B-factory Programme Advisory Committee
Full Report for
Focused Review Meeting

24-25 September 2024, Remote meeting

A. Andreazza* (Milano), P. Collins* (CERN), G. Corti (CERN),
M. Demarteau (ORNL), R. Forty (CERN), S. Gori (UCSC),
G. Isidori (Zurich), P. Mato* (CERN), F. Meijers* (CERN),
N. Neufeld (CERN), N. Pastrone (Torino), A. A. Petrov* (South Carolina),
B. Ratcliff* (SLAC), M. Sullivan (SLAC), H. Tajima (Nagoya),
O. Tajima (Kyoto), M. Titov* (Saclay)
and chaired by T. Nakada (EPFL)
* Expert member.

22 December 2024

1 Short summary

A focused review meeting of the B-factory Programme Advisory Committee (BPAC) took place remotely on 24th and 25th of September, where the committee heard presentations from the accelerator and Belle II groups. This section gives the committee's feedback on the four questions asked by the management of the Institute of Particle and Nuclear Studies. A detailed report on the broader findings of the committee can be found in the following sections.

1. **Optimisation of Short-Term Run Plan: Is the current short-term run plan effectively optimised to meet the immediate needs of data collection and machine performance improvements?**
   The committee considers the sudden beam loss phenomena as currently the most critical issue. Therefore, it supports the plan giving priority to the machine study and development to overcome the problem of sudden beam loss, while operating the detector in a safe way against radiation damage. Once flipping some of the beam pipe sections with electrodes and knocking of the beam pipe show improvement, increasing the beam currents that is mandatory for achieving higher luminosities could become more feasible. Reaching a luminosity of $10^{35}\mathrm{cm}^{-2}\mathrm{s}^{-1}$ during the coming run period would be an important target. Evaluating different strategies with the increase of beam currents and decreasing β^* might be useful

depending on the machine behaviour. The committee supports the plan to flip the remaining beam pipe sections with electrode as much as possible during the coming shutdown periods. Injection efficiency and stability are key elements for achieving not only high instantaneous luminosities but also high integrated luminosities. The machine group has already been introducing many hardware improvements for stable operation of the injector chain, and the committee strongly encourages this effort to continue. Solving the problem of emittance increase and achieving efficient two-bunch injection are among the priority items. All those efforts could benefit from having an organisational structure with clear authority for defining priorities and further consulting experiences from other accelerators. Having such an organisation could also help to obtain the required resources and expertise.

2. **Belle II Contribution to SuperKEKB Operations: Is the contribution of Belle II to the operation of SuperKEKB well organised? Are there specific areas where Belle II's contributions could be enhanced? If so, what suggestions can be made?**
   The committee finds that the machine group and Belle II collaboration have been working together very well. The Belle II team contributes significantly to the hardware and analysis effort to improve the machine operation to achieve the luminosity goal with an acceptable background level. A particularly important area where the joint effort could be strengthened is to develop further ways to detect the start of beam instabilities resulting in sudden beam loss and abort the beams as quickly as possible.

3. **Detector Operation Concerns: Have all issues and concerns related to the operation of the Belle II detector been appropriately addressed?**
   The overall status of the Belle II detector is good. The detector should be able to collect data with the quality needed for the physics goals in the coming run. Concerning PXD2, the committee supports the decision to keep it switched off until the sudden beam loss phenomenon is under control. At the same time, the PXD group should establish beam condition criteria for switching the detector on and develop scenarios for its operation once it is switched on, optimising for the detector safety and physics output. The committee stresses again the importance of retaining crucial technical knowledge of departing personnel. The CDC team should continue the effort to establish a clear picture of the gain loss due to the beam background. This will affect the upgrade strategy of the Belle II tracking system.

4. **Physics Analysis Planning: Are the plans for physics analyses well-structured and on track to achieve significant results in time for the upcoming winter conferences?**
   The Belle II collaboration has been successfully producing interesting results exploiting the unique features of the experiment. The committee appreciates the state of analysis activities, planning and organisation as satisfactory, and is confident that further interesting results will emerge for the forthcoming winter con-

ferences.

The committee congratulates the accelerator group and Belle II collaboration on their achievements and for maintaining the successful data taking at the highest luminosity e^+e^- collider. It must be stressed that providing an adequate running time in the coming years is crucial for the Belle II experiment to remain competitive.

2 SuperKEKB and injection complex

2.1 Status

During the long shutdown 1 (LS1) several repairs and improvements were made to the injection system and to the main rings of the SuperKEKB accelerator. In the main rings, several collimators were repaired that had been damaged by Sudden Beam Loss (SBL) events. In addition, several more diagnostic monitors were installed around both main rings to assist in uncovering the source point for the SBLs. The injection system for the electron beam got a larger aperture beam pipe at the injection point into the main ring. The main injection system was upgraded with new fast pulsed magnets to improve the injection orbit of the second bunch for both beams. These improvements were helpful in improving the injection efficiency for the HER and in reducing the background levels from the injected bunches in both rings. As a result, a new peak luminosity of $4.47 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ has been achieved and over 100 fb^{-1} of integrated luminosity has been added to make the total luminosity integral over 500 fb^{-1}. Maximum beam currents of 1.2 A for the HER and 1.5 A for the LER, although below the target values of 1.4 A and 1.8 A respectively, also helped to achieve the new luminosity peak value.

In the first run after LS1, substantial effort went into searching for the sources of the SBL events that have been observed in both main rings. The detector team and the accelerator team have worked in close collaboration on this very difficult issue. There is now strong evidence that dust particles coming from clearing electrodes mounted on the inside top wall of the beam pipe in the LER wiggler sections is a source of many of the SBL events. During this summer shutdown, 15/50 beam pipe sections with clearing electrodes were turned upside down. In addition, a beam pipe knocker was applied to the remaining chambers with clearing electrodes. This knocker mechanism was successful in producing an SBL event during the run and there is evidence that the SBL rate was reduced after the knocker was applied to several chambers during the run.

The goals for this fall run are to verify that the chamber flipping reduces the SBL event rate, to achieve $1 \times 10^{35}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ luminosity, and to integrate as much luminosity as possible with the help of improved HER injection efficiency, collimator repairs, and the added fast kickers in the Linac.

2.2 Concerns

- While the finding of SBL events coming from dust events in the clearing electrode chambers is a great step forward, there are probably more sources of SBL events in both rings.

- Due to the risk of further damaging the PXD, the PXD team has decided to turn off the detector.

2.3 Recommendations

- Try to find out as soon as possible whether the chamber flipping is successful in reducing the number of SBL events. If successful, develop a plan for flipping the rest of the chambers as soon as possible. Consider extending the winter shutdown in this planning. The committee encourages further studies of injection efficiency in both main rings.

- Continue the efforts to minimize the time between detecting the beginning of an unstable beam and aborting the beam. Finding the sources of, and mitigating the SBL events, may be a long-term endeavour and as the beam currents increase, new sources of SBL events may surface in both rings.

- Develop a plan for how and when to switch on the PXD (see below), in particular by understanding which SBL events results in large radiation doses delivered to the IR and implementing mitigation strategies.

3 Belle II detector

3.1 Vertex detector (VXD)

3.1.1 Pixel detector (PXD)

3.1.1.1 Status

SuperKEKB started running again after LS1 on January 29, 2024 with physics data taking starting on February 20. The physics run, run 2024ab, ended on July 1. Run 2024c is expected to start on 9 October and run until 27 December. During run 2024ab a peak luminosity of $4.47 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ was reached with 2249 bunches in the machine with currents exceeding 1 A in both the HER and LER. The PXD started taking data with 35 out of the 40 modules operating. Module 2.12.2 was turned off because of excessive noise and all modules in ladders 2.7 and 2.8 were turned off because of the bowing issue. Detector operation was mostly smooth and required little intervention. Good performance was shown with hit efficiencies in both layer 1 and 2 exceeding 98%. Unfortunately on April 22 and May 6, there were sudden beam loss events that damaged a small fraction of the PXD. After the first event an increase in the switcher current was observed that increased the PXD temperature and two additional ladders, 2.4 and 2.9, had to be turned off. Moreover, 84 gates died in the accident. 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. How much dose the detector has received is unknown, since the diamond sensors, including the ones used for the abort function, were saturated for some of the samples. To avoid further damage the PXD was turned off on May 7 for the remainder of the run. During the run 2024ab a total

luminosity of about 45 fb^{-1} was taken with the PXD and about 58 fb^{-1} without the PXD.

The increase in switcher current after the first incident led to an increase in temperature, in some areas by as much as 5 °C as measured by the FOS sensors. The glue joints were stable, but are temperature dependent and are known to become soft at temperatures above 45 °C. Ladder bending was also monitored and found to be acceptable, noting that ladders 2.7 and 2.8 were off. Since these two ladders are turned off, no alignment with data is possible and the ladder bending has to be determined from mapping the nuclear interactions in the ladder. This study is ongoing. It was also noticed that the ladder bowing is time dependent and decreases with increasing beam current. The displacements are modest, on the order of 100 μm, but implies that a time-dependent alignment has to be implemented, which is being done.

While running the detector it has been observed that some Data Handling Processor (DHP) links are unstable and introduce a truncation of the data and sometimes even drop the link. This has a negligible effect on the PXD efficiency, but it is noted that this did not occur this extensively for the PXD1 and its origin is unclear. It seems that for some ASICs it is very difficult to find an optimal setting.

The manpower situation to maintain and operate the PXD is marginal and could see a further reduction.

3.1.1.2 Concerns

- The vulnerability of the PXD to sudden beam loss events is a grave concern. The committee fully supports the decision to keep the detector turned off until the source of the sudden beam loss events is understood.

- The personpower to operate, maintain and analyse the data from the PXD has been marginal for a while. With key members of the construction and operation team leaving the collaboration, this situation is exacerbated. Most PXD institutions are funded by the German ministry (BMBF), but even here cuts are being felt and one institution is no longer receiving funding. This is another grave concern that the collaboration has to address.

- The instability of the DHP links results in a minor loss of data and understanding the root cause would increase the overall data taking stability of the experiment.

- The fact that the origin of the sudden beam loss events is not understood, and could take a while before it is adequately addressed, implies that the physics reach of the experiment is adversely impacted given that the PXD is turned off. Furthermore, it is difficult if not impossible to guarantee, even if the utmost care is given, that future damage to the PXD will not occur. The experiment has to date only collected about 530 fb^{-1} of data, a mere 1% of its final goal. The idea that the PXD may stop working prematurely is of serious concern. The collaboration is encouraged to initiate in earnest discussions on developing plans to mitigate this risk.

3.1.1.3 Recommendations

- Define the criteria that have to be met, and develop well-defined policies before the PXD can be turned on again. This should include a minimum number of hours of operating the machine at nominal currents without any incident.
- 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.
- Engage with DESY, where most expertise for the PXD resides, to ensure that the current support for the PXD is maintained, and possibly increased. Engage the collaboration to identify additional support for this critical detector.
- Start developing a realistic schedule for the activities to be carried out during LS2 related to the VXD, including the option for (partial) replacement in close collaboration with the accelerator group and KEK management.

3.1.2 Silicon strip vertex detector (SVD)

3.1.2.1 Status

No dedicated SVD status report was given at this review. It was reported that the SVD operates well at the same level as during Run 1.

3.1.2.2 Concern

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

3.2 Central drift chamber (CDC)

3.2.1 Status

Many improvements for the CDC gas system have been accomplished during LS1: water and oxygen sensors were re-calibrated and are now better controlled, additional sensors for comparison and calibration, as well as extra gas lines to increase gas flow circulation rate, have been installed. At the beginning of Run 2024a, gas conditions were relatively stable and the CDC performance was much better than in June 2022. However, the dE/dx showed strange degradation after late April that cannot be fully attributed to O_2, H_2O content or beam background conditions. In particular, the relation between decrease of dE/dx mean (gain drop) as a function of the HER and LER backgrounds (CDC chamber currents) is similar to Run 2024ab. The change of resistors on HV

lines did not have any significant impact. Understanding the reasons for CDC gain degradation in the latter half of Run 2024ab is still under investigation.[1]

The DAQ instability, due to CDC front-end Single Event Upset (SEU), increases with the beam backgrounds for larger HER and LER beam currents. The CDC error handling in case of SEU is currently under discussion, while run stops due to the CDC in Run 2024b represent approximately 11% of the total cases.

The ageing test with a test chamber have started in June 2024. The total charge of 66 mC was accumulated on the target wire in June and July. Simulation of a beam-spot on a wire is ongoing to study the charge per unit length (mC/cm). The next benchmark will be to irradiate the test chamber up to 1 C/cm.

3.2.2 Concerns

- The relation between the decrease of dE/dx mean (gain drop) as the function of the HER and LER backgrounds (CDC chamber currents) is similar to Run 2024ab and is still not well understood.
- Appearance of Malter-like currents in several CDC layers during the October physics run (after the BPAC meeting) represents a major risk for the CDC future.

3.2.3 Recommendations

- Quantitative separation of gain drop reasons due to the gas composition or the beam backgrounds needs to be further advanced.
- Systematic real-time monitoring and alarms for any sporadic activity in all CDC layers, such as HV trips or increase in leakage current, should be implemented and available for the shift crew.
- Ageing studies with the laboratory test setup should continue to reach the next benchmark of $\sim$ 1 C/cm and to determine the CDC chamber longevity. On a longer term, this might facilitate development and test of different remediation strategies for gain recovery, if anode ageing or Malter-type effects are observed.

3.3 Particle identification system

3.3.1 Barrel system (TOP)

3.3.1.1 Status

During 2023, the TOP group successfully performed the planned LS1 replacement program for the PMTs, and exchanged and repaired the front-end electronics, all within the scheduled time frame. This excellent work was discussed in the June 2023 report. These

[1] Following the BPAC meeting, a major onset of Malter-like current effects have been observed in many CDC layers during the October physics run. This represents a major risk for the future CDC operation and should be followed up and reported at the next BPAC meeting in March 2025.

newly installed PMTs increased the number of photons observed in Run 2 (2024ab) compared to Run 1, especially in the bottom region. As expected, the repair of the electronics substantially increased the number of active channels. No significant drop in PMT quantum efficiency (QE) was seen during runs 2024ab, as expected, since the accumulated charge was small.

Additional DAQ projects to allow more efficient TOP operations and monitoring, that were undertaken during LS1 to improve the data taking efficiency during 2024, performed essentially as expected during runs 2024ab. Further firmware modifications were undertaken during summer 2024 to provide an internal signal veto during injection in the front-end (FE) readout to deal with a large DAQ deadtime when the beam injection quality is bad. Feature extraction is also being moved from the FE firmware to the processor on the readout PC so that the data transmission will not be interrupted even if the processor in the FE crashes due to SEU. Both of these improvements will be deployed before the start of run 2024c.

The PMT 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 flatter in bench tests than predicted from the collision data. Some small degradation (relative to 25°C) may be observed at higher temperatures (50°C) in the bench test but it is probably about an order of magnitude too small to explain the difference between the bench and collision rates of QE loss. The Nagoya group is now planning to measure the B-field effect in bench tests.

3.3.1.2 Concerns

- The literature on ageing is complicated and confusing, but it is unclear that anode charge is the most relevant variable for cathode ageing if other effects, which change the overall gain and counting rates substantially, are changed. Hence, the fact that the bench tests are at much higher gains and different counting rates may be problematic. Ageing studies in a B field at gains similar to those of the collision data could be very helpful.

- Because of the order of magnitude differences in the measured ageing rates between collision and bench data, it is still unclear whether high temperatures might play a substantial role in the collision data.

3.3.1.3 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 much as possible.

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 to date.[2] There was no specific ARICH detector presentation during this review. Some new problematic HAPDs were found at the end of LS1, but no new HAPDs died during run 2024ab. The number of dead channels is stable at 6.6%.

3.4 Electromagnetic Calorimeter (ECL)

3.4.1 Status

ECL operated stably during the last run. Thus, there was no specific ECL review. It took some time to fix a pedestal instability, while it was indicated in the data quality monitor (DQM); this is related to the concern about human resources.

3.4.2 Concern

- There are relatively few human resources (i.e., lack of eyes) for detector maintenance and data quality monitoring compared to other detectors. For example, the pedestal instability shown in DQM might have been more promptly fixed if the resources were sufficient.

3.4.3 Recommendations

- Organize systematically the off-site data quality monitoring data, and provide quick feedback from it.

3.5 K-Long Muon Detector (KLM)

3.5.1 Status

The low efficiency of RPC BB2 layers showed up in June 2021 when the gas circulation was interrupted and the presence of water vapour contamination may have caused damage of the RPCs. Subsequent tests in 2021-2022 revealed that the circulation of fresh gas with high rate did not cure the efficiency problem. However, substantial recovery of RPC BB2 chamber efficiency was observed during the recent Run 2024ab. Therefore, plans for a potentially risky NH_3 injection have been suspended due to the observed recovery.

As the result of the stress from Cu-tubes and possible ageing effect of old plastic tubes, several cracks on polyethylene tubes of the bubbler flowmeter system were identified. Visual inspection for all tubes was performed and all of the tubes in one station, which correspond to 25% of the tubes, were replaced. To avoid this problem in the future,

[2] Following the BPAC meeting some unexpected temperature increase has been seen in regions of the ARICH, which should be followed up and reported on at the next meeting.

an alarm on channel-by-channel flowmeters was implemented to have a fast diagnostics, and strain reliefs were installed for the tubes.

Sudden and frequent, roughly once per day, turn-off of high voltage modules have been observed for BWD RPC's in Run 2024ab. Such events might cause long DAQ downtimes. The nature of the problem is still not understood, but the suspicion exists that it is due to a bad combination between the HV module and the CAEN crate. Note that during shutdown, HV modules were tested with a spare old-type crate for about three weeks, and no sudden HV turn off was observed. This problem is also being investigated with CAEN. The current worst-case scenario is to operate the RPC with the old-type crate in Run 2024c. However, the real issue is the availability of spares for the old crates.

3.5.2 Concern

- The availability of an adequate number of spare crates during the operation is very important for the RPC, as several RPC HV issues were observed.
- Ageing of RPC gas system components remains worrisome and needs to be regularly addressed.

3.5.3 Recommendations

- In a view of the recovery of RPC BB2 chamber efficiency during the recent 2024ab operation, the earlier consideration of potentially risky NH_3 injection should be put on hold.

3.6 Trigger and Online

3.6.1 Status

Steady progress has been made in the Data Acquisition, Online operations and Trigger. In the trigger, UT3 boards were replaced with UT4 boards as scheduled and operated without issues. FPGA resources are increased by a factor 2-4 and the optical transmission latency between UT4 boards is reduced.

The data-taking efficiency has almost reached 90%. Very good monitoring of the individual causes for inefficiency is available. Many factors have been successfully addressed or mitigation plans are in place. In particular, the Stop-Abort-Load-Start procedure to reset (parts of) the readout system and event filter has been significantly accelerated and an improved active injection veto has been deployed.

For the next run, several improvements are ready: the CDCTRG ADC filter was successfully tested (which was made possible by the new more powerful UT4 board), a new 3D fitter for the CDC and improved ECLTRG Bhabha veto have been prepared and are kept "in reserve" in case a DAQ limit will be hit, the HLT has been reinforced by additional cores, a new TOP firmware will suppress the payload for busy events to avoid DAQ dead-time and the feature extraction in the TOP firmware will be moved to the back-end (PCIe40) to make it less vulnerable to single-event upsets (SEU). The

migration of the B2TT to an optical physical layer is ongoing and should, if successful, remove a large fraction of the down-time due to link losses.

If all these measures can be deployed, they will significantly reduce the DAQ inefficiencies, which are related to the detector itself.

3.6.2 Concern

While, as noted above, an impressive number of mitigation plans and improvements have been prepared, many of the problems are related to occupancy and/or beam conditions and it is to be feared that with increasing luminosity they will not be sufficient to protect the data acquisition and trigger from efficiency degradation.

- The injection veto remains a main concern, as one of the main contributors to dead-time.
- Single event upsets will become more frequent with more demanding beam and background conditions.

3.6.3 Recommendations

- The work on an improved injection veto should be continued. The effect on HLT load is to be assessed, and the impact on physics performance to be checked.
- Moving critical logic away from SEU affected areas in the CDC and in the TOP should be followed up as much as possible.
- The new B2TT should be deployed as soon as possible.
- Regarding the modification of the ECLTRG Bhabha veto, attention should be paid to relevant DQMs and information with relevant group shared if necessary.

3.7 Data processing and software

3.7.1 Status

The committee received an update on the run-dependent *MC15rd* and the re-skimming campaign that were first reported at the February BPAC meeting. Both were completed successfully around May, well ahead of schedule. This is particularly impressive for the production of Monte Carlo (MC) run-dependent signal which had been anticipated to take more than 10 months. The fast turnaround has largely been possible by having enabled local production at KEKCC: such productions cope better with many short jobs, as expected. In parallel, the production of multiple signals in a single job is under test. The re-skimming of data (*proc13* and prompt) and MC (*MC15rd*) has also been very efficient, approaching almost a full campaign in number of skims produced in only a few months. The committee warmly congratulates the Belle II collaboration for these impressive achievements.

Re-processing of Run 1 data (*proc16*) was in progress at the time of the meeting. Data is grouped in five chunks processed sequentially. Improved and new calibrations are first run for a given chunk. Each data chunk is then processed and the corresponding run-depend MC (*MC16rd*) produced in parallel. Calibration of the following chunk starts as soon as the previous one is completed. The calibration of the first chunk took considerably longer than expected. This was due to having to run it by hand, as already anticipated at the February BPAC, and to some issues with the new calibrations. The completion of the change in Airflow to enable job submission to KEKCC allowed its use as calibration job scheduler for the other chunks. The resulting automation considerably reduced the achieved vs. expected lag time. The remaining delays are due to unavailability of experts and the KEKCC refurbishment. AirFlow is running stably and will be kept until an alternative is deployed. It was reported that no progress was made toward a replacement with *b2luigi*, with person-power still to be identified for the task. At the time of the meeting the *proc16/MC16rd* campaign was in different stages for each chunk. The committee congratulates Belle II for succeeding to provide for the first time the corresponding MC samples concurrently to reprocessed `mDSTs`. The new skimming strategy to group skims with large overlaps and adding a flag was in the last phases of verification. It was expected to be deployed in production for the first chunk (*proc16, MC16rd* and *MC16ri*[3]) around the end of October.

The committee received an update on the prompt calibration of Run 2 data. Calibration is performed in *buckets* of data spanning about 2–3 weeks each, with a full calibration completed on one bucket while data is collected for the next. The first bucket ($11.2\,\mathrm{fb}^{-1}$) took nearly six times longer than the targeted 20–25 days due to issues with new calibrations and delays from novice detector experts. Improvements in both areas led to faster calibration for the second bucket ($25.4\,\mathrm{fb}^{-1}$), reducing the time to around 40 days. Efforts are underway to further involve the detector community.

Release-08 of *basf2* is used for reprocessing *proc16*, MC production *MC16* and 2024ab data taking. A patch release with minimal changes for DAQ and DQM[4] will be used for 2024c data taking and prompt calibration. Doing so will allow thorough validation of the next *release-09*.

The main focus of the software group continues to be the reduction of the CPU footprint of the software and algorithms. Particular attention has been put in improving the slowest two, that are even slower in high background conditions. The ECL clustering algorithm has been tuned, reducing the HLT execution time by 10% with a negligible effect on the HLT retention rate. Discussions were ongoing about activating it already for 2024c data taking in the case of extreme background conditions. A significant time reduction in CDC track finding algorithms in high background conditions was obtained in *release-09* by introducing a BDT[5] to better identify CDC beam background hits and filter them before reconstruction. Validation studies were ongoing with preliminary indications of a small improvement ($\sim$ 1–4%) in tracking performance.

The committee was pleased that the collaboration will continue to deploy a major

[3]A small fraction of run-independent MC is produced for sanity checks and exploratory studies.
[4]Data Quality Monitor.
[5]Boosted Decision Tree.

release a year. Many new features in the reconstruction and physics analysis software are available in *release-09*. Validation is ongoing with MC samples and the reprocessing of a small data sample. The release was on track to be deployed on the HLT farm for 2025a data taking.

Several MVA-based algorithms have been introduced in reconstruction and analysis software in the last years, improving the physics performance. This has introduced 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. The software group is investigating ONNX[6] as a possible solution. *Basf2* relies on a large number of external software packages. The software group deploys the software for the most common operating systems used in the HEP community. With the proliferation of external packages, the task of building the software stack has become increasingly complex. With no CI[7] pipeline in place for external packages the software group relies on private tests by developers. Dedicated resources would be needed to make the system more robust.

With the transition to RHEL9, it was no longer possible to run *release-06* on a few major sites, including KEKCC. To ensure the portability and reproducibility of this older release, the approach of utilizing the *release-06* RHEL8-based images within *apptainer* was selected. This is already in place at KEKCC and will soon be deployed for Grid distributed jobs.

3.7.2 Concerns

- Fragility of relying on AirFlow as calibration workflow scheduler remains as a concern. This is even more so now, since the expert left the collaboration in October.

- Prompt calibration has been very slow, in large part due to "childhood" issues. While this is understandable after a long shutdown, availability and responsiveness of experienced experts has been a recurring issue.

- Belle II is increasingly relying on external software packages, particularly within the areas of ML tools and the Python ecosystem. Combined with the large number of supported operating systems, this makes the task of building the software stack progressively more complex. This growing demand for human resources is beyond what the collaboration currently has available.

3.7.3 Recommendations

- The committee reiterates the recommendations on Data Processing given in last February's BPAC report, i.e. urging the collaboration to provide a prototype for replacement based on *b2luigi* as soon as possible and recommending to pursue the

[6] Open Neural Network Exchange.
[7] Continuous Integration.

possibility to merge short runs with the same or very similar conditions into a single multi-run.

- For the development of large and increasingly complex software stacks, reduction of the number of supported Linux distributions and utilize containers is recommanded. At the same time, the committee suggests investigating whether the existing effort within the LHC to build the LCG software stack, currently used by ATLAS and LHCb, could be leveraged.

- To prevent software releases from becoming unusable when computing centres, especially KEKCC, migrate to a new OS version, the committee recommends ensuring that each software release is compatible with the current compiler version, as well as the previous and next versions.

4 Physics

4.1 Status

The Belle II collaboration has made significant strides in producing compelling results since the last BPAC meeting in February. In that time, 16 new analyses have been completed, with seven of these presented at the ICHEP conference last July. Notably, some analyses were collaborative efforts between Belle and Belle II.

Among the highlights, Belle II completed the measurement of the branching ratio and CP asymmetry in the $B^0 \to \pi^0\pi^0$ decay using $365\,\text{fb}^{-1}$ of data. The committee commends the collaboration for achieving the uncertainties comparable to those of the Belle and BaBar experiments, with a smaller integrated luminosity. Additionally, Belle II has presented a new search for the decay $B^0 \to K^{*0}\tau^+\tau^-$: an important result that constrains some extensions of the SM addressing the $R_{D^{(*)}}$ anomalies. Belle II has set the most stringent constraint on the branching ratio for this process, surpassing the Belle result by nearly a factor of two with less than half of the Belle dataset. The Belle collaboration continues to advance in hadronic physics, recently producing evidence of observation of the pentaquark state $J/\psi\Lambda$ at 4.459 GeV produced from $\Upsilon(1S, 2S)$ decay.

Looking ahead, Belle II has an ambitious agenda of publishing 40–50 papers per year. Several new analyses are in the pipeline and are expected to be completed soon. Particularly interesting is the measurement of the $B \to \tau\nu$ branching fraction, which is a golden channel for precise SM tests and is highly sensitive to New Physics (NP). Also a long-awaited result of the search for an invisible dark photon is anticipated by next summer. By early next year, the first results are expected from searches for inelastic dark matter, as well as two new axion-like particle (ALP) searches and an update on the light invisible Z' search. The collaboration also plans to finalize the measurement of $\sigma(e^+e^- \to \pi\pi)$ in 2025.

The collaboration will maintain the Physics Week meetings in the new format introduced last year, focusing each meeting on a particular type of physics of interest. This year's focus will be on τ and dark sector physics.

4.2 Concerns

- The physics results will be influenced by the damage to, or absence of, the PXD detector. Fortunately, areas such as tau physics, searches for dark sector particles, and analyses involving missing energy remain unaffected. And, the primary analyses impacted—specifically, time-dependent analyses—will require significantly more data to be competitive.

- The analyses that explore CP quantum coherence of the initial state require large statistics, which is not available until higher luminosity is achieved.

4.3 Recommendations

- The committee recommends to give high priority to clean decay modes with missing energy, in particular $B \to K^* \nu\bar{\nu}$, where the Belle-II collaboration has demonstrated to be able to achieve competitive results, of large impact, despite the limited statistics.

- The committee looks forward to seeing many new dark sector and low multiplicity analyses finalized in the coming months and years.