

CERN OPEN HARDWARE EXPERIENCE: UPGRADING THE DIAMOND FAST ARCHIVER

I.S. Uzun and M. Abbott, Diamond Light Source, Oxfordshire, UK

Abstract

Diamond Light Source developed and integrated the Fast Archiver [1] into its Fast Orbit Feedback (FOFB) communication network in 2009. It enabled synchronous capture and archive of the entire position data in real-time from all Electron Beam Position Monitors (BPMs) and X-Ray BPMs. The FA Archiver solution has also been adopted by SOLEIL and ESRF. However, the obsolescence of the existing PCI Express (PCIe) based FPGA board from Xilinx and continuing interest from community forced us to look for a new hardware platform while keeping the back compatibility with the existing Linux kernel driver and application software. This paper reports our experience with using the PCIe SPEC card from CERN Open Hardware initiative as the new FA Archiver platform. Implementation of the SPEC-based FA Archiver has been successfully completed and recently deployed at ALBA in Spain.

INTRODUCTION

The Diamond FOFB system is operating successfully and achieving the required level of suppression of beam disturbances since 2007 [2]. The system architecture for the FOFB for a single Storage Ring (SR) cell is shown in Fig. 1. Diamond SR consists of 24 cells, and overall FOFB architecture consists of:

- 172 Electronic Beam Position Monitors (BPMs) located around the storage ring.
- 24 VME-based processors, each populated with a PMC receiver, running the feedback algorithm.
- Power supply controller units for 168 slow corrector magnets.
- Communication network to distribute the beam position values.

BPMs are used for electron beam diagnostics. In the BPM FPGA the sampled button intensities are converted to X, Y positions and reduced by decimation and filtering to machine revolution frequency (534 kHz) and then to the “Fast Acquisition” (FA) data rate of 10 kHz and the “Slow Acquisition” rate of 10 Hz for EPICS clients. The FA data is broadcast by Diamond Communication Controller IP (DCC) [3] on the FPGA through a dedicated communication network, which makes this data stream simultaneously available to all PMC receiver nodes on VME processors for fast feedback processing.

Following the successful operation of FOFB system, a need for recording longer periods of synchronous position data from all nodes on the communication network (the whole orbit) at FA data rate was identified. The DCC IP and

dedicated communication network enabled us to additional nodes with little or no impact on the operation of the FOFB; and the FA Archiver is one such node implemented on a PCIe hardware [1].

Two FA Archiver nodes have been running at Diamond, one for Booster and one for Storage Ring since 2009, giving us access to records of up to two weeks of position data across the both rings. We have been using this data for:

- Identification and detailed investigation of any orbit events that happened in the recent past.
- Correlation of orbit events with other measurements by adding X-ray BPMs and RF cavity voltage readings to the FA network.
- Model analysis of orbit motion over longer periods for the optimisation of the fast orbit feedback.

FA Archiver interface was implemented on a commercial Xilinx ML555 PCIe FPGA [4] board as a cost-effective solution so that we could focus our efforts on firmware development and software integration. Following its successful commissioning at Diamond, SOLEIL and ESRF integrated FA Archiver as part of their fast feedback system. This was possible since both facilities built their FOFB data distribution network using DCC giving them the protocol compatibility required.

Recent interest from ALBA as the latest user of DCC, and Brazilian Light Source (LNLS) for evaluation purposes brought the obsolescence of the ML555 platform to our attention. In order to support the community, we took on an upgrade project targeting implementation of FA Archiver on a modern platform with longer term support and availability. Following evaluation of commercially available products and in-house built possibilities, we have decided that SPEC card from CERN Open Hardware initiative would be the best

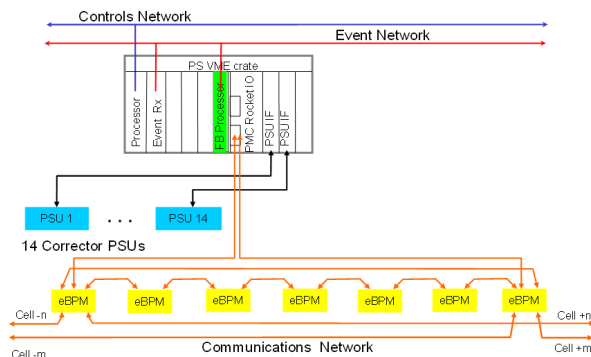


Figure 1: FOFB architecture for single SR Cell. It is repeated for each 24 cells in Diamond SR.

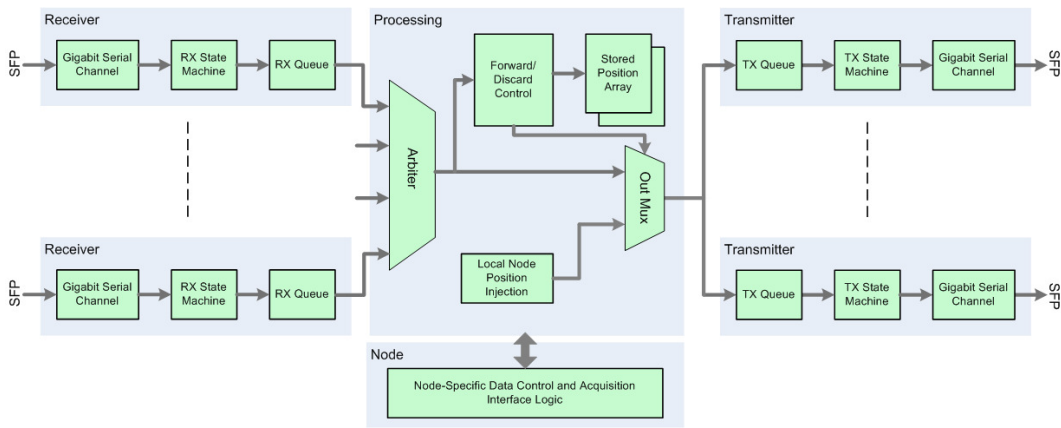


Figure 2: Major building blocks of DCC design.

candidate due to its wide use in the accelerator community for various projects and openness of the design sources.

DIAMOND COMMUNICATION CONTROLLER

To realise the required data transfer (BPMs to computation nodes) at a 10 kHz update rate, DCC IP was designed and implemented in VHDL. The IP uses Multi-gigabit Serial Transceivers (MGTs) on the FPGA achieving communication rate of 2.12 Gbps for Diamond.

DCC utilises a data-forwarding protocol whereby each node forwards its own data, plus each piece of incoming data once. The data is encapsulated in a low-overhead packet including CRC. The architectural block diagram of the DCC is shown in Fig. 2. The IP design consists of three main blocks:

- Receiver block implements up to 4 input channels, and is responsible for de-serialising, decoding, CRC checking and queuing the incoming packets.
- Data Processing (Store, Analyse and Forward) keeps tracks of incoming packets from each node on the communication network, and transmits a specific node's packet only once in a given timeframe. This approach prevents flooding the communication network. It also stores received values for feedback or archiving purposes.
- Transmitter block implements output queues for each channel. It is responsible for adding CRC, encapsulating and encoding position data for transmission through the serial links.

A dedicated fibre network between BPMs was chosen as a cost effective method of achieving low-latency communication. The BPMs have four SFP multi-gigabit serial transceivers, the 24 storage ring cells are wired in a torus network with single-mode fibre while copper patch cables are used intra-cell. The DCC network topology is structured as a 2D torus, with one computation node per SR cell, and gives a degree of resilience to failure of single or combinations

of BPMs or links. The communication network structure is shown in Fig. 3.

DIAMOND FA ARCHIVER

The design of the DCC and communication network makes it easy to add new nodes to the network, both as contributors and as passive listeners. The FA archiver acts as a passive listener storing the data stream in a 30TB ring buffer. The stored data is then available to clients via a set of software tools.

Implementation of FA archiver interface requires a PCIe form factor hardware platform with following on-board resources:

- Minimum 1-lane of PCIe bus interface. This could be either through dedicated IPs blocks on an FPGA device or a bridge chip external to the FPGA.
- At least 1 MGT on the FPGA for DCC implementation.
- On-board SFP connector(s) for physical connection to the communication network.
- Clocking resources required for PCIe, DMA data transfer and DCC FPGA logic.




Figure 3: Communication network topology, with each circle being one cell consisting of seven BPMs, and a PMC receiver in the centre. The coloring represents total distribution time for each node.

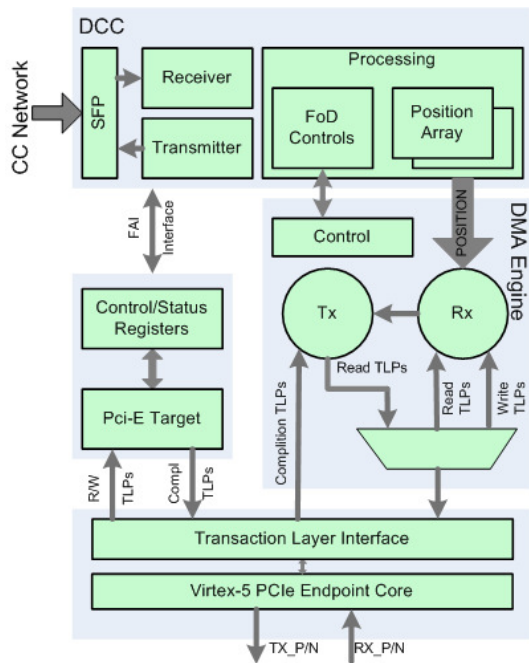


Figure 4: Architecture of the FA Archiver FPGA design on ML555 board. The design directly interfaces to PCIe physical connector.

The FPGA design for the FA Archiver is shown in Fig. 4. It integrates the DCC into a bus master PCIe architecture on the ML555 FPGA board. The design consists of target logic, DMA initiator logic, status/control registers and the Virtex-5 endpoint core for PCIe. Target logic is responsible for capturing single doubleword memory write and memory read PCIe Transaction Layer Packets (TLPs) for control and status register access. The DMA initiator logic generates memory write TLPs to transfer 4K byte frame data from the DCC core to the host's system memory through PCIe endpoint.

CERN OPEN SOURCE HARDWARE

Open source hardware (OHW) is hardware whose design is made publicly available so that anyone can study, modify, distribute, make, and sell the design or hardware based on that design [5]. Open hardware has in recent years established itself as a very effective way for CERN to make electronics designs and in particular printed circuit board layouts, accessible to anyone, while also facilitating collaboration and design re-use. It is creating an impact on many levels, from companies producing and selling products based on hardware designed at CERN, to new projects being released under the CERN Open Hardware Licence (OHL) [6]. Today the open hardware community includes large research institutes, universities, individual enthusiasts and companies. Many of the companies are actively involved in the entire process from design to production, delivering services and consultancy and even making their own products available under open licences.




Figure 5: Commercially available SPEC card.

Open Hardware Repository

Recognizing the benefits that an open source philosophy could provide to hardware development and dissemination, CERN created the Open Hardware Repository(OHR) in 2009. It is a place on the web for electronics designers at experimental physics facilities to collaborate on open hardware designs. Technical specifications, research documents, CAD drawings and HDL code relating to projects are available for download from the OHR version control system. The Open Hardware Repository currently hosts more than 100 projects, ranging from small projects with a few collaborators to bigger projects with multiple contributors from both industry and academia. A dozen companies are actively involved in projects in the OHR and some of them are producing the physical hardware for CERN and other customers.

One of the prominent projects within OHW framework is the Simple PCIe FMC carrier (SPEC) card developed by CERN, as shown in Fig. 5. SPEC is a data-acquisition card that can be plugged in to a PC, allowing the use of different FMCs, such as a digital input/output module, a time-to-digital converter and a fine delay module. The SPEC card was designed with the White Rabbit (WR) timing project in mind and is the most widely-sold card offering WR support for synchronization. Several companies sell the card and are involved in its development. SPEC cards have found their way to many scientific applications, such as fusion energy, tracking of space debris and control systems for accelerator complexes.

The SPEC card hardware specification, availability of design sources on OHR along with its low-cost commercial availability made it an attractive solution as the new hardware platform for upgrading Diamond FA Archiver. It was also our aim to interact and build experience with OHW initiative as a user.

FA ARCHIVER ON SPEC CARD

The initial challenges with the OHW repository at the start of our project were:

- Understanding our way around the OHR since hardware, firmware and software projects for the SPEC card are maintained in separate repositories.

- Finding details of the components and features of the SPEC card due to lack of a Hardware User Guide.

By studying the open-source board schematics and reference firmware design available on the repository, we were able to identify the modifications to be taken towards implementing the FA Archiver on the SPEC card.

PCI Express Interface

The Spartan-6 FPGA on the SPEC card doesn't have an embedded PCIe IP block unlike the Virtex-5 FPGA on the ML555 board. As a result, an external PCIe bridge, Gennum GN4124, is used to handle the host communication on the SPEC card. This hardware difference required a completed re-implementation of FA Archiver's existing PCIe and interrupt interface logic so that the communication is achieved by using the external bridge's local interface rather than the embedded Xilinx IP block.

Clocking Resources

The SPEC card was design specifically targeting the White Rabbit project, therefore the clocking resources around the FPGA was built on WR requirements which wouldn't suit the DCC. To achieve the required clocking frequency of 106.25MHz for DCC, an external on-board programmable oscillator is used. This was achieved by modifying the DCC firmware to route the clock using on-chip global buffers rather than dedicated MGT clocking pins. This scheme introduces unwanted jitter on the clock, but since MGTs on the FPGA are qualified up to 6.8 Gbps, its effect at the DCC line rate of 2.12 Gbps does not create any unwanted bit errors rates. This modification on the clock routing created additional conflicts for the main FPGA system clock input. This was solved by eliminating the use of the dedicated system clock inputs, and using the PCIe bridge's local clock as the FPGA system clock.

On-Chip Communication

The reference firmware design for the SPEC card uses the Wishbone System-on-Chip (SoC) bus for interconnecting IP cores. The FA Archiver design uses ARM's Advanced Microcontroller Bus Architecture (AMBA) which is the adopted solution for interconnecting Xilinx IPs. Towards achieving the maximum re-use of reference design IPs, FA Archiver's Target and DMA Control Logic design blocks handling the status and control registers for operating the board were also modified to adopt the Wishbone bus.

The resulting FPGA design architecture for the FA Archiver on the SPEC card is shown in Fig. 6.

Software

The only software modification was in the device driver to detect and identify the card installed (ML555 or SPEC), since the SPEC card's PCIe chip has to be initialised before bringing the interface up. The archiver processing software and client tools as explained in [1] didn't require any modifications keeping the back-compatibility.

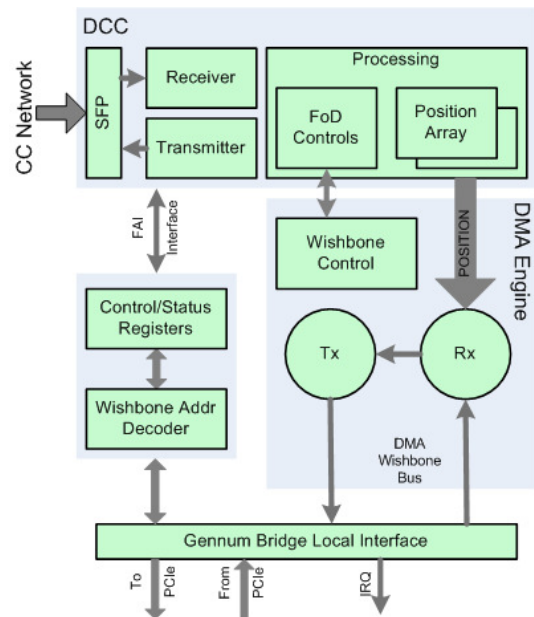


Figure 6: Architecture of the FA sniffer FPGA on SPEC card. The design interfaces to an external PCIe bridge using the parallel local interface.

SUMMARY

The motivation for the new FA Archiver interface has been to address obsolescence of the existing PCIe platform, and move to an open platform for continuing support in the community. We took this opportunity to get familiar with Open Hardware initiative, and build experience with the inner workings of the repository. The upgrade project has been successfully completed with initial challenges due to learning stages involved. The new FA Archiver has been running at ALBA for the last year, and will be evaluated by LNLs with DCC running at an increased line rate of 5 Gbps.

REFERENCES

- [1] M.G. Abbott, G. Rehm, I. Uzun, *A New Fast Data Logger and Viewer at Diamond: the FA Archiver*, Proc. of ICALEPCS'2009, Grenoble (France).
- [2] M.T.Heron, et al., *Diamond Light Source Electron Beam Position Feedback: Design, Realization and Performance*, Proc. of ICALEPCS'2009, Kobe (Japan).
- [3] I. Uzun, et al., *Initial Design of the Fast Orbit Feedback System for Diamond Light Source*, Proc. of ICALEPCS'2005, Geneva (Switzerland).
- [4] Xilinx, *Virtex-5 LXT ML555 FPGA Development Kit for PCI Express*, <http://www.xilinx.com/products/boards-and-kits/HW-V5-ML555-G.htm>
- [5] CERN, *Open Hardware At CERN-Knowledge Transfer*, <http://knowledge-transfer.web.cern.ch/files/OH-web.pdf>
- [6] CERN, *Open Hardware License*, www.ohwr.org/cernohl