

Trigger Timing Interface for the Read-Out Upgrade of the Belle II DAQ

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

Abstract—To improve data throughput of the Belle II data acquisition we are upgrading the CPU-based COPPER system with a PCIe40 board carrying Arria 10 FPGA. Since one of the main functionalities of the new system is event building in FPGA, the read-out system must be synchronized with the front-end electronics. This task is performed by the bidirectional trigger timing distribution system. During system commissioning, we prepared several versions of the interface to this system. In the initial version of the interface, we ported the code from Xilinx FPGAs to Arria 10. This revision also introduces monitoring of the status for multiple channels and a ring buffer to distribute trigger information to all channels in parallel. To improve stability under external noise, we implemented a clock-data recovery using an independent on-board oscillator as a reference clock in the next revision of the interface. We are also developing a version utilizing a high-speed serial transceiver to replace CAT-7 RJ45 cables with optical fibers. The system commissioning started in 2021 with a few detectors and will be completed after the long shutdown 1 of SuperKEKB in 2023. In this paper, we present the architectures of the interface to the trigger timing system implemented in the PCIe40 board and the system performance in the experiment.

I. INTRODUCTION

THE Belle II experiment [1] is an experiment at the electron-positron collider SuperKEKB [2] with the goal to accumulate a world-leading data set for the CP (Charge-Parity) violation measurements in the heavy quark sector and the search for physics beyond the standard model. The experiment consists of seven subdetectors: pixel detector, PXD [3],

double-sided silicon strip detector, SVD [4], central drift chamber, CDC [5], time-of-propagation detector, TOP [6], aerogel ring imaging Cherenkov detector, ARICH [7], electromagnetic calorimeter, ECL [8], and long-lived kaon and muon detector, KLM [9].

Fig. 1 shows the layout of the DAQ (data acquisition) system before the upgrade. Six detectors are read out by the unified read-out system COPPER [10]. This system consists of about 200 boards in the VME (Versa Module Eurocard) form-factor. Each board reads data from four data sources, combines them into an event, and sends them to one of the 43 read-out PCs (personal computers), ROPC, for the first stage of the event building. Only data packaging is performed in the COPPER system. The VME board carries 4 data receiver boards, with each one can receive 2.5 Gb/s of detector data, plus a trigger timing receiver board, and a CPU (Central Processing Unit) board for event building. The COPPER system is installed outside the radiation environment, so radiation tolerance is of no concern for the system.

Due to the high data rate of the PXD in comparison to all other detectors, we adopted a two-step event building scheme. First, we combine data from all detectors except for the PXD. The high-level trigger computing farm, HLT, then extrapolates the tracks to the PXD planes to calculate regions of interest, ROIs, within the PXD. The ROIs are used for online data reduction in the online selection node [11] which removes all data outside the ROIs. Finally, reduced data from the PXD is merged with the remaining event in the event builder 2.

The read-out system has several shortcomings which limit the performance and maintainability of the read-out system. First, while the input to the COPPER system is 4x2.5 Gb/s, the throughput of each COPPER board is limited by 1 Gb/s Ethernet output. Until now, the instantaneous luminosity and the corresponding radiation background were low, and the data rate through the COPPER system was well below 1 Gb/s. Therefore, no data reduction was needed. With the rising luminosity foreseen in the future, the radiation background and the interaction rate will increase leading to an increase in the data flow.

Another source of the increase in the data rate is the plan of certain sub-detectors to switch to the raw data readout to move data processing algorithms outside of the environment with high radiation exposure. For example, the TOP detector suffers from radiation-induced lockups of the processing system in the Zynq SoCs (System-on-a-Chip) which are responsible for pedestal subtraction and feature extraction. These lockups

Manuscript submitted for review August 22, 2022; revised January 6, 2023
D. Levit, S. Yamada, R. Itoh, M. Nakao, S.-H. Park, S.Y. Suzuki, and T. Kunigo are with High Energy Accelerator Research Organization (KEK), Ibaraki 305-0801, Japan (e-mail: dmytro.levit@kek.jp).

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

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.

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

P. Kapusta is with The Henryk Niewodniczanski Institute of Nuclear Physics (IFJ), Polish Academy of Sciences (PAN), Kraków, 31-342, Poland.

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

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

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

Q.-D. Zhou is with Institute of Advanced Research and Kobayashi-Maskawa Institute, Nagoya Univ., Nagoya 464-8601, Japan.

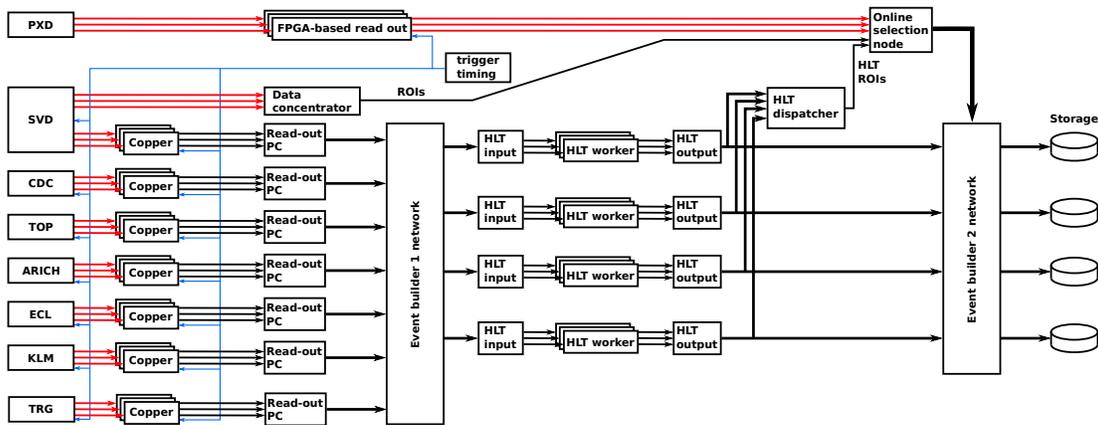


Fig. 1. Layout of the Belle II DAQ before the upgrade. The read-out system is built out of COPPER boards and read-out PCs

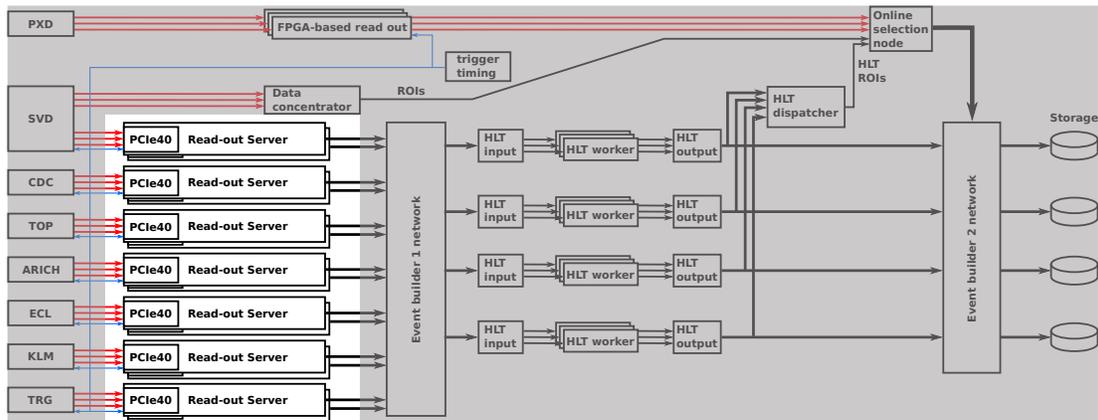


Fig. 2. Layout of the Belle II DAQ after the upgrade. We replaced the COPPER-based read-out system with read-out servers which carry PCIe40 boards. Unchanged components are grayed out.

cannot be mitigated and recovering them reduces the efficiency of data taking. By switching to the raw-data readout, we plan to move the pedestal subtraction and feature extraction to the ROPCs, but this will increase the expected data rate at 30 kHz from approximately 12 MB/s/link to 90 MB/s/link, which will be impossible to transfer out of the COPPER system without data truncation.

Next, the CPU board used in the COPPER system carries an Atom processor which has low data processing performance. And last, the maintenance of the system designed more than 10 years ago is getting more expensive due to the lack of availability of the components.

To remove the bottleneck in the system performance and guarantee system operation in the coming years we decided to upgrade the read-out system with an FPGA (field-programmable gate array)-based board PCIe40 [12]. This board carries an Arria 10 [13] FPGA and provisions up to 48 bidirectional high-speed serial lines for data input and two 8-lanes PCIe Gen3 (Peripheral Component Interconnect Express generation 3) interfaces for data output. Hypothetically, this would give us a throughput of 12 GB/s/board. We plan to replace all the 200 COPPER boards of the read-out system with approximately 20 PCIe40 boards, thus increasing the maximum throughput of the system by a factor 10.

We started to upgrade the system with 2 detectors: TOP and KLM detectors in the autumn 2021 run. We upgraded the read-out electronics of the ARICH in the winter 2022 run. We will complete the upgrade of the read-out electronics for all detectors except for the PXD, which uses an independent read-out system [14], in the long shutdown 1 of SuperKEKB in 2022-2023.

The paper has the following structure. Section II describes the layout of firmware in PCIe40. Section III describes the trigger timing distribution, TTD, system in Belle II experiment. Section IV describes the implementation of the interface to the TTD system in PCIe40. Section V describes the performance and operation of the interface in Belle II experiment.

II. PCIe40 IN BELLE II

Fig. 2 shows the layout of the Belle II DAQ system after the foreseen upgrade. The PCIe40 is installed in the ROPC and transmits data received from the front-end electronics to the ROPC over PCIe interface. Due to the higher density of the high-speed links within PCIe40, the number of ROPCs is reduced to 20.

Fig. 3 shows the clock tree and the layout of the PCIe40 firmware. The PCIe40 receives data using up to 48 receivers for Belle2Link protocol [15], and sends them to the event

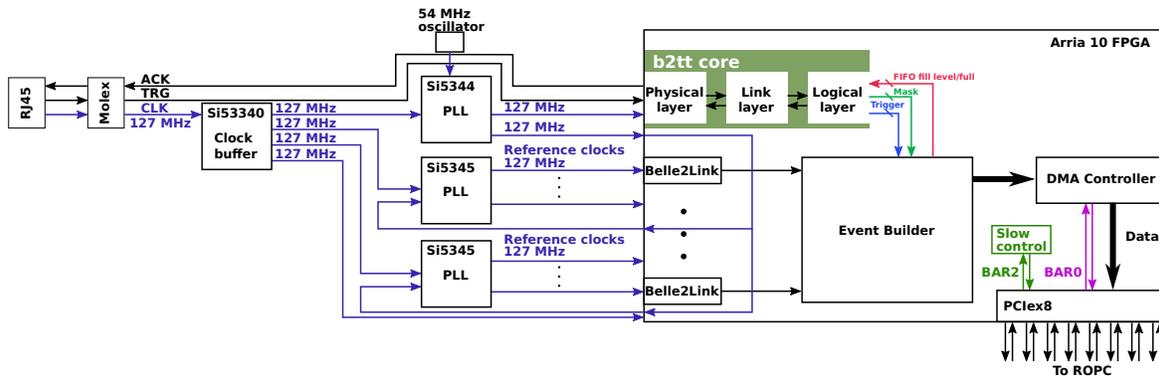


Fig. 3. Trigger timing interface and firmware layout of PCIe40 in Belle II.

builder core. The event builder core combines data from all sources, packages them to an event frame, and sends the event frame to the direct memory access (DMA) controller to transfer data to the ROPC over the PCIe8 interface.

The PCIe interface also provisions two register-based interfaces: base address register (BAR) 0 and 2. BAR0 configures the DMA controller and controls the data flow. BAR2 is used for slow control and configuration of all other firmware components.

In addition to the event building, we use the PCIe40 to configure and monitor the front-end electronics. The registers of the front-end electronics are accessible over the Belle2Link, which offers an in-band stream for the slow control messages. The interface to this stream on PCIe40 is mapped to BAR2. The slow control software for the front-end electronics runs on the ROPC and has direct access to the slow control stream of the Belle2Link over the PCIe interface.

To check data for consistency, for the flow control, and for the monitoring, the system is synchronized with the front-end electronics by using the TTD system [16]. The TTD system provides trigger information to the event builder and in case of data congestion can assert a temporary stop of the trigger distribution in the experiment to avoid buffer overflow.

III. TRIGGER TIMING DISTRIBUTION SYSTEM

The TTD system distributes the synchronization information and the clock signal to the front-end and read-out electronics, and collects monitoring information and flow control status.

The TTD system consists of the front-end timing switch boards, FTSW [16], that provision up to 20 downstream links each. We use three versions of the boards:

- 1) boards with 20 RJ45 (registered jack 45) connectors for downstream electrical cables,
- 2) boards with 12 RJ45 connectors for downstream electrical cables and 8 SFPs for downstream optical cables, and
- 3) boards with 12 RJ45 connectors for downstream electrical cables, four optical output ports, and two optical input ports for receiving clock and trigger information from the upstream FTSW.

Fig. 4 shows the topology of the TTD system. FTSW boards communicate with each other using either electrical

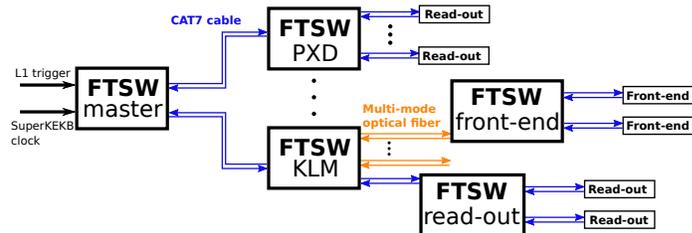


Fig. 4. Topology of the TTD system. Multi-layer separation by subdetectors gives us a flexibility to define the experiment's configuration.

CAT-7 cables or multimode optical fibers depending on the hardware type used. The master FTSW receives the clock signal from the SuperKEKB and the level-1 trigger from the trigger system and distributes this information downstream to the FTSW boards responsible for specific subdetectors which then distribute the information to the front-end and read-out electronics of the subdetector. This layout allows us to change the configuration of the read-out system by including or excluding separate subdetectors in the configuration of the master FTSW.

The physical interface in the TTD system consists of three LVDS (low-voltage differential signaling) signals on the RJ45 connector: clock, trigger, and acknowledgment. With these three signals, FTSW communicates with the front-end and read-out electronics using Belle II trigger timing protocol, b2tt. The clock signal is a 127 MHz clock derived from the 508 MHz accelerator clock. Trigger and synchronization information on the trigger line, TRG, is transmitted as an 8b10b-encoded serial data stream with the data rate of 254 Mb/s from FTSW to the downstream subsystem. The most-significant bit of each decoded byte carries a trigger flag. If the trigger flag is set, then the three least-significant bits carry the precise timing of the event within this byte, and the remaining four bits carry the trigger type. If the trigger flag is not set, the remaining 7 bits of the byte are used to construct a 77-bit frame out of eleven received bytes. The information within the frame can reset the receiver device, synchronize event and time counters on the receiver, or access internal registers of the receiver device.

The monitoring and flow control information on the acknowledgment line, ACK, is sent by a downstream device and received by an FTSW as a serial data stream with the

same parameters as the synchronization stream. Each FTSW decodes this information, multiplexes status from all active links, and transmits it upstream. The master FTSW decides to stop trigger distribution based on the information collected by the downstream FTSWs.

IV. TRIGGER AND TIMING INTERFACE IN THE READ-OUT SYSTEM

Fig. 3 shows the interface to the TTD in the PCIe40 and its connection to the data processing logic in FPGA. The TTD interface in the PCIe40-based read-out system consists of the physical interface and a b2tt core. The physical interface handles electrical signals on the board. The b2tt core consists of a physical, a link, and a logical layers. The physical layer deserializes received data and serializes data to be sent to the FTSW. The link layer takes care of decoding precise trigger information and addressing between FTSW and the receiver, allowing FTSW to communicate with a specific device. The logical layer provides a command and register-based communication interface to the FPGA logic. This section describes the details of the physical interface and the layers of the b2tt core.

A. Physical Interface

PCIe40 board which was originally designed for the LHCb experiment [17] lacks the RJ45 connector necessary for the TTD interface. We use a custom-built adapter cable that has RJ45 on one end for connection to the TTD and an 18-pin Molex connector on the other end for connection to PCIe40. Fig. 5 shows the PCIe40 board with the adapter cable used for interfacing the board with the TTD system. Fig. 3 shows the b2tt signals and clock distribution on the PCIe40. The TRG and ACK signals are connected directly to the FPGA pins. The clock signal, CLK, which enters the board, is received by a clock buffer which distributes the FTSW clock to the external PLLs (phase-locked loop) and to the FPGA inputs.

The external PLL Si5344 [18] delivers the 127 MHz clock for the b2tt core. There are two configurations that can be loaded into this PLL by the slow control software. The first configuration sets the PLL in the jitter cleaner mode which uses the FTSW clock as a reference. This configuration generates the clock signal follows the frequency of the FTSW clock and, therefore, it is the only configuration which can be used to supply clock for the logic-based b2tt core, which is described in the next section in more details. The second configuration sets the PLL into a clock generator mode, which generates the 127 MHz clock from the free-running 54 MHz onboard oscillator. This configuration can only be used by the b2tt core with clock-data recovery.

Other two external PLLs Si5345 [18] deliver a 127 MHz reference clock for the high-speed links to the front-end electronics. There are two configurations that can be loaded into these PLLs by the slow control software. Both configurations set the PLLs in the jitter cleaner mode. The first configuration sets the PLL to use the FTSW clock as a reference, and the second sets the PLL to use the clock from the FPGA. The clock supplied by the FPGA is a 127 MHz output clock of the Si5344 PLL.

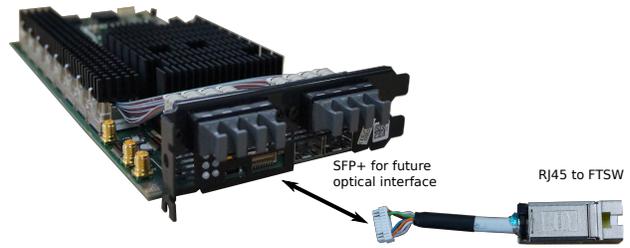


Fig. 5. The PCIe40 board and the adapter cable for the TTD interface.

B. Physical and Link Layers

The first version of the trigger timing interface used in the read-out system was ported to Arria 10 FPGA from the code developed for the Xilinx-based FPGAs of the front-end electronics. The code includes serialization/deserialization interface using logic resources of the FPGA, 8b10b decoder/encoder, and a link layer that decodes packets with trigger information and encodes monitoring and flow control information.

Fig. 6(a) shows the layout of the physical layer of the logic-based b2tt core. The deserializer uses a 508 MHz clock, generated in the internal IOPLL [19] from the incoming 127 MHz clock, to find the phase of the data. Once the phase is determined, the deserializer uses the same clock to receive 2 bits of data per cycle of the 127 MHz clock and performs word alignment by using feedback information from the 8b10b decoder. The 8b10b decoder decodes data and identifies comma words that are used as packet delimiters.

The link layer decodes TTD frames which carry precise trigger timing and supplemental trigger information or allows the FTSW to address a specific b2tt core for access to the internal registers of the system. The internal registers which manage the status and configuration of the system are controlled by the logical layer described later.

The transmitter side of the link layer encapsulates packets provided by the logical layer with comma symbols, encodes data with the 8b10b encoder, and sends them using the DDR (Double-data rate) implementation in FPGA logic. The DDR implementation uses the 508 MHz clock to send data synchronously to the 127 MHz clock.

C. Trigger Timing Core with Clock-Data Recovery

During commissioning, we observed distortion of the electrical TTD signals which is caused by induced noise in the 10 m CAT-7 cable and can lead to desynchronization of the link. We observed 20 link losses during physics data taking in the autumn 2021 run.

Figure 10(b) shows the period distribution of the clock received from FTSW by FPGA and sent back over an unused LVDS line of the CAT-7 cable. Multiple overlaying peaks hint at the crosstalk with the data signal in the cable. The measurement is described in more details in section V.

To mitigate this situation we designed a second version of the core which recovers a clock signal from the incoming data stream thus allowing us to use a free-running clock. The free-running clock is generated on the board and does not suffer

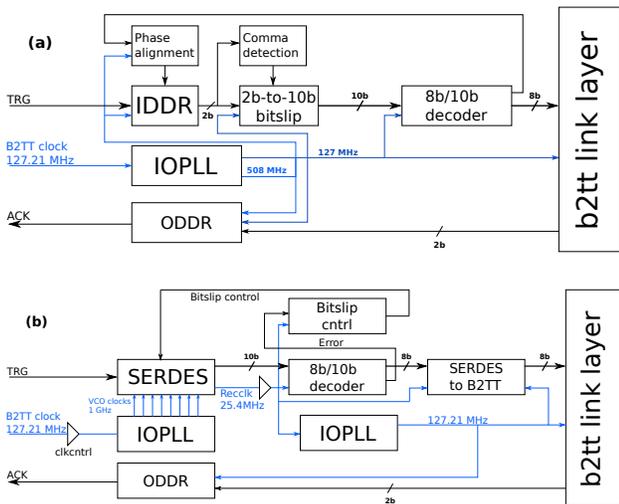


Fig. 6. Architectures of the physical layers of the logic-based b2tt core (a) and the soft-CDR-based b2tt core (b).

from the crosstalk in the cable. This core replaces only the physical layer of the b2tt core but keeps the link and the logical layers.

Fig. 6(b) shows the layout of the new core. This version uses a soft-CDR (clock-data recovery) functionality of the SERDES (serializer/deserializer) [19] in the GPIO (general-purpose input/output) of the Intel Arria 10 FPGAs, and the output DDR registers. The 127 MHz reference clock for the soft-CDR is synthesized from the external PLL which uses the 54 MHz free-running onboard oscillator to generate the reference clock.

The input pins for the reference clock are located in a different IO (input/output) bank than the data pins. Since the IOPLL is required to be instantiated in the same bank as the SERDES, it is impossible to supply the clock to the IOPLL directly from the input pins as recommended by the user manual [19]. We supply the reference clock to the global clock buffer to deliver the clock to the IOPLL in the proper IO bank. While being not optimal in terms of input jitter, the core with this configuration remains functional.

The IOPLL which generates a set consisting of eight 254 MHz clocks, phase-shifted by 45° of the serial clock with respect to each other. The IOPLL is set to a direct mode to minimize jitter at the PLL's output [19]. Fig. 7 shows the phase alignment of the clocks in the soft-CDR receiver. These clocks are distributed to the SERDES over the dedicated clock buffers thus keeping perfect phase alignment at high frequencies. The soft-CDR within the SERDES continuously selects one of these clocks whose phase matches most to the phase of the incoming data to be used as a recovered clock. The SERDES provides the received 10-bit data word synchronous to the parallel clock of 25.4 MHz generated by dividing the recovered clock.

We decode the data in the 8b10b decoder and send the status of the decoded data to the bitslip controller for data alignment. The bitslip controller checks whether the 10-bit data word is a valid 8b10b character. We consider the link established if we receive more than 250 consecutive valid characters. If we receive an invalid character, we activate the bitslip circuit of

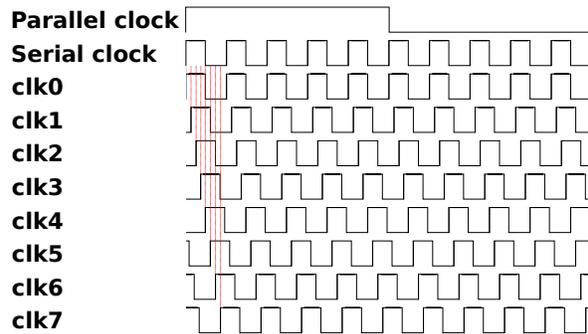


Fig. 7. Phase-shifted clocks used in the soft-CDR receiver. The soft-CDR logic continuously monitors the phase of the data signal and selects the clock signal with the best matched phase.

the SERDES to trigger realignment. The bitslip controller then waits for four clock cycles to ignore invalid data during the transition period before starting to monitor the data decoding status again.

During deployment of the new b2tt core, we discovered that the b2tt system might send a single invalid 8b10b character during the reset of the b2tt link. Therefore, we implemented a debouncing logic which ignores a single invalid character and only triggers realignment if we receive more than one consecutive invalid character.

After decoding data, the core sends them to the link layer which is clocked with a 127 MHz clock through a clock-domain crossing within the SERDES-to-B2TT block. We derive the 127 MHz clock from the recovered 25.4 MHz clock in the internal PLL. Because both clocks are related, we directly sample data on the rising edge of the slow clock in the 127 MHz clock domain.

Due to the use of the SERDES in the b2tt receiver which is synchronous to the slow 25.4 MHz clock, the data reaches the link layer later than the data received by the original b2tt core. This constant latency is 173 ns which corresponds to 22 cycles of the 127 MHz clock.

We have upgraded the output logic of the physical layer to use the DDR component for sending data. This component runs with the 127 MHz clock to send 2 bits of data per clock cycle. Therefore, this implementation eliminates the use of the 508 MHz clock in the design to make achieving timing closure easier.

D. Logical Layer

In addition to the single-channel code, which we ported from the front-end systems, we implemented a logical layer on top of the b2tt core which allows us to monitor the status of multiple channels and use a ring buffer to distribute trigger information to all input channels in parallel. The logical layer decodes commands and corresponding data from the packets. The protocol defines three commands:

- mask,
- data select, and
- register select.

The mask command sets the channel mask to deactivate certain channels. The data select command selects the channel. The

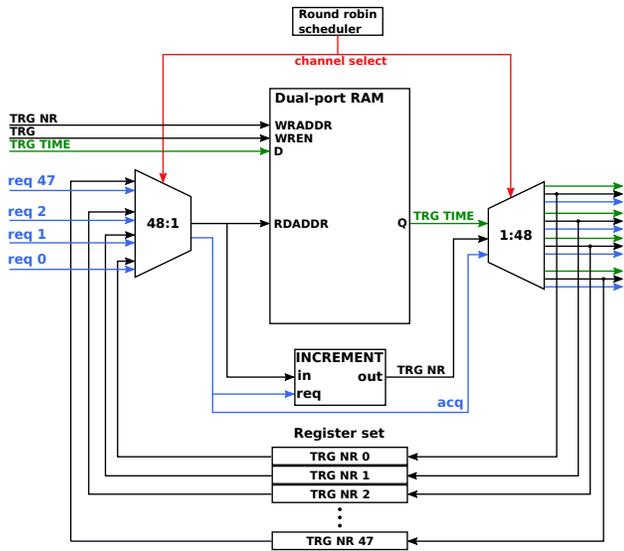


Fig. 8. Layout of the ring buffer. We store trigger timing information in the shared memory to reduce the memory usage of the design.

register select command selects a register of the channel. The data from the selected register are then combined to a packet and supplied to the link layer encoder.

Since there are up to 48 data channels, distributing trigger information to the data processing logic of each channel in parallel would be expensive in terms of memory resources. For example, to store trigger and timing information which uses 96 bits per event for 1024 events and 48 channels would require 589 kB. Therefore, we designed a ring buffer to reduce the memory resources used by the design.

Fig. 8 shows the layout of the ring buffer. The ring buffer stores up to 1024 triggers and trigger timing information and makes the data accessible to the data processing logic of each channel. We store trigger timing information in the dual-port RAM (Random Access Memory) and access these data by addressing the memory by the least significant bits of the trigger number. Therefore, these data are shared by all channels. The trigger number itself is not stored in the RAM, but we keep up to 48 sets of registers with the number of the last read trigger and the number of stored triggers. The number of stored triggers is incremented by b2tt every time a new trigger arrives. This number decrements every time the data processing logic of a channel requests a trigger. When this happens, the last read trigger number in this set increments. The algorithm works due to the sequential nature of the trigger numbers received from b2tt. The ring buffer keeps track of the channels with the maximum and the minimum trigger number to monitor the fill level of the buffer. If the difference reaches a given critical threshold, the busy signal is sent to FTSW to pause trigger distribution until the fill level falls below the threshold.

E. Fiber Optics as a Replacement for the CAT-7 Cable

The future implementation of the b2tt will use optical fibers to deliver the trigger and clock information to the front-end and read-out systems. Optical fibers will help to improve signal

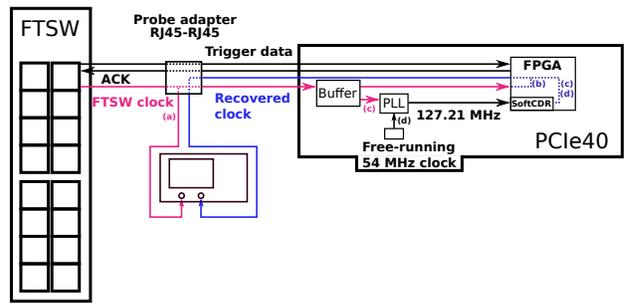


Fig. 9. Layout of the jitter measurements of the b2tt and the recovered clocks.

quality and avoid link desynchronization problems. To use optical signals in PCIe40 we started the development of the physical interface involving high-speed transceivers. Because the data rate of the link is below the operational frequency of the transceiver, we will have to use an oversampling mode of the transceiver and implement decoding logic which will probably further increase trigger latency.

V. CHARACTERIZATION OF THE INTERFACE AND SYSTEM OPERATION IN BELLE II

A. Jitter Measurement

Clock period distortion, also known as clock jitter, is an important parameter that influences the stability of the system. Therefore, we characterized the clock by measuring the clock period distributions at different stages of clock processing to identify noise sources and compare the performance of both b2tt core implementations.

Fig. 9 shows the layout of the setup for the jitter measurement of the clocks in the system. The FTSW provides the 127 MHz clock and encoded trigger data stream, and receives the monitoring and flow control information over a CAT-7 cable. A custom RJ45-to-RJ45 adapter, inserted between the CAT-7 cable from the FTSW and the CAT-7 cable from the PCIe40, provisions connections for oscilloscope probes to spy on the signals in the cable. We use the LeCroy WavePro 7 Zi oscilloscope [20] to measure the clock period distributions.

The FPGA is configured to output the clock on the unused cable pair of the CAT-7 cable. We prepared two firmware versions for this measurement: the first version outputs the FTSW clock which is received directly from the clock buffer without data processing; the second version outputs the clock recovered by the soft-CDR core. In addition, we can reconfigure the PLL to use either the FTSW clock or the external oscillator as a reference clock in the PLL.

Fig. 10 shows the measured clock period distributions. Fig. 10 (a) shows the distribution of the clock period measured at the output of FTSW. The shape of the distribution is Gaussian with a standard deviation of 21.89 ps. Fig. 10 (b) shows the distribution of the FTSW clock period with the signal returned from PCIe40 without any processing. This distribution shows an irregular shape with 5 recognizable modes. This shape points out the interference with the data lines within CAT-7 cable. Fig. 10 (c) shows the distribution of the period of the clock recovered by the soft-CDR core

Table I
FWHM FOR THE MEASURED CLOCK PERIOD DISTRIBUTIONS

Figure	Clock source	FWHM, ps	Visible crosstalk
10, a	FTSW, direct	48	-
10, b	FTSW, returned	370	5 modes
10, c	FTSW, recovered in soft-CDR	264	2 modes
10, d	Oszillator, recovered in soft-CDR	108	-

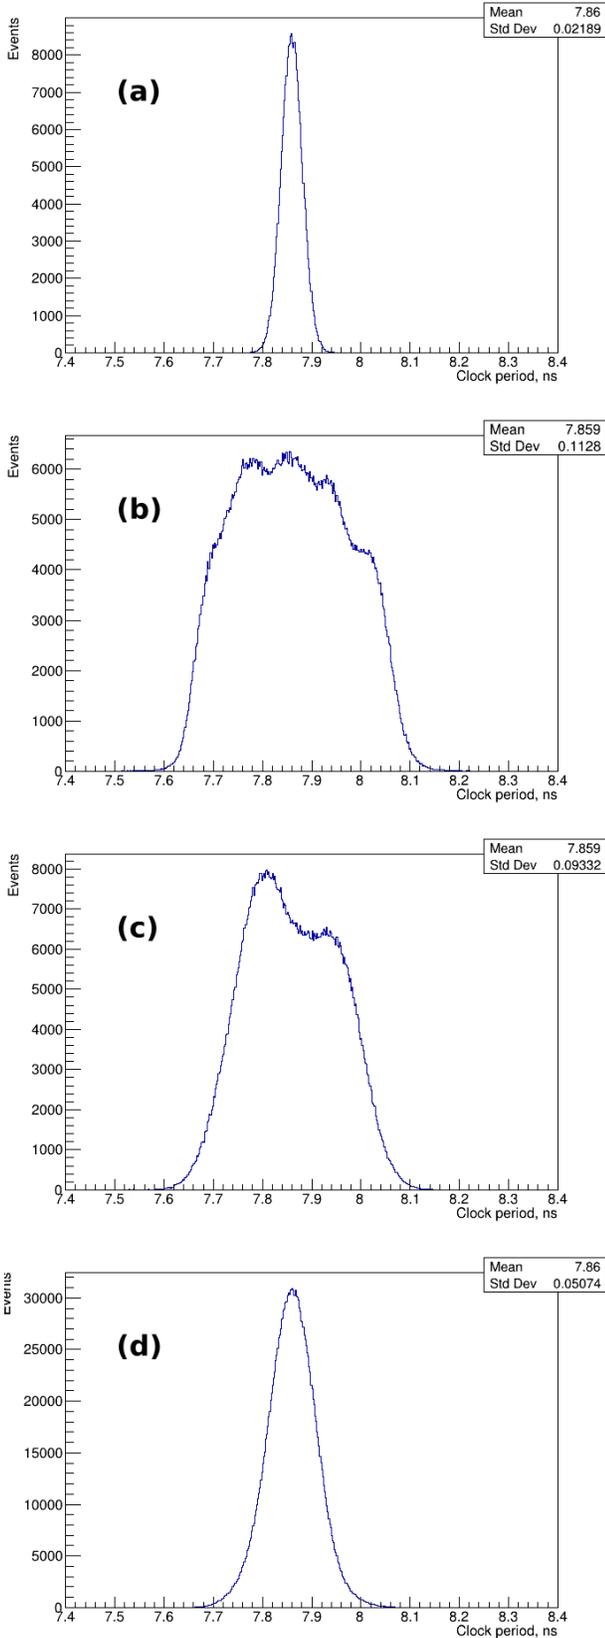


Fig. 10. Clock period distributions for the direct FTSW clock (a), FTSW clock received by FPGA and returned through CAT-7 cable (b), clock recovered from data using FTSW clock cleaned by the external PLL (c), clock recovered from data using clock synthesized in the external PLL from the oscillator (d).

using the reference clock generated in the external PLL from the FTSW clock. This distribution also shows an irregular shape but has only 2 visible modes. Fig. 10(d) shows the distribution of the period of the clock recovered by the soft-CDR core using the reference clock generated in the external PLL from the external oscillator. The shape of this distribution is Gaussian with a standard deviation of 50.74 ps.

To compare these distributions, we will use full width at half maximum (FWHM) as a quantitative value due to the irregularity of the shape for two distributions. Table I summarizes the measurements. The measurement shows that while the external PLL can attenuate the jitter and remove some of the deterministic noise components in the signal, the noise still propagates to the system and can be measured. Using the clock, generated on the board, the recovered clock does not show noise components and, therefore, this clock is preferable for use as a reference clock for the high-speed SERDES in Belle2Link after jitter cleaning in the external PLLs Si5345 of the PCIe40.

This measurement points out the advantage of the soft-CDR-based approach. Because the soft-CDR-based core recovers clock from data, it can use independent clock sources, and thus profits from the low-noise clock. The logic-based core must by design use the clock delivered with data, and, therefore, its performance will be affected by the noise picked up by the signal in the cable.

During the measurements, we have also observed that due to the data being transmitted at 254 Mb/s, the phase difference between the 127 MHz clock derived from the recovered clock and the FTSW clock may be 0° or 180° . Since precise phase synchronization is not important for the trigger timing interface of the PCIe40, we did not implement an additional phase synchronization logic.

B. System Operation in Belle II

To describe the system stability, we count the number of desynchronization events, which we call link drops, during system operation. A link drop refers to the condition when the receiver receives an invalid 8b10b character and goes out of synch with the sender. Each desynchronization causes a stop of the physics data taking and associated dead time of approximately 2 minutes needed to restart the DAQ.

We started to use the logic-based version of the b2tt core in the autumn run in 2021 with the TOP and KLM subdetectors. The system consisted of 3 PCIe40 boards during the physics data taking. The PCIe40-based event builders were running without major issues with the trigger rates up to 12 kHz. We have observed 20 drops of the b2tt links during 38 days of pure data taking. Some of these link drops were correlated

with the maintenance work close to the read-out electronics which points to the external noise as a reason for the link drops.

These problems lead to the development of the soft-CDR version of the b2tt core. We have used this core for 22.4 days of pure data taking at Belle2 during the 2022 run with 5 PCIe40 boards in the TOP, KLM, and ARICH read-out systems. The system operated stably with the trigger rates up to 12.8 kHz and we have not observed any link drops during physics data taking. There were 3 link drops related to the maintenance work in the vicinity of the system.

To compare the performance of both cores, we can compare the cumulative running time T_c :

$$T_c = n_b \times T_t,$$

with n_b is the number of boards in operation, and T_t is the total run time. The T_c of the logic-based core in the autumn 2021 run is 114 days, which is similar to the T_c of the soft-CDR core of 112 days. If we compare the number of link drops for each core version, we can conclude that the soft-CDR version of the core provides better stability than the logic-based version.

The trigger rate during physics runs is limited by the level-1 trigger. We can also operate PCIe40 with both versions at 30 kHz trigger rate, expected when SuperKEKB will reach full luminosity, with an artificially generated trigger signal.

VI. SUMMARY

We started the upgrade of the read-out system of the Belle II experiment using the PCIe40 board. We successfully upgraded three sub-detectors and will finish the upgrade of the read-out system for the remaining three sub-detectors during the long shutdown 1 of the SuperKEKB accelerator, foreseen in from 2022 to 2023. To provide synchronization with the front-end electronics, we implemented a TTD interface in the Arria 10 FPGA. We use clock-data recovery to reduce the number of noise sources in the system to prevent loss of synchronization. We have characterized the performance of the interface and successfully operated the PCIe40-based event builder at the Belle II during physics data taking. In the future, we plan to use optical fibers for TTD signals and have started the development of a new version of the core.

REFERENCES

- [1] T. Abe, I. Adachi, K. Adamczyk, S. Ahn, H. Aihara, K. Akai *et al.*, "Belle II Technical Design Report," Tech. Rep., 2010. [Online]. Available: <http://cds.cern.ch/record/1304162> Accessed on: December 12, 2022.
- [2] K. Akai, K. Furukawa, and H. Koiso, "SuperKEKB collider," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 907, pp. 188–199, 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900218309616> Accessed on: December 12, 2022.
- [3] C. Marinas, "The Belle II pixel detector: High precision with low material," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 731, pp. 31–35, 2013. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900213003136> Accessed on: December 12, 2022.

- [4] L. Zani, K. Adamczyk, L. Aggarwal, H. Aihara, T. Aziz, S. Bacher *et al.*, "The Silicon Vertex Detector of the Belle II experiment," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1038, p. 166952, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900222003977> Accessed on: December 12, 2022.
- [5] N. Taniguchi, "Central Drift Chamber for Belle-II," *Journal of Instrumentation*, vol. 12, no. 06, p. C06014, jun 2017. [Online]. Available: <https://dx.doi.org/10.1088/1748-0221/12/06/C06014> Accessed on: December 12, 2022.
- [6] U. Tamponi, "The TOP counter of Belle II: Status and first results," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 952, p. 162208, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900219307053> Accessed on: December 12, 2022.
- [7] Y. Yusa, "ARICH for Belle II," *Journal of Instrumentation*, vol. 9, no. 10, p. C10015, oct 2014. [Online]. Available: <https://dx.doi.org/10.1088/1748-0221/9/10/C10015> Accessed on: December 12, 2022.
- [8] Belle-ECL, V. Aulchenko, A. Bobrov, A. Bondar, B. G. Cheon, S. Eidelman *et al.*, "Electromagnetic calorimeter for Belle II," *Journal of Physics: Conference Series*, vol. 587, no. 1, p. 012045, feb 2015. [Online]. Available: <https://dx.doi.org/10.1088/1742-6596/587/1/012045> Accessed on: December 12, 2022.
- [9] T. Aushev, D. Besson, K. Chilikin, R. Chistov, M. Danilov, P. Katrenko *et al.*, "A scintillator based endcap K_L and muon detector for the Belle II experiment," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 789, pp. 134–142, 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S016890021500385X> Accessed on: December 12, 2022.
- [10] S. Yamada, R. Itoh, T. Konno, Z. Liu, M. Nakao, S. Y. Suzuki *et al.*, "Common Readout Subsystem for the Belle II Experiment and Its Performance Measurement," *IEEE Transactions on Nuclear Science*, vol. 64, no. 6, pp. 1415–1419, 2017.
- [11] T. Geßler, W. Kühn, J. S. Lange, Z. Liu, D. Münchow, B. Spruck *et al.*, "The ONSEN Data Reduction System for the Belle II Pixel Detector," *IEEE Transactions on Nuclear Science*, vol. 62, no. 3, pp. 1149–1154, 2015.
- [12] J. Cachemiche, P. Duval, F. Hachon, R. L. Gac, and F. Réthoré, "The PCIe-based readout system for the LHCb experiment," *Journal of Instrumentation*, vol. 11, no. 02, pp. P02013–P02013, feb 2016. [Online]. Available: <https://doi.org/10.1088/1748-0221/11/02/p02013> Accessed on: December 12, 2022.
- [13] Intel Corporation, *Intel Arria 10 Device Overview*, v. 2020.10.20.
- [14] S. Huber, I. Konorov, D. Levit, S. Paul, and D. Steffen, "Performance of the Data-Handling Hub Readout System for the Belle II Pixel Detector," *IEEE Transactions on Nuclear Science*, vol. 68, no. 8, pp. 1961–1967, 2021.
- [15] D. Sun, Z. Liua, J. Zhao, and H. Xu, "Belle2Link: A Global Data Readout and Transmission for Belle II Experiment at KEK," *Physics Procedia*, vol. 37, pp. 1933–1939, 2012. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1875389212019104> Accessed on: December 12, 2022.
- [16] M. Nakao, "Timing distribution for the Belle II data acquisition system," *Journal of Instrumentation*, vol. 7, no. 01, pp. C01028–C01028, jan 2012. [Online]. Available: <https://doi.org/10.1088/1748-0221/7/01/c01028> Accessed on: December 12, 2022.
- [17] The LHCb Collaboration, A. A. Alves Jr, L. M. A. Filho, A. F. Barbosa, I. Bediaga, G. Cernicchiaro *et al.*, "The LHCb Detector at the LHC," *Journal of Instrumentation*, vol. 3, no. 08, p. S08005, aug 2008. [Online]. Available: <https://dx.doi.org/10.1088/1748-0221/3/08/S08005> Accessed on: December 12, 2022.
- [18] Skyworks Solutions, Inc., *Si5345, Si5342 Rev. D Family Reference Manual, Rev. 1.3*.
- [19] Intel Corporation, *Intel Arria 10 Core Fabric and General Purpose I/Os Handbook*, v. 2021.07.16.
- [20] Teledyne LeCroy, Inc., *WavePro 7 Zi/Zi-A Oscilloscopes Operator's Manual, 927899 Rev B*.