC1P coils. This would move the front face of the cryostat from the present location of about 60 cm to about 30 cm from the IP. Moving the QC1P coils closer to the IP means the coils will need to become thinner because the beam trajectories are rapidly converging. This is what has started the study of using Nb_3Sn superconducting magnet wire for these new coils. The new design will also use Fe shielding instead of a cancelling coil which is presently used for the leakage field seen by the HER. The Nb_3Sn wire will produce a more robust magnet that will have a higher temperature margin. The ARC review in January encouraged the team to look at a 2-layer design as opposed to a 1-layer design and the preliminary results of this study look encouraging. The temperature margin is increased and the current density is decreased with the 2-layer design although the requirement for the wire thickness become tighter. The harmonic content of the magnet can also be more carefully controlled. Work is ongoing to construct a one-quarter mirror magnet (one of four quadrupole coils) to test the feasibility of the design.

6.2.2 Concerns

- The use of Nb_3Sn wire is relatively new and may have unforeseen difficulties when being used in the construction of these dimensionally small coils.

- The new design could dramatically alter the space available around the collision point.

6.2.3 Recommendations

- Proceed with the testing of the Nb_3Sn wire to quickly learn whether it can be used in the new magnet design.
- The two-layer solution looks very promising and the committee encourages the team to look into making a prototype coil similar to the one-layer coil now under construction.
- As was mentioned in the presentation, the space assignment around the IP will need to be carefully planned together with the detector team and the committee encourages the formation of a subgroup which includes detector and accelerator people to study this issue.

6.3 Vertex detector

6.3.1 Status

A new vertexing and tracking system (VTX) for the Belle-II detector, to be installed during LS2, is being proposed. The new detector targets a higher space-time granularity, lower material budget and higher radiation tolerance. There is in addition interest in achieving individual hit time stamping at the ≈5 ns level and the possibility to provide inputs to a level-1 track trigger. The conceptual design calls for a straight geometry with two inner layers, starting at a radius of 14 mm with a material budget of 0.2% X_0/layer and three to four outer layers with 0.85% X_0/layer with an outer radius of 140 mm. The design is being optimized based on full simulations of benchmark physics channels.

A single sensor, a depleted monolithic active pixel sensor, would be used throughout the new detector. The sensor envisioned is the Obelix sensor, based on the TJ-Monopix-2 design, which has been tested in beam tests at DESY and KEK. A next test beam campaign is being planned for April 2025 at DESY. This will inform the final design of a first full-size prototype sensor, the Obelix-1. This is a 869×464 matrix of 33×33 μm^2 pixels with a trigger unit expected to handle a hit rate of 120 MHz/cm^2 and fine time stamping, at limited hit rate, of 6 ns. A track trigger with strixel size of 4×18 mm^2 is also implemented. Tape-out is expected in the July 2025 time frame.

First ideas for ladder design for the inner and outer layers are being developed. The baseline design for the inner layer calls for water cooling with a flow rate of 0.2 ml/s, giving a maximum sensor temperature of 34 ^{0}C. The outer layers have a carbon fibre and rohacell support structure. A first pass system design envisions an interface board with radiation hard optical transceivers and DC-to-DC converters interfacing through high-speed optical links to the Belle readout architecture. The lpGBT chip, designed for operation at 40 MHz for the LHC experiments, is being studied for operation at 42.4 MHz required for Belle II.

Within the ECFA DRD3 collaboration a 6-layer telescope is being developed based on the Obelix-1 chip, which would benefit the development of the VTX detector. An optimistic schedule calls for the final sensor to be available at the end of calendar year 2027, at which time a full technical design report will be submitted targeting a completion of the new VTX by 2032. If a second iteration of the Obelix sensor is required, the schedule would be stretched by an additional two years to 2034.

6.3.2 Concerns

- The design of the new VTX strongly depends on the design of the new QCS final focus magnets. There is a concern that the VTX design efforts may not be fully synchronized with the final design of the QCS.
- The assumption that only two submissions will be needed for the Obelix chip seems optimistic.
- The fine time stamping requirement of 6 ns seems not compatible with the required hit rate of 120 MHz/cm^2.
- To implement a track trigger, the granularity is currently limited to eight strixels per sensor due to bandwidth limitations. It is unclear if the granularity is fine enough to provide for an efficient track trigger. Moreover, the anticipated fake rate will be very high.

6.3.3 Recommendations

- Fully map the operation of the TJ-Monopix-2 sensor in the upcoming DESY test beam with both unirradiated and irradiated sensors, in particular the cross-talk from time stamping, operating temperature of the sensors, power consumption and overall efficiency.
- Re-evaluate the hit timing resolution and track trigger design and flow down the requirements for the sensor based on simulation of physics data under real (high background) conditions to ensure it will be effective for the physics program.

6.4 Central drift chamber

6.4.1 Status

As described in Section 3.2, the long-term performance of the CDC is a serious concern for various reasons. The current CDC can continue to be operated, but with slowly decreasing efficiency. The evolution of the performance degradation is hard to predict given the large uncertainties in the machine background conditions. The projected operating currents at higher luminosities are in excess of 1 mA per layer, which are beyond the limits of safety margin for long-term CDC operation. Some mitigation measures such as reduction in gain by a factor of two in the innermost layers are being considered at the start of Run 3.

The collaboration is at the moment considering three options. Continued use of the current CDC is feasible if the Malter effect and ageing effects are limited, with possible reduction of the HV. If the inner layers of the CDC are heavily affected, construction of a new drift chamber for the outer layers using a similar technology as the current drift chamber, complemented by the installation of silicon pixel detectors covering the current CDC inner layers, is being evaluated. If the CDC degradation effects cannot be controlled, construction of a new drift chamber may be necessary before the next shutdown. This would necessitate a revision of the long-term schedule, in particular the start of LS2. Hence, it is critical to complete the performance studies and the ageing studies with the two test chambers. The collaboration notes that removing the inner small-cell chamber is not a viable option due to unacceptable risk of damage to the outer volume. Intense effort is required to understand the origins of the observed CDC performance degradation, which will be critical to find a path to follow for the success of the Belle II physics program.

6.4.2 Concerns

- With the new operation plan in which Run 2 will continue for another seven years, the CDC performance may not stay at the level required for a competitive physics programme until the end of Run 2.
- To date no coherent effort has been started to optimise the whole tracking system for the LS2 upgrade.

6.4.3 Recommendations

- The Belle II collaboration is encouraged to start optimising the complete tracking system for the LS2 upgrade. The formation of a working group charged with developing a coherent plan is encouraged.
- A strategy should be developed for the case where a major intervention is needed for the CDC in order to maintain the performance until the end of Run 2. This could include the construction of a new CDC that could fit in both the current geometry and with the QCS modification for the LS2 upgrade. This requires close discussion among the accelerator, VTX and CDC teams.