ISBN: 978-3-95450-221-9 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2021-THBR02

WHITE RABBIT AND MTCA.4 USE IN THE LLRF UPGRADE FOR CERN'S SPS

Tomasz Włostowski*, Andrew Butterworth, Grzegorz Daniluk,
Julien Egli, John Robert Gill, Tristan Gingold,
Juan David González Cobas, Michel Arruat, Grégoire Hagmann, Dimitrios Lampridis,
Maciej Marek Lipiński, Mattia Rizzi, Arthur Spierer,
Maciej Sumiński, Adam Artur Wujek, Karol Adrianek,
Predrag Kuzmanović, Philippe Baudrenghien, Saúl Novel González,
Julien Palluel, CERN, Geneva, Switzerland

Abstract

The Super Proton Synchrotron (SPS) Low-level RF (LLRF) system at CERN was completely revamped in 2020 [1]. In the old system, the digital signal processing was clocked by a submultiple of the RF. The new system uses a fixed-frequency clock derived from White Rabbit (WR) [2]. This triggered the development of an eRTM module for generating very precise clock signals to be fed to the optional RF backplane in MTCA.4 crates. The eRTM14/15 sandwich of modules implements a WR node delivering clock signals with a jitter below 100 fs. WR-clocked RF synthesis inside the FPGA makes it simple to reproduce the RF elsewhere by broadcasting the frequency-tuning words over the WR network itself. These words are received by the WR2RF-VME module and used to produce beam-synchronous signals such as the bunch clock and the revolution tick. This paper explains the general architecture of this new LLRF system, highlighting the role of WR-based synchronization. It then goes on to describe the hardware and gateware designs for both modules, along with their supporting software. A recount of our experience with the deployment of the MTCA.4 platform is also provided.

INTRODUCTION

The High Luminosity LHC (HL-LHC) project at CERN aims to increase the integrated luminosity of the LHC by a factor of 10 in the 10-12 years after its implementation in 2027. As a result of new requirements for the LHC beam, the whole CERN injector complex has undergone an extensive upgrade program. For the SPS, the synchrotron just upstream of the LHC, the new requirement for a beam intensity of 2.3×10^{11} protons/bunch with a bunch spacing of 25 ns resulted in a need to overhaul the complete accelerating system, including cavities, amplifiers and LLRF.

The new LLRF system needed to be ready in a short amount of time and required a significantly higher data transfer rate from the cards to the host PC than the legacy VME platform. A decision was therefore taken to capitalize on the effort of DESY and other institutes, as well as the Commercial-Off-The-Shelf (COTS) components using the MTCA.4 standard. Field Programmable Gate Arrays (FPGA) and System-on-Chip (SoC) solutions were used extensively to speed up the development and provide the flexibility allowing the system to be reused in future applications.

White Rabbit (WR) timing technology plays a key role in several parts of the system:

- It is used to derive a very low-noise base clock signal which is then used to sample the cavity antenna and beam Pick-Up signals, and to synthesize the RF drive signals.
- It is the transmission medium through which the LLRF system receives bending magnetic field information in real time.
- Coupled with Direct Digital Synthesis (DDS) technology, it allows the LLRF system to transmit the RF signals not as physical signals but as messages containing Frequency-Tuning Words (FTW) which allow receivers to replay the RF in synchronism all around the accelerator. This allows distributing the reference RF phase with fixed and very stable latency.

This paper does not aim at describing the upgraded LLRF system in general, but just the synchronization aspects and the use of the WR technology therein. After a quick introduction to the general architecture of the system, we describe the low-noise eRTM14/15 clock signal generation and distribution module. We then move on to the module allowing regeneration of RF signals based on reception of FTWs, the WR2RF-VME. Finally, we share our experience as new users of the MTCA.4 form factor, hoping it can be useful to others in their exploration for optimal platforms to implement LLRF or other types of systems.

GENERAL ARCHITECTURE

The overview of the SPS LLRF system architecture is shown in Fig. 1. Note that it focuses on the usage of WR in the system and omits many details of the RF part.

The renovated system drives six 200 MHz RF cavities and is implemented entirely in the MTCA.4 form factor. Each cavity has a dedicated Cavity Controller (CC), consisting of a Struck SIS8300KU card [3] and a DS8VM1 analog frontend/vector modulator RTM from DESY [4]. The CC's primary role is to maintain stable cavity power and phase, compensating the error introduced by the power amplifier (the Polar Loop) as well as handling the cavity beam loading

* Tomasz.Wlostowski@cern.ch

ISBN: 978-3-95450-221-9 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2021-THBR02

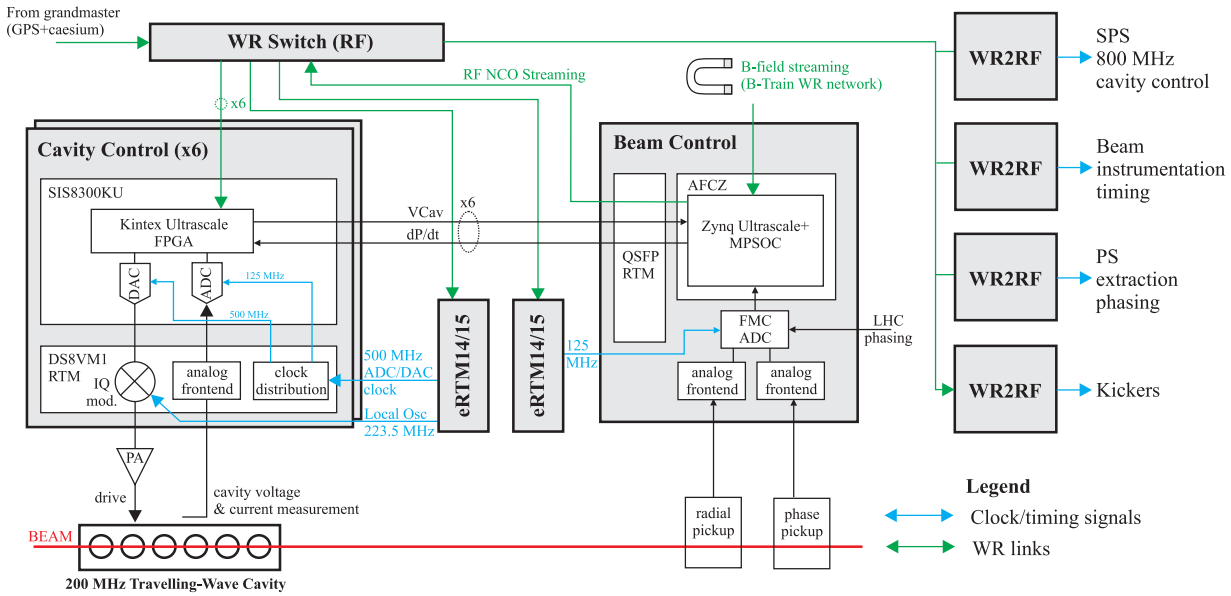

Figure 1: Overview of the WR-based SPS Low Level RF system.

and impedance effects (One-Turn Feedback). The measured cavity voltage information is streamed to the Beam Controller (BC) over dedicated gigabit links. Furthermore, the CC implements a monitoring and protection system which ensures safe operation of the high power RF equipment.

The BC uses the AFCZ board by Creotech [5] and two custom-made FMC mezzanines to interface with the power plant and the beam instrumentation. Its main task is calculating the instantaneous RF frequency and cavity voltage setpoints for a particular beam (particle type, intensity and fill pattern) during the acceleration ramp. The most significant inputs for the BC are the beam's transverse position, phase, current and the magnetic field in the machine's magnets. The latter is provided by a WR-based B-Train system [6], also using a dedicated WR link. The frequency information calculated by the BC is broadcast with every beam turn over the WR network using a protocol similar to the B-Train system. More details are provided in the section about the WR2RF-VME module.

Timing-wise, each CC requires two clock signals: the Local Oscillator for upconversion of the RF drive output (223.5 MHz sine wave) and a 500 MHz ADC/DAC clock. Both clocks have stringent phase noise (PN) and stability requirements: PN of −130 dBc/Hz at 1 kHz (carrier 223.5 MHz) and precision better than 13 ps (1 degree at 200 MHz). These requirements were set to minimize the detrimental effect of the RF noise during the long filling plateau for the LHC ions beams in the SPS. The clocks are provided by the eRTM14/15 module, described later. The BC needs a 125 MHz WR-synchronous clock for its ADCs, with similar noise and stability figures as the CC.

All six CCs share a single MTCA.4 crate (Schroff/nVent model 11890-170) with a single MCH/CPU set from N.A.T. [7] and a single eRTM14/15 module, whereas the BC resides in a separate MTCA.4 crate with its own CPU/MCH/eRTM boards. The crates are also equipped with the CERN General Machine Timing receiver cards (not included in the drawing for the sake of clarity). The low-jitter clock connections between the eRTM board and the CCs are provided by the DESY MTCA.4 LLRF backplane (RFBP) [8].

As the new LLRF system distributes the frequency information in a purely digital form, a need arises for interfacing with the legacy systems (mostly VME-based) that require an analog cavity frequency and RF-related trigger signals (e.g. kicker magnets, beam instrumentation, Proton Synchrotron (PS) extraction synchro and the auxillary 800 MHz SPS cavities). This task is performed by the WR2RF-VME cards, described in detail in the subsequent chapters.

THE ERTM14/15 MODULES

The eRTM14/15 [9] is a MTCA.4-compliant WR receiver providing a variety of ultra-low jitter clock signals which are distributed over the RFBP:

- Two DDS clocks (named LO and REF in this paper), distributed to slots 4..12 of the MTCA crate and with two front-panel outputs for up/down-modulation of the RF analog signals.
- Two digital clocks (distributed to slots 4..12 and the front panel).
- 10 MHz and Pulse Per Second (PPS) input and output for interfacing with non-WR timing devices.

The module is made of two boards connected with a high speed board-to-board cable (Samtec QSH series), as depicted in Fig. 2. The board residing in slot 14 contains the FPGA (Kintex-7 XC7K70T) that implements the WR synchronization stack (the WR PTP Core - WRPC [10]), two redundant SFP uplink ports and basic ("low performance") WR local oscillators. Slot 15 contains the high performance oscillators and analog circuitry synthesizing the clocks listed above. Therefore, the eRTM14 can work alone as a basic

ISBN: 978-3-95450-221-9 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2021-THBR02

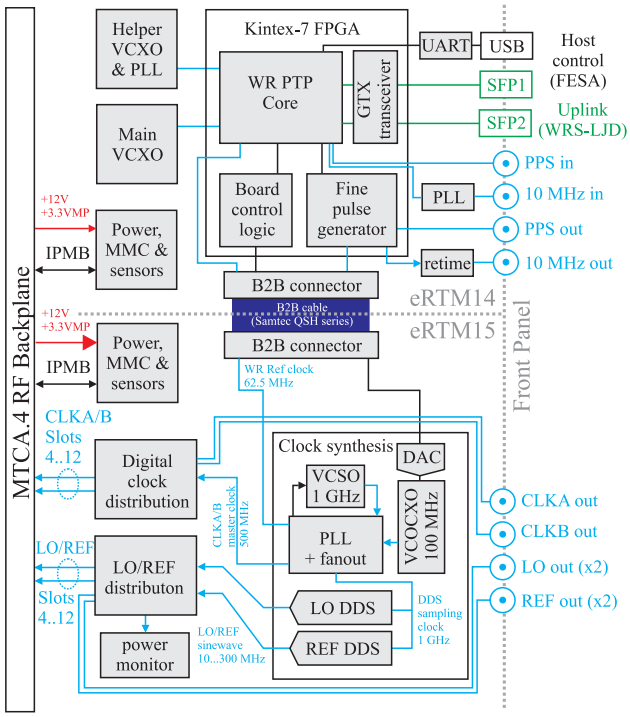

Figure 2: Simplified block diagram of the eRTM14/15 WR Receiver.

WR receiver or together with the eRTM15 with the full set of features.

To achieve the timing performance required by the SPS LLRF system, we had to develop novel techniques described in the two next subsections.

WR Clock Synthesis

The WR clock recovery and synthesis was designed to meet a −130 dBc/Hz phase noise figure at 1 kHz offset (carrier frequency 223 MHz) and less than 100 fs rms jitter (100 Hz - 10 MHz).

Meeting the latter requirement necessitated a local oscillator with sufficiently low PN in the low frequency range to match −100 dBc/Hz, which is the noise floor of the Digital Dual Mixer Time Difference (DDMTD) phase detector used by WR at the offset of 50 Hz, much less than the bandwidth of standard WR PLL implementations (about 300 Hz). An Axtal AXIOM45 series OCXO has been selected, as a compromise between the performance and the price.

The OCXO provides a 100 MHz sinewave signal, fine-tuned within the range of a few ppm by a low-noise DAC (AD5660) controlled by the WRPC. This 100 MHz is multiplied by the subsequent PLL synthesizer (LTC6950) that delivers the 62.5 MHz WR reference clock (thus closing the WR clock recovery feedback loop), a 500 MHz master clock for the CLKA/B outputs and a 1 GHz sampling clock for the LO/REF DDS synthesizers. The PLL uses a discrete voltage-controlled SAW oscillator (Crystek CVCSO-914) to further improve the PN in the high frequency range and provide enough headroom for the additive PN of the DDS and clock distribution circuitry.

The LO/REF clocks are produced by two AD9910 integrated DDS synthesizers, capable of outputting up to 400 MHz with a tuning resolution of 0.23 Hz. The eRTM design allows them to be programmed to any arbitrary (non-modulated) frequency between 10 and 300 MHz. The output of each DDS is followed by a two-stage clock distribution network, each stage consisting of a Mini-Circuits GVA-81 amplifier followed by a passive 1:4 splitter. As the distribution network is partly passive, the outputs are coupled to the RFBP with individually programmable MEMS RF switches, terminating the unused slots. Each LO/REF RFBP output is also equipped with an RF power monitor. The module also allows to adjust (globally) the output power of the LO and REF chains between +9 and +19 dBm, achieved by digitally changing the DDS DAC bias current.

Furthermore, the DDS outputs can be phased together at an arbitrary moment of time through a message sent over the WR network. Such a message contains a TAI timestamp at which the Fine Pulse Generator core in the FPGA generates a pulse to reset the AD9910 phase accumulators, thus ensuring all the chips start producing the LO/REF signals at exactly the same phase. The SPS LLRF uses this feature at the start of each beam cycle to ensure the LO's in all WR nodes are correctly aligned.

The CLKA/B clocks are digital, differential (LVPECL levels) intended for driving ADC/DACs. They are derived from a 500 MHz master clock through two LTC6953 clock dividers/fanout buffers. The user can select from the frequencies of 62.5, 125, 250 and 500 MHz, each output capable of independent frequency, phase and on/off setting. To ensure the fixed phase alignment of the divided-down clocks across different nodes in the network, the FPGA's Fine Pulse Generator core provides PPS-aligned divider synchronization pulses for the LTC6953s.

The PN figures, measured for CLKA and LO channels

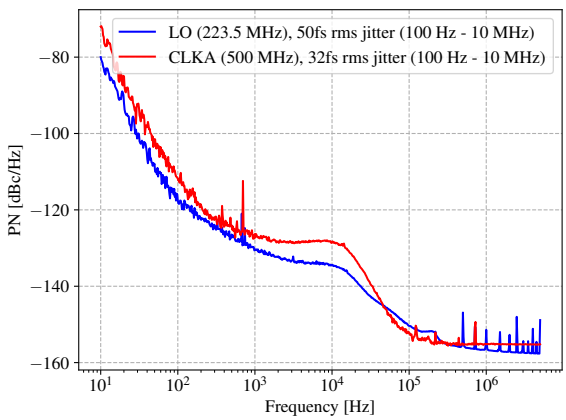

Figure 3: Phase noise of the CLKA (500 MHz) and LO (223.5 MHz) outputs of the eRTM.

Furthermore, for the eRTM to achieve the required performance, the upstream elements of the SPS WR network had to be improved:

ISBN: 978-3-95450-221-9 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2021-THBR02

- Changing the master GPS receiver to Microchip Syncserver S650 [11], as the previous receiver showed small phase jumps (few dozen ps), sufficient to disrupt the operation of the LLRF system referenced to it.
- Upgrading all the upstream WR Switches [12] to the Low Jitter Daughterboard (WRS-LJD) version [13], which offers lower additive PN and improved DDMTD tracking stability. More details can be found in [14].

Phase Stability

The second requirement of the SPS LLRF that resulted in substantial development effort was the phase stability (precision) better than 13 ps. In other words, every time the system is switched on, the clock phases between all RF nodes must not change by more than 13 ps. Standard WR implementations offer a precision of around 100 ps, which we had identified to come from the internal clocking and Clock/Data Recovery (CDR) of the FPGA's Gigabit transceivers. We focused on the Xilinx's GTXE1 and GTXE2 transceivers, used, respectively, in the Virtex-6 and Kintex-7 series of FPGAs and improved their phase stability to better than 5 ps (std). The excessive phase error is caused by multiple effects:

- The RX/TX buffer bypass logic aligning the SerDes high clock with the FPGA fabric clock, which has been implemented as a tapped delay line, with the tap selected during the initialization of the link.
- Various clock dividers (most importantly, the RX SerDes bit clock to parallel word clock).

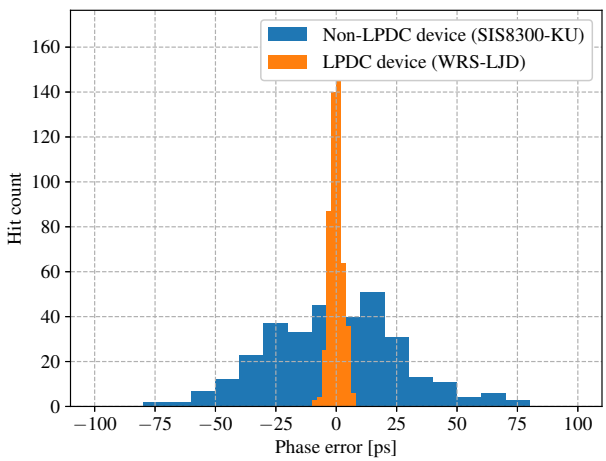

Figure 4: Phase stability comparison between LPDC and standard WR nodes.

These effects have been mitigated by instrumenting the transceiver internal clocks with a DDMTD-based phase measurement unit and repeatedly resetting the transceiver until the clock phases are fixed on previously calibrated values. The detailed information about the procedure is available in [15]. The WR devices which offer such improved phase stability, such as the version 6.0 of the White Rabbit Switch, are marked as Low Phase Drift Calibration (LPDC). Figure 4 compares the phase stability of a standard WR PHY implementation with the LPDC-enabled transceiver of the eRTM board.

Digital and Software

All timing features of the eRTM module are controlled by the Kintex-7 FPGA, encompassing the WRPC, a set of I2C/SPI/GPIO interfaces to control all onboard chips and a Fine Pulse Generator core to produce the PPS output and the counter/phase accumulator synchronization pulses.

The WRPC contains an embedded CPU core (LM32) which runs a customized version of the WRPC software. The key differences between the standard WRPC and the eRTM are:

- Support for eRTM board-specific circuitry (PLLs, DDS, clock distribution).
- Improved regulation loops of the WR PLL (PI gain optimization, gain scheduling).
- Remote control and monitoring of the board through the front panel USB port. A library in C/Python is provided, which allows a remote PC to control the most important run-time parameters of the board (clock frequencies, channel assignment, etc.).

Being an MTCA.4-compliant module, the eRTM implements two MTCA Management MicroControllers (MMCs), one per board. These are STM32F3 series MCUs, running a customized version of the LNLS's OpenMMC firmware [16]. Their key responsibilities are monitoring the board health (voltages, currents and temperatures), power sequencing and MTCA.4 identification of the boards.

THE WR2RF-VME MODULE

The WR2RF-VME card is a VME-compatible WR RF receiver, featuring:

- High performance WR receiver.
- x2 RF channels, each capable of reconstructing the modulated RF signal.
- x2 Programmable trigger outputs per RF channel.
- A number of standard timing I/Os (10 MHz/PPS) and general-purpose digital inputs/outputs.

The simplified block diagram of the WR2RF-VME is depicted in Fig. 5. As the board inherits the clock generation circuitry (OCXO, PLL and DDS) from the eRTM board, we will focus exclusively on the RF-specific features.

RF Streaming and Reconstruction

The SPS Beam Control distributes Ethernet frames known as the "RF-train" over the WR network to Cavity Controllers and WR2RF-VME cards. These frames contain Frequency Tuning Words (FTW) that describe the setpoints for Numerically Controlled Oscillators (NCOs) implemented within the Kintex7 FPGA on the WR2RF-VME card. To ensure all receivers have the same NCO value, two things are necessary:

- Every frame from Beam Control is transmitted using WR Streamers, a feature of the WR-Core that guaran-

ISBN: 978-3-95450-221-9 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2021-THBR02

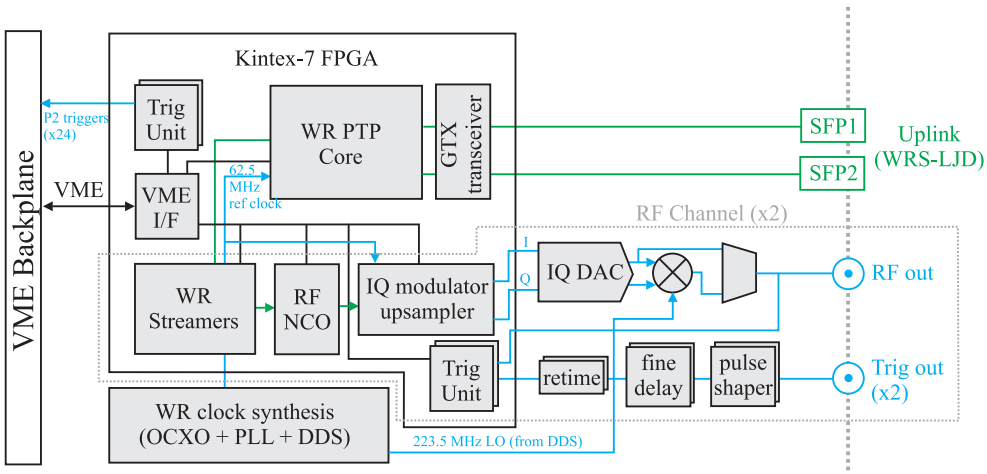

Figure 5: Simplified block diagram of the WR2RF-VME board.

tees a fixed latency delivery (the fixed latency must be larger than the worst case network latency).

- Each NCO is reset prior to starting a new accelerator cycle.

These features ensure that each NCO, in every receiver, has the same value on any WR clock tick.

The RF is reconstructed using a mixer approach. The FPGA uses the output of the NCOs to drive I and Q data lines for an AD9783 Digital to Analog Converter (DAC), with a varying frequency close to 23.5 MHz. The DAC generates the Intermediate Frequency (IF) for the LTC5598 mixer. Meanwhile, the Local Oscillator (LO) input of 223.5 MHz for the mixer is derived from an AD9910 DDS chip, as described for the eRTM modules. The output from the RF mixer is then filtered to accept the downside conversion and attenuate the LO and image frequency leakage close to 200 MHz. This is the reconstructed RF signal.

Trigger Units

In LLRF systems at CERN, a Trigger Unit (TU) [17] is used to generate a pattern of pulses that are synchronized with an RF signal. These TUs, implemented in FPGA logic, are clocked by a digital version of the reconstructed RF signal to create an RF synchronous system. In the WR2RF-VME card there are two cascaded TUs, and typically they create:

- The revolution frequency or orbit clock.
- A bunch clock, which is a divided by 5 digital version of the RF signal.

Applications of these TUs are:

- Precise RF bucket selection to adjust the injection of particle bunches.
- Timing pulses to control systems responisble for dumping the beam.
- Signals for beam instrumention and monitoring.

The output from the TU then drives a discrete pulse shaping and retiming circuit on the WR2RF-VME card, allowing for low jitter trigger signal generation and precise adjustment of trigger delay (10 ps steps).

EXPERIENCE WITH MTCA.4

The SPS LLRF system is the first medium-scale deployment of the MTCA.4 platform in the Accelerator Sector at CERN. While MTCA.4 has easily met the project's requirements regarding the CPU processing power and the PCI Express bandwidth compared to the VME platforms, we identified several issues affecting the reliability and integration of the platform in the CERN Controls framework, in particular the LLRF backplane and the modules which utilize it. The list below applies only to the hardware models used at CERN and not the MTCA.4 ecosystem as a whole:

- Suboptimal thermal design of the MTCA.4 crate and fan speed management: the fans in the front were too powerful with respect to the fans in the rear. As the crate's cooling air inlet is in the front, the stock version of the fan unit would direct almost all of the airflow to the front part of the crate, thus causing overheating of the rear modules such as the Rear Power Modules (RPMs) and the eRTM board. We have observed temperatures of ~ 90 °C on the eRTM at an approximate power dissipation of 36 W (per 2 slots). Furthermore, the eRTM15 slot that dissipates the most heat in our system receives less airflow compared to other RTM slots. This is due to the placement of the fan too far from the slot.
- Unreliable Power Module (PM) firmware. The dedicated RFBP rear PMs would refuse to work with the eRTM module without indicating any error. We continued our development using standard PMs, but the MCH refused to recognize the eRTM boards (no IPMI/I2C communication).
- Bugs in the MCH, such as non-functional "upstream slot power delay" parameter causing PCIe enumeration issues for FPGA-based PCIe cards.
- CPU configuration process. The BIOS/UEFI settings of the CPU module are extremely complex and there is no recovery mechanism ("CMOS reset"). Incorrect settings can render the board permanently inoperable.

ISBN: 978-3-95450-221-9 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2021-THBR02

It is worth noting that after an intense, 4-month collaboration with the vendors, the above issues have been fixed or mitigated by the companies, thus making the MTCA.4 platform technically viable for current and future applications in the Accelerator sector at CERN.

On the non-technical side, the SPS renovation project's relatively short timespan required the RF team to choose from COTS MTCA.4 modules instead of the traditional in-house-developed electronics. Unfortunately, to obtain the basic information required to integrate one of the modules, a Non-Disclosure Agreement was needed. This is undesirable in a scientific environment and the long signature process even impacted the project planning.

CONCLUSION

The MTCA.4-based SPS LLRF system has been in commisioning since April 2021 and has operated with physics beams since July 2021. After 6 months of successful commissioning, the results are very positive. We plan to use this novel solution for the prototype upgrade of the HiLumi LHC Accelerating and Crab Cavity LLRF systems.

The key highlights are:

- Timing performance previously reserved for analog RF distribution systems in a WR-network based solution.
- Proven RF distribution over WR, driving accelerating cavities in an operational machine (SPS).
- Compactness: taking two 9U MTCA.4 crates and one 19-inch rack with custom electronics instead of over ten 19-inch racks.
- Backplane connectivity for low-jitter clocks and RF signals reducing cabling cost and complexity.
- Bandwidth of the PCI Express fabric provided by MTCA.4 outperforming the legacy VME platform by two orders of magnitude and allowing for much more sophisticated beam diagnostics.
- Usage of COTS modules significantly shortened the development time (especially with the SIS8300-KU and AFCZ boards).

The SPS LLRF project also triggered the development of ultra-low phase noise and improved precision White Rabbit equipment: the eRTM14/15, WRS-LJD and WR2RF-VME which offer unprecedented performance among the available WR devices. These boards are available under the CERN Open Hardware Licence [18] and can be easily adapted to other LLRF systems as well as beam instrumentation or experiment electronics.

ACKNOWLEDGMENTS

The authors would like to thank the CERN Electronics Design Office and Assembly Workshop, Creotech, N.A.T., Struck Innovative Systeme, Javier Serrano and Erik van der Bij for their support during the development and commissioning of the SPS LLRF project.

REFERENCES

[1] G. Hagmann, P. Baudrenghien, J.D. Betz, J. Egli, F.J. Galindo Guarch, G. Kotzian, *et al.*, "The CERN SPS Low Level RF upgrade Project", in *Proc. IPAC'19*, Melbourne, Australia, May 2019, pp. 4005–4008, doi:10.18429/JACoW-IPAC2019-THPRB082

[2] White Rabbit project website, https://ohwr.org/project/white-rabbit/wikis

[3] SIS8300KU digitizer, manufacturer webpage, https://www.struck.de/sis8300-ku.html

[4] DESY DS8VM1 up/downconverter RTM module, project webpage, https://techlab.desy.de/products/rtm/drtm_ds8vm1/index_eng.html

[5] AMC FMC Carrier, Zynq Ultrascale variant (AFCZ), project webpage, https://github.com/elhep/AFCZ

[6] M. Lipiński, "Real-Time distribution of magnetic field values using White Rabbit the FIRESTORM project", 2017. https://www.ohwr.org/attachments/5795/BE-CO-TM-WR-BTrain.pdf

[7] NAT MTCA MCH and CPU modules, manufacturer webpage, https://www.nateurope.com/products/NAT-MCH-PHYS.html

[8] MTCA.4 RF Backplane, product webpage, https://techlab.desy.de/products/other_products/rfb/index_eng.html

[9] WR eRTM14/15, project webpage, https://ohwr.org/project/ertm15-llrf-wr

[10] White Rabbit PTP Core (WPRC), project webpage, https://ohwr.org/project/wr-cores/wikis/Wrpc-core

[11] Microchip SyncServer S650 GPS Receiver, product webpage, https://www.microsemi.com/product-directory/gps-instruments/4135-syncserver-s650

[12] WR Switch, project webpage, https://ohwr.org/project/wr-switch-hw/wikis

[13] WRS-LJD Daughterboard, project webpage, https://ohwr.org/project/wrs-low-jitter/wikis

[14] M. Rizzi, M. Lipiński, P. Ferrari, S. Rinaldi, and A. Flammini, "White Rabbit clock synchronization: ultimate limits on close-in phase noise and short-term stability due to FPGA implementation", *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 65, pp. 1726-1737, Sept. 2018. doi:10.1109/TUFFC.2018.2851842

[15] T. Włostowski, "Achieving deterministic phase in Xilinx GTX transceivers ", https://ohwr.org/project/white-rabbit/wikis/deterministic-phase-in-xilinx-gtx-transceivers

[16] H.A. Silva and G.B.M. Bruno, "openMMC: An Open Source Modular Firmware for Board Management", in *Proc. PCaPAC2016*, Campinas, Brazil, Oct. 2016, pp. 94–96. doi:10.18429/JACoW-PCaPAC2016-THPOPRPO04, 2017.

[17] D. Barrientos, J. Molendijk, and G. Hagmann, "A 1 GHz RF Trigger Unit implemented in FPGA logic", https://cds.cern.ch/record/2314652/files/1803.09041.pdf

[18] CERN Open Hardware Licence, project webpage, https://cern-ohl.web.cern.ch