B-factory Programme Advisory Committee
Full Report for
Annual Review Meeting at KEK

19-21 February 2023

A. Andreazza* (Milano), P. Collins* (CERN), G. Corti (CERN),
M. Demarteau (ORNL), R. Forty (CERN), B. Gavela (Madrid),
S. Gori (UCSC), W. Hulsbergen*& (NIKHEF), M. Ishino& (Tokyo),
V. Luth (SLAC), P. Mcbride*& (FNAL) P. Mato*& (CERN),
F. Meijers* (CERN), N. Neufeld& (CERN), B. Ratcliff* (SLAC),
A. Petrov* (Univ. South Carolina), M. Sullivan (SLAC),
H. Tajima& (Nagoya), M. Titov* (Saclay)
and chaired by T. Nakada (EPFL)
& Remote participation
* Expert member.

22 May 2023

1 Short summary

The annual review for the Japanese Fiscal Year (JFY) 2022 by the B-factory Programme Advisory Committee took place from the 19th to the 21st of February 2023 at KEK. This report gives brief responses to the questions presented by the director of the Institute of Particle and Nuclear Studies, including some additional remarks. It was refreshing to have a face-to-face discussion after three years of remote meetings due to the COVID restrictions.

-General: Since this was an annual review of the overall project, reviewers were asked to give general advice to further promote the Belle and Belle II projects. Especially, the plan to join the analysis efforts of the two collaborations is to be evaluated.

It is impressive to note that the Belle collaboration continues to publish papers in high-quality peer-reviewed journals after they completed data taking in 2010. Although the statistics of the Belle II data are still about half of the sample collected by Belle, the Belle II analyses start to produce results competitive with those of Belle, benefiting from the superior detector. It was shown that analysis strategies developed for Belle II data could be applied to the Belle data to improve the results. Furthermore, Belle data have been ported to a format that could be analysed by the Belle II analysis software. Therefore, the committee strongly urges the Belle and Belle II collaborations to merge the two efforts by unifying the analysis plan and development, as well as the review process. There are several possibilities to implement this goal, and it should be left to the scientists to work out a solution. Since the two collaborations are largely overlapping, the committee is confident that a solution acceptable to both parties will be found very soon.

-Progress of LS1 works and Plan: Review the progress of the work to be completed in LS1 on the various detector systems: PXD2, TOP PMT, KLM RPC, CDC, and the online/offline monitoring and control. Point out if there are other possible improvements to be made on this occasion.

-Plan to resume machine operation after LS1: Is the run plan for JFY2023 appropriate to ensure improvements of the machine and detector while securing reasonable operation time?

The two halves of the PXD2 have now been assembled and should be transported to KEK at the earliest possible time. TOP PMTs from one slot have been removed, and laboratory tests confirmed the suspected degradation of the quantum efficiency. A firm plan for the PMT replacement has been accepted. The action being taken now on the KLM-RPC monitoring should avoid irregularities in the detector performance passing unnoticed. Any attempt to recover the chamber performance by changing the gas mixture should only be tried after thorough laboratory tests with spare chambers. The long-term instability of the CDC operation remains a concern since the cause of the gain drift has not been understood. Careful monitoring of the gas properties, water and oxygen content, and the laboratory irradiation test should be reinforced. The systematic consolidation of the HV control and monitoring of multiple subsystems is highly appreciated, and the committee recommends that this effort be extended to other detector control and monitoring processes, including the interlocks and alarm systems.

It was reported that the MEXT budget plan for KEK would allow three months of machine operation during JFY2023. If LS1 can be completed by the end of the calendar year 2023, this will give a sizeable block of running time before the 2024 summer shutdown. This should allow time for commissioning the machine and detector and collecting data for the summer conferences, judging from the luminosity projection. The Belle II LS1 schedule, which is basically driven by PXD2, is currently not quite compatible with this goal. Thus the Belle II collaboration is strongly encouraged to examine options that could allow the machine operation to start at the beginning of 2024 by carefully assessing the risks. For example, the extraction of the VXD could begin already after the initial mechanical and electrical integrity assessments of the PXD2 while further tests of PXD2 would proceed in parallel. The PXD2 group should be prepared for some repairs, including the replacement of ladders if necessary at KEK, and spare ladders should be ready.

-Analysis: Are the planned analyses for the winter conferences ready? Is there anything to be streamlined to improve the output?

The presented analysis plan is good and advancing well. As already pointed out, combined analysis of Belle and Belle II data will enhance the precision of the results and should be pursued. In order to obtain new world-leading results, further strengthening and broadening of the analysis efforts for the dark sector physics is encouraged.

-Computing [to remind] the adequacy of the resource estimate is to be reported by April/May.

A preliminary estimate of the computing resource requirement was shown at this meeting. A dedicated discussion and recommendation will be given once the final document, prepared by the expert group consisting of G. Carlino, W. Hulsbergen, P. Mcbride and P. Mato, and chaired by the BPAC chair, becomes available.

In addition to these responses, there are some comments on the LS2 upgrades. It is understood that some upgrade of SuperKEKB will be required during LS2 to achieve the Belle II goal, i.e. collecting 50 ab^{-1} of data, within a reasonable timescale. This, in turn, calls for upgrades of the detector to fully exploit the upgraded machine capability. The committee was presented with the current ideas for upgrades of various detector systems. Their scope will be strongly influenced by the way the machine upgrade will be implemented, in particular the design of the intersection region. Since the machine group is still developing ideas for the full scope of the upgrade, it is premature to start the actual detailed design of the detector components. The committee recommends the group to remain open for exploring further promising technologies. This would allow the Belle II collaboration to proceed rapidly in defining the scope and definite design of an optimal detector once the extent of the machine upgrade has been fixed.

The committee fully appreciates the effort by the KEK management and the positive response by MEXT to secure funding for adequate running time for the Belle II experiment. However, uncertainty in future data taking due to the increase in electricity cost remains a serious concern.

2 SuperKEKB and injection complex

2.1 Status

Although the SuperKEKB accelerator now has the world record for colliding beams luminosity, many improvements are still needed in order to improve machine performance up to the design luminosity of $4-6\times10^{35}\,\mathrm{cm^{-2}s^{-1}}$. This is a well recognised issue and several things have been identified that will improve SuperKEKB and the injection linac performance. Further improving the injector efficiency into the two SuperKEKB storage rings is a very important part of the improvement plan. Very high injection efficiency with very low backgrounds in the detector will be needed in order to achieve the design high beam currents in the rings. To name a few linac improvements, hardware being built to improve the optics in the J-arc, to make the linac beam tuning more robust and to improve the air conditioning. Unfortunately, the injection complex does not have a long shutdown and is currently supplying beams to the synchrotron rings. Nevertheless, several items for injector improvement have been identified and attempts are being made to either modify or install new hardware when the injector complex becomes available (this is usually during the summer downtimes). The collider storage rings are benefiting from the Long Shutdown and several upgrades are presently being installed. In particular, a larger beam pipe for the HER at the injection point is being installed, a non-linear collimator is being installed in the LER and an improved (more robust) collimator head is being installed in the LER.

The new interaction region beam pipe has been completed and is ready for the arrival of the PXD2. This beam pipe has several improvements that should improve the background levels from synchrotron radiation. In addition, the beam pipe aperture is slightly larger (this is at the Y section of the beam pipe) which may improve injection backgrounds.

2.2 Concern

- The short available time to improve the injector linac complex due to the synchrotron light rings running schedule.

2.3 Recommendation

- A consideration could be made to review the running schedule for the synchrotron light sources to allow for more upgrade time for the linac complex.

3 Machine detector interface

3.1 Status

Beam backgrounds in the Belle II detector have been acceptable partly due to the relaxed limit for the TOP background rate. Currently, the LER beam-gas background is the dominant component at a level of 40% followed by the LER Touschek background 20% and the luminosity background 10%. Although the measured background rates are close to or below the simulated rates for most background components, the measured LER beam-gas and Touschek background rates are three times higher than the simulated rates. The latest background estimates up to $L \sim 3\times10^{35}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ (before LS2) based on the data/MC ratio are acceptable for the Belle II operation. However, the background rates at the target luminosity are expected to be an issue without further reduction. New background shields will be installed around the bellows during LS1 in order to reduce shower leakage from beam-gas backgrounds. Collimator upgrades including the introduction of a non-linear collimator and more robust collimator heads are also planned to reduce beam related backgrounds. Neutron shields will also be installed around the QCS magnet to reduce neutron backgrounds which is a major cause of the DAQ downtime due to FPGA single event upsets (SEU)s. In addition, four dedicated FPGA SEU counters, four thermal neutron detectors (^{3}He tubes and/or LiI scintillators) and six BEAST TPC directional fast neutron detectors will be installed to better understand the neutron backgrounds, in order to improve the modelling of neutron background and to plan further mitigations.

Beam injection backgrounds are another major cause of the DAQ downtime due to long trigger veto window. Solutions for horizontal emittance growth and longitudinal acceptance issues have been found and some of them will be implemented during LS1. Causes of vertical emittance growth are not understood yet.

Understanding and mitigation of the sudden beam losses are critical for safe operations of both Belle II and SuperKEKB, since higher beam currents may cause catastrophic damages to the vertex detectors and the QCS magnet. Extensive efforts have been made to localise the origin of the beam loss by installing fast beam loss monitors at various points. Some candidate causes are identified, but more investigations are required for further understanding. Faster beam aborts are also implemented to minimise the damages due to sudden beam losses.

3.2 Concerns

- The measured LER beam-gas and Touschek background rates are three times higher than the simulated rates and the dominant factor for the TOP degradation.
- It is not clear how much background can be reduced by the planned collimator and injector upgrades.

3.3 Recommendations

- Further efforts should be made to understand and reduce the LER beam-gas and Touschek backgrounds.
- A long-term strategy to refurbish and upgrade the injector complex should be developed to reduce injection backgrounds as well as improve injection stability.

4 Belle II detector

4.1 Vertex detector (VXD)

4.1.1 Pixel detector (PXD)

4.1.1.1 Status

A dedicated BPAC review of the PXD was held in October 2022 triggered by the observation of two serious issues after the two half-shells of the PXD2 had been completed at MPI Munich and transported to DESY for electrical tests. After initial tests of the first half-shell, it was found that two Layer 1 ladders had a visible kink at the glue joint and were considered “broken”. This pointed to two major issues: improper functioning of the ladder gliding mechanism and insufficient rigidity of the adhesive joint between modules. Full details can be found in the October report. The committee made six recommendations as a result of that review. One of the key recommendations was to have an internal review of the PXD2 by a small task force to ensure that the procedures for remedial action and installation are safe and meaningful. It was suggested that more than half the members of this task force should be non-PXD members so that actions could be evaluated from multiple viewpoints. The collaboration assembled an internal review committee charged with assessing if the root cause of the issues observed was fully understood and if the proposed corrective actions are adequate and sufficient. A series of tests with dummy ladders was performed. The cooling performance was studied as function of the applied torque on the gliding mechanism and no strong dependence was observed in the relevant torque range. The ladder gliding mechanism was verified at room temperature with different foil configurations and different torque values. With an additional foil underneath the washer, even under extreme conditions, the gliding mechanism performed nicely. Tests of a dummy half-shell by heating and cooling the dummy beam pipe allowed verification of the Support Cooling Block (SBC) and the ladder gliding mechanism. These tests led to the conclusion to leave the forward SCB screws fixed and rely primarily on ladder gliding to compensate for thermal effects. Some modifications and repairs were made to the half-shells, but the final changes will be implemented at KEK based on continued studies at KEK. Because of the differences of the DESY setup compared to the final configuration, it is difficult to draw firm conclusions and the performance will have to be verified with the real beam pipe at KEK at the foreseen operating temperature of $-20°$C. At the time of the review the PXD was being prepared for shipment to KEK.

A detailed choreography of testing, extraction of the VXD and re-installation of the VXD with the new PXD has been developed. The plan calls for approximately five weeks of testing to ensure that the PXD2 is working well before the extraction of the current VXD system. The LS1 schedule of Belle II is essentially driven by the validation and subsequent installation of the PXD2. A detailed schedule has been developed indicating that it will take about 20 weeks to install the PXD2 and an additional three months to install the endcap, insert the QCS and complete the concrete shield. That results in a possible restart of operations after mid-January 2024. It was reported that the MEXT budget plan for KEK will allow three months of running during JFY23. This puts the PXD2 schedule in tension with maximising the running period before the end of the fiscal year. If LS1 can be completed by the end of the calendar year 2023, this will allow three months of running in JFY23 and give a sizeable block of running time before the 2024 summer shutdown. The experiment has considered the option to continue running with the PXD1 detector.

4.1.1.2 Concerns

- If the installation of the PXD2 must be deferred to a later shutdown, this may result in a permanent delay with the detector never being installed, given the imminent loss of personnel due to retirements.
- The programmatic value of the PXD2 is very high, since it will extend the capabilities of the current detector. At the same time, the PXD1 is taking data satisfactorily and provides good physics results. It would be a mistake to replace

the PXD1 with PXD2 that is not adequately verified and could not match the performance of the PXD1. The overarching goal, however, should remain to install the PXD2 this shutdown.

- It is noted that the spare ladder count is almost inadequate to recover from further accidents in assembly and testing and the team should proceed with utmost caution.

4.1.1.3 Recommendations

- The Belle II collaboration is strongly encouraged to examine options that could allow machine operation to start at the beginning of 2024 by carefully assessing the risks. For example, the extraction of the VXD could begin already after the initial mechanical and electrical integrity assessments of the PXD2, while further tests of the PXD2 would proceed in parallel.

- Verify thoroughly the validity of any modifications to the PXD2 detector, evaluate the changes at the system level, and study in-depth the interplay between the changes.

- Establish a safe storage location for the PXD1 to ensure it can be used as a fallback detector for the duration of the experiment.

- The PXD2 group should be prepared for some repairs, including the replacement of ladders if necessary at KEK, and spare ladders should be made ready.

- Develop a detailed day-by-day schedule for all tasks to ensure that there are absolutely no gaps in coverage, which will also ensure that all the resources needed are scheduled. These needs should be communicated to the collaboration early so that activities and travel can be planned well-ahead of time.

4.1.2 Silicon strip vertex detector (SVD)

4.1.2.1 Status

During LS1, the SVD activity was concentrated on the preparation for the VXD deinstallation and dismounting, necessary for the installation of the full two layer PXD2, and in firmware and track reconstruction improvements needed to cope with the higher luminosities expected after LS1.

The cables and dockbox connections of the whole VXD have already been removed to provide access for the TOP and CDC interventions, and to allow for a prompt extraction of VXD, as soon as the decision of PXD2 installation is taken. This decision is now pending, waiting for the delivery of PXD2 at KEK and its tests. The timeline of the intervention has been presented. In total, 23 weeks are needed for VXD extraction, SVD disassembly and reassembly around the new beam pipe with PXD2, and reinstallation. These include 3 weeks of SVD commissioning after the disassembly and additional preparation time, due to the moving of the power supply and DAQ to B4.

The SVD disassembly and reassembly carry a significant risk on a detector that is vital for the Belle II operation and it is performing well, therefore high attention has been put by the collaboration in reviewing, documenting and practising these operations. This task has also been strongly supported by the BPAC. During the Committee visit to the Tsukuba Hall, the infrastructures for the (dis)assembly operations have been shown. Detailed manuals, illustrating the procedures to be executed, were available next to each working table. The committee expresses its appreciation for the way in which this preparatory work has been carried out.

In parallel, the work on the mitigation measures needed to keep high performance at the high luminosities expected after LS1 has continued. The firmware has been modified to operate with three-fold sampling, instead of the 6-fold sampling used before LS1. This will reduce the SVD dead-time at 30 kHz trigger rate from 3% to 0.6%. The hit timing cut of 50 ns will also be used to reduce the background contamination during the reconstruction.

4.1.2.2 Concern

- Despite the amount and quality of the preparation work, the SVD disassembly and reinstallation remains a risky operation and needs to be carefully carried out. At the same time, this is the driving factor in the definition of the end of LS1 and in the restarting of the Belle II physics running, and that puts pressure on the detector groups.

4.1.2.3 Recommendations

- The installation schedule of the VXD should be optimised, but the priority is to avoid the injection of additional risks due to time pressure. As far as SVD is concerned, that might be achieved by concentrating the experts on the critical tasks that carry high risks and must not be compressed, and by involving additional human resources to accelerate low-risk tasks like the cabling and the moving of components.

- The decision of SVD dismounting could be anticipated based on the verification of the integrity of PXD2, so that the SVD commissioning can proceed in parallel with the PXD2 mounting and calibration.

4.2 Central drift chamber (CDC)

4.2.1 Status

Several studies to improve CDC gas distribution system have been initiated during the shutdown, such as a better understanding of water and oxygen content, possibility to increase the number of gas ports at FWD (outlet) for more uniform flow distribution, and re-design of the gas distribution system. Improvement of water content monitoring is particularly important for CDC operation; 2000 (1300) ppm of water were added since 2019 (2021). However, there was a discrepancy between the measured value and estimated value based on water consumption in the bubbler. The latter was nearly an order of magnitude higher than the measured value. A dew point sensor was installed in the gas circulation line to compare results with a default water monitor and also revealed a discrepancy of one order of magnitude. Therefore, the water monitor used during physics runs has been removed and sent to the company for investigation and re-calibration, and several other water monitors have been installed in different parts of CDC gas system for additional cross-checks.

A large gas leak occurred in the middle of September 2022 and track trigger rate drastically dropped due to the contamination of oxygen from the air. The relative pressure could not be kept positive even after the fresh gas flow rate was increased. The found leak point was a flexible gas pipe at the CDC outlet. Vibration by a gas pump might have caused the flexible tube being rubbed by a touching valve. The plan is to suppress vibration by putting a bumper and mat between the pump, floor and structure. The oxygen monitor did not work properly at the time of the leak and the very large contamination from the air inside CDC has been accumulated and the gain rapidly decreased.

The replacement of resistors on HV dividers from $120\,\mathrm{k}\Omega$ to $1\,\mathrm{k}\Omega$, to decrease the gain drop due to the voltage drop at high CDC currents, was finished in January 2023. No change in performance has been confirmed by studying pedestal data after the resistor replacement. Study with cosmic ray data has to wait until improvement in gas mixture conditions to decrease the oxygen content.

No progress on the laboratory ageing test has been made since the last BPAC meeting. The CDC team are focusing on the tasks which are important while detector is open.

4.2.2 Concerns

- Recent studies revealed that both water and oxygen monitors did not function properly. Precise monitoring of water and oxygen content inside CDC is mandatory for stable operation;
- A slow drift of the gas gain ($\sim 5\%$ gain drift over three months was observed in 2022ab cosmic data without beam) represents one of the major concerns for the CDC long-term operation;

4.2.3 Recommendations

- Water and oxygen monitors have to be re-calibrated on a regular basis to avoid time-dependent drifts. A possibility to install a water filter (silica gel tubes) to reduce water content inside CDC has to be considered, as the water level might be 5-10 times higher than anticipated during physics running.
- Gas condition monitoring for absolute/relative pressure, water/oxygen contents and circulation gas flow rate needs to be included into alarm monitoring system to notify sudden changes.

- LS1 is the right time to consolidate the CDC gas system, i.e. improve gas tubing and filters, better understand flow dynamics inside the chamber, and eventually to modify the system and to install additional gas ports in order to increase the CDC gas flow rate.
- Further studies are required to identify the nature of a slow drift of the gas gain in cosmic ray runs, taking into account that water and oxygen monitors did not work properly.
- The trial to reproduce the gain drop with the laboratory test chamber ageing setup should be pursued. On a longer term, this might facilitate development and test of different remediation strategies for gain recovery.

4.3 Particle identification system

4.3.1 Barrel system (TOP)

4.3.1.1 Status

The quantum efficiency (QE) of the TOP photon detectors has been monitored during the running periods using dimuon events, and larger than expected degradation has been seen, as reported at previous meetings. For the conventional MCP-PMTs (i.e. those without ALD treatment) the ratio of QE at the end of 2022 compared to the beginning of 2021 has an average of 0.84, indicating a 16% loss of efficiency. A good correlation of these results has been seen with the analysis of cosmic ray data. Studying the QE variation across the detector as a function of the angle of photon incidence has revealed that some of the optical connections between the PMT and radiator plate have degraded, ascribed to detachment of the optical filter from the PMTs; this has led to areas with about 10% lower detection efficiency, but independent of the running period, i.e. unrelated to the QE degradation. The effect of increased temperature on the QE has been studied in the laboratory, exposing PMTs to light while raising their temperature to 40°C in a thermostatic chamber. While no difference in aging has been observed for the life-extended PMTs, an indication of faster degradation has been seen for the conventional PMTs. While this change does not explain the total loss seen in the detector, it is being considered to reduce the temperature in the detector environment, e.g. by adjusting the chiller flow rate or temperature.

The major development since the last meeting is that the PMTs from one of the sectors of the detector (Slot 16, corresponding to 1/16 of the full detector) have been removed and tested in the laboratory. This delicate procedure involved disconnecting some CDC cables and dismounting the corresponding CDC front-end, opening the TOP detector window, and dismounting the front-end electronics and 8 PMT modules. The QE drop measured in the laboratory is in agreement with the one observed with the monitoring during data taking, confirming that it is a real loss of efficiency. This has reinforced the importance of replacing the conventional PMTs with life-extended PMTs during LS1, so that the detector can run successfully for years afterwards at the expected background rate of 5 MHz/PMT. The plan for how this replacement will be scheduled has now been chosen, involving a major rearrangement of PMTs from most slots of the detector, so that those in the upper half will be all life-extended – since they will be less easily accessible between long shutdowns. The lower half will be filled with the best PMTs from Slots 3–16 along with the remaining new PMTs: it will be possible to replace these PMTs in future summer shutdowns if needed, and this will minimise the future risk from possible degradation of QE for the ALD-coated PMTs. At the same time, any failing electronics channels will be replaced. Training and testing for this major intervention has been carried out with spare modules, removing and re-installing PMTs within several hours with improved tooling. The schedule of the intervention has been updated following the experience gained with Slot 16, and has been extended a little to allow for steady, careful work over about 10 weeks in total. The work on the front-end, cabling and DAQ testing will be accelerated with the use of three teams. Tests will be made during and after re-installation slot-by-slot, with CCD camera images, electronics checks, laser data taking, etc.

The committee congratulates the TOP team on the careful, systematic studies that have been made to understand the problem of the loss of efficiency in the photon detectors, and in the preparation for their rearrangement or replacement. The choice of the more extensive PMT replacement plan is well justified by the long-term perspective of maintaining excellent performance while operating the detector for many more years.

4.3.1.2 Concern

- The rearrangement and replacement of PMTs across almost the entire detector is a major operation, with associated risks. The highly skilled colleagues who will be responsible for performing the bulk of the work are few in number, and should not be overloaded.

4.3.1.3 Recommendations

- Perform a detailed risk assessment of the PMT replacement program, and ensure that the scheduling maintains contingency to allow for adaptation if any unforeseen problems arise. Extra personnel should be trained if necessary to avoid single points of failure.
- Efforts to understand and mitigate the effects of temperature on the PMT efficiencies for tubes in the detector should continue.

4.3.2 Endcap system (ARICH)

4.3.2.1 Status

ARICH is a stable detector that has generally run well during all data runs to date. There was no specific ARICH detector presentation during this review. Approximately 6.6% of the detector channels were disabled by the run's end. The loss rate for channels has remained reasonable, with the number of masked channels increasing by about 10% during the last year. The ARICH DAQ was upgraded to PCle40 early in run 2022a.

DAQ operation was stable, the number of ARICH problems was reduced, and no major downtime of greater than 30 minutes was due to ARICH. DAQ errors increase in frequency at higher beam backgrounds, thought to be due to SEUs. More automated recovery procedures are needed and are being developed for future operations at high backgrounds.

No hardware intervention inside the ARICH detector is scheduled during LS1 as it will remain attached to the ECL. Some LV cable connections were repaired in August 2022 after the end yokes were opened allowing the recovery of 5 FE modules. Also the chiller system was replaced with a new one, with a backup chiller.

4.3.2.2 Concern

- Continuous recovery from DAQ errors will be crucial for efficient data taking at higher backgrounds.

4.3.2.3 Recommendation

- It would be helpful to further evaluate the performance both for data collection and for physics as a function of backgrounds.

4.4 Electromagnetic calorimeter (ECL)

4.4.1 Status

The ECL is a mature, stable detector that is generally running well during physics runs to date. There was no ECL specific presentation at this review. The main concerns at present are understanding the injection backgrounds, and determining pedestals. Further work is ongoing to improve monitoring as occupancies increase, and to define the appropriate actions to be undertaken by shifters if running conditions deteriorate. Continuing work on analysis level corrections is discussed elsewhere in this report.

No hardware intervention is planned during LS1. Due to the beam background leading to a larger pile-up, there might be a degradation of the energy resolution and of the determination of the photon direction. A study is underway during LS1 to understand whether the background could be effectively subtracted in firmware.

4.4.2 Concern

- As backgrounds grow, the ECL performance may deteriorate with a strong dependence on crystal location.

4.5 K-long Muon detector (KLM)

4.5.1 Status

The low efficiency of RPC BB2 was noticed in April 2022 when a set of monitoring plots of the RPC efficiency for each layer was introduced. Looking back at the history, the trend began in June 2021. The original cause was the interruption of the gas circulation. After the recovery of the gas circulation at the end of April 2022, the efficiency of RPC BB2 was partially recovered, however, its efficiency dropped again. Another cosmic run was taken in February 2023, and the efficiency stayed as low as the one in June 2022.

It is realised that water vapour entering through the polyethylene tube of the gas inlet caused the problem. To recover the operational condition, the standard operating gas was circulated with a higher flow rate. The contamination was reduced in the first month and saturated before the middle of June 2022, but this did not cure the low-efficiency problem of the RPCs.

The degradation of the physics performance due to the low-efficiency problem was evaluated. The muon identification efficiency in this particular region, BB2 and BF2 corresponding to 1/8 of the barrel system, is lower by 10%, while there is no visible impact on the K_{L} identification. This effect has been taken into account to the run-dependent simulation data and the physics analyses that are performed.

There is a plan to recover the efficiency with a gas mixture containing ammonium. The KLM group is preparing a test-stand at the KEK Tsukuba hall to test the recovery procedure using test RPCs before applying to the real chambers. The first trial to reproduce damage on the RPCs by introducing water vapour was reported.

The KLM group is reinforcing slow monitoring and data monitoring to find similar problems as quickly as possible. The gas flow of each inlet and outlet of individual lines will be monitored in a total of 832 channels. The new system will be ready before the end of LS1. The hit rate and the efficiency of each layer have been monitored in the framework of Data Quality Monitoring (DQM). The plots to show trends of relevant parameters are available and an automatic alert system is in place to spot unusual behaviour on these parameters. In addition to these, the infrastructure will be improved. A part of the polyethylene tubes for the BB2 RPC will be replaced with copper tubes. Water sensors for the chimney of the solenoid magnet and additional drain tubes will be installed.

4.5.2 Concerns

- Methods to recover damaged RPCs were studied only a long time ago. It is important now to follow the same method and to check all procedures with the experts of Osaka Metropolitan University.
- Ammonium gas is corrosive to copper and should be treated carefully. The committee considers that introducing ammonium in the gas mixture to recover the damaged chambers appears to be a highly risky operation.

4.5.3 Recommendations

- As proposed by the KLM team, it is important to check the recovery procedure, such as the addition of Ammonium gas, at the test bench before applying the method to the Belle RPCs.

- It may be worth to apply a different algorithm to identify muons in the problematic region to mitigate the efficiency loss of 10% but still retaining a good S/N ratio, e.g. by reducing the number of hit layers.

4.6 Trigger and data acquisition system (DAQ)

4.6.1 Status

During the 2022ab run, the trigger had to cope with a decreasing efficiency of the CDC track trigger, caused by the gain drop in the CDC. The overall L1-trigger rate reached 10 kHz, which, if extrapolated to post-LS1 conditions, would exceed the available capacity of the downstream DAQ. The LS1 work by the trigger group is focusing on reducing the trigger rate, recovering the CDC track-trigger efficiency, taking advantage of the large improvement of the CDC trigger primitive quality and reducing the global inefficiency due to the injection veto, the latter accounting for 5–15% of DAQ dead-time.

One of the main activities of the DAQ team in LS1 was the completion of the transition of the front-end DAQ from the COPPER system to the PCIe40 boards. During the 2022ab run TOP, KLM and ARICH have used the PCIe40 without any major problems. Several other improvements for the efficient operation of the PCIe40 have been implemented, these have been well received by the experts. Cosmic ray tests have been conducted successfully with all systems in LS1. Synthetic high-rate tests using a dummy trigger also showed no problem and achieved up to 27 kHz. Some small issues remain but work-arounds are in place for most of them. A test with more realistic event sizes, achieved by lowering thresholds in the front-end electronics, is planned once the detector is completely assembled again.

Several enhancements and improvements are being tested: sub-event building in the hosting PC, improving the data throughput from the PCIe40 to the host, and using the SFP+ ports for an optical connection of the *b2tt* link.

A new version of the FTSW module is being prepared which can accommodate many optical links. This should significantly improve stability as it removes the electrical connections over Cat7 cables, which have been affected by noise at longer than one meter distances. This noise can lead to link-lost condition, which costs a non-negligible time to recover.

On the back-end, DAQ side several new worker nodes, control and storage units are being added, which should bring the sustainable rate up to 20 kHz by the 2023c run. This rate increase is also possible because of optimisation work done in the HLT reconstruction software.

The output of the DAQ will be switched from SROOT files to standard ROOT files. This saves downstream computing (in particular also I/O) resources. The necessary software changes have been validated on the existing hardware and will be used for the next physics run.

The committee congratulates the DAQ team for the impressive amount of work and improvements achieved.

4.6.2 Concerns

- The personnel situation in several trigger subsystems is very worrying, in particular for the KLM and NeuralNet 3D CDC-triggers.
- The new injection veto algorithm which combines timing from the accelerator with detector activity will affect data quality. The effects are not yet clearly known.
- The L1 trigger has implemented major changes in the hardware and firmware, in particular the CDC chain. The margin in the latency budget appears to be very small ($\sim 0.2\,\mathrm{ns}$).

4.6.3 Recommendations

- The PCIe40 upgrade is proceeding well. As stated in previous review, firmware exports of the sub-detector systems where the PCIe40 migration was done just after the end of the 2022ab physics run, should be available at the start of the next physics run, in case unforeseen problems show up.
- Long-term institutional commitment should be found for the trigger systems at risk of becoming unmaintained.
- Effort to understand if the proposed improved injection veto is a viable option should be increased, so that in case an alternative approach is needed, work on this could start as soon as possible.
- The excellent efforts to improve automatic recovery and increase automation should be continued. They will be very beneficial for increased data-taking efficiency.
- Although not yet critical, the implementation of sub-event building in the PC will improve throughput. This should be thoroughly tested before deployment in the next physics run.

4.7 Slow-control, monitoring, and alarm handling

4.7.1 Status

The detector slow-control, monitoring and alarm handling are provided by the hardware-oriented Detector Control System (DCS) and the performance-oriented Data Quality Monitoring (DQM). Both provide monitoring and alarms (alerts) to the shift crews in the control room as well as remote. The DCS is also the principal system for experiment control.

As a response to past data taking events, when detector issues of different severity were not spotted by the shift crew in the (virtual) control room, a review and consolidation process has started under the auspices of the technical coordination and run coordination. A prioritised plan for review, design, and implementation in 2023 has been established.

The main goals for the DCS are to satisfy new requirements, unify the behaviour of the processes across the sub-systems and to review and improve the monitoring system and alarm handling. The areas addressed, so far, are the HV (high voltage) control, including the injection inhibit system, and the detector environment monitoring. Concerning the HV control, a review took place in January summarising the current system and discussing the requirements of the new system. This is being documented and will be provided to the reviewers for further iterations as well as serving as the official document for future reference. The plan is to have the new system in place for the summer and test the real setup with cosmic runs. Concerning the detector environment monitoring, the items from each sub-system have been reviewed as a first iteration. The next step is to compile it into a hierarchical structure and implement an alarm system. The plan is to have the implementation finished by autumn, leaving some contingency for it to be operational before 2024.

The main goal for the DQM is to reorganise the main control room DQM panels in order to provide the shifter with unambiguous information of when, and in what form, action is needed. Two main panels, one for monitoring the ongoing run and one for monitoring the longer-term (run-dependent) evolution, are foreseen. A review of the DQM content has been done and a revision of the content was requested from all detector groups, at the same time making the approach more uniform across sub-detectors. The technical implementation by the central DQM group and all sub-detector groups is foreseen until end-summer. Full-scale tests can commence prior to the start of data taking using playback of existing raw data.

4.7.2 Concern

- A schedule for the review, implementation and test of the DCS has been established for 2023. This plan appears ambitious, given the limited time for implementation and the need for thorough testing. In addition to the two main tasks of HV control and detector environment, interlocks and detector settings are to be completed in parallel. It is not clear to what degree these tasks compete for resources from the central and/or sub-detector teams.

4.7.3 Recommendations

- Most of the recommendations of the last review remain valid until the programme is completed.

- There is considerable momentum in the consolidation effort. The upgrades of the detector control monitoring should continue to be given a high priority in order to achieve a well-tested new system in advance of the next data-taking period.

- The development of common components for the control and alarm handling would benefit from engaging dedicated experts.

- A peer review process of the implementation would be beneficial, engaging people from sub-detector groups as well as central experts.

- Documentation and consolidation of DCS and DQM components should continue beyond LS1. Major system revisions and upgrades can be installed during future shutdowns after successful tests.

4.8 Detector performance

4.8.1 Status

During the final months of operation before the LS1 shutdown the beam currents were increased to enhance the luminosity. This resulted in damage to some collimators and high beam backgrounds. Newly injected bunches appeared to oscillate and this resulted in significant degradation of the performance of some of the sub-detectors after each injection. These high backgrounds primarily originated during the first 20 ms after the bunch injection. Trigger vetoes have been used to avoid the saturation of the read-out. The data processing uses events with timing after the LER and HER injection to study the detector response as a function of the time difference and suppress the largest backgrounds.

The physics performance group is focusing on comprehensive studies to improve the understanding of the hardware and software, provide data/MC corrections, and uncertainties of the Monte Carlo modelling of the data, to improve the efficiencies and resolutions.

Charged particle tracking and vertexing benefit from the high-precision single-layer PXD and four-layer SVD with stable performance, well reproduced by the run-dependent simulations. The CDC is the primary charged particle tracking device and used to be stable. But it has recently indicated high sensitivity to beam injection background. Recent studies focused on the CDC performance in terms of track reconstruction, resolution, and fake rates, identifying various sources of efficiency losses, totalling 15%, and dependent on particle momenta and angles.

Concerning the performance of particle identification, the effects of the recent high beam backgrounds and the impact of detector problems such as the BB2 issue of the KLM have been studied. A significant drop of the electron ID efficiency by about 20% is seen in the first 10 ms after beam injection, due to degradation of the dE/dx information from the CDC. There is also an apparent slight improvement in the pion-to-electron fake rate in that period, but this is due to the shift of the dE/dx to lower values; the overall hadron separation is largely unchanged, as the TOP and ARICH are relatively tolerant to these backgrounds. However, if only the CDC information is used, very large misidentification rates between pions and kaons are observed, due to the dE/dx shift. Re-calibration of dE/dx as a function of time since the injection is almost ready, with gain and resolution corrections for LER/HER injection time bins produced for each run block.

The electron ID efficiency is seen to decrease somewhat in the later runs of the last data-taking period, when the backgrounds were high. In particular, for $p < 1$ GeV/c the efficiency drops from its usual value of 90% by about 5%. This is ascribed to the dE/dx gain drop in the CDC due to the bad beam background conditions. Similarly, towards the end of the running period the loss of muon ID efficiency is seen from the BB2 issue in the KLM (provoked by the local interruption to the gas circulation). It gives a large effect of order 10% loss when limiting the muon ID to the BB2 region, but over the full acceptance of the detector this effect is less visible, at a few percent only, so there should be small impact on the physics analyses. Nevertheless, this inefficiency should be corrected for future running.

Studies are underway to improve the particle ID performance using machine learning, using neural networks to combine the information from TOP, ARICH, CDC, ECL, KLM, and SVD, as well as exploiting more fully the spatial information from energy deposition in the ECL. The first results look promising for significant performance gains based on simulated data, the performance with real data will now be studied. For physics analysis, the PID group provides efficiency corrections and systematic uncertainties of the efficiency, but this sometimes takes a long time due to a lack of human resources. A framework has now been developed, inspired by the *PIDCalib* framework of LHCb, to compute them automatically; this work should be completed by this summer.

Studies of neutral particles, specifically photon energy and efficiencies, are progressing well. The ECL performance is strongly impacted by the large beam backgrounds and fake photons from cluster splitting. Tight restrictions on photon timing resulted in effective photon background classification by a fast BDT discriminator. Photon and π^0 efficiencies at low energies are dependent on polar angles and have been corrected. K_L^0 classifiers based on ECL and KLM sub-detectors are being developed.

4.8.2 Concerns

- In the future, operation at larger luminosity is expected to require larger beam currents which are likely to cause radiation damage and increased occupancies in the detectors, impacting their performance.

- Automation of the detector performance studies is planned. Currently, particle ID corrections are being tested and plans for similar processes are being prepared for tracking. It remains to be seen how stable the performances are, and to what degree automation can be expected to be feasible. Reinforcing the team with expertise in specific areas will be needed for reprocessing the current data, preparing for future data sets, and also planning for upgrades.

- The loss of particle ID performance is seen for 10 ms after beam injection in runs with high backgrounds. This could limit the performance in the future running at higher luminosity.

4.8.3 Recommendations

- The most recent data were impacted by the large beam injection backgrounds which strongly affected the sub-detector performance. The collimation system should be repaired or upgraded if required, and it is hoped that the machine experts can develop ways to further improve the bunch properties and injection, thereby reducing the detector backgrounds. The future backgrounds affecting the sub-detector performances should be carefully monitored to assess their impact. Re-calibration of the dE/dx performance as a function of time since beam injection should be completed with high priority.

- The efforts and recent results by the physics performance group to further develop the event reconstruction, and determine the data/Monte Carlo corrections and uncertainties are very promising. Efficiencies and resolutions to distinguish signals from backgrounds will need to be further explored using multivariate criteria. While there is good progress in understanding detector performance and algorithms, data/MC corrections with high precision will be required. The implementation of automation for providing PID and charged particle tracking will require the engagement of experts with an in-depth understanding of the requirements.

- Detailed planning and documentation of the complex data processing steps should be implemented, to optimise the data analyses and record variations and problems.

4.9 Computing

4.9.1 Software

4.9.1.1 Status

Work practices for development and management of the Belle II software have become well established. Software validation is also becoming a more routine and automated effort. The *basf2* software is organised in a set of packages covering the whole range of needs of the collaboration from simulation to data taking, to documentation. Each package contains the software specific to a given topic and is managed by one or two librarians. Any contribution to the GitHub software repository has to be approved by at least one of the librarians concerned by the changes. Bi-weekly developers meetings are the forum to ensure overall coherence of new software.

The collaboration has used Atlassian tools to manage the software life cycle for over 7 years. Last year a change of Atlassian's business model, introducing a non-sustainable license cost, prompted DESY, that hosted the tools, to migrate tracking of issues from JIRA to GitLab. A migration strategy was devised and a full scale test carried out at the end of last year. Care has been taken to train the Belle II community on the use of the new tool. The migration was planned for a few weeks after the BPAC meeting.

Software developments continue to be checked with the well-established procedure of unit tests run for each commit contribution and nightly validations with experts checking the distributions produced over night. Validation of a new release branch is carried out with the dedicated candidate build run on a sample of 10k raw data events to check its performance on the on-line system. The upgrade of the Monte Carlo (MC) simulation from Geant4 10.6 to the latest 11.1 required a change in the *basf2* interface and a re-optimisation of the cut on production of delta rays. The use of the new version of Geant4 will be validated following the same steps carried out for the previous change of version, including analysis of high-level variables in a number of representative modes. For the first time tuning of Pythia8 parameters is being carried out with the Professor toolkit[1] showing encouraging results in the agreement of event shape variables between off-resonance data and MC.

Many new features are under development in a variety of software aspects: from improvements in reconstructions (ECL clusters, vertex fit, etc.) to adding new features in data analysis, to exploitation of new Data Quality Monitoring features. Various initiatives are being carried out to further exploit Machine Learning techniques in reconstruction and physics analysis tools, including the use of new Graph Neural Networks (GNN) for Flavour Tagging and offline photon reconstruction.

The committee was pleased to note that a concerted effort has been put in reducing processing time in the reconstruction. This is particularly important for online processing where the L1 trigger rate and background conditions strongly constrains the processing time available for the High Level Trigger (HLT). The most time critical algorithms have been identified and improved. It was shown that with the 2022 software *release-07* combined with an increase of 3 HLT nodes it will be possible to operate with an L1 rate of up to 20 kHz, allowing to record a luminosity of $9 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. Most of the reduction in execution time was obtained with improvements in the tracking software, that is currently responsible for 80% of the reconstruction time. It should be noted that in Belle II tracking is the first step in the reconstruction chain and reconstruction of neutral particles and particle identification rely on good track quality. Effort has also been spent in improving the quality of some type of tracks. Preliminary studies show that restructuring the reconstruction sequence by moving particle identification not used in the HLT decision after the trigger filtering will further reduce by about 10% in the processing time. Careful validation will be carried out to ensure that no hidden dependencies influence the HLT decision. Removing non-essential modules for either HLT or skim production and introducing additional filters between L1 trigger and reconstruction are under consideration. The ability to carry out software optimisation is hindered by the reduced number of software developers available for the task.

4.9.1.2 Concerns

- Tuning with Professor is carried out introducing reference distributions by hand. This is intrinsically more error prone and by not exploiting the use of Rivet plugins the tuning cannot profit from concurrent use of references distributions from other experiments.
- Not enough appropriately qualified personnel is available for software optimisation. The committee was given the impression it is difficult to recruit inside the collaboration.

[1] the de-facto tuning tool in the HEP community.

4.9.1.3 Recommendations

- The use of Rivet plugins as references for the tuning with Professor should be investigated. Plugins with Belle II data could be kept private to the collaboration until publication of the relevant results. Their inclusion in the public Rivet repository would provide valuable input to generators authors and the HEP community at large.
- The considerable effort the software group is spending in algorithmic optimisation is to be commended. Nevertheless a complementary effort should be put in the more technical and general aspects of software optimisation.

4.9.2 Data processing

4.9.2.1 Status

In general the data processing is proceeding smoothly. Over the last months the number of concurrent running productions has increased significantly, with the result that the pledged resources are fully used when the run-independent MC is being produced. The committee is pleased to see that many of the previous concerns on the data production system are being addressed. In particular, to reduce the stress on the Grid due to large amount of small files being produced. Data processing without run boundaries is now partially implemented and is available in production. This means that the output files are fewer and larger, benefiting mainly the scalability of user analysis jobs. However, the run-dependent MC production still needs a separate additional production step for merging the output. The result of the merging campaign was a huge reduction of number of files for old experiments, which had many very short runs. The production of skims is also benefiting from new improved file merging in the production system, which avoids the need of an additional merge step, resulting on the reduction of small files, so improving the efficiency of users jobs. The production of MC15 signal still suffers from the large number of independent productions but is much better than MC14. The need to better streamline and automate the job preparation, submission and bookkeeping remains, mainly because of lack of human resources.

The improvements in the production of the run-dependent Monte Carlo, *MC15rd*, are very visible. It took seven months less than previous production. This is mainly due to better automation, more people joining the team, the possibility to have more concurrent jobs, and better documentation. Nevertheless there is still a list of improvements possible that will make a difference, such as the automation of the production of the payloads (via Airflow), better bookkeeping and, automatic tools to submit new samples on the Grid.

The ongoing development in the core computing is very welcome. Producing directly in the DAQ files in ROOT format, and a better transfer protocol between DAQ and offline, makes the system more scalable to handle the expected increase in data volume after LS1.

A new workflow for data production quality control has been developed to check that no mistakes were made during a production. The system uses *b2luigi* as workflow management control similarly to the analysis validation workflow. Deployment of the system is under discussion.

4.9.2.2 Concern

- The human resources situation is certainly getting better, but in some areas the injection of new resources will make a difference.

4.9.2.3 Recommendation

- The committee encourages the pursuit of the ongoing developments to automate as much as possible the data processing, the production of skims, and the production of both run-independent and run-dependent MC.

5 Physics

5.1 Belle physics

5.1.1 Status

The committee congratulates the Belle Collaboration for continuing to produce world-leading results twelve years after the shutdown of the Belle experiment. Last year the Belle collaboration published a total of 17 papers. Below is a list of highlights:

- Important advances on $|V_{cb}|$ and $B \to D^*$ transitions: Semileptonic $B \to D^* l\nu$ decays of four categories $D^* = D^{*0}$ or D^{*+}, and $l = e, \mu$, have been measured with high precision (clean signatures above modest backgrounds). Exclusive BF's for $b \to c$ transitions, including $B \to D^*\pi(\pi)l\nu$ decays, were obtained. These results have had a substantial impact on the $\{|V_{cb}|, |V_{ub}|\}$ puzzle.[2]

- New constraints on lepton flavour violation (LFV) involving τ leptons: The establishment of the very first limit on $BR(B_s^0 \to e^{\mp}\tau^{\pm})$ of $< 1.4 \times 10^{-4}$ is impressive. New results on $B^+ \to K^+\tau l$ $(l = e, \mu)$ decays have been obtained for both same-sign and opposite-sign lepton pairs.

- $\Upsilon(1S)$ di-lepton decays: New limits have been set for di-lepton decays $\Upsilon(1S) \to l^{\pm}l'^{\mp}$ for $e\mu, \mu\mu, e\tau, \mu\tau$ channels, reducing the PDG limits by a factor of two. The limits on $\tau^+ \to l^+V^0$ transitions have also been improved by $\sim 30\%$.

- Important advances on searches for CPV in D^0 decays to kaons and pions: These recent results are the most precise measurements up to now. They pertain $D^0 \to K_s^0K_s^0\pi^+\pi^-$, $D^+ \to K^+K^-\pi^+\pi^0$ and D_s^+ decays.

[2]Shortly after this meeting, the Belle collaboration announced a very important new results (https://arxiv.org/abs/2303.17309) solving a long-standing discrepancy between exclusive and inclusive measurements of $|V_{ub}|$. The new measurements are compatible and in agreement with SM expectations.

- Studies of B^+ and B_s^0 decays: New measurements exploring purely pionic and pion-ρ channels have for some decays enhanced the sensitivity by orders of magnitude.

- A cusp in the hadron spectrum: Over many years Belle has played a major role in the study of charm baryons and hyperons. Recent studies of charm baryons decays resulted in two new decay modes of the ${\Lambda_c}^+$ baryon: ${\Lambda_c}^+ \to \eta\Lambda(1670)\pi^+$, and $\eta\Sigma(1384)\pi^+$. At the $\Lambda\eta$ threshold, a surprising discontinuity in the derivative of the spectrum function was observed. This cusp is not compatible with a resonance (it does not have the soft shape of a Breit-Wigner resonance), but it is compatible with a re-scattering of final state particles (i.e. $pK^- \to \Lambda\eta$ from ${\Lambda_c}^+ \to \Lambda\eta\pi^+ \to pK^-\pi^+$, with a peak at the pK^- invariant mass). This is the first time ever that such a cusp has been detected in the shape of a hadron spectrum.

5.1.2 Concern

- The slow merging of common activities and analyses with the Belle II collaboration may slow down new physics results stemming from the Belle data set.

5.1.3 Recommendations

- The committee encourages the Belle Collaboration to use their full data set to develop a new analysis that tests the lepton universality in semileptonic B meson decays, i.e. to measure the ratios $R(D^*)$ and $R(D)$, comparing the decay rates involving the heavy lepton τ relative to the rate for electrons and muons. These analyses should overcome the modelling issues which impacted an earlier Belle publication. The results could be combined with Belle II results on these measurements.

- In analogy to their earlier work on $B^0 \to e^\mp\tau^\pm$ decays, a measurement of $B_s^0 \to e^\mp\mu^\pm$ decays would be of interest. Likewise, a study of radiative $\Upsilon(2S)$ decays should be considered.

5.2 Belle II physics

5.2.1 Status

In 2022 the Belle II Collaboration presented 26 new results at conferences: 16 at the ICHEP2022 conference and 10 at the Rencontres de Moriond. This is a much larger number of new results compared to last year, based on substantially larger integrated luminosity: 428/fb compared to 189/fb. At this time, 11 of these results have been submitted for journal publication, two of which have been published. Many more analyses (~ 100) are ongoing. In fact for the 2023 Rencontres de Moriond, 21 new results are expected. In spite of the relatively limited luminosity, the Belle II Collaboration obtained results similar to the corresponding Belle analyses, and in some cases even world-leading results. Few analyses joining Belle and Belle II data have been performed.

The Committee congratulates the Belle II Collaboration for their enriched portfolio of new analyses, many of them based on half the available data sample. In particular,

- Belle and Belle II have been pursuing combined analyses for the branching ratio and CPV asymmetries in $B \to \rho\gamma$ decays, and spectroscopy (the search for the $\Lambda_c(2910)$ resonance). Belle II data would have been statistically limited. The combination with Belle data could result in the world-leading results for the decay $B \to \rho\gamma$ and the $\Lambda_c(2910)$ resonance.

- The update on $B \to K\nu\bar{\nu}$ with both ITA (inclusive tagged analysis) and the more conventional HTA (hadronic tagged analysis) is almost complete. The uncertainty on the branching ratio is comparable to past results from Belle and BaBar.

- The precision measurement of the τ lepton mass is impressive. The analysis based on 50% of the data set led to a world-leading measurement, with an uncertainty that is a factor of ~ 3 better than an earlier BaBar measurement.

- The first Belle II flagship measurement of $R(D^*)$ is in progress to test the lepton flavour universality of the τ versus $l = e, \mu$ channels. In addition, an analysis for $R(X)$ is being pursued. This analysis could result in the most precise test of lepton-universality of this kind.

- One new dark sector analysis was completed: the search for a long-lived particle, X, produced in $B \to KX$ decays and decaying to visible tracks. This analysis produced world-leading results. The results of the search were presented model independent as a function of the mass and lifetime of the long-lived particle.

In general, the Belle II Collaboration has been improving the quality of the analyses.

- There have been substantial improvements in flavour tagging, surpassing the previous 31.7% level in Belle or BaBar and reaching a $36.8 - 38.4\%$ level.

- The decision to base most of the publications on the full data set and run-dependent Monte Carlo simulation will increase the quality, consistency, and precision of the published results.

5.2.2 Concerns

- Of the 26 analyses being prepared for publication, only four use the run-dependent Monte Carlo simulation and a third of them are based on half of the available data.

- The Committee regrets the limited advances in dark sector physics analyses. For example, the Collaboration has not yet pursued the promised mono-photon search for an invisible dark photon. This will be a very important analysis for which the Belle II has unique capabilities. Past simulations showed that even an analysis of a data set as small as 20 fb^{-1} could probe a very sizeable region in new parameter space, well beyond the BaBar search more than a half decade ago.

- Belle needs to adopt (and perhaps has already adopted) the Belle II methods and software for treating and generating the data and MC samples, and to share assessment of the detector efficiencies, resolution, and different backgrounds and coverage, as well as the detector response for signals and backgrounds. The two detectors and operational circumstances are different. The high precision results will require scrutiny and should be consistent for the two experiments, but not strictly identical.

5.2.3 Recommendations

- The Belle II Collaboration should extend the use of run-dependent simulations. The same holds for the work on FEI and beam-energy corrections.
- Among the many research projects, the tests of the lepton universality in semileptonic B decays should include not only the $B \to D^*\ell\nu$ decays but also the $B \to D\ell\nu$ decays, which have a smaller branching fraction but potentially higher sensitivity.
- Joint publications based on Belle and Belle II data have great potential to enhance the precision of important measurements and should be pursued. Up to now, there have been only a few common publications, while the richness of the data call for intense further actions for the benefit of the scientific results.
- The Committee recommends that the Belle and Belle II Collaborations merge to benefit from the large Belle data set and the more recent design of the Belle II detector and large efforts on data processing, analysis software and simulation. This will avoid duplication of effort and streamline the completion of many analyses.
- The Belle II Collaboration is encouraged to seek the help of dark-sector experts to optimise the physics reach of the current and future program. Many dark sector analyses can lead to world-leading results already now; therefore, more human resources will be critical in this area.
- The Committee is pleased by the restart of a structured platform for regular exchange with theorists. More activities and exchange of ideas with theorists are very encouraged.

6 Upgrade beyond Long Shutdown 1

The machine upgrade is focused on delivering a total integrated luminosity of 50 ab^{-1} by the year 2035. This target will require upgrades to the detector and computing capacity to record and effectively process and exploit this very large and unique data sample. The Upgrade Working Group (UWG), which was formed in October 2018, gave an update on their activities on replacing and upgrading those detector components that would not be able to maintain their expected performance for the desired integrated luminosity. The drivers for the upgrade are to improve the detector robustness against background, to increase the long-term radiation resistance, to develop the technologies coping with possible future SuperKEKB and Belle II upgrade paths, and to improve the physics performance and extend the physics reach of Belle II. The BPAC applauds this effort and notes that there remain significant uncertainties driven by the fact that the upgrade plans of SuperKEKb are still being discussed and have not been finalised. The timescale for a possible upgrade is the next long shutdown, LS2, extending at least one year in the time frame of 2027 to 2028. The current discussions are anchored around this time-frame. The scale of the upgrade being discussed is large and time is short. The exact implementation of the accelerator upgrade will significantly influence the design of the detector upgrade, especially the region close to the interaction point. The BPAC therefore has one overarching recommendation for this section: scope, timeline, decision making process and the spatial envelope for the detector upgrade should be decided, preferably before the end of this calendar year. In order to achieve this, Belle II should work closely with the SuperKEKB (SKB) accelerator group and the SKB International Task Force (ITF). The following three subsections will review the implications for the physics, machine-detector interface and the current upgrade plans for the detectors. Their recommendations all reflect a similar sentiment expressed above, but with more specificity with regard to the subsystem.

6.1 Physics

6.1.1 Status

The UWG presented their assessment of the extended physics reach with an upgraded detector, emphasising that the LS2 SuperKEKB upgrade is not yet well-defined and it will be essential to develop a strategy for the period after LS1. The team is working towards a reinforced organisation for CDR preparations, aiming at completing a CDR by this spring and realistic detailed plans by the time of the next B2GM in June. To this end, the team is proposing a very ambitious plan to study the performance of tracking, PID, reconstruction of neutrals, and triggers for low-multiplicity final states. All of this with multiple geometries of upgraded detectors. The studies concentrate on signal Monte Carlo only with small event samples to allow for fast turnaround. A set of benchmarking channels have been identified. For example, the low momentum efficiency is being studied by evaluating the decay $B \to D^* \tau \nu$, the vertex resolution through the decay $B^0 \to J/\psi K_S$ and the π^0 efficiency and resolution is determined through the decay $B^0 \to \pi^0 \pi^0$. Various layouts of a new five-layer vertex geometry have been studied looking at the K_S reconstruction efficiency and purity and the B^0 purity. An improvement in the proper decay time difference resolution is obtained. A very clear improvement in performance is seen for the decay $B^0 \to D^* \ell \nu$, in particular for the soft pion reconstruction. A comprehensive plan to estimate the physics performance of an upgrade of the Belle II detector is being developed.

6.1.2 Concerns

- The committee is very concerned that the physics simulation task supporting the upgrade is too ambitious. It will require strong engagement by the collaboration.

It is not clear that these resources are available.

- To validate detector choices, a wide range of parameters needs to be studied for different geometry models. There is a significant overhead associated with these studies. Full simulations may not always be needed to inform various decisions and full simulations would not be available on the time scale of a CDR.

6.1.3 Recommendations

- The committee prompts the Collaboration to further sharpen the physics case, including the case with beam polarisation. It will be important to have a more defined physics scope beyond the improved $\sin^2\theta_W$ measurement and exploration of $g-2$ of the τ lepton, low-momentum tracking efficiency, and the reconstruction efficiency of neutrals, which will be very important.

- Explore if parametrised fast simulations would be adequate to provide sufficient support for the choice of detector parameters and geometries.

6.2 Machine and machine detector interface

6.2.1 Status

Several interesting proposals were presented as ways of improving the performance of SuperKEKB. These suggestions are designed to enhance the luminosity of the machine up from 2×10^{35} to closer to $6\times10^{35}\,\mathrm{cm^{-2}s^{-1}}$. Two of the suggestions involve modifying the final focus cryostats. The first increases the size of the beam pipe inside the QCS without affecting the outside dimension of the cryostat through the use of Nb_3Sn superconductor. The second adds more to the first suggestion by moving the QCS magnet closer to the IP and modifying the compensating solenoid windings. This second suggestion is a change of the size and location of the cryostat and this does affect the detector by eliminating some of the space currently in use for the PXD and for cables. In addition to these suggestions for the IP, there were several proposals for improving the injector complex.

in addition to replacing some of the older accelerating structures in the main linac, there is a plan to install an electron beam Energy Compressor System in the beam transport line. The compressor will reduce the energy spread of the electron beam which will decrease the dispersion and should make the electron bunch smaller thereby improving injection efficiency into the main ring and reducing detector injection backgrounds. Another proposal is to move the electron transport line into the tunnel currently used for transporting the electrons to the AR ring of the photon factory. This change would reduce the number of tight bends of the current beam transport line and would reduce the coherent and incoherent synchrotron radiation. There was also a proposal to modify the optics for the electron beam dump so that pulse-by-pulse diagnostics could be installed.

6.2.2 Concerns

- All of the above suggestions look promising. However, careful and complete studies are needed to ensure that the time and effort for each proposal will pay off in significantly better performance, as pointed out in the presentation.
- For some of the more complicated options that require significant modifications to the cryostat, time will be very limited, since any decision will have an impact on detector upgrades, especially on the design of the silicon tracking system.
- It is not clear how and when the upgrade plan can be finalised.

6.2.3 Recommendations

- A significant effort should be invested in obtaining simulation results as soon as possible for the proposed changes to the IP cryostats as well as for the different beam transport lines.
- Rank the different proposals as to how much improvement one could expect, in order to guide the decision process.
- The committee recommends that the Belle II collaboration and the SuperKEKB team set up procedures and deadlines to finalise the LS2 upgrade plan. This is particularly needed for upgrades that affect the machine and detector envelopes.

6.3 Detector

6.3.1 Status

The UWG has submitted a whitepaper to the US Snowmass process describing proposed near-term and potential longer-term upgrades of the Belle II detector. From the large range of options, the decision was made to defer a possible upgrade of the ARICH, ECL, and STOPGAP detectors to a date beyond LS2. The team is now working towards a Conceptual Design Report (CDR) to be submitted by the June BPAC meeting. The CDR is expected to specify the various paths on which upgrades can be developed, reducing the number of technology options, but leaving the door open to further R&D depending on the LS2 timescale. A technology table will be built towards that goal, using the timing and length of the LS2 as parameters and will assess the technology readiness level (TRL) of the various options. The CDR will be based on a technically driven schedule; a full cost estimate is deemed premature.

Having a polarised electron beam where the electron helicity can be flipped bunch-by-bunch would enable a unique physics program and could open up the possibility for physics channels with excellent signal-to-background ratios. To enable this, a low-emittance polarised electron source will be required, as well as spin rotators and a Compton polarimeter. This will require significant contributions by both the collaboration and the experienced accelerator team.

For a possible VXD upgrade, the requirements to improve the physics reach are pretty well-established. Besides excellent position resolution and low material budget, with a power dissipation of <200 mW/cm^2, an integration time of less than 100 ns is required. Safety factors for radiation hardness and hit rate can be derived from detailed simulations.

A candidate technology is depleted monolithic active pixel sensors (DMAPS), a technology that has already been extensively studied by the community with devices produced by Tower-Jazz (TJ) within the Monopix family. The design is being advanced with the "Obelix" implementation with features adapted for Belle II. Mechanical layouts for both the inner and outer layers are being developed. An alternative technology is the Silicon On Insulator (SOI) process, with a first prototype ready for testing.

To avoid further deterioration of the performance of the Central Drift Chamber (CDC), a partial or total replacement of the CDC is being considered, as well as new front-end readout electronics to reduce cross talk and provide better timing. A partial replacement of the CDC could be a silicon pixel or strip detector. For a full CDC replacement, a full silicon tracker is being contemplated. Thin (140 μm) double-sided silicon detectors are being studied, read out with the SNAP128 chip, which provides binary readout with a target noise of less than 800 e^-.

For the TOP detector, a full replacement of the conventional PMTs is foreseen with the life-extended ALD PMTs, which will complement the life-extended PMTs currently being replaced in slots 3-8. No deterioration for muon identification is expected with the KLM system and the efficiency of K_{L}-identification is currently being studied in simulations. Some research on SiPMs is being proposed in case of sudden, unexpected drop in efficiency of the current photodetectors.

An upgrade of the Front-end Timing Switch (FTSW) module is also being considered. The Cat7 connections between crates are currently problematic with spurious noise on the inter-crate cables. A first prototype module will be produced this fiscal year. If the decision is to upgrade the system, the production run of the modules would take place in the year 2025.

6.3.2 Concerns

- The BPAC appreciates the efforts by the UWG to date but notes that the physics and performance studies are still not at the level needed to address and compare quantitatively different upgrade options.

- There is a very strong coupling between the upgrade plans of the accelerator, in particular the final-focus magnets, and the detector. The final configuration of the accelerator upgrade will determine the spatial envelope for the upgrade and inform the background estimate. There remain large uncertainties in the exact scope and timeline of the accelerator upgrade. This will have a significant impact on the schedule of the detector upgrade and choice of technologies. As long as the accelerator plans have not been finalised, the Belle II collaboration can only proceed based on their best judgement. This is an impediment, especially given the very short time scale to LS2, if the current schedule of 2027 is to be maintained.

- The committee thinks that it is premature to start the actual detailed design of the detector subsystems at this time, given that the detector upgrade is strongly dependent on the accelerator upgrade scope and in particular the design of the machine-detector interface. The collaboration would benefit from continuing to remain open to promising technologies and advancing them for use at SuperKEKB. This allows the collaboration to proceed rapidly in defining the scope and design of an upgrade of the detector, optimised for the future configuration of the accelerator.

6.3.3 Recommendations

- Work with the accelerator group and the ITF to define the scope of the accelerator upgrade for LS2, judge the machine-detector interface, and determine the boundary parameters for the detector upgrade as quickly as possible.
- Delineate what upgrades can be deferred to a shutdown after LS2 and focus on those upgrades for LS2 that have the most impact on physics capability. Defer any effort on the actual design of upgraded detector subsystems until the accelerator upgrade plans have been finalised and concentrate in the interim on promising detector technologies and advance them for use at SuperKEKB.
- A decision-making process for the different upgrade options should be included in the CDR preparation.
- Evaluation of the need for an upgrade of different subsystems should continue and be checked by simulations. In this context, the simulation effort should be strengthened.
- Work with the accelerator group to understand if the resources are available to carry out a feasibility study of electron polarisation.