

PCI-express based high-speed readout for the Belle II DAQ upgrade

Q. D. Zhou, S. Yamada, P. Robbe, D. Charlet, R. Itoh, *Senior Member, IEEE*, M. Nakao, S.Y. Suzuki, T. Kunigo, E. Jules, E. Plaige, M. Taurigna, H. Purwar, O. Hartbrich, M. Bessner, K. Nishimura, G. Varner, Y. -T. Lai, T. Higuchi, R. Sugiura, D. Biswas, and P. Kapusta

Abstract—Belle II is a new-generation B-factory experiment, dedicated to exploring new physics beyond the standard model of elementary particles in the flavor sector. Belle II started data-taking in April 2018, using a synchronous data acquisition (DAQ) system based on pipelined trigger flow control. The Belle II DAQ system is designed to handle a 30-kHz trigger rate with approximately 1% of dead time, under the assumption of a raw event size of 1 MB. The DAQ system is reliable, and the overall data-taking efficiency reached 84.2% during the run period of January 2020 to June 2020. The current readout system cannot be operated in the term of 10 years from the viewpoint of DAQ maintainability; meanwhile, the readout system is obstructing high-speed data transmission. A solution involving a PCI-express-based readout module with high data throughput of up to 100 Gb/s was adopted to upgrade the Belle II DAQ system. We particularly focused on the design of firmware and software based on this new generation of readout board, called PCIe40, with an Altera Arria 10 field-programmable gate array chip. Forty-eight GBT (GigaBit Transceiver) serial links, PCI-express hard IP-based DMA architecture, interface of timing and trigger distribution system, and slow control system were designed to integrate with the current Belle II DAQ system. This paper describes the performances accomplished during the data readout and slow control tests conducted using a test bench and a demonstration performed using on-site front-end electronics, specifically involving Belle II TOP and KLM sub-detectors.

Index Terms—Belle II, data acquisition (DAQ), high-speed readout system, PCIe40, PCI express, DMA.

I. INTRODUCTION

This work was supported in part by Grants-in-Aid for Scientific Research -KAKENHI- JSPS International Research Fellow Grants No. 18F18324.

Q. D. Zhou is with Institute of Advanced Research and Kobayashi-Maskawa Institute, Nagoya Univ., Nagoya 464-8601, Japan (e-mail: qzhou@hepl.phys.nagoya-u.ac.jp).

S. Yamada, R. Itoh, M. Nakao, S.Y. Suzuki, and T. Kunigo are with High Energy Accelerator Research Organization (KEK), Ibaraki 305-0801, Japan

P. Robbe is with the Univ. Paris-Saclay, CNRS/IN2P3, the Laboratoire de Physique des Deux Infinis Irene Joliot-Curie (IJCLab), Orsay F-91898, France.

D. Charlet, E. Jules, E. Plaige, and M. Taurigna are with the Laboratoire de Physique des Deux Infinis Irene Joliot-Curie (IJCLab), Orsay F-91898, France.

H. Purwar, O. Hartbrich, M. Bessner, K. Nishimura and G. Varner are with the Dept. of Phys. & Astr., Univ. of Hawaii at Mānoa, Honolulu, HI, 96822, USA.

Y. -T. Lai, and T. Higuchi are with Kavli Institute for the Physics and Mathematics of the Universe, University of Tokyo, Kashiwa, 277-8583 Japan.

R. Sugiura is with Tokyo University, Tokyo, 113-0033, Japan.

D. Biswas is with University of Louisville, Louisville, Kentucky, 40292, USA.

P. Kapusta is with Institute of Nuclear Physics (IFJ PAN), Krakow, 31-342, Poland.

THE standard model (SM) of particle physics, which describes the properties and physical behavior of elementary particles, has finally been established by the discovery of Higgs boson in 2012 [1], [2]. Although SM has demonstrated huge success, some phenomena are left unexplained, such as dark matter and material dominance universe. Belle II [3] is an experiment dedicated to exploring new physics beyond the SM in the flavor sector at the beam luminosity frontier. The SuperKEKB accelerator [4], located at KEK, Tsukuba, Japan, is designed to use the “nano-beam” scheme to achieve the instantaneous luminosity of $6.5 \times 10^{35} \text{ cm}^{-2}\text{s}^{-1}$, which is approximately 35 times that of its predecessor KEKB accelerator, by colliding 7 GeV electrons and 4 GeV positrons. It can produce copious amounts of B and D mesons and τ leptons. The large statistics enables precise measurements of rare decays for testing SM with unprecedented sensitivity. The Belle II experiment aims to collect 50 ab^{-1} integrated luminosity by 2031.

The Belle II detector [3] is a general-purpose spectrometer upgraded or replaced with respect to the Belle detector. It is configured around a 1.5 T superconducting solenoid surrounding the interaction region. The Belle II detector consists of a vertex detector (VXD), central drift chamber (CDC), particle identification (PID) detector, electromagnetic calorimeter (ECL), and neutral K_L and Muon (KLM) detector, which are located from the inside to the outside of the detector. The VXD is made of a silicon pixel detector (PXD) and a silicon vertex detector (SVD) used for detecting decay vertices. The CDC located outside of the SVD is used as a tracking detector for precise measurement of momentum and energy loss of charged particles. Using the Cherenkov counter technique, the PID detector is divided into barrel and forward endcap parts, which are the time-of-propagation (TOP) counter and the proximity-focusing Aerogel Ring Imaging Cherenkov detector (ARICH), respectively. The ECL uses CsI(Tl) crystals and CsI in the barrel and endcap regions, respectively. The KLM uses glass-electrode resistive plate chambers (RPCs) for the barrel region and layers of scintillator fibers for the endcap regions and the innermost layer of the barrel region.

II. BELLE II DAQ AND READOUT SYSTEM

The Belle II DAQ [5], [6] system is designed to process the data from the front-end electronics (FEE) to the storage system through several steps based on a highly unified system as shown in Fig. 1. The highly unified timing and trigger

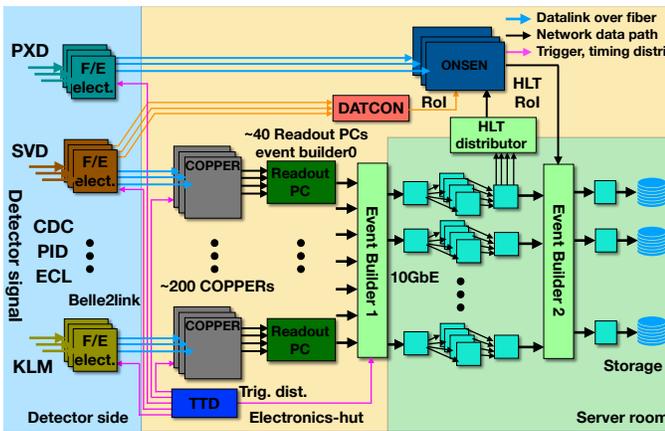


Fig. 1: Block diagram of the current Belle II DAQ system.

distribution (TTD) system [7] first distributes Level 1 (L1) triggers [8] to each FEE board using Frontend Timing Switch (FTSW) modules [9] based on a 254 Mb/s serial link protocol called “b2tt” [9]. The digitized signals from the FEE boards of each sub-detector, except for the PXD, are read out by a unified readout system based on the common module called Common Pipelined Platform for Electronics readout (COPPER) [10], as shown in Fig. 2. The data size of the PXD detector is too large to be read out by the COPPER module, an online event size reduction is necessary [11]. A total of 203 COPPER boards are used in the current Belle II DAQ system. To unify the readout system, a high-speed data transmission based on a custom serial link protocol, called Belle2link [12], is adopted. Xilinx GTP or GTX RocketIO up to 2.54 Gb/s, driven by the 127 MHz system clock, is used as the transceiver. The required bandwidth of each data link is less than 1 Gb/s, while the throughput of each readout PC is also limited to 1 Gb/s. As shown in Fig. 2, each COPPER module hosts four custom High Speed Link Boards (HSLB). They are interconnected by a 32-bit peripheral component interconnect (PCI) bus and a local bus used for slow control with 7-bit and 8-bit widths for address and data, respectively. The data formatting (per link) and module-level event building are processed by an on-board CPU on the COPPER module; subsequently, the data from several COPPER modules are collected to a readout PC through the Gigabit ethernet. The event building and software event selection are performed by high-level trigger (HLT) farms [13]. Finally, the selected data are sent to the storage servers. To handle the large data flow at the beam intensity frontier, the Belle II DAQ system is designed to handle 30 kHz of L1 trigger rate with approximately 1% of dead time, under a raw event size of more than 1 MB.

III. REQUIREMENTS FOR BELLE II DAQ UPGRADE

The current readout system, mainly the COPPER module and its readout PC, is obstructing high-speed data transfer and DAQ maintainability. The motivations for upgrading the Belle II DAQ readout system are as follows:

- 1) As shown in Fig. 2, the COPPER module hosts four HSLB daughter cards used as single-link receivers, a processor PCI mezzanine card (PrPMC), and a custom

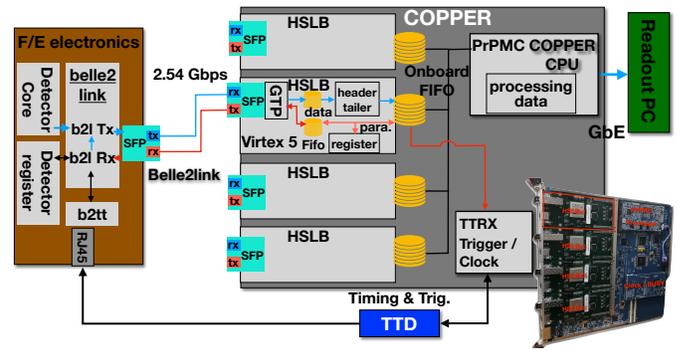


Fig. 2: Schematic view of unified COPPER readout system.

daughter card used as the TTD interface. The broken parts are increasing, such as the chipset of PrPMC, because the COPPER module has been in use since the Belle experiment.

- 2) The data formatting and module-level event building are processed by an onboard CPU. The CPU usage has been observed to reach $\sim 60\%$ under the test condition of 1 kB/event/COPPER at a 30-kHz trigger rate. In addition, the smallest bandwidth of the current Belle II DAQ system is the Gbit ethernet throughput of the COPPER module and the corresponding readout PC [11].
- 3) The current DAQ readout system could turn into a bottleneck when the Belle II experiment takes an aggressive plan of operation in the next 10 years. For instance, the SuperKEKB accelerator can achieve higher luminosity than the designed value, the background level of the Belle II experiment can be higher than expected, and the Belle II experiment can implement a high rate of low multiplicity triggers and even realize a trigger-less DAQ.

To reduce the modifications required for the other systems, all the functionalities of the current COPPER-based readout system must be retained. For the new system and readout module, the following are required:

- 1) The Belle2link protocol needs to be continued so that no further modification is applied to the FEE’s hardware and firmware for each sub-detector.
- 2) The functionalities implemented in the COPPER module, including data handling, slow control, data-formatting for each link and each COPPER, event-building of merging data of 4 links into one event, and error checks, are essential for the new system.
- 3) To integrate into the TTD system, the TTD interface needs to be implemented on the new system. A busy handshake is needed for each link.
- 4) The designed characteristics of the current readout system need to be followed. The trigger rate of 30 kHz with the estimated event size needs to be handled, onboard data corruption needs to be detected and should be reduced to less than 1 error / (100 kB/event \times 30 kHz \times 8 h).

Furthermore, the Belle II experiment will run for at least 10 years, and an upgrade of any or all of its sub-detector

TABLE I: Number of COPPER and PCIe40 required by the sub-detector

Sub-detector	Number of COPPER	Number of PCIe40
SVD	48	4
CDC	75	8
TOP	16	2
ARICH	18	2
ECL	26	2
KLM	8	1
TRG	12	1

is predictable. Thus, the new system should be capable of handling enhanced requirements from the sub-detector side. For instance, when the L1 trigger rate exceeds 30 kHz, the payload for each link will be higher than the bandwidth of COPPER system. In the Belle II Collaboration, we had four proposals for the upgrade of the readout system, all of them had a well-established module used in other experiments. Three proposals have proceeded to the development of the firmware and software to verify that the boards satisfy the basic requirements of the Belle II DAQ upgrade, as listed above. Finally, a PCI-express-based high-speed readout module, known as PCIe40 [14], was adopted by the Belle II Collaboration for DAQ upgrade. This is because PCIe40 fulfills the requirements of the Belle II DAQ upgrade and a large amount of data can be transferred easily from a readout board to a PC server via PCI-express bus.

IV. PCIe40 BASED NEW BELLE II READOUT SYSTEM

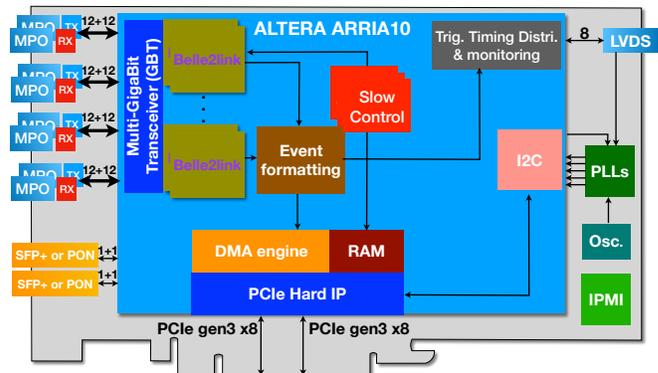
PCIe40 is a new generation of readout modules based on the PCI-express Gen3 solution with a high data throughput of 100 Gb/s [15]. This module will also be used for the DAQ upgrade of the LHCb and ALICE experiments at the Large Hadronic Collider (LHC) [14], [16]. To integrate with the current Belle II DAQ system, we particularly focus on the design of the firmware and software for the PCIe40 board. Since all functionalities of the current readout system are retained, no modification in the firmware and hardware of the sub-detector FEE is required. A single PCIe40 board can connect up to 48 bidirectional optical links, while only 4 bidirectional links can be handled by the COPPER module. Thus, the new readout system will become much more compact by replacing 203 COPPERs with 20 PCIe40 boards. The number of COPPERs used in the current readout system and the number of PCIe40 required for replacement are listed in Table I.

A. PCIe40 module

The PCIe40 module is required to maintain interface compatibility with the currently running Belle II DAQ system. As illustrated in Fig. 3, four pairs of onboard MiniPOD optical transceivers are used to receive or transmit data from or to the FEE through optical fibers. A high-density Intel Altera Arria 10 field-programmable gate array (FPGA) chip (1.15 million cells) provides the board with powerful reconfigurable logic capabilities. Extra effort is required to implement the firmware of the Belle2link protocol on the Altera FPGA, as the current Belle2link logic is implemented on the Xilinx Virtex 5 FPGA. It has been verified that Belle2link is working well on PCIe40



(a)



(b)

Fig. 3: (a) Photo of PCIe40 board. (b) Block diagram of the PCIe40 board illustrating the key features of the board.

by checking the cyclic redundancy check (CRC) errors during the data transmission. PCIe40 is capable of running at up to 10 Gb/s for each link using the Gigabit Transceiver (GBT) architecture [17] with a maximum of 48 bidirectional links. In addition, there is an onboard LVDS connector with 8 links used as an interface to the Belle II TTD system, and if it is necessary each board can be connected to 2 TTD modules (FTSW). Three external phase-locked loops (PLLs) (Si5344 and Si5345) are used for clock jitter cleaning, and a clock source may be selected from the TTD system or from an onboard oscillator. An embedded USB blaster on the front panel connected to the PC server enables remote downloading of the FPGA firmware. There is an over-heating protection system using the intelligent platform management interface (IPMI) for the management. A point-to-point PCI-express Gen3×16 links is used to connect the PC server and the PCIe40 board. A PCIe40 board is presently installed on a 1U rack-mount server (Xeon E5-2640v4 10 Core/2.4GHz), which was prepared as a spare server for the current COPPER readout system. Multiple PCIe40 boards can be installed in a larger server.

B. Firmware development

The Belle2link protocol, b2tt protocol, and slow control functionalities which are implemented on the current COPPER-HSLB system, need to be integrated into the PCIe40 system. Event formatting and event building, as well as PCI-express hard IP-based direct memory access (DMA) are newly

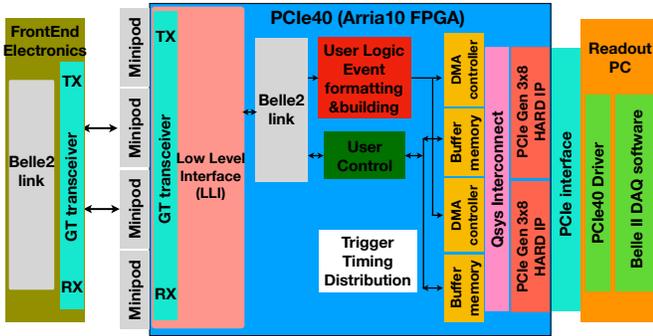


Fig. 4: Block diagram of key features of firmware implemented on PCIe40. Front-end electronics and readout PCs are also shown to illustrate the data flow direction.

designed for PCIe40. As shown in Fig. 4, a generic firmware architecture is designed to provide a unified readout platform for all sub-detectors (except for PXD). A hardware abstraction layer, called low-level interface (LLI) [18], is designed for user control and initialization of the PCIe40 board. The frontend side of the PCIe40 firmware includes the configuration of the GBT and decoding, synchronization, and aggregation of all GBT links. Meanwhile, it provides a simple interface for user application codes, such as the slow control unit. Thus, most user applications can connect with the LLI by simply retrieving from and pushing data to first-in-first-out (FIFO) interfaces.

The state machine of the Belle2link protocol is implemented on an Arria 10 FPGA. It includes both decoding and encoding of the event fragment data and slow control parameters based on the defined Belle2link format. Because the local bus widths of COPPER for address and data are 7-bit and 8-bit, respectively, the slow control feature was originally designed to access 8-bit registers mapped onto a 7-bit address space. This is known as the A7D8 access method for the Belle II experiment. A16D32 and stream file features are upgraded based on A7D8 to satisfy different requirements of the sub-detectors. As listed in Table II, each corresponding sub-detector uses one or two slow control access methods. All three features for slow control access are retained when upgraded to a new readout system. Address and data decoding and encoding for slow control accesses are moved from firmware to software. A FIFO is used as the interface to store addresses and parameters received or need to be transmitted by Belle2link. The FIFO is controlled by signals connected to the registers that are accessible using slow control software. The ARICH detector uses the stream file method to transfer the firmware of the FEE, and KLM uses it to stream the threshold parameters for the RPC detector. To maintain a low error rate during transmission, the file to be streamed is split into packets of a certain length. The data packet length is 100 words for ARICH and 6 words for KLM. This difference mainly depends on the FEE of ARICH and KLM.

Figure 5 illustrates the key features of data processing on the PCIe40 board and readout PC server. A user logic block is designed to handle event building and data formatting, and a

TABLE II: Access methods of each sub-detector for slow control

Sub-detector	A7D8	A16D32	Stream file ^a
SVD	✓	×	×
CDC	×	✓	×
TOP	×	✓	×
ARICH	×	✓	✓
ECL	✓	✓	×
KLM	×	✓	✓

^a Stream file method includes streaming a complete file, such as FEE firmware, and a defined parameter length.

PCI-express hard IP-based DMA architecture offers a fully integrated, flexible, and highly optimized solution for high bandwidth and low latency data transmission. With DMA, data are copied from the FPGA on-chip memory into the memory of the readout PC server without preempting the CPU. The event fragment data are received by each Belle2link and buffered to a corresponding FIFO, and then the generic event formatting and event building logics are implemented to handle the data based on the firmware. The state machine is designed to merge fragments of the 48 links sequentially into an event and to remove redundant header and trailer information of each link, after it was checked. Subsequently, redefined event header and trailer information is added to the built event. The merged event data are then passed to the DMA data FIFO with a size of 32 kB, used as an interface to the PCIe DMA block.

The Arria 10 Avalon-MM (Memory Mapped) DMA interface [19] is used to easily design the PCIe protocol. As shown in Fig. 5, the design includes the following main features: on-chip memory, PCIe Hard IP, and DMA descriptor controller. The built event data from the user logic block are stored in the on-chip memory having a 256-bit width and 8 kB pre-DMA page. Conversely, 128 DMA pages are contained in a DMA super page on the PC server side. The DMA read and write operations are managed by an external DMA descriptor mechanism, which is separate from the DMA engine. The procedure of sequentially copying data from the on-chip memory by reading out the PC server's memory is based on the descriptor table. The descriptor table contains the data source address (a pointer in the on-chip memory mapped fabric of the FPGA), destination PC memory's address (in the PCIe bus address space), and the transferred data size. The software controller first programs the descriptor controller's register with the location, size, and the number of the descriptor table; subsequently, the Read DMA engine issues a memory read request to forward the entire descriptor table to the FIFO controller. When the FIFO controller receives a list of free DMA super pages, the firmware DMA controller fills these DMA super pages without further software intervention. The controller then fetches the table and requests the Write DMA engine to transfer the data from the Avalon-MM to the PCIe domains with one descriptor at a time. When the data transmission is completed, the controller sends the DMA status upstream.

An optimized handshake among the user logic block, PCIe-based DMA, and the corresponding software driver is essential. A certain large buffer capability, careful synchronization

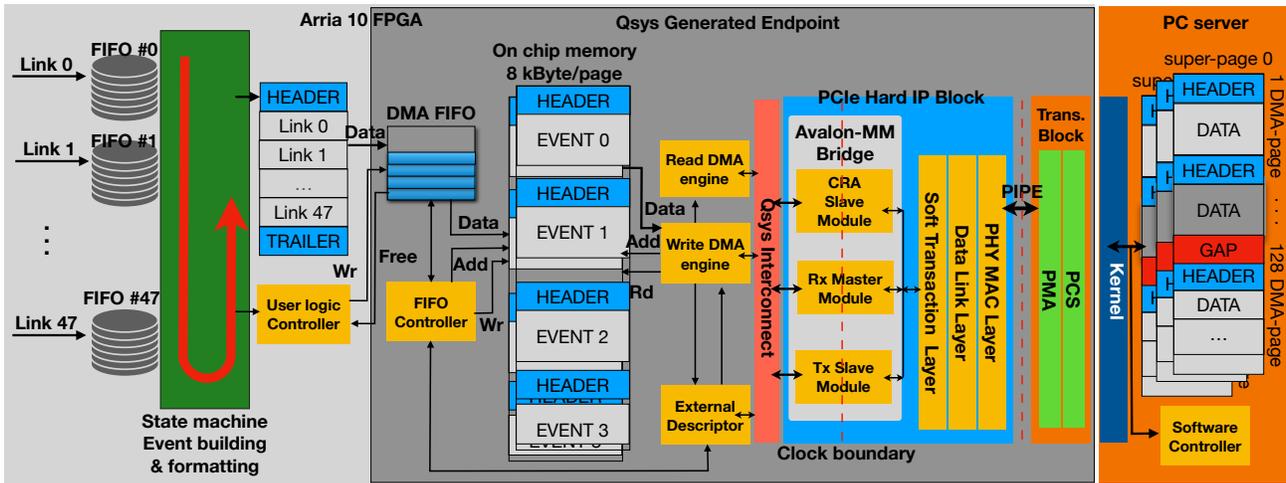


Fig. 5: Block diagram of key features of firmware implemented on PCIe40 for data-flow processing and the software features for DMA. Firmware designs include two main functionality blocks, which are user logic part for event formatting, and speed-up event building and the PCI-express DMA core, which offers high bandwidth and low latency data transmission.

within the DMA controller, and sufficient data processing power of the readout PC server ensure that data can be simultaneously transferred downstream at the highest possible rate with efficient utilization of resources.

The TTD system is designed for trigger and timing distribution and busy handshake control for the Belle II experiment. A unified bidirectional communication based on a 254 Mb/s serial data link, called b2tt, is used as the interface between the PCIe40 board and the TTD system. A TTD interface with one b2tt link is implemented for the PCIe40 board to handle all 48 belle2links simultaneously. Currently, architecture to handle 48 belle2links individually is being developed. Thus, when a link has an error or is busy, the corresponding issue can be successfully collected by the TTD system. A 127-MHz external clock derived from the SuperKEKB accelerator's radio frequency (RF) is distributed to the PCIe40 board and used as the system clock for the firmware.

C. Software development

Data readout and slow control access based PCI-express interface are new features compared to the COPPER-HSLB system. The software driver and library are developed based on C/C++, and they mainly provide the functionality for accessing and controlling the data acquisition or configuration of the PCIe40 board. On the PCIe40 side, the base address register 0 (BAR0) and BAR2 of PCI-express are designed to be used for slow control access and DMA access, respectively. The corresponding driver and library were developed individually. The driver is involved in managing physical memory, and the library provides utilities for accessing registers, performing tests, etc. The driver for DMA is also used to manage the DMA access and provide the functionalities of accessing the availability status of memory for DMA, checking the arrived data, and then applying operations to the arrived data. As shown in Fig. 5, the unit used for individual DMA transfer is the DMA page, which is 8 kB per page, and this size depends

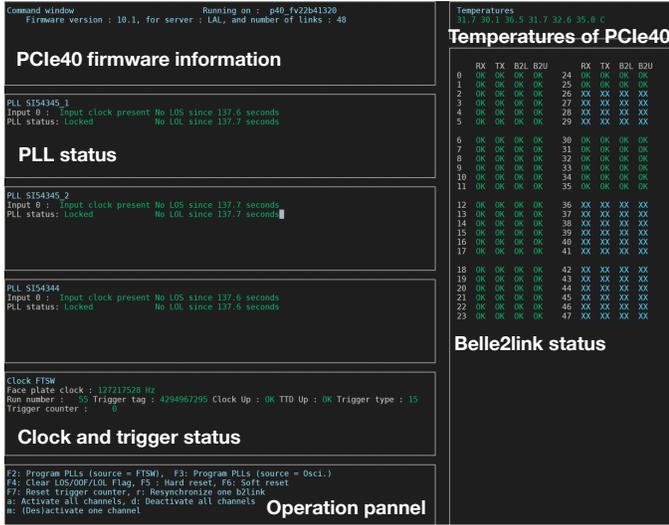
on the Linux constraint. On the PC server-side, 128 DMA pages are contained in a DMA super page of 1 MB. A typically large buffer of several GBs is allocated for each DMA channel. The management of the DMA access and memory allocation is based on the super page.

An application layer, developed based on low-level drivers and libraries, contains individual utilities for high-rate data acquisition, slow control register access, monitoring, and configuration of the PCIe40 board, Belle2links, etc. An example of an application is shown in Fig. 6a. A Python script is designed to show a GUI window with firmware information, PLL status, clock and trigger information, and all Belle2link status. All this information is used for a monitoring purpose, it can help the operator to quickly investigate the readout system if an issue occurred. An operation panel provides commonly used configuration options to rapidly switch the clock source, resynchronize the Belle2link for link establishment, and mask or unmask an individual or all Belle2links.

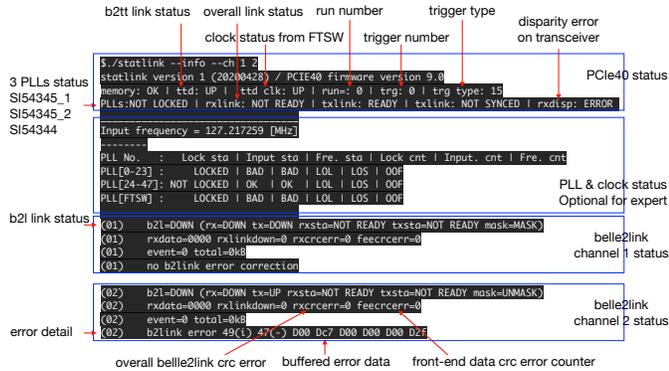
D. Software development for integration with Belle II DAQ system

There are mainly two blocks of software that need to be upgraded for PCIe40 board in order to integrate with the current Belle II DAQ system. One is related to the slow control system, and the other to the data readout system. The slow control feature of the current COPPER system was designed based on the COPPER local bus. Similar features must be developed for the PCIe40 board. As shown in Fig. 6b, a commonly used tool for investigating and debugging the cause of an error occurring in the readout system has been developed. It shows inclusive PCIe40 status, individual link status, counters of CRC errors, and buffered corrupted data. Meanwhile, this output format can be saved into the logging system thus can be easily integrated into the current Belle II logging and monitoring system.

The Belle II DAQ slow control system is designed based on a network interface called network shared memory 2



(a)



(b)

Fig. 6: (a) GUI view of a general monitoring system for PCIe40 board and Belle2link. (b) Output of the monitoring system with specific error details. It can be integrated into the Belle II log system.

(NSM2) [20]. The application nodes are C/C++ daemon programs, which are used to manage the slow control functionalities. Figure 7 illustrates the slow control architecture on the PCIe40 readout PC with newly developed daemon processes, such as pcie40controll and pcie40linkd. pcie40linkd is used to manage and control an individual Belle2link, and pcie40controll is designed to manage and control all pcie40linkd processes running on the readout PC. PCIe40FEE is a sub-detector-specified library, where the sub-detectors can define their functionalities to control their detectors through Belle2link.

The data readout related modifications to be integrated with the current DAQ system are relatively small. Data readout software with the new DMA scheme has already been developed. The data format of the merged event is slightly different and needs changes for data compression. Because the data from 48 links are processed in one readout PC, the upgrade of the related software from 4 links for COPPER to 48 links for PCIe40 is essential.

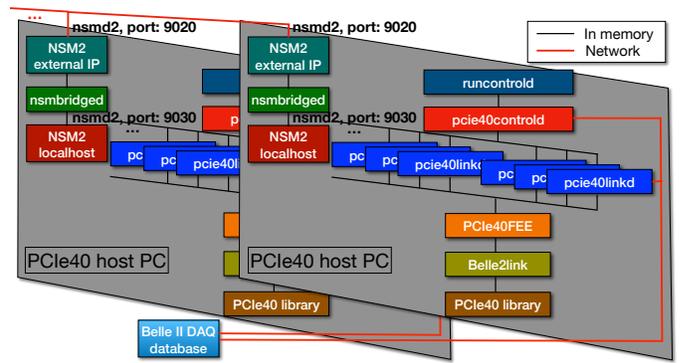


Fig. 7: Block diagram of slow control architecture on readout PC; a NSM2 based system, involving several newly developed daemon processes, such as pcie40controll and pcie40linkd, has been developed.

V. PERFORMANCE OF NEW READOUT SYSTEM

Firmware and software for the new readout system have already been developed. As discussed in Sec. III, it is necessary to confirm that the new readout system fulfills the requirements of the Belle II DAQ system. A test bench was constructed for the demonstration of the readout system. Meanwhile, another readout system was built and connected with the Belle II TOP and KLM sub-detectors at the Electronics-Hut in Tsukuba Hall, at KEK, which is used to host the electronics of the Belle II experiment. Slow control and data readout tests were performed using the TOP and KLM on-site detectors.

A. Test bench

As shown in Fig. 8, we constructed a test bench for the readout system, specifically involving the FEEs of the CDC, SVD, TOP, and ARICH and 46 dummy FEEs based on COPPER-HSLB, for demonstration. The PCIe40 board was installed on a host PC server, and the electric power was provided by the host PC. The firmware was downloaded using an embedded USB blaster on the PC. Forty-eight FEEs were connected to the PCIe40 board using 8 Multiple-Fiber Push-on/Pull-off (MPO) fibers, and each channel could be activated or deactivated by controlling the corresponding register on the PCIe40. The clock and triggers were distributed to PCIe40 and FEEs by the TTD system. The two types of triggers are evenly spaced pulse triggers and random triggers with intervals of pseudo-Poisson distribution, which can be issued by the TTD system with an adjustable input trigger rate.

B. Performance of data readout

Firstly, we measured the throughput of PCIe DMA data-transfer by using a throughput test block shown in Fig. 9, where a pattern generator was implemented as the data source instead of data from the FEEs. A FIFO was added for back-pressure monitoring. Measurements at high event-rates were performed with a pulse trigger control to confirm the capability of the new readout hardware.

- 1) The back-pressure was observed in 10% of events when using a pulse trigger with 470 kHz trigger rate and an

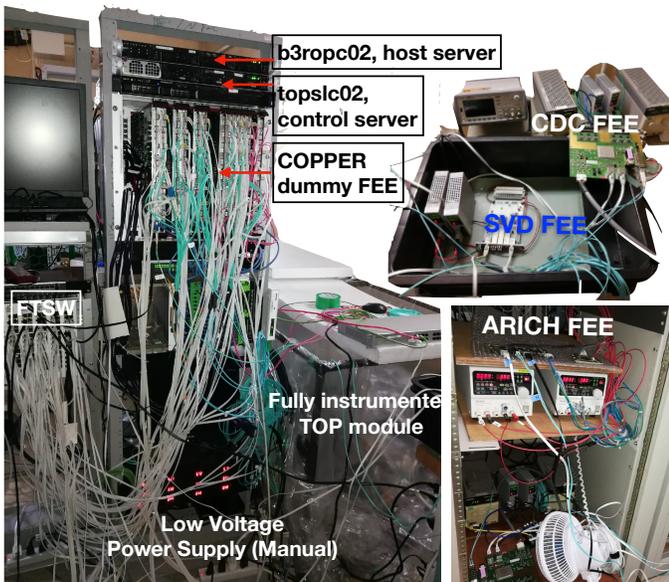


Fig. 8: Overview image of test bench for DAQ readout upgrade. FEEs of TOP, CDC, ARICH, SVD, and KLM and 46 dummy FEEs based on COPPERs were connected to the PCIe40 board for the performance test.

event size of 8 kB, for which the corresponding data transfer rate was 31 Gb/s,

- 2) No event was lost if the trigger rate was reduced to 260 kHz and when the data transfer rate was 17 Gb/s.

It is reported that the performance of a single instance of the DMA controller bonded to a PCIe \times 8 Gen3 interface can consistently reach 50 Gb/s [15], and it can eventually be increased to 100 Gb/s using the PCIe \times 16 Gen3 interface completely. Conversely, the current performance of the data readout system already fulfills the Belle II DAQ upgrade requirements, while the throughput limit of the new readout system is 10 Gb/s, which is determined from the new readout PC.

Secondly, we demonstrated a back-pressure scheme in data-transfer from 46 COPPER-HSLB dummy FEEs as data sources to a PCIe40 board, by using a back-pressure test block in the FPGA of the board shown in Fig. 9. No back-pressure or lost events were observed with the new readout system under a 43 kHz trigger rate and 1 kB / event / link data readout. Here, the trigger control was set with a restriction of

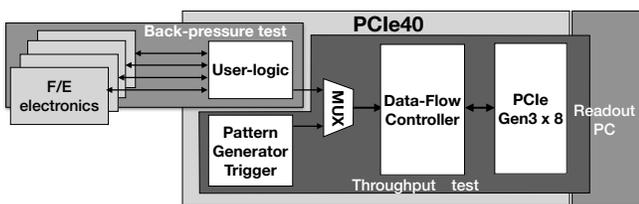


Fig. 9: Schematic view of components (throughput-test block and back-pressure-test block) involved in the two performance tests.

no more than 10 triggers within 130 μ s, which was determined from the SVD FEEs. Meanwhile, 2.5% of the events issued a back-pressure under a 26-kHz trigger rate when a 0.8 ms latency was artificially inserted into the FEE data source. The back-pressure is due to the limited size of the FIFOs, which are implemented to buffer the data from the FEEs for event building. The on-chip memory size of the PCIe40 board is only \sim 70 Mbits. As the memory of PC server is much larger, this issue will be solved by transferring the event building from the FPGA logic to the software in the future.

C. Performance of slow control system

All three slow control access methods using the PCIe interface were tested on the test bench. The single access speed of the A16D32 method was tested by reading the register of the CDC FEE 1×10^6 times. The average access time was 83 μ s/access, while the COPPER-HSLB system took \sim 1 ms. The streaming file method was tested using the FEEs of KLM and ARICH. The streaming speed was \sim 360 kB/s for KLM, which is comparable with the COPPER-HSLB system (\sim 350 kB/s). Both the PCIe40 and the COPPER-HSLB system took less than 2 s to transfer the 3 MB ARICH firmware. A new feature for accessing different individual links in parallel is available for the PCIe40 board, which sufficiently reduces the time required for configuring the FEEs of a sub-detector, while, for the COPPER-HSLB system, the four links need to be accessed sequentially.

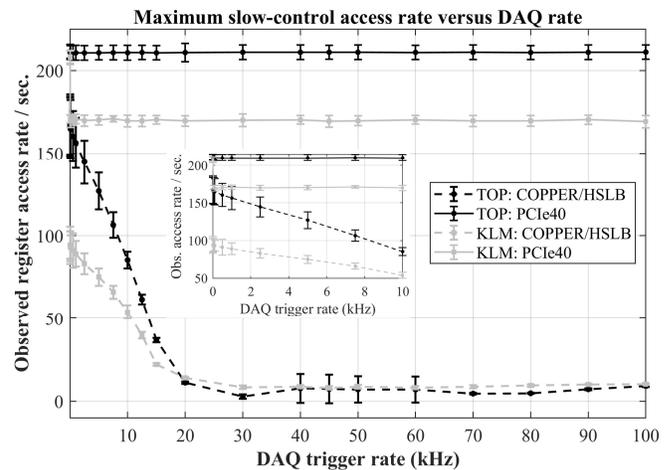


Fig. 10: Register accessing rate when simultaneous data acquisition is performed with different trigger rates. The curve showing a decreasing accessing rate from 100 to 0 corresponds to the access of a KLM FEE register using the COPPER-HSLB system, while the accessing rate retains a constant value of 170 for the PCIe40 board. A similar tendency is observed for TOP in which the accessing rate decreases for the COPPER-HSLB system when the DAQ trigger rate increases and remains constant for the PCIe40 board.

A register access test was performed together with data acquisition under different trigger rates. As shown in Fig. 10, the register accessing rates for TOP and KLM based on the COPPER-HSLB system decrease from 170 and 100 to 0,

respectively, when the data acquisition trigger rate increase from 0 to 30 kHz. This indicates that there is a limitation of bandwidth for slow control and data acquisition for the COPPER-HSLB system. However, the corresponding register accessing rates for TOP and KLM based on the PCIe40 board retain a constant value even when the data acquisition trigger rate is increasing. There was no limitation of bandwidth observed for the PCIe40 board.

Another complete slow control test was performed with the on-site FEEs of TOP and KLM. All 64 FEEs of TOP were connected to 2 PCIe40 boards that were hosted by 2 separate servers, and 32 FEEs of KLM were connected to a single PCIe40 board. The slow control test was performed based on the upgraded software architecture, as shown in Fig. 7. The speeds of configuring FEEs and setting the threshold for both TOP and KLM detectors based on the PCIe40 board are compatible with the COPPER-HSLB board. The configuration for TOP was confirmed by taking a short run and checking the data quality. The threshold for KLM was verified by checking the output raw data file size by varying the threshold value. It is confirmed that the new software architecture for slow control works properly.

VI. PREPARATION AND REPLACEMENT

The mass production of 31 PCIe40 boards for Belle II DAQ upgrade is completed, and the boards have arrived at KEK after they have been successfully tested at CERN. Three boards were installed on 3 servers for the readout system replacement of TOP and KLM. MPO fibers are connected to the patch panels located close to the COPPER-HSLB system of TOP and KLM. A 10 Gb/s network interface card (NIC) is installed on host PC servers for downstream connection. A 40 Gb network switch is installed close to the readout PCs for connection with the network switch located in the HLT server room. The replacement of the readout system for TOP and KLM is scheduled for summer 2021, and the replacement for other sub-detectors will take place afterwards. The current COPPER system will stand by for some time after the replacement. In the case of a serious problem, it can be rolled back to the COPPER system quickly.

VII. CONCLUSION

A proposal based on the PCIe40 board was adopted for the Belle II DAQ upgrade. It is necessary to retain the main features of the current readout system and reduce the modifications required for the sub-detectors when upgrading to the new readout system. The developments of FPGA firmware and software are almost complete. The PCI-express hard IP-based DMA architecture of the new readout hardware is capable of handling a 260-kHz trigger rate with 17 Gb/s data throughput without data loss. No back-pressure or event loss was observed in data-transfer from 46 dummy FEEs to the PCIe40 board under 43-kHz trigger rate and 1 kB/event/link data size. Three accessing methods used for slow control offered similar or better performance with respect to the current COPPER-HSLB system. A new feature that enables accessing individual links in parallel improves the configuration speed of

the sub-detectors. It has been confirmed that the new software architecture for slow control works properly. The replacement of the Belle II readout system is planned to start from TOP and KLM sub-detectors for summer 2021, and the replacement of the COPPER modules of the other sub-detectors will be done afterwards.

REFERENCES

- [1] ATLAS Collaboration, "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC," *Phys. Lett., B* vol. 716, issue 1, pp.1-29, Sep. 2012.
- [2] CMS Collaboration, "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC," *Phys. Lett., B* vol. 716, issue 1, pp.30-61, Sep. 2012.
- [3] T. Abe et al., "Belle II Technical Design Report," *arXiv:1011.0352*, Nov. 2010.
- [4] K. Akai, K. Furukawa, H. Koiso et al. (SuperKEKB Collaboration), "SuperKEKB collider," *Nucl. Instrum. Methods Phys. Res., Sect. A* vol. 907, pp. 188-199, Nov. 2018.
- [5] S. Yamada et al., "Data acquisition system for the Belle II experiment," *IEEE Trans. Nucl. Sci.*, vol. 62, no. 3, pp. 1175-1180, Jun. 2015.
- [6] M. Nakao et al., "Data acquisition system for Belle II," *J. Instrum.*, vol. 5, no. C12004, Dec. 2010.
- [7] M. Nakao, "Timing distribution for the Belle II data acquisition system," *J. Instrum.*, vol. 7, no. C01028, Jan. 2012.
- [8] Y. Iwasaki et al., "Level 1 trigger system for the Belle II experiment," *IEEE Trans. Nucl. Sci.*, vol. 58, pp. 1807-1815, Aug. 2011.
- [9] M. Nakao et al., "Minimizing dead time of the BelleII data acquisition system with pipelined trigger flow control," *IEEE Trans. Nucl. Sci.*, vol. 60, pp. 3729-3734, Oct. 2013.
- [10] T. Higuchi et al., "Modular pipeline readout electronics for the super-belle drift chamber," *IEEE Trans. Nucl. Sci.*, vol. 52, pp. 1912-1917, Oct. 2005.
- [11] S. Yamada et al., "Common Readout Subsystem for the Belle II Experiment and Its Performance Measurement," *IEEE Trans. Nucl. Sci.*, vol. 64, no. 6, pp. 1415-1419, Jun. 2017.
- [12] D. Sun et al., "Belle2Link: A global data readout and transmission for Belle II experiment at KEK," *Phys. Procedia*, vol. 37, pp. 1933-1939, 2012.
- [13] R. Itoh et al., "Data flow and high level trigger of Belle II DAQ system," *IEEE Trans. Nucl. Sci.*, vol. 60, pp. 3720-3724, Oct. 2013.
- [14] J. P. Cacheriche, "The PCIe-based readout system for the LHCb experiment," *Topical Workshop on Elect. for Part. Phys.*, Lisbon, Portugal, Sep. 2015, 10.1088/1748-0211/11/02/P02013.
- [15] P. Durante, N. Neufeld, R. Schwemmer, U. Marconi, G. Balbi and I. Lax, "100Gbps PCI-express readout for the LHCb upgrade," *2014 19th IEEE-NPSS Real Time Conference*, Nara, 2014, pp. 1-3, doi: 10.1109/RTC.2014.7097492
- [16] J. Mitra, S.A. Khan, S. Mukherjee, and R. Paul, "Common Readout Unit (CRU) - A new readout architecture for the ALICE experiment," *JINST* 11, C03021 (2016), doi: 10.1088/1748-0221/11/03/C03021
- [17] P. Moreira et al., "The GBT Project," *CERN-2009-006.342*.
- [18] A. F. Prieto, "Back-end firmware for the LHCb VELO upgrade phase 1," *Topical Workshop on Elect. for Part. Phys.*, Santiago de Compostela, Spain, Sep. 2019, 10.22323/1.370.0101.
- [19] Intel ALTERA "Arria 10 Avalon-MM DMA Interface for PCIe Solutions," *User Guide*, May 2016, www.altera.com.
- [20] M. Nakao and S. Y. Suzuki, "Network shared memory framework for the belle data acquisition control system," *IEEE Trans. Nucl. Sci.*, vol. 47, pp. 267-271, Apr. 2000.