B-factory Programme Advisory Committee
2019 Annual Review Meeting
Full Report

11–13 February 2019 at KEK

A. Andreazza* (Milano), I. Belyaev* (ITEP), G. Carlino* (Naples),
D. Cassel (Cornell), P. Collins* (CERN), G. Corti (CERN)&,
M. Demarteau (ANL), S. Gori (UC Santa Cruz), M. Ishino (Tokyo),
Z. Ligeti (LBL), P. Mato* (CERN), P. Mcbride*& (FNAL),
F. Meijers* (CERN), B. Ratcliff* (SLAC), M. Sullivan (SLAC),
H. Tajima (Nagoya), M. Titov* (Saclay)
and chaired by T. Nakada (EPFL)
*Expert members &Remotely participated in part

11 May 2019

1 Executive summary

An annual meeting of the B-factory Programme Advisory Committee (BPAC) took place from 11th to 13th of February 2019 at KEK, with presentations covering all aspects of the status for the Belle II experiment, including strong and successful ongoing analyses of Belle data. This section is an executive summary focusing only on the most important issues and a complete report is found in the following sections.

The Belle II experiment is about to start the Phase 3 run for physics. There have been impressive achievements toward understanding the machine backgrounds with the Phase 2 data. However, there are still backgrounds that cannot be reproduced by the simulation and their origins have not yet been identified. Further studies to determine the individual background sources are encouraged in order to improve the understanding of the SuperKEKB machine. For Phase 3, additional collimators and new masks have been introduced and eight extra diamond sensors have been placed around the four bellows between the QCS region and the beam pipe. These sensors will carefully monitor the background and will be used to protect the experiment. However, given that the Belle II detector is equipped with the full first layer of pixel sensors located very close to the beam, the Phase 3 commissioning must proceed very carefully. The close collaborative work with the machine group established during the Phase 2 run should continue and be strengthened for Phase 3. The Belle II collaboration should remain prudent in order to avoid damaging the detector in this critical phase.

There is not much time between the start of Phase 3 and the summer shutdown. A very interesting set of measurements has been identified for presentation at the summer conferences. The plan for detector commissioning, including calibration, and for data taking may require careful optimisation and prioritisation in order to achieve those measurements.

The committee was pleased to see that the Vertex Detector system (VXD), although the second layer of the Pixel Detector (PXD) is only partially equipped, is now installed and ready for data taking. Testing the capability of the VXD for taking data at the nominal rate, including gating for the PXD, is now important. The present PXD will be replaced by a new one with two fully equipped layers during the long shutdown in 2020, at the earliest. A detailed planning for the construction of the new PXD shows that the group is making good progress. However, the work must be carefully monitored to ensure the timely completion of the detector.

While the Central Drift Chamber (CDC) is now ready for Phase 3, an occasional high dark current in some outer layers is a serious concern. The committee welcomes the creation of a task force of experienced members of the Belle II collaboration to deal with this problem. For the Phase 3 startup, a set of systematic and structured step by step approaches should be devised addressing different possible causes for the problem. Gas detector experts outside of the collaboration, including those from BPAC, should also be consulted to set up this programme and to analyse the results. In parallel, the committee recommends that the collaboration set up a team who should construct a small test chamber with exactly the same material to observe the long-term behaviour of the chamber.

The barrel particle identification system (TOP) has been operating reliably. Particular care should be taken to monitor that the bar boxes and DAQ operations remain stable. The case of the wrong starting time (T_0) for the propagation time measurement must be understood. Preparations for replacing radiation damaged photon detectors (the non-ALD MCP-PMTS) during the long shutdown in 2020, at the earliest, is progressing well. Although the photon detectors might last longer than currently anticipated, depending on the actual background level, the preparation work should be carried out as planned assuming the 2020 date. The ARICH end-cap particle identification system cooling issue has been resolved. Although some more work is required to complete the full system, the required steps seem to be well understood and the committee sees no particular concern for attaining good operation in Phase 3.

While the Electromagnetic Calorimeter (ECL) is ready for Phase 3, the committee is concerned about the status of the K-Long Muon detector (KLM). Integration of the barrel and end-cap KLM teams to work as a single group seems to be working well. But the overall human resources seem to be still scarce. The well-functioning KLM, including its trigger, is essential for the physics programme planned for the early period of Phase 3. The KLM group, together with the Belle II management, is encouraged to carefully examine and prioritise commissioning tasks in light of the detector elements needed most for the first physics results.

Although not all the subsystems are ready to provide trigger primitives, the overall trigger is adequate for the physics goals of Phase 3.

Experiments cannot operate without well-functioning online and control systems, even if all the subsystems are working well. The committee is glad to observe all the elements for the control and online systems are in place. On the other hand, the committee feels that the overall system is in a fragile state. Operation with the complete Phase 3 detector at nominal rate is still to be demonstrated. Although improvements have been made to increase the core group members, the work still relies on a small number of experts. The committee has already suggested deployment of technical support personnel who would take care of the computing and electronics infrastructure. Another important step is to enlarge the online group by recruiting not only experienced people but also early career people to guarantee long-term support of the system and to explore modern working practices, commonly used by large-scale experiments, for the development and deployment of the system. Using some of the maintenance and operation funds for supporting people at KEK should also be considered. While the high workload of the central online team is readily acknowledged by the committee, the team is urged to keep paying attention to the needs of the sub-detector teams for their commissioning work.

The physics software is now in a good shape. All the important elements of the offline computing infrastructure are steadily making good progress. The committee thinks that the computing team should now shift its attention to the robustness of the complete chain from the raw data storage to the physics analysis. Operation must be automated as much as possible for ensuring smooth running. The committee recommends introduction of Rucio-based distributed data management for long-term sustainability, although the changeover of the system should not disrupt data taking.

The Belle II collaboration recently formed an upgrade working group. We encourage this effort, to explore the physics case and technical feasibility of maximising the scientific return on the investment in SuperKEKB and Belle II.

The BPAC welcomes that the Belle II collaboration recently formed a Diversity Committee. We encourage maintaining data on various minority categories, which should help promote an inclusive and welcoming environment for all collaborators.

2 Physics

2.1 Belle

2.1.1 Overview

The Belle collaboration remained very active in 2018, eight years after the end of Belle data taking. In 2018 the collaboration consisted of 472 scientists from 93 institutions in 21 countries. Utilising the world's largest data sets at the Υ resonances $1S$, $2S$, $4S$, and $5S$, the Belle collaboration maintains a very strong program of physics analyses.

In the past year, 28 new papers were published, an increase from 17 papers in 2017, and increasing the total number Belle publications to 512. This vigorous publication activity has a large impact in terms of citations and entries in the Particle Data Group reviews, and is also reflected in strong presence at conferences. Together with LHCb, Belle is co-leading the Heavy Flavour Averaging Group (HFLAV). Belle physics analyses will continue until the data set is superseded by Belle II. The Belle II analysis framework, b2bii, plays an important role in the analysis of Belle data. The b2bii framework was used to validate the Full Event Interpretation (FEI), developed by Belle II, and to estimate the improvement of flavour tagging in Belle II with the Belle data. It was also used to analyse $B \to D^{(*)}\tau\nu$ and $B \to J/\psi K_S$ decays.

Among the recent 2018 highlights was the PRL and PRD publication of a joint analysis by Belle and BaBar for the time dependent Dalitz-plot distribution in $B^0 \to D^{(*)0}h^0$, $D^0 \to K_S\pi^+\pi^-$. This analysis resulted in the measurements of $\sin 2\phi_1 = 0.80 \pm 0.14 \pm 0.06 \pm 0.03$ and $\cos 2\phi_1 = 0.91 \pm 0.22 \pm 0.09 \pm 0.07$. These results exclude a discrete ambiguity at $\phi_1 > \pi/4$, allowed by constraining ϕ_1 only from the more precise measurements of $\sin 2\phi_1$ in the gold-plated modes.

2.1.2 Spectroscopy highlights

Hadron spectroscopy can be explored at Belle in numerous processes, in particular, in decays of B mesons, charmed baryons, and bottomonium.

The charmed baryon $\Xi_c(2930)^+$ was observed in the decay $\Xi_c(2930)^+ \to \Lambda_c^+ K_S^0$, and its mass and width measured, utilising the process $B^0 \to \Lambda_c^- \Lambda_c^+ K_S^0$ and measuring its branching ratio. In $B^- \to \Lambda_c^- \, \Xi_c(2930)^0$ decay, both this B^- branching ratio and some new $\Xi_c(2930)^0$ branching ratios were measured, and the results are to appear in PRL.

Belle also reported the observation of the doubly-strange baryon, $\Xi(1620)^0$, in its decay to $\Xi^-\pi^+$, by studying $\Xi_c^+ \to \Xi^-\pi^+\pi^+$. They made its most precise mass and width determinations, and also obtained a 4σ evidence of a $\Xi(1690)^0$ state with the same data sample. These results will also appear in PRL.

Other highlights include the observations of $\Upsilon(4S) \to \Upsilon(1S)\eta'$ and $\Upsilon(2S) \to \eta_b\gamma$, a search for $Y(4260) \to J/\psi\,\pi^+\pi^-$ in the decay $B \to J/\psi\,\pi^+\pi^- K$, and measurements of the p_T dependence of π, K, p production rates.

2.1.3 B, D, and τ physics

Flavour physics with B, D, and τ decays using Belle data remains a rich field of investigation. On many hot topics, Belle continues to provide important new results.

In the area of semileptonic B decays, Belle carried out new measurements of $B \to D^{(*)}\pi\ell\nu$ decays, for both the charged and the neutral D and D^* modes, using full reconstruction and the reprocessed data. These decays are important, as they constitute important backgrounds to lepton flavour violation measurements in $B \to D^{(*)}\tau\bar{\nu}$, at Belle, Belle II, and LHCb. They can also help understand $B \to D^{**}\ell\nu$ decays better, where D^{**} denote the four lightest orbitally excited D states, for which there are tensions between past Belle and BaBar measurements. The branching ratios have been measured, and an amplitude analysis is forthcoming.

The $B^\pm \to \ell^\pm\nu\gamma$ decay is important, because the photon lifts the helicity suppression of $B^\pm \to \ell^\pm\nu$ in the SM. The decay rate for energetic photons is calculable from first principles, and is sensitive to λ_B, which is a weighted integral over the b quark distribution function in B mesons. This same parameter also enters predictions for charmless hadronic two-body B decays, of great importance to study CP violation. Using the recently developed FEI algorithm for B tagging, the bound $\mathcal{B}(B^\pm \to \ell^\pm \nu\gamma)|_{E_\gamma > 1\,\mathrm{GeV}} < 3.0 \times 10^{-6}$ was obtained, improving upon earlier limits. The measured ratio $\mathcal{B}(B^\pm \to \ell^\pm \nu\gamma)|_{E_\gamma > 1\,\mathrm{GeV}} / \mathcal{B}(B^\pm \to \pi^0 \ell^\pm \nu) = (1.7 \pm 1.4) \times 10^{-2}$ implies $\lambda_B = 0.36^{+0.25}_{-0.09}$ GeV or $\lambda_B > 0.24$ GeV at 90% confidence level.

Belle performed searches for lepton flavour violating B decays, such as $B \to K^* e^\pm \mu^\mp$, motivated by recent hints for lepton universality violation. They set a new upper limit at 1.8×10^{-7}, which is a factor of 3 smaller than the previous best limit from BaBar.

The inclusive $B \to X_s \gamma$ decay remains an excellent probe of new physics. Belle set new bounds on the isospin and CP asymmetries, and significantly improved constraints on the so-called resolved photon contribution to the decay rate, which has been argued to be larger, and a big part of the uncertainty of the SM prediction.

Belle also performed searches for CP violation in $B^\pm \to K_S^0 K_S^0 h^\pm$ decays, motivated partly by the large CP asymmetry observed by LHCb in limited regions of the $B^\pm \to K^+ K^- \pi^\pm$ Dalitz plot. New and sensitive searches for direct CP violation in $D^0 \to K^+ K^- \pi^+ \pi^-$ were also published.

Searches for very weakly coupled light particles are important parts of the physics program. A new limit on an invisibly decaying CP-odd scalar, A^0, with mass below 9 GeV, was obtained in $\Upsilon(1S) \to \gamma A^0 (\to \chi\chi)$, using Belle's large $\Upsilon(2S) \to \Upsilon(1S)\pi^+\pi^-$ sample.

The Belle collaboration presented well-considered plans for new analyses and publications in 2019, in many areas of B, D, and τ physics.

2.1.4 Concerns and Recommendations

- The committee congratulates the Belle collaboration on continuing to successfully produce exciting physics results. The BPAC encourages continued analysis of Belle data with the b2bii framework, and the ongoing efforts to eventually merge Belle into Belle II. Simultaneously, the Belle roles in HFLAV should be transferred to Belle II.

2.2 Belle II Physics

2.2.1 Overview

The Belle II physics book has been submitted for publication in PTEP and has already received over 100 citations. The first public notes based on the Phase 2 data set (up to 500 pb^{-1}) are underway. These include notes on physics measurements and detector performance. There is the intention to publish these notes (subject to publication committee review). In particular, a few dark sector searches are on track for early publication.

The collaboration is getting ready for the Phase 3 data taking. An important goal will be to demonstrate the Belle II detector performance, and in particular the performance of the VXD, through the analysis of key channels. The luminosity targets for 2019 are 10 to 50 fb^{-1} by July, and 100 fb^{-1} by the end of the year. There are good publication prospects for dark sector searches with the Phase 3 data by the Lepton-Photon 2019 conference. These early searches will likely use a relatively small data set of $\mathcal{O}(1\ \text{fb}^{-1})$, but might be nevertheless competitive with the corresponding Babar and Belle analyses. Even though most of the new heavy flavour publications are likely to start later with the 2020 data set, several analyses are planned for Lepton-Photon 2019. These include analyses across all working groups. The data planned for 2020 ($\sim$500 fb^{-1}) will be smaller than the full Belle data set, but it will offer the opportunity to make the most precise measurements of several processes across all working groups.

Finally, an impressive longer term plan for physics publications has been laid out. This includes the most precise measurement of $B \to D^{(*)}\tau\nu$ in 2021 with $\sim$2 ab^{-1} of data, and measurements of rare $b \to s$ transitions competitive with the corresponding measurements at LHCb in 2022 with $\sim$6 ab^{-1} of data.

At least 10% of data is planned to be taken regularly off-resonance, with the highest centre-of-mass energy depending on RF cavity availability, probably around 11.25 GeV, about 200 MeV above the $\Upsilon(6S)$. Besides the B_s data that $\Upsilon(5S)$ running can provide, the physics case of data taking at higher energies has also been studied.

2.2.2 Dark sector, low-multiplicity, and τ

A few analyses based on Phase 2 data are on track for publication. Particularly, the dark sector working group is pushing forward the publication of a search for a new gauge boson, Z', produced in association with $\mu^+\mu^-$ and decaying invisibly, using 276 pb^{-1} of data. This signature arises in well-motivated models that can address, e.g., B-physics anomalies, as well as the $(g-2)_\mu$ anomaly. This analysis might be the very first physics publication based on Belle II data. At the same time, the Belle collaboration also plans to search for the same signature with its 1 ab^{-1} of data, with the goal of completing the analysis by the spring or summer, before the Phase 2 Belle II publication. This Belle II search is expected to start probing the $(g-2)_\mu$ favoured region with $\mathcal{O}(100\ \text{fb}^{-1})$ of Phase 3 data.

The search for $e^+e^- \to a\gamma$, $a \to \gamma\gamma$ utilising the Phase 2 data (472 pb^{-1}) is also on track for publication: this search will likely give access to unexplored regions of parameter space of models for axion-like particles coupled to photons. Other searches for dark particles are planned for early Phase 3: examples are the search for lepton flavour violating Z', and the search for an invisible dark photon, $e^+e^- \to \gamma A'$ with $A' \to$ invisible. The latter search needs an efficient KLM veto to be competitive (this was not available for most of Phase 2).

A Phase 2 public note on the tau mass measurement using 1 and 3 prong decays is envisioned. Performance and data challenge studies are also planned. The tau working group has laid out a clear path towards the completion of several early publications. Possible publications include $\tau \to ea, \tau \to \mu a$ with a a Goldstone boson, and $\tau \to \omega h\nu$, $\omega h\pi^0\nu$.

2.2.3 (Semi)leptonic and electroweak penguin B decays

Several analyses for both (semi)-leptonic and missing energy B decays, and electroweak penguin decays are planned with Phase 2 data. As a highlight, the collaboration plans to complete analyses for $B \to D^* \ell \nu$ and for $B \to X e \nu$ with ~ 490 pb^{-1} of data. Studies of the performance of FEI hadronic tagging in Phase 2 data are also planned. This is particularly important for channels with missing energy. As an example, thanks to the implementation and application of the FEI in the Belle II Analysis Software Framework, Belle II should be able to observe the decay $B \to K^{(*)} \bar{\nu} \nu$ with 4 ab^{-1}.

An early goal for Phase 3 will be the rediscovery of $b \to s$ penguins. The collaboration also plans further data challenge studies, and to continue preparing the tools required by many B decay analyses (FEI calibration and validation, fitting tools, etc.).

2.2.4 Other physics results and performance studies

The Belle II collaboration made impressive progress using the Phase 2 data to commission the detector, understand its performance, and to establish the data preparation and software tools to make the data available for the collaboration for physics analysis. Belle II has observed and re-discovered many light mesons and baryons, charm and charmonium states, and a number of key B decay modes, with good reconstruction efficiency and resolution.

As early Phase 3 goals, the BPAC was shown plans based on acquiring a data set of $(10 - 50)\,\text{fb}^{-1}$ by Summer 2019, $(100 - 200)\,\text{fb}^{-1}$ by Winter 2020, and $\sim 500\,\text{fb}^{-1}$ by Summer 2020. Flavour tagging and fitting displaced vertices at Belle II have been improved beyond Belle's performance, and will more than compensate for the reduced $\Upsilon(4S)$ boost in the lab frame, compared to Belle. By this Summer, the collaboration plans to continue improving detector performance, B reconstruction, flavour tagging, continue the rediscovery of many B and D decays, and maybe even make initial time-dependent CP violation measurements. The impact of the much smaller beam-spot compared to Belle has to be carefully considered. With these initial data sets the most novel measurements will be dark sector searches. By Summer 2020, Belle II may have a data set between BaBar's and Belle's, at which point new analyses could start seriously impacting world averages.

2.2.5 Concerns

- The collaboration has ambitious and well-considered goals to obtain physics results from the initial Phase 3 data taking, by the 2019 Summer Conferences, based on an anticipated $\sim 10\,\text{fb}^{-1}$ of data set. This requires good performance of SuperKEKB and improved understanding of remaining issues with Belle II detector performance.

2.2.6 Recommendations

- The Phase 3 schedule shown to BPAC at this meeting anticipates a slower luminosity ramp-up than earlier plans. It is important to avoid further delays, to maximise competitiveness with LHCb and to provide the best opportunities for discoveries.
- Although the Belle II physics book has recently been completed, continuing the close interaction between the collaboration and the theory community is encouraged. This can be particularly relevant for having a broad program for dark sector searches.
- The current run plan contains seven months of data taking in 2019 and eight months in the later years. The collaboration and KEK management are encouraged to try reaching the original goal of nine months of data taking per year in the coming years.

3 Machine and related issues

3.1 Status

The accelerator demonstrated significant progress during Phase 2 commissioning achieving a luminosity of 5.5×10^{33} cm^{-2}s^{-1} with beam currents of about 750 mA in both low energy and high energy rings (LER and HER, respectively). The nano-beam scheme was validated. The beams were colliding with a significantly shorter β_y^* than the bunch length of 5-6 mm and the luminosity increased as β_y^* decreased. The detector was able to operate at a luminosity of about 3×10^{33} cm^{-2}s^{-1} and collected about 500 pb^{-1} of data. The SuperKEKB goal for the first part of Phase 3 is to achieve luminosities above 10^{34} cm^{-2}s^{-1} with beam currents of $\sim$ 1.2 A (LER) and $\sim$ 1.0 A (HER). This is with the 3mm β_y^* lattice they had when Phase 2 ended. If the initial Phase 3 commissioning goes very well they will try to achieve a luminosity of $>1.5\times10^{34}$ with beam currents of $\sim$ 1.4 A (LER) and $\sim$ 1.0 A (HER) and a higher tune shift. In addition, significant accomplishments have been made in the Linac. Several issues were cleared up during the fall commissioning time and the committee looks forward to improved and cleaner injection during the Phase 3 run. Further injection tuning is planned for the Phase 3 run and continuous injection will be worked on in the spring.

Significant improvements have been made in getting the background predictions from the simulation to match the data collected during the Phase 2 run. The QCS cryostat geometry and materials have been corrected and this has led to an increased background rate prediction that is much closer to the observed background rates. There are still rather large differences between the predictions and the HER background data (factors of 10 to 50 higher for the data values) and these are primarily coming from the beam-gas and Touschek sources. These backgrounds will be investigated early in the Phase 3 running. New collimators will be added and these should help to improve both the HER and LER background rates. Luminosity backgrounds, which will eventually be the largest background for Belle II, have not been well measured yet since the luminosity is still low.

3.2 Concerns

- Background estimations from the simulation are now closer to the measured Phase 2 values but that also means the actual backgrounds are too high. The single beam background rates need to come down about a factor of two in the short term and ten in the long term.
- Two collimators were found to be damaged during the Phase 2 run and have been replaced. This is a concern as the beam currents have not gone much above 1 A and the design currents are near or above 3 A.
- In addition to the collimator damage, it was found that the bellows closest to the IP have been damaged due to arcing. Arcing can cause beam instabilities and/or backgrounds by injecting small (micron sized) blobs of material into the beam.

3.3 Recommendations

- Continue the effort presently planned to understand the Phase 2 and early Phase 3 backgrounds and get the predictions to match the measured values.
- The committee recommends a careful monitoring of collimator effectiveness to try to spot when damage occurs. The arcing in the IP bellows will probably be noticed by large background spikes or unexpected beam instabilities. Continuing study of unexpected beam aborts due to background spikes may help to identify this particular issue if it persists.

4 Detector System

4.1 Vertex detector (VXD)

4.1.1 Overall status

The collaboration is to be congratulated for the successful installation of the vertex detector system (VXD) on November 21, 2018 with the full strip detector system (SVD) and all Layer 1 and two Layer 2 ladders of the pixel detector system (PXD). The first cosmic rays were observed and commissioning of the detector has started. Initial results are very encouraging and indicate that the design resolutions can be achieved. The committee appreciates that the collaboration has verified the spare parts at KEK to ensure that repairs can be completed promptly. The spare count seems fully adequate.

The VXD group has installed a few more diamond sensors than originally foreseen to provide better protection for the pixel detector. There are a total of 28 sensors installed, eight on the beam pipe, 12 on the SVD support cones and eight around the four bellows between the QCS region and the beam pipe. The readout electronics of the diamond detectors has been updated and the system is being commissioned. It is proposed to start Phase 3 operations with abort-threshold values at the same level as at the end of Phase 2 of the four sensors mounted on the beam pipe. The VXD is also protected by a VXD Local Hardwired Interlock (VLHI) system. The programmable VLHI is based on hard-wired inputs from environmental monitors and other systems, and interlocks the power to the detector in a fail-safe mode. The issues observed during Phase 2 with the dew point sniffers were addressed for Phase 3.

4.1.2 Concerns

The background situation of the VXD remains worrisome. A factor of up to ten in background reduction is required based on scaling the rates from Phase 2 running for reliable operation of the detector. A plan is being developed in collaboration with the machine group to achieve reasonable background rates by mid-April, 2019. Good progress is made in the area of simulations. The new results indicate that the new collimators can address the background issues, but available personpower to run the simulations remains a concern.

4.1.3 Recommendations

- The committee voices again its opinion to proceed very cautiously and closely monitor the detector during the critical initial phase of running the detector.

4.1.4 Pixel detector (PXD)

4.1.4.1 Status

The PXD group continues to make very detailed studies of the background. The backgrounds from the HER and LER are much higher than expected and their origins are not well understood. As noted above, a factor of up to ten in background reduction is required based on scaling the rates from Phase 2 running and a plan is being developed in collaboration with the machine group to study and mitigate the high background rates.

First measurements of the alignment of the PXD has revealed that the gap between the lower and upper half-shell of the PXD is larger than anticipated by about 0.5 mm due to an unforeseen interference of the kapton circuit with the cooling block mounts. Additionally, some small distortions are observed in the upper half-shell after installation of the VXD. In particular, the data shows a slight increase in the bowing of some Layer 2 ladders by about 80 μm. The origin is being studied.

After installation of the PXD, the performance of a Layer 1 module, which is fortunately backed by a Layer 2 module, has degraded significantly and is no longer providing useful data. This module was a grade B module which, due to time constraints, was not replaced before installation. Many operational improvements are being implemented, notably the firmware to take data at 30 kHz. The Gated Mode operation of the PXD has not been exercised and should be a high priority for Phase 3. Significant improvements have also been implemented to make operation of the PXD more user-friendly and the pool of experts has been increased.

With the current PXD construction completed, with 20 out of 40 modules installed, the group has shifted its attention to ensuring that a full PXD can be installed during the summer shutdown of 2020 (PXD2020). The group has decided to build a complete replacement of the current detector, a move that is supported by the committee. An update of the preparations for the PXD2020 project was presented. In the previous report, a history of the module production was provided, which showed an overall ladder assembly yield of 77% and 50% for Layer 1 and Layer 2, respectively. A thorough ladder inventory shows that there are no good Layer 1 modules left for ladder production; for Layer 2 there are 15 good forward and 16 good backward modules left. From the previous PXD9 sensor productions, there are six good Layer 1 and two good Layer 2 sensors left, which would give three Layer 1 and one Layer 2 ladders, assuming 100% yield. Thus, based on the current inventory and assuming that the issues with the ladder production are resolved, three new Layer 1 and 16 new Layer 2 ladders can be built. In addition, there are three possible spare Layer 1 ladders and one possible spare Layer 2 ladder, for a sum total of six Layer 1 and 17 Layer 2 ladders. This would allow for construction of one full Layer 2 and one half-shell for Layer 1, including spares. The complement of required Layer 1 ladders is provided by the two batches of sensor production that started in March and July 2018, respectively. Each batch has 12 wafers and can provide six to eight ladders, depending on yield. The ladder assembly procedure has been completely revisited and an internal review is scheduled for the end of February. Unless there are insurmountable issues with the ladder assembly, the production schedule for completing the PXD2020 is sound. Availability of ASICs is not a concern, nor is the bonding of the ASICs since it is now done at the wafer level. Use of a new version of the Switcher chip is proposed, which has a slightly lower power consumption.

4.1.4.2 Concerns

- As noted before, it is a bit of concern that the 30 kHz gated mode could not be tested during Phase 2.
- With the start of the Phase 3 run, expertise will be required to commission and operate the full PXD detector and analyse its data. The construction of a complete PXD2020 detector requires the continuous availability of additional resources for nearly another two years. To keep the expertise and enthusiasm of the current PXD crew will require good resource management. Furthermore, the modified assembly procedures might require some modification of the current infrastructure for the assembly, requiring additional technical resources.
- Although the availability of components does not seem to be of concern, the schedule remains demanding for the PXD team, especially in the early part of 2020.

4.1.4.3 Recommendations

- Development and testing of the firmware for 30 kHz operation as early as possible is encouraged.
- The committee applauds the internal review of the ladder production and encourages that ladder assembly for the PXD2020 be started as soon as possible after the review to exercise the assembly procedures and the quality control process to ensure maximal yield and create schedule contingency.
- The committee encourages continued detailed simulation of the VXD with the latest beam pipe design, backgrounds at various levels and full performance evaluation of the VXD with PXD equipped with zero, one and two layers.
- With the addition of new collimators and masks, the committee reiterates its recommendation that significant effort be directed towards understanding the various background sources with the new configuration.

4.2 Strip detector (SVD)

After the successful VXD insertion on 21 November 2018, the SVD group has completed the detector installation and it is now commissioning for the Phase 3 run. The installation showed a fragility with the hybrid cables, resulting in cuts of the wires that were repaired. The tight space for cables and pipes also caused damage of the cable for the diamond detector, which was fixed in place. Minor rearrangements were made in order to position the QCS magnets, which was completed by 25 January. SVD has been taking part in cosmic-ray and high-rate calibration runs since January. The availability of spare parts for prompt repairs during operation has also been checked.

Analysis of cosmic-ray data shows a good response to minimum ionising particles, with a signal-to-noise ratio in the 15–18 range for the p-side and > 20 for the n-side. The efficiency is measured on sensors with an incident angle similar to particles in the e^+e^- annihilation events. It averages above the 99% level, in agreement with the Phase 2 Bhabha scattering data analysis. Nevertheless, the p-side clusters show splitting if the threshold is increased with respect to the nominal value $S/N > 3$.

The committee appreciates the large amount of work on the online software, that is now ready for Phase 3, and the serious planning of shift organisation, allowing for both local and remote shifters with defined roles.

The SVD takes part in high-rate tests, but it has not reached yet the 30 kHz trigger rate and already at 20 kHz it shows a dead time of 20%. The bottleneck is attributed to the partially deployed HLT system and to the COPPER readout PCs (ROPC). Mitigation actions are planned: increasing the number of HLT units and upgrading the ROPC network connection.

Extrapolation of the background rate for Phase 3 has been performed. Early Phase 3 operation should not be affected, but a Layer 3 occupancy up to 4.4% is expected at nominal luminosity, dominated by single beam background. No quantified effect on performance has been presented, but, besides the readout limitations, track reconstruction should be affected for occupancy larger than 2-3%.

4.2.1 Concerns

- The observed bandwidth limitation on the HLT and ROPCs may hamper Phase 3 operation if the background extrapolation is confirmed. It may also hide other limitations of the system that have not been identified yet.
- The extrapolation of Phase 2 background measurements to Phase 3 conditions suggests the SVD may operate with a background occupancy rate significantly higher than the design value.

4.2.2 Recommendations

- In order to cope with the expected data rate for Phase 3, possible solutions to the bandwidth limitation, such as the installation of more HLT units and upgrade of ROPCs should be implemented. Increasing threshold would be a less desirable solution.
- As already mentioned in the VXD section, with the addition of new collimators and masks, a significant effort in the early Phase 3 should be directed towards understanding the various background sources and updating the extrapolation to nominal luminosity.

4.3 Central drift chamber (CDC)

4.3.1 Status

The CDC group is making significant progress on the outstanding hardware and software issues identified in the October 2018 BPAC. However, some important challenges remain.

Front-end electronics problems were already traced to DC power supply issues. The DAC problem was mitigated by increasing a DC voltage and the performance improvement remains stable. A DC power supply was replaced to address the FADC problem, but it malfunctioned. Further progress is awaiting repair of the malfunctioning unit.

Software upgrades led to the improvements in understanding CDC tracking and resolution. Reconstructed track parameters and their resolutions from the dimuon events largely agree with those obtained from the simulated data. Further work is ongoing to understand the remaining small discrepancies.

Crosstalk was identified as a source of tracking inefficiency. The times of crosstalk hits are close to those of real hits. However, the ADC values for crosstalk hits are significantly lower than those for real hits, so an ADC cut improves tracking efficiency. So far, the source of the crosstalk is not yet understood.

Common noise in outer layers appears to be related to operation of the SVD Front-end Timing Switch (FTSW) crate. The SVD is now fully installed and there is hope that the improvements in cable shielding may mitigate this common noise issue.

The CDC group has developed a plan for shift personnel. There will be an on-site shifter at KEK, who will be on call to fix operational problems, take local runs, and attend daily run meetings. An off-site shifter will be responsible for offline data quality monitoring.

Occasional high dark currents in the CDC outer (L53-L55) layers during Phase 2 is a major concern, both for Phase 3 start-up operation and for Belle II data taking beyond 2019. Just after Phase 2, the situation became worse because dark current increases were also observed in several middle layers. CDC HV studies in November 2018 revealed that the dark current appearance in one of the most affected outer layers (L54) could trigger sporadic currents in neighbouring layers, confirming that the nature of the observed phenomena is the Malter effect. Observation that the dark current increases after HV is turned ON and then decreases within an hour is also consistent with the Malter hypothesis. A well-known intermediate solution to mitigate the Malter effect has been tried, i.e. increasing the CDC water content up to 2000 ppm. This suppressed the appearance of occasional dark currents in all layers. Still, this remedy does not guarantee the long-term survival of the CDC.

4.3.2 Concerns

- Even if no further degradation of CDC performance will be observed during the Phase 3 startup, taking into account that the luminosity still has to be increased by about factor of 50, the long-term behaviour of CDC beyond 2019 represents a very serious concern for successful data taking.

- Some relatively minor hardware problems, i.e. FADC performance, crosstalk hits, and common noise in outer layers, still need serious attention.

4.3.3 Recommendations

- The committee welcomes the creation of a CDC Task Force to deal with the dark current problem (both for the short-term and the long-term) and to devise a set of systematic and structured step by step investigations that address possible causes of the problem and their potential mitigations. Gas detector experts inside and outside of the collaboration should be consulted to set up the programme and to analyse the results.

- Even though addressing the dark current challenge is by far the major priority, significant effort on the other hardware issues is still required.

- Augmentation of the CDC team with new members in order to ensure adequate personnel for addressing all of the hardware challenges is a high priority.

4.3.4 Special remarks on the CDC dark current issue

After the start of beam operation on March 21, the CDC behaved reasonably well. No persistent dark currents were observed in the CDC recently with 2000 ppm of water. This subsection summarises a list of issues to be addressed in order to develop short-term and long-term coherent strategies for CDC operation in Phase 3. These recommendations are based on the outcomes of the BPAC review in February and the CDC Task Force weekly meetings.

4.3.4.1 Short-Term Program (The Phase 3 start-up)

Based on the outcomes of the recent studies, the committee encourages pursuing the following steps during early Phase 3 running:

- Until further notice, continue to operate with the following CDC layers HV settings during beam operation and collisions: L0–L52 nominal HV, L53 & L55 nominal HV minus 25 V or 50 V, L54 nominal HV minus 1000 V (OFF).
- Monitor continuously any sporadic or irregular current behaviour, not directly related to SuperKEKB luminosity and/or beam backgrounds, as well as the maximum and average current for each "potentially problematic" layer separately. Keep a record of accumulated charge for all layers in Phase 3.
- Since no persistent dark currents and/or spikes have been observed recently, with water content kept at 2000 ppm for an extended period, stabilise and operate the CDC with a water content of about 2000 ppm throughout Phase 3 start-up.
- Optimise procedures to stabilise the water content in the CDC, i.e. estimate the necessary rate from the slope of the concentration increase or decrease and correct it iteratively. Caveat: If the SVD is off for more than several hours, special precautions are necessary to avoid dew condensation in CDC.
- Study effect of water content on the r–t relation and tracking performance.
- Optimise and increase the gas-flow circulation rate through the CDC during beam operation, including purchase of an additional flowmeter, as soon as possible, to have the flexibility to increase the flow rate, if needed. Note: Simulation of the gas flow distribution inside the CDC revealed that the gas exchange might not be sufficient in the corners and outer layers of the CDC at a flow rate of 2 L/min. Conditions in those regions might be significantly worse for high-mass hydrocarbon components created in gas avalanches.
- Increase dry air circulation to the surface of endplates, since there was an indication that some CDC surface effects were observed during chamber construction.
- Develop a procedure (that was missing during Phase 2 operation) to minimise the duration of CDC operation at the maximum current limit level, defined by HV power supply protection settings. Develop a "dark current duration limiting" procedure, i.e. trip the HV if a current stays at a maximum level for a few tens of second; optimise thresholds, time constants and software trip conditions; and – if possible – run at lower HVs during machine studies.

- Verify that the current HV settings – (L53 & 55 – nominal HV minus 25 V or 50 V, L54 – OFF), and if CDC conditions worsen (HV is OFF on L53–L55) – do not significantly degrade physics performance, including: Modify trigger firmware as needed. Study the angular and position resolution of tracks extrapolated to the TOP detector, track matching with ECL/TOP/KLM, etc. Monitor tracking efficiency (gain along the wire length) vs. time to confirm that there is no significant gain drop along the z-coordinate. Check individual layers, looking for dE/dx per layer and, if possible, look at signal shapes (wave-forms) for problematic layers when the CDC is drawing excessive current.

- Perform CDC gas analysis to search for any pollution from high-mass hydrocarbon or Si-O-based molecular components in the mass range from 1 to 300 amu. Reference studies of CDC gas connected with the TOP gas analyser did not reveal any significant excess for high-mass fragments. Studies of the gas mixture using pre-mixed gas from bottles and from the CDC mixing system showed identical results, i.e. fractions of He, Ethane and a 28 amu peak are consistent between pre-mixed gas and CDC gas. It is important to repeat measurements once the chamber is operated with beam as the production rate of such components will depend on the current in the chamber.

4.3.4.2 Long-Term Program (Beyond the Phase 3 start-up)

The committee recommends that the collaboration set up a team, which will address the long-term survival of the CDC. Based on the outcome of the recent studies, the committee encourages that this team pursue the following steps:

- Investigate the possibility of removing some Al field wires from CDC layers L53–L55 during the 2020 shutdown, in order to subject them to SEM analysis. Note: SEM analysis of non-irradiated Al field wires, used for CDC construction, has revealed some pollutants (carbon-like deposits) on their surfaces.

- Construct a small test chamber, consisting of the same sense and field wires, endplate material, feedthroughs, and glue. Irradiate that chamber with heavily ionising particles; e.g. alphas from ^{241}Am or betas from ^{90}Sr or ^{106}Ru to reproduce CDC ageing effects and to validate long-term behaviour. If dark current or gain drop are observed, these data will serve as the basis for mitigating ageing effects in the current CDC or can be used for better understanding of the construction of a new CDC. Using the small test chamber, optimise the amount of water and, if necessary alcohol, to be added to the CDC gas mixture.

- Study the possibility of operating the CDC with a hydrocarbon-free gas mixture, e.g. He/CO_2, based on the outcome of experimental tests with a test chamber and simulation results. Optimise the CO_2 fraction in the mixture by studying if the lower and non-saturated drift velocity in He/CO_2 will deteriorate tracking and physics performance.

- In case CDC running conditions worsen significantly, evaluate the necessary resources and time to build a new CDC.

Depending on CDC performance during Phase 3 start-up and the outcome of the studies with a small test chamber, adjust the long-term strategy and CDC operation conditions for the run in the Fall of 2019.

4.4 Barrel particle identification (TOP)

4.4.1 Status

The TOP group has made substantial progress since the last BPAC, concentrating on the evaluation of TOP performance and the lessons learned from the Phase 2 data taking, including a number of important electronics/DAQ issues. They have also continued to do R&D and monitoring to understand long-term system integrity and mechanical issues, especially the PEEK delamination. The committee congratulates the TOP group for their hard work and progress and wishes them well as they move towards the Phase 3 data taking.

The TOP group appears on track to fully participate in the Phase 3 data taking. Phase 2 running was successful with major problems reasonably rare. The system hardware and calibration components are operating reliably and the TOP front end electronics and DAQ performances are stable. Operation manuals are being constructed. There is good progress on the slow controls. A number of chronic problems currently being addressed should not inhibit good operations. The TOP is nearly able to operate reliably at a 30 kHz trigger rate with physics occupancy. Work continues to make data taking more robust at high trigger and hit rates. TOP experts have been in high demand to fill vital experiment wide positions, so that finding and training replacements is a high priority.

The machine backgrounds in Phase 2 were much higher than estimated by the simulation and can not be tolerated long term. Therefore, background mitigation in the Phase 3 running is a major concern for the photon detectors of the barrel particle identification system, as discussed later. The group has appointed a dedicated Background Liaison to coordinate operations and analysis as the Phase 3 running proceeds.

The TOP software/analysis group has made significant progress in understanding the Phase 2 data, though some mysteries remain. Overall particle separation is significant but also substantially worse than expected. Though the timing resolution is worse than anticipated and appears to need an extra 100 ps added to the overall timing, this should not drive the present performance gap. Just before the meeting, the group discovered an issue with finding the correct bunch, but there are still unknown additional components driving the poor performance. Module alignment is also not fully understood and needs to be corrected, but this again should not dominate the overall performance. It was puzzling that there has been little performance improvement with improved reconstruction and better time calibration, but this may be, at least partially, related to failures in the bunch finding. Due to the connection with tracking, the group has appointed a tracking liaison.

The "partial" delamination of bar-box glue joints at the readout end, driven by thermal stresses at the PEEK frame to quartz glue joint, continues to be carefully studied. The inlet/outlet gas is being monitored for leaks and contamination using continuous dew point measurements box by box and an RGA gas analysis. No smoking guns have been yet identified that suggest box deterioration but the gas piping may need to be upgraded in order to understand the RGA studies better. Further off-detector R&D continues in order to understand how the reflectance deteriorates with contaminated gas, and how the PEEK delamination proceeds under thermal stresses.

4.4.2 Concerns

- Operations and upgrades to the LV/electronics/DAQ will continue to require substantial investment during the next few months from an expert crew. The loss of expert manpower from TOP and its renewal must be carefully managed.
- The plan to understand and replace photon detectors that have low QE lifetimes seems reasonable, as is discussed in the upgrade section.

4.4.3 Recommendations

- Continued studies of the data from Phase 2/Phase 3 remain crucial in order to understand and commission the detector and software to ensure against possible lurking hardware/DAQ problems and to obtain required performance for physics. The strongest possible effort should be maintained to ensure that all of this occurs.
- Understanding and mitigating machine backgrounds must be a very high priority of the Phase 3 running.

4.4.4 Photon detector replacement

4.4.4.1 Status

The 224 conventional MCP-PMTs, filling 7 slots, have inadequate quantum efficiency (QE) lifetimes and will need replacement soon. To maximise PMT lifetime, signal amplitudes from the PMTs during the Phase 2 running have been kept as small as possible, even though this may lead to some timing resolution degradation. Even with this mitigation, the machine background at present requires that these tubes be replaced during the 2020 shutdown. To survive until then, the upper limit to the PMT hit rate can not exceed 4 MHz/PMT in 2019-2020. Furthermore, in order for the ALD MCP-PMTs to survive until 2027, the backgrounds must be mitigated leading to an average rate for 2019-2020 of less than 1 MHz/PMT. If this target is attained, it may be possible to delay replacing the conventional tubes until the 2021 summer shutdown. Tube ordering, production, testing and assembly schedule is consistent with a Summer 2020 replacement date. Tube replacement process has been studied and requires two weeks, including some contingency.

4.4.4.2 Concerns

- Hamamatsu has been struggling recently to produce ALD MCP-PMTs and have attained poor yields. The cause(s) is(are) thought to be understood, with recent samples from the improved process being tested at the time of this meeting. No QE degradation has yet been seen up to 1 C/cm^2. These tests should be completed in March 2019, and, assuming success, it is hoped will lead to a stable production environment, and timely conclusion to production.

- The machine background rates are not yet understood, but are much larger than predicted before Phase 2.

- Long term survival of the ALD tubes requires successful mitigation of the present (Phase 2) backgrounds.

4.4.4.3 Recommendations

- Understanding and mitigating backgrounds must be a very high priority during the initial Phase 3 running.

- Production and testing of the new ALD MCP-PMTs must be closely monitored, and installation plans carefully optimised.

4.5 Endcap particle identification (ARICH)

4.5.1 Status

The ARICH group installed a new cooling system to address insufficient cooling of the front-end board. Tests after the installation verified that the new cooling system worked as expected although some edge parts showed somewhat higher temperature than the rest of the system due to the lack of cooling plates. This is not an issue since the temperatures of the affected FEBs are still less than 40°C. The heat power extracted by the cooling system is consistent with the power consumption of the electronics.

The ARICH had some issues in 10 out of 72 merger boards during the Phase 2 data taking. The root causes were traced to faulty CAT7 cables for JTAG and timing and connector pins for the FEB-merger data cables. These issues were resolved during the shutdown.

However, a new issue developed after the ARICH was moved into position in Belle II where one merger board suffered occasional difficulties in uploading the FPGA firmware. Also, 4 out of 420 HAPDs (1%) still have issues with high voltage (reduced from 11 during Phase 2).

In the previous BPAC, discrepancies between the data and simulation were observed in the average number of photons and the standard deviation of the Cherenkov angle distribution. One of the causes was traced to the dispersion in the thresholds. The ARICH group plans to improve the threshold adjustments in Phase 3. They also made progress in understanding the pion and kaon identification performance and further reduction of the discrepancies between the data and simulation by including Cherenkov photons from the air gap and the mirror glass, although they still have a long way to go before fully understanding the physics performance of the detector.

4.5.2 Concerns

- The ARICH group has not yet fully verified the physics performance of the detector.

4.5.3 Recommendations

- The ARICH group should verify the physics performance of the ARICH as soon as and as detailed as possible.

4.6 Electromagnetic calorimeter (ECL)

4.6.1 Status

Successful operation of the ECL in Phase 2 was reported in the last BPAC held in October 2018, and further progress on the data analyses and several efforts to fix the problems found in Phase 2 were reported in the review.

Excess noise unrelated to beams was found in 30 endcap channels during the Phase 2 period, and new grounding belts were being installed for all the noisy counters. It is expected that the noise level will be reduced by about a factor of two to three, and the ECL team is now checking the effect of the new grounding in situ.

The beam-background monitoring is one of the important functions of the endcap ECL, and the PMT current is converted to the radiation dose. Relatively large fluctuations of the dose measurements were found in Phase 2. These measurements are stabilised by applying a correction factor depending on the temperature of the ECL system.

Calibration of the ECL is progressing well. In addition to the energy calibration using the mass peaks of neutral mesons, π^0 and η, single photons from the process $e^+e^- \to \mu^+\mu^-\gamma$ are used for both energy and position calibration. Good agreement between data and simulation is demonstrated and further studies to unfold the error of the muon tracking are being performed to understand the ECAL performance better. Timing of the ECL was calibrated using Bhabha scattering events, and unexpected shifts of 8, 4, and 2 ns were observed in some electronics crates. For the time being, such timing shifts are monitored and corrected, but the cause of the shifts should be understood.

A new noise covariance matrix is obtained from the experimental data. It allows reduction of pileup noise, and it will be important to cope with the increased luminosity in Phase 3 and further physics data taking.

4.6.2 Concerns

- The unexpected amount of beam-related noise in the ECL, which was found in Phase 2, is still a concern.

4.6.3 Recommendations

- The ECL group should continue to work on the beam-background with the accelerator team to improve the beam-related noise. It is important to monitor the beam background, since new collimators and beam-masks will be installed in Phase 3.
- Efforts to understand the fine details of the calibration and simulation, in particular of energy and spatial resolution, and of the timing, should be continued to achieve their ultimate potentials.

4.7 K-Long muon detector (KLM)

4.7.1 Status

The KLM group merged the DAQs for the barrel part (BKLM) and the end-cap part (EKLM). The DAQ system works up to 40–50 kHz with a trigger hold off of 100–200 ns and with simplified firmware that does not take advantage of the waveform of the scintillator readout. Human resources were reallocated to accelerate the firmware development. Simplified firmware with buffer is now under test to remove the trigger hold off. Fully functional firmware with waveform analysis is expected to be implemented during Phase 3.

The slow control system for the RPC readout is not complete and still needs more work.

Some of the hardware issues reported in the previous BPAC were solved by replacing cables and power supplies. However, the KLM still has many hardware issues like ageing VME crates, issues with the Trigger-Timing Distribution system and issues with scintillator modules.

The committee is concerned about the current status of the KLM for the Phase 3 run. Integration of the barrel and end-cap KLM teams to work as a single group seems to be working well. But the overall human resources seem to be still scarce. A well-functioning KLM, including its trigger, is essential for the physics programme planned for the early period of Phase 3.

4.7.2 Concerns

- The committee is concerned that the KLM still has many hardware issues and firmware/software tasks.
- The committee is further concerned that working plans are not well focused towards the physics goals of Phase 3.

4.7.3 Recommendations

- The KLM group, together with the Belle II management, is encouraged to carefully examine and prioritise commissioning tasks in light of the detector elements needed most for the first physics results.

5 Online system

5.1 Data acquisition (DAQ)

5.1.1 Status

Overall, the DAQ has worked well during the Phase 2 run, thanks in particular to the extraordinary dedication of the team of core and junior experts. The monitoring of the system availability provides very useful insights. The committee was pleased to learn that the Belle II management is committed to increase the core DAQ (and slow-controls) team with at least two full time equivalents. However, even with this increase the size of the core team for DAQ and slow-controls remains rather small compared to other systems of similar size and complexity. Some of the committee members had an interesting and informative guided visit to the control room where various operation scenarios were demonstrated.

5.1.2 Concerns

- No test has been done confirming simultaneous data acquisition of *all* subsystems at the nominal maximum readout rate of 30 kHz, which leaves the possibility of still having uncovered problems.
- Several sub-detectors have shown robustness problems when confronted with unusually / unexpectedly large size events in very close sequence.
- Human resources of the core DAQ team are still very limited, despite the reinforcement by junior "experts". There seems to be a risk of over burdening the core-experts
- As a consequence of limited personpower, several consolidations and improvements have not been implemented or deferred. Those consolidations and improvements, in the view of the committee, would have a potential to increase the ease and efficiency of operation, as well as the efficiency of data-taking.
- In the case of the SVD, bandwidth limitations in the readout system have affected detector commissioning negatively.
- It was not clear how and where the configuration information for the front-end is stored. There seems to be no central location or generic mechanism for this, even if some sub-systems have shown that they have mechanisms in place.
- There appear to be still robustness issues with the timing and trigger distribution system and some of the CAT7 cables and connections.
- So far, HLT has run (mostly) in pass-through mode. Once it will need to run in selection mode, it will be necessary to have a well-established strategy to validate and check HLT performance. How this would be done was not shown.

- In the Belle II experiment, the data-transfer between online and offline appears to be the responsibility of the offline group. However the offline group reported problems with managing the received data due to issues with the online-managed book-keeping.
- The recovery of the DAQ system in certain situations requires intervention by the experts and these errors may cause serious concerns on the sustainability of the operations, depending on the rates.

5.1.3 Recommendations

- A fully integrated test at nominal rate should be conducted as soon as possible and any limitations found should be carefully traced back to the subsystems causing this.
- DAQ inefficiency should be recorded over delivered luminosity and consistently monitored to identify the sub-system causing the dead-time. Finer grained classification by root-causes within a subsystem is an advantage but not crucial at this stage. A simple "BUSY" monitor should allow the shift-crew in real-time to drill down from a GUI to find out which element of the front-end is causing the dead-time and this information should also be recorded.
- Reproducibility of software configurations is critical for consistent system behaviour and thus eases debugging. Version management should be introduced and enforced for the operating system, DAQ and HLT software.
- Version control and management are a part of modern best practices in software engineering and should be implemented. This should allow, for example, individual developers to contribute to the code using their own accounts, while quality assurance can be implemented by the core experts in their role as release managers.
- Automation, which is crucial for long-term efficiency and for freeing up valuable experts, should be introduced. It is also a way to share valuable expert knowledge. An example is the data-transfer between Online and Offline.
- Recovery and reset procedures should be tested and bench-marked. This will allow to identify weaknesses which could lead to losing data taking time. An example is the reboot of critical servers or Copper crates. Reboots and reset should bring the units into operational state without human or at least expert intervention.
- The committee got the impression that the bandwidth limitations affecting the SVD could be easily mitigated using existing hardware only. In that case this should be implemented as soon as possible.
- The procedures between offline and online for data-transfer need to be more clearly defined and the responsibilities for the implementation should be agreed by all concerned parties.

5.2 Slow-controls, run-control and interlocks

5.2.1 Status

The slow-control and run-control have been demonstrated to the committee. Operation of the entire experiment by non-expert shifters is possible, which is an important achievement. A comprehensive interlock implementation has been presented.

5.2.2 Concerns

- The personpower of the core slow-control team is extremely limited. The team is also responsible for the management of the computing infrastructure and also one of the key developers recently left the team for a new job.
- Documentation of many parts of the system is very scarce, which is a consequence of the limited personpower.

5.2.3 Recommendations

- The committee recommends that the collaboration considers engaging a dedicated technician for system management.
- If additional people can be made available by the collaboration, a fraction of their time should be dedicated to documenting the description of the system and its implementation.
- Wherever possible, duplication of tools and mechanisms should be avoided and uniformity should be aimed for. Examples are environmental and other data acquired by the slow-control system which are used in reconstruction and/or physics analysis. These should be stored in the EPICS archiver and transferred to the conditions data base.
- In the same way as the DAQ, the slow-control system should introduce version control at all levels
- Basic infrastructure management such as rebooting of computers or power-cycling of crates should be made possible through the slow-control.
- A specific part of the previous recommendation is the systematic configuration of IPMI for the management of the server hardware

5.3 Trigger

5.3.1 Status

The trigger has made significant progress since the end of Phase 2 and seems to be ready for Phase 3.

5.3.2 Concerns

- Some important elements of the trigger are not validated yet, for instance, the 3D-track trigger or the TOP timing.
- The personpower situation in many trigger subsystems remains critical.
- The KLM trigger is not ready and there is no input to the central trigger. If this situation persists, this will likely have an impact on the physics.

5.3.3 Recommendations

- For Phase 3, the unusual variation patterns observed in the 2D-CDC trigger performance should be watched for and additional monitoring processes should be introduced, which might help to determine the root-cause if the problem persists.
- Integration of missing trigger components into the slow-control and central run-control should be pursued vigorously
- Integration of the trigger into the data quality monitoring process should continue.
- Getting the KLM trigger to work should be a priority for the KLM team.

5.4 Upgrade of readout system

The purpose of the upgrade is to replace a layer in the DAQ chain currently consisting of the Copper boards and the ROPCs. This upgrade should preserve the following existing functionality:

- data acquisition and formatting,
- front-end configuration (slow control),
- back-pressure detection and signalling (fast control).

It should address problems due to component obsolescence and insufficient performance of the existing hardware in view of possible higher Level-1 rates in order to deal with backgrounds and future improvements in SuperKEKB luminosities.

The committee was presented with a very well worked out and compelling procedure for selecting the best fitting solution out of three proposals. The BPAC is convinced that this will lead to an optimal decision for Belle II

5.4.1 Concerns

None.

5.4.2 Recommendations

- The technical specification and performance criteria of the system for the proponents should be spelled out in great detail.
- All proposed solutions have potential to provide functionalities significantly beyond just replacing the current system. Those functionalities could enhance the capability of the experiment at very high luminosities, for example, by
  - making the readout system completely uniform where the PXD DATCON and ONSEN would be replaced by a software driven ROI selection in the HLT,
  - introducing data preprocessing / reduction in the FPGAs of the future readout-cards to simplify and accelerate subsequent HLT processing for more elaborate event selections.

  Extended capability of the upgraded DAQ system is most likely able to accommodate possible future evolutions of the detector as well, and the collaboration should be aware of this.

6 Offline software and computing

6.1 Offline software

6.1.1 Status

The committee was informed of some minor changes of the software project with new conveners now in place. The status of software release version 3 is in a good shape and deemed ready for Phase 3. It includes the consolidation of the analysis tools such as the combinatorial Kalman filter for SVD to CDC extrapolation for the tracking, and, for the ECL, the energy and first timing calibration methods as well as pulse-shape discrimination and a new track-to-cluster matching for the particle identification. The simulation and reconstruction has been updated for SVD. A long list of bug-fixes and improvements have also been included. Adjustment of the simulation and reconstruction to match the beam background and detector performance observed in real data still needs to be done. Tracking performance in data are within 10% agreement with those in simulation. The collaboration is also improving the overall release process to minimise last minute requests, to be stricter with the validation procedures and to have a better test coverage. Memory and CPU resources needed by the software are continuously monitored and provide input for the computing resource estimates. Version 4 release is in preparation with a long list of features to be deployed.

Work has started to streamline and automate setting up the software and conditions in the online system. Many of the operations, in particular for the full validation, are still done manually and development of automatic procedures are ongoing. The development of the HLT trigger menu is well advanced, suiting well the physics aims of Belle II.

The conditions associated to a set of runs are implemented via the concept of a global tag that associates a list of payloads and their intervals of validity. Work has been done to define the status of the global tags, and developers' roles with the right to modify them. Different procedures for specific workflows, e.g. software development, online processing, offline (re-)calibration, analyses, have also been established

6.1.2 Concerns

- The conditions database infrastructure at BNL is in place and is ready for Phase 3. The use of the conditions database and the management of global tags has been progressing but there are still a number of open questions on how to manage the different workflows, e.g. data taking and HLT, re-processing, simulation data production, etc. Further work is still needed on the software framework tools to correctly interface to the conditions consumers and to handle overlapping intervals of validity.

- The management and procedures for the global tags are better understood and are being put in place but are still far from being very practical for the different workflows. Many of the procedures still rely on too much manual intervention.

6.1.3 Recommendations

- The committee recommends to investigate the distribution of the conditions data base payload files, which are in ROOT format, to be done using the CVMFS file system that is already used for the deployment of the software to all the computational sites, therefore reusing existing infrastructure.

- The committee recommends to focus the priority on the aspects of usability of the conditions at the software interface and at the user level, rather than on the aspects such as authentication and authorisation for the database access.

6.2 Computing

6.2.1 Status

The production system for Belle II is based on the DIRAC system, which provides the basic and common functionalities, while Belle II develops experiment specific features on top, such as the production system, the user interface, the raw data transfer, etc. The collaboration started to evaluate Rucio as the engine for the distributed data management system (DDM). The initial evaluation didn't identify any show-stoppers but the plan had minimal contingency to be ready in time for Phase 3. Since a number of developments need to be done to adapt the Rucio engine with the rest of the system, in particular the integration with DIRAC and the fabrication system, the collaboration concluded that it was not possible to make the transition in time for Phase 3. Therefore the problem of lack of automation for replication and deletion of data files remains. Despite this problem of automation that poses some extra load on the people running the computing system, the production system components are stable and performing well. The Phase 2 collision data was reprocessed multiple times with additional improvements and new features, e.g. changes in the software, modifications to the fabrication system, calibration data format and algorithms, introduction of calibration skims, etc. Several MC production campaigns have been done successfully, ie. MC9 and MC10 completed, MC11 ongoing and MC12 in preparation. There is ongoing work for the preparation of the run-dependent MC production, with samples prepared at KEKCC upon request, and addressing the difficulties of distributing the background samples to the Grid sites.

6.2.2 Concerns

- The data movement and conversion between online (output of the HLT) and offline continues to be a weak link in the full chain: the protocols are not fully robust since discrepancies in number of triggers and data sizes have been observed; many operations are done manually, and the format conversion times are very long and not done in parallel. Therefore, if the data rates increase with the expected increase of luminosity, it may become a serious problem.

- Readiness for the run is based on MC production and Phase 2 experience. The computing group will have a large operations load in Phase 3 since automation is missing in many areas. There is a concern that the operations load will severely limit the resources available for development for the automation of various tasks, for example.

- The storage at the sites have been close to their full capacity and the collaboration does not have an efficient way to delete unneeded files.

- The problems with the KEK storage element was worrying, but the computing group seem to have found a workaround that increased their dependence on other sites. The experience of the BNL team should reduce their risk.

6.2.3 Recommendations

- The committee urges the collaboration to develop an automated procedure for the RAW data copy from online to offline, a set of monitoring and data registration tools, and means to perform the data format conversion in parallel.

- The computing resources are pledged based on number of authors. The sites vary greatly in size and probably in level of support. The collaboration should consider implementing some sort of site monitoring and some standard of availability - at least for the sites needed for production activities.

- The committee is pleased to note that the collaboration has recently implemented a CPU accounting system based on information extracted from the EGI portal or, for the sites not exporting to EGI, from the KEK Dirac portal. The CPU power used in each site has been reported and compared with the CPU provided, defined as the sum of pledged and opportunistic resources, with a monthly breakdown. This accounting system provides a detailed breakdown of each site's performance, as long as the CPU usage is concerned. However, the choice not to group the sites using the Belle II hierarchical site classification (Raw Data Centres, Regional Data Centres and MC Production Centres) and without any distinction between sites offering pledged resources and sites offering only opportunistic resources could be a source of confusion. The committee recommends to provide separate accounting information for sites offering pledged resources using the Belle II site classification as is the case for Tier0, Tier1 and Tier2 LHC centres. Furthermore, the amount of pledged resources provided by each site or federation of sites should be explicitly indicated.

- We recommend to increase the priority on evaluation and development of the missing parts for integrating the Rucio tool as the engine for the DDM of Belle II.

- According to the evaluations and the experience from the Rucio experts, the tool seems to be suitable for Belle II. The committee judges that the use of a tool like Rucio will ensure better sustainability because already in use by several experiments and under evaluation by others. Therefore, if the final goal for the distributed data management is to use Rucio, the collaboration should not be investing heavily in evolving the incomplete current solution. Only minimal developments should be required to keep working for Phase 3 and migrate to the new tool as soon as possible afterwards, and eventually run both systems in parallel to ease the migration.