ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-MOPHA049

TEST-BENCH DESIGN FOR NEW BEAM INSTRUMENTATION ELECTRONICS AT CERN

M. Gonzalez-Berges, J. O. Robinson, M. Saccani, V. Schramm, M. A. Stachon
CERN, Geneva, Switzerland

Abstract

The Beam Instrumentation group has designed a new general-purpose VME acquisition board that will serve as the basis for the design of new instruments and will be used in the renovation of existing systems in the future. Around 1200 boards have been produced. They underwent validation, environmental stress screening and run-in tests to ensure their performance and long term reliability. This allowed to identify potential issues at an early stage and mitigate them, minimizing future interventions and downtime. A dedicated test-bench was designed to drive the tests and continuously monitor each board functionality. One board has more than 45 functions including memories, high speed serial links and a variety of diagnostics. The test-bench was fully integrated with the CERN Asset Management System to allow lifecycle management from the initial production phase. The data captured during these tests was stored and analyzed regularly to find sources of failures. This was the first time that such a complete test-bench was used. This paper presents all the details of the test-bench design and implementation.

INTRODUCTION

The CERN Beam Instrumentation (BI) group is in charge of designing, implementing and maintaining the instruments for the CERN accelerator complex. There are different types of instruments, including Beam Position Monitors (BPMs), Beam Loss Monitors (BLMs), Wire Scanners, etc. Currently there are several thousand instruments deployed in the different accelerators, some of which have been installed for several decades, with many using dedicated electronics. During the construction of the LHC the majority of instrumentation systems [1] were designed to use a common VME acquisition platform, the Digital Acquisition Board (DAB)-64x, which speeded-up developments, provided a standardised software interface and greatly improved maintainability. The largest instrumentation systems using this board are the BLM [2] and BPM [3] systems for the Large Hadron Collider (LHC), with over 4000 and 2500 devices acquired respectively.

Looking towards the future, the group intends to standardize the new electronics systems with the development of another general purpose VME board, the VFC-HD (VME FMC Carrier High Density) [4] Over 1200 units have been produced so far. This board is already being used in the development of new systems, such as the new BPM electronics for the CERN Super Proton Synchrotron, the renovation of the BLM systems in the LHC and its injectors, and is foreseen to be the default acquisition platform for several other new or upgraded instruments.

Since the VFC-HD will be used in many critical systems, the quality and reliability of the production are extremely important. This paper describes the design and implementation of a test-bench that was used to validate the VFC-HD production in terms of both functionality and reliability.

THE VME HIGH DENSITY FMC CARRIER BOARD (VFC-HD)

The VFC-HD (Figure 1) is a generic acquisition board for processing data either from an on-board FMC mezzanine, or via dedicated SFP fibre-optical inputs. Once processed in the on-board FPGA, this data can be stored in an on-board memory or read-out directly via Ethernet or through the VME bus to a host CPU. The high density FMC connector allows instrument specific or standard industrial modules to be hosted, providing great flexibility. The card also has on-board temperature and voltage sensors for self-diagnostics and provides inputs for timing and debugging links such as White Rabbit [5] or Beam Synchronous Time (BST) [6].

The VFC-HD board will be used for the main upgrades of many beam instrumentation systems at CERN in the coming years, offering increased performance, maintainability and reliability. Around 1200 boards have been produced and tested with the following main characteristics:

- Intel ArriaV GX Field Programmable Gate Array (FPGA)
- High Pin Count FPGA Mezzanine Card (FMC-HPC) slot, fully ANSI/VITA compliant:
  - 3 Low/High User Banks A and B (LA, HA & HB)
  - 10 gigabit lanes connected to FPGA transceivers
  - Programmable power supply voltage (Vadj).
- 6x Small Form-factor Pluggable Plus (SFP+) transceiver slots
- 4x GP (General purpose) SFP+ with up to 6.5Gbps
- 1x GP SFP+ with optional clock for White Rabbit Ethernet
- 1x GPSFP+ with optional clock data recovery (CDR) for Beam Synchronous Time (BST) reception
- 2x 4GB DDR3 memories
- Flexible clocking resources: adjustable and programmable Voltage Controlled Oscillators (VCO) and Phase Locked Loop (PLL).
- 30 single ended connections to VME64x P0 to support clock & trigger distribution in (custom) BI LHC VME crates
- 40 single ended (or 20 LVDS) connections to VME P2 available for rear transition modules
- 4 LEMO connector for general purpose input/output on the front panel

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-MOPHA049

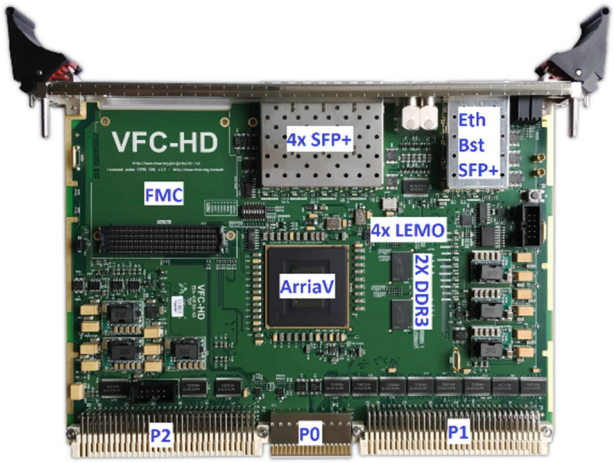

Figure 1: The VFC-HD board.

TEST-BENCH DESIGN AND IMPLEMENTATION

The test procedure for the VFC-HD validation is composed of three main steps.

The first step occurs directly after production and is performed by the manufacturer at its premises on a specific test-bench designed by CERN. This setup allows each feature to be checked one at a time on a single board and generates a complete report. Most of these tests are automatic and take a few minutes per board. This test-bench is not described in this paper.

Then, in a second step, the boards are delivered to CERN where they are visually checked before undergoing a screening sequence of 24 hours in a climate chamber with temperature and humidity cycling. This is intended to trigger any impending faults due to things such as bad soldering or incorrect cleaning that can lead to short circuits.

Finally the boards are installed in a VME crate for a run-in phase of 2 to 4 weeks to detect any infant failures which typically appear at the beginning of the lifetime for most of their electronic components.

For the two types of tests performed at CERN a specific test-bench including hardware, firmware and software was designed and implemented. It consists of the parts described in the following sections.

Hardware

The VFC-HD test bench hardware is similar to the operational hardware used in beam instrumentation systems. It is composed of a VME crate, equipped with standard power supply module and fan tray, a CERN VME backplane and a MEN-A25 CPU board (Intel Pentium D1519, 1.5GHz, 8GB RAM). A total of nine crates were prepared to test the full VFC-HD production over six months. One crate was used for the burn-in procedure and the eight others for the run-in procedures. A full crate can be seen in Figure 2.

Three custom boards were designed in order to test the VME64x user defined connector (P0 connector). These boards were mainly used to pull-up, propagate and read back the two daisy chained signals from all the VFC-HD boards under test.

In order to test the front panel connections, loopbacks were installed on each VFC-HD connector i.e. on LEMO, SFP (3.5dB electrical attenuation) and on the FMC slot with the loopback mezzanine (LA, HA, HB, and clock lines, plus a custom JTAG loopback)

All the other VFC-HD features were tested internally by the FPGA and didn't need any additional hardware. Only the VME64x P2 rear transition modules (40 lines) were not tested by this test-bench, because it would have involved too many extra hardware connections. They are nevertheless tested right after production by the manufacturer.

Figure 2: Full VME crate equipped with VFC-HD boards in the climate chamber.

Firmware

A specific ArriaV firmware (Figure 3) was designed for this test-bench. The goal was to be able to test every component on the board and as many hardware lines as possible in an autonomous way and all in parallel, so that power consumption, noise and temperature would create worst case scenario of operation conditions.

Thus the FPGA firmware was built in such a way that it generates patterns for all its interfaces in parallel and constantly checks their behaviour.

More than 20 hardware components and their links are tested by this firmware such as all high speed serial links, GPIO, DDR3, EEPROM, Flash, external PLL and oscillators, ADC, all FMC lines, VME bus including IRQs and custom lines (i.e. P0), temperature and voltage sensors, etc. The unused resources of the FPGA; multipliers-accumulators, Block RAMs and Logic, have also been filled by arithmetic functions and comparators to operate the chip under full stress conditions. All the results, including status, errors and test counters, are stored in a simple 1kB array of registers quickly readable by the software in one multiplexed block transfer access, protected with a checksum code.

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-MOPHA049

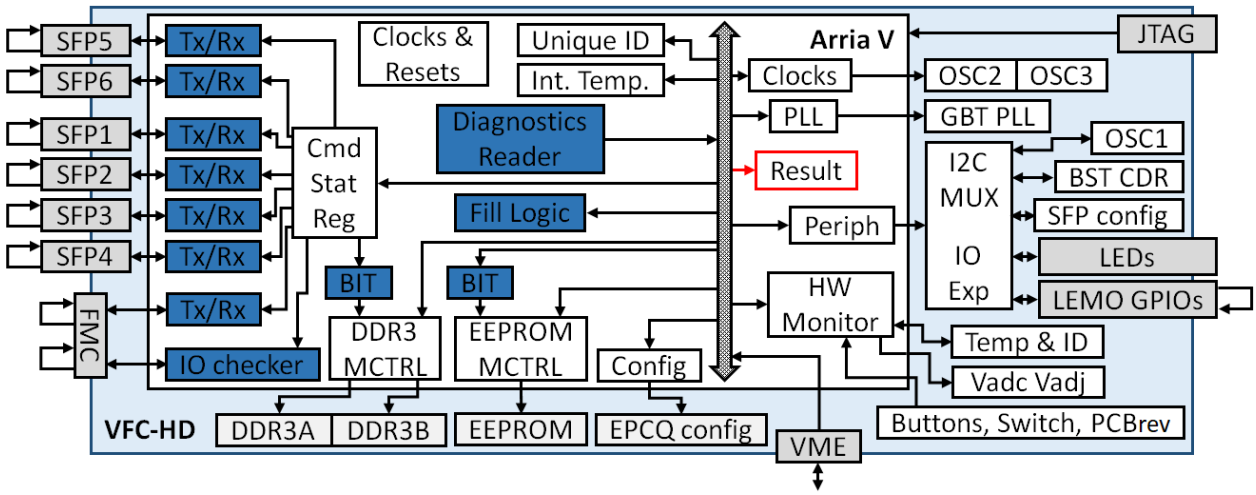

Figure 3: VFC-HD built-in-test (BIT) firmware overview.

Real-Time Software

A simple real-time (RT) program was written to read the status and the internal test results of all the boards under test. The real-time behaviour was needed to guarantee that the results of each test iteration were read and processed before the next one. The program also checked the functioning of the VME interfaces (e.g. data transfer, interrupts) C++ was used for the development, together with the CERN Front-End Software Architecture (FESA) Framework [7].

The RT program runs in a CPU in the same crate as the boards under test. In an ideal setup, we would have fully decoupled the devices under test from the system running the tests. However, to simplify the design and keep the cost at a reasonable level, we decided to have the test programs in the same VME system inserted in the climate chamber.

The frequency of the read-out of the status and hardware test results was adapted to the type of tests being run. A one second read-out rate was used for the burn-in phase, while one minute was used for the long-term tests.

User/Expert Application

A Java application was developed to drive the execution of the different tests and to monitor their proper execution. The engineer in charge of the test had a full overview of the tests that were running at any one time. This included nine VME crates with up to 16 boards each (144 boards in parallel). This was very useful to identify cases where the system was not running correctly allowing it to be immediately reconfigured, potentially saving up to two weeks of test run time.

The application also offered integration with the CERN Asset Management and Maintenance System. This is explained in detail in the next section.

An overview of the graphical display of the tool can be seen in Figure 5.

Data Storage and Analysis

The results of the tests were stored in the new CERN logging system, NXCALS, based on Apache Spark. This allowed for direct analysis of the time series data with Python scripts. We also used Jupyter notebooks intensively to facilitate the sharing of results (Figure 4).

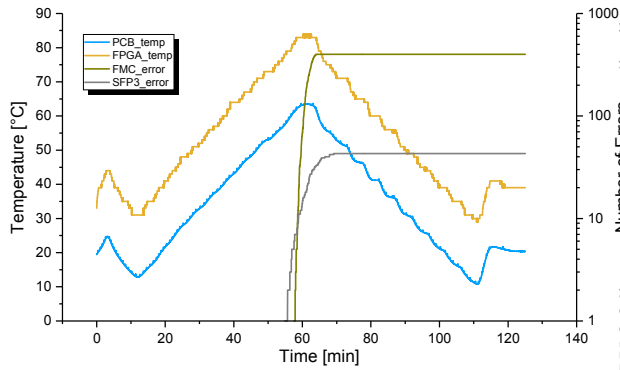

Figure 4: Analysis of correlation of transmission errors with temperature.

ASSET MANAGEMENT INTEGRATION

We wanted to trace the boards from a very early stage with the testing at CERN after the manufacturing. This would ensure higher reliability and the possibility to trace back any issues found during future operations.

The User Application interfaced with the CERN Asset Management System [8] through web services. This allowed checks on whether a board present in the test-bench was already registered in the system. If not, then the system allowed the engineer in charge of the test to register the board automatically.

The results of the different phases of testing were stored for each board in the CERN Manufacturing and Test Folder (MTF) system [9]. Non-conformities were also entered in MTF along with their correspondent details.

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-MOPHA049

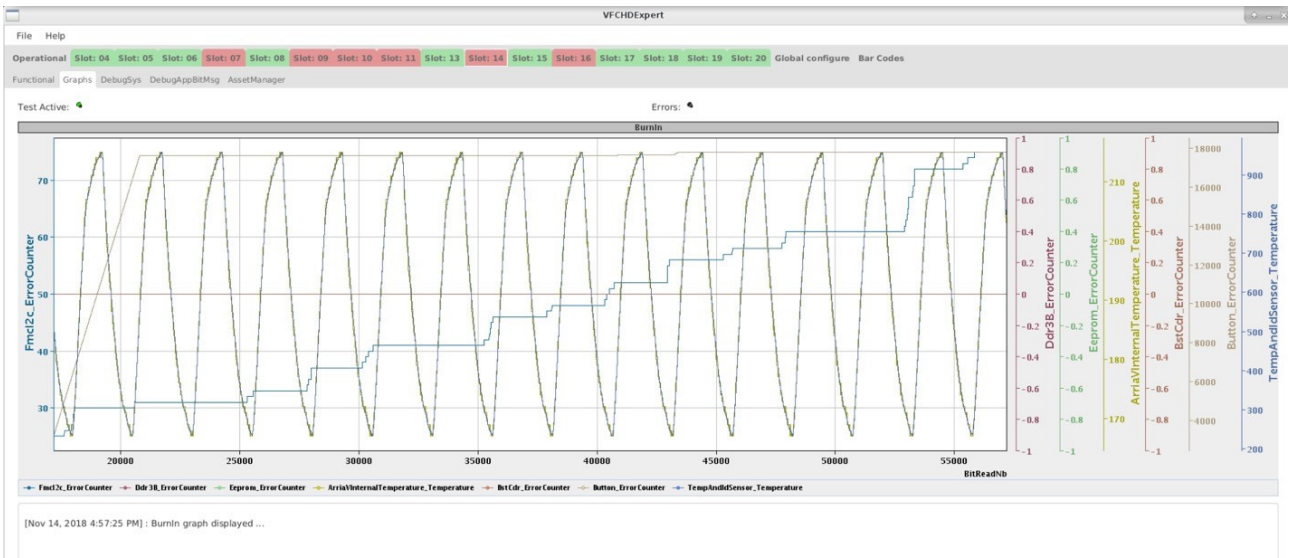

Figure 5: Overview of the user application during tests. We can see the errors appearing with the temperature cycling.

RESULTS

The validation tests started at the end of 2018 and lasted until the summer of 2019, when the nine-month long screening and run-in was completed for the 1200 boards. During that time, the test setup was constantly improved to address any issues found resulting in an advanced test-bench which is now available for any future test campaigns. The tests successfully validated the devices in different environmental conditions, determined the high temperature tolerance [10] and screened early life failures of the VFC-HD [11]. In this manner, the overall device dependability was determined, with a Mean Time To Failure (MTTF) above 200 000 h estimated following the run-in tests.

At the same time all the boards are now in the CERN Asset and Maintenance Management system, guaranteeing that we can keep track of the full history of the boards, from any early non-conformities in the manufacturing process to future operational issues.

MAIN CHALLENGES

The approach that was followed to build the test-bench was to use the same tools as in operation to minimize the overall effort. This turned out to be very useful in general but we faced issues linked to the nature of the execution of the tests and the fact that many of CERN's operational systems were not fully available during the accelerators shutdown period. As an example, the burn-in tests required the hardware configuration to be changed every day, with some of the services involved needing up to several hours to take into account the new configuration.

The test-bench evolved during its lifetime to cover more test cases as we got familiar with the behavior of the VFC-HD, and also to fix issues. We had more than 20 versions of the setup. We could only handle these changes following a data driven approach. The interface to the tests results were described in several spreadsheets that were processed and imported in the system with a suite of Python scripts.

CONCLUSION

We have presented the full design and implementation of a test-bench for the new acquisition electronics platform that will be used for many beam instrumentation systems at CERN. This is the first time such a complete and intensive testing methodology has been used for such systems.

The test-bench was very useful to identify early failures before installing the boards in the accelerators which should reduce the number of interventions required to fix issues in the future, with a clear positive impact on the availability.

While such a complete test-bench may not be appropriate for all projects in terms of cost-benefit, for large scale productions it is essential to ensure future reliability. Now that the initial investment has been made, the test-bench itself can be reused for other types of devices with only small adaptations required.

ACKNOWLEDGEMENTS

We would like to thank our colleagues from the different sections in the CERN Beam Instrumentation group that have supported the design and development of the electronics test-bench described in this paper.

REFERENCES

[1] O. R. Jones, "LHC Beam Instrumentation", in *Proc. PAC'07*, Albuquerque, NM, USA, Jun. 2007, paper THXC01, pp. 2630-2634.

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-MOPHA049

[2] B. Dehning *et al.*, "The LHC Beam Loss Measurement System", in *Proc. PAC'07*, Albuquerque, NM, USA, Jun. 2007, paper FRPMN071, pp. 4192-4194.

[3] E. Calvo-Giraldo *et al.*, "The LHC Orbit and Trajectory System", in *Proc. 6th European Workshop on Beam Diagnostics and Instrumentation for Particle Accelerators (DIPAC'03)*, Mainz, Germany, May 2003, paper PT08, pp. 187-189.

[4] A. Boccardi, M. Barros Marin, T. E. Levens, B. Szuk, W. Vigano, and C. Zamantzas, "A Modular Approach to Acquisition Systems for Future CERN Beam Instrumentation Developments", in *Proc. 15th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'15)*, Melbourne, Australia, Oct. 2015, pp. 1103-1106. doi:10.18429/JACoW-ICALEPCS2015-THHB2002.

[5] J. Serrano *et al.*, "White Rabbit Status and Prospects", in *Proc. 14th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'13)*, San Francisco, CA, USA, Oct. 2013, paper THCOCA02, pp. 1445-1448.

[6] D. Domínguez, J. J. Gras, J. Lewis, J. J. Savioz, J. Serrano, and F. J. Ballester, "An FPGA Based Multiprocessing CPU for Beam Synchronous Timing in CERN's SPS and LHC", in *Proc. ICALEPCS'03*, Gyeongju, Korea, Oct. 2003, paper MP531, pp. 113-115.

[7] L. Fernandez *et al.*, "Front-End Software Architecture", in *Proc. 11th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'07)*, Oak Ridge, TN, USA, Oct. 2007, paper WOPA04, pp. 310-312.

[8] CERN Asset & Maintenance Management platform, https://espace.cern.ch/cmms-service/default.aspx

[9] P. Martel *et al.*, "Manufacturing and Test Folder: MTF", in *Proc. 8th European Particle Accelerator Conf. (EPAC'02)*, Paris, France, Jun. 2002, paper TUPDO026, pp. 2688-2690.

[10] V. Schramm *et al.*, "Combined Testing and Validation Strategy for the new LHC BLM Processing Module," in Reliability and Maintainability Symposium (RAMS) 2019, Orlando, FL, US, 2019.

[11] S. Eitelbuß, "Screening and Reliability Testing of Beam Loss Monitor Electronics at CERN", Geneva, CH: CERN-THESIS-2019-050, 2019.