UPGRADE OF THE LHC LOW-LEVEL RF BEAM-CONTROL FOR THE DIGITAL FREQUENCY DISTRIBUTION OVER WHITE-RABBIT

A. Spierer*, T. Gingold, A. Goujon, G. Hagmann, M. Jaussi, L. Kaffka, H. Timko,
CERN, Geneva, Switzerland

Abstract

In preparation for the High-Luminosity LHC (HL-LHC) era, a significant upgrade to the LHC Low-Level RF (LLRF) Beam-Control system is being implemented. The installation of crab cavities during Long Shutdown 3 (LS3, 2026-29), together with the obsolescence of existing components, necessitates a modernization of the LHC LLRF architecture.

The upgraded Beam-Control system – based on MicroTCA and VME hardware – replaces the frequency program and reconstructs a digital master RF. These frequencies are transmitted over a White-Rabbit (WR) network in the form of Frequency Tuning Words (FTW) to maintain synchronization among the distributed RF devices. This approach enables a scalable RF distribution infrastructure across the accelerator complex and experiments, with improved phase stability and reduced noise.

This paper presents the motivation and design of the upgraded system, including the method used to achieve robust phase synchronization over WR. Finally, it shows results obtained from a dedicated test with beam acceleration in the LHC.

MOTIVATION

In the LHC, LLRF signals are distributed throughout the machine to synchronize equipment with the circulating beam. This equipment includes experiment triggers, beam instrumentation, and RF cavity controllers. For the HL-LHC, crab cavities and their controllers will be installed around the interaction points, several kilometers away from the LLRF Beam Control system located at LHC Point 4.

To support this upgrade, a digital RF distribution over WR has been deployed, based on a first implementation in the SPS [1]. The concept is to send frames containing FTWs, reference phases, and control signals over the WR network. These are received and applied to a Numerically Controlled Oscillator (NCO) in the RF nodes with deterministic latency.

This new distribution is installed in parallel with the existing analog-over-fiber system, allowing performance comparisons prior to LS3. This project involves not only the RF group but also the experiments, controls, and beam instrumentation groups, as they benefit from the improved performance and flexibility of the new system.

This architectural choice offers several advantages: built-in mechanisms to compensate for temperature related phase drifts, cable length compensation, a common infrastructure for RF and machine timing, scalability through Ethernet networking, and progress toward a unified distribution across the accelerator complex. In the case of the LHC, it also provides an opportunity to replace obsolete components.

The main challenges are ensuring a high system availability, phase recovery after power cycles, and preservation of phase noise and latency.

* arthur.spierer@cern.ch

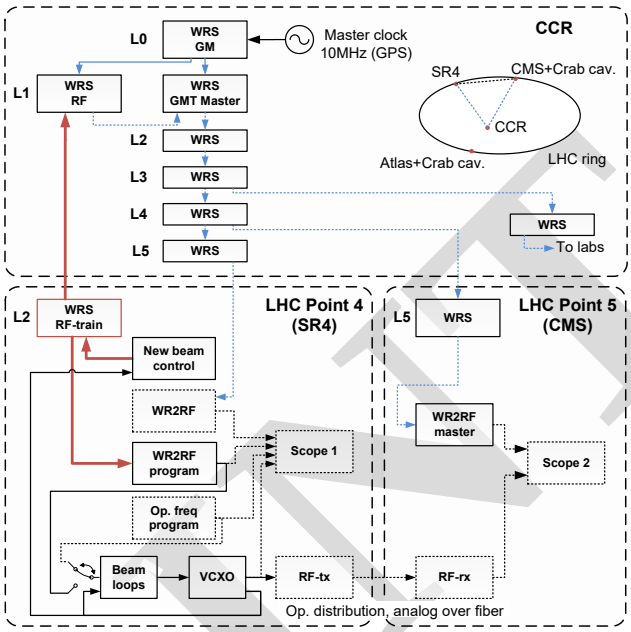

Figure 1: Test setup in the LHC. In red: the latency critical path, in blue: RF distribution over the general timing network. All dashed line equipment will be decommissioned after LS3.

IMPLEMENTATION

Hardware

The Beam-Control upgrade is based on the AFCZ MicroTCA board and WR2RF VME board [2]. They replace several VME-based modules for the function generator (FGC), frequency program, analog and optical distribution.

The function generator is now embedded in the AFCZ FPGA and plays the predefined frequency program of the LHC (around 400 MHz). The frequency is converted to an FTW and transmitted within the WR Ethernet frame (Table 1) once per beam revolution. The RF receivers such as the WR2RF board, reconstruct the analog RF from the received frame.

The master oscillator (VCXO) is locked to the reconstructed analog frequency program via the existing synchro loop and serves as the reference for the accelerating cavities. In the AFCZ, a digital PLL is locked to the VCXO using a 125 MS/s ADC and its frequency correction is used to

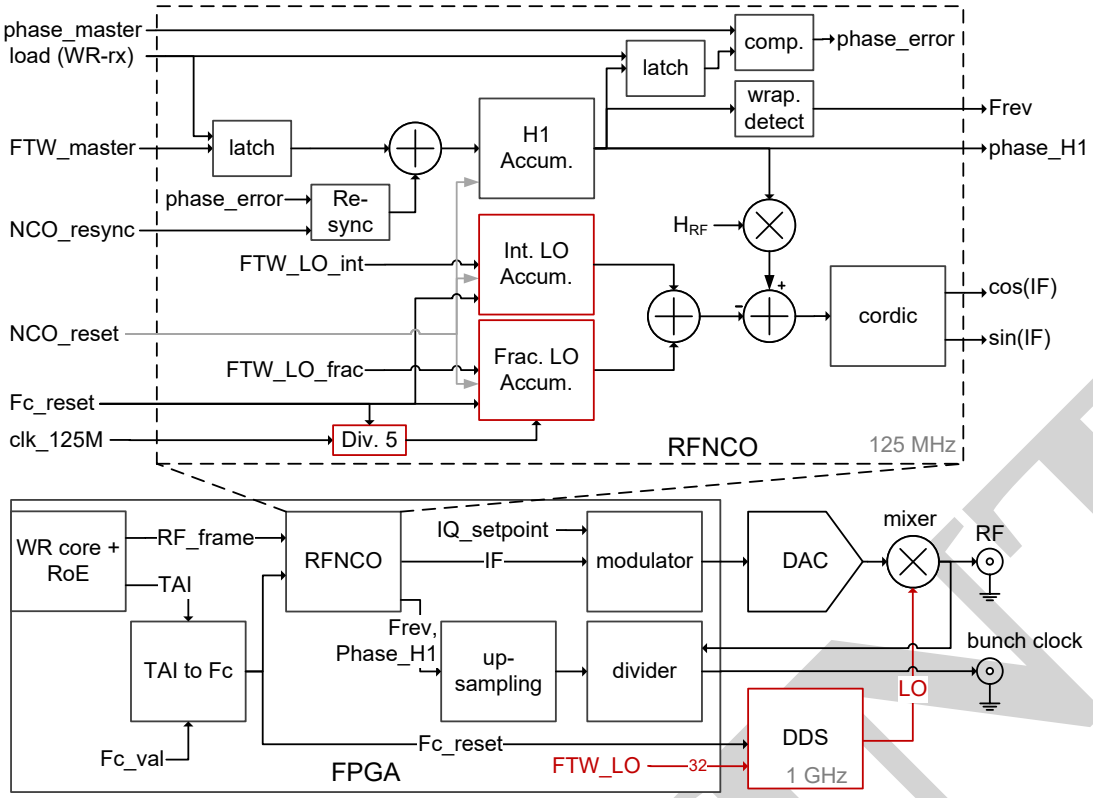

Figure 2: Synchronization features implemented in the WR2RF board. In red, the fixed frequency part.

reconstruct the master FTW. This is also transmitted in the WR frame, allowing any device to generate either the master or program RF independently for both LHC rings.

Table 1: WR Frame Fields Description, * : One per Beam

Field	Description
FTW_prog*	48 bit program frequency
FTW_master*	48 bit program + beam loops
controls	NCO_reset, NCO_resync, DDS_resync
phase_prog*	48 bit reference program phase
phase_master*	48 bit reference master phase
	Total Payload 50 bytes

Network and Infrastructure

The distribution network consists of multiple layers of WR switches (WRS). Layer 0 (L0) is the top layer, locked to the GPS master clock, while end nodes are connected at layer 5.

The RF frame generated by the Beam-Control is injected at layer 1, where the Ethernet traffic is merged with other accelerators' frames and the timing network. An exception is made for the Beam-Control loops and the distribution to the accelerating cavities, where latency constraints prevent traversal of the full network (shown in red in Fig. 1).

The RF-over-Ethernet (RoE) protocol [3] is adopted for the first time in an accelerator and will serve as a standard.

Each WRS adds approximately 5 us latency, in addition to fiber propagation delay. In the test setup, latency-critical paths are set to 10 us, while general network nodes are set to 128 µs, accounting for multiple switch stages and over 10 km of fiber.

Function Generation and Firmware

An embedded function generator has been developed to replace the legacy hardware. It is fully integrated into the LHC settings management system.

It uses tables of 4096 vectors of time intervals and slopes with interpolation, supporting the 30-minute LHC ramps with a frequency resolution of 16 µHz at 400 MHz.

Orbit feedback, previously handled by the VME FGC, is now applied directly in the Beam-Control firmware. A UDP stream carrying frequency corrections derived from beam position monitor measurements is received by the MicroTCA front-end computer (FEC) and written directly to the FPGA.

Synchronization Mechanisms

Devices generating RF from WR must provide reproducible phase relative to the master VCXO, regardless of startup time or location. Since RF clocks must remain continuous, global resets cannot be applied as they would cause phase discontinuities. Synchronization mechanisms are therefore implemented to recover phase automatically.

Synchronization is ensured by the following principles, also shown on Fig. 2:

1. All devices use WR-recovered clocks, which are inherently phase-aligned.

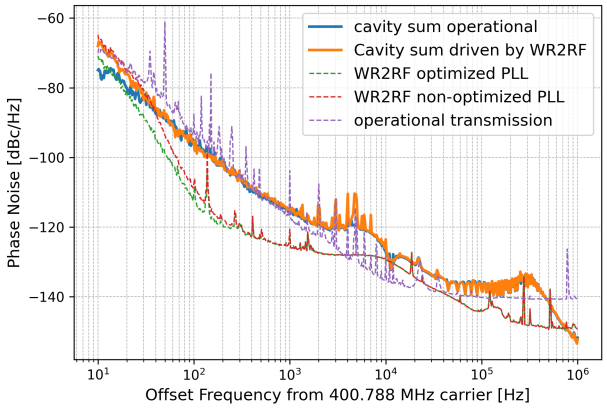

Figure 3: Comparing the phase noise of the current and new RF reference generation. Tuning the PLL shows an improvement in phase noise at the output of the WR2RF (dashed traces).

2. Fixed-latency transmission ensures that FTWs are applied synchronously across receivers, preserving relative phase from startup.
3. Reference phase from the Beam-Control NCO is transmitted and compared to the local NCO phase; the error is corrected using a ramp to avoid glitches.
4. Local oscillator (LO) frequencies are chosen to be phase-synchronizable using a cyclic pulse fc, derived from International Atomic Time (TAI).
5. Sub-harmonics, such as bunch clocks and revolution frequency are synchronized with the RF using high-resolution pulses derived from NCO phase information, achieving ~78 ps resolution.

Frequency Choices

Different devices impose different RF constraints. For example, the WR2RF uses:

- a 32-bit DDS at 1 GHz for analog LO (~413 MHz),
- a 48-bit NCO at 125 MHz for digital LO and RF.

Frequencies are defined using the least common multiple (LCM) of their resolution steps (232 mHz). This process is extended across all systems.

To reduce synchronization time, frequency granularity is coarsened, reducing the period of fc. In cases where frequencies cannot be matched directly, fractional accumulators are used to extend resolution.

Delay Compensation

For each receiver, known transmission delays and instantaneous RF frequencies allow compensation of phase shifts during ramps. The device's azimuthal position is also considered to compensate for beam time-of-flight. Once parametrized, the compensation is done automatically in the RFNCO IP core.

RESULTS

Phase Reproducibility and Drifts

Phase reproducibility was validated in CMS by comparison with beam phase measurements and the existing distribution system. With resets at each machine cycle, which are only necessary at the current development stage, phase variations between the two systems remain within RF noise.

Phase recovery after power cycling and glitch-free resynchronization were successfully demonstrated in the laboratory.

Measurements comparing local and global distributions showed slow drifts (up to ~60 ps pk-pk) [4], correlated with temperature variations. Improvements are proposed.

Phase Noise

Figure 3 shows the close-in phase noise of the analog master RF distribution, the reconstructed Master RF with the WR2RF board, and the accelerating cavity sum of one beam.

- above 50 Hz the WR2RF significantly improves the phase noise of the RF distribution, except with a small degradation between 4 and 30 kHz. Above 50 KHz, the operational RF distribution phase noise floor is limited by the analog-over-fiber transmission, a limitation not present in RF over WR.
- below 50 Hz, the phase noise is degraded by the WR recovered clock noise. A 7 dB improvement at 20 Hz is expected by optimizing the PLL settings, and from the next WR2RF hardware generation.

CONCLUSION

The upgraded LLRF system and WR-based RF distribution were tested in parallel with the operational system over one year. A dedicated test using the upgraded system for beam acceleration was also performed.

Results show:

- slight increase in low-frequency phase noise (mitigable),
- improved phase noise and temperature stability for long-distance distribution,
- successful phase recovery across RF nodes.

These results validate the design choices for the HL-LHC and provide a basis for further improvements. Continued measurements in 2026 will support future upgrades and consolidation of the LHC LLRF system.

ACKNOWLEDGEMENTS

The authors would like to thank the LHC operations team, and particularly M. Hostettler, for the integration of the system and support during machine development, as well as all colleagues involved from Electronic Systems for Experiments and Controls Electronics who contributed to the successful deployment and measurements.

REFERENCES

[1] A. Spierer *et al.*, "The CERN SPS Low Level RF: The Beam-Control", in *Proc. IPAC'22*, Bangkok, Thailand, Jun. 2022, pp. 895–898. doi:10.18429/JACoW-IPAC2022-TUPOST021

[2] G. Hagmann *et al.*, "HL-LHC RF Distribution over White-Rabbit: The WR2RF and eRTM modules". https://indico.cern.ch/event/1470887/

[3] "RoE structure-aware direct digital synthesis mapping using an OUI subType". https://roe-mapping.web.cern.ch/documents/RoE-DDS-mapping-v1.4.pdf

[4] L. Kaffka *et al.*, "Assessing the performance of future White Rabbit RF and timing distribution system for HL-LHC under real conditions", *J. Instrum.*, vol. 21, p. C04004, 2026. doi:10.1088/1748-0221/21/04/C04004