B-factory Programme Advisory Committee
Focused Review Meeting

25–26 June 2017 at KEK

D. Cassel (Cornell), M. Frank* (CERN), N. Neufeld+ (CERN),
B. Ratcliff* (SLAC), M. Sullivan (SLAC), H. Tajima (Nagoya),
and chaired by T. Nakada (EPFL)

* Expert members

+ Participated remotely.

Full Report, 7 September 2017 (Rev 8 Sept 2018)

1 Executive Summary

The main purpose of this meeting was to review the readout integration status of the Belle II experiment. In addition, recent incidents for the two subsystems, i.e. the Silicon Strip Vertex Detector (SVD) and barrel particle identification system (TOP) were reported. This short summary with essential conclusions of the review is followed by a full report providing more comprehensive information. It should be noted that the review for the integration of the Vertex Detector System is planned to be held later in October, 2017.

Readout Integration Status

The back-end implementation of the systems related to the data handling of the Belle II experiment, such as the Data Acquisition (DAQ) and run control, condition and configuration database, high level trigger and data storage, detector control and interlock system and detector and data monitoring, are well in place and ready to be used by the subsystems for their commissioning work. Although further optimisation and refinement are needed, all the necessary functionalities for the Phase 2 and Phase 3 data taking are in place. The central DAQ framework should ensure stable running up to a rate of 30 kHz so that the commissioning and debugging work of subsystems can be carried out smoothly to reach the design performance. The committee appreciates successful effort by the central DAQ team to attract outside groups for well identified tasks, such as the high level trigger and data quality monitoring. This has allowed the central DAQ team to concentrate on the core issues and to provide essential expert advice to all the subsystem groups. The committee has no major concern with the overall progress to date.

Readout integration of the subsystem is at various stages. All the subsystems, except the end-cap particle Identification system (ARIC) and the end-caps of the K-Long muon detector, can now be included in the Belle II central DAQ system and they participated in the global cosmic ray data taking tests. The electromagnetic calorimeter and central drift chamber system (CDC) are in particularly good states. While the central DAQ team must ensure that the commonly used Belle2links and COPPER readout boards are flawless, the subsystem teams should continue their effort in the firmware development of their front-end electronics. Guidance to the subsystem teams provided by the central DAQ team is crucial. Sharing experience among the subsystem groups is a good practice and effort being done along this direction is appreciated. The committee is looking forward to hearing about further progress on the integration of subsystems in future meetings.

The first level trigger is affected by a stability problem in the high-speed serial links, which currently does not allow the links to transfer all required data for the CDC trigger. Finding a solution may require very specialised knowledge on the FPGAs used in the system and consulting the vendor could be useful. In parallel, the amount of CDC data required for the triggering could be revisited.

Concerning the delamination of PEEK frames of the TOP modules

A problem with the unexpected movement of photon detectors (PMTs) for the barrel particle identification system (TOP) due to the solenoidal magnetic field of the Belle II detector was first reported to the committee during its meeting in October 2016. This problem seems to be now well understood and under control. The vacuum tightness of the PMTs is also shown to be unaffected.

On the other hand, apparent delamination of the PEEK frames glued to the prism, newly reported in this meeting, is alarming. Not only is the cause of this phenomenon unknown, the picture taken by the CCD camera mounted inside of the module does not provide unambiguous identification of the delamination spots and their extent. There is some evidence that the optical silicone oil has entered into the gap caused by the delamination which makes it difficult to identify the delamination area and to know exactly when it occurred. This also generates a concern that the oil may have entered into the prism side of the module and that the bar box is open to the electronic readout side. In this case, the quartz bar box may be exposed to the external atmosphere and to the oil. The committee thinks that the highest priority for the TOP group is to understand the exact situation of the delamination. They must develop a reliable method which can unambiguously identify the affected areas and determine how far delamination has progressed. They should also examine whether the quartz bars are now exposed to the external atmosphere. Analysing data on the N_2 flow, i.e. the flow rate, its composition and especially dew point, might provide further information. By correlating them with the apparent moment of delamination, one might gain more insight on the process and the degree of contamination inside the module through the delamination. In parallel, available prototypes and purpose made mockups, including Module 01 when appropriate, should be used urgently to understand the cause of the delamination, e.g. whether it is due to external forces, which may or may not be due to the detector magnetic field, thermally induced forces from the warm readout side or intrinsic to the design and material choices. The effect of contamination due to the external atmosphere and the oil vapour on the quartz bar surface should be investigated, as well as the effect on the physics performance. Although it is still premature to launch any direct intervention on the detector now, every effort must be made to avoid any possible further damage and possible repair methods should be worked out and practiced.

Concerning the delamination of Kapton flex cables of the SVD

In this meeting, reoccurrence of delamination of the Kapton pitch adapters was reported. The forward module of a newly made Layer-4 ladder again developed delamination of pitch adapters. Further reinforcements are introduced and tests show very good results. In order to better identify infant mortality, the additional waiting period is added to the production procedure, which is also re-optimised to minimise the delay. The quick responses by the SVD group to the problem and all the actions being taken are fully appreciated. However, the committee strongly recommends performing full thermal cycle tests of the reinforced ladders. For this purpose, class-C modules should be sufficient.

2 Status of the Belle II Central DAQ & Control System

2.1 COPPER, Belle2link and Readout Hardware

2.1.1 Status

Combined runs with CDC, TOP, ECL and Barrel-KLM have been successfully conducted. While succeeding with those runs was a big step forward for which the collaboration and in particular the central and detector DAQ teams are to be congratulated, these tests also revealed a number of outstanding issues. The following problems have been reported to the committee.

1. Spontaneous loss of the Belle2link: Although, this has been observed in several subsystems, neither common cause nor common circumstance related to the loss could be identified.

2. Mismatch of event numbers at the input to the COPPER board: This seems to be caused by the front-end firmware that is not yet ready for high-rate (30 kHz) running and was observed by TOP and KLM. Evidence that the problem is related to the rate is that the problem can be cured by introducing a (significant) trigger latency of 10 ms.

3. Data corruption in the receiver of the Belle2link (HSLB) or in the FIFO in the COPPER.

2.1.2 Concerns

While some problems are clearly caused by the individual elements of subsystems, such as the link-loss for the CDC that was shown to be related to three specific front-end modules, there are some issues appear only in global combined running. Examples are the link-loss in the TOP and the data-corruption in the COPPER that is a generic part of the readout system. These global issues have to be dealt by the central online team, which is quite small.

2.1.3 Recommendations

- Global combined runs for an extended period must continue even if work-arounds are needed. Long runs are important to uncover not yet identified problems which might exist. Moreover, making long runs stimulates automatisation of procedures, which will be beneficial for the future Belle II operations.
- It appears that many versions of firmware are used in parallel without documentation describing all the characteristics of different versions. Such inventory should be created with any of the suitable standard software tools already used by the collaboration.
- For the data corruption, it is important to make sure that this is not a hardware issue. The current worst-case of 0.2‰ corrupted event rate per COPPER is still high.
- The timing distribution with FTSW must be stress-tested to make sure that the link-loss problem does not come from there.

2.2 Event Builder

2.2.1 Status

The first stage of event building (event builder 0) is performed in the COPPER boards and consists of a merge of maximal four input links into one output stream. This setup was tested extensively both when operating individual subsystems and in the global mode during the subsystem integration. During the last review some problems started to manifest themselves, which at the time were thought to be relatively simple to be resolved. Unfortunately this is not the case and details have been already discussed in the section above.

The second level of event building (event builder 1) happens at the input stage of each High Level Trigger (HLT) worker node and assembles the data originating from one particle collision of all subsystems except PXD. The final size of the HLT with 6400 cores or $O(100)$ physical CPUs gives us an event fragment rate of 90 kHz per node.

For all events accepted by the HLT filter, the actual raw data are sent to the storage layer and a descriptor containing the decision is sent to the ONSEN box, where a temporary buffer for the PXD data is waiting for the HLT decision to either be discarded or sent to the 3^{rd} layer event building in storage layer.

The third level of event building takes place in the storage layer when merging the HLT accepted events received from the HLT facility with PXD data received from the ONSEN box.

The requirements for the event builder 1 seem to be met with reassuring margins. For event builder 1, a system with 1/4 of the envisaged size was able to copy data with a rate of 1440 MB/sec where 3 GB/s is the design value of the complete system.

The third level of event building was tested earlier this year during the VXD test using the DESY test beam where numerous tests were performed, which confirmed the feasibility and the expected performance of the chosen architecture. From the VXD test beam experience, the committee concludes that the chosen solutions are appropriate.

Performances of the second and third level depend strongly on the number of the HLT hardware racks and the performance of the HLT software. Since both evolve according to expectations (please see Section 2.5) no problems in event building should arise. In particular the experience gained with the VXD test beam activities gives confidence that the system can be extrapolated by a factor of ten to reach the required scale.

2.2.2 Concerns

- The committee sees problems within the area of the event builder 0, which is likely originated from the COPPER.
- Although the VXD test with the DESY test beam was successful, the system still needs to be scaled up by a factor of ten for the final system, which is not a small factor.

2.2.3 Recommendations

- It is absolutely essential to understand the problems with the COPPER boards, which have direct consequences in the first event building step.
- There is no immediate action required, but the VXD and HLT subgroups should stay aware of possible scaling problems.

2.3 Run Control

2.3.1 Status

The run control is responsible for the coherent configuration and steering of all elements of the Belle II detector participating in the data taking process. These elements may be hardware based such as individual subsystems or sub-components thereof or software components such as the HLT. In a simplified picture, this configuration must be able to handle the following discrete steps:

1. First, the desired configuration is retrieved from the DAQ database.

2. Then, the hardware is configured according to the retrieved parameters. In the data-flow, data sinks are prepared and activated.

3. Data sources are prepared and connected to the corresponding data sinks.

4. The trigger is enabled and data taking starts.

During the previous BPAC-DAQ review, the individual finite state machines of the readout components and its steering entity and the run control was discussed in detail and updated during this review. In the meantime, the local flavour of the run control was tested successfully during various occasions notably during the test of the VXD system at DESY.

A real-time demonstration of the functionality of the global run-control in action was given in the Belle II control room. The occurrence of a DAQ error condition followed by the appliance of a proper recovery procedure was successfully demonstrated.

2.3.2 Concerns

- The DAQ errors demonstrated would unlikely be the only possible failures of the readout system.

- The current graphical design is very complete and well suited for experts steering the data taking activity of the experiment. It was not clear if such an elaborate interface is also well suited for less experienced shift crews.

2.3.3 Recommendations

- Extensive testing of the DAQ system should continue and be an essential part of the subsystem integration procedures.

- Development of the DAQ user interface should also take less experienced regular shift crews into account.

2.4 Online Databases

2.4.1 Status

There are four major databases, which provide the basic information to steer and configure the experiment hardware components and the corresponding software components such as the HLT. They also provide information obtained during the data taking activity necessary for the later interpretation of the event data:

- DAQ DB. Repository of readout configurations. Accessed by the run control to configure subsystems. This database is interfaced to the run control software. The data are used to steer the data taking activity and command line tools which allow users to extract this information from the database exist. The committee concluded that the DAQ DB is a well designed module describing subsystems.

Since this should be a rather small database $O(10 \text{ GB})$, no problems are expected. Though enhanced testing during the commissioning phase of the subsystems will help to identify possible hidden problems.

- Conditions DB. This database is used by HLT and the Express Reconstruction (ER). In the online system this database is read-only and entirely imported as a snapshot.

- Run/File DB. A repository of all runs/files created while taking data in the experiment. This is logically a part of the data storage. The database is implemented and interfaced to the storage applications. Users can retrieve information from the database either using a web based GUI or a command line application.

- Logger DB. The purpose of this database is to save logging messages of all components participating in the data taking activity for later analysis to understand the detector performance and data taking environment. The components populating this database are the run control, the front-end hardware as well as software based components such as the slow control toolkit based on EPICS or the HLT using the data processing framework basf2.

During the live demonstration, the committee could see all these databases in action including a display of the logger database.

2.4.2 Concerns

During the previous review, the committee expressed some concerns about the performance of the error logging. In this review, this point was not addressed.

2.4.3 Recommendations

For the performance of the Logging DB, it would still be useful to have an estimate of possible error message rates under anomalous circumstances such as power cuts etc., as well as the behaviour of the database access over the years with growing database table sizes. At least a qualitative assessment should be made.

2.5 High Level Trigger, Storage, Express Reconstruction and Data Quality Monitoring

2.5.1 Status

The committee congratulates the HLT and VXD subgroups for the successful operation of the HLT and VXD subsystems during the VXD test with the DESY test beam.

The test showed that the chosen architecture can provide the required functionality and the implementation is flexible enough to relatively easily reconfigure individual components such as the data quality monitoring to copy with unexpected loads.

The test at DESY confirmed the feasibility of merging the data of the PXD subsystem with HLT accepted data in the ONSEN box using regions of interest obtained by the real time tracking in the other subsystems. The results showed that the implemented communication protocol among the HLT, the ONSEN box and the storage works. The fact that the data quality monitoring was attached and could be tuned to match the needs gives confidence that the components are collaborating as expected and are flexible enough to cope with an increasingly hostile processing environment. A 16 hour long run at a trigger rate of 1.6 kHz showed that the HLT works reliably and – given the processing rates presented – sufficiently fast. The committee appreciates the attempt to control different versions of the HLT configuration using a software versioning system, git. Such an approach is necessary to ensure reproducibility in an environment where many developers are working concurrently. Such a system also helps to understand the HLT conditions and behaviour a posteriori.

The committee believes that the data transport protocol between the storage and the Express Reconstruction cluster shall not pose problems, since the same protocols and software are used for the data transportation from the HLT to the ONSEN and from the HLT to the storage. The disk throughput for one storage unit of roughly twice the required data rate of 300 MB/s gives a comfortable margin on the hardware side.

The assembly and installation of the remaining missing HLT units develops according to plans. Currently, three HLT units are operational and two new units will become operational soon.

2.5.2 Concerns

A concern on the data format for raw-data files was expressed during the last review that the ROOT-like files, which do not contain streamer information, are not formally supported by ROOT and may pose a compatibility problem in the future. No presentation was given on this point during this review.

2.5.3 Recommendations

A clear strategy for maintaining the compatibility of the stored data with the future ROOT versions should be developed. In the currently foreseen approach, all raw data will be converted to ROOT trees in the KEK computer centre for the further processing, but kept for possible future needs. In this case, re-readability of the original raw data should be ensured.

2.6 Trigger

2.6.1 Status

The trigger hardware installation of CDC, ECL and KLM is already very advanced and that of TOP is progressing well. Verification of function is ongoing. The firmware validation of the full subsystem trigger components is progressing rather slowly and there are a lot of FPGA technical issues, such as timing constraint violations and logic resources, to be looked at.

Also an issue with the optical link was discovered that currently limits the amount of data from the CDC to the trigger. The impact on the trigger performance, which this would have if the problem could not be fixed, remains unclear.

2.6.2 Concerns

- Most of the firmware issues are probably related the manpower situation, since FPGA experts are overly booked in most subsystems. Training of new experts takes time and this expertise will still be needed until the experiment is fully commissioned
- The optical link issue seems to come from a rather specific use of the link, which may well require support from the vendor to solve

2.6.3 Recommendations

- Training of additional firmware experts should continue. A good example is the TOP as reported in the present review. Sharing of information among different subsystems should be encouraged as much as possible.
- For the link issue, the support from the vendor (Xilinx in this case) should be sought.

3 Detector Operation

3.1 Slow Controls, Monitoring and Interlock

3.1.1 Status

The slow control, monitoring and interlock systems consist of two frameworks, NSM2 and EPICS, running on a common network dedicated to slow control and monitoring. NSM2 is the primary program for detector monitoring and run control and EPICS is used by the slow-control of the VXD. The systems are connected to each other and to the accelerator control system (also based on EPICS) using gateways.

The run-control and GUI are based on Control System Studio. The committee was very pleased to hear reports and see in action the combined operation of several detectors using the control-system. All important elements are in place and are actively used by the collaboration. There are now numerous contributions by subsystem experts outside the small central core-team.

The central interlock system is also well advanced. The major components are designed and are working. New interlocks will be now added as needed. The alarm system will be ready for the cosmic ray runs.

3.1.2 Concerns

- There are still several interlock and alarm sources, in particular from subsystems, which are not fully developed or not connected.

- Feedback to the monitoring system seems to be missing or underdeveloped
- In spite of the significant progress, the committee did not receive a detailed list of replies to its previous recommendations for the the interlock system.

3.1.3 Recommendations

- Subsystem alarms and interlocks should now be integrated as quickly as possible. This is very important for the safe long-term operation of the detector.

4 Status of the Subsystem Readout Integration

Since the committee's review of Belle II progress in February, there have been two major highlights. The first was the roll-in in April. The second was the successful operation of the CDC, TOP, ECL, and BKLM in a barrel Global Cosmic Ray Test (GCRT) using the global DAQ system. The committee congratulates the collaboration on achieving these two vital milestones.

Test runs have exposed some front-end electronics (FEE) issues and some issues with the integration of individual subsystems with the global DAQ. The DAQ issues or their causes tend to be somewhat different for each detector. This section is devoted to discussion of the results of and lessons learned from the integration of subsystems with the global DAQ and plans for addressing the issues.

4.1 CDC

4.1.1 Status

Following the roll-in of the Belle II detector, cosmic ray CDC data were taken with a magnetic field of 1.5 T, complementing the earlier data taken without magnetic field. The CDC was integrated with the global DAQ system for the runs with and without magnetic field.

Significant progress has been made in calibrating CDC position resolution with and without the magnetic field. Reconstruction resolution and accuracy were studied by considering the segments of a cosmic ray track on either side of the centre of the CDC to be two separate tracks. Preliminary results from comparison of their impact parameters and coordinates on the z axis are gratifyingly close to the results obtained without magnetic field and from Monte Carlo estimates. These results indicate that the FEE and the DAQ integration are generally working very well.

Low-rate (< 10 Hz) CDC runs triggered by a scintillation counter in the VXD region tended to last more than a few hours without difficulties. Higher-rate (> 100 Hz) runs triggered by the CDC and ECL tended to last only about 10 to 20 min, due to "b2llost" errors and data corruption. The b2llost errors (where the Belle2link is lost or disconnected) have been traced to 3 FEE boards. When cables are switched, this problem occurs with the same boards. These boards can be replaced during a short shutdown in September. There has also been a problem with high voltage in layer 19, which will be investigated at the same time.

The data corruption problem may be in the Belle2link receivers (HSLBs) or COPPERs. A recent HSLB firmware update may resolve this problem.

4.1.2 Concerns

The b2llost failures of 3 FEE boards may be indicative of marginal failure that could eventually affect other FEE boards as they age.

4.1.3 Recommendations

Ultimately understanding the reasons for the b2llost failures is as important as replacing the 3 boards that have failed.

4.2 TOP

4.2.1 Status

The TOP is fully integrated with the global DAQ, and was included in GCRT runs, with and without magnetic field. A total of 61 of the 64 board stacks are fully operational, and the remaining 3 are masked currently. A timing issue that can be resolved with tuning is thought to be the source of the problem with the 3 board stacks.

Two principle limiting issues with the current firmware were identified in the GCRT runs and workarounds were installed. A minimum trigger interval of 1.5 ms was introduced as a temporarily workaround for high trigger-rate issues. These rate issues should be resolved in future firmware upgrades. DAQ throughput and storage were overwhelmed with ~ 1 MByte event sizes, since all channels were reporting a minimum of 60 Bytes, whether or not there was a signal. This problem was resolved by introducing 4 Byte headers for empty channels.

TOP data are also affected by some instabilities in the Belle2link connections. The source is thought to be incorrect global commands in FEE firmware and debugging efforts are ongoing.

The committee is pleased to learn that, despite these issues, several 100 k events were recorded in the GCRT runs. These events provide data for developing calibration and reconstruction software.

Substantial firmware updates based on experience in standalone and GCRT runs are underway, and milestones with dates for crucial issues have been established for getting the TOP ready for the Phase 2 run.

Firmware development for the TOP, as well as for other subsystems, is limited by the availability of firmware programming experts. To begin to address this problem, a TOP Firmware Bootcamp with 20 participants was initiated the week of the review. A TOP hot spare module has been set up to provide a platform for local firmware tests. Although the short-term benefits are not clear, this program is clearly a step in the right direction.

4.2.2 Concerns

- TOP firmware is complex and many issues still require resolution for reliable Phase 2 running.
- The firmware milestones established for Phase 2 running are very important, but achieving them may be challenging.

4.2.3 Recommendations

- Experts with knowledge and understanding of TOP firmware should devote as much effort as possible to realising the development milestones to ensure that the TOP will be ready for Phase 2 running.
- With data starting to arrive, significant effort must be made promptly to ensure that all critical calibration software is operational.

4.3 ECL

4.3.1 Status

Integration of the barrel ECL and the backward endcap ECL into the global DAQ is complete. Installation of the forward endcap ECL awaits ARICH completion, but DAQ integration will be tested sooner. The committee is very gratified with the ECL performance described at this review.

All 6,624 counters in the barrel ECL are active, although 6 have only half signals due to diode or preamplifier failures. Only one of the 960 backward endcap counters gives no signal and this issue can be investigated in the September shutdown. Noise levels of all counters are acceptable, although noise levels in 7 barrel and 3 endcap counters are higher than desirable.

There were some problems with event sizes received by the HSLBs. These were resolved with firmware updates. Now operation of the ECL in tests is quite stable. Cosmic ray runs at ~ 1 kHZ last longer than "overnight," while runs with a dummy 30 kHz trigger last more than a few hours.

Preliminary results from energy and time calibrations with cosmic ray data are encouraging, indicating that the counters, FEE, and integration with the global DAQ are all working well.

4.3.2 Concerns

- There are no known concerns with the barrel and backward endcap ECLs, since ECL/DAQ integration issues identified so far have been resolved.
- Integration of the forward endcap ECL with the global DAQ has not yet occurred.

4.3.3 Recommendations

- Continue monitoring the ECL to ensure that data corruption issues found in other detectors will not affect ECL performance at some low level.
- The planned integration and testing of the forward endcap ECL should be completed as soon as possible in order to have time for addressing any issues that might arise.

4.4 KLM

4.4.1 Status

For the GCRT, enough RPC FEE boards were available to populate nearly all of the RPCs in six forward octants of the Barrel KLM (BKLM). Scintillator FEE boards were installed previously. The scintillators and RPCs are connected to Data Concentrators, which include RPC event builders and RPC/Scintillator combiners, as well as slow control. The results of these modules are passed to the global DAQ and the trigger. In the GCRT, tracks in the BKLM matching tracks in the CDC were observed.

An earlier test run with only the ECL lasted for more than 7 hours. However, later runs have generally lasted only a few tens of minutes due to "fifoerr" and "ttlost" errors. Fifoerr errors occur when data are not sent from the FEE and ttlost errors occur when a b2tt (trigger) link is lost. It appears that fifoerr and ttlost errors are due to trigger intervals that are too short for the current BKLM firmware.

DAQ integration of the endcap KLMs (EKLMs) was setback by a mechanical problem, the vertical scintillators in two modules sank under their own weight and pushed their preamplifier boards into the frame. These two modules were replaced before roll-in and now all EKLM channels are active.

However, attempts to readout EKLM hits have not been successful, although the readout architecture and firmware are identical to those of the scintillators in the BKLM. The source of this problem was not understood at the time of the review.

4.4.2 Concerns

- Identifying the ultimate firmware origins of the fifoerr and ttlost errors and eliminating them is crucial for the high-rate operation of the BKLM.
- Understanding and resolving the problem with reading out the EKLMs is an important element in the DAQ subsystem integration.

4.4.3 Recommendations

- The committee endorses the ongoing effort of the BKLM FEE experts to address the outstanding firmware issues.
- A serious effort to understand and resolve the EKLM DAQ integration problems is required.

4.5 ARICH

4.5.1 Status

The committee was pleased to learn that substantial progress has been achieved on outstanding issues in ARICH construction and DAQ integration.

Mechanical tests on a HAPD indicate that magnetic torques due to the presence of a magnetic material (Kovar) in the HAPDs are unlikely to pull them out of the sockets of the front-end boards (FEBs). Studies reported in February indicate that a 1.5 T magnetic field could produce torques as large as the equivalent of about a 1 kg mass hanging on one end of a HAPD. In an extended mechanical test, the force of a 17 kg weight applied to a HAPD for a month did not pull the HAPD out of its sockets.

The long-awaited short inner HV cables arrived in April and installation is now underway. Final delivery of the outer HV cables should occur in September. The HV power supplies have been ordered. The delivery of the first is expected momentarily, half should be available in September, and the rest in November.

ARICH FEBs are connected through Merger Boards (MBs) to Belle2links and the DAQ system. A problem with high trigger rates reported earlier was traced to a FEB firmware problem. When the time interval between two events was too short (less than the 2.6 μs length of the data stream of an event) the two events were merged into one event of twice the length and sent to the MB. This problem was fixed by a firmware update and, in a test of six FEBs coupled to a MB, data rates as high as 50 kHz were sustained.

Issues with trigger synchronisation between the FEBs and MBs have been addressed and handshaking between the FEBs and the trigger has been upgraded. After these improvements, a system with 4 MBs integrated to a COPPER in the global DAQ was tested in a cosmic ray run. The system ran stably for more than 12 hours at a rate of about 0.1 Hz. A number of clear ring images were observed with about 13 photo electrons and Cherenkov angles near 17°, consistent with muons in the momentum range of 0.5 – 4.0 GeV/c.

Preparations are underway for the next step in testing DAQ integration. This will be a test of a sector that is 1/6 of the ARICH with a copy of the global DAQ. Beyond that there is a schedule with milestones for completing construction and DAQ integration of the ARICH in time for Phase 2.

4.5.2 Concerns

- So far, integration of the ARICH with the global DAQ has not been possible, and tests have included only a few FEBs and MBs. The full sector test being prepared is the next crucial step in verifying ARICH/DAQ integration.

4.5.3 Recommendations

- Achieving full sector test as soon as possible should be a very high priority for the ARICH group.

4.6 VXD

4.6.1 Status

The committee is gratified with the successful test of the Phase 2 slice of the PXD and SVD in an electron beam at DESY. This was a major milestone in VXD integration. A pocket DAQ was used for SVD readout and an ONSEN chain was used for the PXD. This system functioned well. The SVD reconstruction software provided tracks, which mapped into regions of interest (ROIs) in the PXD. Then, these ROIs were successfully matched to PXD hits.

In principle, the only firmware item missing for Phase 2 running is a high-rate trigger veto in the master FTSW for 30 kHz operation. This veto will emulate FIFO occupancy and veto triggers to prevent overflows. It is currently being developed.

The next major steps for Phase 2 operation are: testing VXD/DAQ integration without detectors at KEK; installing, commissioning, and testing the VXD slice; and final tests of the VXD/DAQ integration with the detector. The schedule for this effort leads to completion in December.

Completion of Phase 3 VXD readout hardware is largely paced by fabrication and testing of the V4 FADCs for the SVD. Samples of these new FADCs with improved noise performance are due momentarily in HEPHY. After acceptance tests, the comparisons of V3 and V4 performance will occur at KEK. If all goes well, mass FADC board production should be complete by December, with final testing and delivery to KEK completed by March 2018. The latter date coincides with the anticipated completion of L6 production in the same month.

4.6.2 Concerns

- There do not appear to be any issues that are likely to affect the readiness of the VXD slice for Phase 2 operation.
- The schedule for completion of Phase 3 SVD readout hardware is tight, matching the anticipated completion of the SVD ladders.

4.6.3 Recommendations

- Progress in producing the Phase 3 SVD readout hardware should be monitored closely, and preparations for installation and integration with the SVD ladders should be well-organised.

4.7 Hardware Monitoring and Slow Control

Hardware monitoring and slow controls for individual detectors are in a variety of states of installation and development. The success of the GCRT indicates that the most essential components are working at least under expert supervision.

The CDC environmental monitoring and interlock system are well advanced. CDC water leak, temperature, and humidity monitors are connected to the local CDC interlock, and displays of their statuses are available for monitoring and control. The major components that are missing are the monitoring of the status of the FEE low voltage supplies and the the high voltage status.

Environment monitors for the ECL detectors are in various stages of development. For the barrel, temperature and humidity monitoring hardware installation is finished and software development is in progress. The same monitoring for the backward endcap is ready and running. Temperature and voltage monitoring for the VME electronics are working. For the voltage and current measurement of the bias supplies, hardware installation and software development are underway.

Integration of the TOP slow control into the global Belle II slow control framework is being developed. The system was tested locally and global tests are progressing.

Operation of the slow control of the high voltages and bias voltages for a section of the ARICH is being tested. The hardware interlocks are not yet ready.

High voltage control and run control hardware of the BKLM are ready. Some software items, including interlock and user interfaces are still being developed.

The state diagram of the slow control for the SVD DAQ has been developed and implemented and the necessary GUIs are ready. Network and database development remain to be finished.

4.7.1 Concerns

- From the presentations, it is not clear that monitors and interlocks for all conditions that can lead to damage of the detectors are in place.
- It is not clear how ready individual detectors are for operation without expert monitoring and control.

4.7.2 Recommendations

- It is essential that all monitors and interlocks necessary to prevent detector damage be operational by Phase 2.
- By Phase 2, slow control and monitoring software should be developed to the level that subsystems can function well under the supervision of non experts.

5 Particular Detector Issues Discussed

5.1 Barrel Particle Identification System

5.1.1 Status

TOP module installation was completed nearly one year ago. Since that time, three mechanical problems have been observed that raise substantial concerns, so they were discussed at this "readout integration status" review as special issues of interest.

The first two arise due to the use of a magnetic material (Kovar) in the construction of the PMT tubes, which can lead to a rotation of the PMT modular units, and the individual PMTs, under the influence of the detector's magnetic field, as reported in October 2016. These rotational forces lead to the possibility of performance loss due to optical de-coupling, and a variety of mechanical concerns, including putting forces on the pins or potting that could break the pin seal or the PMT body to window indium seal. These problems have now been studied in depth in mock-ups and R&D, are quite well understood, and appear to be under control. The performance concerns due to the possibility of optical decoupling had already been demonstrated to be modest in February using Monte Carlo simulation. Further confirmation of this conclusion can be obtained soon in Phase 2 running. A substantial number of studies that put forces on tube elements similar to those expected in situ due to the magnetic fields has yet to demonstrate any vacuum failures. The tentative conclusion is that the vacuum seals are rather robust, and repairs are not urgent. Of course, this is a situation that bears watching as there could be an as yet unseen time element in play. Two repair methods have been developed that can mitigate this issue in part, but they will not be invoked unless performance degradation is observed in Phase 2 running.

The third, much more recent concern (in the last month before the review), is more alarming. Photos taken with cameras located within the bar box modules show evidence of at least partially decoupled PEEK to fused silica wedge joints on nearly all modules. The joints appear to partially open and close with time, perhaps with temperature changes, or mechanical detector activities like detector roll-in or magnetic field. It seems plausible (though uncertain) that open joints fill with the optical coupling fluid via capillary action. Assuming the conventional explanation that there is really decoupling followed by filling in with oil, as seemingly demonstrated by the photos, it seems likely that many if not all joints have delaminated at least in part. One can even see some evidence of small bubbles in the refilled joints. The mechanical strength of the bar/module system is probably OK, in the sense that the readout side is mechanically stable with respect to the optical side, and the forces are not too large. However, the frame to wedge seal, if broken, allows a leak path from the dirty readout region into the very clean bar/wedge region. It is unclear that any joints have fully failed yet in depth, nor that leak paths exist as yet from the readout side to inside the box. We were not shown photos of all joints. More photos may give worse news. The gas system instrumentation (especially gas flow, pressure, and dew point on a module by module level) should also provide further information about leaks. Optical oil is likely to be a major problem if it enters the box, but even air entering from the readout side is quite problematic over a long time period. Leaks into the optical side must be avoided!

It is unknown at this point how strong or compliant the glue joint is between the PEEK frame and the wedge. But since many joints appear to have failed at least in part, there must be an overall design issue here. Changes in the PEEK to wedge coupling have been noted between mechanical events like magnetic turn-on or detector roll-in, but these events may also be associated with thermal cycling due to the warm readout side switching off and on. Unfortunately, PEEK has a much larger thermal expansion coefficient than that of fused silica, so the forces across the glue can be large, and the thermal gradients across the PEEK interface in depth can also be large, since the module quartz side is cool. This could be saving the situation at the moment.

Most of the concerns and recommendations noted here have been discussed in the short summary and somewhat more throughly explored informally in a note to the proponents right after the review.

5.1.2 Concerns and Recommendations

- Try to keep the system stable magnetically, mechanically, and especially thermally to prevent the system from further damage.
- As immediate mitigation, raise the nitrogen flow rates on the modules to help withstand impacts from small leaks pending better understanding.
- Analyse the nitrogen flows in the bar boxes, pressure, flow rates, and especially dew point on a continuous basis to search for leaks, if any, and understand changes as they occur. If this is not possible or requires too much intervention with the present hardware, upgrade the system so that these data can be recorded and studied continuously.
- Expand efforts to understand what is happening to the PEEK/wedge joints on the modules in detail. This should include understanding the photos better, and having a consistent program to follow the situation of joints from photos with time and changes in conditions, so that these can be correlated.
- In parallel, construct mock-ups with wedge/PEEK joints, mocked-up thermal input and loads, and temperature measurement to understand the joints and their failure mechanisms.
- Use available prototypes whenever appropriate to understand the cause of the delamination.
- If moving optical oil is implicated, it is important to understand how much impact it could have on the bar reflectivity. It is not as time urgent as trying to understand and prevent the delamination from occurring. It can be studied with a scalable reflectivity test using a prototype bar in a box with a gas flow system and a reflectance measuring system.
- Collect and analyse cosmic ray data as soon as possible to be able to compare with simulations and understand the number of photons as a function of meaningful variables such as photon dip angle, length down the bar, and bounce number.

5.2 Silicon Strip Vetrex Detector

5.2.1 Status

The SVD group successfully built one additional good ladder for layer 4, six additional good ladders for layer 5 and five additional ladders for layer 6, since the last review in February. In particular, IPMU reduced the ladder assembly time to 18 days from 21 days by optimising the shift schedule. In addition, the ladder throughput is reduced to

16 days by staggering the last two days of preceding assembly and the first two days of following assembly.

However, the ladder assembly encountered another incident of pitch adapter delamination from the forward (FW) sensors of a layer-4 ladder even though a glue reinforcement was in place. There was no delamination issue for layer 5 and layer 6 ladders since the reinforcement. This is consistent with different levels of the expected stress on the joint of the pitch adapter among different layers, worst for layer 4 and least for layer 5. It appears that the strength of the reinforced glue is still marginal for layer 4.

The SVD group came up with a new reinforcement method using a mylar strip with flaps which mechanically hold the pitch adapters onto the sensor. Measurements show that the improved reinforcement holds 2–2.5 times more stress. This method will be applied to all existing ladders as well as unassembled parts.

In addition to the above reinforcements, the SVD group introduced a new waiting period of one week after the pitch adapter is in place and glued. The waiting period is determined based on the past delamination timing after the gluing. In order to minimise the impact on the ladder assembly duration, FW sensor gluing is performed for two ladders, effectively reducing the loss to 3.5 days per ladder.

Assembly of layer-6 ladders is still on the critical path and barely on time for the second half of the ladder mounting. New assembly procedure to reduce the risk of the delamination more or less canceled out the time gain by staggering the two ladder assembly.

5.2.2 Concerns

The committee is concerned that new assembly schedule for layer 6 is quite complex and tight, which is very demanding and stressful for both operating personnel and managing personnel.

5.2.3 Recommendations

- Although the new reinforcement seems quite solid, the committee strongly recommends performing full thermal cycle tests of the reinforced ladders. For this purpose, class-C modules should be sufficient.
- Once the robustness of the new reinforcement is established, the committee recommends to revisit the necessity of the one-week waiting period after the FW sensor gluing. A trade-off of the risks and consequences between the delamination and the mishandling of the assembly operation due to complex procedures should be carefully examined.